



**HAL**  
open science

## Explainable Artificial Intelligence: Concepts, Applications, Research Challenges and Visions

Luca Longo, Randy Goebel, Freddy Lecue, Peter Kieseberg, Andreas Holzinger

► **To cite this version:**

Luca Longo, Randy Goebel, Freddy Lecue, Peter Kieseberg, Andreas Holzinger. Explainable Artificial Intelligence: Concepts, Applications, Research Challenges and Visions. 4th International Cross-Domain Conference for Machine Learning and Knowledge Extraction (CD-MAKE), Aug 2020, Dublin, Ireland. pp.1-16, 10.1007/978-3-030-57321-8\_1 . hal-03414756

**HAL Id: hal-03414756**

**<https://inria.hal.science/hal-03414756>**

Submitted on 4 Nov 2021

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

# Explainable Artificial Intelligence: concepts, applications, research challenges and visions

Luca Longo<sup>1</sup>[0000-0002-2718-5426], Randy Goebel<sup>2</sup>[0000-0002-0739-2946], Freddy Lecue<sup>3,4</sup>[0000-0003-2763-7856], Peter Kieseberg<sup>5</sup>[0000-0002-2847-2152], and Andreas Holzinger<sup>2,6</sup>[0000-0002-6786-5194]

<sup>1</sup> School of Computer Science, Technological University Dublin, Ireland

`luca.longo@tudublin.ie`

<sup>2</sup> xAI-Lab, Alberta Machine Intelligence Institute

University of Alberta, Edmonton, Canada

`rgoebel@ualberta.ca`

<sup>3</sup> INRIA, Sophia Antipolis, France

<sup>4</sup> Thales, Montreal, Canada

`freddy.lecue@inria.fr`

<sup>5</sup> JRC Blockchains, University of Applied Sciences St. Pölten, Austria

`Peter.Kieseberg@fhstp.ac.at`

<sup>6</sup> Human-Centered AI Lab, Institute for Medical Informatics, Statistics & Documentation, Medical University Graz, Austria

`andreas.holzinger@medunigraz.at`

**Abstract.** The development of theory, frameworks and tools for Explainable AI (XAI) is a very active area of research these days, and articulating any kind of coherence on a vision and challenges is itself a challenge. At least two sometimes complementary and colliding threads have emerged. The first focuses on the development of pragmatic tools for increasing the transparency of automatically learned prediction models, as for instance by deep or reinforcement learning. The second is aimed at anticipating the negative impact of opaque models with the desire to regulate or control impactful consequences of incorrect predictions, especially in sensitive areas like medicine and law. The formulation of methods to augment the construction of predictive models with domain knowledge can provide support for producing human understandable explanations for predictions. This runs in parallel with AI regulatory concerns, like the European Union General Data Protection Regulation, which sets standards for the production of explanations from automated or semi-automated decision making. Despite the fact that all this research activity is the growing acknowledgement that the topic of explainability is essential, it is important to recall that it is also among the oldest fields of computer science. In fact, early AI was re-traceable, interpretable, thus understandable by and explainable to humans. The goal of this research is to articulate the big picture ideas and their role in advancing the development of XAI systems, to acknowledge their historical roots, and to emphasise the biggest challenges to moving forward.

**Keywords:** Explainable Artificial Intelligence, Machine Learning, Explainability;

## 1 Introduction

Machine learning is often viewed as the technology belonging to the future in many application fields [47], ranging from pure commodities like recommender systems for music, to automatic diagnosis of cancer or control models for autonomous transportation. However, one fundamental issue lies within the realm of explainability [62]. More precisely, most of the existing learning algorithms can often lead to robust and accurate models from data, but in application terms, they fail to provide end-users with descriptions on how they built them, or to produce convincing explanations for their predictions [7]. In many sensitive applications, such as in medicine, law, and other sectors where the main workers are not computer scientists or engineers, the direct application of these learning algorithms and complex models, without human oversight, is currently inappropriate. The reasons are not only technical, like the accuracy of a model, its stability to decisions and susceptibility to attacks, but often arise from sociological concerns, practically settling on the issue of trust. In fact, one of the principal reasons to produce an explanation is to gain the trust of users [13]. Trust is the main way to enhance the confidence of users with a system [68] as well as their comfort while using and governing it [42]. Trust connects to ethics and the intensity of regulatory activities, as for instance the General Data Protection Regulation in the European Union, leads to many legal and even ethical questions: responsibility for safety, liability for malfunction, and tradeoffs therein must inform decision makers at the highest level. Many methods of explainability for data-driven models have emerged in the years, at a growing rate. On the one hand, a large body of work have focused on building post-hoc methods mainly aimed at wrapping fully trained models, often referred to black-boxes, with an explainability layer [38]. A smaller body of research works, on the other hand, have concentrated on creating self-explainable and interpretable models by incorporating explainability mechanisms during their training, often referred to as the ante-hoc phase [7]. Despite the fact that all this research activity is the growing acknowledgement of the topic of explainability [70], by now referred to as Explainable Artificial Intelligence (XAI) [56], it is important to recall that it is also among the oldest fields of computer science. In fact, early AI was retraceable, interpretable, thus understandable by and explainable to humans. For these reasons, many scholars have tried to review research works in the field [53, 22, 74, 1, 3]. These reviews reveals the needs for a variety of kinds of explanation, for the identification of methods for explainability and their evaluation as well as the need to calibrate the tradeoffs in the degree or level of explanation appropriate for a broad spectrum of applications.

