Artificial Intelligence Research Flanders

challenge-based research with **demand-driven** impact

CHALLENGE BASED RESEARCH

RESEARCH CHALLENGES

PROOF-OF-CONCEPTS (POCS)

WITH DEMAND-DRIVEN IMPACT

FOCUS DOMAINS and USE CASES

May 2020

Flanders AI Impulse Program

The Flemish Ministry of Economy, Science and Innovation has made 32 million euros per year available to get Flanders to the head of the pack on AI. The AI-impulse program will be centered around three main pillars:

1. Strategic Basic Research (12 million euros)
2. Implementation: Technology Transfer and Industrial Applications (15 million euros)
3. Supporting activities (awareness, training, ethics…) (5 million euros)

Flanders AI Research Program

Challenge-based research with demand-driven impact

The program focuses on domains of strategic importance in Flanders that have a large market potential for which Flanders can be uniquely positioned as early adopters.
For the coming years, the program will **focus** on **use cases** within **three application domains**: Healthcare, Industry 4.0, and Government & Citizens.

Review of the strategic needs of Flanders in the coming years, might lead to the addition of or shift to new domains, using AI for the challenges in society such as climate, mobility, the aging of people.

The analysis of the needs in the focus domains has led to defining the **four research themes or research challenges** the program will focus on.

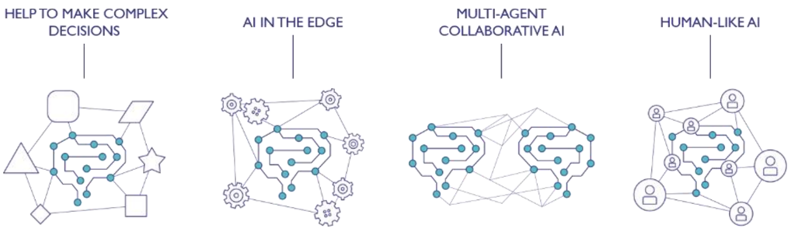

For each of the above domains and research challenges, concrete **proof-of-concepts** will be developed as demonstrators of the match of the research results with the needs from the application domains.

Validation of this challenge-based research with demand-driven impact will be done by the Flemish industry and an International Advisory Board. While the focus in the coming years is on the above-mentioned domains, the results will be more widely applicable in other domains as well. It is the aim that the demonstrators or proof-of-concepts serve as the entry point for the Technology Transfer and Industrial Application pillar of the AI impulse program, in which companies can tailor the results of the research to towards their specific needs.

The research program builds on existing strengths in Flanders: the Flemish government wants to further develop the current centers demonstrating excellent AI research. The **partners** in the AI research program are 5 Flemish universities (KU Leuven, UGent, VUBrussel, UAntwerpen, UHasselt), the four strategic research centers (imec, VIB, VITO, Flanders Make) and Sirris. Agoria and VOKA support the program and facilitate the interaction with the Flemish Industry.

The 4 research challenges

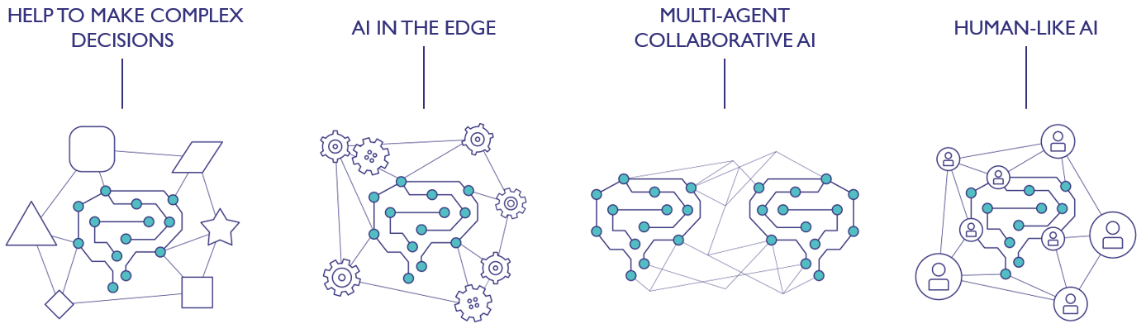

Data science: hybrid, automated, trusted, actionable

HELP TO MAKE COMPLEX DECISIONS

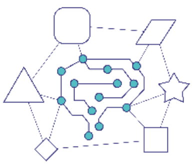

Decision makers for industrial processes and societal systems face an ever more daunting task. Every choice they make needs to be based on:

- knowledge and knowhow from experts such as doctors, engineers or market analysts;
- vast amounts of unstructured and structured data;
- numerous rules, guidelines and regulations on safety, ethics and privacy.

