Human-Feedback for AI in Industry

Izaskun Fernandez
Intelligent Information Systems
TEKNIKER, member of BRTA
Eibar,Spain
email:izaskun.fernandez@tekniker.es

Kerman Lopez De Calle
Intelligent Information Systems
TEKNIKER, member of BRTA
Eibar,Spain
email:kerman.lopezdecalle@tekniker.es

Eider Garate
Intelligent Information Systems
TEKNIKER, member of BRTA
Eibar,Spain
email:eider.garate@tekniker.es

Regis Benzmuller
IT MES
CONTINENTAL France, Sarreguimes
Sarreguemines, France
email:regis.benzmuller@conti.de

Melodie Kessler
R&D vision industrielle et Deep Learning
CONTINENTAL France, Sarreguimes
Sarreguemines, France
email:melodie.kessler@conti.de

Marc Anderson
LORIA
Université de Lorraine, Inria and CNRS
Vandœuvre-Lès-Nancy, France
email:marc.anderson@inria.fr

***Abstract*—Artificial Intelligence (AI) offers a wide variety of opportunities to the manufacturing industry. However, there are still gaps and challenges to be solved before it can be successfully applied, with data availability and quality being one of the critical factors. The latter highlights the necessity of developing AI systems that can continually learn (from one or more domains) over a lifetime, starting from limited sets of data. This work presents research done on human reinforced learning approaches on small training data sets of open dynamic environments. Beginning this way allows the development of AI models able to learn over time, while taking advantage of a data driven approach along with a knowledge-based approach considering human-feedback as a key enabler.**

***Index Terms*—Human-feedback, Artificial Intelligence, Data Quality, Data Annotation, User-Centric**

I. INTRODUCTION

One of the most significant and challenging open problems in Artificial Intelligence (AI) is that of developing systems that can continually learn (from one or more domains) over a lifetime. Although new approaches are appearing to bridge the gap of continuous learning [1], for many years the dominant Machine Learning (ML) paradigms have adopted isolated learning. The latter runs a ML algorithm on a given dataset to produce a model, without any attempt to retain the learned knowledge and use it in the future. The isolated ML approach, has been very successful, but it requires many training examples, and is only suitable for well-defined and narrow tasks in closed environments. This ideal situation is not common in real industrial environments of the small data regime type, i.e., where the amount of available data is scarce. Data may be scarce not only in terms of a limited volume of data, but also due to some environments having highly unbalanced data or slow dynamics which prevent extracting the underlying pattern from the training data. Moreover, to solve many real industrial problems through ML, temporal sequence and/or sensor data should be dealt with, which implies an additional challenge. Accordingly, there is a clear consensus regarding the importance of the quality of the data for the development and deployment of accurate AI based models [2].

Human-feedback could be a valuable source of information for improving ML models and systems by improving data quality and annotations. Collecting and using human-feedback in ML is not a trivial task however. It requires careful design, implementation, and evaluation of different methods and techniques. Indeed, using human-feedback in ML can be expensive, time-consuming, or impractical, especially for large-scale or complex problems. It can also be noisy, inconsistent, or biased due to human errors, preferences, or motivations. It is crucial to take into consideration these challenges and limitations when using human-feedback in ML.

This work presents an overview of different human-feedback mechanisms for human-oriented AI models reinforcement, explored in various real industrial scenarios within the AI-PROFICIENT project. These mechanisms are based on different types and levels of technologies, but all can be classified according to two main groups: implicit and explicit feedback. Implicit feedback is tied to a user action that they would perform (or not perform) regardless of their desire to influence the results given by the AI. Explicit feedback is when the user performs an action specifically designed to enable them to give feedback to the system. Regardless of the feedback type, all the strategies have been designed and implemented following an ethics by design approach, which has contributed to an efficient and good quality data collection, while also promoting user engagement. The most relevant ethical aspects are also presented in this work as good practises/guidelines to be considered when dealing with human-feedback approaches.

The rest of the article is structured as follows. Section II presents the related work. The different human-feedback approaches for AI in industry are detailed in Section III and the main ethical aspects to be considered are detailed in Section IV. Section V describes a successful implementation and deployment of human-feedback for AI in a real use case scenario. And finally, conclusions and lessons learned are summarized in Section VI.