The goal of this research is to articulate the big picture ideas and their role in advancing the development of XAI systems, to acknowledge their historical roots, and to emphasise the biggest challenges to moving forward. The reminder of the paper focuses on relevant notions and concepts for explainability in section 2. It then continues in section 3 with descriptions on the applications of methods for XAI and on domains and areas in which these can have a significant impact.

A discussion on the research challenges surrounding XAI is presented in section 4. Eventually, recommendations and visions follow by presenting what we believe scholars should focus on in the development of future explainable AI systems.

## 2 Notions and related concepts

A serious challenge for any attempt to articulate the current concepts for XAI is that there is a very high volume of current activity, both on the research side [68, 74, 22], and in aggressive industrial developments, where any XAI functions can provide a market advantage to all for profit applications of AI [23, 62]. In addition, there remains a lack of consensus on terminology, for example as noted within, there are a variety of definitions for the concept of *interpretation*, but little current connection to the formal history, that means in formal theories of explanation or causation [31]. One recent paper [4] provides an organizing framework based on comparing levels of explanation with levels of autonomous driving. The goal is to identify foundational XAI concepts like relationships to historical work on explanation, especially scientific ones, or the importance of interactive explanation as well as the challenge of their evaluation. Note further that the need for a complex system to provide explanations of activities, including predictions, is not limited to those with components created by machine learning (example in [55]). Pragmatically, the abstract identification of a scientific explanation that enables an explainee to recreate an experiment or prediction can arise in very simple circumstances. For example, one can evaluate an explanation by simply noting whether it is sufficient to achieve an explainee’s intended task. For example, in figure 1, the pragmatic value of an Ikea visual assembly “explanation” is whether the assembler explainee can achieve the assembly using the diagram.

Overall and within this broad spectrum of ideas related to explanation, there is some focus on the foundational connection between explanation and that of abductive reasoning. For example, the historical notion of *scientific explanation* has been the subject of much debate in the community of science and philosophy [72]. Some propose that a theory of explanation should include both scientific and other simpler forms of explanation. Consequently, it has been a common goal to formulate principles that can confirm an explanation as a scientific one. Aristotle is generally considered to be the first philosopher to articulate an opinion that knowledge becomes scientific when it tries to explain the causes of “why.” His view urges that science should not only keep facts, but also describe them in an appropriate explanatory framework [15]. In addition to this theoretical view, empiricists also maintain a belief that the components of ideas should be acquired from perceptions with which humans become familiar through sensory experience. The development of the principles of scientific explanation from this perspective prospered with the so-called Deductive-Nomological (DN) model that was described by Hempel in [24–26], and by Hempel and Oppenheim in [27].

There is a more pragmatic AI historical research thread that connects scientific explanation to AI implementations of abductive reasoning. One such thread, among many, begins with Pople in 1973 [61], Poole et al. in 1987 [60], Muggleton in 1991 [54], to Evans et al. in 2018 [14]. Pople described an algorithm for abduction applied to medical diagnosis. Poole et al. provided an extension to first order logic which could subsume non-monotonic reasoning theories and also identify explanatory hypothesis for any application domain. Muggleton proposed a further refinement referred to as inductive logic programming where hypotheses are identified by inductive constraints within any logic, including higher-order logics. Finally, the adoption of this thread of reasoning have been generalised to explanation based on inductive logic programming by Evans et al. [14]. This most recent work connects with information theoretic ideas used to compare differences in how to learn probability distributions that are modeled by machine learning methods.

Interpreting and explaining a model trained from data by employing a machine learning technique is not an easy task. A body of literature has focused on tackling this by attempting at defining the concept of *interpretability*. This has led to the formation of many types of explanation, with several attributes and structures. For example, it seems to human nature to assign causal attribution of events [23], and we possess an innate psychological tendency to anthropomorphism. As a consequence, an AI-based system that purports to capture *causal relations* should be capable of providing a causal explanation of its inferential process (example in [57]). Causality can be considered a fundamental attribute of explainability, especially when scientific explainability carries a responsibility to help the explainee reconstruct the inferential process leading to a prediction. Many have noted this role on how explanations should make the causal relationships between the inputs and the outputs of a model explicit [53, 17, 42, 31].

Despite the fact that data-driven models are extremely good at discovering associations in the data, unfortunately they can not guarantee causality of these associations. The objective of significantly inferring causal relationships depends on prior knowledge, and very often some of the discovered associations might be completely unexpected, not interpretable nor explainable. As pointed by [1], the decisions taken considering the output of a model should be clearly explainable to support their *justifiability*. These explanations should allow the identification of potential flows both in a model, enhancing its *transparency*, the knowledge *discovery* process, supporting its *controllability* and *improvement* of its accuracy. Although the importance of explainability is clear, the definition of objective criteria to evaluate methods for XAI and validate their explanations is still lacking. Numerous notions underlying the effectiveness of explanations were identified from the fields of Philosophy, Psychology and Cognitive Science. These were related to the way humans define, generate, select, evaluate and present explanations [51].

