information management unit
School of Electrical and Computer Engineering
National Technical University of Athens
imu.ntua.gr

Human-AI Intelligence for Industry 5.0

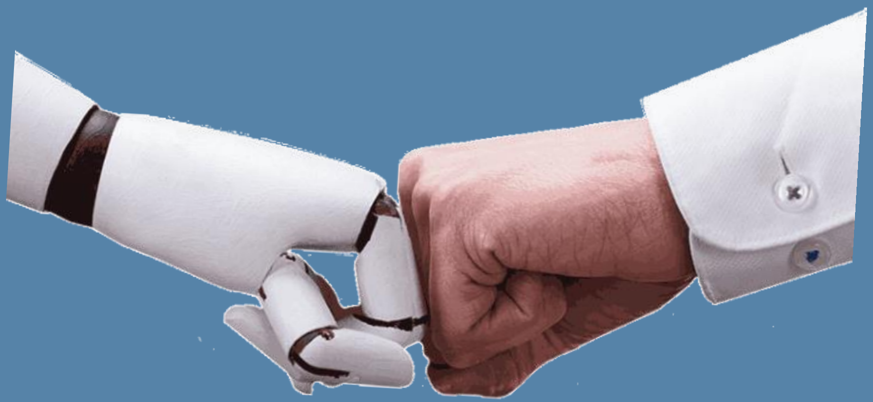

Prof. Gregoris Mentzas
School of Electrical and Computer Engineering
Institute of Communication and Computer Systems
National Technical University of Athens

Some words about me

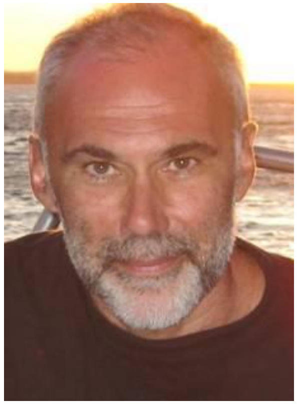

- Professor at the School of Electrical and Computer Engineering of the National Technical University of Athens (NTUA)
- Director of the Information Management Unit (IMU)
- Led or contributed in > 50 R&D projects
  - Attracted research funding > 16 M€
- Google Scholar estimates approx. >7,000 citations (h>40)
  - Research has led to 3 internet technology companies
  - Published 4 books and more than 350 papers
  - (co-)Chair or PC Member in > 60 international conferences
  - 5 best papers awards
  - Associate Editor of 5 international journals
- Always a learner (and a beginner)

IMU - Information Management Unit

IMU is a **multi-disciplinary unit** engaged in research and technology development activities in **Information Technology Management**.

IMU carries out research activities within ICCS of NTUA

Our mission is **to enable the development of knowledge-driven organisations**

53
Research Projects

16M€
in Funding

> 350
Publications

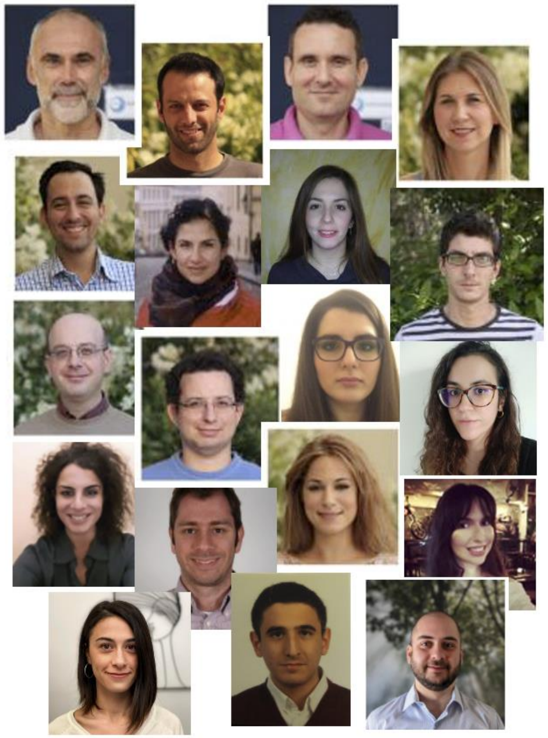

[AI, Data and Robotics Association (DAIRO)]

[Gaia-X Association for Data & Cloud]

[International Data Spaces Association (IDSA)]

[ETP Software and Services]

[Made in Europe (formerly EFFRA)]

[Alliance for Internet of Things]

Research areas

Analyse

from Data to Insight

We mine and combine data and discover useful knowledge by using AI, machine learning and prescriptive data analytics methods

Compute

from Cloud to Edge

We address issues like resource allocation across hybrid environments (cloud, fog, edge) and context-aware security and privacy

Decide

from **Intelligence to Prescription**

We develop multi-criteria optimisation and collective intelligence methods that proactively recommend the appropriate decisions

Application domains

Factories of the Future

Electronic Governance

Digital health & well-being

Digital Innovation

Sustainable & Green Economy

Outline

- What is Human-centered Artificial Intelligence?
- Our approach: Human-augmented Analytics
- Industry 5.0 Business Cases
- Conclusions and further work

What is an AI System?

- An AI system is a computer-based system that is capable of influencing the environment by producing an actionable output (predictions, recommendations or decisions) for a given set of objectives.
- It uses machine data and/or **human** inputs to:
  - **perceive** real and/or virtual environments;
  - **abstract** these perceptions into models through an automated or manual manner; and
  - use model inference to formulate options for outcomes..

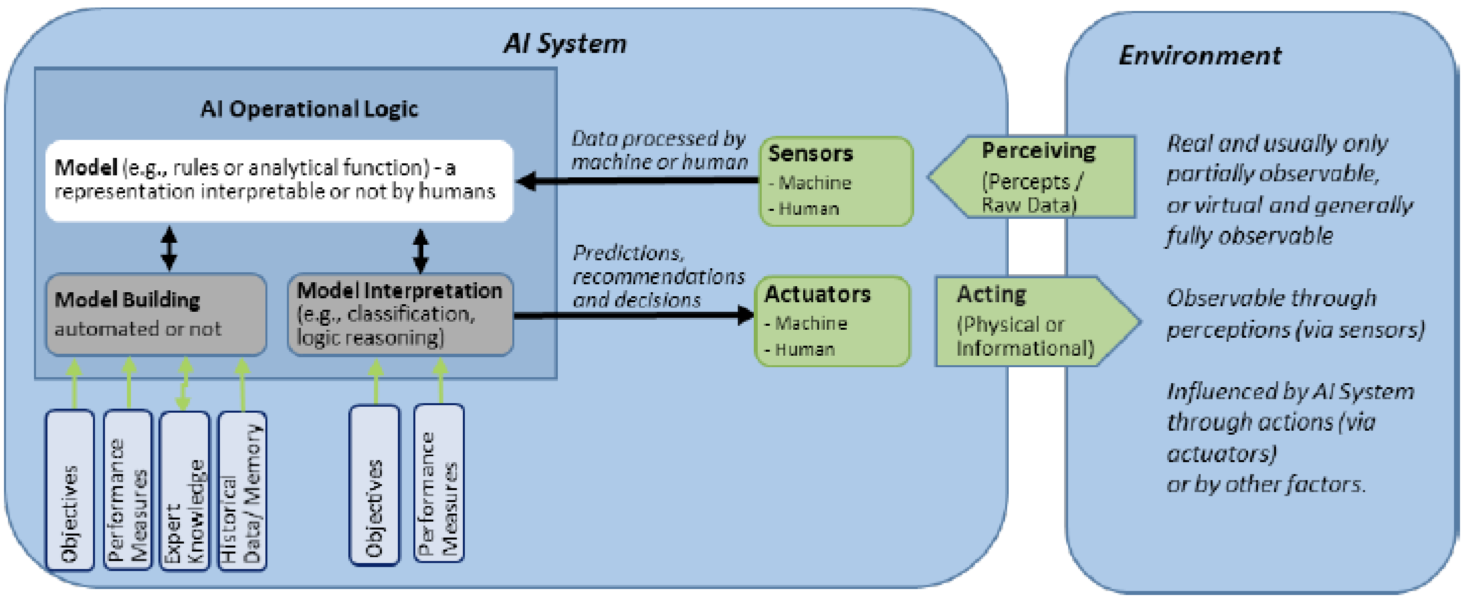

Organisation for Economic Co-operation and Development. (2019). Scoping the OECD AI Principles: Deliberations of the Expert Group on Artificial Intelligence at the OECD (AIGO). OECD Publishing.

What is Human-Centered AI?

- Human-Centered AI (HCAI) is an emerging discipline intent on creating AI systems that **amplify and augment rather than displace human abilities**.
  - HCAI seeks to preserve human control in a way that ensures AI meets our needs
  - while also operating transparently, delivering equitable outcomes, and respecting privacy
- The HCAI strategy is to rigorously investigate and design new forms of human-AI interactions and experiences
  - that enhance & extend human capabilities for the good of products, clients, and society.

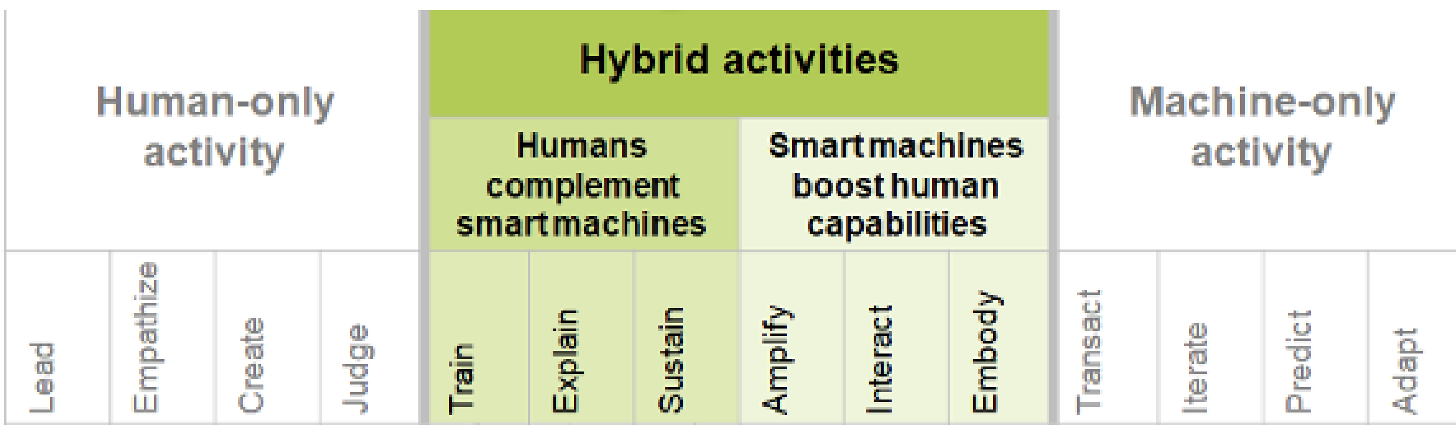

Daugherty, P. R., & Wilson, H. J. (2018). Human + machine: Reimagining work in the age of AI. Harvard Business Press.