II. Related Work

Data for AI is recognized as an innovation ecosystem in the European AI, data, and robotics framework [4], and data sharing is a critical enabler for competitive AI solutions. Data spaces are a key element of the European Data Strategy, fostering Secure and controlled environments, necessary for eliminating distrust of companies and people when sharing their data. They ensure that data exchanges take place in a safe and secure manner, and in an interoperable manner. The availability of these large interoperable datasets will help to develope more robust and reliable AI based systems, however, incorporating user knowledge into the system is considered as a complementary and indispensable path to improve the AI systems still further, and set the path to cognitive AI systems.

Researchers are defining new types of interactions between humans and AI generically called Human-in-the-loop [5], although definitions of the term vary quite widely as [6] have shown. According to one view, Human-in-the-loop aims to train an accurate prediction model with minimum cost by integrating human knowledge and experience. This enables a significant data requirement reduction, increases reliability and robustness of the AI, and creates explainable AI systems [7], by making humans more effective and more efficient.

There are different approaches implementing Human-in-the loop during different phases of AI system lifecycles: development involving humans in data preparation (including data processing and labelling) [8]; training through interactive ML approaches [9]; data labelling to get explainable AI systems [10]; and reinforcement-oriented approaches [11] using end-users feedback to adjust the AI system to the target user preferences while keeping the model objective optimum.

But, to create an effective Human-in-the-loop system, it is important: to understand how humans interact with machines and to focus on creating natural and easy to use mechanisms that can be wielded through human interaction. It is also important to avoid high cost strategies for human-feedback collection and to research strategies to translate such feedback into exploitable information for AI systems. Such feedback could be provided by humans consciously and explicitly, or inferred from other actions not necessarily linked to feedback.

III. Human-feedback AI

When implementing AI in the industrial context, there are certain challenges that industrial plants face. Two of the most common challenges of AI adoption in industry are data scarcity and the human reluctance to accept AI systems. Perhaps there is no information management system in place to gather vast quantities of data through a multitude of sensors. Or, perhaps, there are enough sensors in place, but the data has never been recorded, labelled and collected. Whatever the reason, beginning to design and develope AI based models with limited annotated data, even partial data, in the sense that whole cases are not represented, or are represented in an unbalanced way, is a real challenge that should be engaged in order to successfully adopt AI models in industry.

In this work, we focus on the use of human-feedback to overcome this data scarcity and data quality challenge - a challenge which usually impacts negatively in the AI model's accuracy - by generating quality data automatically. We present different paths to gather this human-feedback, making a special effort toward reducing as much as possible the human intervention with feedback intention, and putting in place friendly and natural interfaces to facilitate as much as possible the intervention whenever it is necessary. This minimizes the increase of human workload in providing feedback, and fosters the acceptance of the reinforcement AI approaches.

A. *Implicit feedback*

Implicit feedback is considered the best feedback gathering approach in terms of impact upon human workload, since it is collected from an action, optionally done by a human, that has nothing to do with feedback. For example, let's assume there is an AI model running which suggests an optimal value for a certain process parameter in a given moment. Implementing a workflow that registers if a suggestion given by a model has been adopted or not, matching automatically the suggested and the applied values, can strengthen the initial dataset, and thus the AI-model, when retraining by, for example, a reward-punish strategy, although no explicit feedback has been given for that action.

This automatic data strengthening, however, could be negatively affected by erroneous information if the real intention behind the action triggering the feedback is not considered by measuring the degree of compliance of the human [12] with respect to the action. This also happens if the contextual information is not properly managed.

Thus, it is crucial to include in the implicit feedback managing strategy both intention and context management. There are several works in the literature that have proposed methods to differentiate between genuine and unintentional disagreements in implicit feedback [13]. These methods render the information sufficiently accurate, prevent noise creation in the AI model, and successfully use implicit feedback, while minimizing impact on human workload.