### 3 Applications and impact areas

Explainable artificial intelligence has produced many methods so far and it has been applied in many domains, with different expected impacts [21]. In these applications, the production of explanations for black box predictions requires a companion method to extract or lift correlative structures from deep-learned models into vocabularies appropriate for user level explanations. Initial activities focused on deep learning image classification with explanations emerging as heat maps created on the basis of gaps in probability distributions between a learned model and an incorrect prediction [5]. However, the field has become so diverse in methods, often determined by domain specific issues and attributes, that it is scarcely possible to get in-depth knowledge on the whole of it. Additionally, one major aspect though is the problem of explainable AI, where lot of problems have been emerged and illustrated in the literature, especially from not being able to provide explanations. While all of these topics require long and in-depth discussions and are certainly of significant importance for the future of several AI methods in many application domains, we want to focus on the benefits that can be reaped from explainability. This means not focusing on the issues of incomplete and imperfect technologies as a stopping point for applications, but discussing novel solutions provided by explainable AI. A discussion of some, partly prominent and partly surprising examples follows, with arguments on why a certain amount of explainability - as a reflection - is required for more advanced AI. There are many sectors that already have fully functional applications based on machine learning, but still serious problems in applying them exist. These are often caused by failing to be capable to explain how these methods work. In other words, it is known that they work, but the concrete results cannot be explained. Many of these applications either come from safety critical or personally sensitive domains, thus a lot of attention is put on explanations of the inferences of trained models, usually predictions or classifications.

**Threat detection and triage** - The detection of threats and efficient triage have been core topics in the area of IT-Security for at least the past three decades. This started with research in the area of code analysis and signature based AntiVirus-Software, moving towards automated decompilation and code analysis, as well as supporting the automated analysis of network monitoring information for triage. Currently, fully automated threat detection and triage is not available in real life systems due to the complexity of the task and the problem with false positives, even though several different approaches exist. These also include strategies that do not try to detect actual threats, but rather filtering out all known legit network travel and thus drastically reducing the amount of information requiring manual analysis [58]. Still, a major problem without explainability lies in the opaque nature of these methods, thus not being able to fully understand their inner functioning and how an inference was reached. Explainability could greatly enhance the detection capabilities, especially since dynamic effects, such as changing user behavior, could be modelled and introduced earlier into the algorithms without generating a large set of false positives.

**Explainable Object Detection** - Object detection is usually performed from a large portfolio of artificial neural networks (ANN) architectures such as YOLO, trained on large amount of labelled data. In such contexts, explaining object detections is rather difficult if not impossible due to the high complexity of the hyperparameters (number of layers, filters, regularisers, optimizer, loss function) of the most accurate ANNs. Therefore, explanations of an object detection task are limited to features involved in the data and modeled in the form of saliency maps [11] or at best to examples [41], or prototypes [36]. They are the state-of-the-art approaches but explanations are limited by data frames feeding the ANNs. Industrial applications embedding object detection, such as obstacles detection for trains, do require human-like rational for ensuring the system can be guaranteed, even certified [40].

**Protection against Adversarial ML** - In adversarial machine learning, attackers try to manipulate the results of learning algorithms by inserting specifically crafted data in the learning process [33], in order to lead a model to learn erroneous things. Detection of such a manipulation is not trivial, especially in contexts with big data, where no model exists before the analysis phase. While there are several proposals on how to deal with this issue [16], some of them employ neural sub-networks for differentiating between malicious and benign input data like [50]. In this specific circumstance, explainability would have a great impact as it will support the task of uncovering such a manipulation far more quickly, efficiently and without actually finding the examples that have been manipulated, thus greatly enhancing trust in machine learning inferences [32].

**Open Source Intelligence (OSINT)** - In Open Source Intelligence [19], information retrieval is purely reduced to openly available information, as contrary to Signals Intelligence (SIGINT). However, there are several major issues surrounding OSINT, especially referring to context, languages and the amount of information available. Similarly, another problem lies in deciding how much a source is trusted, and what level of impact news of sources shall have on the result of their aggregations. This is especially important when considering adversarial attacks against OSINT methods and systems [12]. Explainability could provide means for detecting these attacks, with an impact on mitigating their influence. Furthermore, the information that an attack against an intelligent system was launched is also a valuable input from an intelligence perspective, so explainability might lead to additional valuable information. However, not all false information exists due to malice, especially when reporting very recent events: information particles might be wrong, misleading or simply unknown at the time of reporting. OSINT becomes especially complex in case of ongoing events, where facts change every minute, either due to knew intelligence, or simply because of changes in the event itself. Explainability would allow to estimate the effects of incorrect information particles on the overall machine learning outcomes, thus allowing, for instance, to give error margins on reported numbers.

**Trustworthy (autonomous) medical agents** - Several architectures for integrating machine learning into medical decision making have been devised in the past. These are based upon a doctor-in-the-loop approach whereby doctors act as input providers to machine learning algorithms. These can lead to suggestions related to diagnosis or treatment that can be subsequently reviewed by the doctors themselves, who, in turn, can provide feedback in a loop to further enhance modeling [35]. Additionally, the mechanism can also introduce external knowledge to support decision making aimed at incorporating the latest findings in the underlying medical field.

**Autonomous vehicles** - While certainly being developed within machine learning, explainability would be beneficial for the area of autonomous vehicles, especially considering autonomous cars. In cases of car accidents, explanations can help trace the reasons why an autonomous vehicle behaved in a certain way and took certain actions. Consequently this can not only lead to safer vehicles, but it also can help solve issues in court faster, greatly enhancing trust towards these novel ML-based technologies and especially the resulting artifacts [20].


