Artificial Intelligence Research Flanders

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

Prof. Dr. Bart DE MOOR

ESAT-STADIUS KU LEUVEN
bart.demoor@kuleuven.be

AI PROGRAM FLANDERS

PROGRAM STRUCTURE AS APPROVED BY THE FLEMISH GOVERNMENT

12 MILJOEN EURO	15 MILJOEN EURO	5 MILJOEN EURO
↓	↓	↓
Top Research	Digitization and implementation in industry	Ethics, education and training

FLANDERS AI RESEARCH

FLANDERS AI RESEARCH PROGRAM
WHERE TO MAKE THE DIFFERENCE ?

1. Applied AI – Complex Decision Making
   - *Big Data* is not always *Good Data*
   - adding *Domain Knowledge* is crucial

imec

FLANDERS AI RESEARCH PROGRAM
WHERE TO MAKE THE DIFFERENCE ?

1. Applied AI – Complex Decision Making
   - *Big Data* is not always *Good Data*
   - adding *Domain Knowledge* is crucial
2. AI at the Edge
   - Central cloud-based AI is not sustainable
   - AI will be distributed and federated
   - Edge computing needs a boost

FLANDERS AI RESEARCH PROGRAM
WHERE TO MAKE THE DIFFERENCE ?

1. Applied AI – Complex Decision Making
   - *Big Data* is not always *Good Data*
   - adding *Domain Knowledge* is crucial
2. AI at the Edge
   - Central cloud-based AI is not sustainable
   - AI will be distributed and federated
   - Edge computing needs a boost
3. Autonomous Agents
   - Central Control is not sustainable

FLANDERS AI RESEARCH PROGRAM
WHERE TO MAKE THE DIFFERENCE ?

1. Applied AI – Complex Decision Making
   - *Big Data* is not always *Good Data*
   - adding *Domain Knowledge* is crucial
2. AI at the Edge
   - Central cloud-based AI is not sustainable
   - AI will be distributed and federated
   - Edge computing needs a boost
3. Autonomous Agents
   - Central Control is not sustainable
4. Communicate & Collaborate with Humans
   - Stories & Speech will be the interface

FLANDERS AI PROGRAM STRUCTURE

4 GRAND CHALLENGES

HELP TO **MAKE COMPLEX DECISIONS**

70

6 M

EXTRACT AND **PROCESS** INFORMATION AT THE **EDGE**

30

2.6 M

INTERACT AUTONOMOUSLY WITH OTHER **DECISION-MAKING** ENTITIES

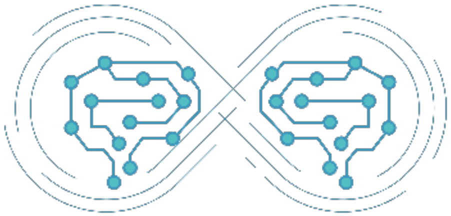

10

0.8 M

COMMUNICATE AND **COLLABORATE SEAMLESSLY** WITH HUMANS

30

2.6 M

Help to make complex decisions through data science

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.

DATA WRANGLING, INTEGRATION & QUALITY HANDLING: TREAT THE INPUT DATA

- Data acquisition
- Labeling and annotation
- Integration structured / unstructured data
- Imperfect data handling
- Data quality estimation
- …

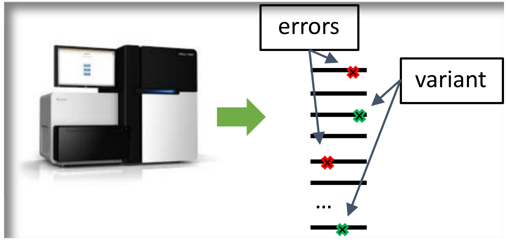

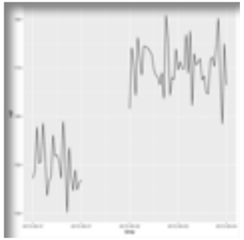

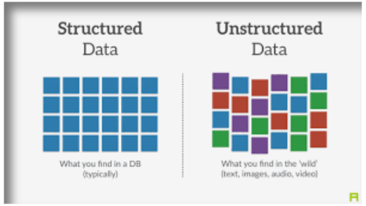

AGORIA · voka · sirris · imec · Flanders Make · VIB · vito · VIL · VUB · KU Leuven · Universiteit Antwerpen · Universiteit Gent · UHasselt

KNOWLEDGE MODELS & REASONING: BRING IN THE EXPERTS

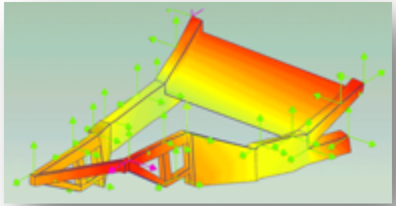

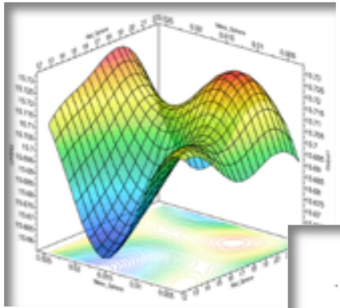

- Graphs
- Simulation models
- Physical models
- Ontologies
- Digital twins
- …

DATA EXPLORATION: YOU DON'T KNOW WHAT'S IN THE DATA

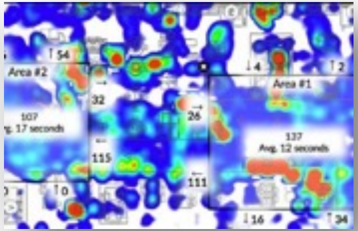

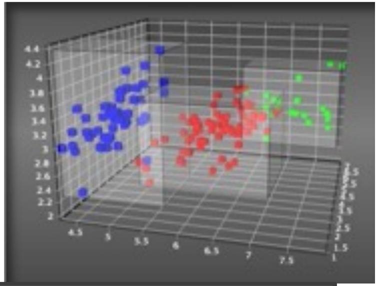

- Dimensionality reduction
- Topological data analysis
- Subjective interestingness
- Personalized visual analytics
- …

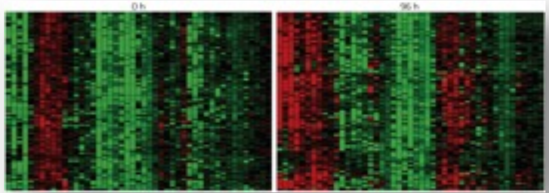

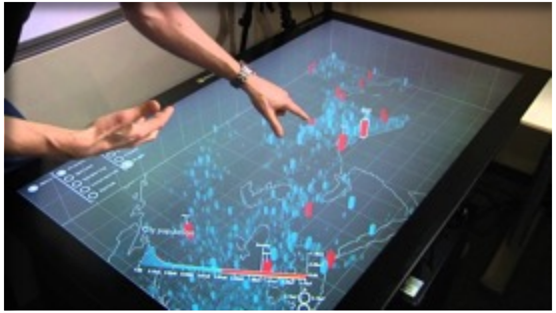

AGORIA · voka Vlaams netwerk van ondernemingen · sirris · imec · FLANDERS MAKE · VIB · vito · VIL EMPOWERING LOGISTICS · VUB VRIJE UNIVERSITEIT BRUSSEL · KU LEUVEN · Universiteit Antwerpen · UNIVERSITEIT GENT · UHASSELT

AUTOMATED MACHINE LEARNING: HOW SELECT THE BEST ML APPROACH

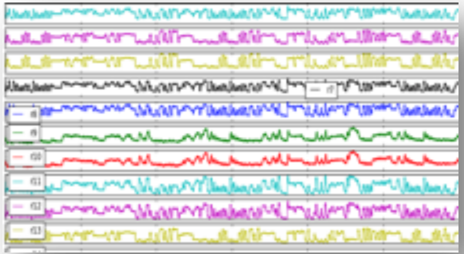

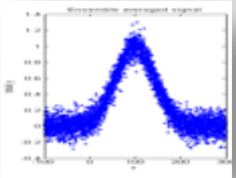

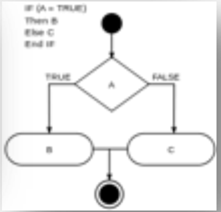

- Data selection
- Model selection
- Feature engineering
- Hyperparameter tuning
  - Neural architecture search
  - Bayesian optimization
- Hybrid modeling
- Transfer learning
- Surrogates/digital twins
- …

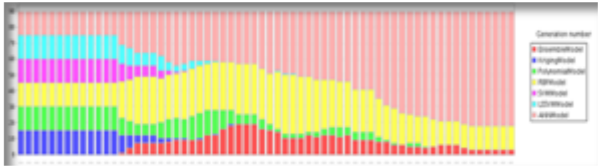

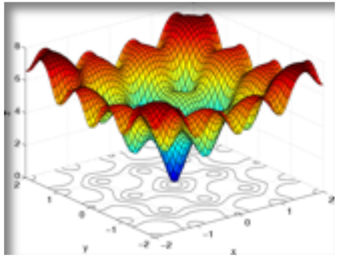

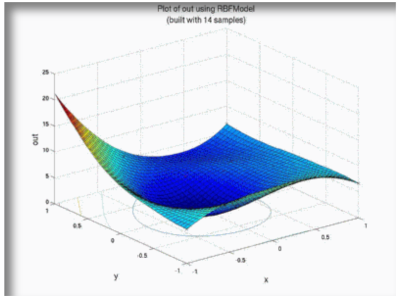

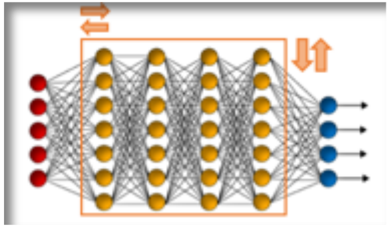

TRUSTWORTHY AI: IS IT ROBUST, EXPLAINABLE AND PRIVATE ?

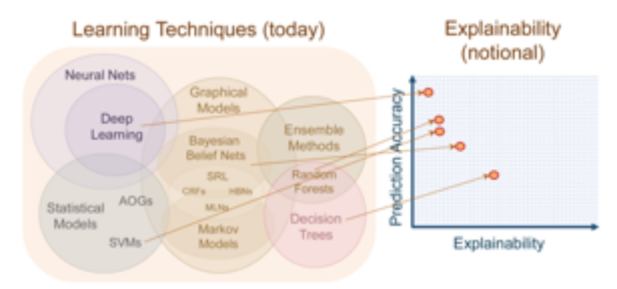

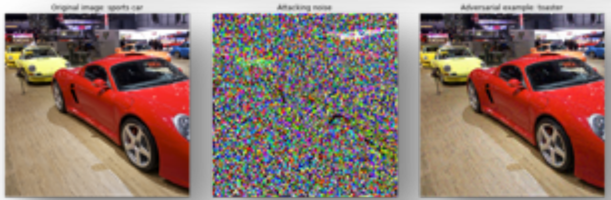

- Fair and Explainable
- Robust and hardened
- Transparent, traceable, reproducible
- Federated AI
- Privacy-preserving AI