Human-AI decision making with augmented intelligence

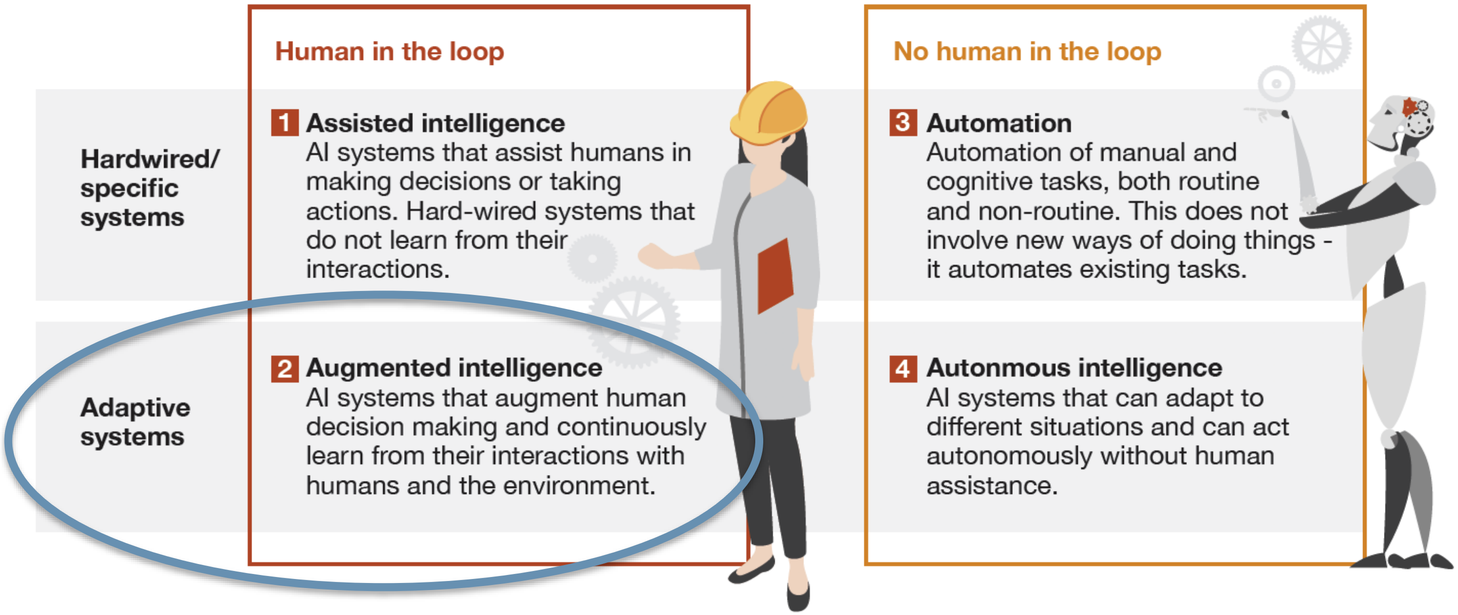

Source: PWC (2017) An introduction to implementing AI in manufacturing, Global Manufacturing and Industrialization Summit.

National Science Foundation

Where Artificial Intelligence typically entails using code and algorithms to replace human reasoning, **Augmented Intelligence** entails using knowledge to <u>empower</u> human reasoning… the science of Augmented Intelligence puts human benefit at center stage.

Gartner's Hype Cycle for AI , 2021

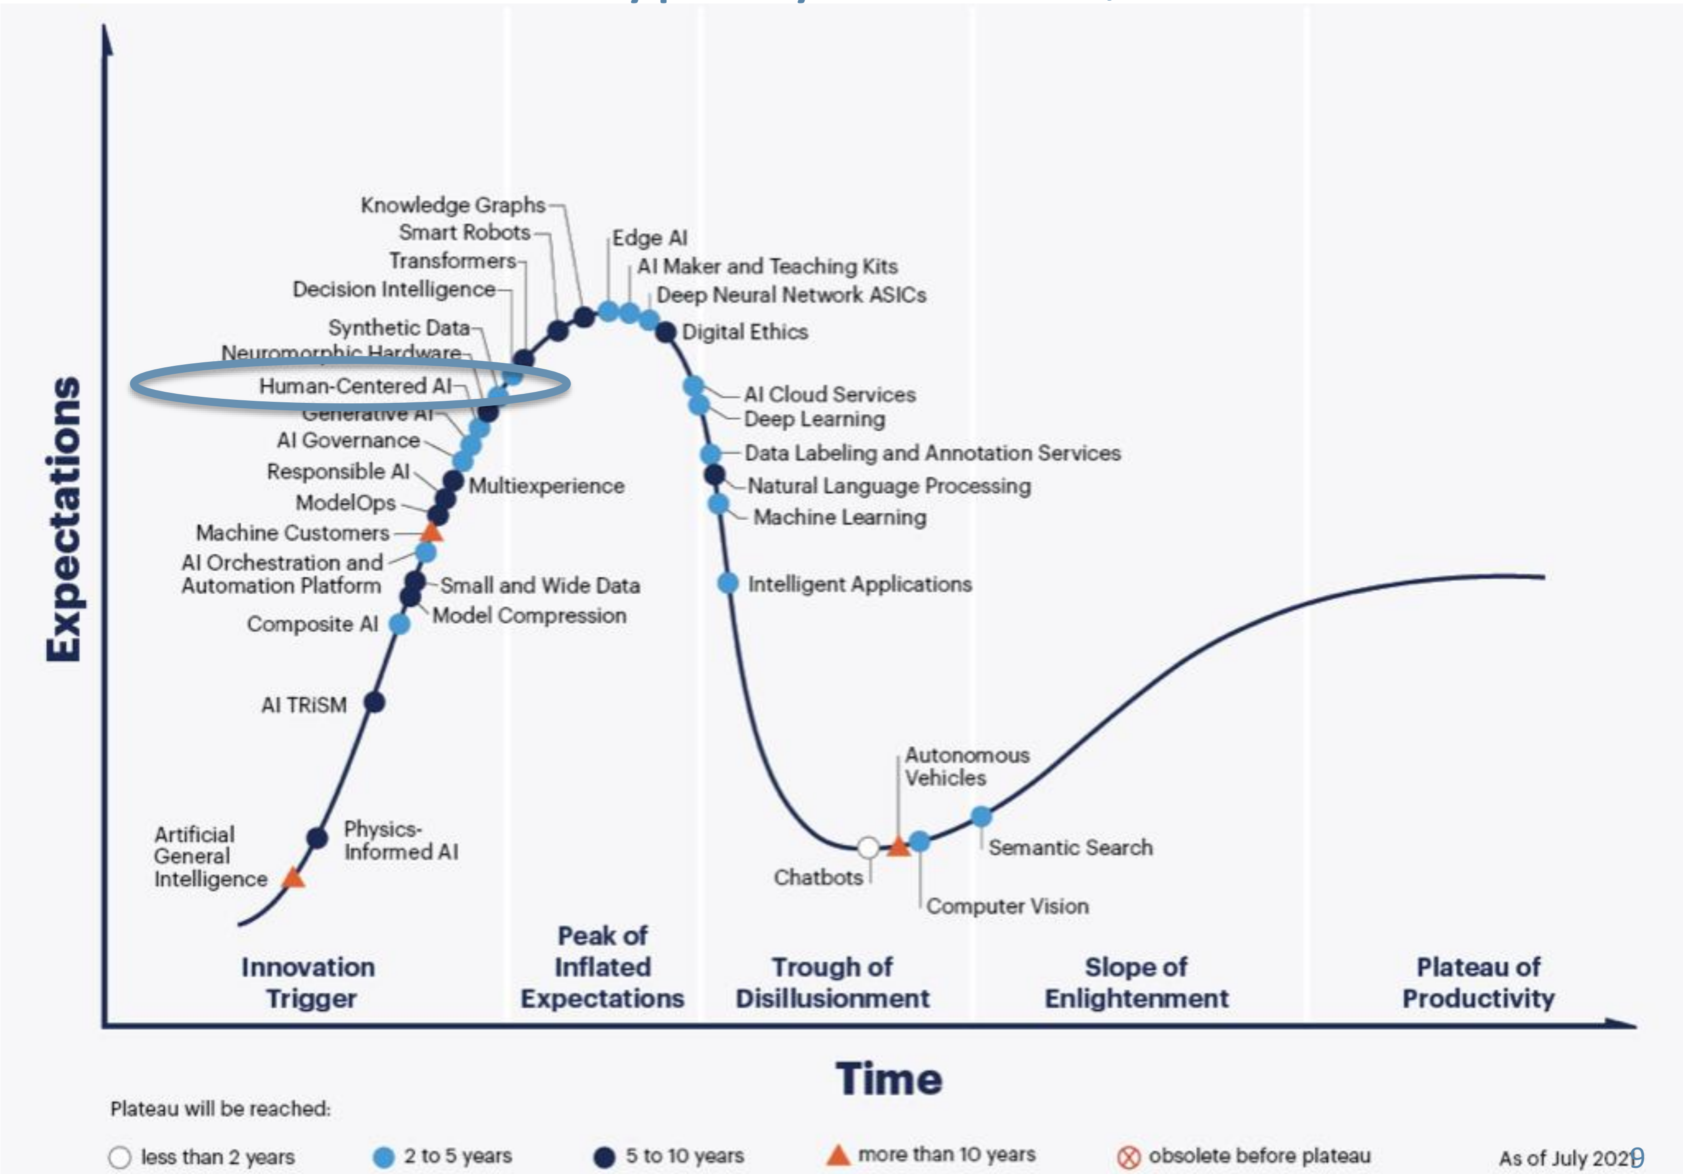

Outline

- What is Human-centered Artificial Intelligence?
- Our approach: Human-augmented Analytics
- Industry 5.0 Business Cases
- Conclusions and further work