## 4 Research challenges

A number of research challenges surrounding the development of methods for explainability exist, including technical, legal and practical challenges.

### 4.1 Technical challenges

**XAI systems evaluation** The comprehensive study of what explanation means from a sociological viewpoint [52] begs a difficult issue that is both technical and non-technical: *how does one evaluate the quality of an explanation?* It is not a surprise that the quality or value of an explanation is at least partly determined by the receiver of an explanation, sometimes referred to as the “explainee”. An easy way to frame the challenge of evaluating explanations, with respect to an explainee, arises from observing the history of the development of evaluation techniques from the field of data visualization [37]. A simple example of “visual explanation” can frame the general evaluation problem for all explanations as follows. Consider the IKEA assembly diagram, rendered in Figure 1. A simple list of requirements to assess explanation quality emerges from considering the IKEA assembly instructions as a visual explanation of how to assemble the piece of furniture. In this case, the visual explanation is intended to guide all explainees, and not just a single individual, to the successful assembly of the furniture item. One measure of quality is simply to test whether any individual explainee can use the visual explanation to complete the assembly. Another measure is about whether the visual explanation is clear and unambiguous, so that the assembly is time efficient. In the case of figure 1, the sequencing of steps might be misinterpreted by an explainee, and that the simple use of circular arrows to indicate motion may also be ambiguous.





**Fig. 1.** An IKEA visual explanation for furniture assembly

Overall, and as anticipated in the general evaluation of explanation systems, one could design cognitive experiments to determine, over an experimental human cohort, which portions of the explanation clearly lead to correct inferences, and those which are more difficult to correctly understand. This means that XAI system requirements should include the need to produce an explicit representation of all the components in a way that supports the appropriate interpretation of the visual classification of components. One can generalize visual explanations to the full repertoire that might obtain for a general XAI system. This means a set of representations of the semantics of an underlying domain of application that can provide support to construct an explanation that is understood by a human explainee.

**XAI Interpretation** Even though XAI systems are supposed to expose the functioning of a learning technique as well as a set of justification of a model’s inferences, it remains rather difficult for a human to interpret them. Explanations are not the final words of an intelligent system but rather the intermediate layer that requires knowledge expertise, context and common-sense characterization for appropriate and correct human interpretation and decision-making [45, 66]. Semantics, knowledge graphs [39] and their machine learning representations [6] or similar technical advancements are interesting avenues to be considered for pushing the interpretation at the next right level of knowledge expertise. These might also include the addition of argumentative capabilities, as applied in [46, 63] to produce rational and justifiable explanations [64].

## 4.2 Legal challenges

While the theoretical ground work in AI stays on the very theoretical side and is thus typically considered to be not problematic from a legal point of view, the actual application of XAI methods in a certain domain can have serious legal implications. This is especially important when considering working with sensitive information. Here, it has yet to be researched whether the explainability related to a model might be used to infer information about individuals, for instance, by using it with slightly different data sets. This technique has been used in many variations in IT-Security, especially considering anonymized data sets or

partially released sensitive information as a basis to gather more intelligence on the people involved [9]. Similar attacks have already been proposed and carried out against machine learned models [67] and these allowed to produce a great amount of information, hidden correlations and causalities that were used to infer sensitive information.

Concepts like federated machine learning are built on the notion of executing machine learning algorithms locally on sensitive data sets and then exchanging the resulting feature sets in order to be combined centrally. These are in contrast to more traditional approaches that collect all sensitive data centrally and then run the learning algorithms. One challenge for federated machine learning is to achieve model robustness but greatly focus on protective sensitive inferences. This justifies the need for more applicable anonymisation techniques, as many of the current methods are unsuitable for many application scenarios, either due to performance or quality issues [49, 48]. In addition, other legal challenges exist such as the right to be forgotten [69]. This ‘reflects the claim of an individual to have certain data deleted so that third persons can no longer trace them’. This fair right is accompanied by technical difficulties ranging from the issue related to the deletion of entries in modern systems, to the problem of inferring information on individuals from aggregates and especially the removal of said individuals from the aggregation process.

Despite the aforementioned challenges, positive benefits can be brought by explainability to the area of machine learning and AI as a whole with respect to legal issues. While the issue of transparency, a key requirement in the General Data Protection Regulation (GDPR), can be a rather a hard issue to tackle, this could change with explainability providing detailed insight, where, when and to what extent personal data of a single individual was involved in a data analysis workflow [71]. While this is currently not a binding requirement to provide that level of details [71], this could be a game changer regarding acceptance of AI, as well as increasing privacy protection in a data driven society. Furthermore, a significant problem currently tackled in machine learning is bias [73], especially since simple methods for tackling the issue have shown to be ineffective [34]. Explainability could support this combat and thus provide a better legal standing for the results derived from data driven systems, especially when used for socio-economic purposes.

