Explainable Artificial Intelligence (XAI) Supporting Public Administration Processes – On the Potential of XAI in Tax Audit

Professor Dr. Peter Fettke
Saarland University and German Research Center for Artificial Intelligence (DFKI)

Towards eXplainable Artificial Intelligence (XAI) in Taxation: *The Future of Good Tax Governance*
University of Amsterdam, Amsterdam Centre for Tax Law

Keynote, Amsterdam, Netherlands, March 10, 2023

Agenda

1. Artificial Intelligence and taxation
2. Explainable Artificial Intelligence (XAI)
3. XAI in tax audit
4. Conclusion and outlook

AI as a plot of Hollywood blockbusters

friendly

K.I.T.T.
Knight Rider
1982

David
A.I. – Artificial Intelligence
2001

technoid

humanoid

HAL 9000
2001: A Space Odyssey
1968

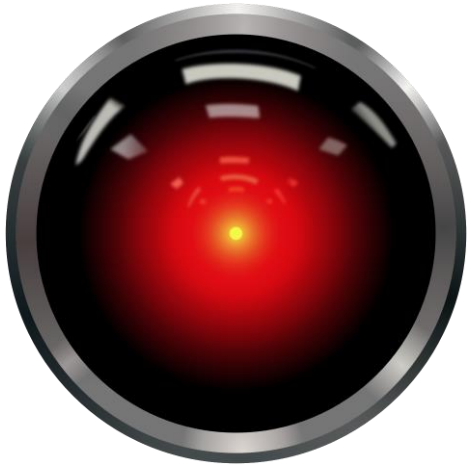

Red Queen
Resident Evil
2002

unfriendly

AI as an academic discipline

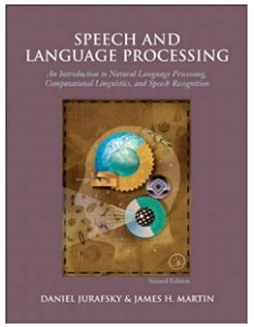

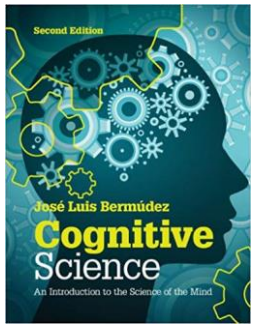

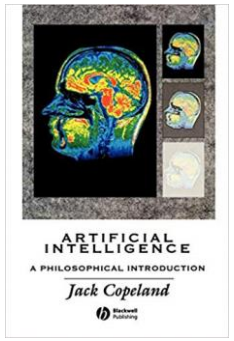

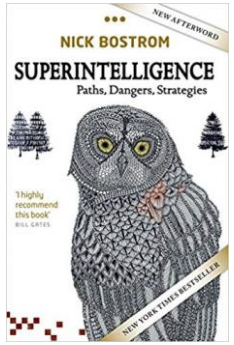

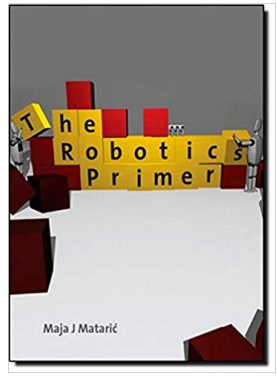

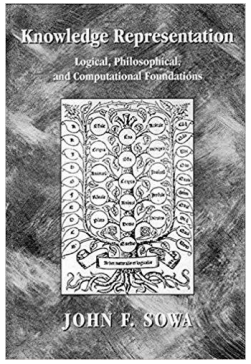

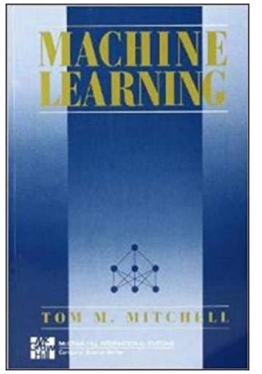

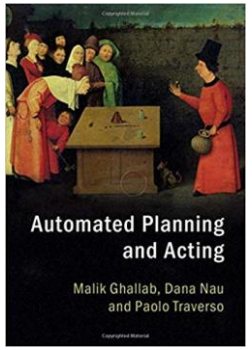

Super-AI

General AI

Narrow AI

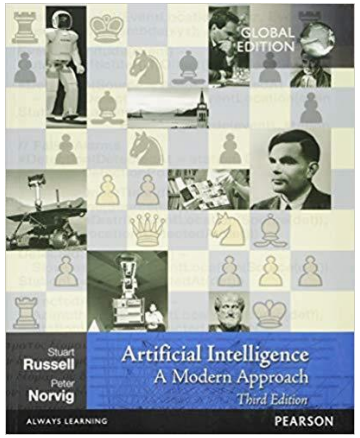

Robotics

Deep Learning

Machine Learning

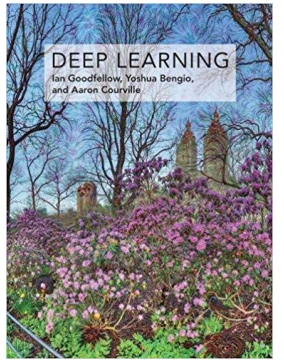

AI as state of the art and avantgarde informatics (Dilemma of AI: „As soon as it works, no one calls it AI anymore ...“*)

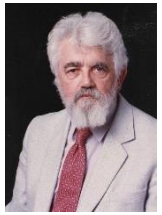

* John McCarthy (1927-2011), one founding father of AI

AI as state of the art: document intelligence

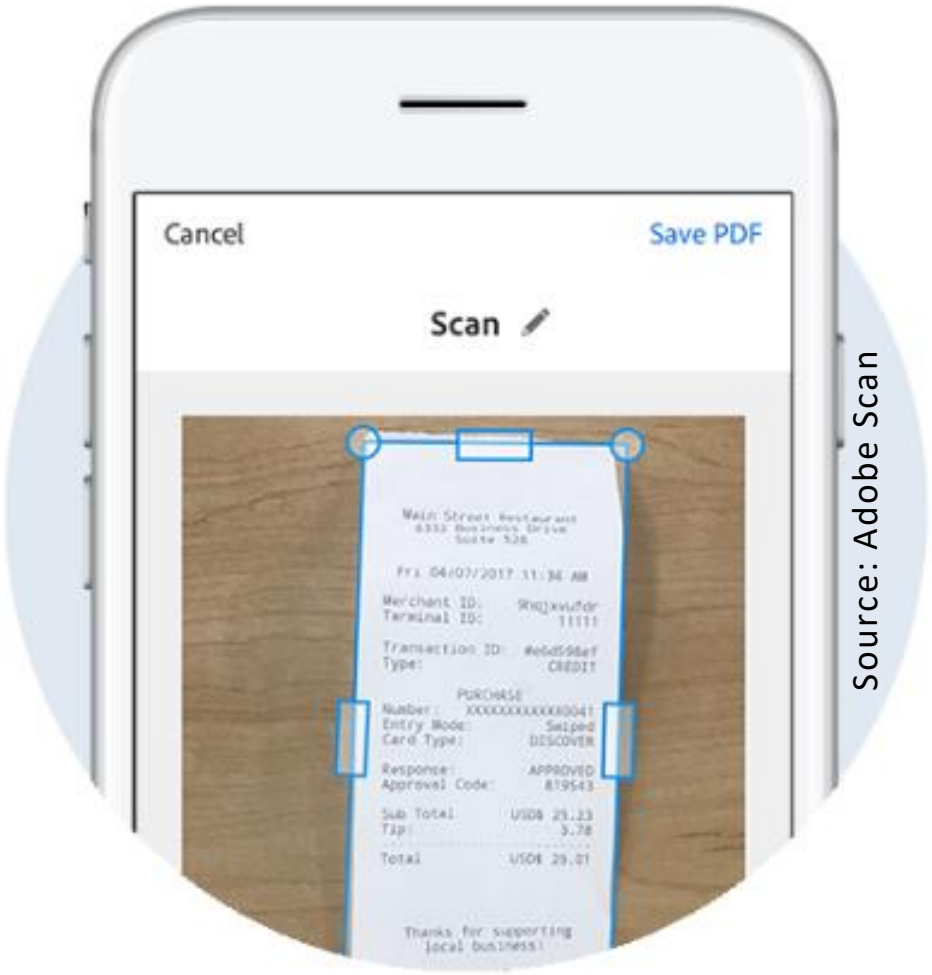

- scan corrections (clipping, rotations, distortions, shadows)
- recognition of document type
- transformation of photo pixels to digital text
- extract relevant information, document classification (business cards, invoices etc.)
- plausibility checks (spelling, anomalies etc.)

WTS-DFKI -Study: “AI in the Taxation World” (2017)

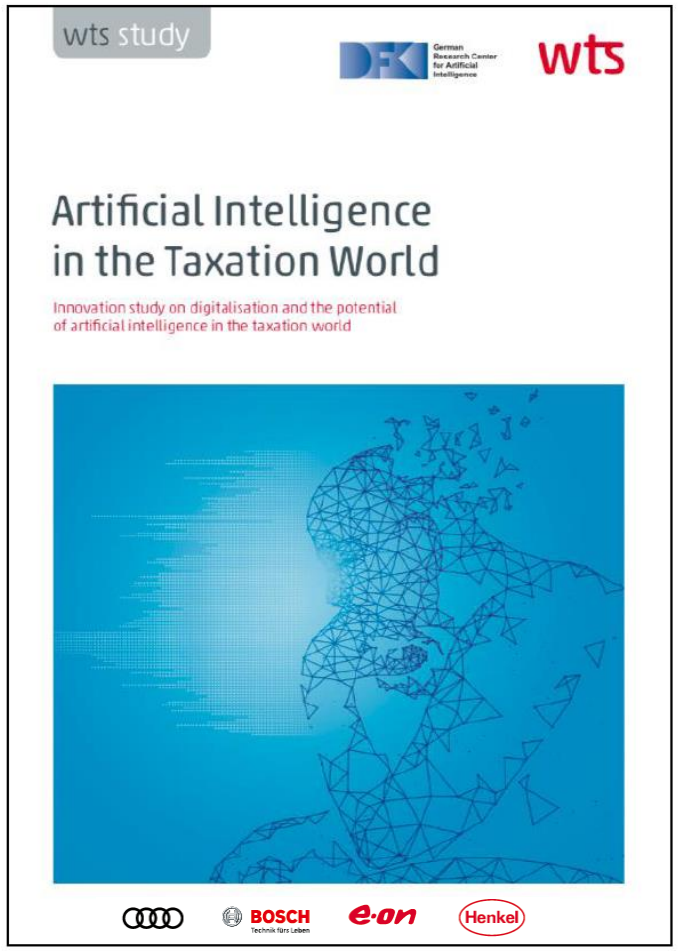

Machine learning

- Classification
- Clustering
- Reinforcement learning
- Decision trees
- Regression analysis
- Decision trees
- Logistic regression
- Support Vector Machines
- Deep learning
- Autoencoder
- Feature extraction

Process mining

- Process identification
- Process prediction
- Anomaly detection
- Business model optimisation
- Process diagnosis
- Process analysis
- Process optimisation
- Process adaptation
- Conformance testing

Information extraction

- Sentiment analysis
- Trend analysis
- Data mining
- Optical character recognition (OCR)
- Correlation analysis
- Argumentation mining

Knowledge management

- Organisational memory
- Semantic technologies and ontologies
- Context identification
- Case based reasoning
- Knowledge-based document analysis
- Knowledge graph
- Expert systems
- Semantic desktop
- Data mining

Speech processing

- Speech recognition
- Speech generation
- Speech synthesis
- Speech comprehension
- Text mining
- Machine translation
- Q & A systems

Multimodal systems

- Adaptive user interfaces
- Intelligent user interfaces
- Multilingual systems
- Assistance systems
- Recommendation systems
- Dialogue systems
- Personalisation

