Cognitive Architectures for Robot Agents

Current Capabilities, Future Enhancements and Prospects for Collaborative Development

TransAIR

CONTENT

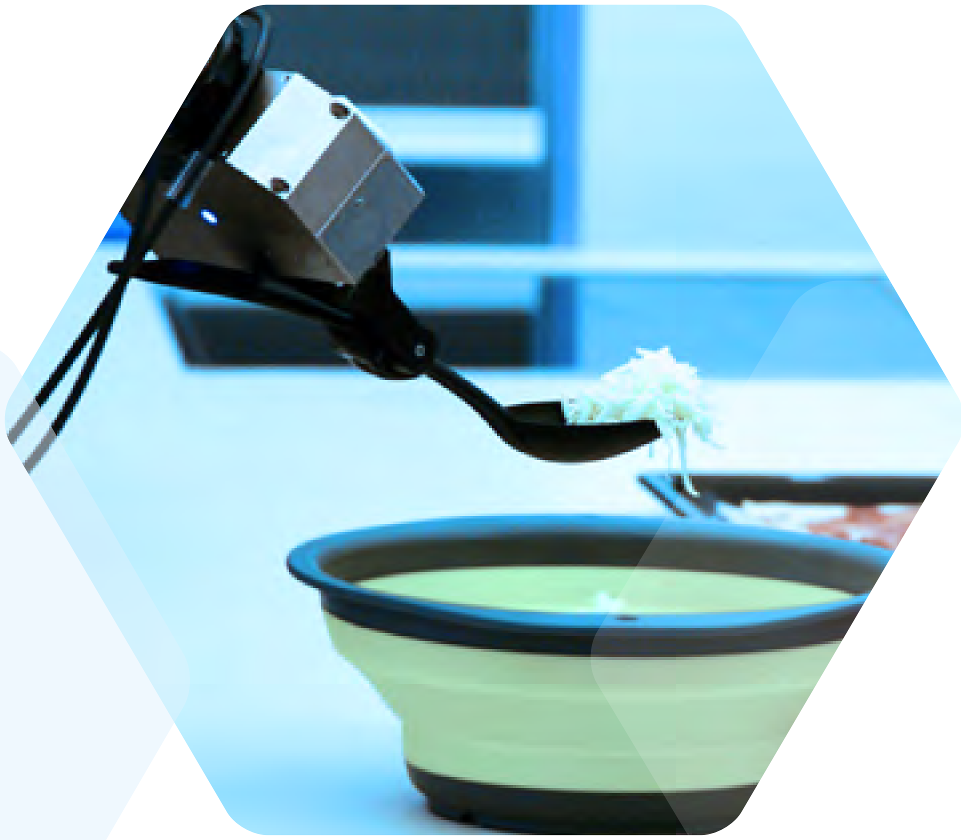

WELCOME

From March 22 to March 28, 2020, researchers from Europe, the United States and various countries around the globe met in a virtual venue to discuss the current capabilities and future prospects of cognitive architectures.

As we are building increasingly powerful AI systems and robot agents, we also feel more pressure to improve the orchestration of representations and computation processes required for achieving competent agency. Probably one of the biggest lessons we have learned over the last 30 to 50 years is the role that actions play in intelligent agency. In the beginning, many considered actions to be atomic entities that can be modeled in a fairly simple way. We tried to realize competent agency through reasoning about these action models.

I think this view has shifted significantly. Many of us feel that the key to competent agency can be found inside the actions. One essential aspect and open research challenge is understanding the reasoning and decision-making which is needed to translate under-determined action tasks, such as "put the coffee on the table," into body motions that achieve the desired effect and avoid the unwanted ones.

In this workshop, we saw a number of different approaches and perspectives. We had lively exchanges in the Q&A sessions and particularly during panel discussions, where audience members weren't shy to ask challenging questions. Despite some diverging views, a number of common themes emerged, including concepts such as artificial episodic memory, internal simulation, and hybrid symbolic or subsymbolic reasoning.

The workshop showcased the huge potential and large synergies that we can unlock if we work together on these topics. While the TransAIR project concludes with the publication of this workshop documentation, please keep the workshop's inspiring and collaborative spirit alive. We tried to capture it on the following pages.

Prof. Michael Beetz
Initiator and project director of TransAIR and head of the Institute for Artificial Intelligence (IAI) at the University of Bremen.

PANEL DISCUSSIONS

Learning and Adaptation in the Minimalist Cognitive Architecture

One of the main ideas discussed during the TransAIR Workshop on Cognitive Architectures for Robot Agents is the concept of minimal cognitive architectures. What should or shouldn't be a component of a minimal cognitive architecture is a question that led to a lively debate on the virtual podium. Different approaches were favored by the various experts on the panel, which eventually led some to declare themselves maximalists as opposed to minimalists.

The unifying goal of all panelists: "create cognitive agents that are capable of understanding the effects of their own actions on the environment", as host Ana Tanevska phrased it. Trying to "explain human behavior computationally" is how John Laird described the motivation for his research, which led him to explore "how you build computational systems that have all that capability." Matthias Scheutz has been intrigued by a similar quest, creating a cognitive architecture that allows to "really control robots in real time." Ron Sun uses computational psychology to understand the human mental processes and structures.

HUMANS AS A MODEL

But to what extent should cognitive architectures for robotic agents mimic human cognition? In what regard do they need to differ significantly? These questions proved to be surprisingly controversial as they set the theme for the rest of the discussion.

"Many things, we cannot just map from cognitive science to technical systems", said Tamim Asfour. He added that it might not always be useful to copy biology: "In robotics we can maybe do things in a much better way than how it's done in biology." John Laird suggested to differentiate between a robot performing a specialized task and a more general artificial intelligence. "As we go to more general systems, then I would expect that there will be an overlap." For a system that performs tasks similar to how humans do them, Laird said it is an open question whether or not the robot's reasoning will also be similar to the way humans do it.

NO BIOLOGICAL CONSTRAINTS

Alternative architectures have one big disadvantage, however: there is nothing to model them after. "Human and animal intelligence is the only one we really know," said Yiannis Aloimonos. "We should really strive to design our robots according to principles that follow this kind of intelligence." The natural evolution of intelligence might have left many pathways unexplored, though, as Matthias Scheutz reminded the panelists and the audience: in nature, a cognitive architecture has to be fully functional and operational anytime – including during steps towards the next design iteration. "We don't have those constraints on our agents. And as a result, we might be able to have very different designs that would not be biologically viable, but much better at certain tasks."