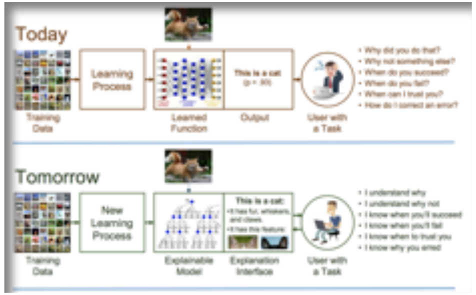

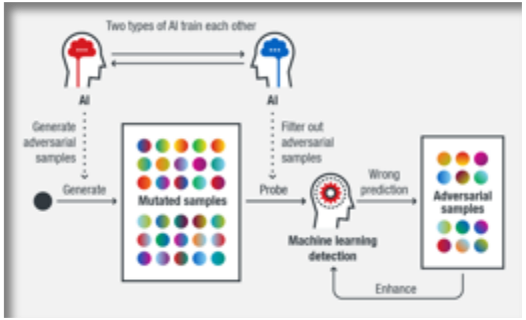

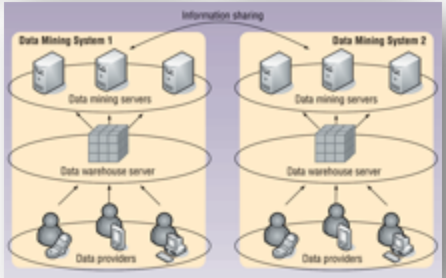

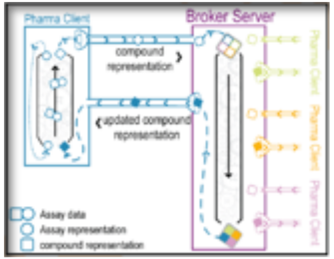

THE BIG PICTURE

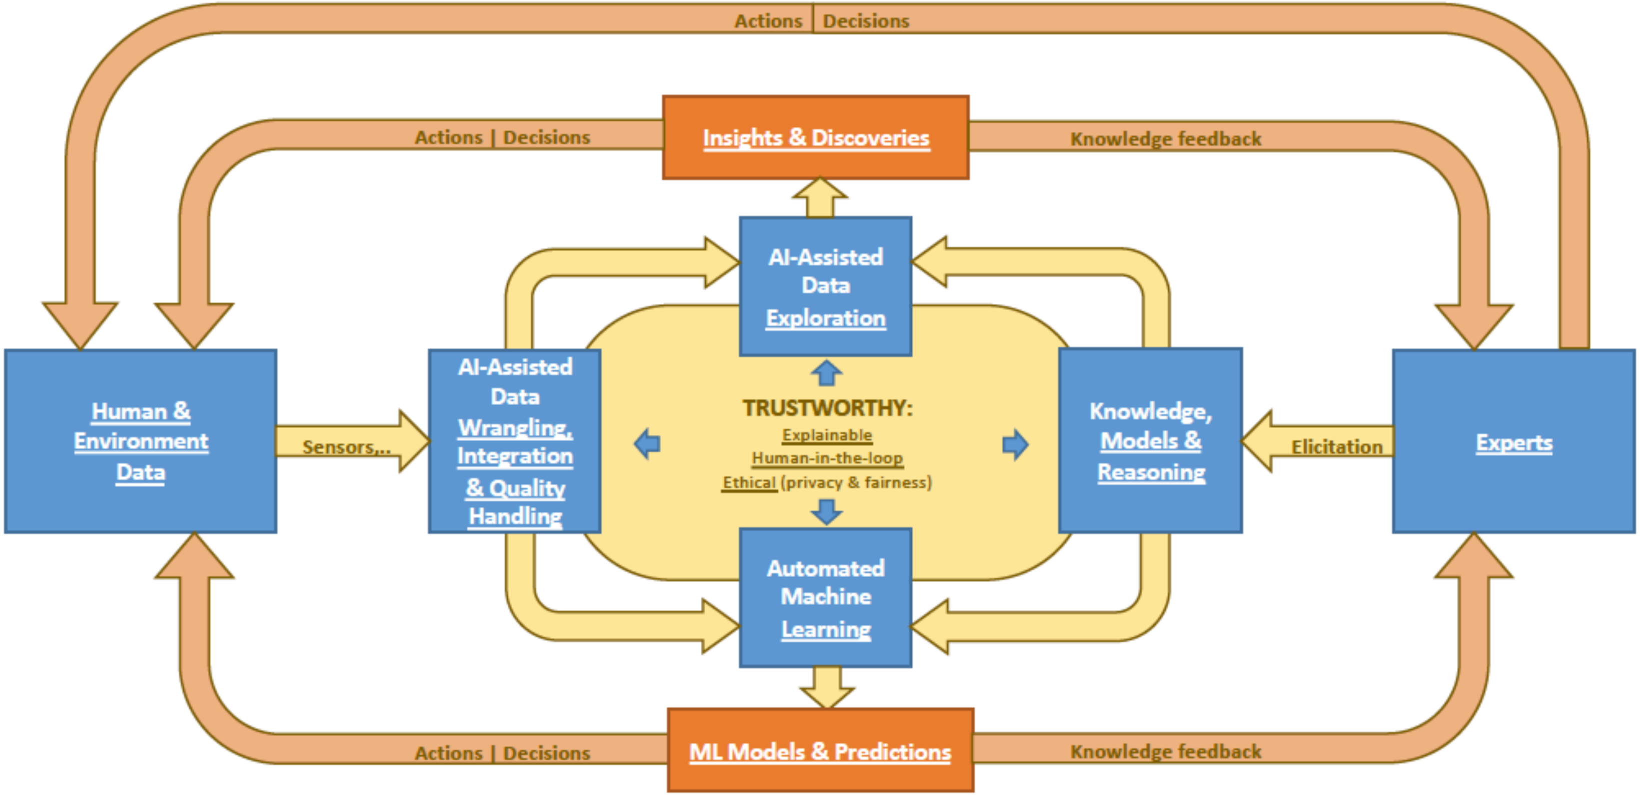

USE CASE EXAMPLE: HEALTHCARE

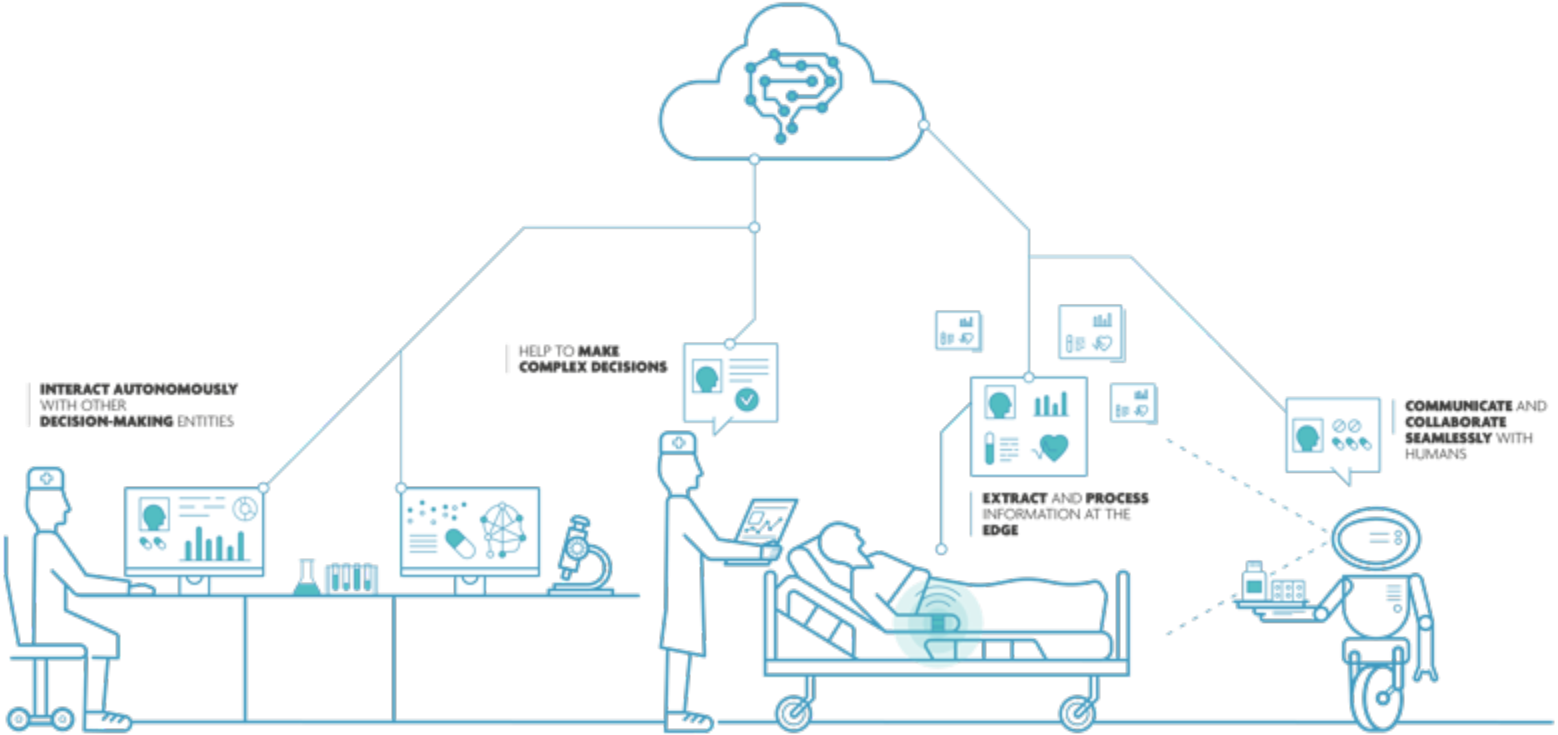

WHO IS IN DEMAND?

PATIENTS

POLICY MAKERS

PROFESSIONALS

IF WE CARE ABOUT THE FUTURE OF CARE...

- PATIENT HEALTH RECORD
- HEALTH DATA ANALYTICS
- TELEMEDICINE & -MONTORING
- WEARABLES & MHEALTH
- ...OMICS (genomic, proteomics, metabolomics, interactomics,...)
- DECISION SUPPORT SYSTEMS

... AI WILL BE KEY


4 P's OF MEDICINE
Personalized
Customized diagnosis and treatment
Preventive
Better than curation
Predictive
Determine risk profiles & predict outcome
Participative
Involve the patient

Data tsunami

Computer Tomography

GS-FLX Roche
Applied Science 454

Sequencers

ACACATTAAATCTTATATG
CTAAAACTAGGTCTCGTTT
TAGGGATGTTTATAACCAT
CTTTGAGATTATTGATGCA
TGGTTATTGGTTAGAAAAA
ATATACGCTTGTTTTTCTTT
CCTAGGTTGATTGACTCAT
ACATGTGTTTCATTGAGGA
AGGAACTTAACAAAACTGC
ACTTTTTTCAACGTCACAG
CTACTTTAAAAGTGATCAA
AGTATATCAAGAAAGCTTA
ATATAAAGACATTTGTTTC
AAGGTTTCGTAAGTGCACA
ATATCAAGAAGACAAAAAT
GACTAATTTTGTTTTCAGG
AAGCATATATATTACACGA
ACACAAATCTATTTTTGTA
ATCAACACCGACCATGGT
TCGATTACACACATTAAAT
CTTATATGCTAAAACTAGG
TCTCGTTTTAGGGATGTTT
ATAACCATCTTTGAGATTA
TTGATGCATGGTTATTGGT
TAGAAAAAATATACGCTTG
TTTTTCTTTCCTAGGTTGA

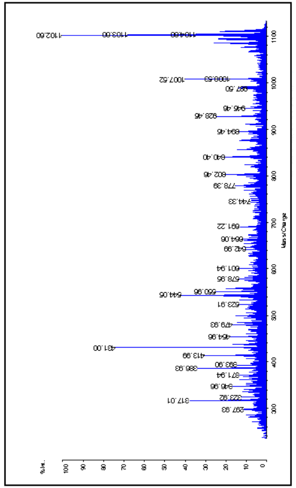

Magnetic resonance

maf VIB
MicroArray Facility

Microarrays
(DNA chips)

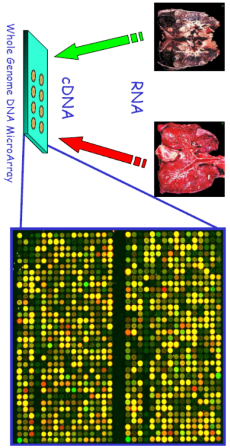

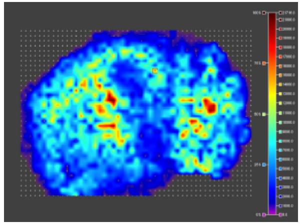

Mass spectrometry

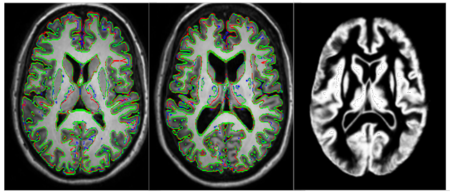

Example: Genomic markers for Leukemia

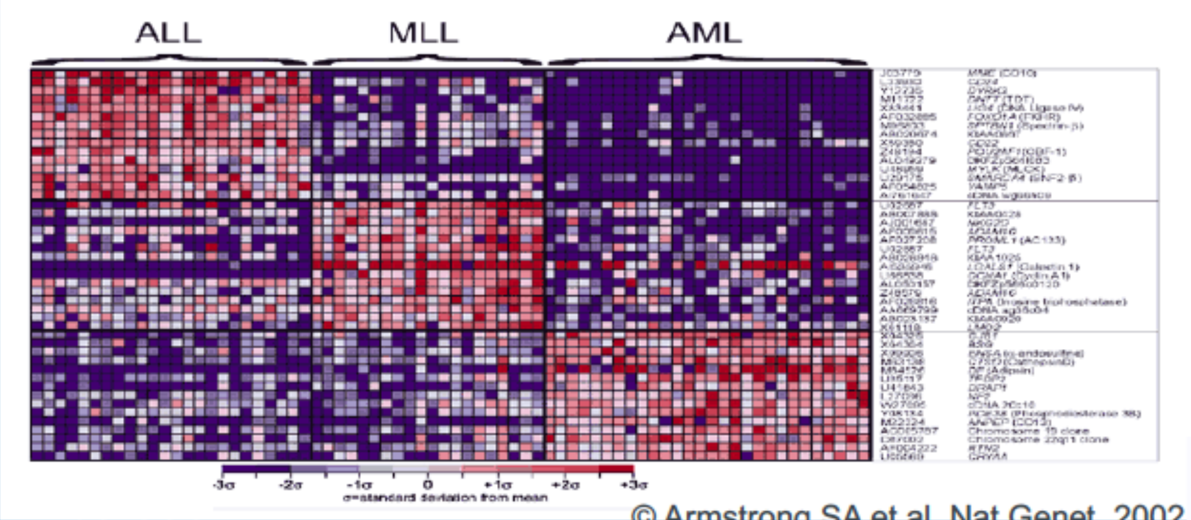

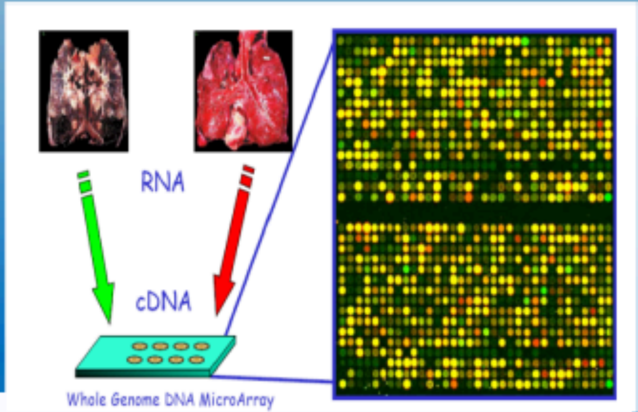

© Armstrong SA et al. Nat Genet. 2002 Jan;30(1):41-7.

