AI·PROFICIENT

Artificial intelligence
for improved production efficiency,
quality and maintenance

Deliverable 6.5

D6.5: Best practices and lessons learned

WP6: Use case evaluation and ethical considerations

Task 6.5: Impact assessment and lessons learned

Version: 1.0

Dissemination Level: PU

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

Table of Contents

Table of Contents 2
Disclaimer 4
Executive Summary 6
1 Introduction 7
2 CONTI-2 UC Specification: Restart Setup 8
2.1 Direct Impact / lessons learned 8
2.1.1 Use of AI 8
2.1.2 End user 9
2.1.3 Ethics 9
2.2 Indirect Impact / lessons learned 10
2.2.1 Comments on the requirements Developers Perspective 10
2.2.2 Comments on the requirements Final User Perspective 10
3 CONTI-3 UC Specification: Released extrusion optimization 11
3.1 Direct Impact / lessons learned 11
3.1.1 Use of AI 11
3.1.2 End user 12
3.1.3 Ethics 12
3.2 Indirect Impact / lessons learned 13
3.2.1 Comments on the requirements Developers Perspective 13
3.2.2 Comments on the requirements Final User Perspective 13
4 CONTI-5 UC Specification: Tread blade wear 14
4.1 Direct Impact / lessons learned 14
4.1.1 Use of AI 14
4.1.2 End user 14
4.1.3 Ethics 14
4.2 Indirect Impact / lessons learned 15
4.2.1 Comments on the requirements Developers Perspective 15
4.2.2 Comments on the requirements Final User Perspective 15
5 CONTI-7 UC Specification: Tread alignment 16
5.1 Direct Impact / lessons learned 16
5.1.1 Use of AI 16
5.1.2 End user 16
5.1.3 Ethics 16
5.2 Indirect Impact / lessons learned 16
5.2.1 Comments on the requirements Developers Perspective 16
5.2.2 Comments on the requirements Final User Perspective 16
6 CONTI-10 UC Specification: Quality analysis tool 17
6.1 Direct Impact / lessons learned 17
6.1.1 Use of AI 17
6.1.2 End user 17
6.1.3 Ethics 17
6.2 Indirect Impact / lessons learned 18
6.2.1 Comments on the requirements Developers Perspective 18
6.2.2 Comments on the requirements Final User Perspective 18
7 INEOS2 UC Specification: Image recognition at Geel plant 19

7.1 Direct Impact / lessons learned 19
7.1.1 Use of AI 19
7.1.2 End user 19
7.1.3 Ethics 19
7.2 Indirect Impact / lessons learned 20
7.2.1 Comments on the requirements Developers Perspective 20
7.2.2 Comments on the requirements Final User Perspective 20
8 INEOS1 UC Specification: Reactor stability at Geel plant 21
8.1 Direct Impact / lessons learned 21
8.1.1 Use of AI 21
8.1.2 End user 21
8.1.3 Ethics 21
8.2 Indirect Impact / lessons learned 21
8.2.1 Comments on the requirements Developers Perspective 21
8.2.2 Comments on the requirements Final User Perspective 22
9 INEOS3 UC Specification: Rheology drift at Cologne plant 23
9.1 Direct Impact / lessons learned 23
9.1.1 Use of AI 23
9.1.2 End User 23
9.1.3 Ethics 23
9.2 Indirect Impact / lessons learned 23
9.2.1 Comments on the requirements Developers Perspective 23
9.2.2 Comments on the requirements Final User Perspective 23
10 Conclusion 25
11 Acknowledgements 25

Disclaimer

AI-PROFICIENT has received funding from the European Union's Horizon 2020 research and innovation program under grant agreement No 957391.

Title: D6.5: Best practices and lessons learned

Lead Beneficiary:	Ineos Services Belgium
Due Date:	31/10/2023
Submission Date:	13/11/2023
Status:	Final
Description:	A detailed description of the best practices identified, and lessons learned throughout the project.
Authors	Regis Benzmuller (CONTI), Alexander Vasylchenko (TF), Christophe Van Loock (INEOS), Vasillis Spais (INOS), Giorgos Triantafyllou (ATC), Alexandre Voisin (UL), Marc Anderson (UL), Dea Pujic (IMP), Katarina Stankovic (IMP), Sirpa Kallio (VTT)
Type	Report
Review Status	PC + TL accepted
Action Requested	None – final version approved

VERSION	ACTION	OWNER	DATE
0.1.	First Draft	INEOS	20/09/2023
0.2.	Second Draft	INEOS	20/10/2023
0.3.	Final version for approval	INEOS	06/11/2023
1.0	Final version approved	INEOS	10/11/2023

Executive Summary

This deliverable provides an extensive description both from developers as from end user's perspective on the best practices and lessons learned identified throughout this 3-year project.

