Artificial Intelligence and Social Computing, Vol. 122, 2024, 21–31
https://doi.org/10.54941/ahfe1004638

Human-Centered Explainable-Artificial Intelligence (XAI): An Empirical Study in Process Industry

Yanqing Zhang[1], Emmanuel Brorsson[1], Leila Methnani[2], Nilavra Bhattacharya[3], Elmira Zohrevandi[4], Andreas Darnell[5], and Rasmus Tammia[6]

[1]ABB AB Corporate Research Center, Västerås, 72226, Sweden
[2]Umeå University, Umeå, 90736, Sweden
[3]ABB AG Corporate Research Center, Ladenburg, 68526, Germany
[4]Linköping University, Norrköping, 60221, Sweden
[5]Södra Skogsägarna, Växjö, 35251, Sweden
[6]New Boliden, Boliden, 93632, Sweden

ABSTRACT

This paper presents an empirical study on the explainability of transformer models analyzing time series data, a largely unexplored area in the field of AI explainability. The study is part of an ongoing EU-funded project which applies a human-centered approach to developing explainable AI solutions for the process industry. Here, we investigate the choice of explainer mechanisms and human factor needs when developing eXplainable Artificial Intelligence (XAI) for operators of two industrial contexts: copper mining and paper manufacturing. On-site evaluations were conducted in these settings involving control room operators to test the prototype developed in the project. The results indicate that the method of feature importance alone was not sufficient to provide explanations that are tailored to individuals and situations, as required by users. Overall, our empirical data supports insights from previous research on human centered XAI and demonstrates the value of involving end users in the design process of effective XAI solutions. We also provide design implications which address human factor needs for such solutions in industrial settings.

Keywords: Explainable AI, Industrial application, Explainability, Empirical study, Human factors

INTRODUCTION

In recent years, there has been a growing interest in the use of Artificial Intelligence (AI) across various domains. This trend has also strongly influenced the technological advancements in the process industry. AI systems in this industry have increasingly improved production efficiency, reduced energy consumption, and ensured safer operations. Despite the high levels of automation, human involvement and decision making remain important for the industry's functioning and regulatory compliance. To mitigate the potential risks caused by black-box nature of many high-performing techniques used, e.g. distrust and misuse of the technology, eXplainable AI (XAI) techniques, i.e. making AI models interpretable and understandable to a wide range of stakeholders, has been researched and promoted in multiple domains, such as the medical field (Chakrobartty and El-Gayar, 2021) and the process industry (Kotriwala et al. 2021). The main body of work for XAI methods has been done for image classification models, however, the complex nature of time series data with an added temporal aspect that image data lacks (Siddiqui et al. 2019), calls for enhancing explainability techniques in domains such as industrial operations. Crucially, determining what constitutes a satisfactory explanation for practical deployment for XAI becomes paramount in this context. Most explainability methods produce technical explanations for data scientists and system developers, without consideration for other stakeholder needs (Miller et al. 2017).

In this study, we aim to fill in this gap and investigate what explanations are needed for operators in the process industry and what human factors considerations are needed when developing XAI for specific industrial use cases. The focused use cases are: UC1. Flotation: predicting copper concentration in the flotation process of copper mining; UC2. Digester: predicting the quality of the pulping process (quantified by a parameter called Kappa) in paper pulp manufacturing. On-site evaluations were conducted in these settings involving real control room operators to test the prototype developed in the project.

Overall, this research offers real-world empirical evidence in industrial settings and demonstrates the value of involving end users in the development process of effective XAI solutions. Additionally, we present design considerations that cater to the human factors essential for effective XAI solutions in industrial settings.

BACKGROUND

This section briefly summarizes the state-of-the-art research in the field of XAI and highlights the importance of a human centered XAI.