This widens the field of possible architectures worth exploring. Something that John Laird agreed is necessary and not happening enough. Perhaps because "it takes a huge amount of work to develop each of these cognitive architectures."

IQ TESTS FOR ROBOTS

Nevertheless, new architectures will emerge without a doubt, while existing architectures will be further developed, raising the question how different architectures can be compared with one another. "We don't know what it means for one thing to be more intelligent than the other," said Yiannis Aloimonos, "and it would be an interesting question to do this in a technical sense, to establish some complex hierarchy for intelligence." Ron Sun suggested a system analogous to IQ tests. There was a broad consensus that it would be useful to give different robots the same task, a "complex, open-ended task," as Tamim Asfour noted, ideally expressed not in code but in natural language, and then not only evaluate their execution but also question them about their actions. A robot, said Matthias Scheutz, "should be able to talk about what it's doing, maybe while it's doing it give you an assessment of where it's doing, what it's doing next, maybe how likely it is that it will succeed."

It seems like for Sean Kugele, robots verbalizing their "thoughts" is too far out. He shared that he is concerned about "too much focus on human-level intelligence." Instead, Kugele favors to first "build systems that can solve the simpler tasks and then work on these higher-level cognitive abilities." Adding language capabilities to a system so that it can express its own state would raise the benchmark substantially beyond basic functionalities.

UNIQUE ARCHITECTURES

The future of minimal cognitive architectures for robotics should be shaped by one guiding question. "What is unique or what is different about manipulation tasks that we need to take into account?" This, hopes John Laird, would lead to architectures that are different from cognitive architectures that are used outside of robotics. Today, robots are often not good at analyzing given tasks in order to determine the necessary steps and capabilities needed. A lot of work has to be done on that, according to John Laird.

The trade-off between abstracting from the individual system and taking into account different forms of embodiments was briefly discussed, including the relationship between episodic memory and the body. "A cognitive architecture for a robot cannot be separated from sensory motor loops which actually describe interaction with the world," said Tamim Asfour, and Matthias Scheutz emphasized the importance of "the human level of timing and the human level of expectations," which have to be met by a robot that interacts with people.

THE MAXIMALIST APPROACH

Regarding minimality, "we need to consider the cost benefit trade-off, with a cost being the model complexity and benefit being the range of functionalities," said Ron Sun. But two panelists weren't sure if minimality is the right way to go in the first place. "My bias is that I like the idea of having a comprehensive theory," said Sean Kugele. "From my perspective, you want a theory that can potentially model everything and you may choose parts of that theory to apply in your work."

Minimalism and being comprehensive are not necessarily mutually exclusive, according to John Laird, who shared some practical experience: "I've never felt that adding more components had a bad effect on the system, when we don't use those components." In fact, he sees learning opportunities in observing how an added component "interacts with and maybe shapes the rest of the system." Laird advocated for incorporating more components as long as no sacrifices have to be made in terms of efficiency. "So, I'm a maximalist," he concluded to the amusement of the panel.

Yiannis Aloimonos is the Director of the Computer Vision Laboratory and Professor for Computer Science at the University of Maryland, USA.

Tamim Asfour is a professor of Humanoid Robotics at the Institute for Anthropomatics and Robotics, High Performance Humanoid Technologies at the Karlsruhe Institute of Technology, Germany.

Sean Kugele is a PhD candidate in the Department of Computer Science at the University of Memphis, USA.

John Laird is the John L. Tishman Professor of Engineering in the Computer Science and Engineering Division of the Electrical Engineering and Computer Science Department at the University of Michigan, USA.

Matthias Scheutz is a professor of Cognitive and Computer Science as well as director of the Human-Robot Interaction Laboratory at Tufts University, Boston, USA.

Ron Sun is a professor of Cognitive Sciences at Rensselaer Polytechnic Institute, NY, USA.

Ana Tanevska is a postdoctoral researcher at the Italian Institute of Technology, Genoa, Italy.

Understanding Human Cognition through Modelling the Human Brain

Panelists did not shy away from big questions when the relationship between the human brain and cognitive models was discussed during the workshop. A fruitful discussion around the role of the body, measures of intelligence and intentionality developed.

Host Maria Hedblom kicked the panel off by getting down to the foundations of the field: "Will it be possible to accurately model human cognition with computational means?" From an engineering perspective, Kazuhiko Kawamura didn't see an obvious way to measure the accuracy of a model or even define what accuracy means in this case. Tomaso Poggio hinted at insufficient hardware "The connectivity in present computers is still pretty small, at the level of transistor and gates there are typically no more than three or four wires coming in and out," he said, "whereas in cortical neurons you have around 10,000 or so synapses per neuron."

SIMULATING THE BRAIN

At the same time, there are unanswered questions about software as well. Even though it is possible to simulate the brain on a certain level, "we don't know which kind of algorithm the brain is using," Tomaso Poggio said.

Even if these problems could be solved someday, the body and the brain of a robot will not be exactly like a human's. "So will it be human cognition or will it be some other cognition?" Jeffrey Krichmar asked rhetorically. Agnieszka Wykowska agreed that embodiment is a crucial factor. But she called for a clear distinction depending on the goals behind the various efforts to model the human brain: "Do we want to model in order to understand how the human cognition works or do we want to model in order to do things the way humans do?"

Some may just desire to create a companion, others a simulation of ourselves. Creating a robot body and brain that is on par with humans is a giant feat of engineering. "It's an open frontier, and we like open frontiers," said Jeffrey Krichmar. "And along the way, we're going to learn a lot about human cognition." But, he said he sometimes cringes when the conversation revolves exclusively around human cognition, while animal cognition is sometimes neglected despite good data availability. "We should be tapping into that and just in general looking into cognition itself."

AN INTELLIGENCE TEST FOR MACHINES

Tomaso Poggio asked whether a machine that can do certain things better than a mathematician is intelligent. "I think there are many forms of intelligence," he said, thereby suggesting that defining intelligence remains a challenge in itself. "I stick, for now anyway, to Turing definitions," he declared. Agnieszka Wykowska, dismissing IQ tests as insufficient, agreed that „probably the closest we can get to having some sort of criterion of intelligence is indeed the Turing test." She suggested to build on it, partly to test whether an intelligent machine is able to "do things in the real-life environment and do things in an adaptive way", and also to test whether people ascribe intentionality to the machine. Following her approach, an intelligence test for machines would look at "whether people explain and predict behaviors of machines in intentional terms and under which conditions they do."