Specific attention is also given to the ethics considerations for the use cases.

A distinction is made between direct and indirect learnings. Direct being linked to the AI-PROFICIENT specific use cases (UCx), indirect being more generic.

1 Introduction

This deliverable describes the best practices identified, and lessons learned captured, throughout the project.

A distinction is made between direct and indirect learnings.

Direct relates to learnings specific to the use cases (CONTI-2, CONTI3, CONTI-5, CONTI-7, CONTI-10; INEOS-1, INEOS-2, INEOS-3). Indirect is broader and targets learnings captured throughout the entire project, including the research and development phase, and hence describes more generic best practices and learnings not necessarily linked to the specific use cases.

Direct learnings are as a consequence mainly described for use cases that found deployment.
Indirect learnings are described for all use cases, also those where eventually it was decided not to deploy the service on the shop floor.

The structure of the deliverable is that for the use cases, under direct learnings first a description is given on the use of AI in the specific use case, followed by feedback from the end user and the ethics team.

Under indirect learnings, the developers view is captured, followed by final views from the end user.

2 CONTI-2 UC Specification: Restart Setup

2.1 Direct Impact / lessons learned

2.1.1 Use of AI

In this UC AI has been used to solve two of the problems related to extrusion:

1) The identification of optimal conditions to launch extrusion.
2) The identification of the optimal extruder settings that minimize time and rework.

AI is nourished with the historic data records stored at CONTINENTAL so that the data-driven models can replicate the behavior of the extrusion system. This way, it is possible to simulate the outcomes of the extrusion without having to perform real extrusions.

In this way, it is possible to build a traffic light that warns the operator of the readiness situation of the extrusion line. In addition, a second model can be used by an optimization engine to identify the optimal setting for each recipe, knowledge that can be transferred to the operators to improve the performance of the extrusion system.

The aim of UC2 is to develop a solution that would be able to assist operators in determining the best settings for the extrusion system restart. In that sense, this problem could be considered an optimization problem. However, in order to allow easily iterating over the input settings without interfering with the production, data-driven models have been employed instead of directly tweaking the machine parameters. For that purpose, in contraposition to a physical model, data-driven models can be an appropriate approximation for system that are too complex, however, extensive amounts of data are required for their development. Given that there already was a three year long signal database, the approach was better suited than starting a physical model from scratch.

Prior to modelling the data for the analysis of the extrusions it was required first to extract extrusion information from the signals. This implied characterizing each of the extrusion from the past so that it would reflect the conditions under which it occurred and the outcome of the extrusion. This step known as Feature extraction, is paramount and as it somehow the way the data-analyst defines the problem. In order to succeed, besides spending sufficient time understanding all the aspects involved in the production line, revising feature extraction has been needed. That is, going back and forth with the final user, checking the extracted features and improving them or creating new ones in order to properly capture the behavior of the system. This means that a good understanding of the process is needed to enable a satisfactory use of the AI model. At the same time, it reduces the generalizability of the approach, as any change from the production system would affect the feature extraction making the approach non-replicable. However, all the domain-knowledge acquired during this project would ease the modelling of similar extrusion processes given that they would also have historical data.

Once having a proper set of features, different data models have been tested to determine which one was better suited for the modelling of the extrusion. At the end, Gradient Boosted Trees have been used, however, it is noteworthy to mention that instead of the model used for the modelling, the revision of the features had greater impact on the goodness of fit of the model instead of the model itself. Another important aspect to note is the fact that the historic data had a huge amount of variability regarding the representativity of the different recipes that were produced in the line. Consequently, the model did not fit properly some of the existing recipes. To handle such cases, those recipes that where scarce in the dataset were handled with simpler methods to identify optimal parameters (identifying the best set of parameters from the history in such recipe).

With the surrogate data-driven model as a substitute of the real plant it has been possible to carry out the optimization of the speed settings. The optimization has required again to carefully revise the restrictions to which the system is limited, after all, the search space must be constrained so that the solutions are physically feasible having to iterate again with the end user and refine.

Finally, deployment has taken far more than it was originally expected. Once with the models validated locally, it was expected that putting them in production would not take long. However, the fact that the data that was used initially used locally came now from different data pipes made it more complex to provide a successful inference pipe. In addition, special emphasize was needed to build a robust system that would handle exceptions, outliers, unseen data etc. Given that the project timeline planned to start the deployment of the platform relatively late, those unexpected obstacles have forced the final deployment stages to go a bit fast. If possible, deployment should have been started earlier to allow analyst to understand better the complexities of the deployment and allow iterating from earlier.