### 4.3 Practical challenges

One of the most crucial success factors of AI generally and XAI specifically, is to ensure effective human-AI interfaces to enable a usable and useful interaction between humans and AI [2]. Such goals have been discussed in the HCI community for decades [10], but it was not really seen as important in the AI community. Now the needs and demands of XAI for ‘explainable user interfaces’ may finally stimulate to realise advanced human-centered concepts similar to the early visions of Vannevar Bush in 1945 [8]. Here, the goal is to explore both the

explainability side, that means the artificial explanation generated by machines, as well as the human side, that means the human understanding. In an ideal world, both machine explanations and human understanding would be identical, and congruent with the ground truth, which is defined for both machines and humans equally. However, in the real world we face two significant problems:

- the ground truth cannot always be fully defined, as for instance when concerned with medical diagnoses [59] when there is high uncertainty;
- human models such as scientific, world, problem solving models, are often based on causality, in the sense of Judea Pearl [57], which is very challenging as current machine learning does not incorporate them and simply follows pure correlation.

Practically speaking, current XAI methods mainly focus on highlighting input-relevant parts, for example via heat-mapping, that significantly contributed to a certain output, or the most relevant features of a training data set that influenced the most the model accuracy. Unfortunately, they do not incorporate the notion of human model, and therefore there is a need to take also into account the concept of causability [30]. In detail, in line with the concept of usability [28], causability is defined as ‘the extent to which an explanation of a statement to a human expert achieves a specified level of causal understanding with effectiveness, efficiency and satisfaction in a specified context of use and the particular contextual understanding capabilities of a human’. Following this concept, it becomes possible to measure the quality of explanations in the same terms as usability (effectiveness, efficiency and satisfaction in a specified context of use), for example with a measurement scale [29]

## 5 Recommendations and visions

Machine learning, as a solid research area within artificial intelligence, has undoubtedly impacted the field by providing scholars with a robust suite of methods for modeling complex, non-linear phenomena. With the growing body of work in the last decade on deep learning, this impact has significantly expanded to many applications areas. However, despite the widely acknowledged capability of machine and deep learning to allow scholars to induce accurate models from data and extract relevant patterns, accelerating scientific discovery [18], there is the problem of their interpretability and explainability. For this reason, the last few years have seen a growing body of work on research in methods aimed at explaining the inner functioning of data-driven models and the learning techniques used to induce them. Currently and generally recognised as a core area of AI, eXplainable Artificial Intelligence (XAI) has produced a plethora of methods for model interpretability and explainability. Hundred of scientific articles are published each month in many workshops, conferences and presented at symposium around the world. Some of them focus on wrapping trained models with explanatory layers, such as knowledge graphs [40]. Other try to embed the concept of explainability during training, and some of them try to merge learning

capabilities with symbolic reasoning [44]. Explainability is a concept borrowed from psychology, since it is strictly connected to humans, that is difficult to operationalise. A precise formalisation of the construct of explainability is far from being a trivial task as multiple attributes can participate in its definition [63]. Similarly, the attributes might interact with each other, adding complexity in the definition of an objective measure of explainability [43, 65]. For these reasons, the last few years have seen also a growing body of research on approaches for evaluating XAI methods. In other words, approaches that are more focused on the explanations generated by XAI solutions, their structure, efficiency, efficacy and impact on humans understanding.

The first recommendation to scholars willing to perform scientific research on explainable artificial intelligence and create XAI methods is to firstly focus on the structure of explanations, the attributes of explainability and the way they can influence humans. This links computer science with psychology. The second recommendation is to define the context of explanations, taking into consideration the underlying domain of application, who they will serve and how. Ultimately, explanations are effective when they help end-users to build a complete and correct mental representation of the inferential process of a given data-driven model. Work on this direction should also focus on which type of explanation can be provided to end-users, including textual, visual, numerical, rules-based or mixed solutions. This links computer science with the behavioural and social sciences. The third recommendation is to clearly define the scope of explanations. This might involve the creation of a method that provide end-users with a suite of local explanations for each input instance or the formation of a method that focuses more on generating explanations on a global level aimed at understanding a model as a whole. This links computer science to statistics and mathematics. The final recommendation is to involve humans, as ultimate users of XAI methods, within the loop of model creation, exploitation, as well as the enhancement of its interpretability and explainability. This can include the development of interactive interfaces that allow end-users to navigate through models, understanding their inner logic at a local or global level, for existing or new input instances. This links artificial intelligence with human-computer interaction.

The visions behind explainable artificial intelligence are certainly numerous. Probably the most important is the creation of models with high accuracy as well as high explainability. The trade-off between these two sides is well known, and usually, increments in one dimension means decrements in the other dimension. Creating interpretable and explainable models that are also highly accurate is the ideal scenario, but since this has been demonstrated to be a hard problem with current methods of learning and explainability, further research is needed. One possible solution is the creation of models that are fully transparent at all stages of model formation, exploitation and exploration and that are capable of providing local and global explanations. This leads to another vision, which

is the use of methods that embed learning capabilities and symbolic reasoning. The former is aimed at generating models and representations with high accuracy for predictive and forecasting purposes, while the latter to explain these representations in highly interpretable natural language terms, aligned to the way human understand and reason.