12 600 genes
72 patients

- 28 Acute Lymphoblastic Leukemia (ALL)
- 24 Acute Myeloid Leukemia (AML)
- 20 Mixed Linkage Leukemia (MLL)

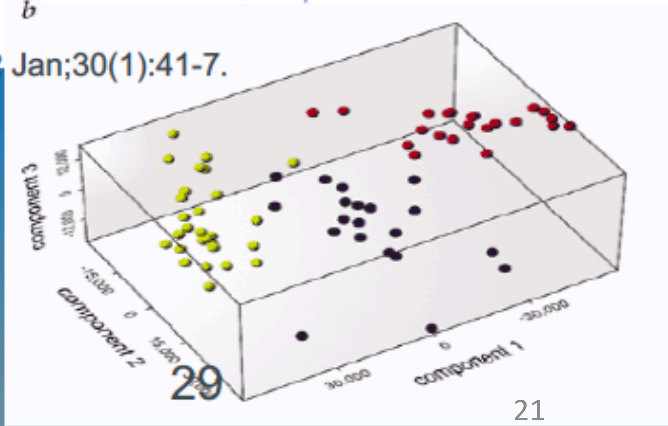

29

Automating Data Science

Sleep apnea detection from data streams

Automated detection of sleep apnea events

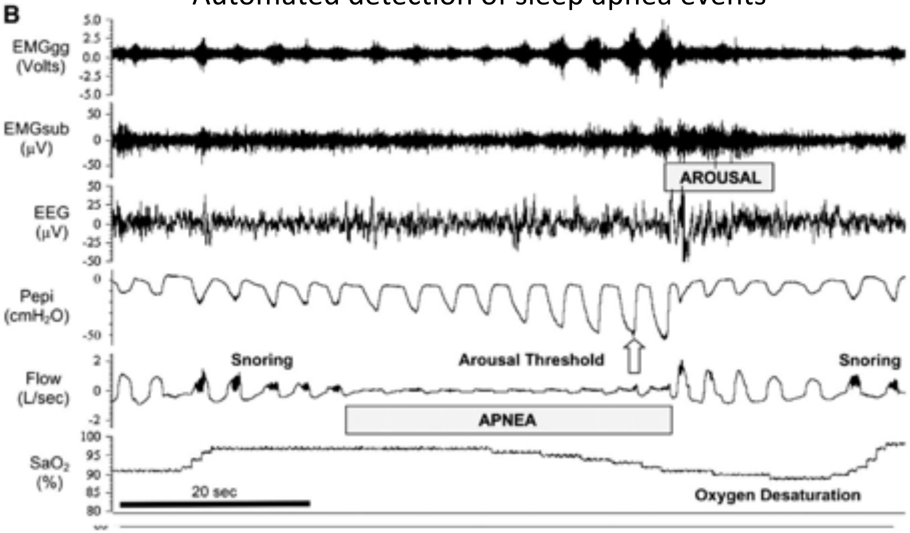

Data annotated by nurses according to AASM guidelines

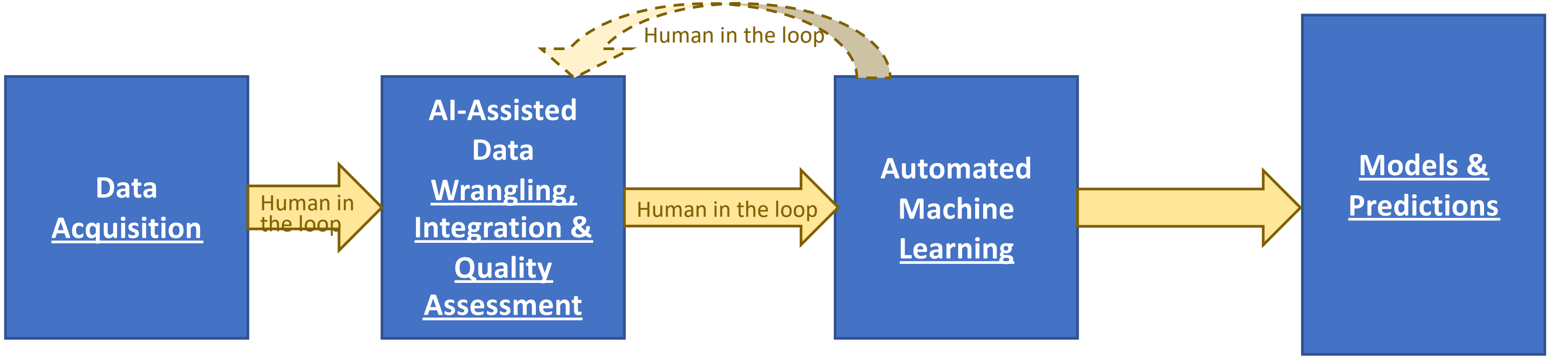

Example: Glycemia control in ICU

- 10 mio adult ICU patients / year (EU + US) (1-2 b$ market)
- ‘Tight Glycemic Control (TGC) in intensive care unit lowers mortality’
  - implement through LOGIC-Insulin: semi-automatic control system that advises nurse on insulin dosage and blood sampling interval aiming at TGC and avoiding hypoglycemia
- LOGIC-I randomized clinical trial (single-centre): compared with expert nurses, LOGIC-Insulin showed improved efficacy of TGC without increasing rate of hypoglycemia
- LOGIC-II randomized clinical trial (multi-centre): Start February 2014