Explainable Artificial Intelligence (XAI)

Explainable AI is a key element of trustworthy AI which, in turn, plays a critical role in the industry's adoption of AI solutions (Arrieta et al. 2018). A system is considered *explainable* if it can provide *explanations*. An explanation offers an "interface between humans and a decision maker that is both an accurate proxy of the decision maker and comprehensible to humans" (Guidotti et al. 2018, p. 5).

State-of-the-art explainers are often categorized based on whether they offer *post-hoc* or *ante-hoc* explanations (Theissler et al., 2022). Simple models like decision trees or linear models produce ante-hoc explanations—they are intrinsically interpretable by design. On the other hand, complex models like deep neural networks are not inherently clear in their decision-making, and therefore require post-hoc methods for understanding. Many post-hoc methods are *agnostic*, meaning they can be applied to any prediction model (Molnar, 2020). The most popular post-hoc and model-agnostic techniques offer an understanding of opaque models by quantifying the extent to which specific feature inputs influence the model outputs (Munn and Pitman, 2022). Whether to select post-hoc or ante-hoc depends on many factors including *who* requires an explanation (Arrieta et al. 2018). While considerations for explainee requirements often come first, many stakeholders remain unaccounted for in this process. Most explainability methods produce technical explanations for data scientists and system developers, without consideration for other stakeholder needs (Miller et al. 2017). This realization has resulted in a push for more interdisciplinary Explainable AI research. By incorporating knowledge from non-technical areas such as sociology and human factors, we are seeing a shift towards Human-centered Explainable.

Human-Centered XAI

The importance of considering explanations as *social* has been increasingly recognized (Miller et al. 2017; Molnar, 2020). Insights from the social sciences have revealed several properties of explanations that make them effective in promoting understanding (Miller, 2019). For instance, Miller (2019) underscores four insights from existing literature: 1) explanations are *contrastive*, focusing not on why an event occurred but rather why it occurred instead of another possibility; 2) explanations are *selected*. Human explanations are biased selections, rarely aiming to provide exhaustive causes but rather picking one or two that are influenced by cognitive biases; 3) While probabilities hold importance, emphasizing statistical relationships in explanations often falls short; *causal* explanations hold greater weight for understanding events; 4) explanations are inherently *social*, involving a transfer of knowledge within conversations or interactions, shaped by the explainer's understanding of the explainee's beliefs. Overall, these stress that explanations in AI transcend mere causal attribution; they're contextually nuanced, subject to selection and allow for interaction between explainer and explainee.

Moreover, providing clear explanations for AI-driven decisions is crucial for enhancing experts' trust and ensuring successful adoption of these systems (Kotriwala et al. 2021). To generate clear explanations, designers should understand the experts' task flows and include them in validation of XAI methods, which also need to be tailored for industrial data (Kotriwala et al. 2021).

METHODS

In this project, we applied a user-centered design approach to develop XAI solutions. Using the framework of contextual inquiry to question and observe users in their natural environment (Duda et al. 2020), field visits were carried out to gather an in-depth understanding of the processes, the needs, and the pain points of process operators. The insights gained were synthesized into initial user requirements. Subsequently, a multidisciplinary team collaborated on functional prototypes that comprised of an explainable time series forecasting model, and a dashboard connected to the model output, which could be applied to both application domains. Two Temporal Fusion Transformer (TFT; Lim et al. 2021) architectures were used for forecasting two different variables in the two use cases, trained on historical data from each plant. As TFTs inherently provide an overview of variable weights for local explainability, the dashboard, which was deployed using Grafana, was able to display feature importance and attention scores associated with each variable (see Figure 1). Finally, the dashboard and its explainability components were evaluated on-site with eight control room operators for UC1 (seven male, one female), and six (all male) for UC2.

In our evaluation, three user-centric dependent measures were adopted; 1) *forward simulatability,* 2) *subjective comparison*, and, 3) *subjective satisfaction*. These were combined with qualitative interviews. Such a combination was proved to be effective in evaluating the utility and satisfaction of explanation methods. For a more detailed overview of the methodology, see Brorsson et al. (2024). In this study, we focus on the qualitative data gathered on site. All interviews were captured by video and audio. Notes taken during the sessions were transferred into a shared digital workspace, analyzed and categorized into various themes using thematic analysis (Braun and Clarke, 2006). During the analysis phase, recordings were also revisited by another coder to further detail the identified themes.