2.1.2 End user

In a positive way, the AI has provided innovative suggestions for settings that would typically not have been explored during the start of the Combiline. These suggestions were coherent and have the potential to enhance the starting process.
Concerning the scrap limitation, the proposed solution still requires online testing to validate its effectiveness.

Considering the project's time constraints, there are specific areas that could be further refined to enhance the solution's efficiency. One such enhancement involves implementing a machine service capable of seamlessly transitioning between AI parameters and nominal process parameters for regular production. Additionally, there is a need for automatic proposals to be generated when the machine is not in an optimal state to initiate the production process.
Another point is that the user and machine feedback should be integrated to the model to make it more robust by time and the model should be retrainable via the HMI.

2.1.3 Ethics

Overview: The majority categories of ethical recommendations for Conti-2, were in General AI/Operator(s) interaction (GAI) and Ethics by Design Developer and Industrial Partner Engagement (EtbD). The least successful recommendations in the UC were those which attempted to make use of operator experience as a means of helping the work team adapt to the AI integration. Also, delays in functional design decisions prevented full implementation of some recommendations. The most successful recommendations in the UC were those aimed at minimizing additional cognitive load for the operators by working from the base of existing technologies and adding a non-use option for the operators regarding the AI suggestions.

Lessons learned: 1) the value of making use of operator experience in terms of value for the manufacturer needs to be better researched and the benefits explained to the managing engineers; 2) the ethical recommendations should specify – with the help of the developer partners – and focus upon conditions that will hold irrespective of the final functional design; 3) the workable parameters of *design time* for ethics by design, should be better researched, i.e., what are the windows of opportunity in an AI system integration in manufacturing, when ethics recommendations related to functional design can be anticipated, implemented, and then assessed?

Best practices: 1) ethics teams should provide an additional ‘what is our ultimate goal’ viewpoint and corresponding recommendations, regarding the technology to be added; 2) ethics teams should push for a technical ‘non-use’ option for the shop-floor user or get formally clarified whether the use of new AI systems/suggestions is obligatory.

2.2 Indirect Impact / lessons learned

2.2.1 Comments on the requirements Developers Perspective

This system has answered the requirements that were identified in the beginning of the project. In this way, the solution provides a set of optimal parameters that, according to the surrogate model, will ensure the fastest setup. At the same time, the system "listens" to the process so that, when it is not extruding, it estimates the readiness for restarting the extrusion. In addition, means for gathering user feedback have been implemented so that the system keeps improving over time. The only requirement that is complex to fulfill (or better said, to quantify) is the reduction of rework. Theoretically, this is a side effect of the implementation of the system, however, measuring the exact amount of rework that is caused by an extrusion is complex, hence, it is not possible to have an accurate estimate of the rework that was caused due to an extrusion (nor of its increase/reduction). Nevertheless, considering the creation of rework is linked to the stabilization time (which can be measured) improving the stabilization time would lead to a consequential decrease of the created rework.

In addition, thanks to the AI-PROFICIENT project, and to the necessary meetings that have enabled users and analyst understand each other, certain awareness on how AI works and what AI can and can't do has been transferred to the final users. This will indeed lead to a smarter digitization of the assets so that better and more accurate models can be created on the future by gathering data of greater quality and interest.

2.2.2 Comments on the requirements Final User Perspective

The first lesson learned concerns the difficulty of mutual understanding of the Use Case needs and its implementation.

The different extrusion types and production phase is complex and needs to be assimilated to give good Use Case results. This requires spending many hours on site which was hindered by the COVID-19 pandemic.
For future projects it will be helpful to plan an on-site phase of several days/weeks with the partners in order to fully understand the issues that will need to be addressed in the Use Case.

The second lesson learned is related to the challenge of understanding one another. The industrial and research partners often use different vocabularies, making mutual comprehension at times difficult. It will be good to offer brief training sessions for both parties to facilitate more effective communication on future project.

The third lesson learned concerns the disposition of process/product data for which it was first necessary to find an offline and then online solution. A catalogue of offline/online solution used by other project could be suggested for future projects

3 CONTI-3 UC Specification: Released extrusion optimization

3.1 Direct Impact / lessons learned

3.1.1 Use of AI

In this UC, the AI has been used to solve the problem of constraint that may remain in the thread due to the extrusion process. The extrusion process is a complex process as an AI driven approach has been considered. Two main expectations have been expressed by the end-user:

- AI must alarm operators if the hot area isn't relaxed.
- AI must provide which parameters impact the relaxation level.