Kazuhiko Kawamura and Jeffrey Krichmar, on the other hand, were cautious. "There's a whole class of problems that cannot be solved with a Turing machine, and a lot of those things we would call intelligence," said Jeffrey Krichmar. The two of them also emphasized that intelligent robots would have to demonstrate intelligent behavior over an extended amount of time. "You have to consider a long time frame," Kazuhiko Kawamura said, "and that type of concept is missing in the current robotics research." Jeffrey Krichmar pointed out that extending the time frame is also needed for the learning phase of models. Instead of life-long learning, "we train them over a short period and freeze them, because that's the best we can do."

BEYOND DEEP LEARNING

There may be a need for an altogether new approach. "I don't think that Deep Learning is the answer," said Tomaso Poggio with regard to the key questions posed by the aim to connect neuroscience with computational models. "In neuroscience, I think there is a rush by many researchers to jump on the Deep Learning bandwagon. And I think it's probably going to be a waste of time." While he used to push neuroscientists to do more computational work and simulation, he now sees the technology being used too excessively and without a full understanding of it. Some researchers, he says, "are losing the concept of what a model, which is biologically interesting, is. They conveniently forget that this piece of code has to correspond to something in the brain."

And it doesn't stop there. From the brain to cognition to social cognition, interaction and embodiment, there are many aspects that are often neglected by models, according to Agnieszka Wykowska. "I think we're not getting any closer to bridging those gaps. And this is very, very dangerous for science altogether." She misses a dialogue between research communities that are focused on the various aspects and each have their own conferences and journals.

On a positive note, closing those gaps and creating integrated models seems more than just theoretically possible. "We have good neuron models. We have good synapse models and learning models that are plausible. We have tons of data on the anatomy," said Jeffrey Krichmar. "A lot of the pieces are there. The thing that blows me away is, when you put it together it's completely unstable, whereas we're, as biological organisms, operating over a wide range."

THINKING ABOUT A ROBOT UTOPIA

In her final question, Maria Hedblom invited the panel to imagine a distant future, in which all the big research questions that were previously discussed have found their answers. "In your absolute utopia, how are we interacting with robots and artificial intelligence?" Jeffrey Krichmar appreciated the outlook on robots with human cognition that we could interact with naturally and which could bring many societal benefits. "But I'm an engineer by training," he said, "so I'm making sure everything that I'm working around is safe for myself and anyone that's interacting with it." When machines make their own decisions which impact humans, that also involves new challenges. "At some stage, probably now, you have to start thinking about ethics," Jeffrey Krichmar said.

Tomaso Poggio thought it would be great if humans could dedicate themselves more to creativity and sports – "you know, a soccer game between humans rather than between robot teams." But he wondered, if people actually need jobs to lead a happy life, and what happens "when they know that what they do is not relevant to their survival."

Without leaning towards the utopian or a more cautious view, Kazuhiko Kawamura, who retired five years ago, said: "It's good to have a dream. That's how I started in robotics. Without a dream, I don't think you will become a good robot researcher."

Kazuhiko Kawamura is an emeritus research professor of Electrical Engineering, Computer Engineering, and Engineering Management at Vanderbilt University, Nashville, USA.

Jeffrey Krichmar is a professor in the Department of Cognitive Sciences and the Department of Computer Science at the University of California, Irvine, USA.

Tomaso Poggio is a professor at the Department of Brain and Cognitive Sciences and Director of the Center for Brains, Minds and Machines at MIT, Boston, USA.

Agnieszka Wykowska is the leader of the unit Social Cognition in Human-Robot Interaction and the senior researcher tenure track - principal investigator at the Italian Institute of Technology, Genoa, Italy.

Maria Hedblom is a postdoctoral researcher at the Institute for Artificial Intelligence at Bremen University, Germany.

Developmental Cognitive Robotics as a Gateway to a More Natural HRI

Human-robot interaction (HRI) is a metaphorical coin with two sides: "Allowing humans to understand their robot companions and allowing robots to understand humans," as Gayane Kazhoyan phrased it in her introduction to the panel that she hosted. The discussion revolved around mutual trust, transparent intentions and the challenges of developmental robotics.

Yiannis Demiris started with a personal anecdote: "I remember having some really bad attempts at learning how to dance." However, dance school seems to have changed his life. "I realized that we were learning by trying to imitate an instructor." This inspired his career in researching how robots can learn by imitation and from social interactions. More recently, he said, he has looked at the flip side: "Can robots help humans develop better?"

For both physical and social interactions, it would be beneficial for robots to have a model of their human counterparts. According to Alessandra Sciutti, this is true for autonomous cars as well as robotic vacuum cleaners: "Being aware, being able to predict or anticipate and being legible and understandable to the human will be of great benefit for establishing a natural way of living together with this kind of agents."

A NATURAL WAY TO INTERACT

Coming from the field of physics, Helge Ritter became fascinated by manual interactions. But he also realized that social interaction is important. Both are embedded in cognitive interaction. "Hands are an entry point to that," he explained his research focus. "We still say 'I grasped something', 'to come to grips with something', 'to touch on a subject.'"

Angelo Cangelosi, whose career started in experimental psychology, came to similar conclusions. For people and robots working together, a natural way of interaction is needed. "Natural interaction is to ask a robot to do something," Angelo Cangelosi said, "or to ask a robot to explain what it is doing and why." This could also help avoid or clear misunderstandings, which are common in natural interactions between people.

"A robot is a reflection of ours," said Minoru Asada. "We do some sort of mirroring, we just project ourselves on the robot." This leads people to expect certain human-like capabilities of the robot. Filling this gap is a problem that Minoru Asada is concerned about. It becomes more pressing as robots are being introduced to society. "Therefore, it's not just an issue of the technology, but also of ethics." He called for involving more ordinary people in robot development.

"A ROBOT IS LIKE A WASHING MACHINE"

Helge Ritter agreed that "society is a very important substrate for developing robots." As exciting as research on autonomous agents may be, Helge Ritter urged to make sure that autonomy needs to be shaped so that a robotic agent's interests don't deviate from the interests of society. "Can we envisage intrinsically safe architectures that make robots intrinsically well-behaved?" he asked, warning of an "entirely uncontrolled development of what can happen in a robot."

Apparently less concerned, Angelo Cangelosi answered: "A robot is like a washing machine." In his view, they are tools that are made to help people, while he also acknowledged challenges concerning autonomy and ethics. Alessandra Sciutti preferred the analogy of cars as we have them today: "You're keeping control, but you have to make all the decisions." While offering an advantage – in this case in travel speed – it adds to the driver's cognitive load. "On the other hand, if you want a collaborator, not everything is driven by you," she continued. As some of the decision-making responsibility gets transferred to another agent with a certain degree of autonomy, goals need to be negotiated, which requires shared perception and intentions, according to Alessandra Sciutti.