Analytics Value Escalator

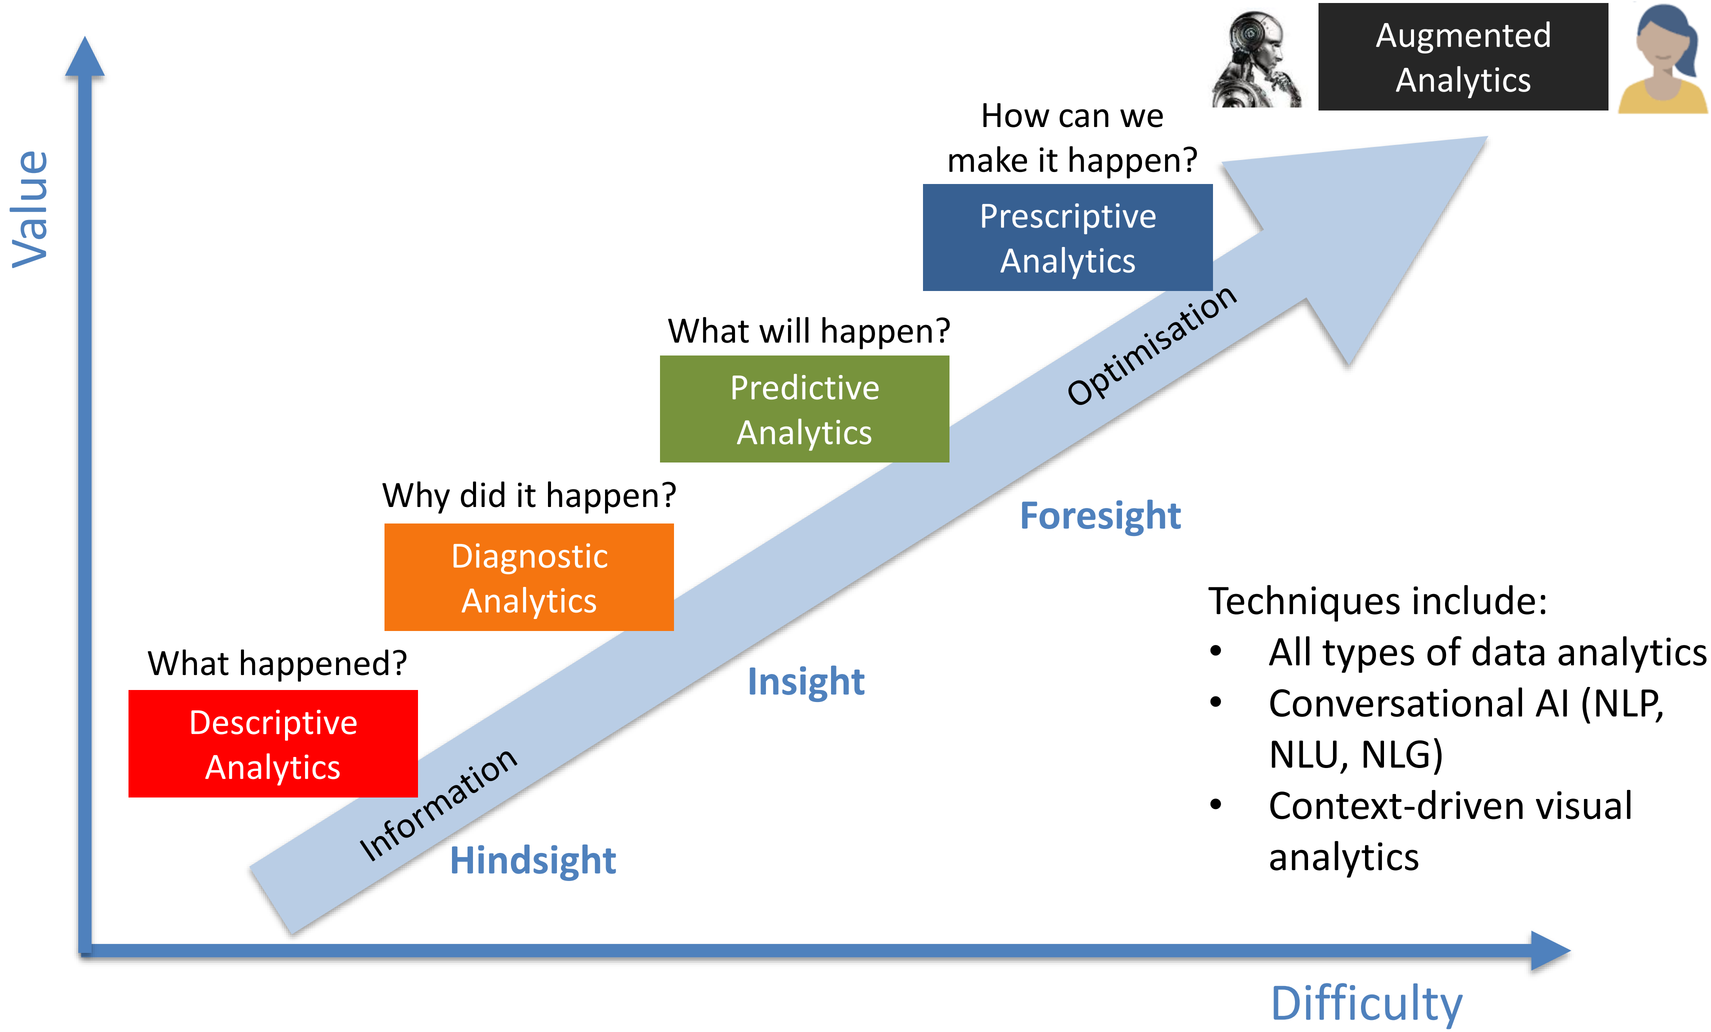

Techniques include:

- All types of data analytics
- Conversational AI (NLP, NLU, NLG)
- Context-driven visual analytics

J

Lepenioti, K., Bousdekis, A., Apostolou, D., & Mentzas, G. (2020). Prescriptive analytics: Literature review and research challenges. International Journal of Information Management, 50, 57-70.

Our approach: Human-augmented analytics

- Human-augmented data analytics is the use of enabling AI technologies to assist with data preparation, insight generation and insight explanation to **augment how humans explore and analyze data**.

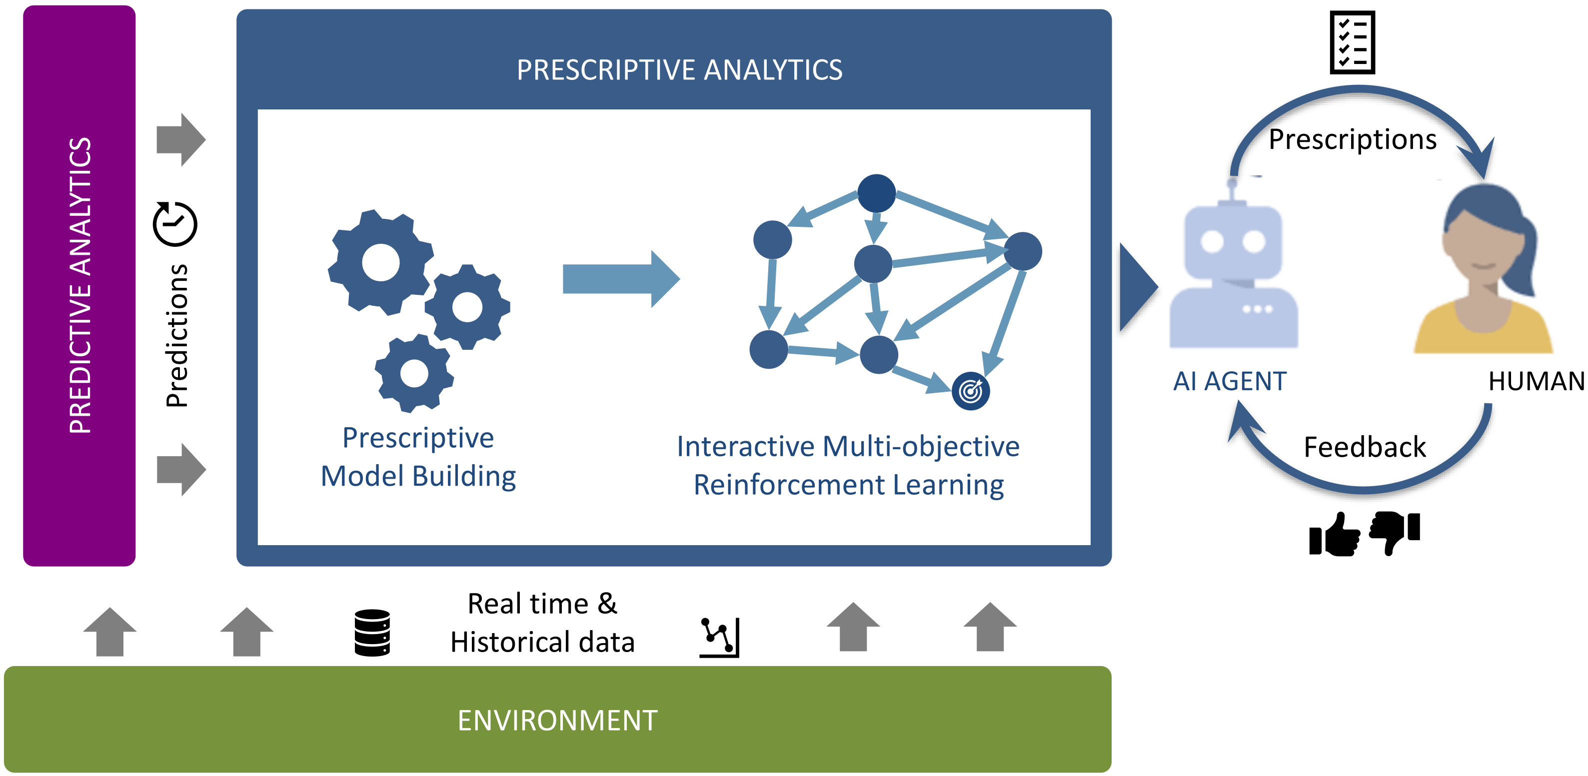

C Mentzas, G., Lepenioti, K., Bousdekis, A., & Apostolou, D. (2021, September). Data-Driven Collaborative Human-AI Decision Making. In Conference on e-Business, e-Services and e-Society (pp. 120-131). Springer, Cham.

Framework for human-augmented analytics

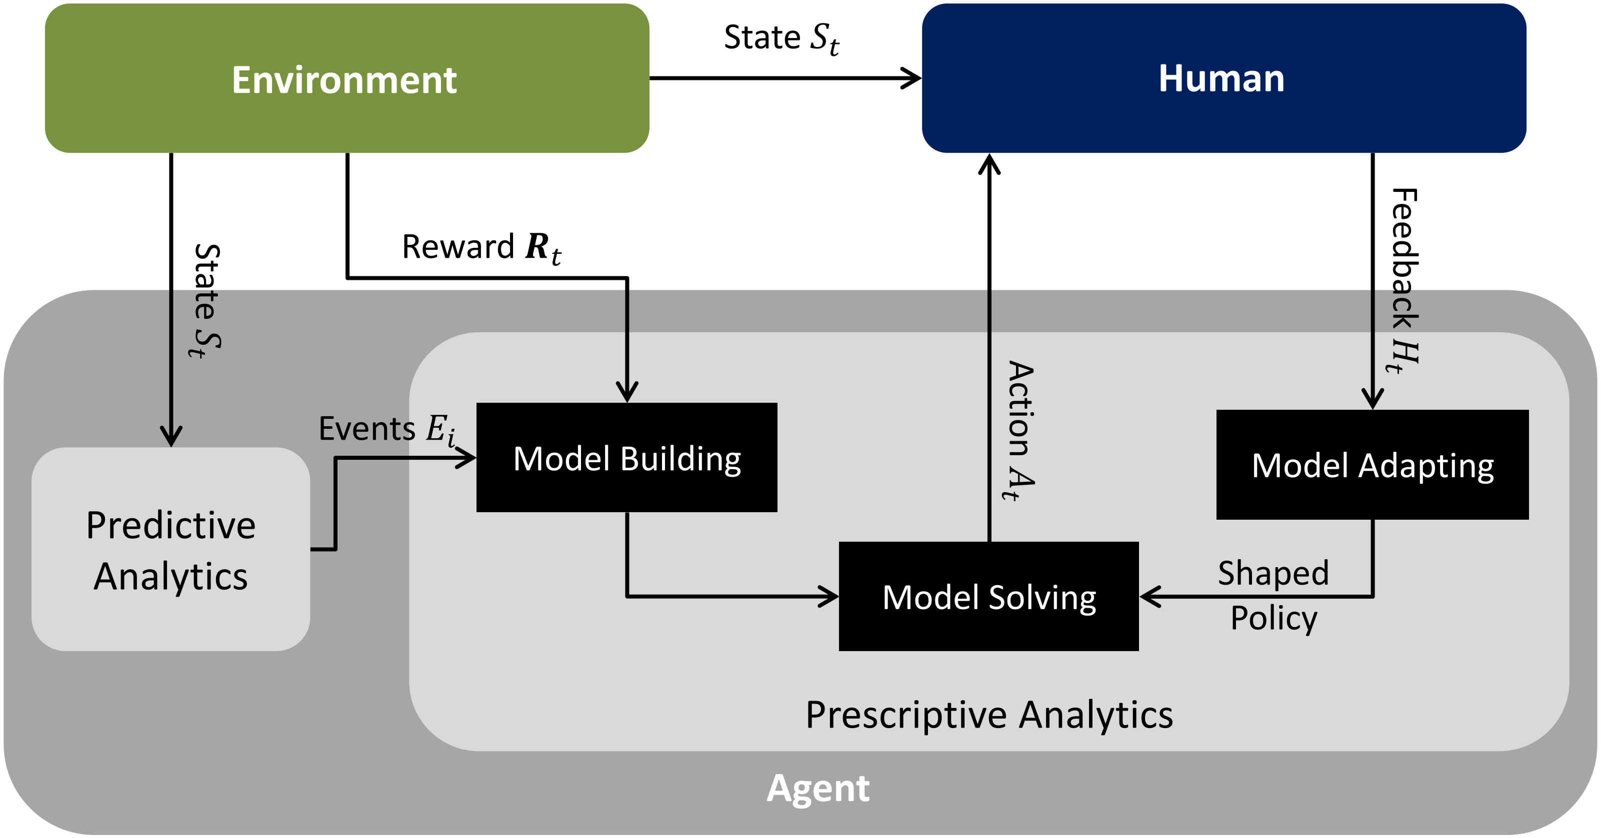

Lepenioti, K., Bousdekis, A., Apostolou, D., & Mentzas, G. (2021). Human-augmented prescriptive analytics with interactive multi-objective reinforcement learning. IEEE Access, 9, 100677-100693.