As we have considered a data driven approach, the data considered in the learning phase, are the data uploaded by CONTINENTAL on PETA repository. The data covers the period from 2019 to 2022. The amount of data is huge: an average of 280 Mo ranging from 170 Mo to 400 Mo.

A first lesson learned is about data and getting relevant data from the entire dataset provided by CONTINENTAL. The whole CONTINENTAL data have been uploaded on PETA as time series for each month, i.e., 1 csv file for 1 month for 1 sensor. Nevertheless, all the data are not in relation to the use case. As such, in collaboration with the end-user and the combiline engineers, the definition of the condition to extract the relevant data has been a first step. Such a task may seem easy but indeed is not. All people don't have the same view and, despite standard exists, some subjective considerations could be encountered.

The model considered is a deep learning model that has been developed in the project for prognostics purpose. Its architecture is quite simple, but it provides good results. As the model is quite simple, the hyperparameters optimization remains reasonable. Despite we consider an end-to-end date driven approach, deep learning model development and training remain a handcraft process. Indeed, model development and training require business knowledge. Several businesses and use case knowledge can be incorporated in the model and the learning process. Such distillation of knowledge enables, but cannot guarantee, the training will reach the user's needs.

A second lesson learned is about the clear definition of the end-user needs and their translation into proper model development and training parameters. To that end, the end-user and the combiline engineers should have been involved in the development process since the beginning.

After being able to cluster the data based on some input values provided by CONTINENTAL and separate the production/non production phases, we have had a lot of data which is a very positive advantage when trying to develop deep learning models. At first, we tried to use all of the data available to have a better coverage of all seasonality and variability in the data, which was conceptually sound, however we were faced with unexpected challenges arising from our decision to use more than two years' worth of data. In practice this proved to be resource-intensive and time-consuming. A strategic pivot was made to focus on a three-month period to narrow down the use of data. Remarkably, the smaller dataset enabled a similar generalization of performance, although with some reservations about its ability to adequately represent varied conditions, particularly with regard to seasonality and weather factors.

A third lesson learned: was that while having access to a large volume of data can be perceived as advantageous for deep learning applications, it is crucial to assess the trade-off between data comprehensiveness and computational efficiency. Our initial approach of utilizing an extensive dataset led to resource-intensive challenges. Instead of adhering strictly to the philosophy of 'more data is better,' it may sometimes be beneficial to focus on a more targeted subset as it can save time, resources, and speed up the development process.

For implementation of this service, we had two containers, one (Python3.8) for the AI service and one (MySQL) to store the outputs of the service. We used containerization to mitigate compatibility issues and facilitate smooth integration and deployment. The inputs used for this service are more than 75 sensor values and are being read from an influxDB, which resulted in constraining real time deployment. However, by optimizing the service's code, we were able to process data and provide an output in a time of 17-22 seconds, which was approved by CONTINENTAL. One other challenge appeared as we used some values to cluster production phases during our first phases of development, but weren't available on the influxDB, so we had to make an approximation using other sensor values based on CONTINENTAL experts' knowledge.

A fourth lesson learned is about Implementation in real-world scenarios, especially with containerized services and live databases, this requires thorough understanding and anticipation of data availability and compatibility issues. It's essential to validate the availability of all necessary data sources in the deployment environment beforehand. Relying on certain data attributes without verifying their accessibility can lead to unforeseen challenges in deployment. Hence, continuous communication with domain experts, regular checks on data availability, and flexible model architecture can be invaluable for smooth service deployment and achieving desired response times.

3.1.2 End user

AI has offered innovative suggestions for settings, which might not have been explored traditionally. When analysing historical data, AI provides information at an optimal frequency, but validation through online tests is necessary.

The Human Machine Interface comprises two views: one designed for operators and another tailored for process technicians. While the operator view provides sufficient detail, there is room for improvement in the explanatory information on the technician view.

Promising offline data exist and the top-ranked suggested setting, believed to influence relaxation extrusion significantly, needs online testing for validation.

3.1.3 Ethics

Overview: the majority category for ethical recommendations in Conti-3, was General AI/Operator(s) interaction (GAI). The least successful recommendation in the UC was aimed at helping the operator(s) adapt to the anticipated AI model, by setting functional limits for AI/operator interaction. The most successful recommendations were those aimed at clarifying the main goal of the UC, and consequently the role of specific shop floor workers (die makers) with respect to the AI integration.