Luckily, future **decision support systems** will come to the rescue. To maximize their impact on the Flemish economy, we must make sure that they are:

- **automated** – By automating aspects of the data science process – such as raw data processing – we unlock its potential to all stakeholders, regardless of their technical data science skill level.
- **hybrid** – We need to unify the power of generated data with domain and expert knowledge. For example, by combining medical science with data from patient records, personal health monitoring sensors and clinical test targets.
- **actionable** – We have to turn data and knowledge into models that readily provide insights and inspire reliable decisions. These models must also give feedback to human experts, e.g. with interactive visual interfaces.
- **trustworthy** – All this has to be done with regard to the human in the loop and with respect for the data subjects' privacy and right to fair treatment.

Real-time and power-efficient AI in the edge

AI IN THE EDGE

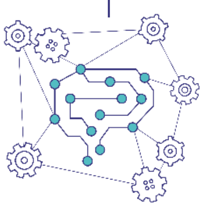

Smartphones, drones, robots on the manufacturing floor, electric vehicles, ... Devices at the edge come with ever more performing and power-efficient AI processors. That enables them to take on **advanced edge computing and distributed machine learning** tasks, driven by three factors:

- increased real-time performance;
- enhanced power-efficiency;
- greater need for data security.

It gives rise to an entirely new set of AI use cases based on **intelligent, low-power (often battery- powered) devices**, as well as cases requiring on-the-spot, **real-time and secure decision support**.

The challenge is to create:

- distributed and hierarchical AI systems;
- advanced signal processing;
- algorithms and technologies for extracting actionable information directly at the edge.

This move to the edge is **technically feasible and very relevant** for many use cases. Edge inference is forecasted to occupy about one third of the total market in 2023.

Interact autonomously with other decision-making entities

MULTI-AGENT COLLABORATIVE AI

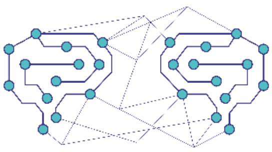

Autonomous decision-making entities each have their own goals and intentions. In **multi-agent systems**, they need to interact with each other. Multi-agent systems are radically different from distributed systems. In multi-agent systems:

- No agent knows the whole system.
- No agent directly controls all the other agents.

Multi-agent systems can be anywhere on the spectrum between cooperative and competitive. And you'll find them in the real as well as the virtual world. Examples in the world of information are trading systems, routing systems and privacy-sensitive systems – where agents can't share certain information with each other. A lot of cyber-physical systems are also multi-agent. Think about smart power systems, traffic and fleet control systems and autonomous vehicles. All this poses a unique set of challenges.

Multi-agent systems need to:

- **adapt rapidly** to unpredictably changing environments;
- **adhere to constraints, rules and regulations**, even in the absence of central control;
- **be accountable and manageable** by their creators;
- **interact with humans**, by understanding their intentions and explaining their own behavior;
- **be open-ended**, so new agents, users and technologies can join at any time.

Communicate and collaborate seamlessly with humans

HUMAN-LIKE AI

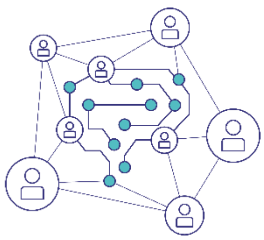

Can an AI system really equal human performance when it comes to complex tasks? Or have we merely created good pattern matching techniques up to now? Many industrial applications need to go beyond such pattern matching.

They have to be **capable of complex reasoning** in a way that is autonomous, intelligent and trustworthy. This requires them to:

- communicate in ways that are effortless to humans, such as natural language;
- perform multi-step, human-like reasoning that entails perception and understanding of a complex environment.