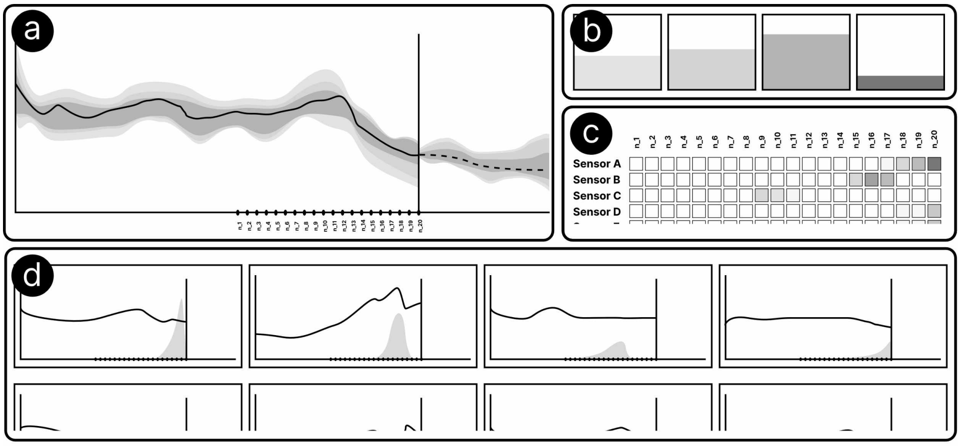

Figure 1: A schematic representation of the designed dashboard that was used in the study. The dashboard contained four views. An overview window (a) which showed the predicted time series graph considering the sensor data and the recovery profile for the key process parameter. A precision-intervals view (b) which showed how often the key performance parameter would lie within each interval. A feature-importance view (c) which represented the sensor weight values and an on-demand representation view (d) which visualized the time series data of all sensors for the time-period of interest selected by the operator.

RESULTS AND DISCUSSION

This section presents the extent to which the designed explanations were suited to the domain experts' needs with respect to *expectation alignment, decision-making alignment, trust,* and *prediction accuracy.* We have chosen to focus on the qualitative data gathered in the interviews.

Expectation Alignment

For both contexts, most participants misunderstood that the model adjusts important features to improve the results on its own. A few participants mistakenly believed that the model could impact the process directly by adjusting reagents to improve capabilities (UC1) or adjusting the H-factor (UC2).

Misunderstandings may arise when explanations insufficiently bridge this gap between what the model's capabilities are believed to be and what the model is actually doing. Given the participants' limited experience with AI, it is likely that their expectations exceeded the capabilities of the models that were tested. Various automated control functions are prevalent in these process industries both in simpler forms such as automated valves and more complex forms such as functions calculating and implementing optimal dosages of a liquid. Participants were found to draw parallels between existing control functions and the evaluated ML model, indicating that expectations might be transferred from the former to the latter.

In existing literature, researchers have found that users with high expectations exhibit over-reliance behavior and users in the neutral condition tend to appropriately adjust their behavior (Mayer et al. 2006). Mismatches between expectations and reality can result in user frustration, distrust, and disengagement with the system (Meurisch et al. 2020). In addition, expectations can influence how explanations are *selected* due to cognitive bias, as Miller also highlights (2019). Therefore, it is needed to consider and manage these expectations to create successful interactions with AI systems. In our case, it is highly important to clarify that the model developed for testing only has the capability to analyze and predict the process without acting upon it. Otherwise, we may encounter passiveness caused by false over-reliance during the interaction. We may also see induced fear of giving feedback based on false assumptions that it will have a direct negative impact on the process.