in collaboration with ICU UZ Leuven

UZ LEUVEN

USE CASE EXAMPLE: INDUSTRY 4.0

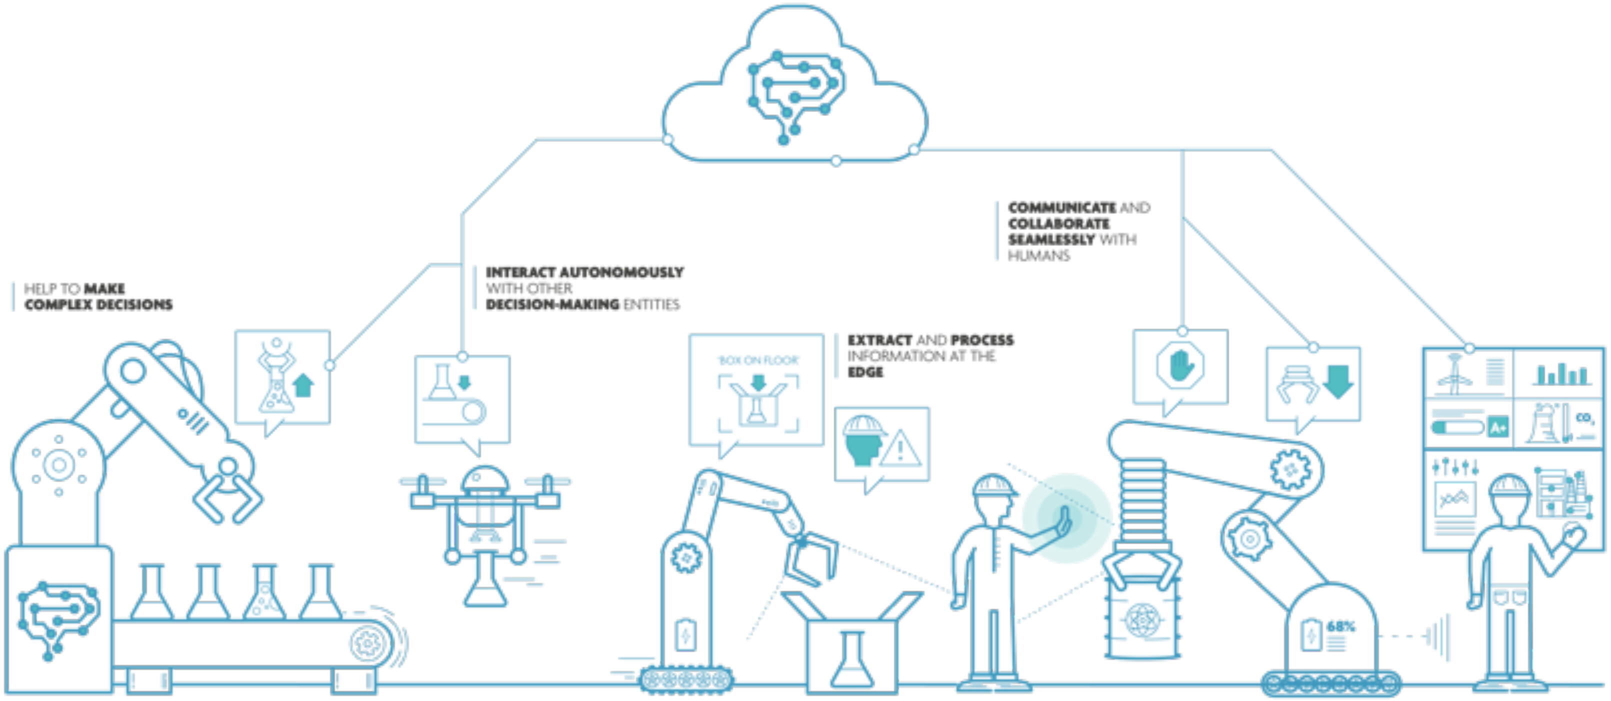

Automating Data Science

Automated ML model building for predictive maintenance

Predictive maintenance in manufacturing

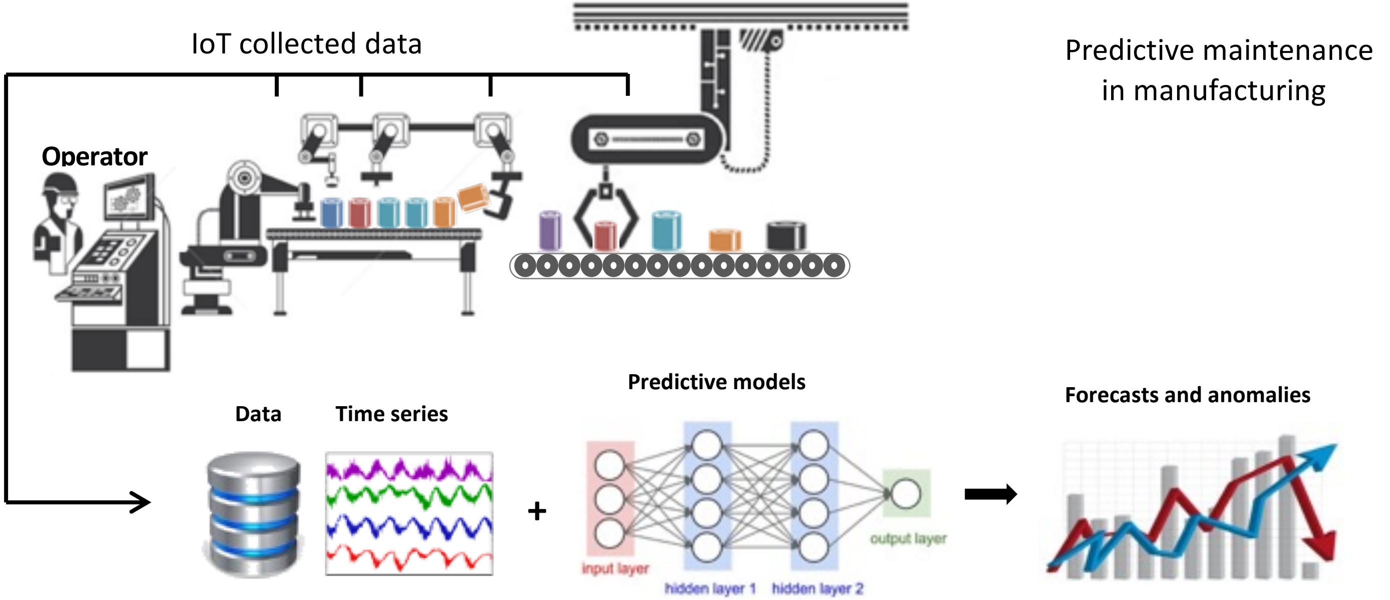

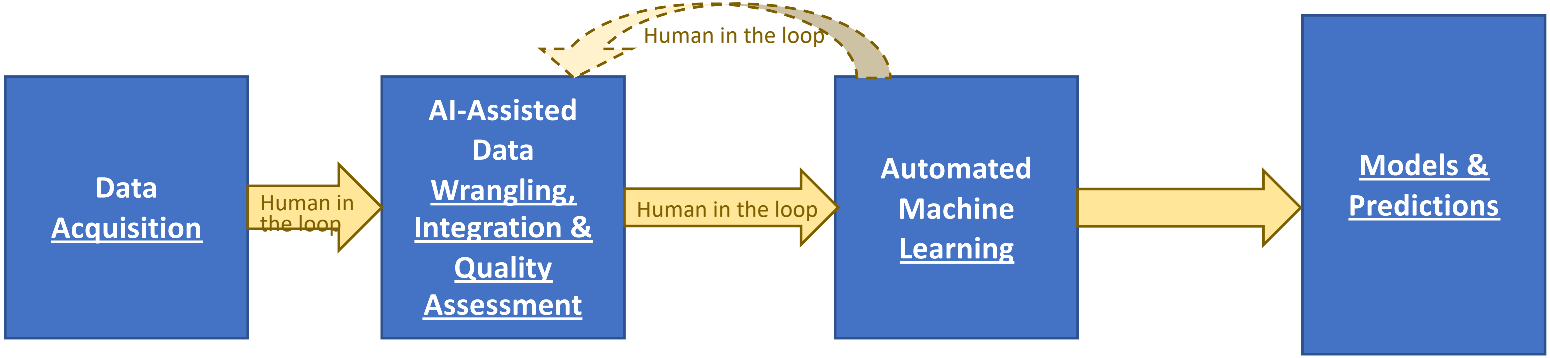

Mechanical structure monitoring DSS

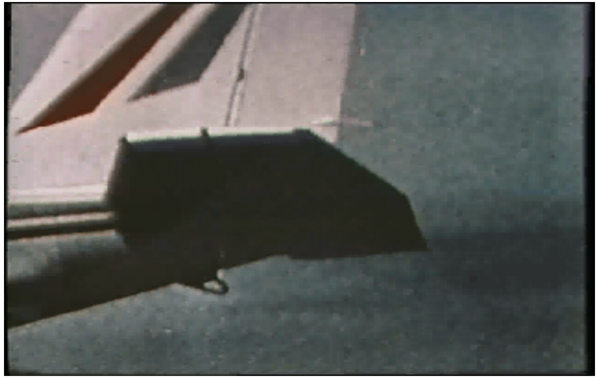

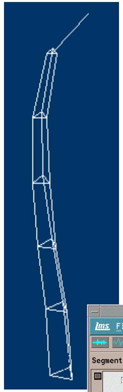

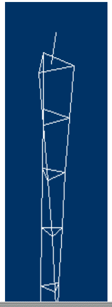

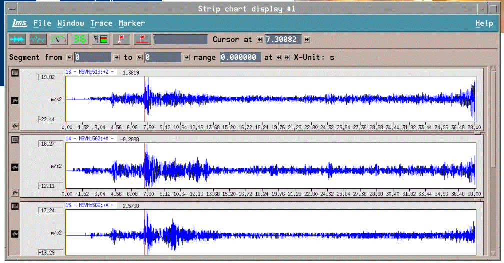

STADIUS
Center for Dynamical Systems, Signal Processing and Data Analytics

LMS
A Siemens Business

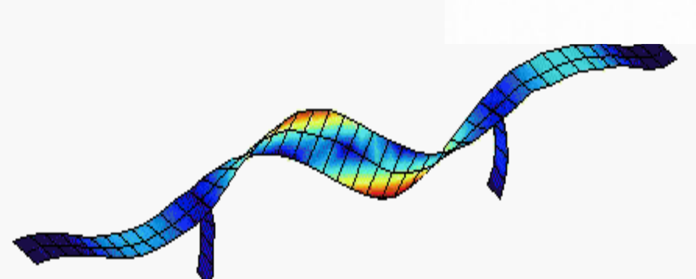

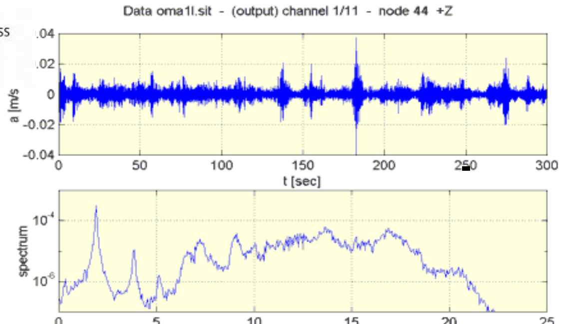

Chemical process DSS

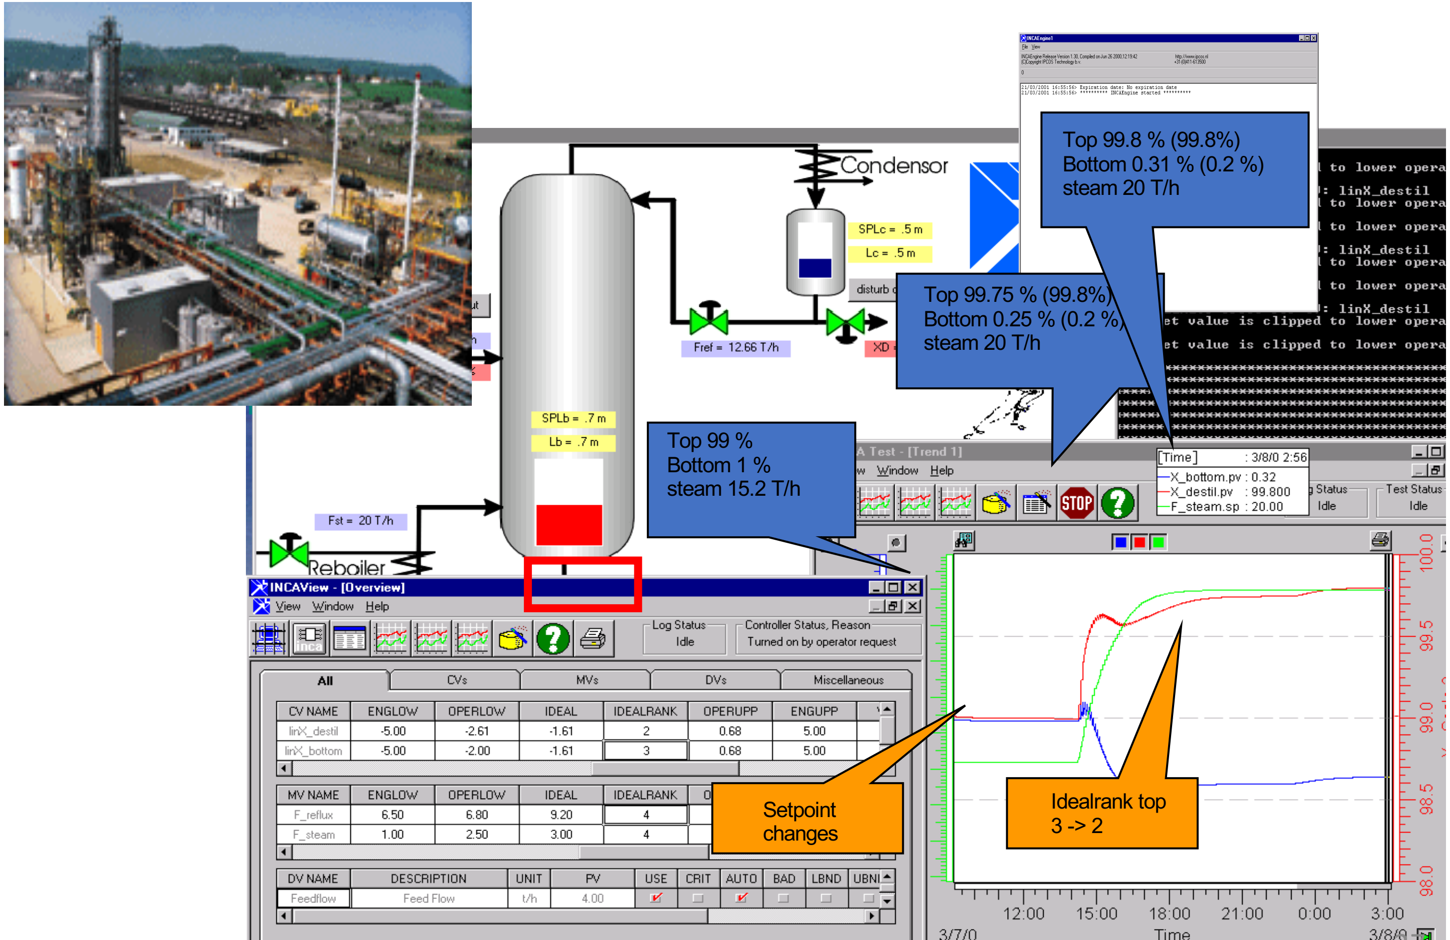

USE CASE EXAMPLE: MOBILITY AND LOGISTICS – SMART CITIES

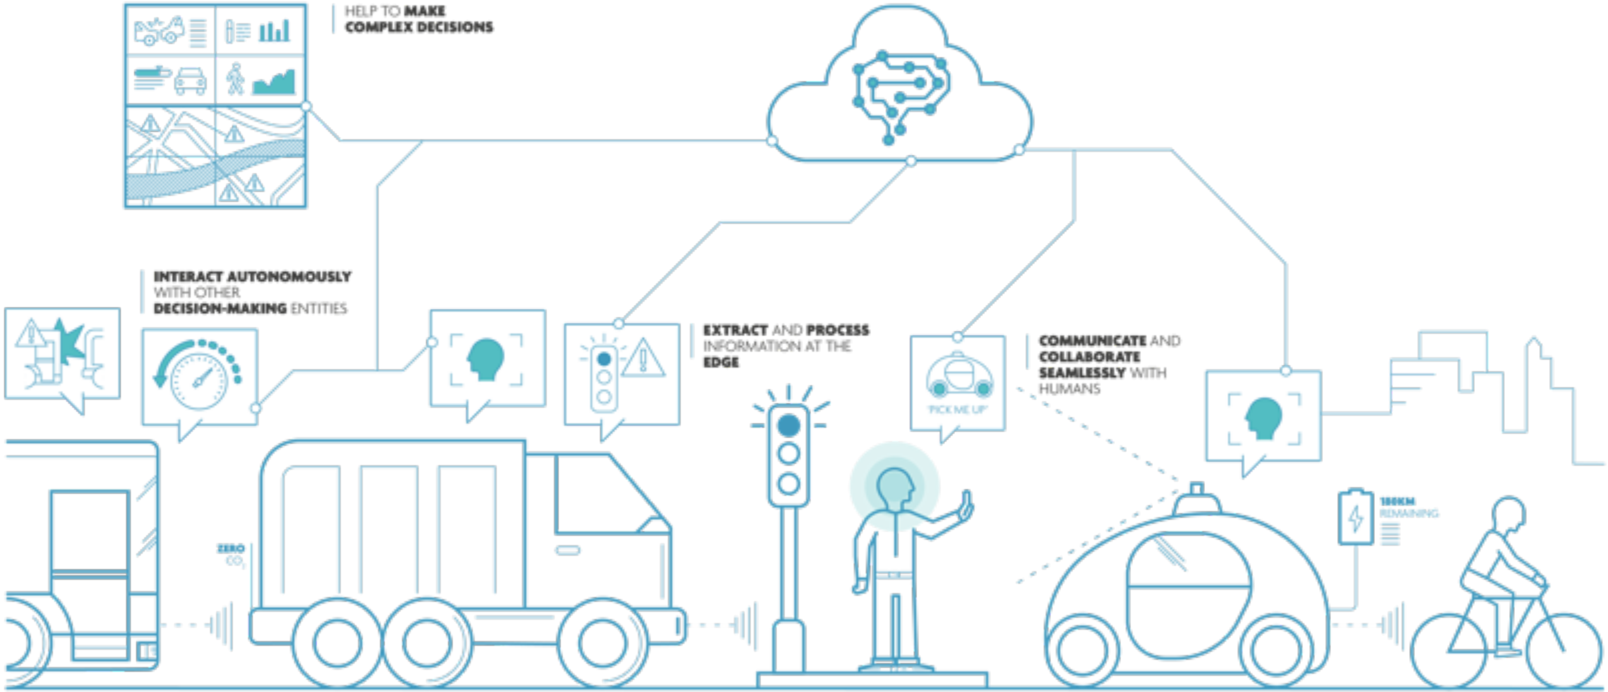

Knowledge discovery

Opportunistic sensing

Propagation of uncertainty

opportunistic data
uncertain conditions

Data Acquisition

Human in the loop

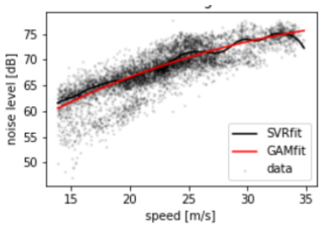

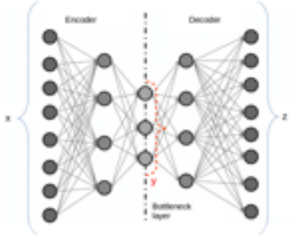

removing known
dependencies
location bas

AI-Assisted Data Wrangling, Integration & Quality Assessment

Human in the loop

Human in the loop

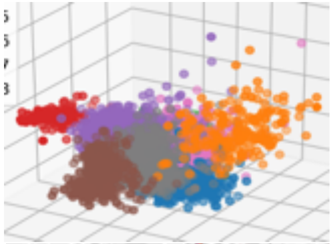

cluster guided
spatial exploration

AI-Assisted Data Exploration

Human in the loop

Road degradation
Safety
Air pollution
Noise

ML models
→ external indicators

Human in the loop

Insights & Discoveries

Models & Predictions

Better understanding

Traffic & Mobility DSS

Detector technology: inductive loops, Gatso-meters, camera's

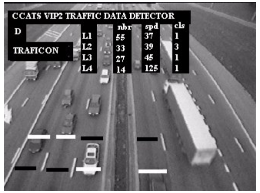

Density – Flow

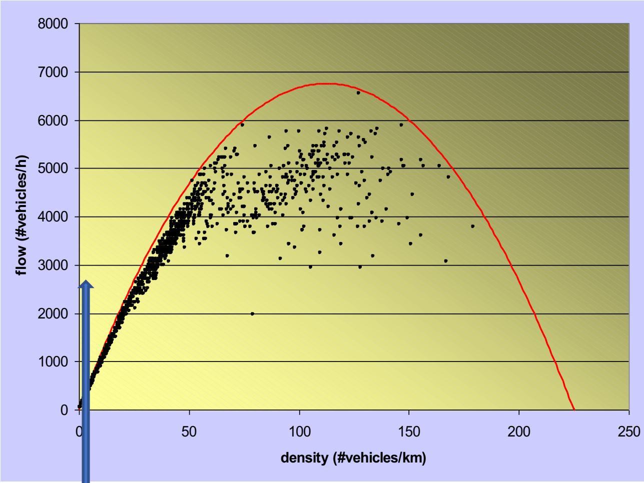