In the context of AI-PROFICIENT, the approach is based on a monitoring system to collect the real value and the context information for ensuring the data quality. This feedback management has contributed especially in increasing the data quality, by correcting potential biased or incorrectly annotated data. According to the different typologies of the AI-based models most associated with parametrization and the status of different agents in production plants, in AI-PROFICIENT two different implicit-feedback strategies have been distinguished:

- Predictive AI-based models: the predicted values are compared automatically with the real value registered by the automatism/agent that affects and creates the reinforcement information, indicating if the prediction was correct or not, and if not including the correct value.
- Recommendation (including optimization) AI-based models: the recommended value will be compared

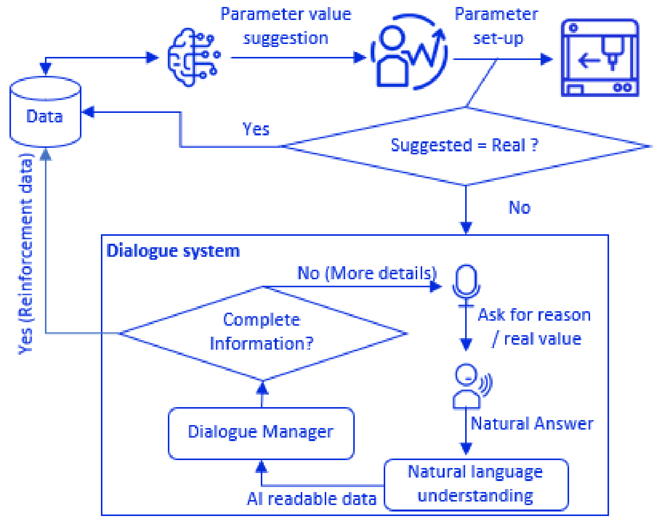

Fig. 1. Natural voice dialogued human-feedback approach.

with the adopted real value introduced in the automatism/agent. When equal it will produce reinforcement positive information and when different, a negative reinforcement information including the recommended and adopted real value.

But in both cases, during the information registration process, contextual information such as product type, timing, and so on, coming from the monitoring system will be registered to ensure data quality assessment.

B. Explicit human-feedback

Reinforcement learning without human intervention is not always feasible. For instance, in a vision recognition scenario, when the AI-based automatic recognition is not the correct one, specific intervention is required to get the right information.

Moreover, explicit human-feedback can overcome implicit feedback drawbacks. In that way, recording the reasons why a decision is taken can make it possible to record contextual information to enrich the model and measure the compliance degree between the human and the AI model.

Since industrial environments are becoming more automated over time, Human-Machine Interfaces (HMI) have increasingly evolved in the last years with the development of new mobile techniques and new gadgets such as smartphones, tablets, or Augmented Reality (AR) glasses. Many solutions have been developed, especially in collaborative robotics, with human-machine interaction capabilities in different degrees, which allow a more intuitive communication with industrial systems: interaction through gestures; programming by demonstration; even dialogue systems, which allow workers to interact with industrial systems in a similar way as they would do with their fellows.

In the following section, different strategies to easily collect explicit human-feedback through advanced Human-Machine Interfaces are presented, taking into consideration the current and future trends in human-machine interaction. This type of feedback could contribute particularly to the improvement of the model by adding to the data collection new previously unobserved types of data which are rarely represented or even unrepresented in the collection, so as to obtain a more balanced data set.

1) Natural voice interaction: The possibility of communicating with industrial systems through natural language is highly encouraged since it triggers acceptance from humans, and dialogue systems are a powerful technological solution to deal with this necessity. Among dialogue systems, task-oriented dialogue systems –as opposed to conversational dialogue systems, which try to emulate regular conversations– are designed to perform specific actions upon a user request and, for this, are especially relevant in industrial contexts. In this sense, task-oriented dialogue systems are powerful technologies that allow workers to work on multiple tasks at once by delegating secondary assignments through communicating with the target system, usually with voice commands. The use of voice instructions to interact with these systems allows workers to use them from a safe distance if necessary, and in a way that they do not need to interrupt their current tasks, leaving the quality of their work unaffected. Furthermore, enhancing these systems with the capability of interacting with users in natural language releases workers from having to learn specific commands to use them.

For feedback gathering purposes, a task-oriented dialogue system should focus on generating the proper dialogues when the AI system fails, asking the human (operator) the necessary questions to obtain the right information and translate it in order to create new insights to reinforce the current model. Fig. 1. shows a generic workflow of this approach.