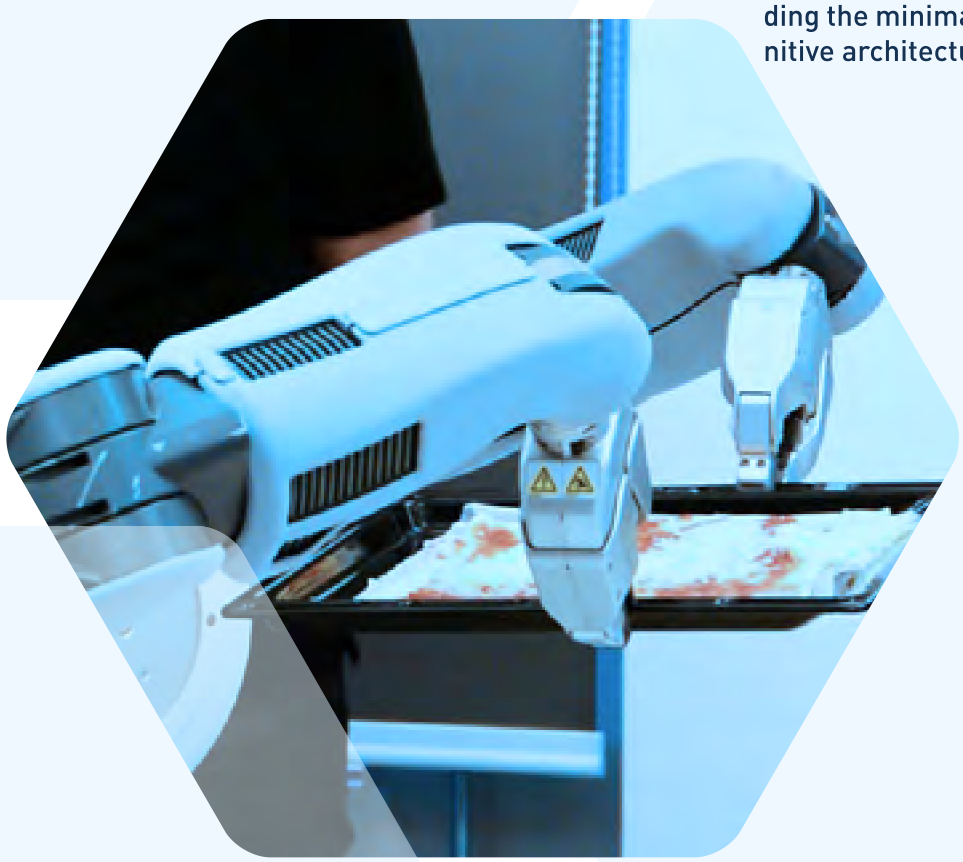

DEVELOPMENTAL ROBOTICS TRADE-OFFS

When it comes to the question of how to develop useful robots, the panel found common ground in developmental robotics – as the title suggested. "Developmental robotics will be able to generalize and be more flexible than a hardwired approach," Yiannis Demiris suggested. "But how quickly do you want your results?" he asked. Certainly, it would not make sense to educate a robot like a child for 20 years or more before it becomes useful. Yiannis Demiris sees a trade-off between the developmental approach and priming systems with knowledge so they can be useful from day one.

Minoru Asada shared his experiences with developmental robotics just looking at the early stages of learning: "In just one year a baby learns so many kinds of behaviors that we cannot design a robot who can obtain the same, because there are so many mysteries in baby development." Following his argument, these mysteries need to be solved first, preferring the type of HRI research that looks at understanding humans. He added that introducing robots into society can provide useful feedback to improve the robots, but will also change human behavior.

The "symbiotic society" that Minoru Asada predicts would mean that the agents' behavior would develop based on past interactions. "Moving away from one-shot interactions between the human and a specific robot is a great challenge," said Alessandra Sciutti. When people interact with each other, they both change, she said. To her, the key point is understanding the minimal elements required in a robot's cognitive architecture to facilitate this development."

DEVELOPING HUMAN-ROBOT RELATIONSHIPS

There was broad agreement on the idea that both the human and the robot should be learning from an interaction. But Yiannis Demiris wasn't sure how to represent this in a cognitive architecture. When it comes to modeling others and then acting on it, it is unclear how humans do it, Demiris said. "How do we understand others? Do we understand them in relationship with ourselves or are they a separate entity that demands their own model?" This, he said, is what he is most curious about.

Helge Ritter steered the discussion towards trust: "When an agent allows itself to be modeled by me and my prediction stays true, then this is a basis for building trust. On the other hand, if the agent tries to evade being modeled, kind of camouflage, then this is opposing trust." He gave the example of a human allowing a robot to hold their arm because there's trust on an emotional level.

Trust also depends on the body of the robot and its movements. Alessandra Sciutti described how the way a robot executes a movement is a means of communication. A fast movement could be perceived as aggressive or nothing to worry about, depending "on very peculiar regularities in terms of the acceleration profile." If a robot is aware of this, it could help build trust. Alessandra Sciutti shared another example, where a robot makes its movements more predictable and easier to anticipate by looking at a target point with its eyes before moving its arm.

The panel identified many open questions regarding the relationships between bodies and brains, emotions and knowledge, trust and privacy. More opportunities to find answers should be given to young researchers, said Yiannis Demiris. Alessandra Sciutti set the direction: "The time has come for the social component of cognition to have more importance in cognitive architectures."

Minoru Asada is a professor at the Department of Adaptive Machine Systems at the Graduate School of Engineering at Osaka University, Japan.

Angelo Cangelosi is a professor of Machine Learning and Robotics at the University of Manchester, UK.

Yiannis Demiris is a professor of Human-Centred Robotics at the Faculty of Engineering at the Imperial College London, UK.

Helge Ritter is the head of the Neuroinformatics Group at the Faculty of Technology and professor at the Department of Information Science at Bielefeld University, Germany.

Alessandra Sciutti is a tenure track researcher and head of the COgNiTive Architecture for Collaborative Technologies Unit at the Italian Institute of Technology, Genoa, Italy.

Gayane Kazhoyan is a PhD student at the Institute for Artificial Intelligence at the University of Bremen, Germany.

Minimalist Cognitive Architectures

“People tend to add a lot of components to cognitive architectures,” says Yiannis Aloimonos. He takes the opposite approach, focusing on the essential. Aloimonos proposes a programming language that takes advantage of the fact that actions follow a certain grammar.