AI in the world of taxation: developed prototypes at DFKI

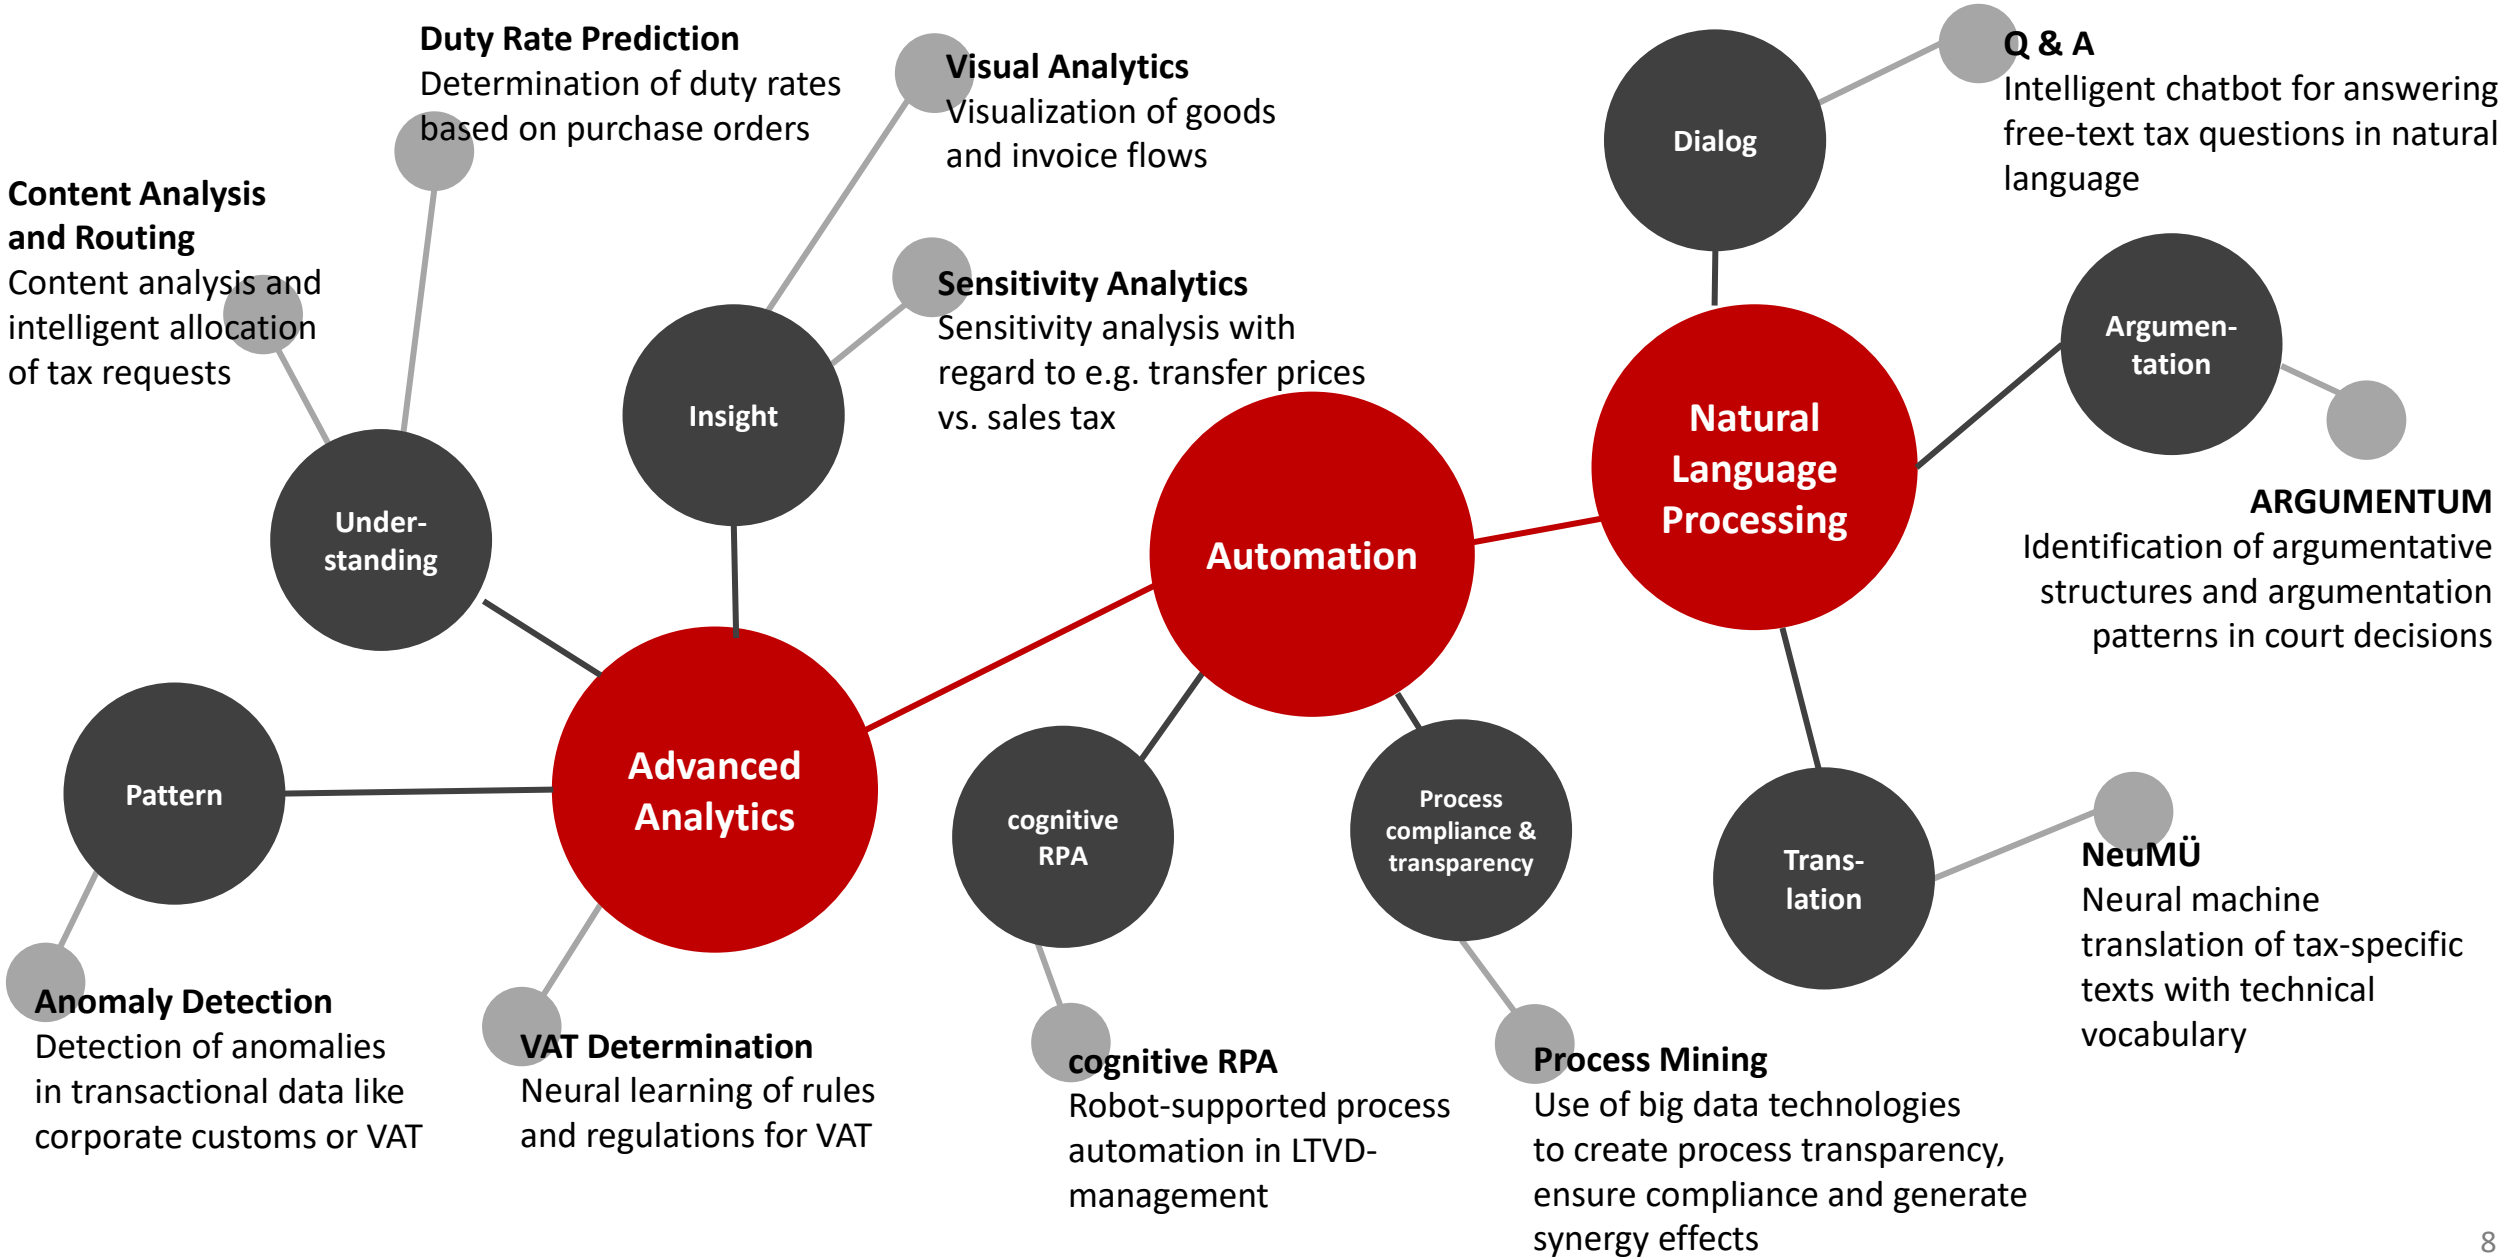

Agenda

1. Artificial Intelligence and taxation
2. Explainable Artificial Intelligence (XAI)
3. XAI in tax audit
4. Conclusion and outlook