2) Augmented Reality based interfaces: Augmented Reality (AR) is a technology field that involves the seamless overlay

 ISBN: 978-1-68558-100-8

of computer-generated virtual images on the real world, in such a way that the virtual content is aligned with real world objects and can be viewed and interacted with in real time. AR research and development has made rapid progress in the last few decades, moving from research laboratories to widespread availability on consumer devices. Augmented Reality through wearable devices such as Google's Augmented Reality glasses can bring numerous advantages for information visualization, such as displaying relevant information to the driver of a fork-lift through the glasses. Such devices can also be a very useful tool for machinery repair and maintenance, mitigating human errors and possible accidents [14]. Common approaches for AR interaction include tangible User Interfaces (UIs) and free-hand gesture-based interaction. These could be used, not only for advanced information visualization in industrial scenarios, but, in addition, to facilitate the user in providing corrections or suggestions, thus enabling a new channel for the collection of human-feedback for reinforcement purposes.

3) Shop-floor interfaces: New HMIs need to be more sophisticated for enhanced efficiency and remote service operations, especially when workers are interacting with intelligent agents in dusty, humid, or dark, industrial environments. Since operators become involved in the manufacturing process for critical decision-making, the HMI system should allow commands to be easily and rapidly entered in order to increase the accuracy, safety and speed of problem-solving. But not all the industrial scenarios are ready to go for a voice or augmented reality driven HMI.

So, not only the most advanced mechanisms should be considered for feedback management, but extending currently available or newly developed dashboards and interfaces must be also considered to ensure human-feedback gathering in scenarios with different levels of digitalization in terms of interfaces: excel sheets, web-based dashboards, etc. The feedback methods can also benefit from the State-of-the-art techniques applied for uses in the industrial environments, such as active noise suppression, AI-supported and trained Optical Character Recognition, etc.

IV. Lessons learned from human-feedback mechanisms in use

In the context of AI-PROFICIENT project, a total of 6 industrial use cases (from Continental and INEOS industrial partners) involving AI models and human-feedback oriented reinforcement strategies have been developed to face different industrial problems. Table I shows a summary of the different human-feedback approaches, according to the classification presented above.

The main common challenge set by the industrial partners in all the scenarios has been not to change the working procedures and to minimize new interfaces and devices in the workplaces for AI models and feedback mechanisms. Taking into account this restriction, the 6 use cases have set up at least one AI model supporting the target problem, and all except one, which has only implemented an explicit human-feedback approach, have combined both, implicit and explicit feedback. Two of them have even combined two different explicit mechanisms.

For implicit feedback, monitoring systems have been set up in the scenarios without affecting existing procedures. While for the explicit feedback, the most common approach has been to extend the interfaces developed for interacting with AI models. For instance, an AR application in INEOS plant that uses a computer vision model to recognize labels has been extended to gather feedback. More specifically, new AR screens have been added to the application enabling users to introduce the correct information by augmented keyboard, and even voice, when the AI model fails. But not only industrial restrictions have been considered. During the design and implementation of the human-feedback mechanisms, an expert team has assessed ethical aspects relative to human acceptance. Following, we list the most remarkable and most often met with aspects that should be considered in order to develope a successful human-feedback AI approach ethically.

TABLE I

Overview of number of use cases implementing feedback mechanisms, per type

Feedback type	# use cases
Implicit Human-feedback	5
Explicit Human-feedback - Voice	4
Explicit Human-feedback - AR	1
Explicit Human-feedback - Shopfloor HMI	3

- When generating human-feedback based data, regardless of whether it belongs to explicit or implicit interaction, indicate that the data is generated by human-feedback
- Given that voice data, because it is biometric, is inherently sensitive data, when feedback mechanisms make use of it, it is recommended that ethical best practices similar to the measures of Article 5(b) GDPR be implemented, regardless of legal compliance requirements, and that demonstrable consent be obtained.
- In those cases where natural voice and language interaction is part of the human-feedback, take into consideration the operator's mother tongue and if the former is not considered when implementing the mechanism, then, develop a clear, detailed, and practical plan for how the language gap and difficulties related to operator language and HMI use will be bridged.
- The feedback mechanisms should not increase the user workload and if it does, strategies should be implemented in order to try to minimize it for ensuring adoption and acceptance.