The hope is for Action Language (AL) to become a universal language that allows execution of the same actions on different systems. Furthermore, it can serve as the basis for a cognitive architecture that acts more like a compiler or interpreter: given an input (a goal, a task, a problem to be solved), the architecture generates the program that will solve the problem.

Aloimonos envisions not only programmers to create AL programs. He works on implementing systems such as the Visual AL Compiler (VALC) that is intended to translate an observed action executed by a human into AL code. The Visual AL Debugger (VALD) will work the other way around: given an AL program it will guide users through the execution of a task using visual instructions and corrective feedback in an augmented reality environment.

Yiannis Aloimonos is the director of the Computer Vision Laboratory and professor for computer science at the University of Maryland, USA. Since the early 2000s, he has been working on the integration of sensorimotor information with the conceptual system, bridging the gap between signals and symbols. His research is supported by the European Union, the National Science Foundation and by the National Institutes of Health, USA.

Contact information:
https://www.umiacs.umd.edu/people/yiannis

Video presentation:
https://youtu.be/YMjghTvJF4M

Collaborating on Architectures: Challenges and Perspectives

When robots interact with humans, it requires machines to be flexible. As Helge Ritter points out, this has been a major challenge that takes experts from a wide range of disciplines to solve: computer science, robotics and mathematics meet linguistics, psychology, neurobiology and even sports science.

“If a robot comes into a new environment, usually the behavior is not very sophisticated,” states Helge Ritter. Part of the reason is a lack of flexible cognitive behavior as opposed to habitual behavior. While the latter can be highly automated and works rapidly with low effort, adding a layer of flexible behavior involves attention and high effort as well as reasoning in order to cope with novel situations.

Interdisciplinary work at Bielefeld University has led to a situation model framework that is supposed to enable flexible behavior by resembling human attention and learning mechanisms, among other components, in artificial agents. The work has been further advanced in a collaboration with the Everyday Activity Science and Engineering (EASE) project at the Bremen Collaborative Research Center.

The group now strives to “combine fast, episode-based learning with model-based and model-free learning to create an architecture that combines the strengths of explicit and implicit knowledge representation in order to reconcile explainability and performance,” as Helge Ritter said, while also raising awareness for new challenges, including how to square increased flexibility with specifications and safety.

Helge Ritter has headed the Neuroinformatics Group at the Faculty of Technology since 1990 and is a professor at the Department of Information Science at Bielefeld University, Germany, one of the directors of the Bielefeld Institute of Cognition and Robotics and coordinator of the excellence cluster “Cognitive Interaction Technology”. In 1999, he was awarded the SEL Alcatel Research Prize and in 2001 the Leibniz Prize of the German Research Foundation DFG.

Contact information:
https://ni.www.techfak.uni-bielefeld.de/people/helge/

Video presentation:
https://youtu.be/9qn8H4BheZY

Cognitive Architectures for Assistive Robot Agents

HAMMER is short for Hierarchical Attentive Multiple Models for Execution and Recognition. It is an architecture that uses a simulation theory of mind approach for perceiving and representing actions and intentions. Yiannis Demiris and his team have developed it at Imperial College London.

"How like me are they?" is a question that HAMMER tries to answer when observing a human. The robots in Demiris' lab learn a representation of their own bodies by recording sensory data as a result of random motor commands. They can then compare the movements of a user or a machine to their own abilities.

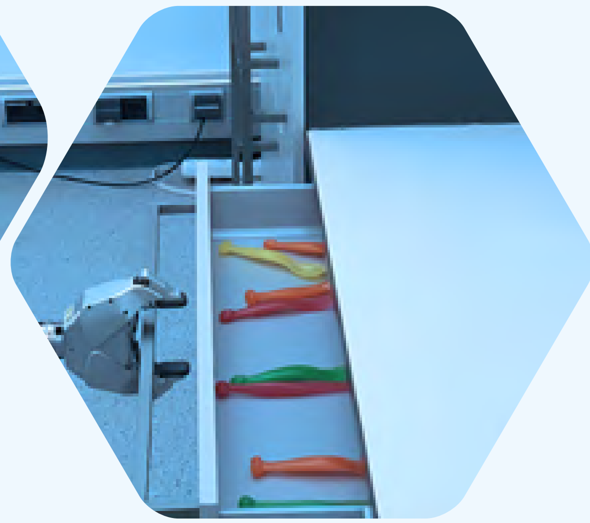

Developed with a focus on assistive technology, HAMMER can determine whether a user needs help. It understands external actions by simulating them internally, allowing for a principle approach into intention prediction, according to Demiris. The team has implemented HAMMER in diverse scenarios, ranging from assisted mobility to assisted dressing.

Yiannis Demiris is a professor in human-centred robotics at Imperial College London, UK, where he holds a Royal Academy of Engineering Chair in Emerging Technologies. He established the Personal Robotics Laboratory at Imperial in 2001. He is currently a Fellow of the Institute of Engineering and Technology (FIET), Fellow of the British Computer Society (FBCS) and Fellow of the Royal Statistical Society (FRSS).

Contact information:
https://www.imperial.ac.uk/people/y.demiris

Video presentation:
https://youtu.be/cqdX2lN5zBU

ArmarX – A Robot Cognitive Architecture

Cognitive robotic architectures, according to Tamim Asfour, need two key capabilities: combining experience-based learning and generative knowledge extension, and processing symbolic and subsymbolic information. AmarX follows this hybrid approach and has been developed and implemented over the past two decades. Its goal is to enable robots to learn from human observation and from experience, while they communicate and interact using natural language.

Recent additions to ArmarX, which also serves as a software development environment, include the ability to recognize a need for help and to provide help to humans proactively. "It's a very difficult task, also for us humans, to recognize that a partner is in need of help," Tamim Asfour said. "We can do that, of course, based on our understanding of the task." He demonstrated how a robot and a maintenance worker collaboratively completed several tasks.

Additional recent work has led to an episodic memory component, which Tamim Asfour says is crucial for interactions with humans. It allows the robot to make predictions, for instance.

ArmarX is available as open source.

Tamim Asfour is full Professor at the Institute for Anthropomatics and Robotics, where he holds the chair of Humanoid Robotics Systems and is head of the High Performance Humanoid Technologies Lab (H2T) at the Karlsruhe Institute of Technology (KIT), Germany.

Contact information:
https://h2t.anthropomatik.kit.edu/21_2372.php

Video presentation:
https://www.youtube.com/watch?v=Ew3yuOMeFpU

Affective Architecture: Pain, Empathy, and Ethics

"Robots can discriminate between touch and pain," says Minoru Asada, who believes that the affective aspect is important to support cognitive computing. His team has demonstrated this with a robust soft skin and a pain-sensitive nervous system embedded into a robot. What is "still a big mystery," according to Minoru Asada, is how the memory of pain works in humans. "Constructive approaches are necessary to reveal and to realize it in robots," he says.