Density per hour / day of the week

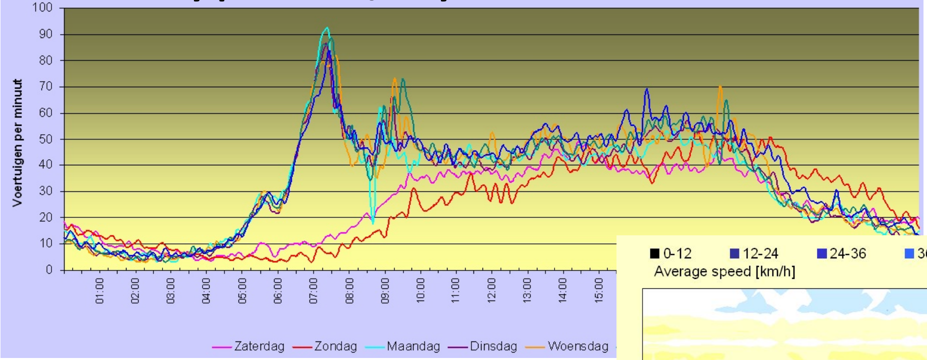

Traffic jam prediction

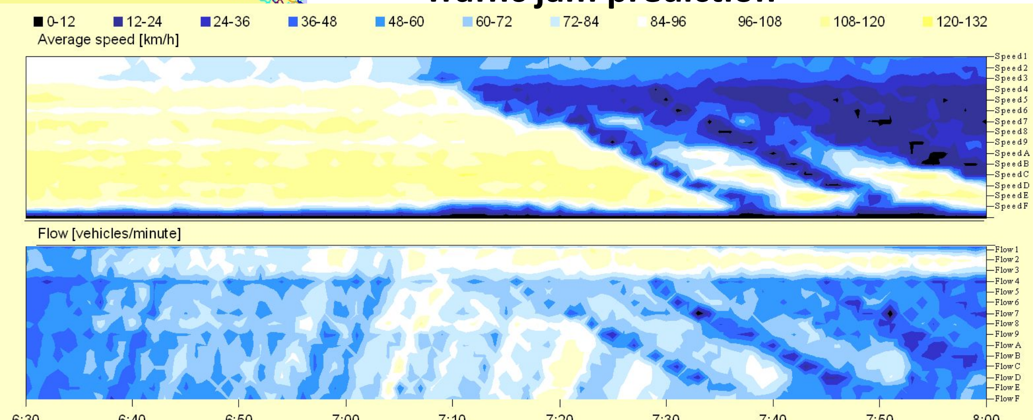

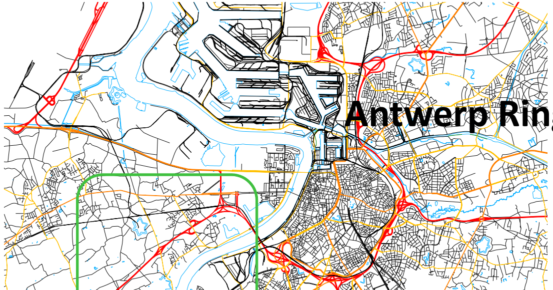

Antwerp Ring

Traffic & Mobility DSS: control

Speed harmonisation

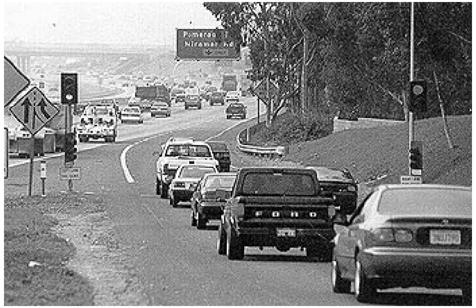

Ramp metering

DRIP

Vlaams Verkeerscentrum

Vlaamse overheid

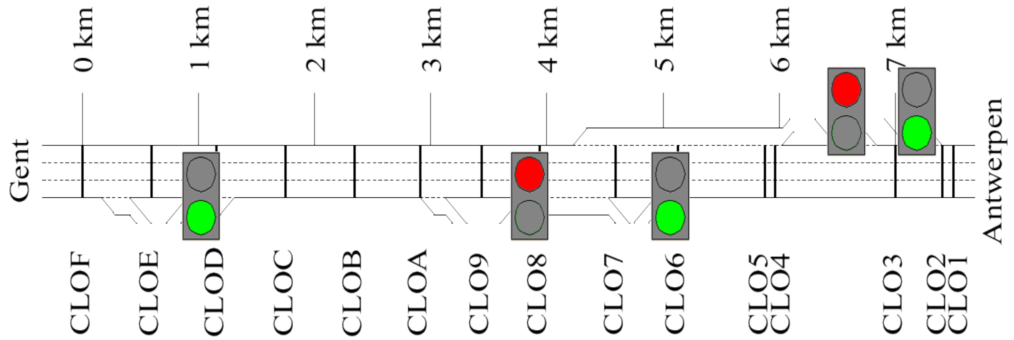

STADIUS
Center for Dynamical Systems, Signal Processing and Data Analytics

TRANSPORT & MOBILITY LEUVEN

Flanders O3/fine particle DSS

O_3 air-quality stations

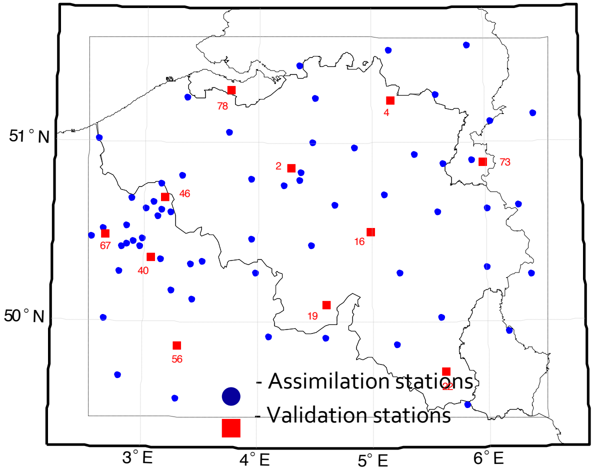

Average of the O_3 concentration over the validation stations

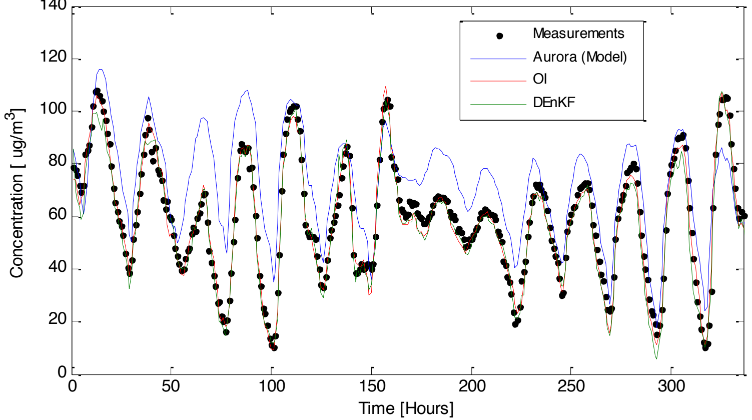

Starting date: May 28th, 2005 at midnight

Average PM_{10} concentration

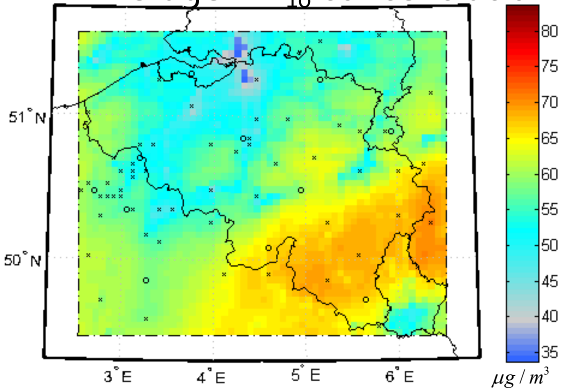

Average PM_{10} concentration

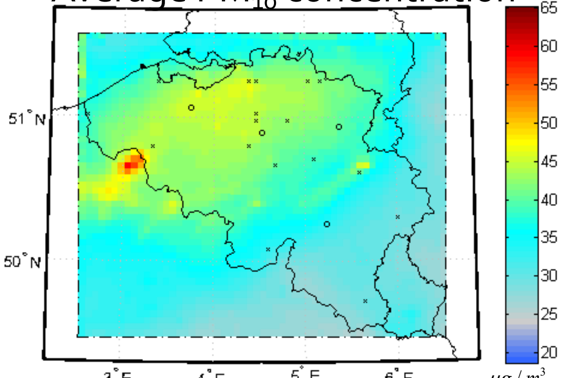

Demer Flood Regulation DSS

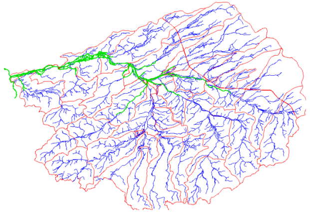

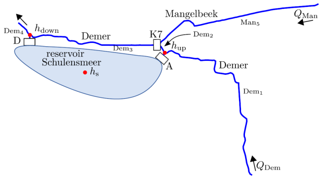

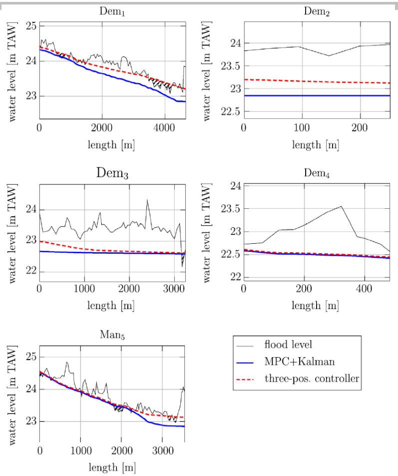

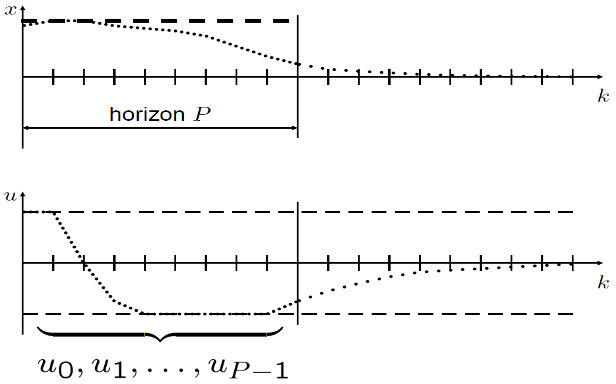

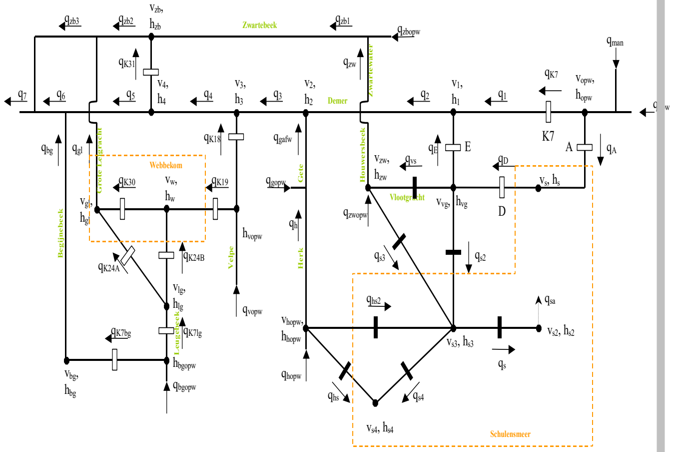

Belgian smart electricity grid DSS

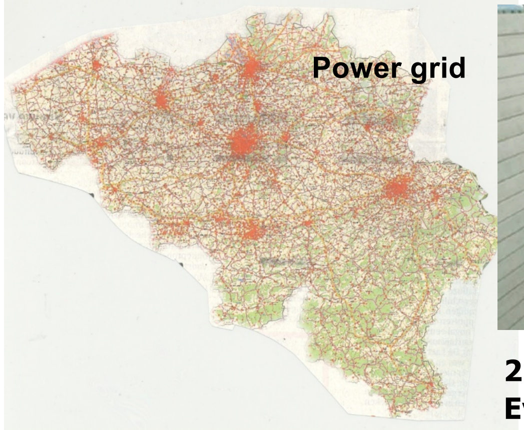

250 transformer substations
Every 15 min, 5 years

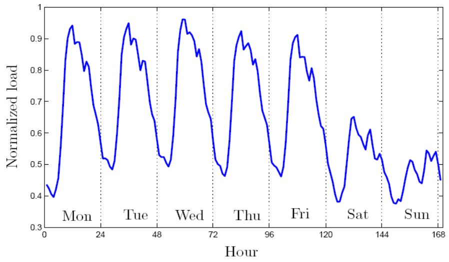

1 post, 1 week

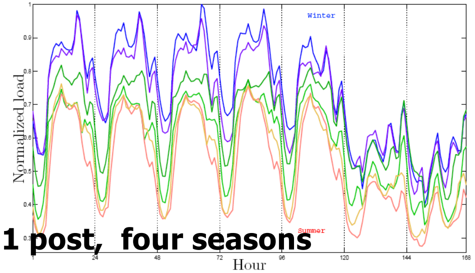

1 post, four seasons

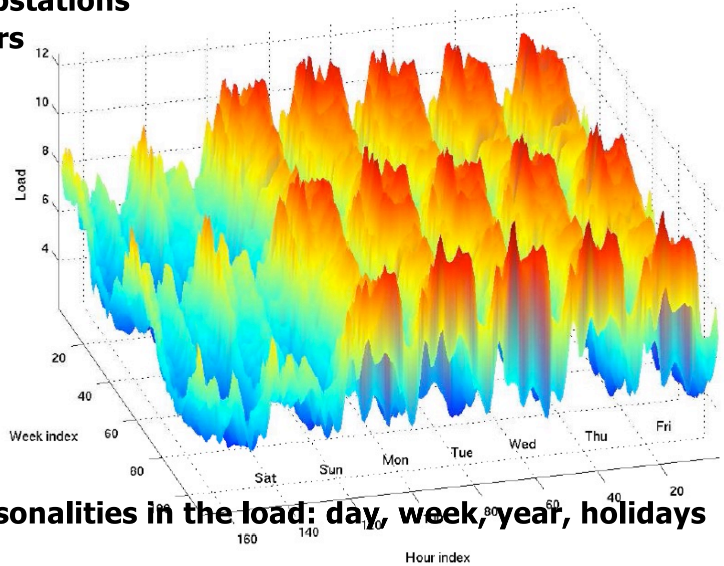

Seasonalities in the load: day, week, year, holidays

STADIUS
Center for Dynamical Systems, Signal Processing and Data Analytics

elia

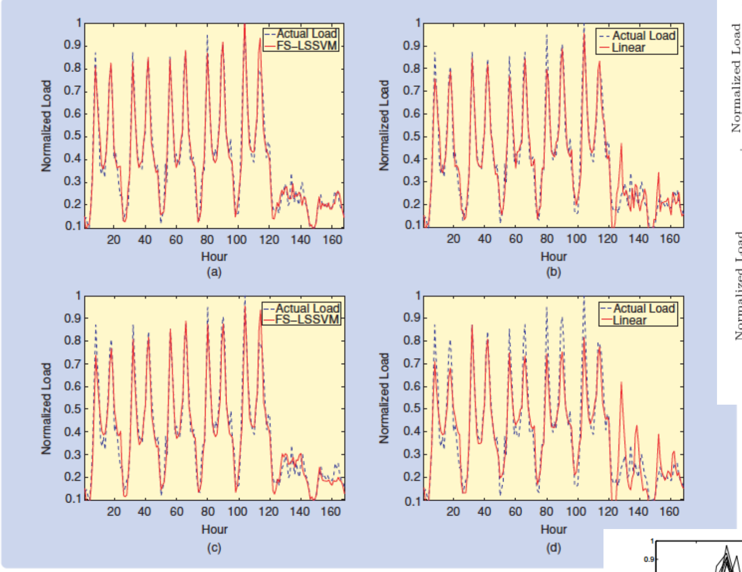

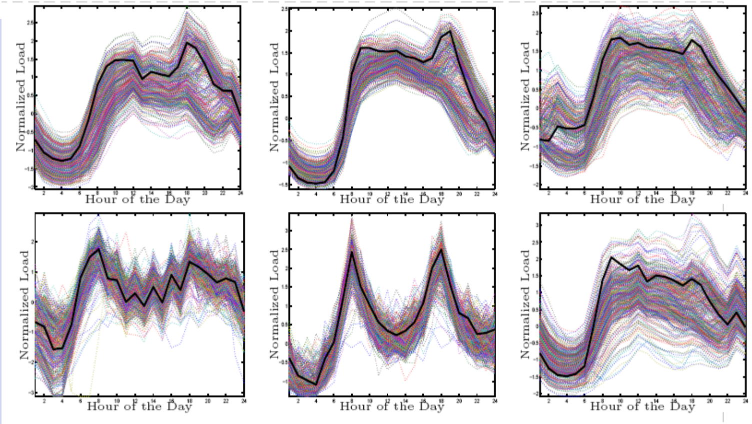

6 posts, 1 year
Seasonalities, calender holidays !