Lessons learned: 1) recommendations aimed at deliberate functional limitations, need to consider and get better estimates of model implementation and training timelines in order to guide the recommendations through to implementation.

Best practices: 1) ethics teams should help locate and clarify primary goals of UCs and guide the manufacturing partner to remove or simplify any shop floor user interactions with anticipated AI systems which cannot be clarified.

3.2 Indirect Impact / lessons learned

3.2.1 Comments on the requirements Developers Perspective

The high-level requirements for use case 3 of CONTINENTAL were clear: “alert when there is a drift leading to not released thread” and “provide the cause of the drift”. Indeed, CONTINENTAL explained what the process is and what is a released product in a qualitative way from the product “point of view”. Nevertheless, when going into more specific technical requirements, it was not so clear. Indeed, some standards exist to monitor the relaxation of the tread. But, when we have translated into sensor and deep learning technical requirements some misunderstanding raised. This has led to the development of a first model that was not relevant despite it predicting what has been discussed and agreed on. Hence, after further exchange and meetings, new sensor and deep learning requirements and objectives have been defined. Finally, a relevant model has been developed and implemented.

3.2.2 Comments on the requirements Final User Perspective

The first lesson learnt concerns the difficulty of mutual understanding of the Use Case needs and its implementation.
Relax extrusion is complex process which requires spending many hours on site which was hindered by the COVID-19 pandemic.
For future projects it will be helpful to plan an on-site phase of several days/weeks with the partners in order to fully understand the issues that will need to be addressed in the Use Case.

The second lesson learned is related to the challenge of understanding one another. The industrial and research partners often use different vocabularies, making mutual comprehension at times difficult. It is recommended to offer brief training sessions for both parties to facilitate more effective communication.

The third lesson learned concerns the disposition of process/product data for which it was first necessary to find an offline and then online solution. A catalogue of offline/online solution used by other project could be suggested for future projects.

4 CONTI-5 UC Specification: Tread blade wear

4.1 Direct Impact / lessons learned

4.1.1 Use of AI

UC5 has severely suffered from the microchip shortage, this fact has some unavoidable consequences in the UC5 leading to a reduction of the original scope of the potential solution. In such situation, the use of AI has been quite limited as it was meant to be employed in the vision system as a means of processing the images. This shortcoming was tackled by employing a profilometer that already existed in CONTINENTAL, but, in any case, the final solution has employed simpler algorithms than the ones that were originally expected.

Anyhow, data-driven models have been employed to assess the wear status of the blade. Developing such models encountered some problems related to the lack of digitization of certain data on the plant, which lead to a low quality of the data. That meant, for example, that the exact timestamp reflecting the specific reasons for changing a blade were recorder by a free text with inaccurate timestamps as the records were filled later than the replacement occurred.

Thanks to the research made during the project, it has been clearly seen that in order to be able to provide a valid estimate, it is paramount that the data used for both training and deploying the model is of good quality. In that direction, interfaces have been created so that user will know be involved in the data acquisition system and the data will be much more accurate and reliable than before.

Similarly, to UC2, the deployment has required longer and much effort than originally expected. The same recommendation would apply, starting the deployment sooner so that more iterations could take place and allow analysts better understand the limitations and conditions of the final environment.

4.1.2 End user

The AI system provides valuable real-time information on blade weariness, generating detailed reports for operators, and maintenance teams. It incorporates operator feedback and offers two levels of information through the Human Machine Interface, ensuring relevance and clarity.

Online tests should still be done to validate the accuracy of the AI-generated data. If the weariness measurement is accurate, a shift from curative to predictive maintenance will be done, optimizing the timing for blade replacement.

4.1.3 Ethics

Overview: the majority category for ethical recommendations in Conti-5, was Facilitate interaction/engagement with the AI system (IN). The least successful recommendations in the UC were aimed at getting the developer partners and the manufacturing partner to engage other ethical recommendations earlier in the design process and at clarifying a timeline for operator abandonment of older techniques. The most successful recommendations were those aimed at shifting responsibility for a task from operators to maintenance, who were better placed to deal with it.

Lessons learned: 1) more research needs to be done on methods for guiding partner developers to a timelier implementation of ethical recommendations, e.g., a) whether the main design options can be mapped out in collaboration with the developer partners before the functional design is finalized, or b) whether prioritizing recommendations would be useful in getting them implemented.

Best practices: 1) ethics teams should make clear and specific recommendations for reallocation of task responsibility, with regard to work tasks directly or indirectly related to AI system integration, when analysis of the UC context indicates a need for such reallocations.