Decision-Making Strategy Alignment

In both use cases, participants commonly spent a fair amount of time finding and comparing different sensors to understand the prediction. Since not all sensors could be displayed simultaneously on the screen, they had to locate the relevant sensors first, check the values for a specific time step, remember those values, and then scroll to compare them with other sensors. This process is quite cumbersome for making comparisons within the current navigation of the prototype. The interview results indicate that operators are strongly influenced by their existing work practices. In both use cases, control room operators rely on their control systems to monitor, analyze and act upon the process. This has caused them to develop certain strategies which are influenced by the systems they use. For example, operators commonly plot several variables in one trend package, which allows them to compare various types of sensor data and infer relationships between them in their current way of working. Overall, these strategies involve monitoring a graphical representation of an overview of the whole process step that they are responsible for, and when specific information is needed, they interact with the system to produce additional details regarding the sensors or components of interest. Notably, while troubleshooting why the process key performance parameter was progressing in a certain direction, operators rely on *causal inference* (Miller, 2019), i.e. patterns in sensor values which, according to their expertise of the process, are adopted as explanations for the deviations they are observing.

Thus, to explain and understand the model forecast, the operators rely on similar strategies of exploring sensor relationships and browsing sensor data gained from their previous experience. This indicates that if we design a system that can align with their decision-making models (Zohrevandi et al. 2023), it will bring great value to the users. However, operators' individual experience of the process and their unique toolkit of strategies might have a great impact on what knowledge gaps they expect to be filled. This can make it challenging to develop a generic interactive interface that could comply with the wide range of strategies and experience levels that might exist within a single working crew.

Trust Calibration

Trust is considered as fundamental element in the interactions between human and automated systems. Without trust, users will not rely in automated systems, especially to conduct critical tasks (Rojat et al. 2021). Our user tests have illuminated three key facets of fostering trust in AI models.

First, the accuracy of the model forms the foundation for users to build trust with the models (Yin et al. 2019). A few users mentioned that their trust in the model would be reduced if the prediction was not in line with what happened after the fact. This is in line with what Grice's maxim of *Quality* stating that information displayed on the explainer should be of high quality, accurate and well-grounded (Grice, 1975). Model performance provides a foundation to assess the quality of the output of the model. This performance must attain a level of proficiency that instils confidence in its ability to deliver reliable predictions. Users' willingness to entrust critical tasks to automated systems hinges on the assurance that these systems can consistently and precisely fulfil their intended functions.

Second, aligning the model with the operators' mental model can augment trust. For example, one said if the dashboard enables comparison among different variables, learning and teaching, like what they are currently doing, trust to the system would increase. When the model's explanations resonate with the variables that operators deem essential for predictions—even if these variables do not align with the model's internal considerations—it establishes a foundation for operators to assess the quality of predictions. Such mental model alignment cultivates a sense of familiarity and comprehension, contributing significantly to the augmentation of trust in the automated system.

Finally, meaningful feedback is crucial for trust. Participants mentioned that it is important to be easy to give feedback and also be acknowledged about the effects of their feedback on the model. This indicates that interactive features should not only enable users to tailor information to their needs but also allow them to influence the model. This fosters mutual understanding between humans and AI (Chander, 2018), empowering operators and enhancing transparency. This supports Miller's argument of explanations as inherently being *social*, involving a transfer of knowledge within conversations (2019).

Selection of Explanation Mechanism

Understanding what explanations are needed, which types of explanations users can comprehend under time-pressure and the appropriate modalities for the explanations are important aspects when designing XAI systems for industrial uses. This section will present the findings related to these questions.