In his view (and citing Ben Seymour), pain is a precise and objectifiable control signal that can be used for reinforcement learning. For the system to direct behavior away from harm, Minoru Asada suggests a mirror neuron system. "This enables a robot to recall its own motor experiences while observing others' actions as well as to produce the action," he explains.

Mirror neurons could equip robots with the ability to feel pain in others, which Minoru Asada views as the beginning of empathy, eventually leading to sympathy and compassion. Minoru Asada proposes the term of silicopathy. In combination with new ethics for a symbiotic society, he sees significant potential towards creating universal moral agents.

Minoru Asada is a professor at Osaka University, Japan, serving as president of the Robotics Society of Japan and as vice president of the Japanese Society of Baby Science. His research focus includes Robotics, Artificial Intelligence, and Cognitive Developmental Robotics. He is a co-founder of RoboCup.

Contact information:
http://www.er.ams.eng.osaka-u.ac.jp/asadalab/

Video presentation:
https://youtu.be/ns0SDIjj0Zk

Cognitive Robotics and Control

The development of cognitive robots at Vanderbilt University started in 1985 with the creation of the Center for Intelligent Systems – long before humanoid robots were readily available. Ever since, Kazuhiko Kawamura and his team have worked towards building robotic companions for humans, with an early focus on disabled people.

Many research questions remain open, as Professor Kawamura stated, consequently leading to a lack of key abilities in cognitive robots, including the ability to develop cognition through sensorimotor association. In order to make progress in this field, the team at Vanderbilt developed a multi-agent-based control architecture and implemented modules such as a working memory system.

The robot ISAC (Intelligent Soft Arm Control) can be trained to store a small number of chunks of information in its working memory, which then influences the selection of actions the robot takes. "This is an important skill for cognitive robots," Kazuhiko Kawamura concluded his talk.

Kazuhiko Kawamura is the emeritus research professor of electrical engineering, computer engineering, and engineering management at Vanderbilt University, Nashville, USA. From 1990 to 2013, he served as the director of the Center for Intelligent Systems, Nashville, USA. Dr. Kawamura is a life fellow of IEEE and has published over 150 research papers.

Contact information:
https://engineering.vanderbilt.edu/bio/kazuhiko-kawamura

Video presentation:
https://youtu.be/7i_l80w2mtg

Circuits for Intelligence

Some 20 years ago, in a collaboration with a German car manufacturer, Tomaso Poggio put computers in the trunk of a vehicle. The system was programmed to detect pedestrians. "It only made three mistakes per second," Tomaso Poggio jokingly says. Current systems are about one million times more accurate, with roughly one error per 100,000 kilometers. While Poggio acknowledges the great success of machine learning, which has become ubiquitous in everyday life, he points out its limited scope: "A program that plays superhuman chess would not notice a fire in a building."

There are still more breakthroughs yet to happen and Tomaso Poggio believes that they, like deep learning and reinforcement learning, will have their foundation in neuroscience. "We first need the natural science of intelligence, cognitive science, in order to get to the engineering of intelligence," he says. "Understanding how the brain makes the mind" is therefore part of the mission of a new Institute for the Science and Engineering of Intelligence, which he is working to establish.

Figuring out the "circuits underlying human-level intelligence" that equip humans with language and logic and the evolutionary steps behind them could lead to computational systems that go beyond what Tomaso Poggio calls "souped-up look-up tables."

Tomaso Poggio is a professor in the Department of Brain and Cognitive Sciences, a member of the Massachusetts Institute of Technology (MIT) Computer Science and Artificial Intelligence Laboratory (CSAIL) and director of both the Center for Biological and Computational Learning at MIT and the Center for Brains, Minds, and Machines headquartered at the McGovern Institute for Brain Research, Boston, USA.

Contact information:
https://cbmm.mit.edu/about/people/poggio

Video presentation:
https://youtu.be/SBw5P-dQe8Q

The LIDA Cognitive Architecture – An Introduction with Robotics Applications

LIDA stands for Learning Intelligent Decision Agent. It models complete cognitive systems. "This aspect is critical for robotics, as it facilitates the creation of autonomous agential software systems," says Sean Kugele. LIDA is inspired by biology and implements several psychological and neuropsychological theories. "It is based on and partially constrained by our knowledge about natural agential systems."

Sean Kugele emphasizes that LIDA does not model brains, but minds. "We define minds as control structures for autonomous agents," he explains. Simply put, this control structure is the mechanism by which an autonomous agent answers the question: What do I do next?

Hence, the LIDA team focuses on explaining "how minds support the selection and execution of action." The model's cognitive cycle is split into three phases: perception and understanding, attention, and learning. Looking at the architecture in more detail, it stands out that LIDA contains a number of different memory models. Sean Kugele argues that "different kinds of knowledge structures – for example perceptual, procedural, episodic, spatial, semantic – have different representational formats and are supported by distinct cognitive processes, for example learning processes."

Sean Kugele is a PhD candidate in the department of Computer Science at the University of Memphis, USA. A computer scientist and professional software developer turned cognitive scientist, he has worked with Stan Franklin on the LIDA cognitive architecture since 2012. Sean was a Technical Principal at FedEx and has also worked for Northrop Grumman as a software engineer.

Contact information:
http://ccrg.cs.memphis.edu/people.html

Video presentation:
https://youtu.be/U7ofPzjMeqE

Clarion: A comprehensive, Integrative Cognitive Architecture

"To me, a cognitive architecture is a broadly-scoped, domain-generic, computational psychological model, capturing the central structures, mechanisms and the processes of the mind," says Ron Sun. "It represents psychological theories in a computational form, which might otherwise be difficult to capture."

With Clarion, Ron Sun suggests a comprehensive cognitive architecture that includes several subsystems to cover essential psychological processes: an action-centered and a non-action-centered subsystem, a motivational subsystem and a metacognitive subsystem. "Together, they address action, skill learning, memory, concept, reasoning, motivation, metacognition, personality, emotion and so on," says Ron Sun, "and more importantly, they address the combination and interaction of these things."

Clarion follows a dual-process theory of mind by distinguishing between explicit and implicit knowledge and processes. "The interaction among these different types of processes is very important to understanding the human mind," he says.

He invites researchers to apply his cognitive architecture to robots, which has not been done, as his research focuses on designing Clarion to closely resemble human psychology.