1 month predictions
depending
on day, season
and weather prediction

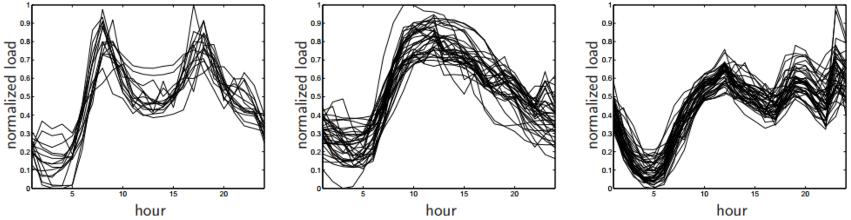

Electricity load: 245 substations in Belgian grid (1/2 train, 1/2 validation)
$x_i \in \mathbb{R}^{43.824}$: spectral clustering on **high dimensional data** (5 years)

3 of 7 detected clusters:
- 1: *Residential profile:* *morning and evening peaks*
- 2: *Business profile:* *peaked around noon*
- 3: *Industrial profile:* *increasing morning, oscillating afternoon and evening*

Customer profiling:
Residential, business, industrial

Fraud Detection DSS (phones, credit cards, tax declaration,...)

	Short Duration	Long Duration	High Frequency	International	Same Destination	Off Peak	Call Forwarding	Behaviour Change
Direct call selling		X	X	X			X	
PABX fraud	X		X		X	X		X
Freephone fraud	X		X		X			X
Premium rate fraud		X	X		X			X
Subscription fraud			X					
Handset theft		X	X	X	X			X

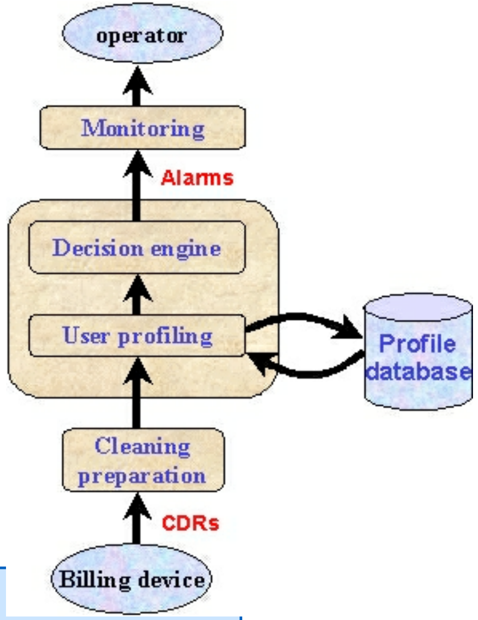

TAXonWEB

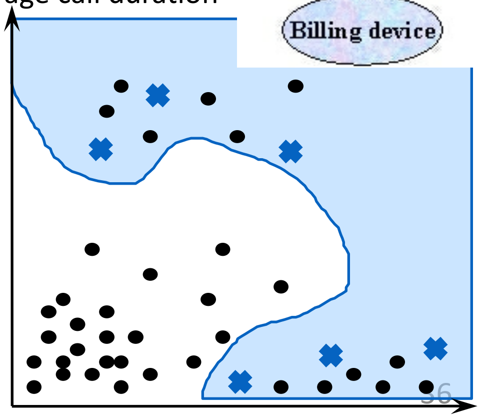

Flanders AI program structure

4 GRAND CHALLENGES

HELP TO MAKE **COMPLEX DECISIONS**

70
6 M

EXTRACT AND **PROCESS** INFORMATION AT THE **EDGE**

30
2.6 M

INTERACT AUTONOMOUSLY WITH OTHER **DECISION-MAKING** ENTITIES

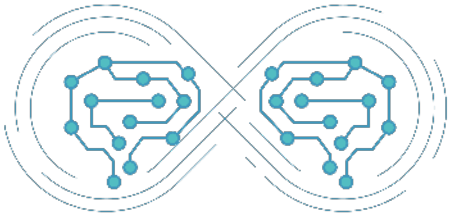

10
0.8 M

COMMUNICATE AND **COLLABORATE SEAMLESSLY** WITH HUMANS

30
2.6 M

Deliver artificial intelligence to the edge

EXTRACT AND **PROCESS** INFORMATION AT 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.

DISTRIBUTED SYSTEMS

DATA VERSUS COMPUTE TRADE-OFF

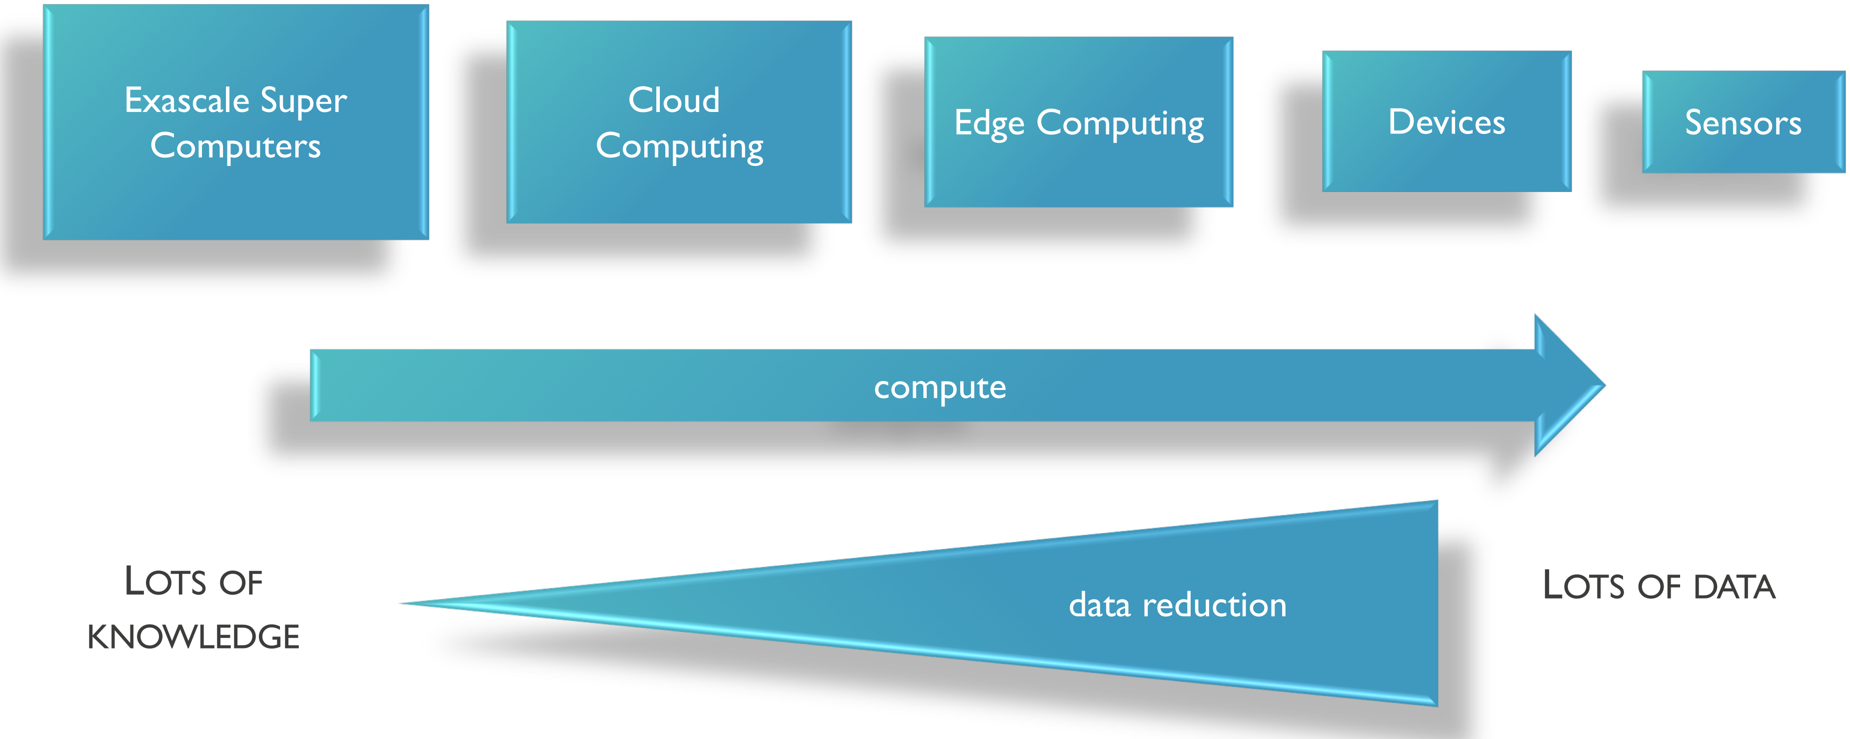

REAL-TIME DISTRIBUTED AND HIERARCHICAL AI FROM EXTREME EDGE TO CLOUD

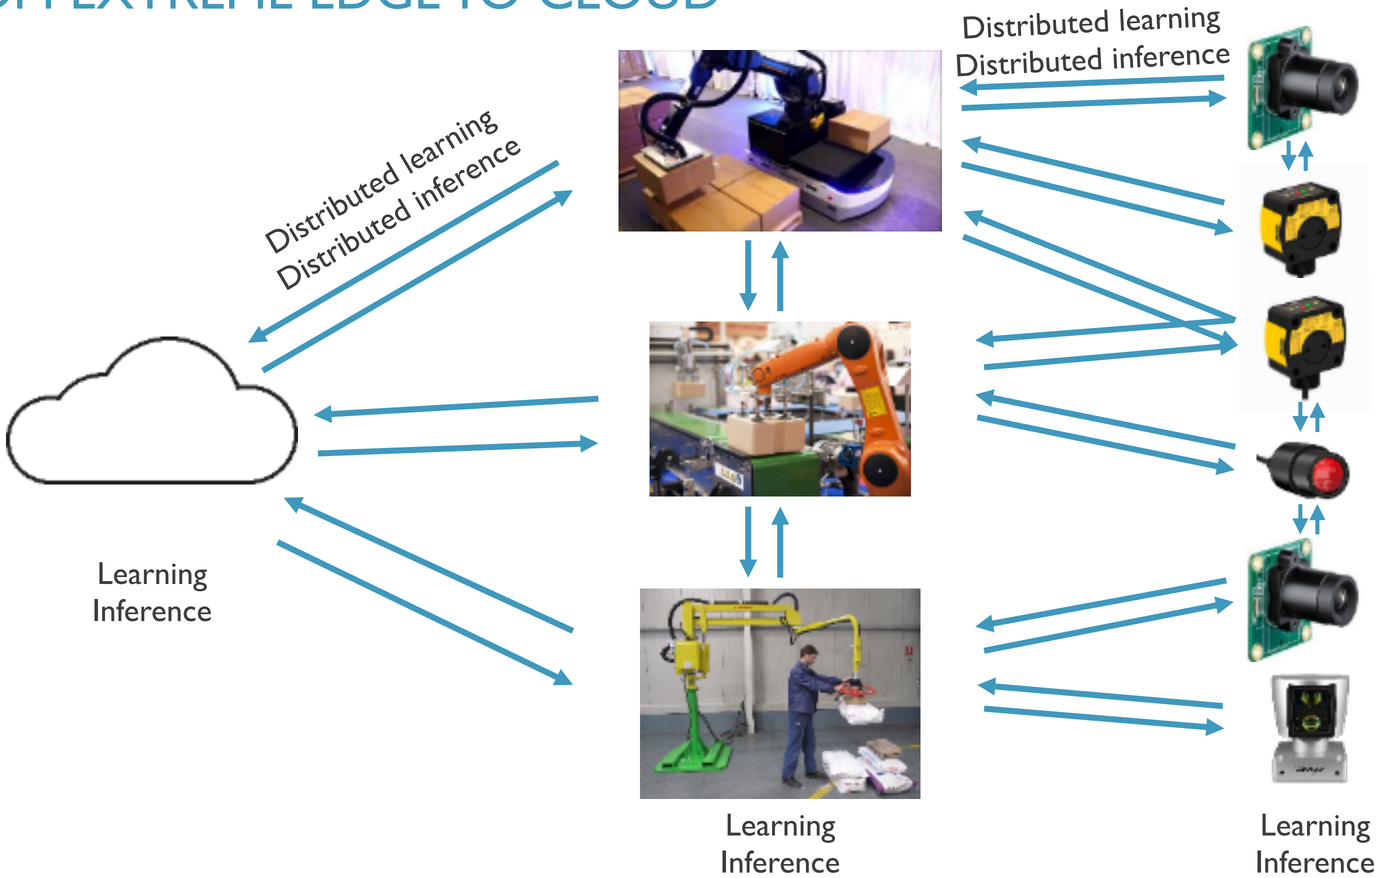

AGORIA · voka · sirris · imec · FLANDERS MAKE · VIB · vito · VIL · VUB · KU LEUVEN · Universiteit Antwerpen · Universiteit Gent · UHASSELT