4.2 Indirect Impact / lessons learned

4.2.1 Comments on the requirements Developers Perspective

As mentioned before, initially the UC would rely on a vision system that would be able to systematically measure the quality of the cut, which as, per today, it is not measured. Unfortunately, due to the mentioned delays developing a vision system for quality measurements has not been possible. As a consequence, some of the KPIs have not been fulfilled (detection of not wear related causes of failure, the improvement of the quality and the detection of deviations on the quality). Yet the approach that has been deployed could have improved the quality of the cut treads, but there is no systematic way of validating such assumption.

Considering the starting point of this UC, it is clear that the grounds for the predictive maintenance of the cutting system have been paved. The data acquisition system has been refined so that the users, with minimal effort, tell the system the exact motivations for the changes in the blades. Such improvement will allow to improve the wear models as the quality of the data will be improved over time as the quality of the data improves. At the same time, additional sources of potential harm are now recorded, which could lead to an improvement of the model used for estimation as it could reflect additional factors that are not considered nowadays. Regarding the capability to estimate the actual wear estate of the blades, without the use of current or vision sensors the system solely relies in the number of cuts, which is known to be inaccurate as there are many existing sources of variation. Yet, with the simulation tool that has been developed based on the historical data recorded, it is possible to estimate to which extent it is economically reasonable to exhaust a blade’s life or if it is best to be conservative and change it before the breakdown occurs.

4.2.2 Comments on the requirements Final User Perspective

In contrast to other use cases, there was no machine sensor available to provide information about blade changes. Initially, feedback was conveyed through a report, but its accuracy was insufficient. Subsequently, a dedicated Human-Machine Interface (HMI) was implemented, allowing operators to select the reason for opening the cover. This led to an improvement in feedback accuracy. To enhance efficiency in future projects, it is advisable to develop use cases around information already available and accessible within the machine.

5 CONTI-7 UC Specification: Tread alignment

5.1 Direct Impact / lessons learned

5.1.1 Use of AI

The UC was not implemented as originally planned. The AI was to be used to associate equipment wear with position drift of tire treads. A method based on model convergence speed was tried but did not deliver the expected prediction. Next, a traditional time series prediction model was tried but it was not implemented in the factory due to issues in integration of the software in the factory setting.

5.1.2 End user

As the UC was not implemented as originally planned, no user input was obtained.

5.1.3 Ethics

Overview: the majority category for ethical recommendations in Conti-7, was Facilitate interaction/engagement with the AI system (IN). The least successful recommendation in the UC was aimed at clarifying Human-in-Command (HIC) expectations. The most successful recommendations were those aimed at clarifying responsibility and scope for the added image labelling work needed to train the AI and centering the alarm solution upon the needs and time constraints of the operators. Overall, this UC had a high rate of implementation of ethical recommendations.

Lessons learned: 1) from the beginning of the design phase the ethics team should guide a review, in collaboration with developer partners, of any technical terms which describe human AI relations, to make sure they are clear, defined, and consistent across all project UCs and work packages.

Best practices: 1) ethics teams should ask for UC clarifications in quantitative terms with regard to AI system training and feedback, whenever the context allows. Quantitative clarification brings out assumptions regarding the non-overt aspects and inputs to the systems.

5.2 Indirect Impact / lessons learned

5.2.1 Comments on the requirements Developers Perspective

The team did not manage to implement a system capable of collecting the required data and the UC had to be abandoned. There are no further specific items to be shared.

5.2.2 Comments on the requirements Final User Perspective

As for the Use Case 5, machine data for this specific Use Cas was absent at the project's initiation. The installation of new sensors, cameras, Industrial Computers was required, consuming a substantial amount of time that could not be dedicated to algorithm development. To enhance efficiency in future projects, it is preferable to formulate use cases around existing and readily accessible machine information.

6 CONTI-10 UC Specification: Quality analysis tool

6.1 Direct Impact / lessons learned

6.1.1 Use of AI

This use case aimed to develop AI-based quality analysis and assurance tool, that would assist the factory personnel in case of occasional anomalies in process setup, that could result in the production of pieces of tread out of the desired scope. To determine what caused the problem when it occurs, the quality manager does a thorough manual inspection of the process parameters along the entire process line. The use case solution, including early anomaly detection, root cause identification, and generative process optimization, assisted by predictive surrogate models of system responses (product characteristics indicating production quality), should automate this activity.