XAI as a “hot topic” in core AI research

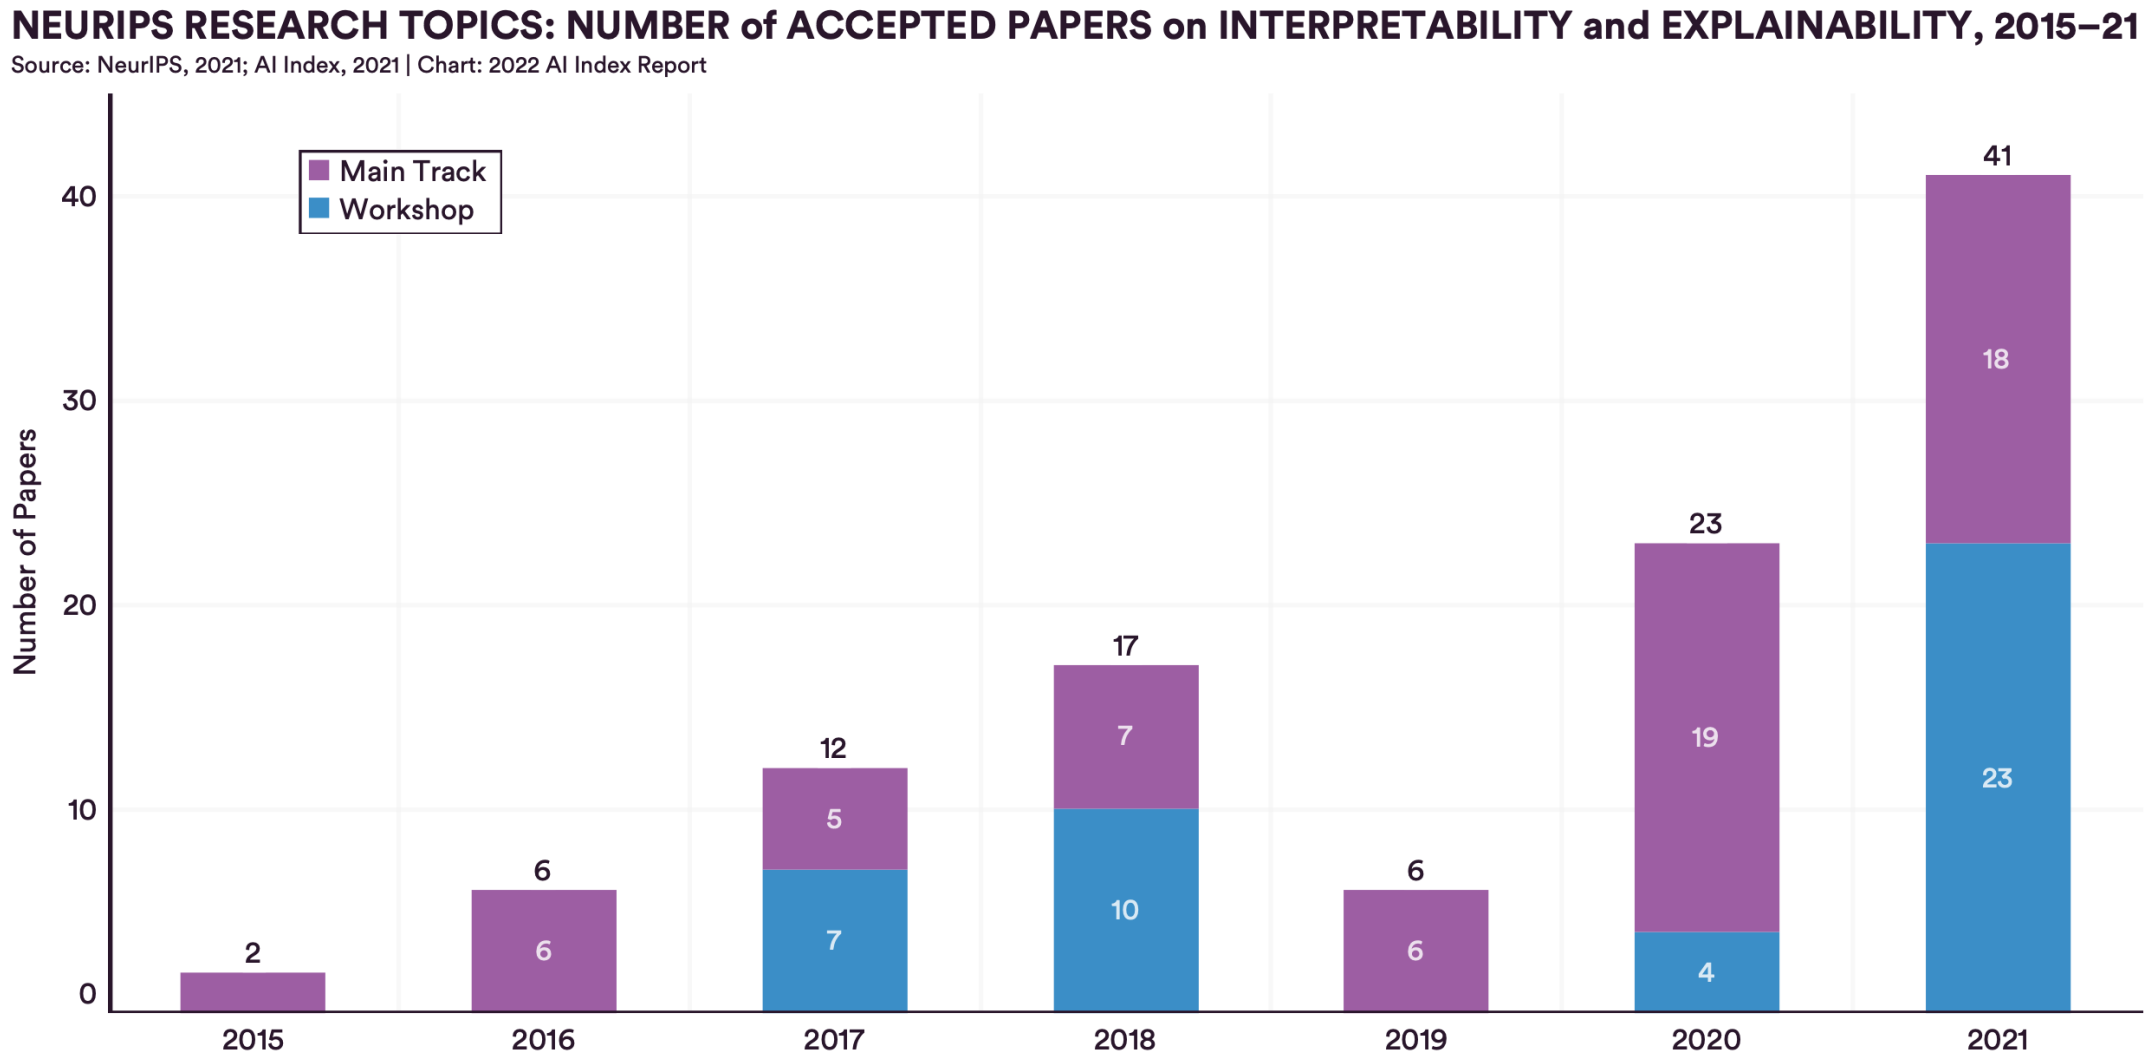

source: Artificial Intelligence Index Report 2022, Stanford University

Trustworthy AI

Trustworthy AI has three components, which should be met throughout the system's entire life cycle:

1. it should be ***lawful***, complying with all applicable laws and regulations
2. it should be ***ethical***, ensuring adherence to ethical principles and values and
3. it should be *robust*, both from a **technical and social perspective**

XAI is defined as a technical method to ensure that Trustworthy AI principles can be incorporated in the design, development and use phases of an AI system

source: https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai

XAI history: the concept of explanation is a multi-faceted, non-monolithic

explanation for expert systems

explanations for advanced black-box AI-based systems

1965

1970s-1980s

1990s-2000s

2010s-now

theoretical and practical foundations for explanations from intelligent systems (driven by IS community and others)

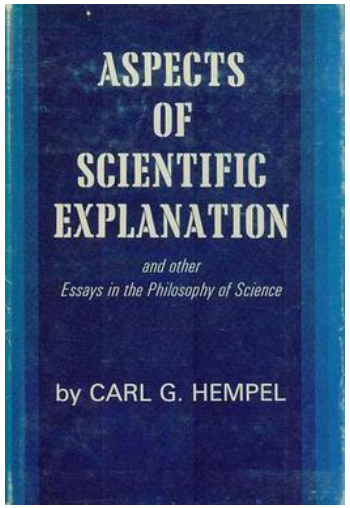

Classical architecture (early XAI systems)

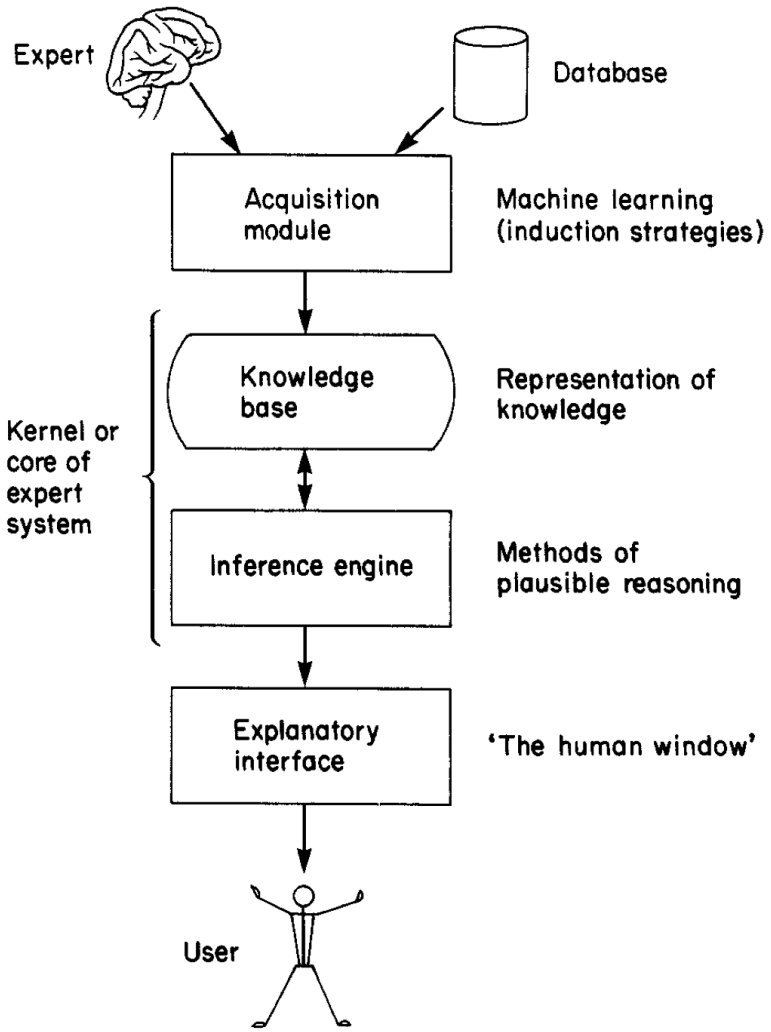

source: Forsyth (1984)

Source: Wikipedia (2023)

XAI idea for machine-learning models

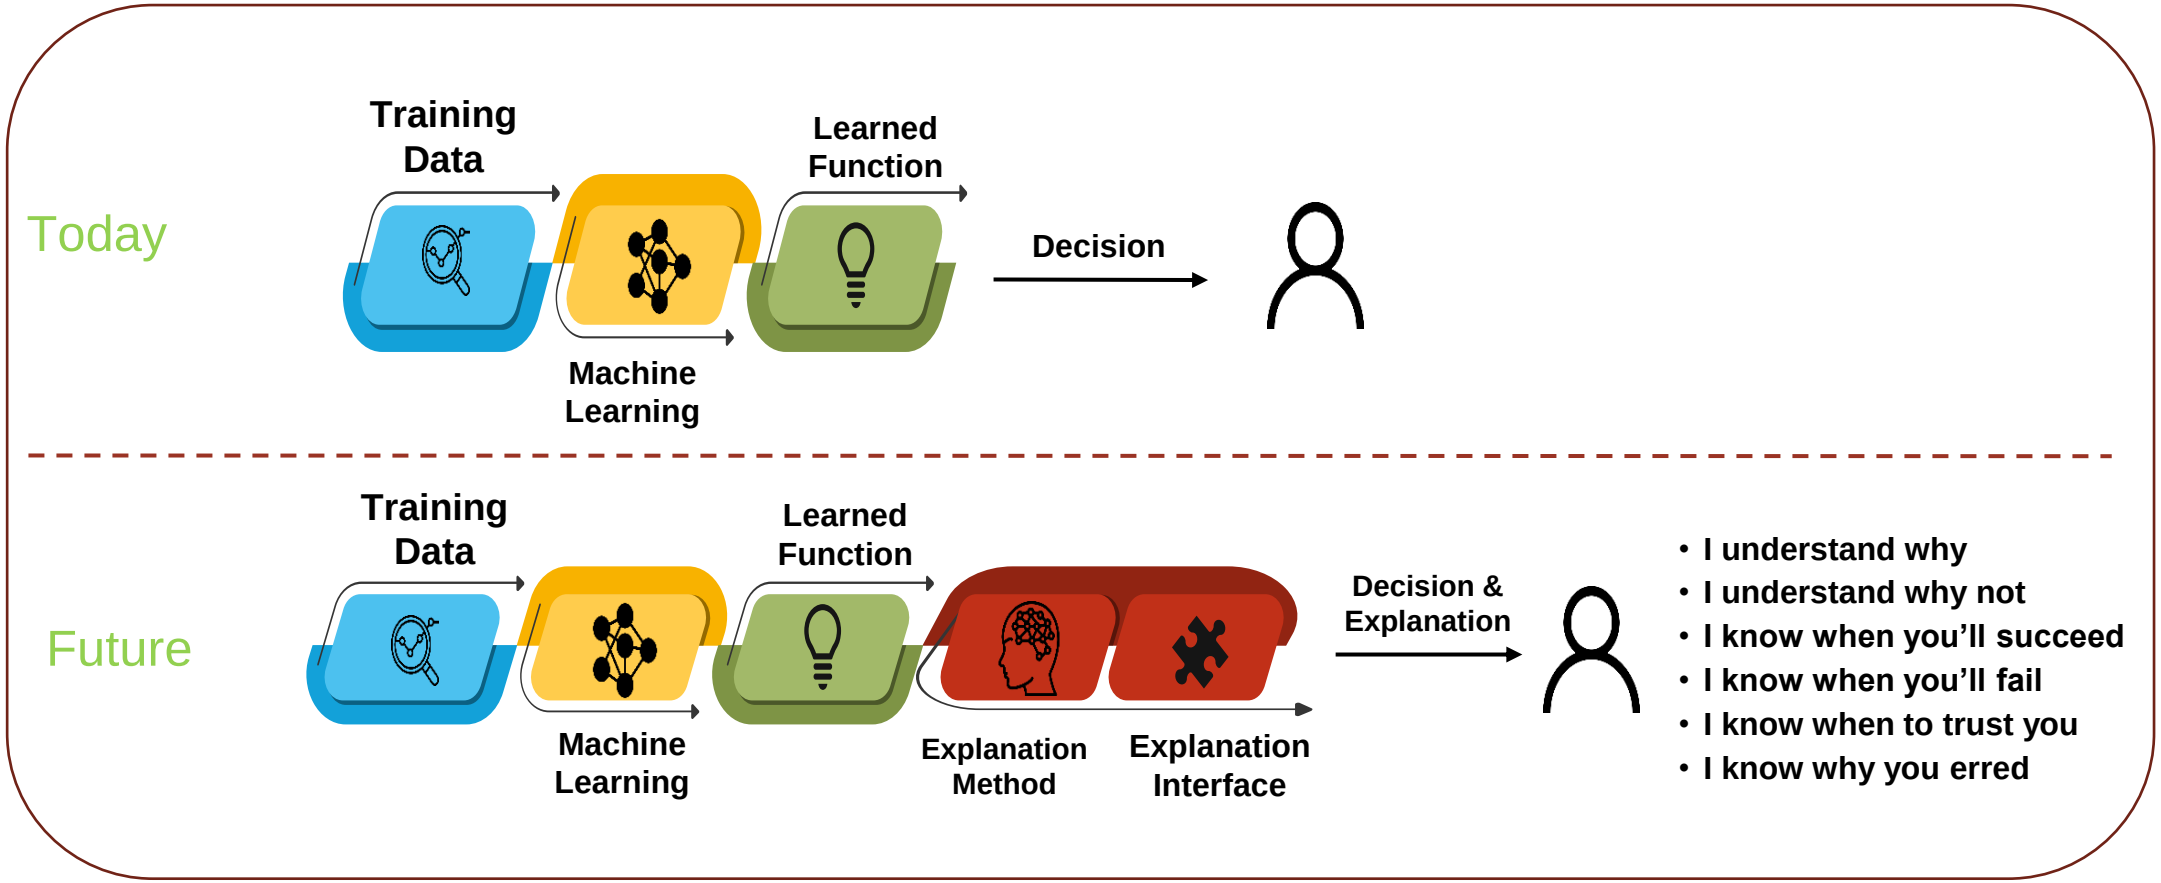

based on Gunning (2017)