REAL-TIME DISTRIBUTED AND HIERARCHICAL AI FROM EXTREME EDGE TO CLOUD

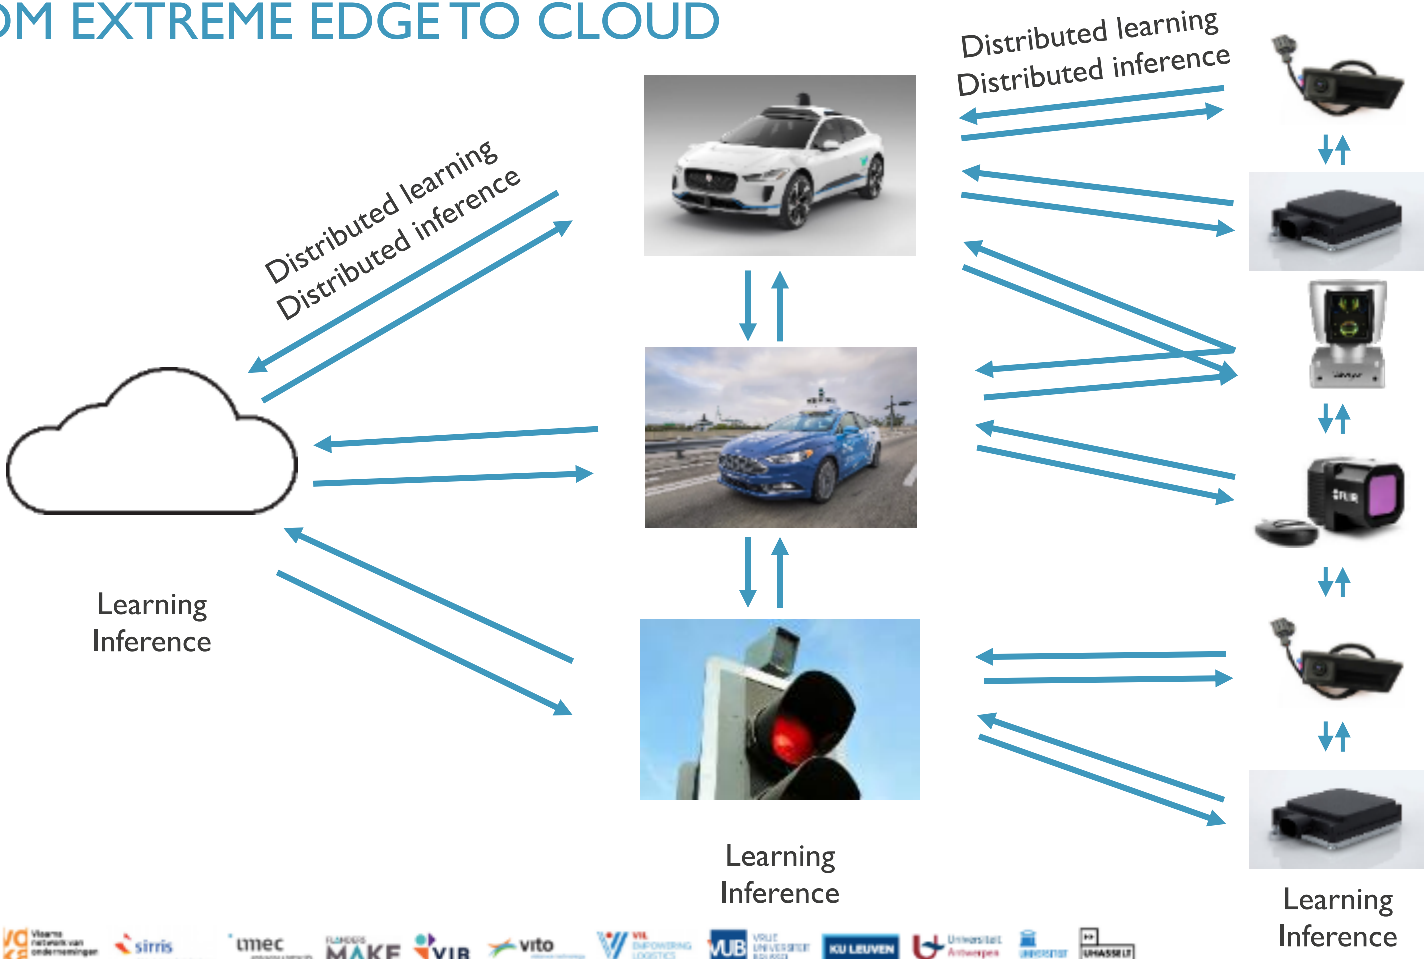

EXTREME EDGE

SENSORS WITH HARDWARE ACCELERATION FOR MACHINE LEARNING

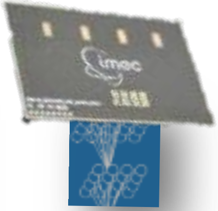

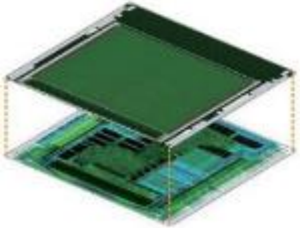

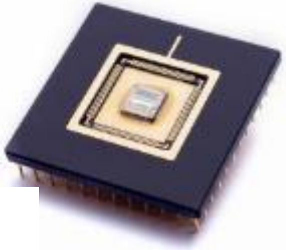

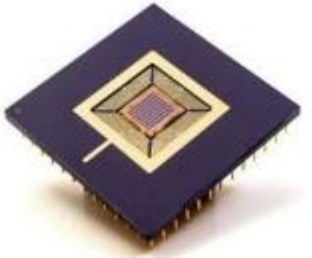

SOLVING THE ENERGY EFFICIENCY BOTTLENECK (ANALOG IN-MEMORY COMPUTE)

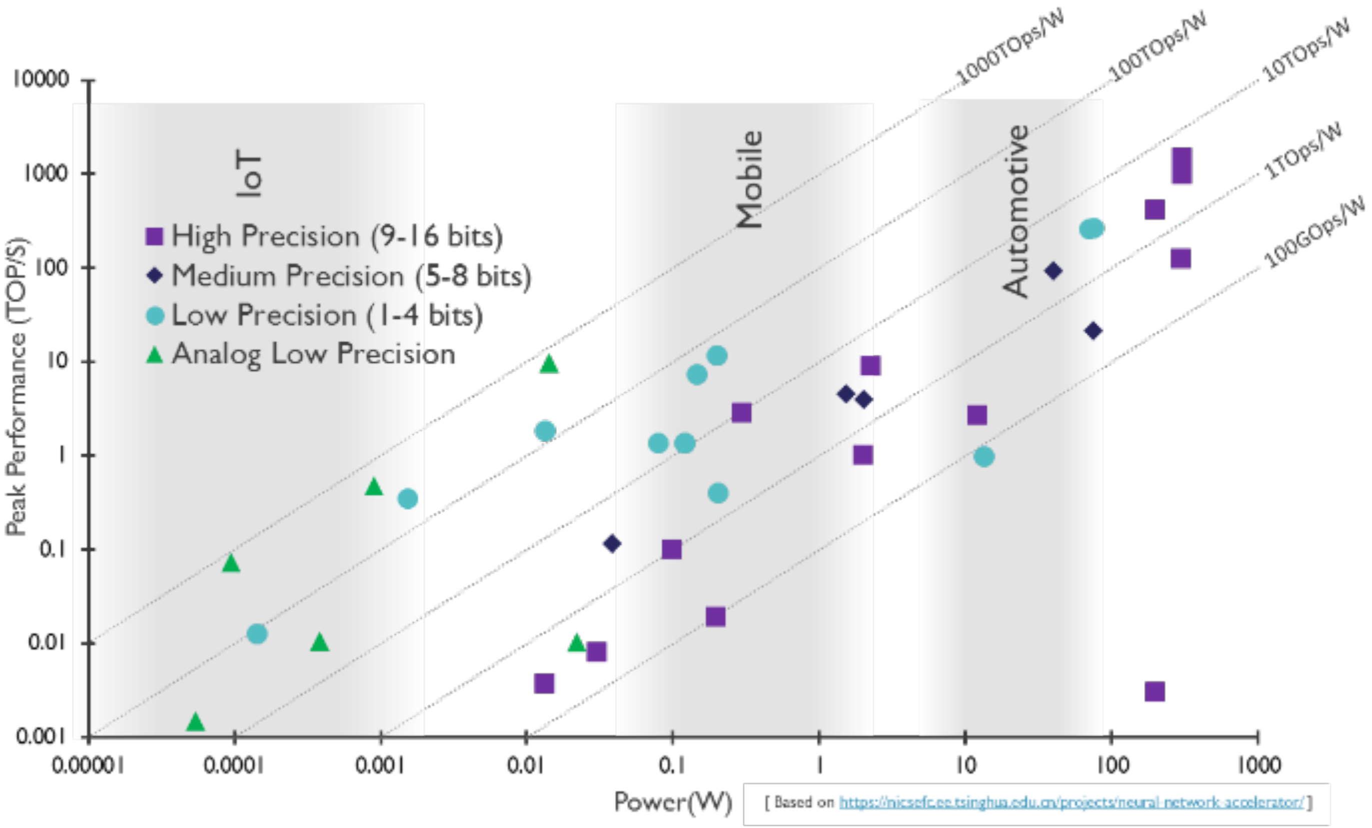

[Based on https://nicsefc.ee.tsinghua.edu.cn/projects/neural-network-accelerator/]

SOLVING THE ENERGY EFFICIENCY BOTTLENECK (ANALOG IN-MEMORY COMPUTE)

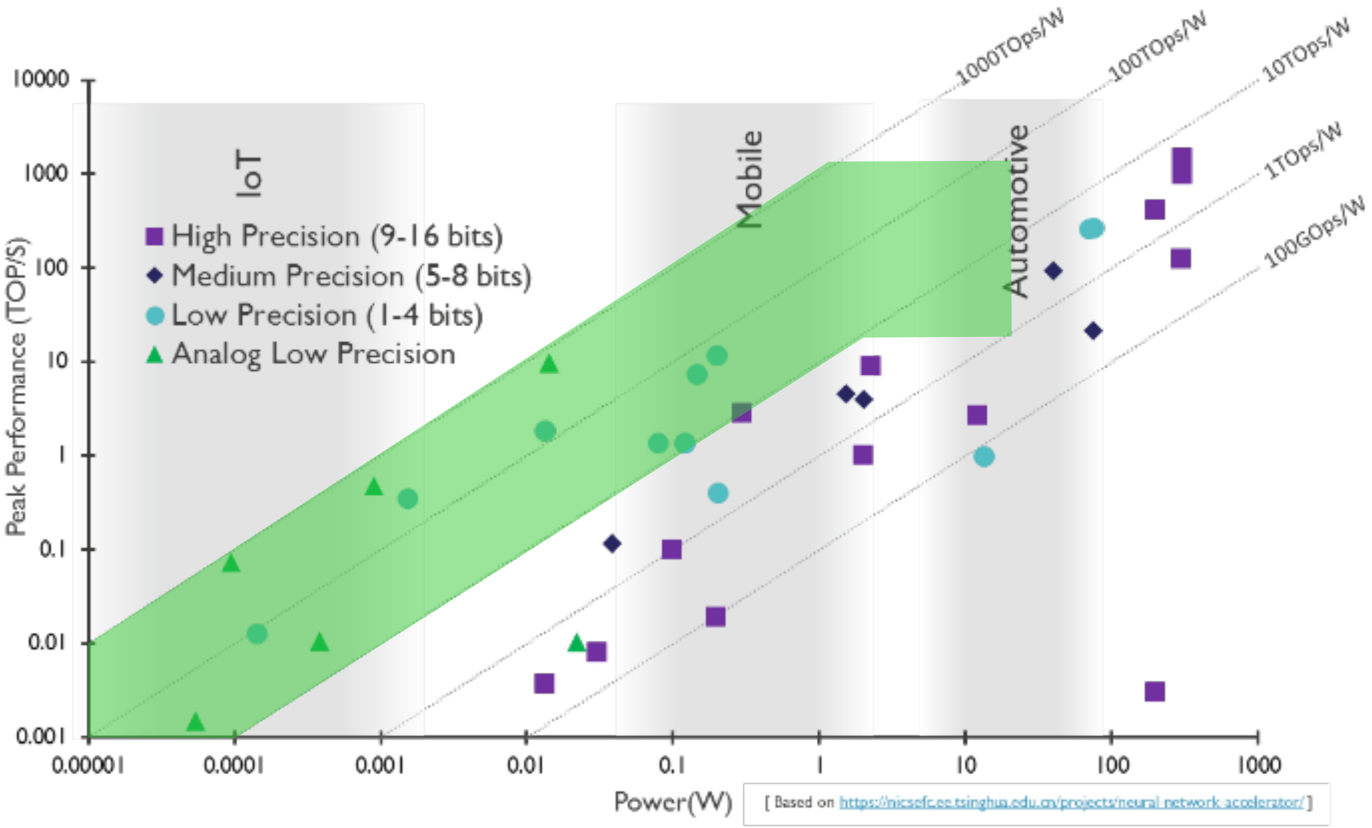

[Based on https://nicsefc.ee.tsinghua.edu.cn/projects/neural-network-accelerator/]

AGORIA · VOKA Vlaams netwerk van ondernemingen · sirris · imec · FLANDERS MAKE · VIB · vito · VIL · VUB · KU LEUVEN · Universiteit Antwerpen · Universiteit Gent · UHASSELT

FLANDERS AI PROGRAM STRUCTURE

4 GRAND CHALLENGES

HELP TO **MAKE COMPLEX DECISIONS**

70
6 M

EXTRACT AND **PROCESS** INFORMATION AT THE **EDGE**

30
2.6 M

INTERACT AUTONOMOUSLY WITH OTHER **DECISION-MAKING** ENTITIES

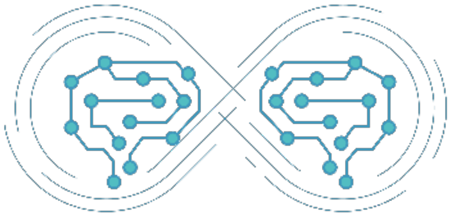

10
0.8 M

COMMUNICATE AND **COLLABORATE SEAMLESSLY** WITH HUMANS

30
2.6 M

Interact autonomously with other decision-making entities

INTERACT AUTONOMOUSLY WITH OTHER **DECISION-MAKING** ENTITIES

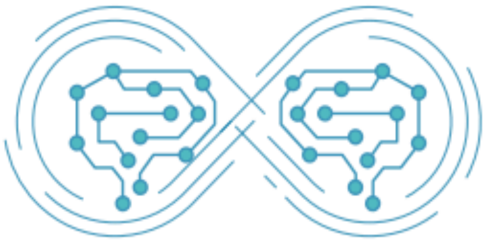

Autonomous decision-making entities each have their own goals and intentions. In **multi-agent systems**, they need to interact with each other. Multi-agent systems are radically different from distributed systems. In multi-agent systems:

- 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.