## Acknowledgements

R.Goebel would like to acknowledge the support of the Alberta Machine Intelligence Institute, which is one of the three Pan Canadian AI Centres. A.Holzinger would like to acknowledge the support of the Austrian Science Fund (FWF), Project: P-32554 “Explainable AI - A reference model of explainable Artificial Intelligence for the Medical Domain”

## References

1. Amina Adadi and Mohammed Berrada. Peeking inside the black-box: A survey on explainable artificial intelligence (xai). *IEEE Access*, 6:52138–52160, 2018.
2. Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N. Bennett, and Kori Inkpen. Guidelines for human-ai interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, 2019.
3. Leila Arras, Ahmed Osman, Klaus-Robert Müller, and Wojciech Samek. Evaluating recurrent neural network explanations. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 113–126, Florence, Italy, 2019. Association for Computational Linguistics.
4. S. Atakishiyev, H. Babiker, N. Farruque, R. Goebel, M-Y. Kim, M.H. Motallebi, J. Rabelo, T. Syed, and O. R. Zaïane. A multi-component framework for the analysis and design of explainable artificial intelligence. (*arXiv:2005.01908v1 [cs.AI]*), 2020.
5. Housam Khalifa Bashier Babiker and Randy Goebel. An introduction to deep visual explanation. *NIPS 2017 - Workshop Interpreting, Explaining and Visualizing Deep Learning*, 2017.
6. Federico Bianchi, Gaetano Rossiello, Luca Costabello, Matteo Palmonari, and Pasquale Minervini. Knowledge graph embeddings and explainable AI. *CoRR*, abs/2004.14843, 2020.
7. Or Biran and Courtenay Cotton. Explanation and justification in machine learning: A survey. In *IJCAI 2017 Workshop on Explainable Artificial Intelligence (XAI)*, pages 8–13, Melbourne, Australia, 2017. International Joint Conferences on Artificial Intelligence, Inc.
8. Vannevar Bush. As we may think. *Atlantic Monthly*, 176(1):101–108, 1945.
9. Zhipeng Cai, Zaobo He, Xin Guan, and Yingshu Li. Collective data-sanitization for preventing sensitive information inference attacks in social networks. *IEEE Transactions on Dependable and Secure Computing*, 15(4):577–590, 2016.
10. Stuart K. Card, Thomas P. Moran, and Alan Newell. *The psychology of Human-Computer Interaction*. Erlbaum, Hillsdale (NJ), 1983.

11. Chun-Hao Chang, Elliot Creager, Anna Goldenberg, and David Duvenaud. Interpreting neural network classifications with variational dropout saliency maps. *Proc. NIPS*, 1(2):1–9, 2017.
12. Sean M Devine and Nathaniel D Bastian. Intelligent systems design for malware classification under adversarial conditions. *arXiv preprint arXiv:1907.03149*, 2019.
13. Mary T Dzindolet, Scott A Peterson, Regina A Pomranky, Linda G Pierce, and Hall P Beck. The role of trust in automation reliance. *International journal of human-computer studies*, 58(6):697–718, 2003.
14. Richard Evans and Edward Greffentette. Learning explanatory rules from noisy data. *Journal of Artificial Intelligence Research*, 61:1–64, 2018.
15. Andrea Falcon. Aristotle on causality. *Stanford Encyclopedia of Philosophy*, 2006. (<https://plato.stanford.edu>).
16. Reuben Feinman, Ryan R Curtin, Saurabh Shintre, and Andrew B Gardner. Detecting adversarial samples from artifacts. *arXiv preprint arXiv:1703.00410*, 2017.
17. Maria Fox, Derek Long, and Daniele Magazzeni. Explainable planning. In *IJCAI 2017 Workshop on Explainable Artificial Intelligence (XAI)*, pages 24–30, Melbourne, Australia, 2017. International Joint Conferences on Artificial Intelligence, Inc.
18. Yolanda Gil, Mark Greaves, James Hendler, and Haym Hirsh. Amplify scientific discovery with artificial intelligence. *Science*, 346(6206):171–172, 2014.
19. Michael Glassman and Min Ju Kang. Intelligence in the internet age: The emergence and evolution of open source intelligence (osint). *Computers in Human Behavior*, 28(2):673–682, 2012.
20. Jon Arne Glomsrud, André Ødegårdstuen, Asun Lera St Clair, and Øyvind Smogeli. Trustworthy versus explainable ai in autonomous vessels. In *Proceedings of the International Seminar on Safety and Security of Autonomous Vessels (ISSAV) and European STAMP Workshop and Conference (ESWC) 2019*, pages 37–47. Sciendo, 2020.
21. Randy Goebel, Ajay Chander, Katharina Holzinger, Freddy Lecue, Zeynep Akata, Simone Stumpf, Peter Kieseberg, and Andreas Holzinger. Explainable ai: the new 42? In *International Cross-Domain Conference for Machine Learning and Knowledge Extraction*, pages 295–303. Springer, 2018.
22. Riccardo Guidotti, Anna Monreale, Salvatore Ruggieri, Franco Turini, Fosca Gian-notti, and Dino Pedreschi. A survey of methods for explaining black box models. *ACM computing surveys (CSUR)*, 51(5):93:1–93:42, 2018.
23. Taehyun Ha, Sangwon Lee, and Sangyeon Kim. Designing explainability of an artificial intelligence system. In *Proceedings of the Technology, Mind, and Society*, page 14:1, Washington, District of Columbia, USA, 2018. ACM.
24. Carl G Hempel. The function of general laws in history. *The Journal of Philosophy*, 39(2):35–48, 1942.
25. Carl G. Hempel. The theoretician’s dilemma: A study in the logic of theory construction. *Minnesota Studies in the Philosophy of Science*, 2:173–226, 1958.
26. Carl G Hempel. Aspects of scientific explanation. *Free Press New York*, 1965.
27. Carl G Hempel and Paul Oppenheim. Studies in the logic of explanation. *Philosophy of science*, 15(2):135–175, 1948.
28. Andreas Holzinger. Usability engineering methods for software developers. *Communications of the ACM*, 48(1):71–74, 2005.
29. Andreas Holzinger, Andre Carrington, and Heimo Müller. Measuring the quality of explanations: The system causability scale (scs). comparing human and machine