Human-augmented analytics – 1 of 3

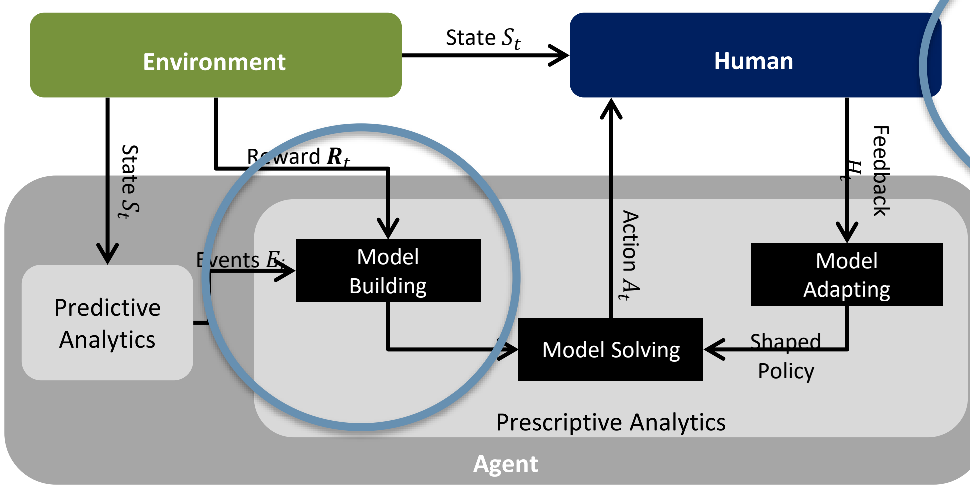

1. Define current state, objectives and desired and undesired future events $\boldsymbol{E}$, $O_i, i \in I$, $\{A_{m,j}\}, m \in J, J \in \mathbb{N}$

2. Observe objectives from historical data

$$\boldsymbol{O}_i(t) = (o_1^i, o_2^i, \dots, o_n^i), \qquad n \in \mathbb{N}$$

3. Receive predictions for upcoming future events

$$Pred = Pr((E_i, \dots, E_j, \dots) | (E_1, E_2, E_3, \dots, E_{i-1}))$$

4. Prescriptive Model Building

(i) Define state space $\{S_i\}, i \in I$, action space $\{a_j\}, j \in J$, starting state S_N, goal state S_G and optional undesired future state S_F that will constitute the MDP model

(ii) Use historical objectives' observations $\boldsymbol{O}_i(t)$, $n \in \mathbb{N}$ to define the reward function $\boldsymbol{R} : S \times A \times S \to \mathbb{R}_n$

5. Prescriptive Model Solving

(i) Calculate optimal policy $\pi_{O_i}(s, a)$ for each **objective** $O_i, i \in I$, with an actor-critic algorithm

(ii) Calculate the multi-objective optimal policy $\pi_{MOopt}(s, a)$, given that the objectives are independent to each other:

$$\pi_{MOopt}(s, a) = \prod_{i \in I} \pi_{O_i}(s, a)$$

6. Prescriptive Model Adapting

(i) Calculate the feedback policy $\pi_{Fopt}(s, a)$ with given feedback:

For each RL action calculate the probability:

$$Pr_H(a) = C^{\Delta_{s,a}} + (1 - C)^{\sum_{j \neq a} \Delta_{s,j}}$$

(ii) Shape multi-objective policy:

$$\pi_{opt}(s, a) = \pi_{MOopt}(s, a) * \pi_{Fopt}(s, a)$$

7. Return prescription $Presc = (A_1, A_2, \dots, A_f)$

8. Collect feedback for the generated prescription

$$H_{Presc} = \begin{cases} H^- \in \mathbb{R}^-, \textit{negative feedback} \\ 0, \textit{neutral} \\ H^+ \in \mathbb{R}^+, \textit{positive feedback} \end{cases}$$

Human-augmented analytics – 2 of 3

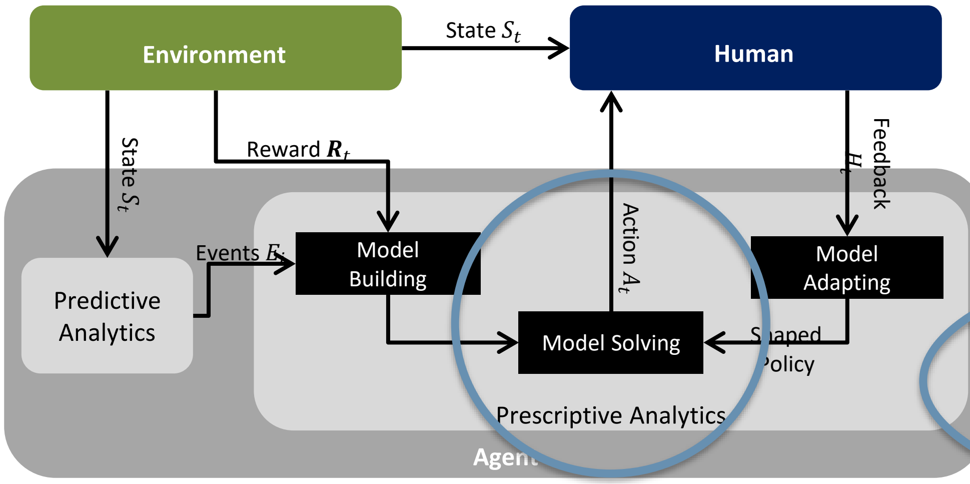

1. Define current state, objectives and desired and undesired future events E, $O_i, i \in I$, $\{A_m\}, m \in J, J \in \mathbb{N}$

2. Observe objectives from historical data

$$\boldsymbol{O}_i(t) = (o_1^i, o_2^i, \dots, o_n^i), \quad n \in \mathbb{N}$$

3. Receive predictions for upcoming future events

$$Pred = Pr((E_i, \dots, E_j, \dots)|(E_1, E_2, E_3, \dots, E_{i-1}))$$

4. Prescriptive Model Building

(i) Define state space $\{S_i\}, i \in I$, action space $\{a_j\}, j \in J$, starting state S_N, goal state S_G and optional undesired future state S_F that will constitute the MDP model

(ii) Use historical objectives' observations $\boldsymbol{O}_i(t)$, $n \in \mathbb{N}$ to define the reward function $\boldsymbol{R} : S \times A \times S \to \mathbb{R}_n$

5. Prescriptive Model Solving

(i) Calculate optimal policy $\pi_{O_i}(s, a)$ for each **objective** $O_i, i \in I$, with an actor-critic algorithm

(ii) Calculate the multi-objective optimal policy $\pi_{MOopt}(s, a)$, given that the objectives are independent to each other:

$$\pi_{MOopt}(s, a) = \prod_{i \in I} \pi_{O_i}(s, a)$$

6. Prescriptive Model Adapting

(i) Calculate the feedback policy $\pi_{Fopt}(s, a)$ with given feedback:

For each RL action calculate the probability:

$$Pr_H(a) = C^{\Delta_{s,a}} + (1 - C)^{\sum_{j \neq a} \Delta_{s,j}}$$

(ii) Shape multi-objective policy:

$$\pi_{opt}(s, a) = \pi_{MOopt}(s, a) * \pi_{Fopt}(s, a)$$

7. Return prescription $Presc = (A_1, A_2, \dots, A_f)$

8. Collect feedback for the generated prescription

$$H_{Presc} = \begin{cases} H^- \in \mathbb{R}^-, \textit{negative feedback} \\ 0, \textit{neutral} \\ H^+ \in \mathbb{R}^+, \textit{positive feedback} \end{cases}$$

Human-augmented analytics – 3 of 3

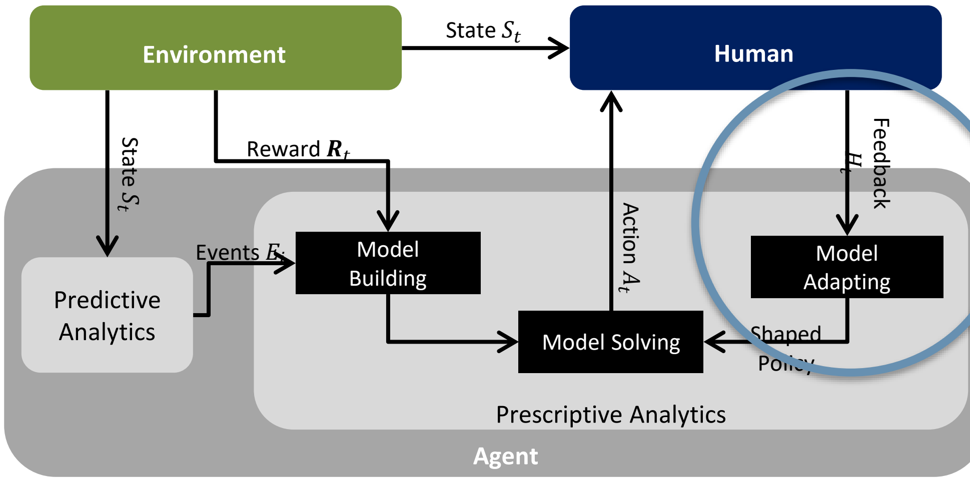

1. Define current state, objectives and desired and undesired future events E, $O_i, i \in I$, $\{A_m\}, m \in J, J \in \mathbb{N}$

2. Observe objectives from historical data

$$\boldsymbol{O}_i(t) = (o_1^i, o_2^i, \dots, o_n^i), \qquad n \in \mathbb{N}$$

3. Receive predictions for upcoming future events

$$Pred = Pr((E_i, \dots, E_j, \dots)|(E_1, E_2, E_3, \dots, E_{i-1}))$$

4. Prescriptive Model Building

(i) Define state space $\{S_i\}, i \in I$, action space $\{a_j\}, j \in J$, starting state S_N, goal state S_G and optional undesired future state S_F that will constitute the MDP model

(ii) Use historical objectives' observations $\boldsymbol{O}_i(t)$, $n \in \mathbb{N}$ to define the reward function $\boldsymbol{R} : S \times A \times S \to \mathbb{R}_n$

5. Prescriptive Model Solving

(i) Calculate optimal policy $\pi_{O_i}(s, a)$ for each **objective** $O_i, i \in I$, with an actor-critic algorithm

(ii) Calculate the multi-objective optimal policy $\pi_{MOopt}(s, a)$, given that the objectives are independent to each other:

$$\pi_{MOopt}(s, a) = \prod_{i \in I} \pi_{O_i}(s, a)$$

6. Prescriptive Model Adapting

(i) Calculate the feedback policy $\pi_{Fopt}(s, a)$ with given feedback:

For each RL action calculate the probability:

$$Pr_H(a) = C^{\Delta_{s,a}} + (1 - C)^{\sum_{j \neq a} \Delta_{s,j}}$$

(ii) Shape multi-objective policy:

$$\pi_{opt}(s, a) = \pi_{MOopt}(s, a) * \pi_{Fopt}(s, a)$$

7. Return prescription $Presc = (A_1, A_2, \dots, A_f)$

8. Collect feedback for the generated prescription

$$H_{Presc} = \begin{cases} H^- \in \mathbb{R}^-, \textit{negative feedback} \\ 0, \textit{neutral} \\ H^+ \in \mathbb{R}^+, \textit{positive feedback} \end{cases}$$

Outline

- What is Human-centered Artificial Intelligence?
- Our approach: Human-augmented Analytics
- Industry 5.0 Business Cases
- Conclusions and further work