However, the introduction of an AI-based tool, created from scratch as a result of research activity, in a conventional, well-established manufacturing control procedure may cause some resistance in the factory personnel. Having that in mind, the solution has been designed as the co-creation of good traditional modelling and control concepts on one side, and SoA deep-learning and self-learning concepts on the other, boosted with explainability and flexibility, which is a key point for achieving the trust of the user. Further, the tool has been created following the ethical recommendations, and adapted to meet end-user preferences, as much as it was possible.

6.1.2 End user

The service provides operators with corrective parameters for current or anticipated deviations, aiding daily tasks. Historical data confirms accurate deviation predictions, but the relevance of suggested machine settings remains unverified. AI-PROFICIENT offers when deviations are detected up to 5 optimized machine settings for operators. The Human Machine Interface displays suggested parameters clearly. Historical data evaluation boosts confidence, awaiting validation through online tests. AI-driven actions prevent further corrective measures, pending online validation. AI generates innovative suggestions, needing confirmation with online data.

Automatic feedback from machine and user feedback from Human Machine Interface should be integrated to the service to make the model more robust.
A retraining model option should be available from the Human Machine Interface.

6.1.3 Ethics

Overview: the majority category for ethical recommendations in Conti-10, were Identification and minimization of (additional) workload (WkL) and Facilitate interaction/engagement with the AI system (IN). The least successful recommendation in the UC was aimed at estimating in advance the increase in operator and quality manager workload due to feedback. The most successful recommendations were those aimed at de-anthropomorphizing discussion of AI systems in early deliverables and getting some early design outlines of human centered formats for AI system guidance to operator. Overall, this UC had the highest rate of implementation of recommendation for the project.

Lessons learned: 1) To prevent recommendations being given early but *effectively* forgotten by the time the related development process arrives, Ethics by Design needs to be *development time aware*. Early recommendations for estimates regarding system elements whose design will come relatively late in the project, should a) set flexible timelines for their implementation which take into account usual development times and delays, and expressly remark that their implementation is expected later, or b) be retained by the ethics team alone until the actual development catches up and then presented to the partners.

Best practices: 1) ethics teams should use recommendations and discussion in ways which remind project partners of the main goals which need to be satisfied by the UC, in order to keep solutions simpler. 2) Ethics team should participle more directly to help developer and manufacturing partners implement recommendations.

6.2 Indirect Impact / lessons learned

6.2.1 Comments on the requirements Developers Perspective

Quality of data, data availability and developers' understanding of data all have a big impact. Good understanding of the related process was necessary in order to implement the UC solution.

In general, the developers' understanding of a given manufacturing process is not at the greatest level, necessitating knowledge exchange and occasional briefing sessions with the domain experts, particularly process engineers. Because of this, several meetings with the tool's developer partners and pilot representatives, as final users of the tool, have been held, and the development process has been made easier by the prompt delivery of all necessary information.

While working on the tool, supposed to comprise several modules, developed by different partners, good communication between them is required to reach a consensus on how work on components development and final integration should be carried out. Before integrating the components, thorough planning should be done, including defining the functions that integration should support, the format of data to be exchanged, and the technology to be in service of integration. To make work easier for current and potential future contributors, documentation including information about the conducted work should be made accessible. For that purpose, the UC leader provided the UC partners with cleansed, pre-processed training data, making those files available on joint sharing repository, while all the relevant information related to work on the use case, has been described within the relevant deliverables, by the UC partners in charge of corresponding tasks, and thus made available for whole consortium.

Concerning the final result, or particular functional requirements of the tool (to ensure the good quality of the production, to detect deviation for the bad quality treads, to identify the cause of the quality deviation, to optimize the current process parameter settings), they have been accomplished, by the use case services implementation, and the level of success will be assessed, through the validation process. Anyhow, to have modules ready to be validated by end-users, modules need to be moved from the development environment to the production environment, which had been challenging for various reasons, mainly due to security measures that made integration non-trivial.

6.2.2 Comments on the requirements Final User Perspective

The first lesson learned concerns the difficulty of mutual understanding of the Use Case needs and its implementation.
Combiline process complexity requires spending many hours on site which was hindered by the COVID-19 pandemic.
For future projects it will be helpful to plan an on-site phase of several days/weeks with the partners in order to fully understand the issues that will need to be addressed in the Use Case.

The second lesson learned is related to the challenge of understanding one another. The industrial and research partners often use different vocabularies, making mutual comprehension at times difficult. It is recommended to offer brief training sessions for both parties to facilitate more effective communication.

The third lesson learned concerns the disposition of process/product data for which it was first necessary to find an offline and then online solution. A catalogue of offline/online solution used by other project could be suggested for future projects.