If we achieve this goal, we're able to enrich our society and workplaces with artificial entities that can identify and solve problems, **take on unseen tasks with the same agility as humans** – all while interpreting their social and physical environment and involving, informing and supporting their human colleagues.

Will we ever be able to equip technology with **real human intelligence**? Despite recent AI advancements, that goal is still far in the future.

We need systems that can **integrate and interpret, represent and understand their complex environment** in multiple styles and domains, over large timescales and in shared human-machine contexts. Therefore, we can identify two main objectives. Machines need to:

- seamlessly understand humans and interact with them;
- mirror the human capacities for learning, adapting, complex reasoning and decision-making across tasks, contexts & time.

Artificial Intelligence for efficient healthcare

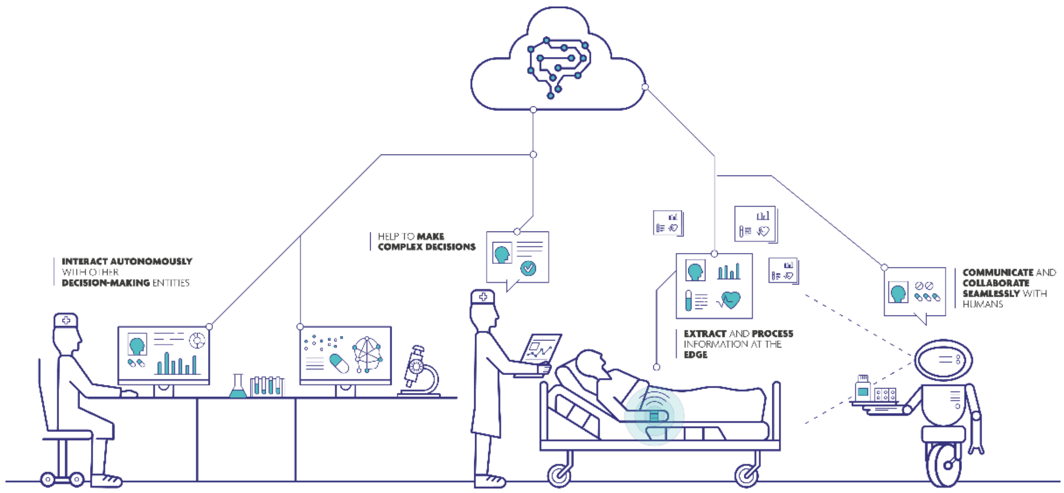

Help to make complex decisions through data science

Do we want to deliver on the promise of **early disease detection, effective prevention and personalized therapies**? Then health professionals need access to decision support systems that allow them to exploit multi-level and real-time **health data** such as environment and lifestyle, and medical and molecular information. These systems must be:

- **hybrid** – The data of evidence-based medicine are combined with systems medicine and expert knowledge.
- **automated** – Raw data are automatically processed and unified with other knowledge sources. That allows non-experts to tackle more complex problems. With an exploratory approach, for example in genome analysis, they can detect unexpected new patterns.
- **actionable** – With predictive methods, medical professionals can improve preventive and therapeutic actions.
- **trustworthy** – Privacy and reliability of data are important concerns in a medical context. Medical staff, caregivers and patients need correct and transparent information that helps them to take the right actions and conscious decisions.

Soon, we will encounter these decision support systems in several medical subdomains: clinical treatment, drug discovery, patient monitoring, care automation, trend analysis, drug support, ...

AI in the edge

There's a promising future for intelligent, low-powered AI devices in the medical sector. Particularly in two domains:

- **life sciences** – In field operations such as emergency care, there is often a need for **quick, accurate and cost-effective analysis** of substances such as blood. Imagine a typical requirement of being able to analyze 10 million DNA fragments within 60 minutes at a cost below € 300. This is only possible with highly intelligent, battery-powered cell-sorting and sequencing devices.
- **connected health** – Today, 24/7 patient monitoring is only possible with wired sensors within a hospital setting. If we want to switch to comfortable and low-cost **ambulant patient monitoring**, we need battery-powered, medical-grade and wireless sensors. These sensors need to operate at least 24 hours on one battery charge. And within a wearable and implantable health sensor network, they have to make collaborative health predictions such as risk at heart failure or septic shock.