Feature Importance Matrix & Sensor Attention (c & d in Figure 1)

The primary explanation mechanism used was a matrix-based visualization of feature importance values (see c in Figure 1), which is a popular post-hoc and model-agnostic technique that quantifies the extent to which feature inputs influence the model outputs (Munn and Pitman, 2022). The feature importance was supported by detailed model attention for specific time steps of each sensor (see d in Figure 1). The results show that for both use cases, most users did not understand the sensor importance matrix at first glance. But the feature importance matrix was found to guide users to explore values for high-importance sensors first, similar to Shneiderman's information seeking mantra "Overview first, zoom and filter, then details-on-demand" (1996, p. 2). However, a few challenges became apparent. Firstly, participants with less knowledge about AI commonly misinterpreted important features as factors that were predicted to have a causal effect on the forecasted variable. This aligns with how operators normally think of the process, as a set of causal relationships between various variables. This has caused confusion as to why Kappa was an important feature for forecasting Kappa (UC2). Some participants mentioned that they understand importance matrix but do not agree with what sensors the model prioritizes.

Secondly, it was challenging for participants to understand why a certain feature importance score was calculated for a feature, even after browsing individual sensor readings and the attention scores associated with important time steps in the retrospective period which was marked as a grey zone in the trend graph. A possible reason for this might be that people use combinations of features and their underlying causal relationships when assessing categories of objects (Rehder, 2006). Rehder (2006) demonstrates that people use feature combinations to categorize objects, particularly when features seem incompatible due to underlying causal mechanisms. Seen through the lens of explainability, this implies that if the feature importance matrix would mainly refer to sensors that align with what features are important for participants when assessing the process, they would likely have a better foundation for assessing the model output.

Third, for UC1, meta-variables (i.e. exclusive to the model such as encoded day of month) were included in the variables that the model accounted for when making its forecast. These variables are not included in the process operators normally work with. This led to the participants not understanding why they were included in the model in the first place. Explanations need to be tailored to the mental models of its viewers (Kulesza et al. 2013), and meta-variables such as these can thus be inappropriate for users with less ML experience (Liao et al. 2022).

Precision of Historical Predictions (b in Figure 1)

In both use cases, the quartiles used for representing forecast accuracy were sometimes misinterpreted as acceptable thresholds of where the key process variable (copper recovery or Kappa) should optimally stay in. Even after facilitators explained that this was not the case and the historical precision should provide an indication of how probable the current forecast might be, participants, in general, did not use the precision indications for assessing the quality of the explanation. This is in line with the findings reported by Miller that people do not tend to assess explanation quality based on probability but rather on usefulness, relevance, and causal behavior (2019, p. 45). However, as noted by Dodd and Bradshaw (1980), explanations are interpreted according to the intent of the systems they are associated with, indicating that if a model output is unreliable or inaccurate, the user risks having reduced trust in the explainer as well. This connects to the previous subsection on trust where the accuracy of the model is stated as the foundation for building trust with the system. Therefore, having some indication of precision or accuracy might provide a foundation to calibrate trust not only for the model but also for associated explainers.

Design Implications

Previous sections have discussed different human factors that influence the selection of explainer mechanisms and interaction design, based on our user study results. The following table summarizes key design implications.

Many of these implications concern operators' mental models and individual differences. Here, we also want to highlight that the contexts play a vital role in designing a suitable explainer solution for process operators. The nature of the industrial process is highly dynamic; the time and effort required from operators to tend to it varies greatly in turn. This, together with individual differences in mental models, calls for dynamic forms of explainability that can fill in knowledge gaps in a wide range of situations for individual differences. To provide such forms of explainability, the interfaces should follow Shneiderman's information seeking mantra (1996) by providing an overview of general explainability tailored to most individuals and situations and supporting *fluid interactions* for operators to fill in their individual knowledge gaps for the situation they are in (Chander et al. 2018).

Table 1. A summary of design implications and connected human factor requirements.