XAI example: computer vision (1/2)

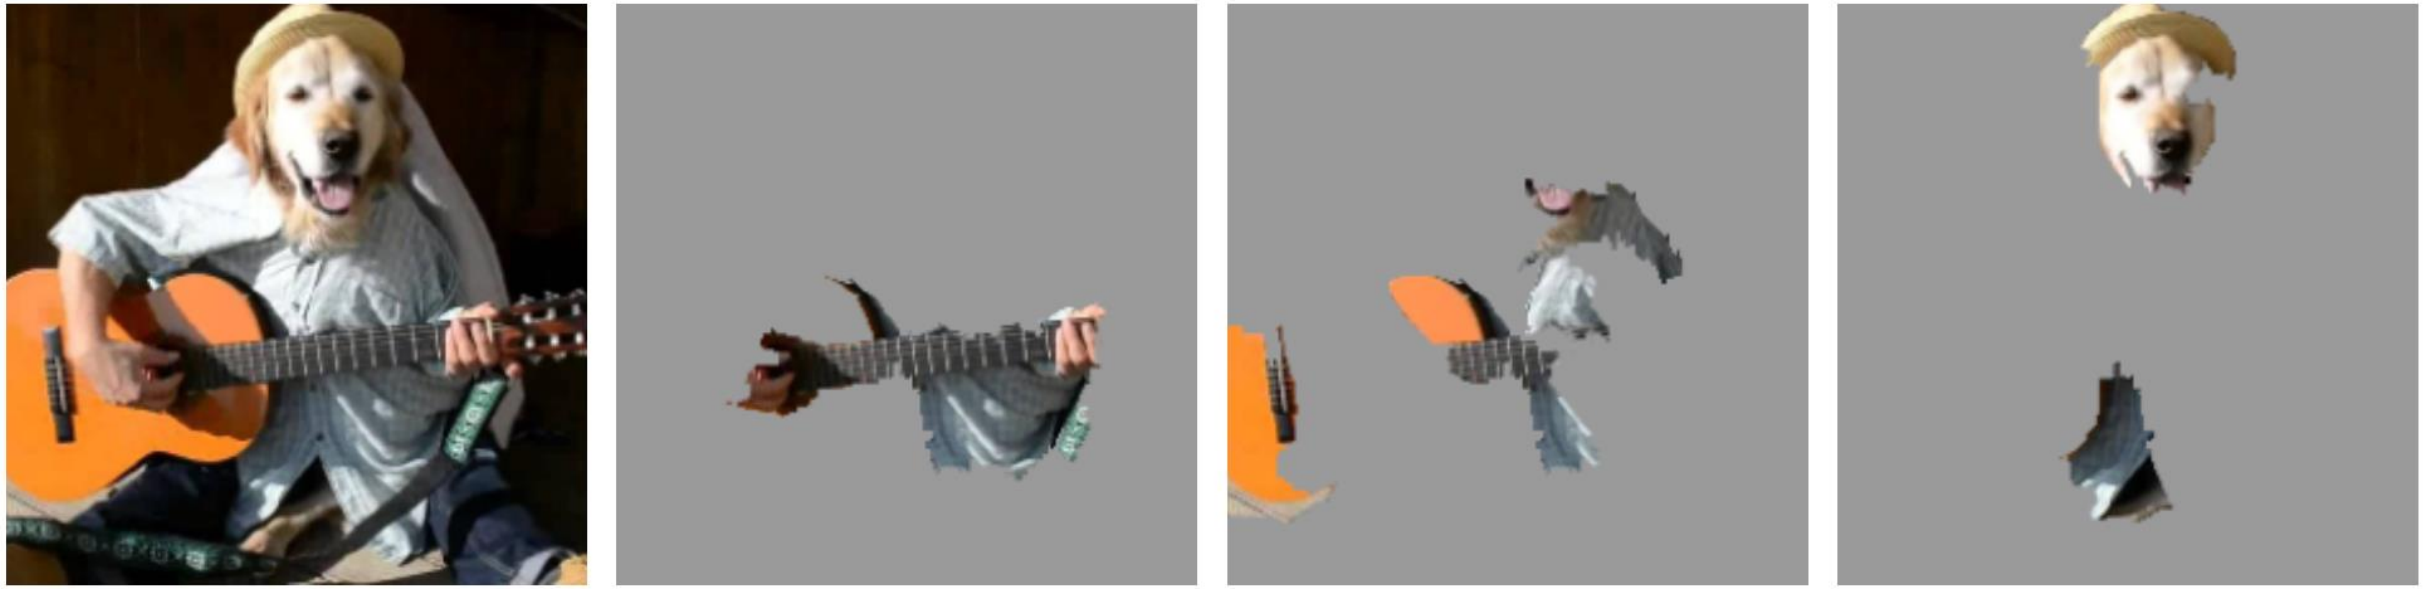

(a) Original Image (b) Explaining *Electric guitar* (c) Explaining *Acoustic guitar* (d) Explaining *Labrador*

Figure 4: Explaining an image classification prediction made by Google's Inception neural network. The top 3 classes predicted are "Electric Guitar" ($p = 0.32$), "Acoustic guitar" ($p = 0.24$) and "Labrador" ($p = 0.21$)

source: Ribeiro et al. (2016)

XAI example: computer vision (2/2)

(a) Husky classified as wolf (b) Explanation

Figure 11: Raw data and explanation of a bad model's prediction in the "Husky vs Wolf" task.

source: Ribeiro et al. (2016)

Agenda

1. Artificial Intelligence and taxation
2. Explainable Artificial Intelligence (XAI)
3. XAI in tax audit
4. Conclusion and outlook