Interact autonomously with other decision-making entities

In the pharmaceutical and medical sectors, data is often owned, explored and valorized by different entities. Huge strides forward could be achieved if these entities are **willing to share their data and models**, or to build new models cooperatively. Because of the sensitive nature of the (often patient-related) data, this must be implemented through **privacy-preserving cross-partner machine learning** that rules out unintentional leaks.

This will lead to a novel data economy model where contextual integrity relies on:

- **Compliance:** users of data need to follow rules and regulations based on beliefs and expectations.
- **Transparency:** in order to establish trust, citizens have to be able to see what happens to their data in the real world.
- **Contestability:** design and change rules must be based on user beliefs and expectations.

Communicate and collaborate seamlessly with humans

The healthcare sector is one of the sectors with a plethora of bottleneck jobs. Existing staff is often overburdened and recruiting skilled people is a huge challenge. Could AI **ease the workload of caregivers** by assisting them in social healthcare? Imagine AI systems that enable comfortable and independent assistant living, preventive care and rehabilitation. They do this by naturally interacting with their users and providing answers to complex problems. To make that happen, we need to focus on these **key research questions**:

- Can AI systems have **meaningful conversations** with patients? This should go beyond simple request-response interactions. AI systems need to interpret the human intention and behavior in open environments – including (non-)verbal aspects and situational contexts. If they succeed, AI-driven cognitive feedback loops can also prove useful in therapist trainings.
- Can AI systems **match therapists** in their interaction with patients? If they're confronted with a complex question, they should be able to provide a truly useful answer – thanks to their understanding of the user's explicit and implicit needs and intents, and of the environment.

Artificial Intelligence for Industry 4.0

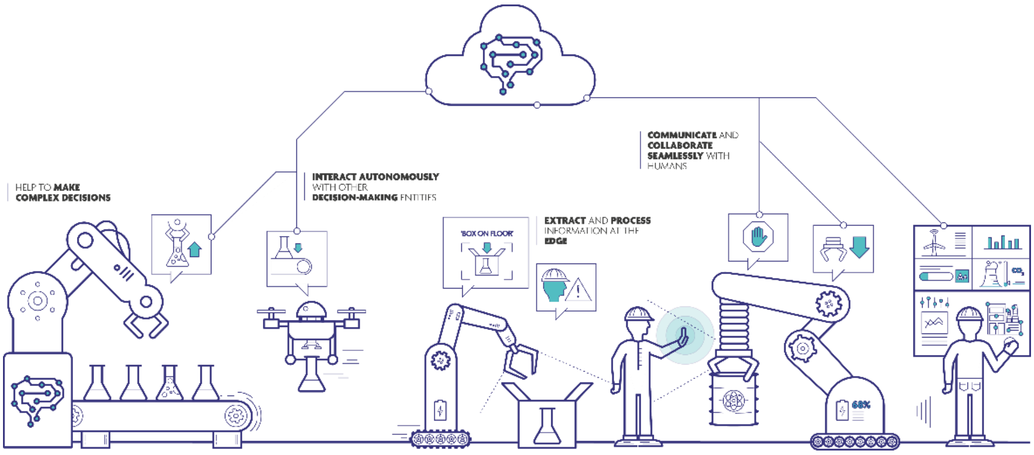

Help to make complex decisions through data science

Today, our companies invest a lot of time in the design of complex physical systems. In **Industry 4.0**, active learning and AI-assisted data exploration will speed up those processes and make it cheaper to do simulations or experiments.

Operators, designers and plant managers can count on **hybrid analytics** to facilitate manufacturing processes such as real- time asset monitoring, process control and reliability management – often in the form of a **digital twin** of the factory. It's important that these actors have access to an overall AI-based system that they can trust and understand, and that helps them through automation and interactive guidance.

AI in the edge

Our vision of the **factory of the future** is a floor filled with robots that are able to perform complex tasks like putting together an unfamiliar item on a conveyor belt in real time. Or a flexible assembly cell that consists of one or two mobots with 3D vision and sensing capabilities that drive around to pick up components and perform assembly operations – assisted by additional visual and cognitive sensors on the fixed infrastructure.

Science fiction? Not at all. But such **closed-loop robotic control systems** do require edge technologies such as high-resolution sensors (cameras, cognitive radars, LIDAR systems, ...) with proven low-latency processing.