Design implications	Human factor requirements
Clear indication of the model capabilities	Expectation alignment
Display model confidence or accuracy	Expectation alignment
Consider different explainer mechanisms (e.g. causal effects) to fit different purposes	Decision making strategies (mental models) alignment
Provide an overview of crucial information, then allow digging deeper into details	Individual differences (AI competence, experience etc.) Decision making strategies (mental models) alignment
Design interface that can allow users to customize and navigate process data to match the current mental models of users	Decision making strategies (mental models) alignment Individual differences (AI competence, experience) Trust calibration
Provide meaningful feedback mechanism that can take in feedback strategically and improve the model	Trust calibration

Limitations

The evaluations of this paper were carried out with limited and intended end users, and not in their natural environment because of safety risks of interfering with the industrial process. Instead, a separate system in a separate room was used. In addition, participants were presented with all explainers simultaneously during subjective comparison. Although this approach enables insights into holistic usage, it may compromise fidelity in determining which individual explainer is most beneficial for participant.

CONCLUSION

In this empirical study, we examined what explanations are needed for end users in the process industry and what human factors should be considered when developing XAI. Our key findings indicate that relying solely on the method of feature importance will not suffice to deliver personalized and situation-specific explanations as demanded by users. Designing XAI for industries should involve studying potential users' expectations of their interactions with AI systems, considering their existing work behaviors and mental models, prioritizing the development of highly accurate models, and providing avenues for feedback to enhance trust. Overall, this research offers empirical evidence from real-world industrial settings, highlighting the significance of involving end users in the design and development of effective XAI solutions.

Future work could investigate other types of explainer mechanisms and how to design them for process industries. For instance, it could explore whether one or several mechanisms should be employed to offer explainability to users in industrial settings and address the human factor needs identified in this study.

ACKNOWLEDGMENT

The present study is funded by VINNOVA Sweden (2021-04336), Bundesministerium für Bildung und Forschung (BMBF; 01IS22030), and Rijksdienst voor Ondernemend Nederland (AI2212001) under the project *Explanatory Artificial Interactive Intelligence for Industry*[1] (EXPLAIN).

REFERENCES

Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A.,... & Herrera, F. (2020). Explainable Artificial Intelligence (XAI): Concepts, Taxonomies, Opportunities and Challenges toward Responsible AI. Information Fusion, Volume 58, 82–115.

Braun, Virginia and Clarke, Victoria (2006). Using Thematic Analysis in Psychology. Qualitative Research in Psychology, Volume 3 No. 2.

Brorsson, E., Zhang, Y., Zohrevandi, E., Bhattachaya, N., Theodorou, A., Darnell, A., Tammia, R., van Driel, W. (2024). "Evaluating Explainability in Time Series Models: A User-Centered Approach in Industrial Settings". In Proceedings of Applied Human Factors and Ergonomics, Nice, France, Upcoming.

Chander, A., Srinivasan, R., Chelian, S., Wang, J. and Uchino, K., (2018). "Working with Beliefs: AI Transparency in the Enterprise". In IUI Workshops, Volume 1.

Chakrobartty, S., & El-Gayar, O. (2021). Explainable artificial intelligence in the medical domain: a systematic review. In proceedings of Americas Conference on Information Systems. 1.

Dodd, D. H., & Bradshaw, J. M. (1980). Leading Questions and Memory: Pragmatic Constraints. Journal of Verbal Learning and Verbal Behavior, Volume 19 Issue 6.

Duda, S., Warburton, C. and Black, N. (2020). "Contextual Research: Why We Need to Research in Context to Deliver Great Products". In Proceedings of Human-Computer Interaction. Design and User Experience: Thematic Area, Copenhagen, Denmark.

Grice, Paul H. (1975). Logic and Conversation in Syntax and Semantics 3: Speech Acts, ed. by Peter Cole and Jerry L. Morgan. New York: Academic Press 1975, 41–58.

Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., & Pedreschi, D. (2018). A Survey of Methods for Explaining Black Box Models. ACM Computing surveys (CSUR), Volume 51, Issue 5, 1–42.