Objective: Explaining ML-based tax decision models

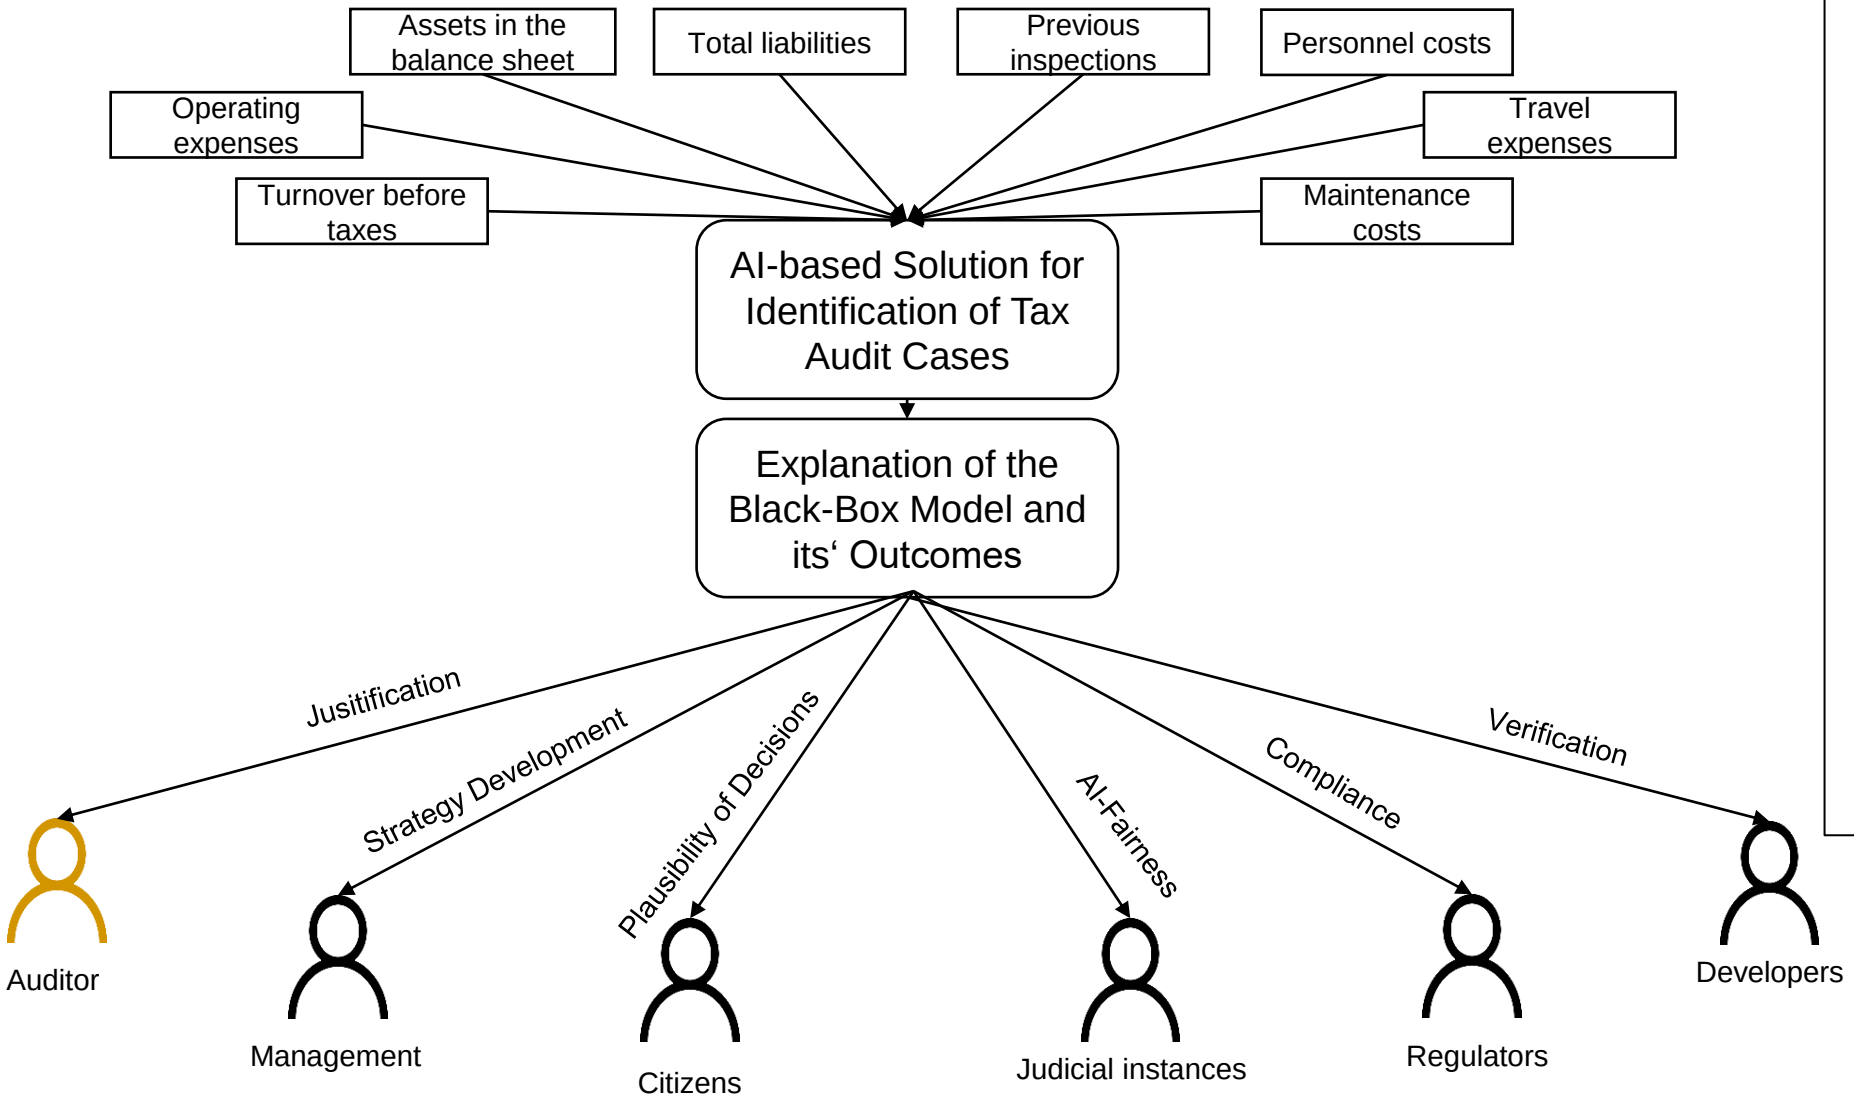

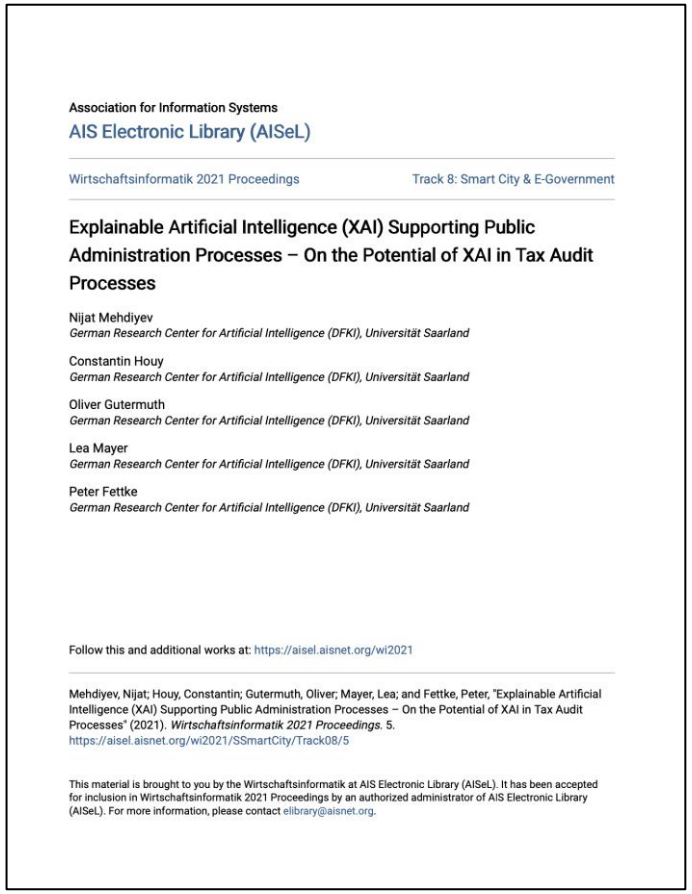

Association for Information Systems
AIS Electronic Library (AISeL)

Wirtschaftsinformatik 2021 Proceedings — Track 8: Smart City & E-Government

Explainable Artificial Intelligence (XAI) Supporting Public Administration Processes – On the Potential of XAI in Tax Audit Processes

Nijat Mehdiyev
German Research Center for Artificial Intelligence (DFKI), Universität Saarland

Constantin Houy
German Research Center for Artificial Intelligence (DFKI), Universität Saarland

Oliver Gutermuth
German Research Center for Artificial Intelligence (DFKI), Universität Saarland

Lea Mayer
German Research Center for Artificial Intelligence (DFKI), Universität Saarland

Peter Fettke
German Research Center for Artificial Intelligence (DFKI), Universität Saarland

source: Mehdiyev et al. (2021)