Edge intelligence is also essential for **distributed predictive maintenance, condition monitoring and health awareness**. If high-frequency raw sensor data from a large number of machines and vehicles are processed locally, it reduces the need to communicate all the data to a central platform. That enables machine learning across a complex distributed network, where precise context-specific models are trained at lower levels and integrated and generalized at higher levels.

Finally, real-time and power-efficient edge technology is essential when machines and people work together. For instance, machines should be able to **continually measure and interpret the health and well-being** of operators, in order to optimize the operation speed and maintain general safety.

Interact autonomously with other decision-making entities

A **fleet** consists of interacting agents that are very similar but not identical – because they operate in other environments, are differently affected by wear and tear, etcetera. Such a fleet can consist of wireless networking nodes, hybrid vehicles, industrial machines, windmills, ...

AI systems can assist the management of such fleets by providing predictive maintenance. And if we go one step further, it should be possible to engineer a fleet to act as a body of autonomous decision-making entities that **jointly learn and optimize their performance**.

In some situations, various parties have their own goals but could benefit from cooperation. In such **a horizontal system**, real-time communication and collaboration lead to considerable optimizations.

Finally, in an **intelligent internet of intelligent things**, devices are able to take decisions on behalf of their users. In order for them to do this, they need to communicate with them, understand their goals and create a relationship of trust.

Communicate and collaborate seamlessly with humans

Will the factory of the future be a factory without humans? Not at all. On the contrary: AI will help **machines and humans to work together** in an industrial process that realizes the possibility of mass customization and personalization. We will encounter human-like AI in different forms:

- **operator support in the manufacturing industry** – Operators will work naturally and in real time with AI systems, thanks to speech recognition, sensors and visual aid through augmented or virtual reality. Of course, the AI systems have to be on a par with the humans in comprehension, interaction and complex reasoning. They need to know when to intervene, based on the operator's behavior, intent and cognitive load. And they have to offer humans complex and creative solutions by understanding problems and reflecting on them.
- **seamless human-machine interaction –** Cobots and humans will work together in dynamic environments. Within such mixed assembly cells, cobots need to understand human behavior, language and (implicit) intentions, flexibly interact and learn new skills. Also, they have to react fast to unforeseen human behavior or changes in the environment, to avoid obstructions and collisions.
- **complex control systems with minimal human input** – Humans should wholeheartedly entrust complex tasks to AI-enabled machinery. Ideally, they demonstrate or explain the task to the machine, which then fulfills it immediately and independently. Because of the trend towards mass customization, it's important that industrial machines are capable of quickly adapting to new tasks and even lifelong learning. More and more, the AI systems will be in the lead of operations, and humans will assist when needed.

Selected Use Cases

3 domains	Applications		Use Cases
Health	Precision Medicine	Single Cell Technologies	Visualisation of single Cell Data and Segmentation of 3D electron microscopy images
		Multiple Sclerosis	Improving MS treatments
	Clinical Decision Support	Epilepsy	Automated detection of epileptic seizures
		Medical Imaging	Segmentation and Classification for Radiation Oncology and Radiology
	Hospital Decision Support	Hospital treatment decisions	Prediction of Length of Stay in Hospitals
Industry 4.0	Smart Machines	Smart Monitoring	Smart Monitoring: Prognostic Health Management
			Smart Monitoring: Power-efficient anomaly detection on high-fequency sensor data
		Smart Control	Smart Control: Closed-loop robotic control
			Smart Control of Machine Fleet
			Smart Control: Optimisation of Manufacturing Processes
	Smart Factories		Smart Control of cooperating robots/arms
			Smart Control: Optimisation of Flexible Production Lines
		Smart Industrial Spaces	Smart Spaces: AI assisted operator
			Smart Spaces: People Detection and Tracking
			Smart Spaces: Embodied agent (industrial AGV) performing tasks assisted by operator
	Smart Energy	Smart Energy Distribution	Smart Energy distribution
			Self-sustainable micro-grids
Government&Citizens	Public employment services	Public employment services	Employment services for Citizens and for Employers and public advice
	Medical & Personal Data	Personal Data	Exchanging Data and knowledge in a healthcare setting
Smart Spaces			Conversational Agents
			Recommender Systems