Society 5.0 is…

CHART 2
Society 5.0

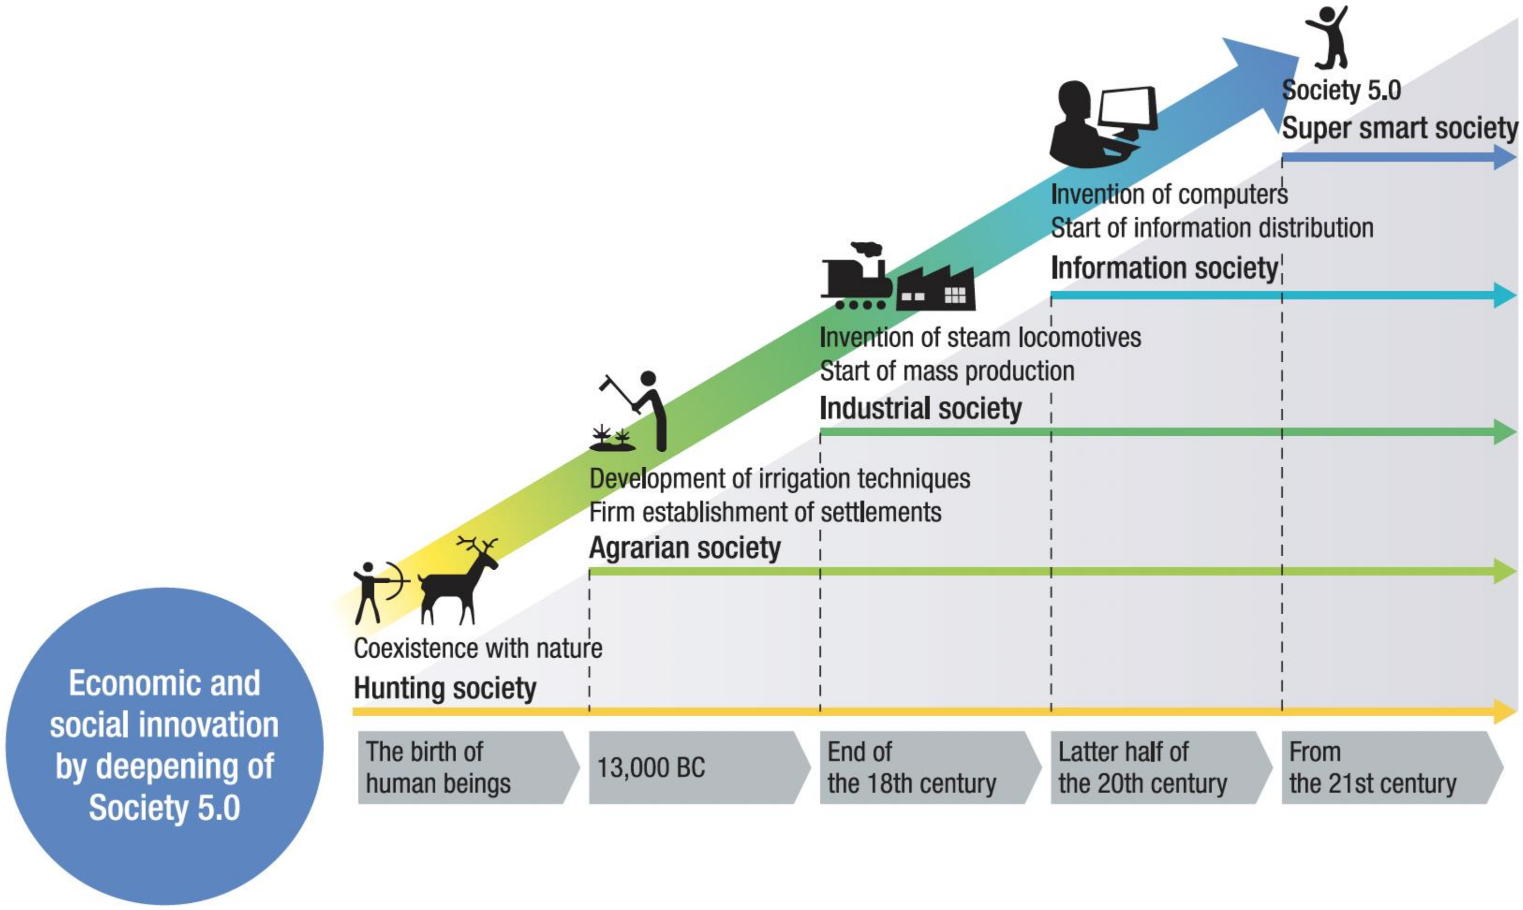

Fukuyama, M. (2018). Society 5.0: Aiming for a new human-centered society. Japan Spotlight, 27(Society 5.0), 47-50.

Industry 5.0 is…

European Commission Directorate-General for Research and Innovation (2022) Industry 5.0, a transformative vision for Europe: Governing systemic transformations towards a sustainable industry

Case in Point: Predictive maintenance (1 of 2)

Estimate Remaining Useful Life (RUL)
Predict future failure modes

Feature extraction
Online Bayesian changepoint detection
Weibull fitting
Long Short-Term Memory (LSTM)

Prescribe the optimal proactive actions
Provide explainable visualization

Markov Decision Process (MDP)
Reinforcement Learning
Multi-Objective Optimization

PREDICTIVE ANALYTICS

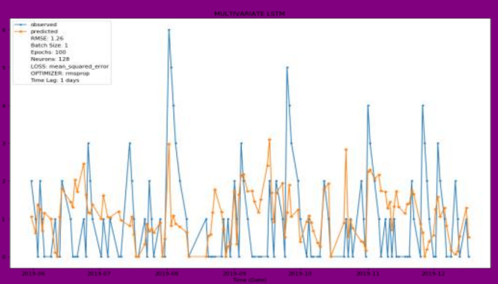

PRESCRIPTIVE ANALYTICS

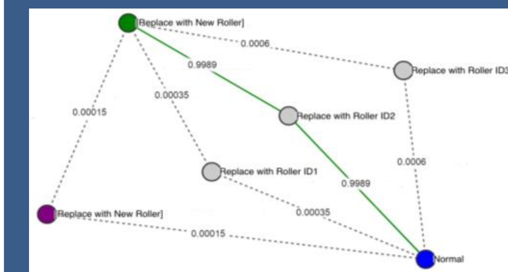

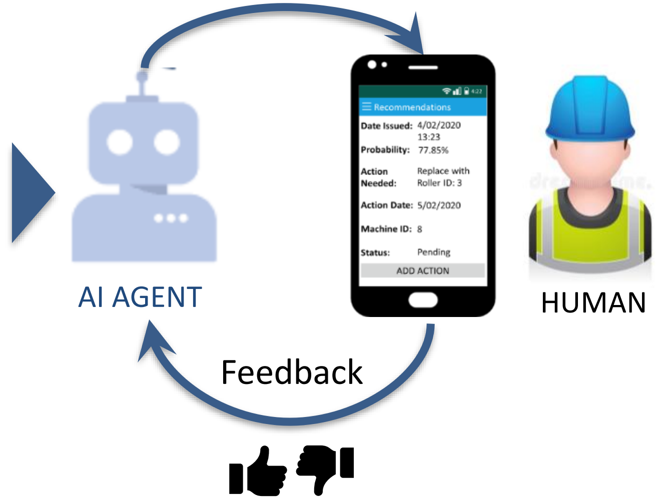

Real time &
Historical data

ENVIRONMENT

Real-time data
Accelerometers
Tachometer

Enterprise data
CMMS
ERP

– Overall aim to integrate predictive and prescriptive analytics using real-time sensor and operational data for shop-floor predictive maintenance

J

Wellsandt, S., K. Klein, K. Hribernik, M. Lewandowski, A. Bousdekis, G. Mentzas, K.-D. Thoben (2022) Hybrid-augmented intelligence in predictive maintenance with digital intelligent assistants, Annual Reviews in Control, 2022.

Case in Point: Predictive maintenance (2 of 2)

- **Data**
- Years 2017, 2018, 2019
- **4-layer LSTM**
- *(When do we expect the next interruption?)*
  - 21 features (e.g. real operational time, time of interruptions, duration of the breakdown events)
- *(Which is the expected interruption duration for the following day?)*
  - Features: Availability, Performance, Minutes of Breakdown, Real Gross Production, Number of breakdowns, date
- **Multi-Objective RL**
- Actor-critic algorithm consisting of a Boltzmann actor and a TD-Lambda critic

- Criticality of maintenance:
  - failures and delays consume 31.1% of operational time
  - Change of Work Rolls and Backup Rolls, related to the cold rolling mill hold a significant percentage of operational time

Actions	Cost (Euro)	RUL (days)
Replace with new roller	8,000	12
Replace with repaired roller ID1	5,284	5
Replace with repaired roller ID2	3,510	4
Replace with repaired roller ID3	4,831	6

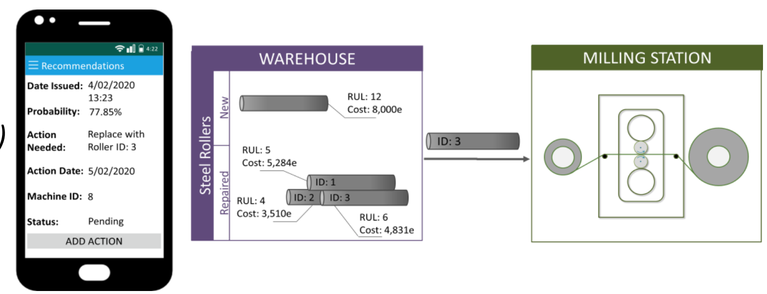

C Lepenioti, K., Pertselakis, M., Bousdekis, A., Louca, A., Lampathaki, F., Apostolou, D., G. Mentzas, & Anastasiou, S. (2020). Machine learning for predictive and prescriptive analytics of operational data in smart manufacturing. In International Conference on Advanced Information Systems Engineering (pp. 5-16). Springer.

Case in Point: Stock Trading (1 of 2)

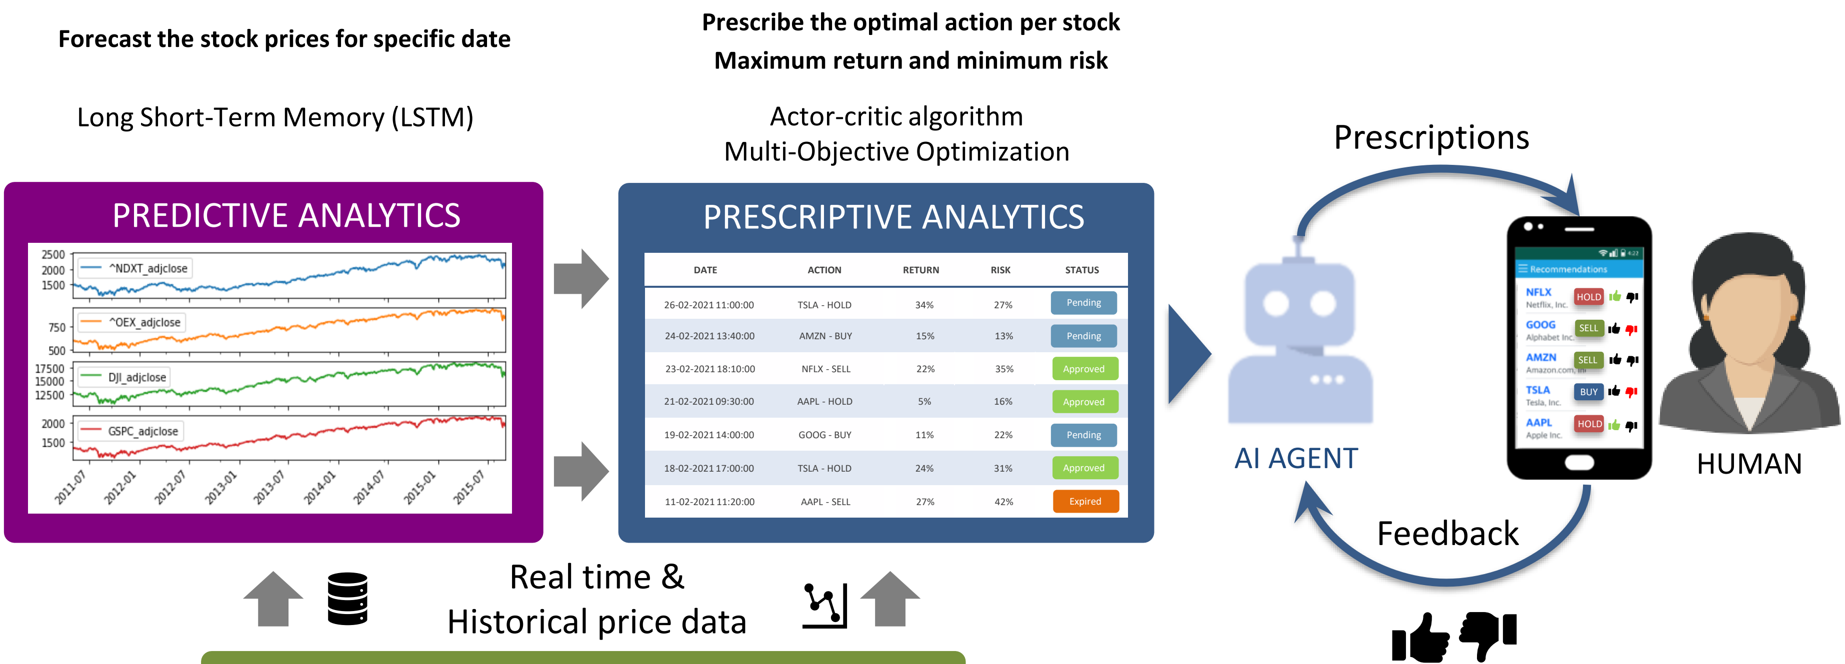

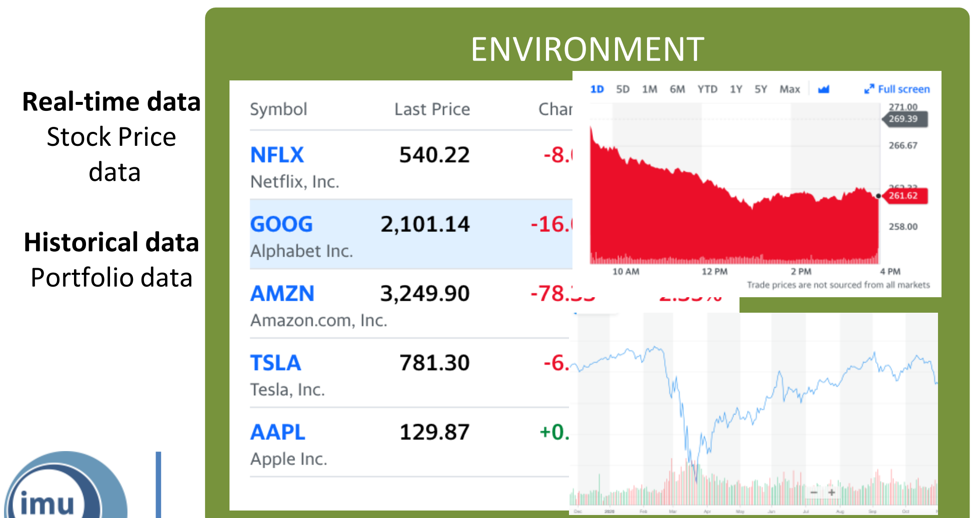

– Overall aim: to integrate stock trading predictions into the decision-making process of the portfolio construction problem

C

Karlis, V., Lepenioti, K., Bousdekis, A., & Mentzas, G. (2021, July). Stock Trend Prediction by Fusing Prices and Indices with LSTM Neural Networks. In 2021 12th IEEE International Conference on Information, Intelligence, Systems & Applications (IISA) (pp. 1-7).