AUTONOMOUS AGENTS

NO ONE SIZE FITS ALL

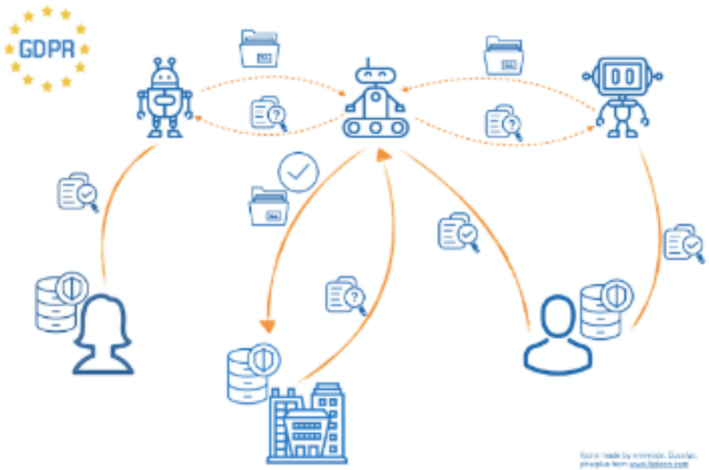

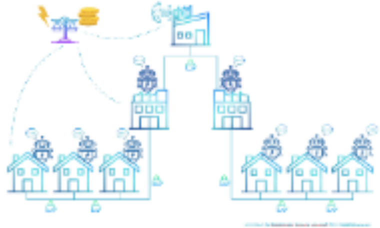

Data Economy

users decide to when to share data with whom when for which service

FLANDERS AI PROGRAM STRUCTURE

4 GRAND CHALLENGES

HELP TO **MAKE COMPLEX DECISIONS**

70

6 M

EXTRACT AND **PROCESS** INFORMATION AT THE **EDGE**

30

2.6 M

INTERACT AUTONOMOUSLY WITH OTHER **DECISION-MAKING** ENTITIES

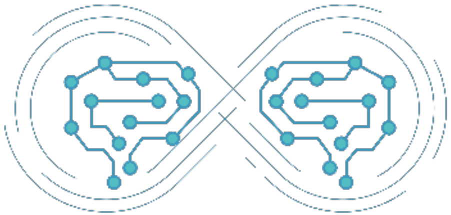

10

0.8 M

COMMUNICATE AND **COLLABORATE SEAMLESSLY** WITH HUMANS

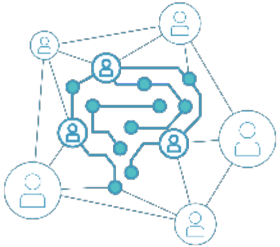

30

2.6 M

Communicate and collaborate seamlessly with humans

COMMUNICATE AND **COLLABORATE SEAMLESSLY** WITH HUMANS

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.

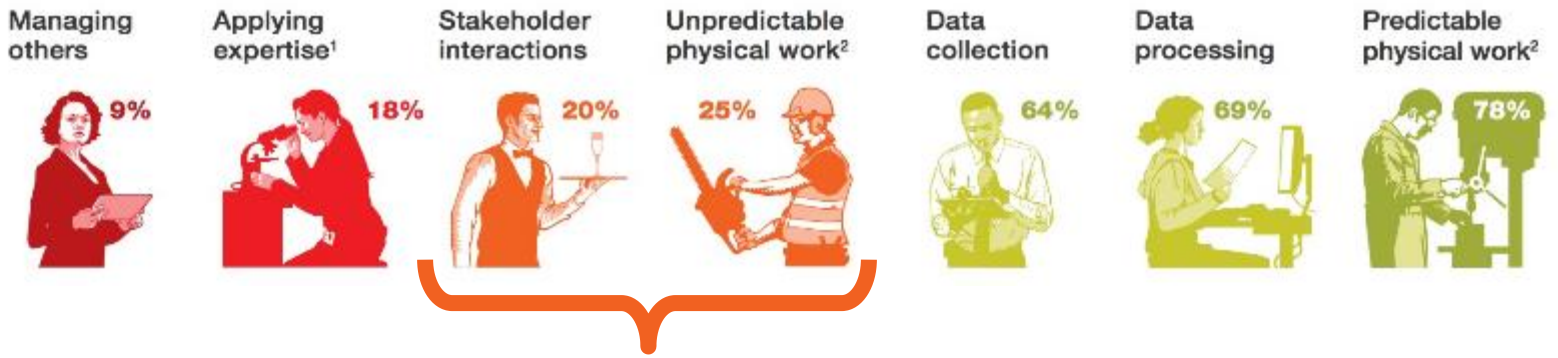

INTERACTION WITH THE ENVIRONMENT & HUMAN

UNDERSTANDING OF "WHAT'S GOING ON"
(beyond simple pattern matching)

FAST AND PERSONALIZED ACTIONS

AI assisted operator

Complex interaction with AI system
(cobot or conversational agent)
Scene understanding & real-time feedback
Incremental learning

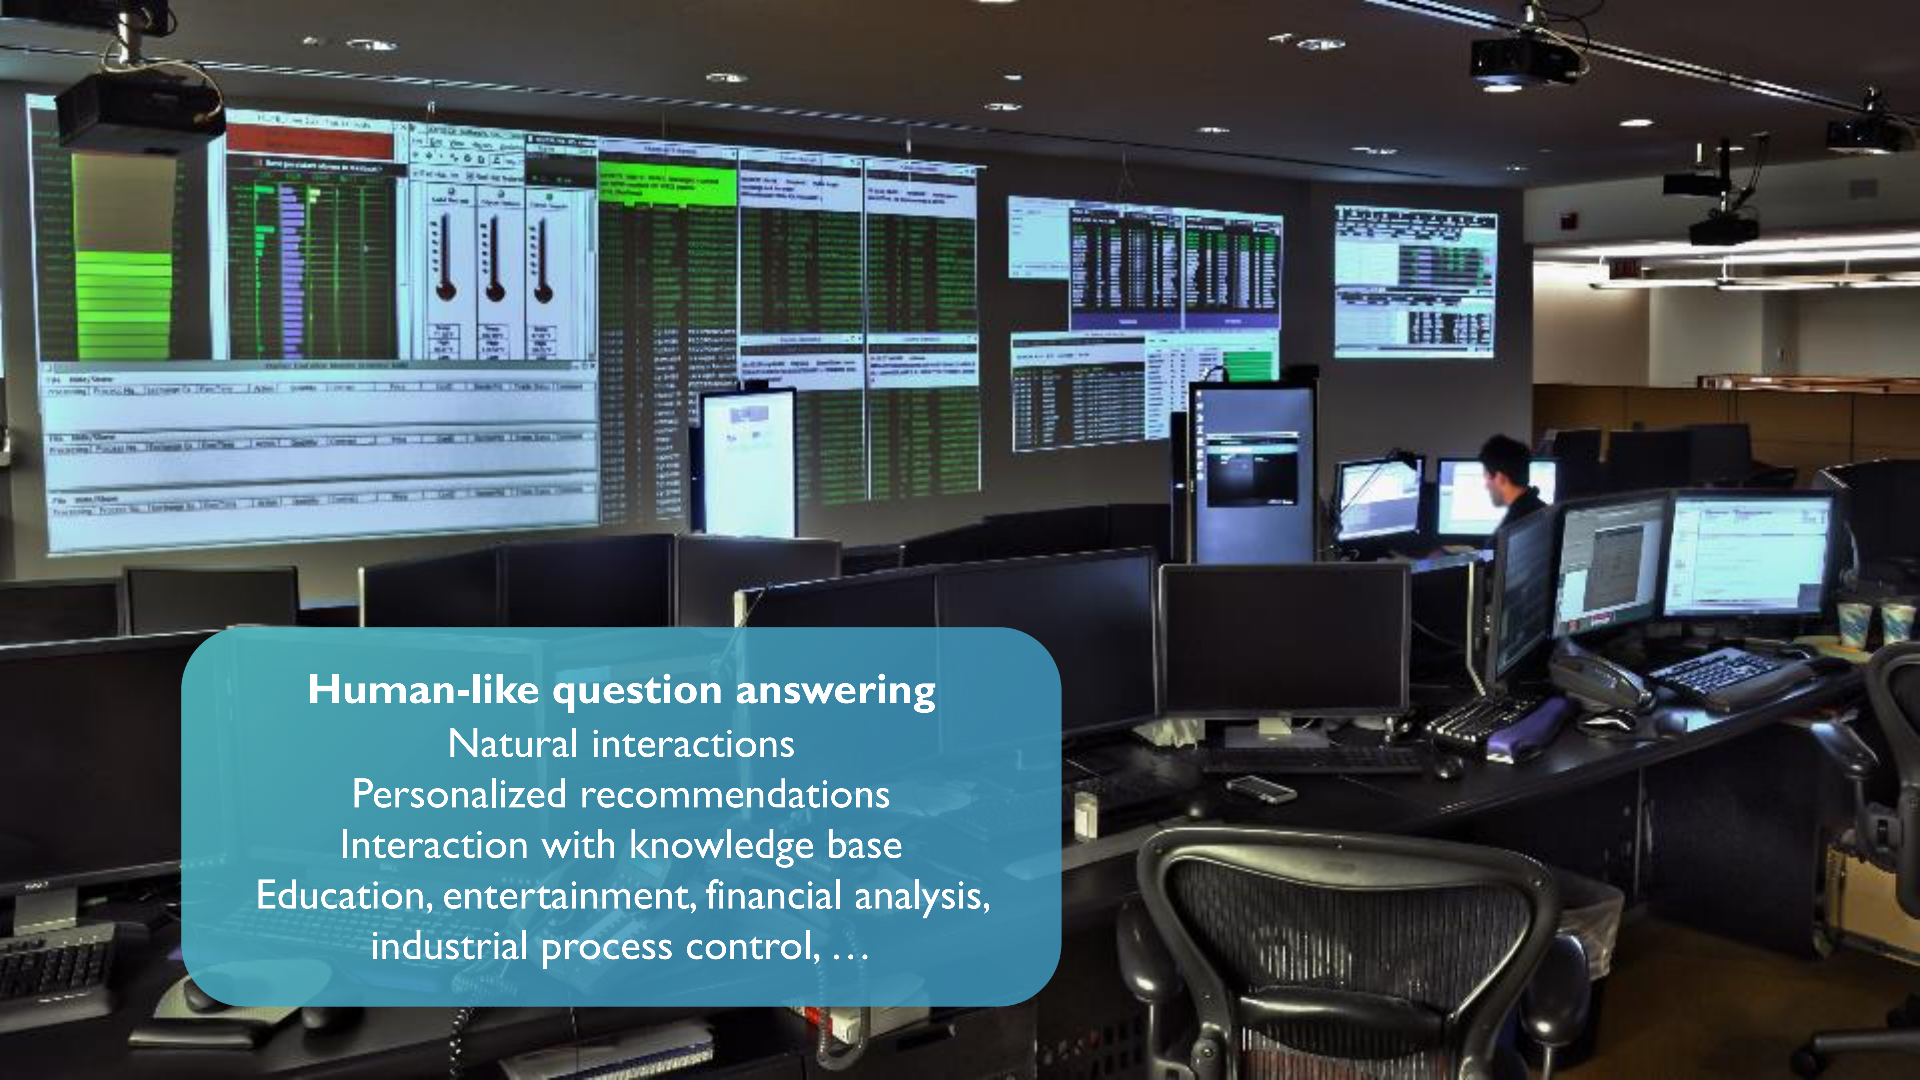

Human-like question answering

Natural interactions

Personalized recommendations

Interaction with knowledge base

Education, entertainment, financial analysis,
industrial process control, …


Understanding Intent
Gaze estimation
Visual Scene Understanding
Depth Estimation
All kinds of mobile actors
WILL TURN LEFT IN 2.3 SEC
DRIVES AT 25.5 KM/H
CMA CGM
ZIM
Hapag-Lloyd
YANG MING
SISU

Understanding of context & relations

Autonomous learning of tasks

System-wide optimizations and scheduling

Decision making & recommendations

FLANDERS AI PROGRAM STRUCTURE

4 GRAND CHALLENGES

HELP TO **MAKE COMPLEX DECISIONS**

70

6 M

EXTRACT AND **PROCESS** INFORMATION AT THE **EDGE**

30

2.6 M

INTERACT AUTONOMOUSLY WITH OTHER **DECISION-MAKING** ENTITIES

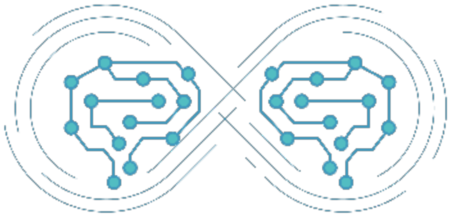

10

0.8 M

COMMUNICATE AND **COLLABORATE SEAMLESSLY** WITH HUMANS

30

2.6 M

AI PROGRAM FLANDERS

Data ethics and society 1 mio €/year
Training and education 3 mio €/year
Public Outreach 1 mio €/year

12 MILJOEN EURO	15 MILJOEN EURO	5 MILJOEN EURO
↓	↓	↓
TOP RESEARCH	DIGITIZATION AND IMPLEMENTATION IN INDUSTRY	ETHICS, EDUCATION AND TRAINING

Flanders AI Research