Decision prediction

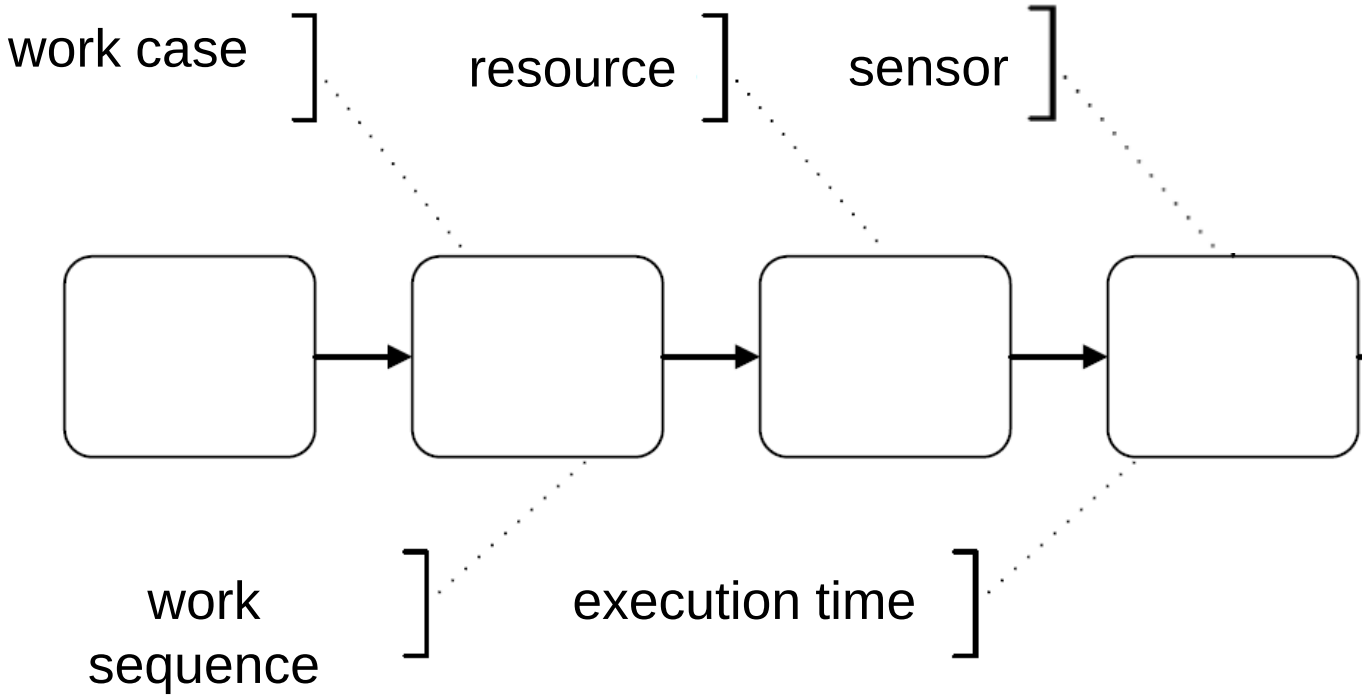

Decision prediction

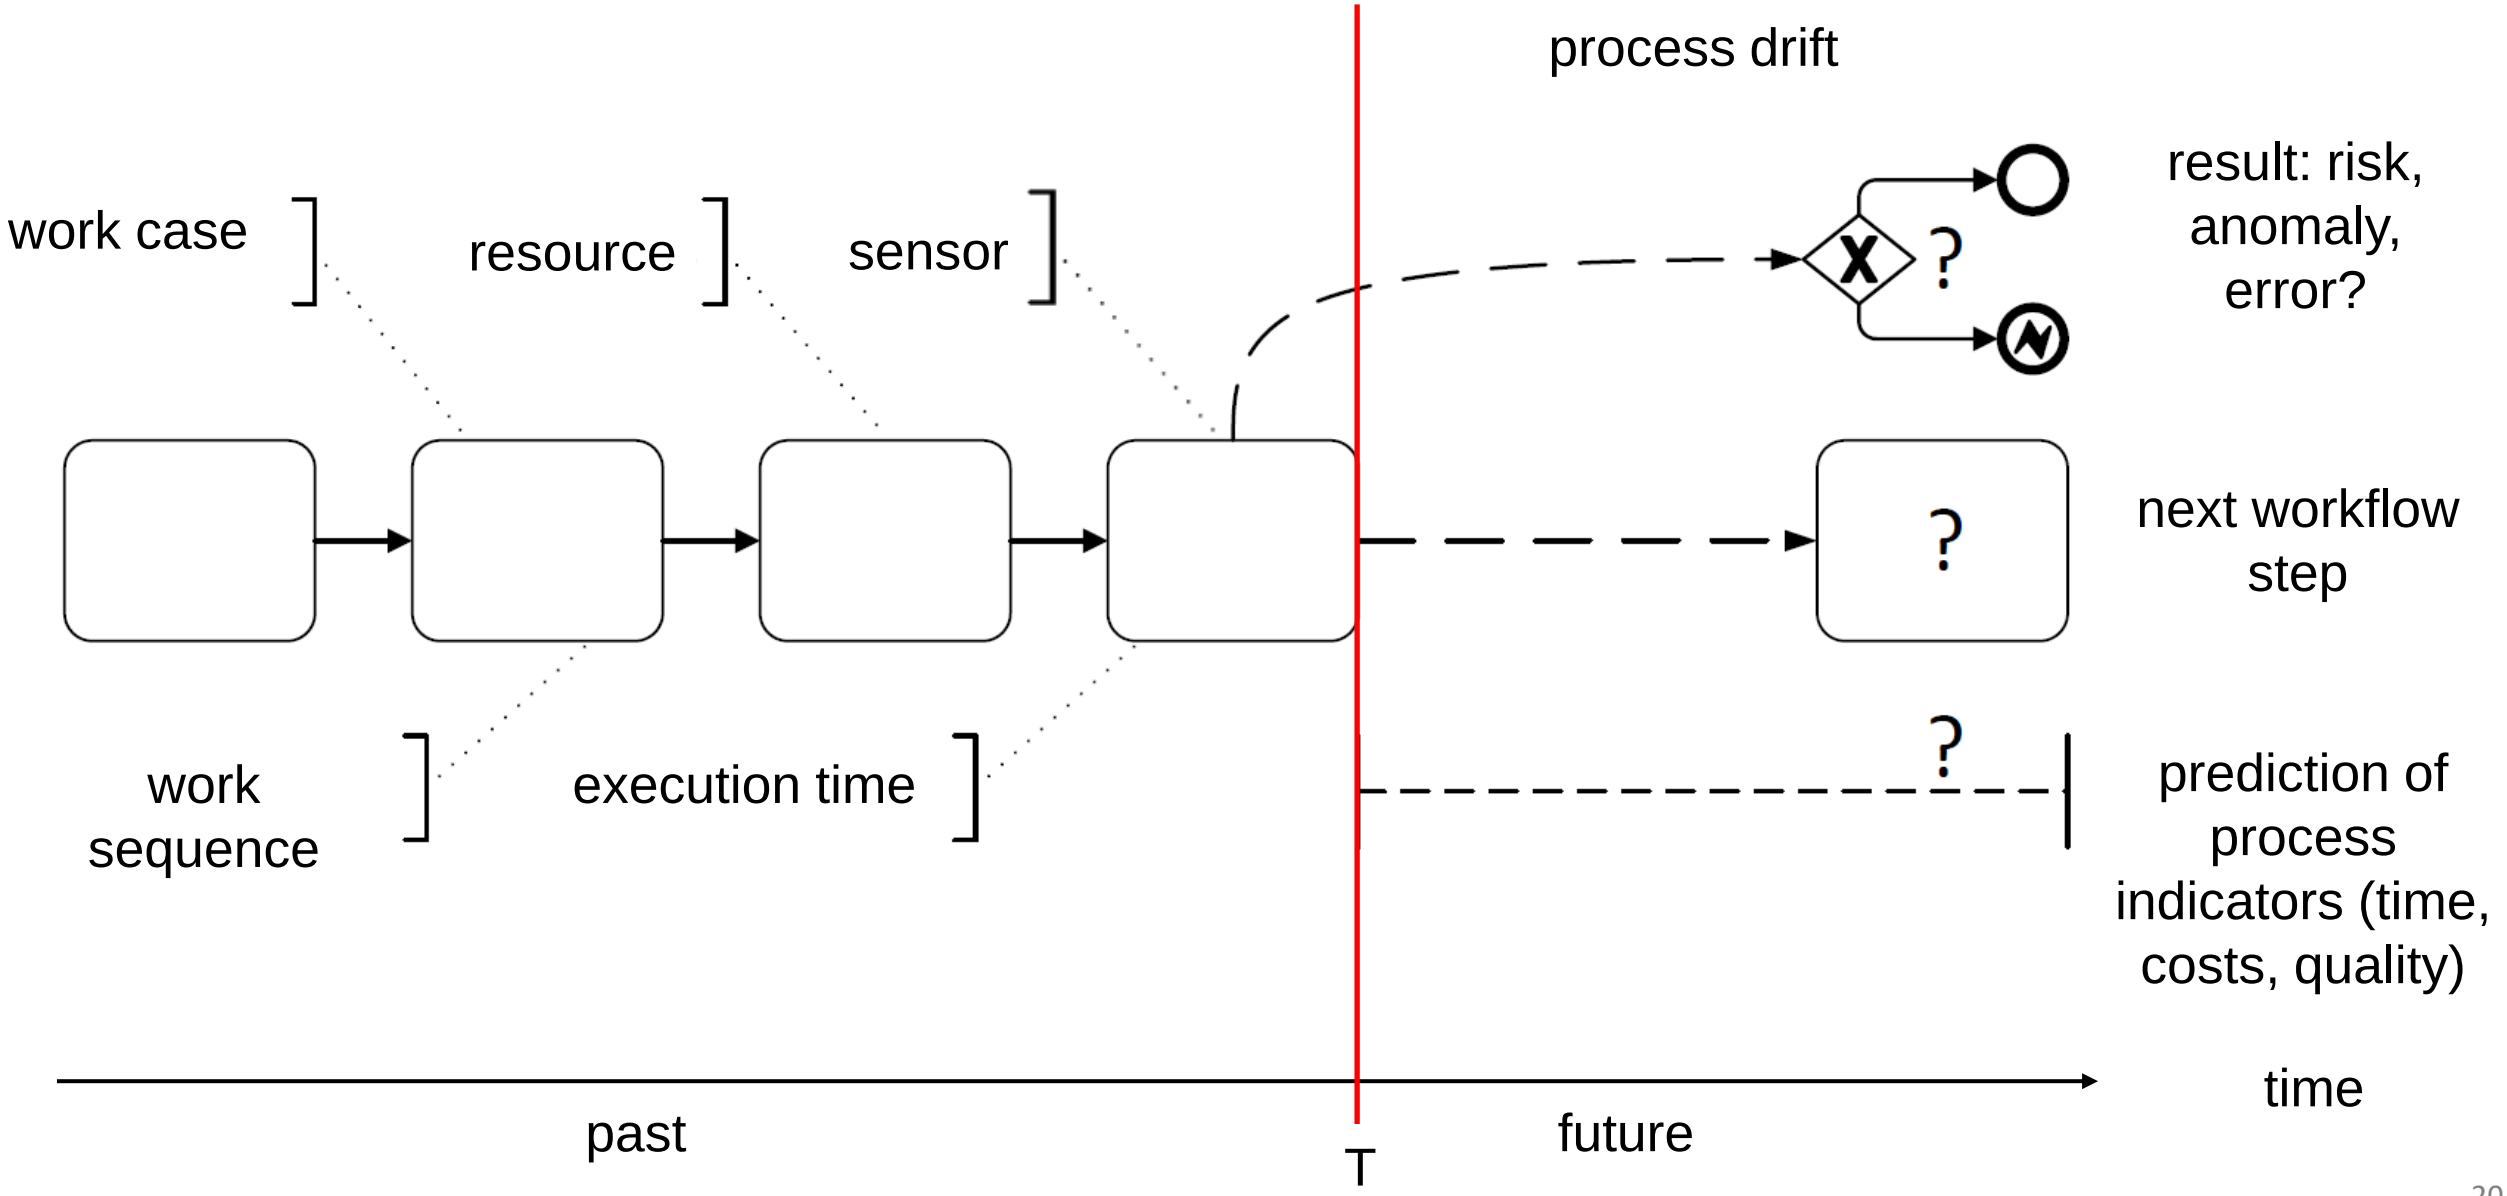

Decision prediction using deep learning

Decision Support Systems 100 (2017) 129–140

Contents lists available at ScienceDirect

Decision Support Systems

journal homepage: www.elsevier.com/locate/dss

Predicting process behaviour using deep learning

Joerg Evermann[a,*], Jana-Rebecca Rehse[b,c], Peter Fettke[b,c]

[a] Memorial University of Newfoundland, St. John's, NL, Canada
[b] German Research Center for Artificial Intelligence, Saarbrücken, Germany
[c] Saarland University, Saarbrücken, Germany

ARTICLE INFO

Article history:
Received 8 July 2016
Received in revised form 22 March 2017
Accepted 5 April 2017
Available online 17 April 2017

Keywords:
Process management
Runtime support
Process prediction
Deep learning
Neural networks

ABSTRACT

Predicting business process behaviour is an important aspect of business process management. Motivated by research in natural language processing, this paper describes an application of deep learning with recurrent neural networks to the problem of predicting the next event in a business process. This is both a novel method in process prediction, which has largely relied on explicit process models, and also a novel application of deep learning methods. The approach is evaluated on two real datasets and our results surpass the state-of-the-art in prediction precision.

1. Introduction

Being able to predict the future behaviour of a business process is an important business capability [1]. As an application of predictive analytics in business process management, process prediction exploits data on past process instances to make predictions about current ones [2]. Example use cases are customer service agents responding to inquiries about the remaining time until a case is resolved, production managers predicting the completion time of a production process for better planning and higher utilization, or case managers identifying likely compliance violations to mitigate business risk.

We present a novel approach to predicting the next process event using deep learning. While the term "deep learning" has only recently become a popular research topic, it is essentially an application of neural networks and thus looks back on a long history of research [3]. Recent innovations both in algorithms, allowing novel architectures of neural networks, and computing hardware, especially GPU processing, have led to a resurgence in interest for neural networks and popularized the term "deep learning" [4]. Our approach is motivated by applications of neural networks to natural language processing(NLP), more specifically the prediction of the next word in a sentence [5–7]. By interpreting process event logs as text, process traces as sentences, and process events as words, these techniques can be applied to predict future process events. The contribution of our research is threefold:

1. We improve on the state-of-the-art in process event prediction. Our results show our method has considerably better precision on next-event prediction.
2. We demonstrate that an explicit process model is not necessary for prediction. Deep learning models, where the process structure is only implicitly reflected, can perform as well as explicit process models.
3. We contribute to process management in general by showcasing the useful application of an artificial intelligence approach, illustrating that business process management can benefit from the application of smart approaches.

Our research is located at the intersection of business process management, in particular process mining, and artificial intelligence (AI) and machine learning. We bring together historic process data with an AI learning technology to leverage real-time case management, opening new perspectives into process execution, monitoring, and analysis. Extending existing solutions to novel problems ("exaptation") is a recognized and valid way to make a contribution in design science [8], which is the research approach we apply here. We not only provide a new approach, rooted in AI, to predicting the next

* Corresponding author.
E-mail address: jevermann@mun.ca (J. Evermann).

http://dx.doi.org/10.1016/j.dss.2017.04.003
0167-9236/

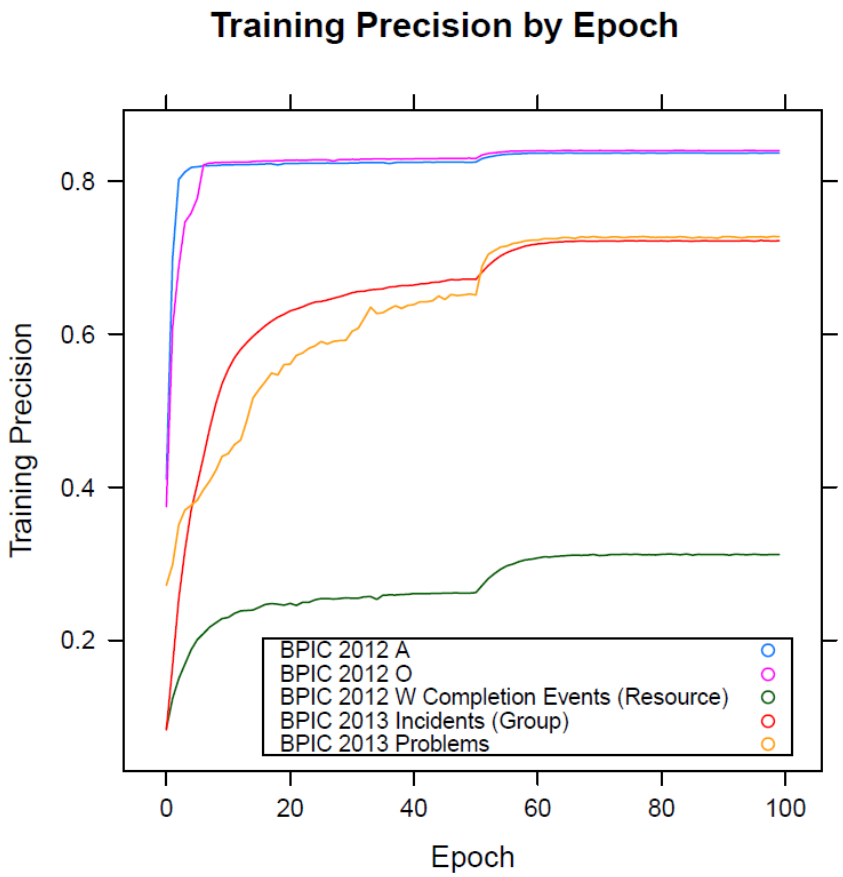

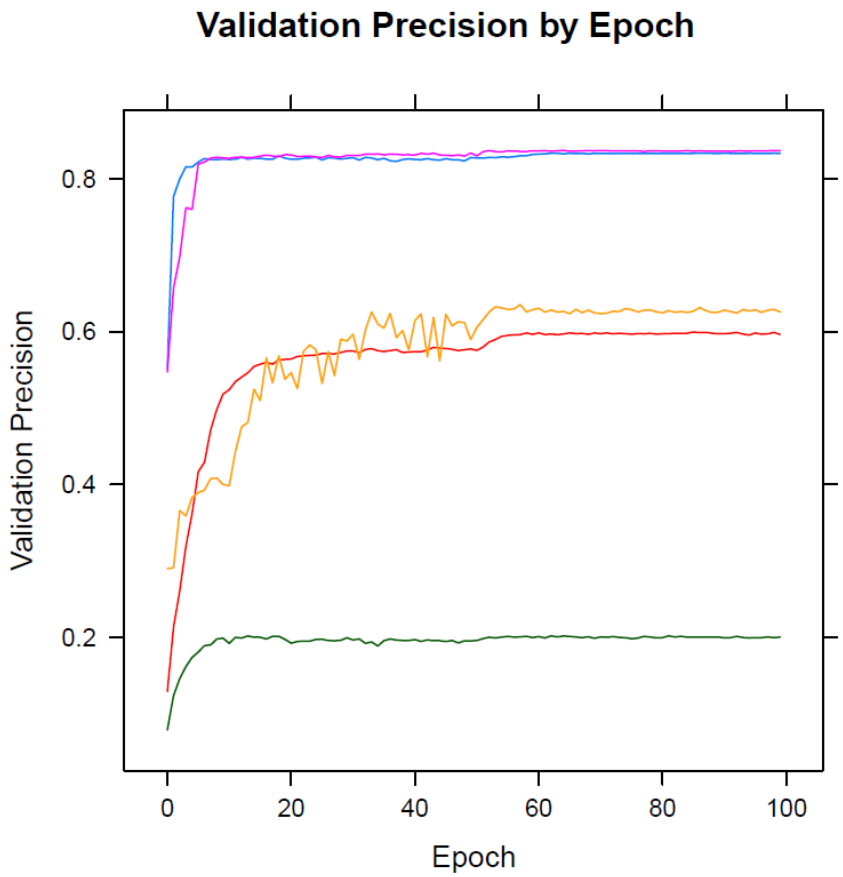

source: Evermann et al. (2021)

XAI-framework for explaininig ML-based tax decisions

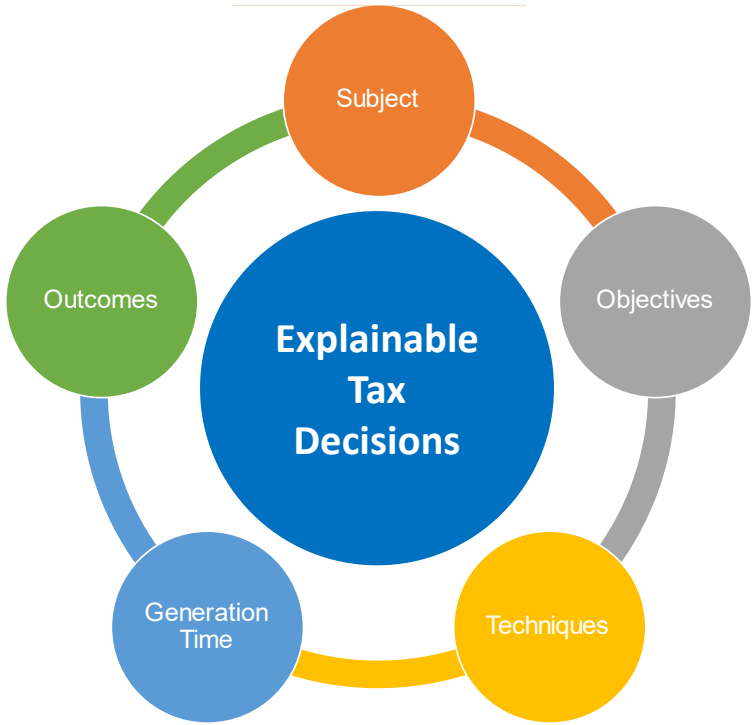

	XAI Scenario 1	XAI Scenario 2	XAI Scenario 3	XAI Scenario 4
Audience	Tax Auditors	Management	Regulatory Authorities	AI Developers
Goal	Justification	Strategical Process Enhancement	Validation	Verification
Approach	Local Post-Hoc Explanation Methods	Global Post-Hoc Explanation Methods	Global and Local Post-Hoc Explanation Methods	Intrinsic Explanation Methods
Outcome	Understanding the model decisions on identified cases	Understanding Influences of Important Input Parameters	Understanding Algorithmic Fairness Violations	Debugging and Improving the Model Performance
Time	Post-Model	Post-Model	Pre-Model and Post-Model	In-model

based on: Mehdiyev, Fettke (2009) and Mehdiyev et al. (2021)