Case in Point: Stock Trading (2 of 2)

- Aim to **assist the portfolio manager** with the capability of formulating the portfolio management problem according to their preferences
  - it utilizes the insight from both the trading expert and the available data
  - provides trading decisions that will lead to maximum return and the minimum risk

- Public dataset available at Kaggle
  - We focused on Apple stock closing price
  - Daily price data (07-07-2004 to 31-01-2008)
- We decouple risk and return into two different reward functions
  - Able to estimate the risk involved in a trading action by capitalizing both on historical and predicted data
- **Negative feedback evolution** as an evaluation metric across six different user profiles

Predicted Event	E_i	Future price after X days, $price^*_{t+X}$
Actions	A_1	Buy shares of the stock, if the stock is not in the current portfolio
	A_2	Sell the shares in possession on the predicted date
	A_3	Hold the shares in possession, until the predicted date or later
Objectives	O_1	Return R_{A_i}
	O_2	Expected Risk $Risk_{A_i}$

	Beginner	Intermediate	Experienced
Risk-averse			
Min	71.4%	30.7%	12.9%
Max	100%	100%	75%
Convergence point	88.9%	31.5%	13%
Risk-seeker			
Min	33.3%	37.7%	21.3%
Max	83.8%	66.6%	100%
Convergence point	83.2%	38.1%	21.7%

Case in Point: Quality Control (1 of 3)

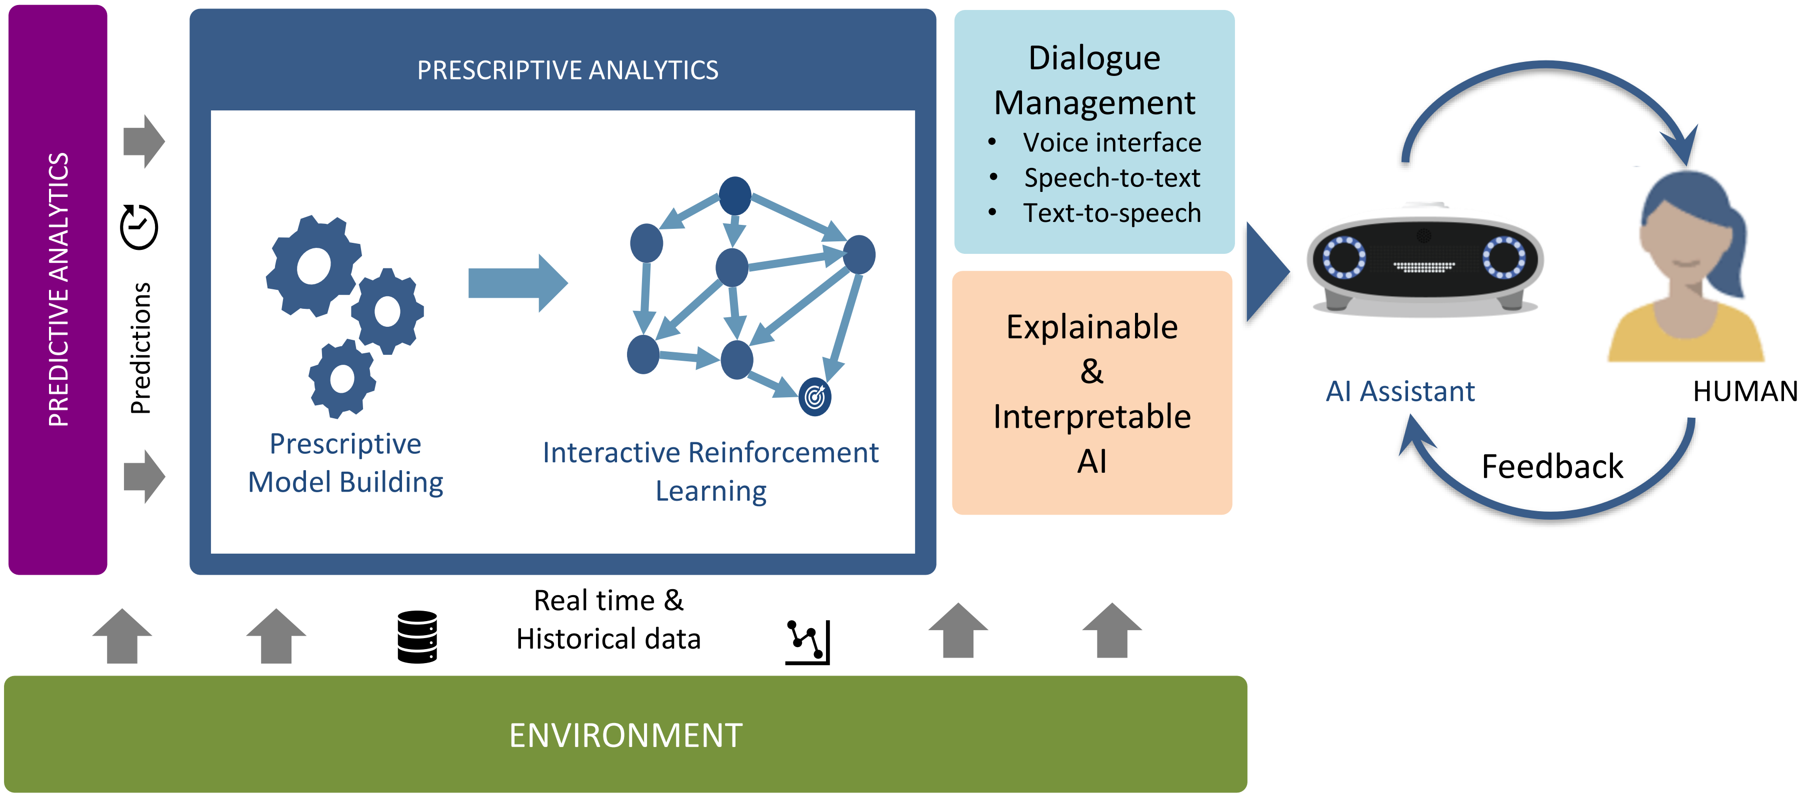

C

Bousdekis, A., Wellsandt, S., Bosani, E., Lepenioti, K., Apostolou, D., Hribernik, K., & Mentzas, G. (2021, September). Human-AI Collaboration in Quality Control with Augmented Manufacturing Analytics. In IFIP International Conference on Advances in Production Management Systems (pp. 303-310). Springer.

Case in Point: Quality Control (2 of 3)

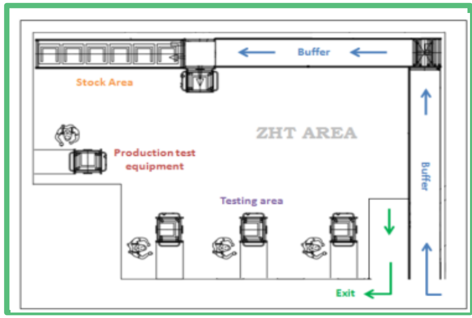

- Aims to adopt a predictive quality strategy
  - Link the quality control of the finished product with design stage and shop floor
  - Reconfigure its production line and facilitate root cause analysis

- **Augmented Manufacturing analytics**
  - Provide novel prescriptive quality analytics service
  - Simplify access to analytics through a conversational interface
  - Explain prescriptions to improve trustworthiness

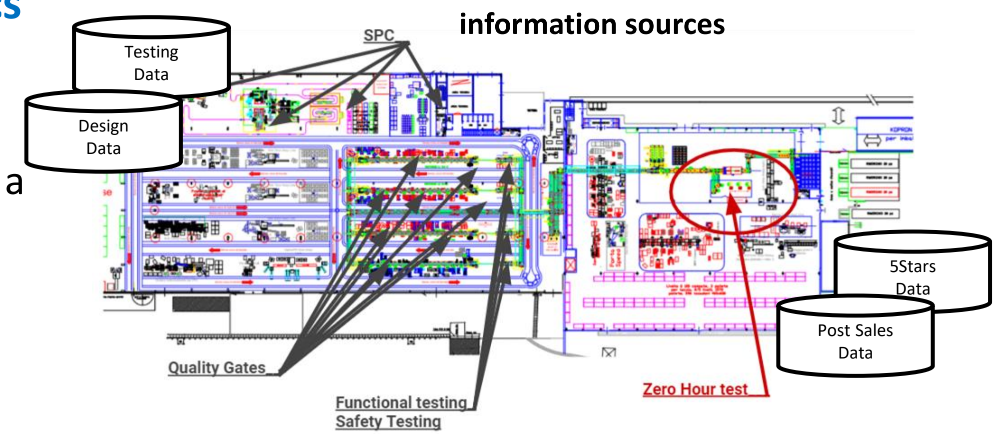

Case in Point: Quality Control (3 of 3)

- Enables the development of several independent analytics processes with different scope (descriptive, predictive, prescriptive) from multiple data sources