Ron Sun is a professor of cognitive sciences at Rensselaer Polytechnic Institute, NY, USA. His research interests center around the study of cognition. He has published award-winning research papers as well as ten books, including „Anatomy of the Mind" and „Cambridge Handbook of Computational Psychology". He is a fellow of IEEE, APS, and other societies.

Contact information:
https://faculty.rpi.edu/ron-sun

Video presentation:
https://youtu.be/HLFijuMhJWQ

The DIARC Architecture for Autonomous Interactive Robots

"While classical cognitive architectures have been used on robots," says Matthias Scheutz, "they were not originally designed for the control of embodied agents." Thus, they lack a number of capabilities that Scheutz thinks are required for the autonomous long-term operation of robots. DIARC is meant to fill these gaps. "The Distributed Integrated Affect Cognition Reflection Architecture was designed in the early 2000s to work on embodied agents in real time, interactive settings in a fault-tolerant manner," he explains.

What makes DIARC unique is a deep integration of natural language understanding, access control and ethical reasoning, as well as component-sharing across different agents, among other features. Matthias Scheutz demonstrated some of the architecture's capabilities during his talk. The robots explained why they couldn't execute certain tasks, for instance because of a lack of trust in one operator versus another or because the task appeared dangerous.

Regarding ongoing work on DIARC, Matthias Scheutz says that it focuses on improving task-based dialogues, more monitoring for better resilience, and human-machine teaming.

Matthias Scheutz is a professor of cognitive and computer science as well as director of the Human-Robot Interaction Laboratory at Tufts University, Boston, USA. He has published more than 400 peer-reviewed papers in artificial intelligence, cognitive science and related fields. His current research focuses on complex ethical cognitive robots with natural language interaction and instruction-based learning capabilities in open worlds.

Contact information:
https://engineering.tufts.edu/people/faculty/matthias-scheutz

Video presentation:
https://youtu.be/RTol2DnIet0

The Soar Cognitive Architecture: Current and Future Capabilities

Work on the Soar cognitive architecture started in 1981 with John Laird being one of the original developers. He says the team was inspired by both psychology and computer science. "We focused on complex behavior and tasks, and often longer time scales, more than just a few seconds." Soar has seen a large variety of implementations and applications and is currently used in autonomous driving.

John Laird highlights Soar's interactive task learning capability, which allows the system to learn a new task from a single interaction: "This is real-time, online, one-shot learning." In their current research, the team is trying to improve reasoning about experience, perceptual reasoning, and motor reasoning. It's also working to equip Soar with more knowledge from the start. "Maybe we can pre-program this basic common sense, core knowledge about space, objects, agents and time," says John Laird.

However, with innate knowledge created from various sources, including existing knowledge bases and pre-trained neural networks, he sees a new challenge arise: "How do you get coherent meaning and reasoning across these different sources?"

John E. Laird is the John L. Tishman Professor of Engineering in the Computer Science and Engineering Division of the Electrical Engineering and Computer Science Department of the College of Engineering at the University of Michigan, USA. He is a fellow of AAAI, AAAS, ACM, and the Cognitive Science Society. With Paul Rosenbloom, he is the winner of the 2018 Herbert A. Simon Prize for Advances in Cognitive Systems.

Contact information:
https://laird.engin.umich.edu

Video presentation:
https://youtu.be/BUiWk-DqLaA

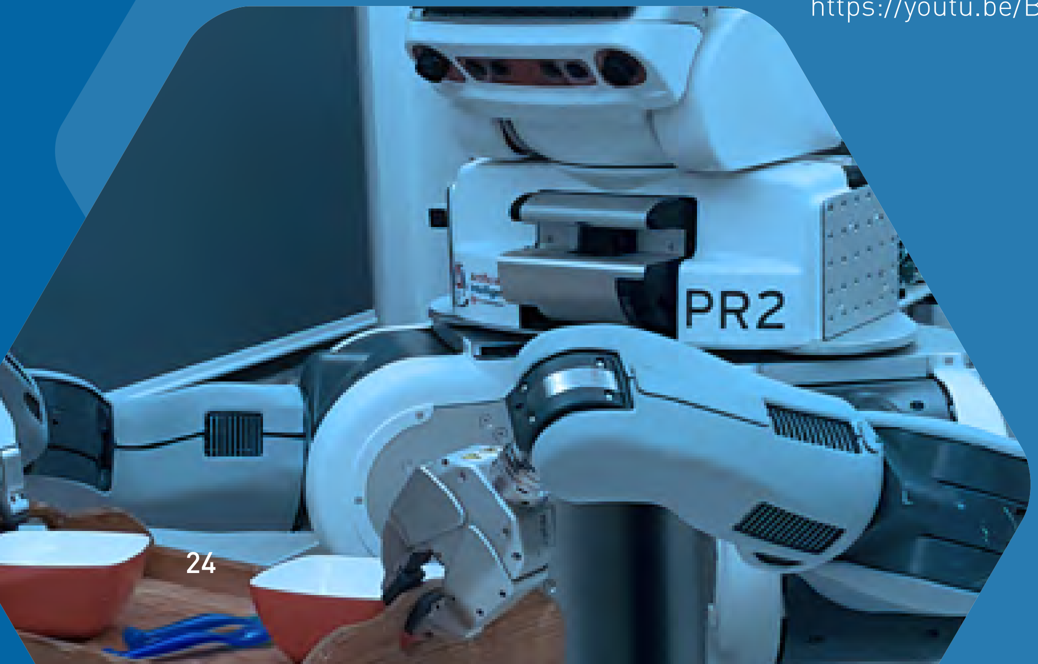

Mechanisms of Human Cognition in Interaction

Agnieszka Wykowska and her team use robots to study cognition and cognitive mechanisms in humans. The humanoids serve as a proxy for real-life scenarios, meaning that they help make study setups less abstract and artificial, while offering more experimental control than a human facilitator. "We can ask the robot to repeat the same movement over many, many, many trials," says Agnieszka Wykowska, "and also we can manipulate specific parameters of robot behavior in order to see what impact it has on human cognition."

Mechanisms the group explores using robots include attention and theory of mind or intentional stance. In one experiment they studied the effects on participants' attention, depending on whether or not a robot made eye-contact with them. "Attentional orienting can be modulated by social signals, such as a mutual gaze. Therefore traditional models of attention might need to be complemented by social components," says Agnieszka Wykowska.

Her group also studied whether humans predict and explain robot behaviors with reference to the assumed mental states of the machine. Some were more likely to adopt intentional stances, others had a more mechanistic point of view. Interestingly, the researchers were able to predict from EEG activity what stance study participants would take towards the robot.