V. Human-feedback AI: In Continental Use Case

Following the ethical recommendations and taking into consideration Continental's need to improve a trade blade change process, an AI-based reinforcement approach has been implemented in the Continental plant.

In the beginning of the research, there was no AI-based model deployed in the Continental plant, but at the end of the

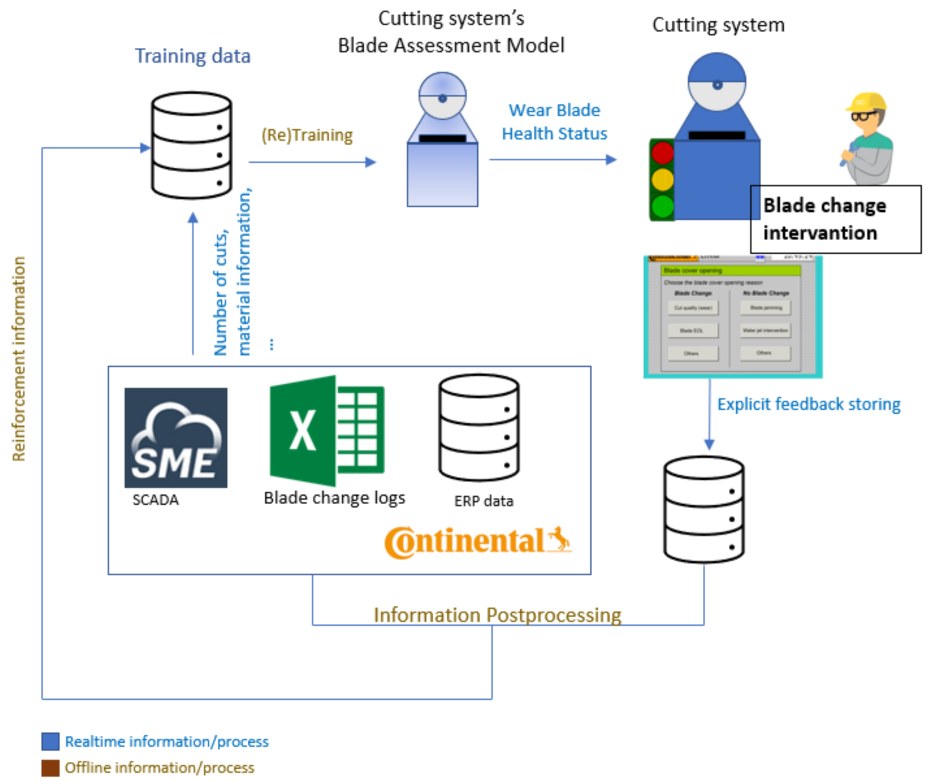

Fig. 2. Continental trade-blade change human-feedback management flow.

workday, the craftsmen and operators used to fill an excel-sheet indicating, among other activities, the blade changes done during the workday, and their approximate time. This information was stored together with the other information associated to the machine with the trade blade, such as material types and compositions, performed cuts, or monitored signals.

Taking advantage of this historical data (from the three last years), necessary actions have been undertaken to develop and deploy an initial AI-based model, aimed at predicting the optimal moment to change the blade, considering all the above mentioned information and the appearance of the blade.

However, due to the imperfection of the AI-model, mainly caused by the scarce and not exact training data, the prediction does not always provide the best moment for the blade change, and so the operator does not always follow the suggestion provided by the model. In order to correct these deviations and train a more robust model as it is used, a human-feedback management approach, based on both implicit and explicit feedback has been implemented and deployed in real plant as summarized in Fig. 2.

A. Implicit feedback in use

AI-based model estimation is shown to the operator in terms of a traffic-light visualization, with green indicating a healthy blade status and red indicating the end of life of the blade, and thus a trade-blade change action needed. Currently, adopting the AI-based estimation or disregarding it, when necessary the operator directly performs the blade change or asks the craftsmen to do it.

The current solution in the Continental plant includes in the trade-blade safety-cap, a sensor which records the exact time in the database whenever the cap is opened. The cap can be opened due to different reasons:

- Blade change
- Adjusting of the blade without change
- ...