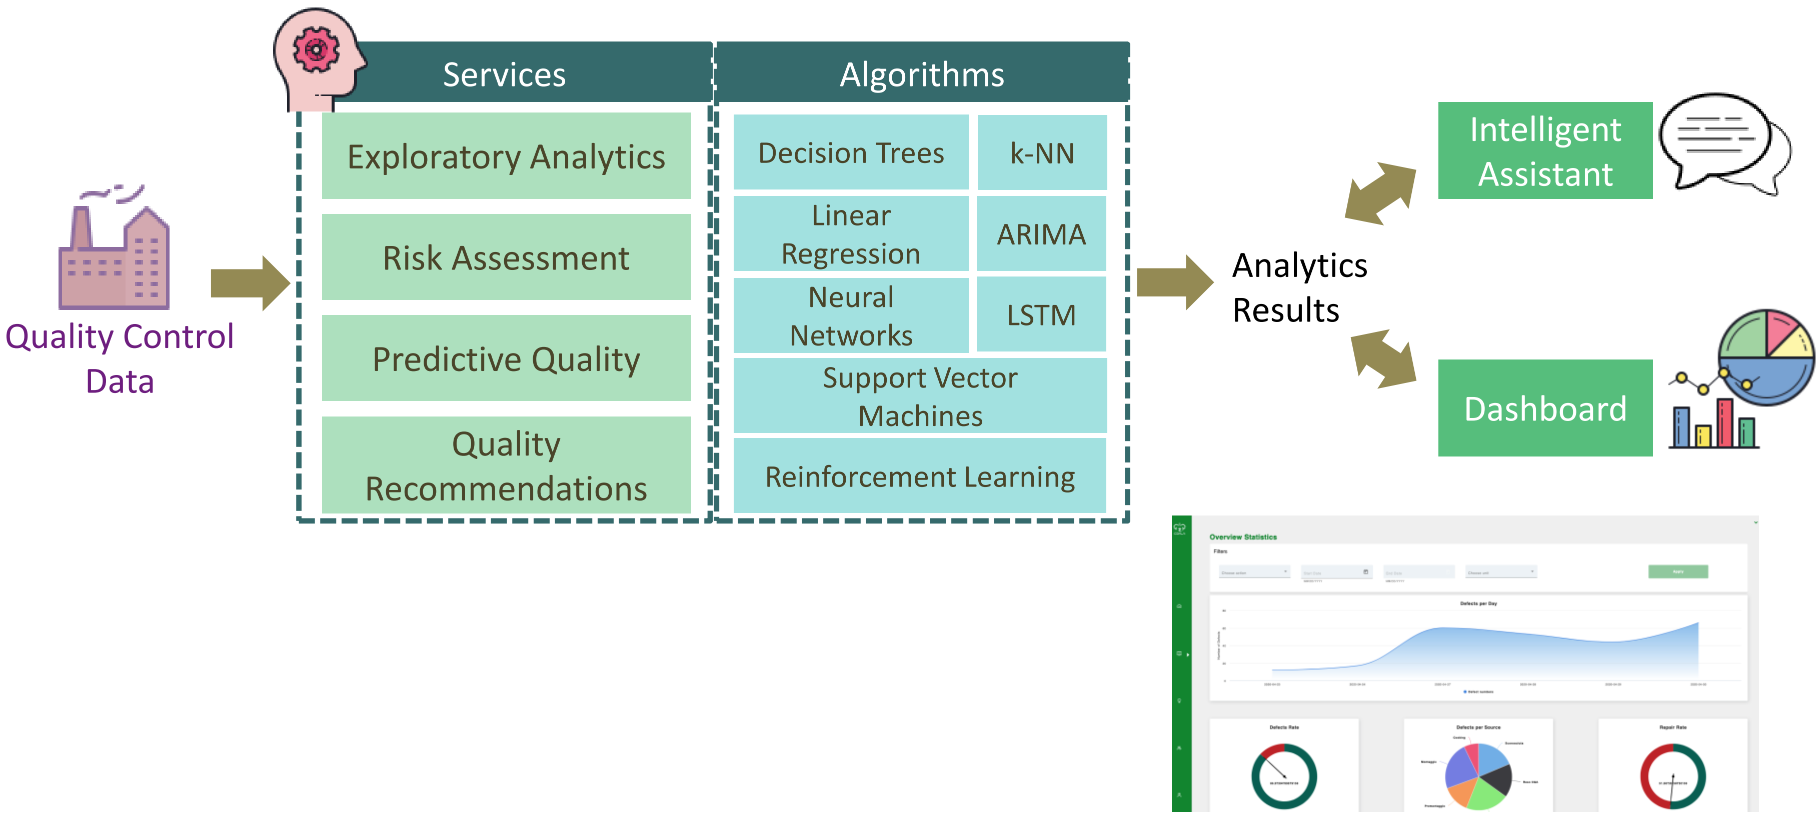

C Bousdekis, A., Wellsandt, S., Bosani, E., Lepenioti, K., Apostolou, D., Hribernik, K., & Mentzas, G. (2021, September). Human-AI Collaboration in Quality Control with Augmented Manufacturing Analytics. In IFIP International Conference on Advances in Production Management Systems Springer, Cham.

Evaluation Method

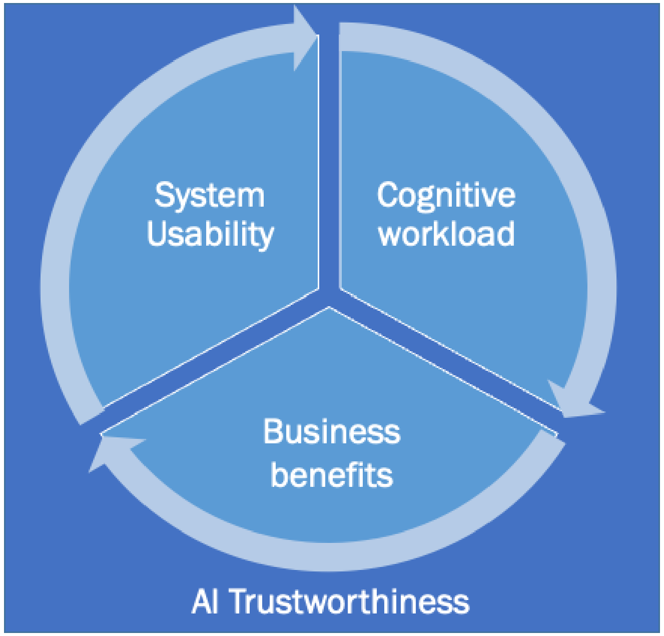

- System Usability is evaluated with the **Voice Usability Scale (VUS)** – an extension of SUS for voice usability
  - Tailored to voice-assistants (Zwakman et al., 2021)
- Cognitive workload quantified with the **NASA Task Load Index (TLX)** (6 subscales)
  - Mental, Physical, Temporal, Overall Performance, Effort, Frustration Level
- Business benefits: examined under various categories of KPIs
  - Organizational, Financial, Business, Operational, Technology, Health and safety, Environmental sustainability
- AI Trustworthiness evaluated with the **Assessment List of Trustworthy AI (ALTAI)** (developed by the High-Level Expert Group on Artificial Intelligence - AI HLEG)
  - Dimensions of ALTAI: Human Agency and Oversight, Technical Robustness and Safety, Privacy and Data Governance, Diversity, Non-discrimination and Fairness, Societal and Environmental Well-being, Accountability

Participation in international initiatives

- **RECITE** is an international, open-source initiative to standardize a formal description of dialog modeling for conversational AI agents
  - RECITE is an an **OASIS Open Project**
  - OASIS Open nonprofit standards consortium is the home of widely adopted standards including the OpenDocument Format, Baseline, SAML, STIX, UBL, and ebXML.

- **ChatBots RESET** (Reveal, Escalate, Substitute, Explain and Track) is an international initiative to design a governance framework for the responsible use of chatbots
  - Part of the **World« Economic Forum**'s Shaping the Future of Technology Governance: Artificial Intelligence and Machine Learning Platform

Outline

- What is Human-centered Artificial Intelligence?
- Our approach: Human-augmented Analytics
- Industry 5.0 Business Cases
- Conclusions and further work

Conclusions

- Developed a framework for Human-Centered Artificial Intelligence in Industry 5.0
- Developed *human augmented analytics* for collaborative human-AI decision making that
  - integrates data-driven predictions with the decision-making process of prescriptive analytics
  - takes into account human experience using mechanisms of interactive RL
- Implementations provide encouraging results
  - financial management (stock trading)
  - predictive maintenance in steel industry
  - quality analytics in white appliances
- Issues to work on:
  - Quality and cost of data
  - Design of human-AI interactions
  - Human-RL agent collaboration (inverse RL)
  - Trust, explainability and transparency

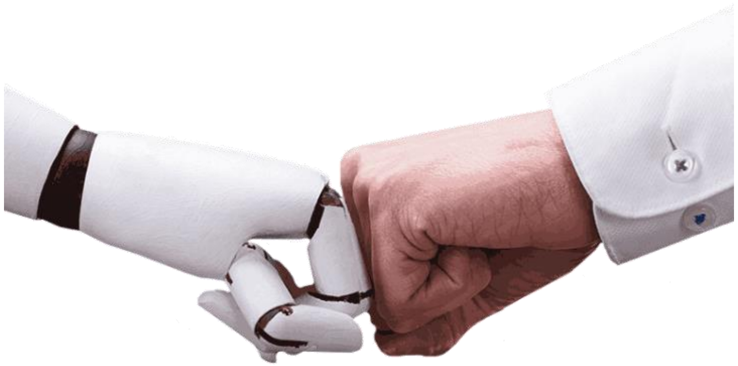

Two main points for further research:

- "There is no AI without data"
  - Need for **data space ecosystems**
- "Don't handicap AI without explicit knowledge"
  - Need for **neuro-symbolic AI**

Data Spaces: European Strategy for Data

The **European Strategy for data** (2020) aims to make the EU a leader in a data-driven society.

The **Data Governance Act** (2020) facilitates data sharing across sectors and Member States.

The **Data Act** (2022) clarifies who can create value from data.

Ten **European common data spaces**, ranging from industry to mobility, from European Green Deal to energy and health.

EC aims to promote the development of common European data spaces in strategic economic sectors and domains of public interest.