Explanation of feature importance

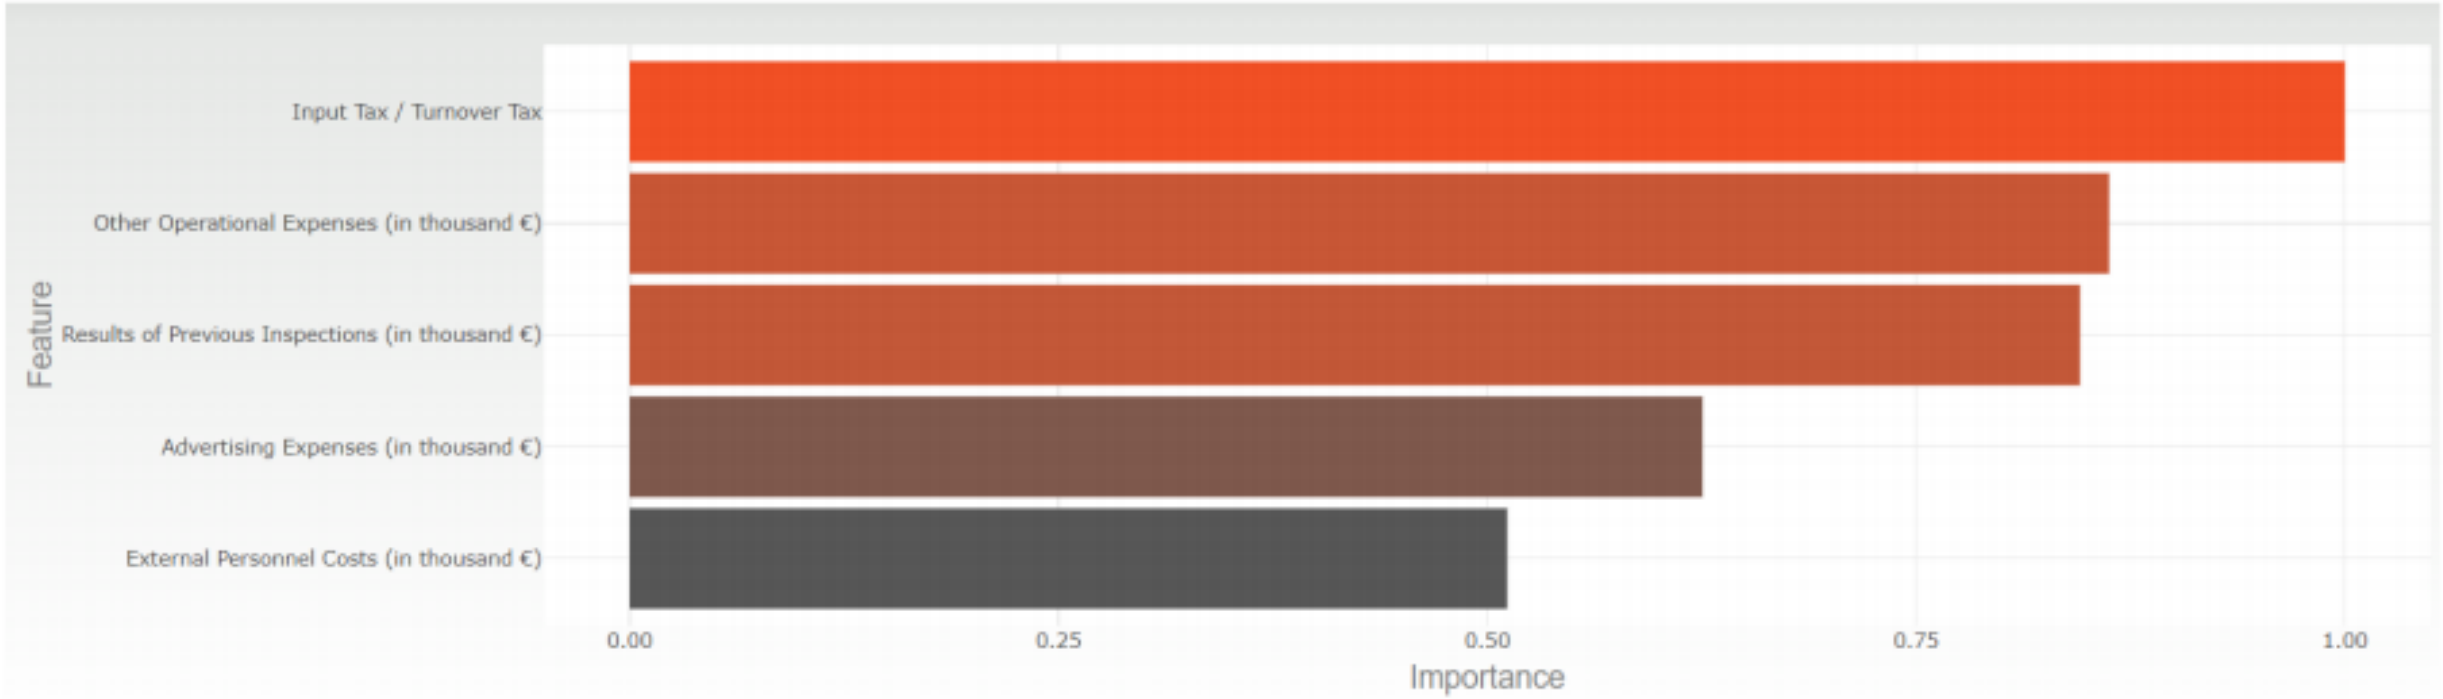

source: Mehdiyev et al. (2021)

More explanation approaches

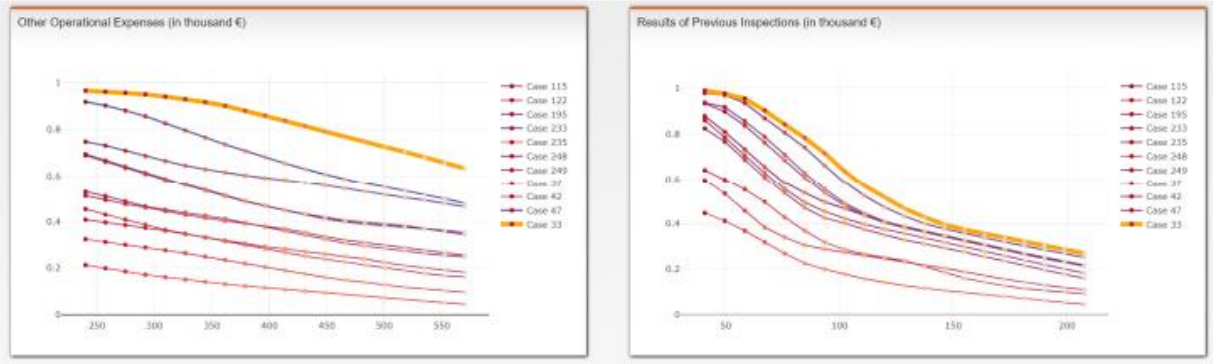

Individual Conditional Expectation (ICE) Plots

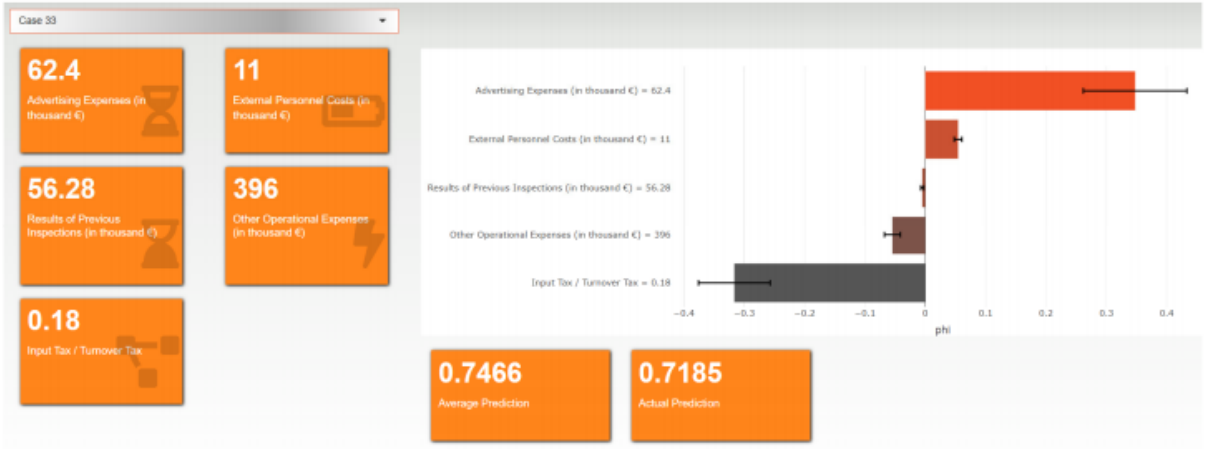

Shapley Values

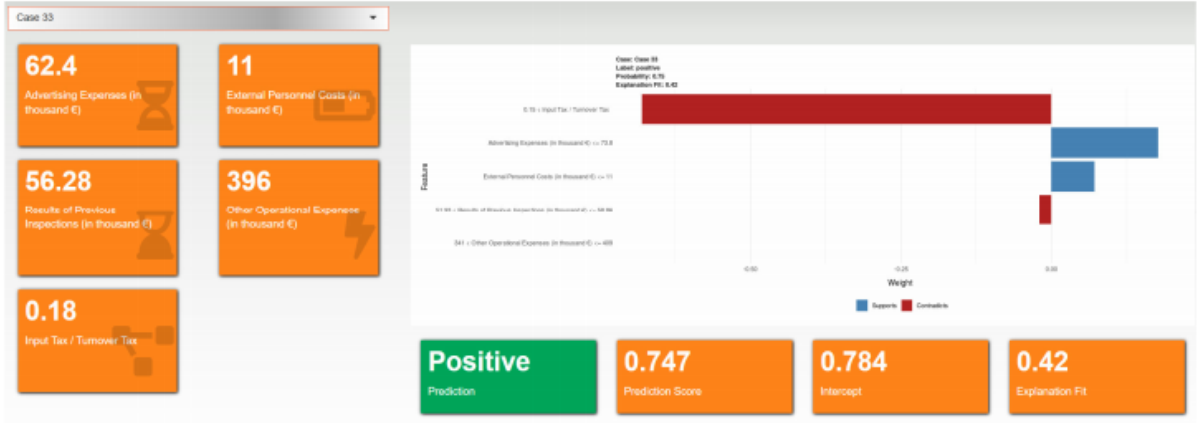

LIME

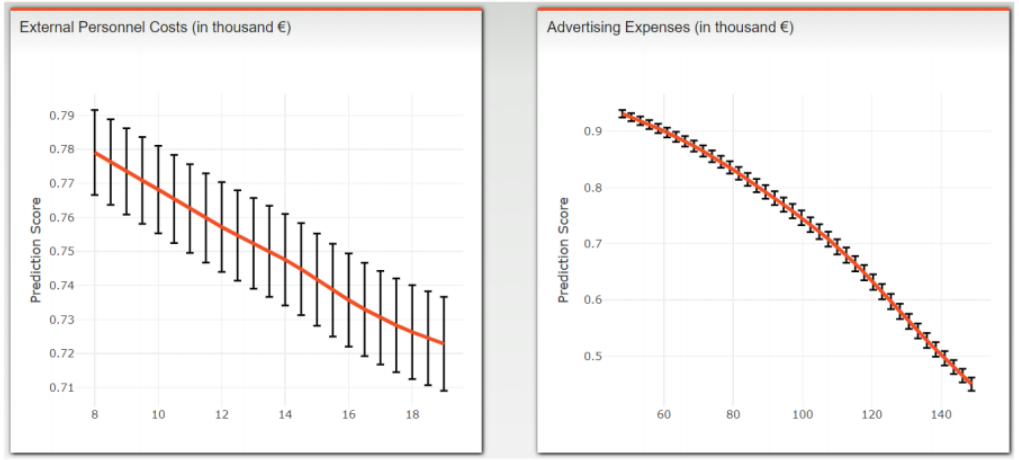

Partial Dependence Plots (PDP)

source: Mehdiyev et al. (2021)