- explanations. *KI - Künstliche Intelligenz (German Journal of Artificial intelligence), Special Issue on Interactive Machine Learning, Edited by Kristian Kersting, TU Darmstadt*, 34(2), 2020.
30. Andreas Holzinger, Georg Langs, Helmut Denk, Kurt Zatloukal, and Heimo Mueller. Causability and explainability of artificial intelligence in medicine. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 9(4), 2019.
  31. Andreas Holzinger, Georg Langs, Helmut Denk, Kurt Zatloukal, and Heimo Müller. Causability and explainability of artificial intelligence in medicine. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 9:e1312, 2019.
  32. Katharina Holzinger, Klaus Mak, Peter Kieseberg, and Andreas Holzinger. Can we trust machine learning results? artificial intelligence in safety-critical decision support. *ERCIM NEWS*, 112:42–43, 2018.
  33. Ling Huang, Anthony D Joseph, Blaine Nelson, Benjamin IP Rubinstein, and J Doug Tygar. Adversarial machine learning. In *Proceedings of the 4th ACM workshop on Security and artificial intelligence*, pages 43–58, 2011.
  34. Toshihiro Kamishima, Shotaro Akaho, Hideki Asoh, and Jun Sakuma. Fairness-aware classifier with prejudice remover regularizer. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 35–50. Springer, 2012.
  35. Peter Kieseberg, Bernd Malle, Peter Frühwirt, Edgar Weippl, and Andreas Holzinger. A tamper-proof audit and control system for the doctor in the loop. *Brain informatics*, 3(4):269–279, 2016.
  36. Been Kim, Oluwasanmi Koyejo, and Rajiv Khanna. Examples are not enough, learn to criticize! criticism for interpretability. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 2280–2288, 2016.
  37. Heidi Lam, Enrico Bertini, Petra Isenberg, Catherine Plaisant, and Sheelagh Cpendale. Empirical studies in information visualization: Seven scenarios. *IEEE Transactions on Graphics and Visual Computing*, 18(9):1520–1536, 2012.
  38. Thibault Laugel, Marie-Jeanne Lesot, Christophe Marsala, Xavier Renard, and Marcin Detyniecki. The dangers of post-hoc interpretability: Unjustified counterfactual explanations. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, (IJCAI)*, pages 2801–2807, Macao, China, 2019. International Joint Conferences on Artificial Intelligence Organization.
  39. Freddy Lécué. On the role of knowledge graphs in explainable AI. *Semantic Web*, 11(1):41–51, 2020.
  40. Freddy Lécué and Tanguy Pommellet. Feeding machine learning with knowledge graphs for explainable object detection. In Mari Carmen Suárez-Figueroa, Gong Cheng, Anna Lisa Gentile, Christophe Guéret, C. Maria Keet, and Abraham Bernstein, editors, *Proceedings of the ISWC 2019 Satellite Tracks (Posters & Demonstrations, Industry, and Outrageous Ideas) co-located with 18th International Semantic Web Conference (ISWC 2019), Auckland, New Zealand, October 26-30, 2019*, volume 2456 of *CEUR Workshop Proceedings*, pages 277–280. CEUR-WS.org, 2019.
  41. Oscar Li, Hao Liu, Chaofan Chen, and Cynthia Rudin. Deep learning for case-based reasoning through prototypes: A neural network that explains its predictions. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 3530–3537, 2018.
  42. Zachary C Lipton. The mythos of model interpretability. *Commun. ACM*, 61(10):36–43, 2018.