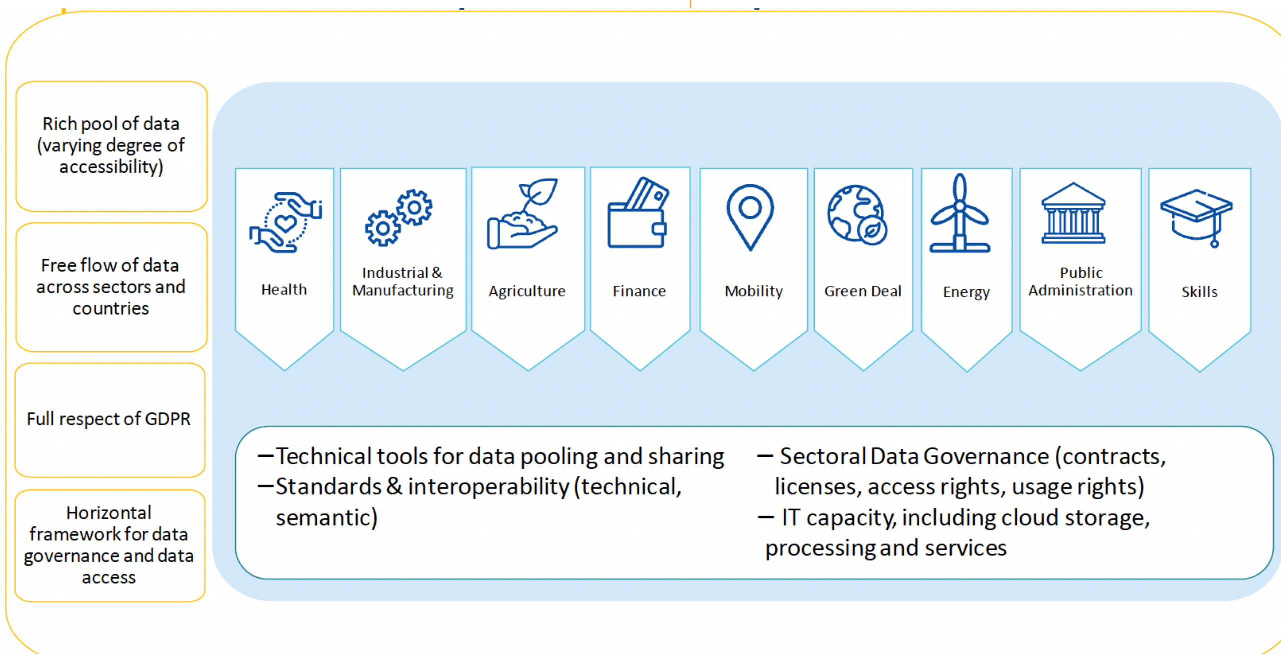

Data Spaces: Architecture and reference models

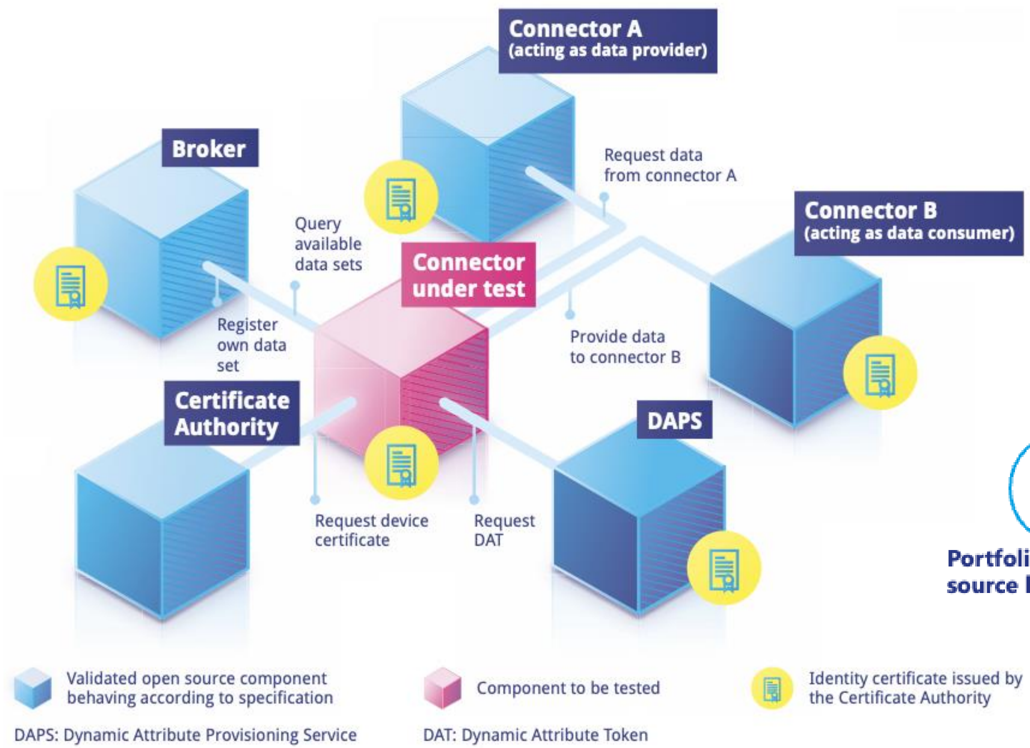

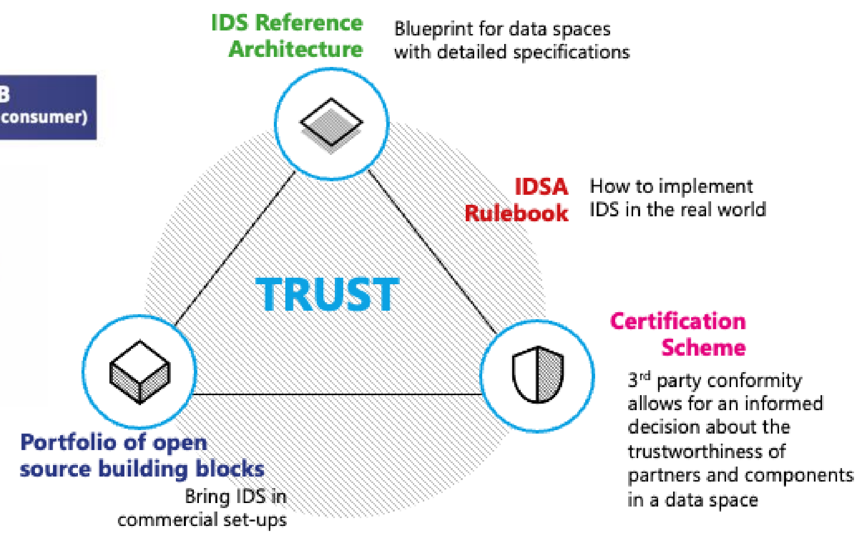

Smart Autonomous Factories

Product Service Factories

Hyper Connected Factories

Otto, B., Hompel, M. T., & Wrobel, S. (2019). International data spaces. In Digital Transformation (pp. 109-128). Springer Vieweg, Berlin, Heidelberg.
Otto, B. (2022). A federated infrastructure for European data spaces. Communications of the ACM, 65(4), 44-45.
Gröger, C. (2021). There is no AI without data. Communications of the ACM, 64(11), 98-108.

Need to integrate explicit and latent semantics

- Semantics is an abstraction of the world we want to describe
  - … consisting of concepts and relations among them
  - …and abstracted into an ontology (world model)

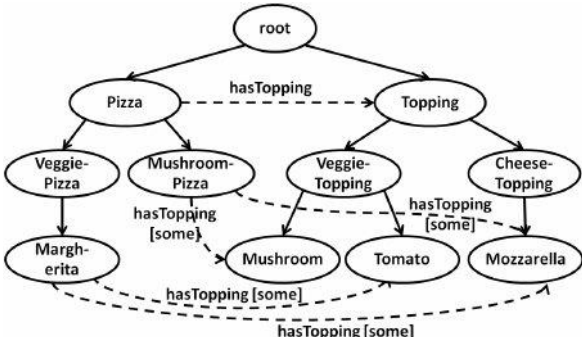

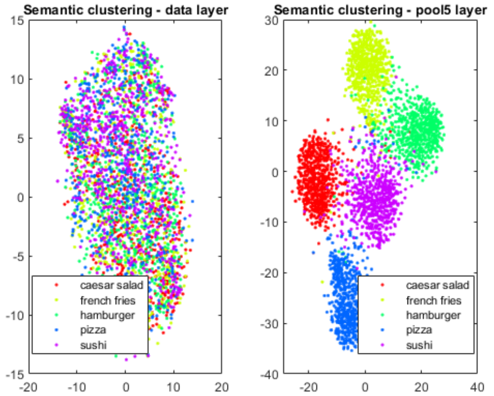

- More concepts then humans can name and identify
  - … 'explicit' concepts (human understandable)
  - … 'latent' concepts (hidden, typically not human understandable)

- Recent success of AI is mostly due to latent semantic representation (deep learning, reinforcement learning)
- Need to integrate symbolic and data-driven (sub-symbolic) technologies in order to bridge human and machine representations

Sheth, A. (2021, September). Don't Handicap AI without Explicit Knowledge: Keynote 3. In 2021 IEEE World Congress on Services (SERVICES) IEEE.

Towards Neuro-Symbolic AI

- New class of AI systems that combine neural networks with symbolic representations of problems and logic.

- In a neurosymbolic system
  - neural networks offer the machinery for efficient learning and computation,
  - symbolic knowledge representation and reasoning enables the use of prior knowledge and explainability.

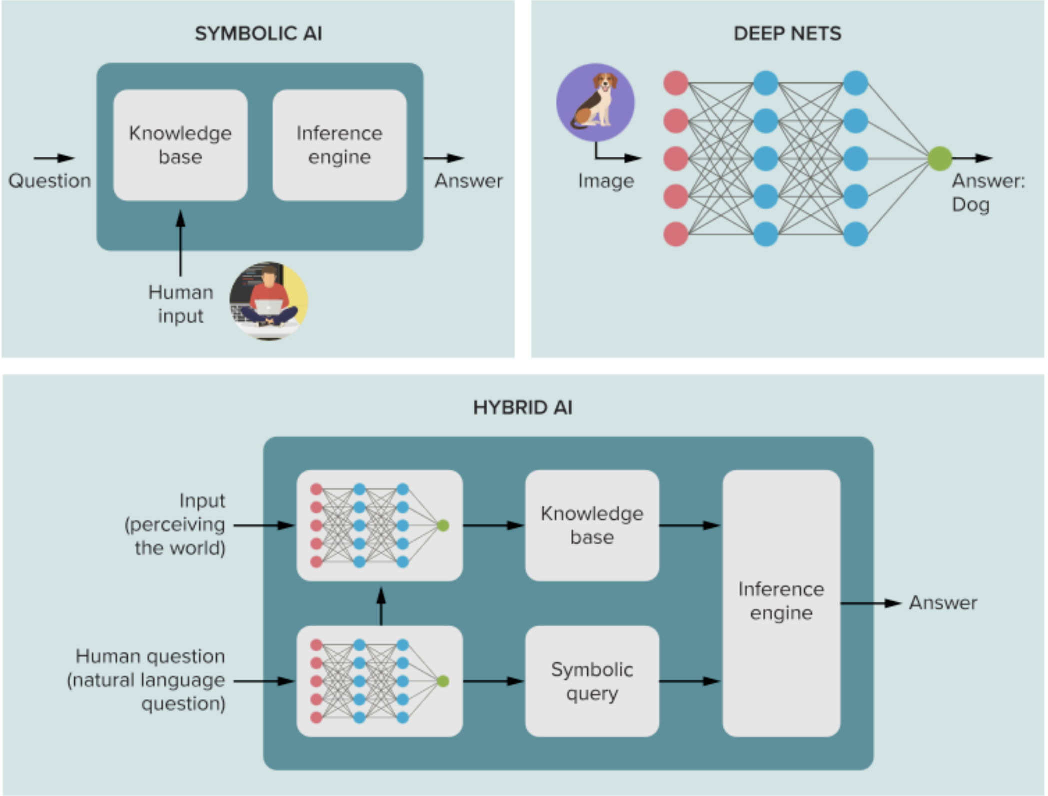

SOURCE: REPORTING BY A. ANANTHASWAMY KNOWABLE MAGAZINE

- ***We believe these systems will usher in a new era of human-centered AI.***

Sarker, M. K., Zhou, L., Eberhart, A., & Hitzler, P. (2021). Neuro-symbolic artificial intelligence. AI Communications,, 1-13.
Hitzler, P., & Sarker, M. K. (Eds.). (2022). Neuro-Symbolic Artificial Intelligence: The State of the Art., Frontiers in Artificial Intelligence and Applications, IOS Press.

More info?

Bousdekis, A., & Mentzas, G. (2021). Enterprise Integration and Interoperability for Big Data-Driven Processes in the Frame of Industry 4.0. *Frontiers in big Data*, 4, 22.

Bousdekis, A., Apostolou, D., & Mentzas, G. (2020). A human cyber physical system framework for operator 4.0–artificial intelligence symbiosis. *Manufacturing letters*, 25, 10-15.

Bousdekis, A., Papageorgiou, N., Magoutas, B., Apostolou, D., & Mentzas, G. (2020). Sensor-driven learning of time-dependent parameters for prescriptive analytics. *IEEE Access*, 8, 92383-92392.

Bousdekis, A., Wellsandt, S., Bosani, E., Lepenioti, K., Apostolou, D., Hribernik, K., & Mentzas, G. (2021), Human-AI Collaboration in Quality Control with Augmented Manufacturing Analytics. In *IFIP International Conference on Advances in Production Management Systems* (pp. 303-310). Springer.

Fouka, A., Bousdekis, A., Lepenioti, K., & Mentzas, G. (2021). Real-Time Equipment Health State Prediction with LSTM Networks and Bayesian Inference. *In International Conference on Advanced Information Systems Engineering* (pp. 155-166). Springer.

Hribernik, K., Cabri, G., Mandreoli, F., & Mentzas, G. (2021). Autonomous, context-aware, adaptive Digital Twins—State of the art and roadmap. *Computers in Industry*, 133, 103508.

Karlis, V., Lepenioti, K., Bousdekis, A., & Mentzas, G. (2021, July). Stock Trend Prediction by Fusing Prices and Indices with LSTM Neural Networks. In 2021 *12th IEEE International Conference on Information, Intelligence, Systems & Applications* (IISA) (pp. 1-7).

Lepenioti, K., Bousdekis, A., Apostolou, D., & Mentzas, G. (2020). Prescriptive analytics: Literature review and research challenges. *International Journal of Information Management*, 50, 57-70.

Lepenioti, K., Bousdekis, A., Apostolou, D., & Mentzas, G. (2021). Human-augmented prescriptive analytics with interactive multi-objective reinforcement learning. *IEEE Access*, 9, 100677-100693.

Lepenioti, K., Pertselakis, M., Bousdekis, A., Louca, A., Lampathaki, F., Apostolou, D., G. Mentzas, & Anastasiou, S. (2020). Machine learning for predictive and prescriptive analytics of operational data in smart manufacturing. In *International Conference on Advanced Information Systems Engineering* (pp. 5-16). Springer.

Mentzas, G., Lepenioti, K., Bousdekis, A., & Apostolou, D. (2021, September). Data-Driven Collaborative Human-AI Decision Making. In *Conference on e-Business, e-Services and e-Society* (pp. 120-131). Springer, Cham.

Wellsandt, S., K. Klein, K. Hribernik, M. Lewandowski, A. Bousdekis, G. Mentzas, K.-D. Thoben (2022) Hybrid-augmented intelligence in predictive maintenance with digital intelligent assistants, *Annual Reviews in Control*, 2022.

With many thanks to: Katerina Lepenioti, Alexandros Bousdekis, Dimitris Apostolou, Afroditi Fouka, Mattheos Fikardos, Nikos Papageorgiou, Yannis Verginadis, Stefan Wellsandt, Konstantin Klein, Karl Hribernik, Marko Lewandowski, Klaus-Dieter Thoben.

For more information and to contact us…

VISIT OUR SITE — http://imu.ntua.gr

Connect with us on LinkedIn — https://www.linkedin.com/company/2929404/

FOLLOW US ON twitter — https://twitter.com/imu_ntua

Watch Us On YouTube — https://www.youtube.com/channel/UCenUYJveI5WHb92x-jiH8Zg