This in-situ time registry, facilitates an exact time recording but since the cap opening is not necessarily related to a blade change, an explicit feedback mechanism is needed to complement the information and be able to select and use only the correct data for model reinforcement.

B. Explicit human-feedback in use

In comparison to the initial situation, the current solution in place in the Continental plant eliminates the operator's need to report blade changes and their approximate time at the end of workday, but on the other hand it requires their intervention each time the monitored blade security-cap is opened to clarify the action reason. The intervention is simplified to a popup in the associated machine HMI, which presents the potential options (see Fig. 3.) that have motivated the cap opening. This pop-up is presented to the operator/craftsmen each time the cap is opened and he/she should only select one of the options. When, and only when, the selected option is related to a blade change action, a data compilation phase is activated,

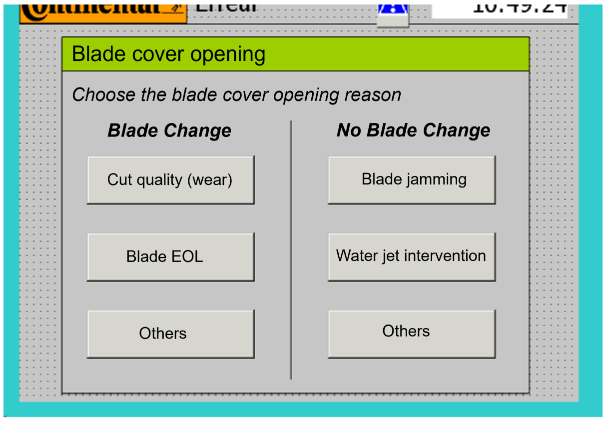

Fig. 3. Continental Shopfloor Interface for trade-blade change explicit-feedback gathering. The language of the operators has been taken into account, so the interface in place is in French, but for better understanding it is translated to English here.

enabling, in real time, the running model to be fed with current real status and so improve further estimations and to update the data to reinforce the model with more accurate data, including a flag indicating the data derived from explicit-human intervention.

The shopfloor HMI was deployed in the Continental plant in 2023-03-15 and is currently working. During this period, the operators have been using it to report the interventions. They have found the mechanism user friendly and no significant increased workload has been identified until now. Fig. 4 shows the annotations that the operators have performed by the explicit feedback mechanism, every time a blade intervention has occurred.

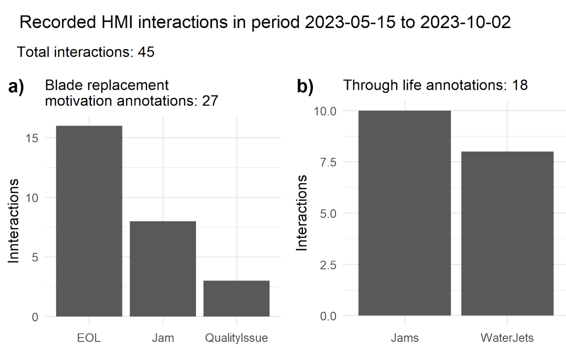

Fig. 4. Summary of annotations provided by users: a) At the removal of the blade b) During the life of the blade.

The reasons that are not related to blade change are also stored and linked to the time they happen but the processing of this data is out of the scope of this work, and will be part of further work. In the period of the 5 months, to date, that the HMI has been in place, a total of 18 interactions related to blade maintenance without blade change actions, have been registered as depicted in the graphic b in Fig. 4..

The annotated data related to blade changes have been used to retrain the model estimating the blade health status. A complete evaluation is a work in progress, but some initial assessments have already been done. Observing the $Rsquare$ metric used for the validation of the fitting[1] of the AI model, the inclusion of the 27 new datapoints obtained through the operator-feedback reporting blade change reason (see details in Fig. 4. a) has improved the value from 0.9683 to 0.9880. This improvement, although small, could be significant on the long run and confirms the human-feedback value for reinforcing AI models.