Agnieszka Wykowska is a senior Researcher tenure track at the Italian Institute of Technology, where she leads the unit Social Cognition in Human-Robot Interaction, Genoa, Italy. She is the editor-in-chief of the International Journal of Social Robotics and serves as a board member and president-elect of the European Society for Cognitive and Affective Neuroscience (ESCAN).

Contact information:
https://www.iit.it/it/people-details/-/people/Agnieszka-Wykowska

Video presentation:
https://youtu.be/rewJUzs2qqw

Neurorobotics: Connecting the Brain, Body and Environment

"Neurorobotics is a holistic approach," says Jeffrey Krichmar. "It combines the brain, body and behavior." Not only does it allow for testing theories of neuroscience in ways that are impossible to pursue in a wet lab, it also enables testing outside of the lab, in real-world situations. Jeffrey Krichmar sees big potential in neurorobotics: "It may be a means to develop autonomous systems with some level of biological intelligence that is not shown by artificial intelligence now."

In his talk, he briefly explained a set of neurorobotic design principles ranging from actions and reactions to adaptive behavior to behavioral tradeoffs. He criticizes that it is common in neuroscience to separate sensory and motor domains, and within the sensory system to separate vision from audition from touch. "The brain does not have these concrete lines, they're blurred," he says.

Following these principles, like sensory-motor integration, could be a pathway toward an artificial general intelligence (AGI), according to Jeffrey Krichmar. However, he identified a number of needs for creating AGI, some near-term, some long-term, including a need for interdisciplinary talent, who must want to conduct field work: "Like a real biologist, testing your robot in the wild."

Jeffrey Krichmar is a professor in the Department of Cognitive Sciences and the Department of Computer Science at the University of California, Irvine, USA. He has over 100 publications, holds seven patents and is a senior member of IEEE and the Society for Neuroscience.

Contact information:
https://cnlm.uci.edu/krichmar

Video presentation:
https://youtu.be/rb2OQH7ghW8

Developmental Robotics – Language Learning, Trust and Theory of Mind

AI assistants as they are common in many households are pre-programmed with a big vocabulary, but lack a full understanding of language. Children, on the other hand, start with just a few words and are slow language learners, but very efficient. "Maybe pre-programming a robot with the full knowledge of the whole English dictionary and grammar might not be a good idea," says Angelo Cangelosi. He also points out that children use their body for situated interaction and learning, for instance counting and calculating with the help of their fingers. Furthermore, they develop a Theory of Mind for social interaction.

In collaboration with child psychologists, Angelo Cangelosi and his team were able to implement child-like learning in robots as well as use robots to predict the behavior of children in later studies.

"Embodiment is very important in many aspects of development, not only language, but also in social cognition", Angelo Cangelosi concludes.

Angelo Cangelosi is a professor of machine learning and robotics at the University of Manchester, USA. He was the founding director at the Centre for Robotics and Neural Systems at Plymouth University, UK. Cangelosi has produced more than 250 scientific publications and is the editor of multiple journals. His latest book is titled "Developmental Robotics: From Babies to Robots".

Contact information:
https://www.research.manchester.ac.uk/portal/angelo.cangelosi.html

Video presentation:
https://youtu.be/jQHQg7VnAaU

A Social Perspective on Cognitive Architectures

„Robots lack the natural ability of establishing mutual understanding with us," says Alessandra Sciutti. Humans, on the other hand, even at a very young age, "are very good at interacting and understanding each other." The goal of Alessandra Sciutti's research is to port some of these human abilities onto robots, for instance the capacity of one-year olds to understand if someone needs help in achieving a goal.

The group's work includes a robot understanding from visual cues whether a human is handling an object with care or not, and a robot expressing attitude by executing motions in a certain way. "There is a lot of information that a properly designed robot motion can convey," says Alessandra Sciutti, "and it can also evoke important changes in the way that a human partner behaves in response." She suggests that even basic learning processes might benefit from the inclusion of social components.

Integrating research on perceptual and motor skills with work on cognitive processes is, according to Alessandra Sciutti, "necessary, if we want to build robots that are more considerate of the human."

Alessandra Sciutti is a tenure track researcher and head of the COgNiTive Architecture for Collaborative Technologies Unit at the Italian Institute of Technology, Genoa, Italy. In 2018, she was awarded the ERC Starting Grant wHiSPER (www.whisperproject.eu), focused on the investigation of joint perception between humans and robots.

Contact information:
https://www.iit.it/it/people-details/-/people/Alessandra-Sciutti

Video presentation:
https://youtu.be/LCkOjR_cvxI

AI AND ROBOTICS IN BREMEN

Bremen's appeal as an AI hub is based on its extensive research network, which is embedded in an area with a strong manufacturing and trade tradition.

The state of Bremen, located on the river Weser near the North Sea, has long been the main industrial and trade center of northwestern Germany. Among the largest employers are Daimler (Mercedes), which builds electric cars in its local plant, and Airbus. Beginning with the late 20th century, Bremen also developed excellent strengths as a city of science and research. More than 50 technology research institutions are based here – they represent all major German research powerhouses. About 37,000 students are enrolled in eight universities and colleges.

At the University of Bremen, the Institute of Artificial Intelligence, the Robotics Group and the Collaborative Research Center EASE lead the way in AI and robotics. Bremen's other large players include DFKI's Robotics Innovation Center, the logistics institute BIBA, Fraunhofer Mevis, Jacobs University, and major IT companies such as Neusta and HMMH. They are joined by a growing list of promising start-ups. Researchers and private companies have started Bremen.AI, a community focusing on strengthening the region's AI ecosystem.

Current research topics these institutions and companies are working on include:

- Autonomous driving on earth and the moon (AO-Car, CC AD)
- Learning household robots (EASE)
- Smart technology in logistics (BIBA)
- Smart technology in retail (Knowledge4Retail)
- Humanoid robot design (Robot AILA)
- Robots that play soccer (six-time RoboCup world champions in Standard Platform League)
- Study of human emotions (Emote, CyberEmotions)
- Smart government (chatbots in the Bremen Citizen Service)

SUPPORT

AN INITIATIVE OF THE

Federal Ministry
of Education
and Research

IMPRINT

TransAIR
https://transair-bridge.org

Institute for Artificial Intelligence
University Bremen

Am Fallturm 1
28359 Bremen Germany

Tel.: +49 421 218 64000/64038
Fax: +49 421 218 64047
E-Mail: contact[AT]transair-bridge.org

Text
Thomas Reintjes, Axel Kölling

Design
eventfive GmbH

TransAIR

Artificial Intelligence

University of Bremen