43. Luca Longo. Argumentation for knowledge representation, conflict resolution, defeasible inference and its integration with machine learning. In *Machine Learning for Health Informatics*, pages 183–208. Springer, 2016.
44. Luca Longo and Pierpaolo Dondio. Defeasible reasoning and argument-based systems in medical fields: An informal overview. In *2014 IEEE 27th International Symposium on Computer-Based Medical Systems*, pages 376–381. IEEE, 2014.
45. Luca Longo and Lucy Hederman. Argumentation theory for decision support in health-care: a comparison with machine learning. In *International Conference on Brain and Health Informatics*, pages 168–180. Springer, 2013.
46. Luca Longo, Bridget Kane, and Lucy Hederman. Argumentation theory in health care. In *Computer-Based Medical Systems (CBMS), 2012 25th International Symposium on*, pages 1–6. IEEE, 2012.
47. Spyros Makridakis. The forthcoming artificial intelligence (ai) revolution: Its impact on society and firms. *Futures*, 90:46–60, 2017.
48. Bernd Malle, Peter Kieseberg, and Andreas Holzinger. Do not disturb? classifier behavior on perturbed datasets. In *International Cross-Domain Conference for Machine Learning and Knowledge Extraction*, pages 155–173. Springer, 2017.
49. Bernd Malle, Peter Kieseberg, Edgar Weippl, and Andreas Holzinger. The right to be forgotten: towards machine learning on perturbed knowledge bases. In *International Conference on Availability, Reliability, and Security*, pages 251–266. Springer, 2016.
50. Jan Hendrik Metzen, Tim Genewein, Volker Fischer, and Bastian Bischoff. On detecting adversarial perturbations. *arXiv preprint arXiv:1702.04267*, 2017.
51. Tim Miller. Explanation in artificial intelligence: insights from the social sciences. *Artificial Intelligence*, 267:1–38, 2019.
52. Tim Miller. Explanation in Artificial Intelligence: Insights from the Social Sciences. *Artificial Intelligence*, 267:1–38, 2019.
53. Tim Miller, Piers Howe, and Liz Sonenberg. Explainable ai: Beware of inmates running the asylum or: How i learnt to stop worrying and love the social and behavioural sciences. In *IJCAI Workshop on Explainable AI (XAI)*, pages 36–42, Melbourne, Australia, 2017. International Joint Conferences on Artificial Intelligence, Inc.
54. Stephen Muggleton. Inductive logic programming. *New Generation Computing*, pages 295–318, 1991.
55. Ingrid Nunes and Dietmar Jannach. A systematic review and taxonomy of explanations in decision support and recommender systems. *User Modeling and User-Adapted Interaction*, 27(3):393–444, 2017.
56. Andrés Páez. The pragmatic turn in explainable artificial intelligence (xai). *Minds and Machines*, 29:1–19, 2019.
57. Judea Pearl. *Causality: Models, Reasoning, and Inference (2nd Edition)*. Cambridge University Press, Cambridge, 2009.
58. Martin Pirker, Patrick Kochberger, and Stefan Schwandter. Behavioural comparison of systems for anomaly detection. In *Proceedings of the 13th International Conference on Availability, Reliability and Security*, pages 1–10, 2018.
59. Birgit Pohn, Michaela Kargl, Robert Reihs, Andreas Holzinger, Kurt Zatloukal, and Heimo Müller. Towards a deeper understanding of how a pathologist makes a diagnosis: Visualization of the diagnostic process in histopathology. In *IEEE Symposium on Computers and Communications (ISCC 2019)*. IEEE, 2019.
60. David Poole, Randy Goebel, and Romas Aleliunas. Theorist: A logical reasoning system for defaults and diagnosis. *The Knowledge Frontier. Symbolic Computation (Artificial Intelligence)*, pages 331–352, 1987.



61. Harry Pople. On the mechanization of abductive logic. In *IJCAI'73: Proceedings of the 3rd International Joint Conference on Artificial Intelligence*, pages 147–152. Morgan Kaufmann Publishers, 1973.
62. Alun Preece. Asking “why” in ai: Explainability of intelligent systems—perspectives and challenges. *Intelligent Systems in Accounting, Finance and Management*, 25(2):63–72, 2018.
63. Lucas Rizzo and Luca Longo. Inferential models of mental workload with defeasible argumentation and non-monotonic fuzzy reasoning: a comparative study. In *Proceedings of the 2nd Workshop on Advances In Argumentation In Artificial Intelligence, co-located with XVII International Conference of the Italian Association for Artificial Intelligence, AI<sup>3</sup>@AI\*IA 2018, 20-23 November 2018, Trento, Italy*, pages 11–26, 2018.
64. Lucas Rizzo and Luca Longo. A qualitative investigation of the explainability of defeasible argumentation and non-monotonic fuzzy reasoning. In *Proceedings for the 26th AIAI Irish Conference on Artificial Intelligence and Cognitive Science Trinity College Dublin, Dublin, Ireland, December 6-7th, 2018.*, pages 138–149, 2018.
65. Lucas Rizzo and Luca Longo. An empirical evaluation of the inferential capacity of defeasible argumentation, non-monotonic fuzzy reasoning and expert systems. *Expert Systems with Applications*, 147:113220, 2020.
66. Lucas Rizzo, Ljiljana Majnarić, and Luca Longo. A comparative study of defeasible argumentation and non-monotonic fuzzy reasoning for elderly survival prediction using biomarkers. In *International Conference of the Italian Association for Artificial Intelligence*, pages 197–209. Springer, 2018.
67. Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference attacks against machine learning models. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 3–18. IEEE, 2017.
68. Nava Tintarev and Judith Masthoff. A survey of explanations in recommender systems. In *IEEE 23rd international conference on data engineering workshop*, pages 801–810, Istanbul, Turkey, 2007. IEEE.
69. Eduard Fosch Villaronga, Peter Kieseberg, and Tiffany Li. Humans forget, machines remember: Artificial intelligence and the right to be forgotten. *Computer Law & Security Review*, 34(2):304–313, 2018.
70. Giulia Vilone and Luca Longo. Explainable artificial intelligence: a systematic review. *CoRR*, abs/2006.00093, 2020.
71. Sandra Wachter, Brent Mittelstadt, and Luciano Floridi. Transparent, explainable, and accountable ai for robotics. *Science Robotics*, 2(6), 2017.
72. James Woodward. Scientific explanation. *Stanford Encyclopedia of Philosophy*, 2003. (<https://plato.stanford.edu>).
73. Adrienne Yapó and Joseph Weiss. Ethical implications of bias in machine learning. In *HICCS 2018, Proceedings of the 51st Hawaii International Conference on System Sciences*, 2018.
74. Quan-shi Zhang and Song-Chun Zhu. Visual interpretability for deep learning: a survey. *Frontiers of Information Technology & Electronic Engineering*, 19(1):27–39, 2018.