VI. Conclusions and Further Work

This work presents the research done in exploring the implementation in industrial environments of a reinforcement AI approach by human-feedback. Different technological approaches have been adopted in different industrial scenarios involving AI-based models, solving different real problems, reinforced by implicit and explicit human-feedback gathering workflows. Although the solutions make use of very different technologies, starting from the adaptation of traditional shopfloor interfaces to more sophisticated augmented reality or voice based approaches, human empowerment should be at the center of all of them to ensure a successful adoption of the solution. Accordingly, using the native language of the target industrial scenario, and taking care not to increase the workload of the humans involved, are some of the critical aspects that should be taken into account during the design phase of the Human-feedback AI approach to ensure acceptance and adoption from both the industrial and human side. Furthermore, for industrial acceptance, a non-intrusive improvement of interfaces and mechanisms in place is preferred. Such an approach enables the capture of those (explicit/implicit) operator interactions which directly affect the accuracy of the model, in order to improve data quality, without modifying - or only slightly modifying - the current procedures in the plants.

Further work includes a quantitative evaluation of the impact of the reinforcement approach, with regard to which initial evaluations show promising results, as well as exploring new paths in the human-feedback workflows inline with the identified human and industrial restrictions.

Acknowledgment

This project has received funding from the European Union's Horizon 2020 research and innovation programme project AI-PROFICIENT under grant agreement no. 957391.

References

[1] L. Wang, X. Zhang, H. Su, and J. Zhu, "A comprehensive survey of continual learning: Theory, method and application," arXiv preprint 2302.00487, 2023.

[2] K. López de Calle, E. Garate, and A. Arnaiz, "Towards a Circular Rotating Blade Wear Assessment Digital Twin for Manufacturing Lines," IFAC-PapersOnLine, vol. 55, n. 2, p. 566, 2022.

[3] Meritxell. Gómez, B. Sierra, y S. Ferreiro, "On the Evaluation, Management and Improvement of Data Quality in Streaming Time Series," IEEE Access, vol. 10, pp. 81458-81475, 2022,

[1]Model fitting is a measure of how well a machine learning model generalizes to similar data to that on which it was trained.

[4] S. Zillner, D. Bisset, M. Milano, E. Curry, C. Södergård, T. Tuikka, et al., “Strategic research, innovation and deployment agenda: AI, data and robotics partnership,” 3 ed., 136p., 2020.
[5] X. Wu, L. Xiao, Y. Sun, J. Zhang, T. Ma, and L. He, “A survey of human-in-the-loop for machine learning,” Future Generation Computer Systems, vol. 135, p. 364-381, 2022.
[6] M. Anderson and K. Fort, “Human Where? A New Scale Defining Human Involvement in Technology Communities from an Ethical Standpoint,” International Review of Information Ethics, 2022.
[7] C. Deng, X. Ji, C. Rainey, J. Zhang, and W. Lu, “Integrating machine learning with human knowledge,” Iscience vol. 23, no. 11, Elsevier, 2020.
[8] R. Munro and R. Monarch, “Human-in-the-Loop Machine Learning: Active learning and annotation for human-centered AI,” Simon and Schuster, 2021.
[9] N. A. Wondimu, C. Buche, and U. Visser, “Interactive machine learning: A state of the art review,” arXiv preprint arXiv:2207.06196, 2022.
[10] M. Johnson, A. Albizri, A. Harfouche, and S. Fosso-Wamba, “Integrating human knowledge into artificial intelligence for complex and ill-structured problems: Informed artificial intelligence,” International Journal of Information Management vol. 64, p. 102479, 2022.
[11] S. S. Shuvo and Y. Yilmaz, “Home energy recommendation system (hers): A deep reinforcement learning method based on residents’ feedback and activity,” IEEE Transactions on Smart Grid, vol. 13, no. 4, p. 2812–2821, 2022.
[12] W. Wang, F. Feng, X. He, L. Nie, and T. Chua, “Denoising Implicit Feedback for Recommendation,” Proceedings of the 14th International Conference on Web Search and Data Mining,p. 373-381, 2021.
[13] J. Lie, Y. Ren, and K. Deng, “FairGAN: GANs-based Fairness-aware Learning for Recommendations with Implicit Feedback,” Proceedings of the ACM Web Conference 2022, p. 297-307, 2022.
[14] A. Martinetti, M. Rajabalinejad, and L. Van Dongen, “Shaping the future maintenance operations: reflections on the adoptions of augmented reality through problems and opportunities,” Procedia CIRP vol. 59, p. 14–17, 2017.