Kotriwala, A. M., Klöpper, B., Dix, M., Gopalakrishnan, G., Ziobro, D., & Potschka, A. (2021). "XAI for Operations in the Process Industry - Applications, Theses, and Research Directions". Proceedings of the AAAI 2021 Spring Symposium on Combining Machine Learning and Knowledge Engineering (AAAI-MAKE 2021) – Stanford University, Palo Alto, California, USA.

Kulesza, T., Stumpf, S., Burnett, M., Yang, S., Kwan, I. and Wong, W. K., (2013) "Too Much, Too Little, or Just Right? Ways Explanations Impact End Users' Mental Models," IEEE Symposium on Visual Languages and Human Centric Computing, San Jose, CA, USA.

[1] https://explain-project.eu/

Liao, Q. Vera, and Varshney, Kush R. (2022) Human-Centered Explainable AI (XAI): From Algorithms to User Experiences, arXiv:2110.10790v5.

Lim, B., Arık, S. Ö., Loeff, N., and Pfister, T. (2021). Temporal Fusion Transformers for Interpretable Multi-horizon Time Series Forecasting. International Journal of Forecasting, Volume 37 No. 4, 1748–1764.

Mayer, Andrew & Sanchez, Julian & Fisk, Arthur & Rogers, Wendy. (2006). "Don't Let Me Down: The Role of Operator Expectations in Human-Automation Interaction". Proceedings of the Human Factors and Ergonomics Society 50th Annual Meeting. San Francisco, CA, United States.

Meurisch, C., Mihale-Wilson, C. A., Hawlitschek, A., Giger, F., Müller, F., Hinz, O. and Mühlhäuser. M. (2020). Exploring User Expectations of Proactive AI Systems. Proceeding of ACM Interact. Mob. Wearable Ubiquitous Technol. Volume 4, No. 4

Miller, T., Howe, P., & Sonenberg, L. (2017). Explainable AI: Beware of inmates running the asylum or: How I learnt to stop worrying and love the social and behavioural sciences. arXiv preprint arXiv:1712.00547.

Miller, Tim (2019). Explanation in Artificial Intelligence: Insights from the Social Sciences. Artificial Intelligence, Volume 267, 1–38.

Molnar, Christoph. (2020). Interpretable Machine Learning: A Guide for Making Black Box Models Explainable. Lulu. Com.

Munn, M., & Pitman, D. (2022). Explainable AI for Practitioners. O'Reilly Media, Inc.

Rehder, Bob, (2006). When Similarity and Causality Compete in Category-based Property Generalization. Memory & Cognition, Volume 34.

Rojat, T., Puget, R., Filliat, D., Ser, J. D., Gelin, R., & Díaz-Rodríguez, N. (2021). Explainable Artificial Intelligence (XAI) on Time Series Data: A Survey. ArXiv, abs/2104.00950.

Siddiqui, S. A., Mercier, D., Munir, M., Dengel, A. and Ahmed, S., (2019). Tsviz: Demystification of Deep Learning Models for Time-Series Analysis. IEEE Access, Volume 7, pp. 67027–67040.

Shneiderman, Ben (1996). "The Eyes Have it: A Task by Data Type Taxonomy for Information Visualizations". In Proceedings 1996 IEEE Symposium on Visual Languages (pp. 336–343). IEEE. Boulder, CO, USA.

Theissler, A., Spinnato, F., Schlegel, U., & Guidotti, R. (2022). Explainable AI for Time Series Classification: A Review, Taxonomy and Research Directions. IEEE Access.

Yin, M., Vaughan, J., and Wallach, H. (2019). "Understanding the Effect of Accuracy on Trust in Machine Learning Models". CHI 19: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems. 1–12. Glasgow, Scotland UK

Zohrevandi, E., Brorsson, E., Darnell, A., Bång, M., Lundberg, J. and Ynnerman, A. (2023). "Design of an Ecological Visual Analytics Interface for Operators of Time-Constant Processes". 2023 IEEE Visualization and Visual Analytics (VIS). Melbourne. Australia. 131-135, doi: 10.1109/VIS54172.2023.00035.