Agenda

1. Artificial Intelligence and taxation
2. Explainable Artificial Intelligence (XAI)
3. XAI in tax audit
4. Conclusion and outlook

Guidelines for understanding and designing XAI tax systems

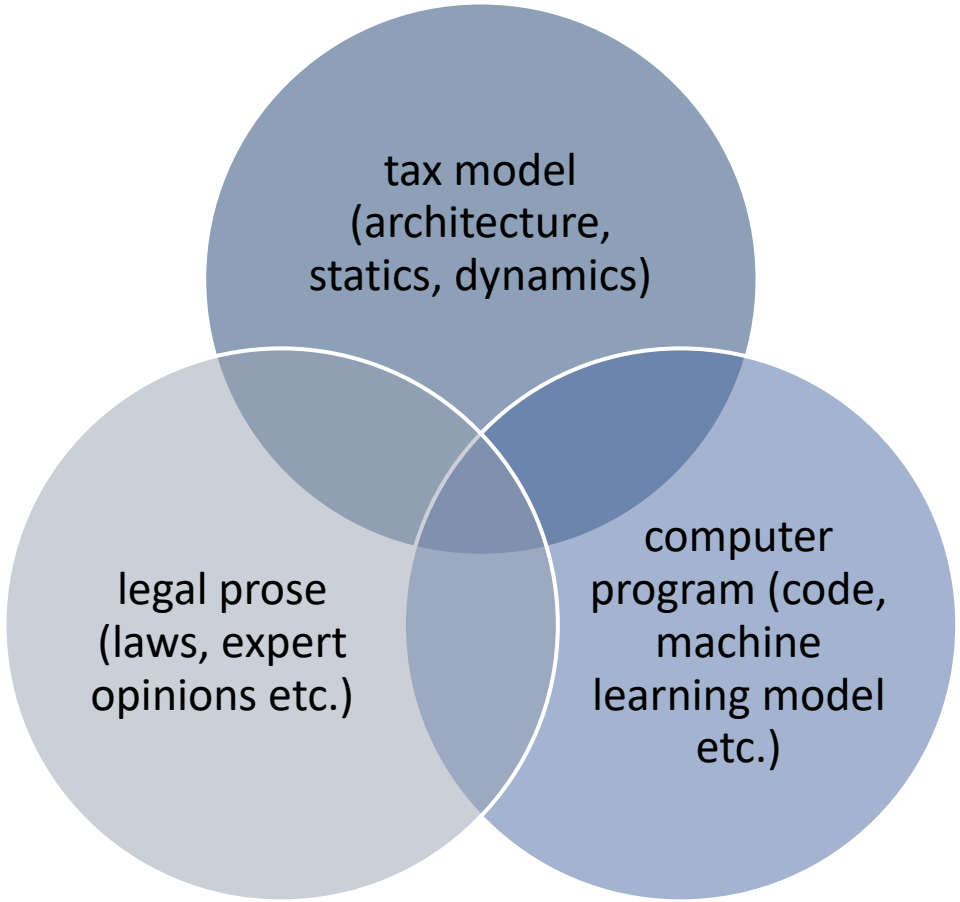

Guideline I: Develop tax model

- Model architecture of tax system
- Model tax data
- Model tax processes

Guideline II: Communicate explanations

- Design specific user interfaces for XAI-powered solutions
- Incorporate findings from cognitive sciences to transfer the generated explanations effectively
- Explore relevant provision mechanisms considering process characteristics

Guideline III: Generate and evaluate explanations

- No “one fits all” XAI solution
- Explanation generation and evaluation should be approached more holistically, considering users’ mental models, situation context, and other relevant aspects

Guideline IV: Develop and evaluate machine learning model

- Interpretable models should be first explored and used if they are capable of delivering relevant outcomes
- Need for a black-box model should be checked in terms of business and technical success criteria

based on: Fettke, Mehdiyev (2022)

TaxTech – The fourth discipline of taxation and the TaxTech-House of the digital tax function for *good tax governance*

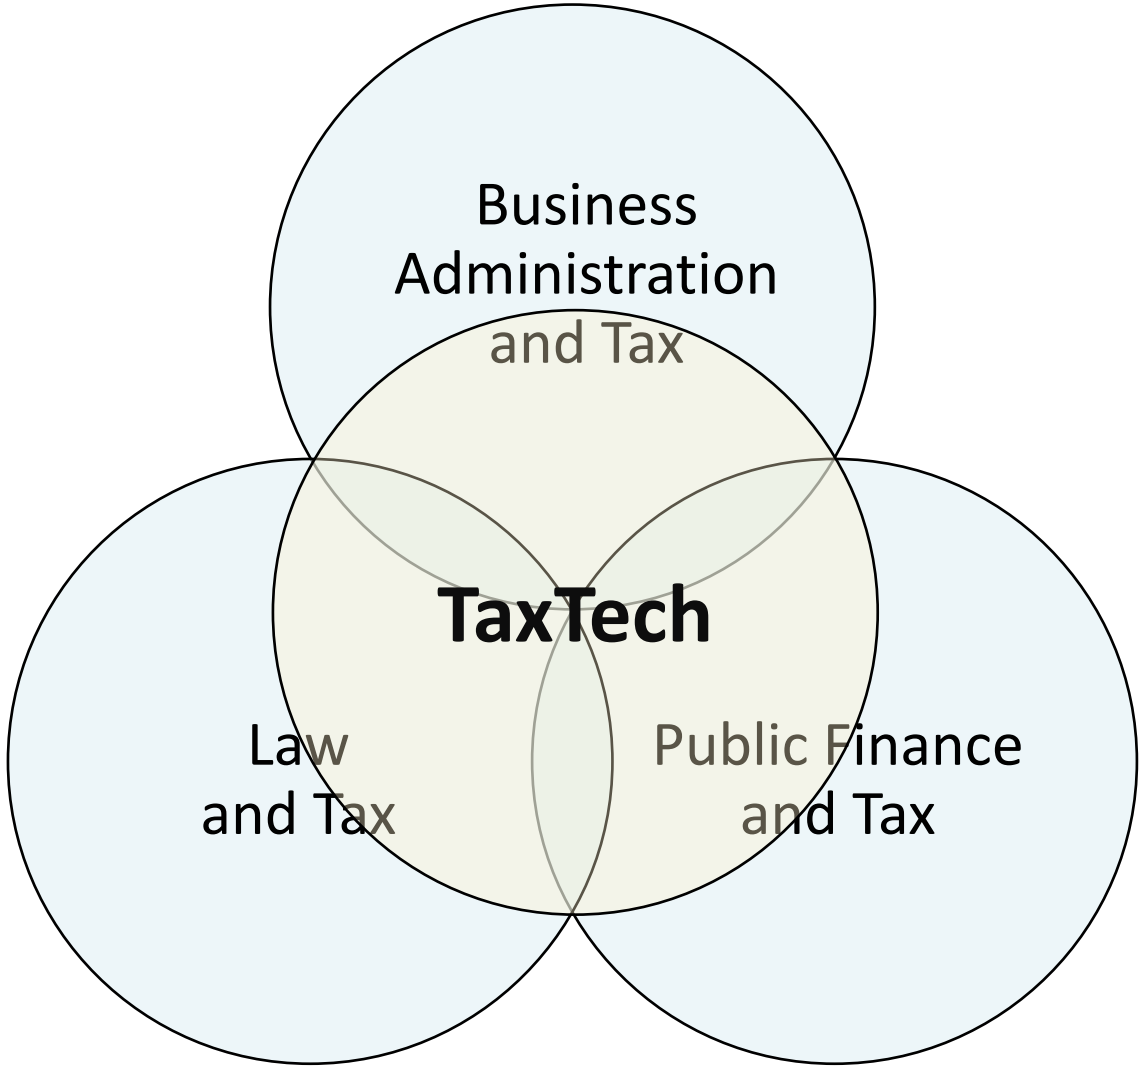

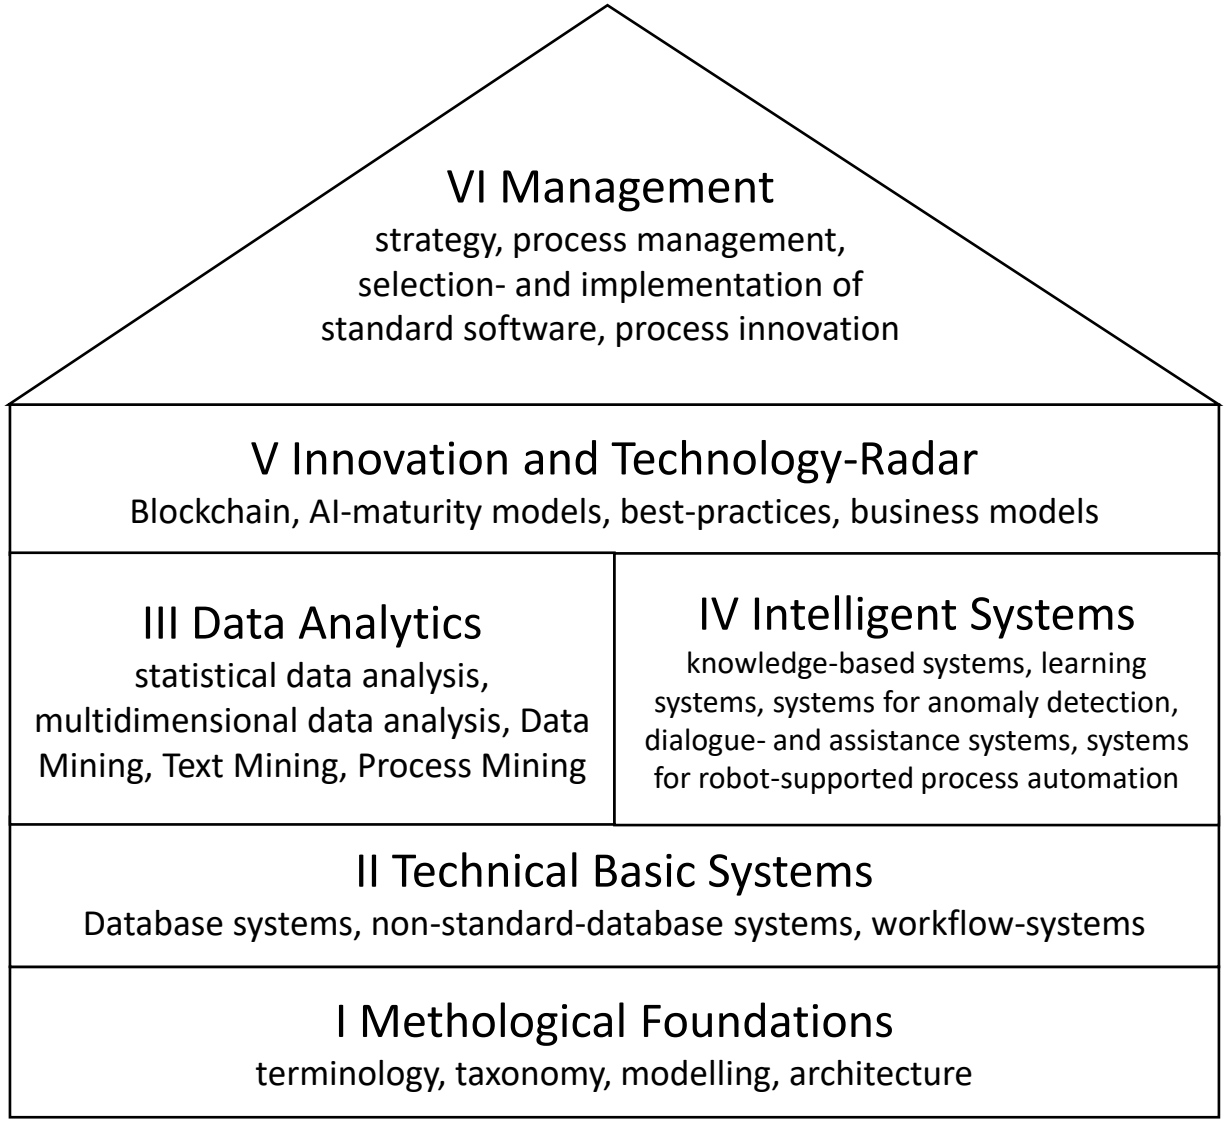

source: Fettke (2019)

Thank you for your attention!

Contact
Professor Dr. Peter Fettke
German Research Center for Artificial Intelligence (DFKI) and
Saarland University
Campus D 3 2
66123 Saarbruecken, Germany
peter.fettke@dfki.de, http://www.dfki.de
Phone: +49 681 85775-5142, Fax: -3696

UNIVERSITÄT DES SAARLANDES

DFKI