7 INEOS2 UC Specification: Image recognition at Geel plant

7.1 Direct Impact / lessons learned

7.1.1 Use of AI

GoogleOCR is used to convert an image taken of a label into text. This saved a lot of time writing a handrolled solution. Many other OCR programs and cloud services were explored, including Azure Recognize, Cognex, Merlix, PaddleOCR, and PyTesseract. Google CloudVision was selected for outstanding accuracy and speed, with no finetuning necessary. We suspect that the open-source solutions can be brought to a similar accuracy with case specific finetuning, at higher speeds because they don't have to run in the cloud.

Another AI module based on the FuzzySharp library is used to find known product names and numbers in the parsed text. Fuzzy matching is a simple algorithm which could be handwritten, but the library's API made experimenting with different pipelines exceedingly easy. While fuzzy matching can be used out-of-the-box to cope with small discrepancies between parsed text and the given data, it does need to be augmented with expert knowledge about the data, e.g., abbreviations, different ways of formatting the same number etc. Because these variations seem obvious to shop floor operators, continuous testing and feedback is necessary.

7.1.2 End user

An analysis of the impact was done by the production engineer, and a survey to capture feedback on the use of the tool/service was performed amongst users/operators (detailed description in D6.3).

On the service: As the service was only deployed in April 2023, the time window on measuring direct impact is limited within the scope of this AI-PROFICIENT project. Nevertheless, both from an impact as from a usability point of view, the end user is happy as performance is in line with the objectives.

There is further room for improvement in the usability of the tool as the obligatory gloves do not always work reliably with the ATEX tablet. A voice-controlled device could potentially provide a solution.

A clear lesson learned is that despite incomplete data, AI can generate satisfying results. More specifically the quality of the images to be recognized in an industrial environment can be rather poor (plastic fools, dust, different shapes and layout), but by training the algorithm satisfying results can be obtained.

7.1.3 Ethics

Overview: the majority category for ethical recommendations in Ineos-2, was Ethics by Design Developer and Industrial Partner Engagement (EtbD). The least successful recommendations in the UC were aimed at establishing formal protocols in case of error by the AI system and getting clear and definite reliability rates for the system. The most successful recommendations were those aimed at better considering the operators' and engineers' needs (human centering) through rethinking which off the shelf technologies would be used and reconsidering the operator's physical context.

Lessons learned: 1) efforts to get metrics related to AI system reliability and thus potentially the formal abandonment of lines of research or development are difficult to implement. More research needs to be done on alternative, e.g., positive rather than prohibitive, ways of framing the use of metrics and benchmarks. Perhaps that can be done by outlining metric levels and redefining those levels as public relations value.

Best practices: 1) ethics teams can profitably bring attention to particular issues in the operator (user) context and to the differences between proposed technology solutions relative to that context and recommend changes. Developing a practical layperson's understanding of the context and discussing with the industrial partner – to better understand the context – are essential steps in this process.

7.2 Indirect Impact / lessons learned

7.2.1 Comments on the requirements Developers Perspective

The main requirements of the INEOS2 UC serve to streamline and enhance the additives check by operators on the factory floor. The additives check is when an operator must connect and additive to a feeder, he must make sure the correct additive goes to the corresponding feeder.

The application streamlines this action by allowing the operator to scan the feeder at which he is located and scan a label with a mobile device. Matching the additive and the feeder and receiving feedback from the quality control system at the location of the feeder. The application allows the operator to receive feedback from the quality control system immediately after scanning the additive and feeder. Entirely removing the need for the operator to manually type the additive and feeder in a terminal.

It is possible that the factory floor has intermittent internet connection issues. Therefore, a recurring request repeater is implemented in the application to ensure that when a request doesn’t immediately go through that the application does not halt its flow but repeats the request at certain time intervals.

The user interface had to be customized by keeping in mind that an operator would be on a factory floor with full equipment on, with gloves and eye protection. The application was introduced to the relevant partner for testing and feedback. Their feedback was incorporated into each iteration of the application. After each round of iteration, the application and user interface would be improved and refined.

An important requirement is that the operator must be the one in full control of the application. The AI services serve to enhance operators' decision making. This is accomplished by having the operator check, confirm and allow them to manually make changes when needed to override the system’s results.

7.2.2 Comments on the requirements Final User Perspective

This project has grown confidence within the company that AI can act as an enabler to improve processes. Whereas AI has been ‘a buzz word’ with little known concrete use cases in the past, this view has evolved.

Close cooperation between developer and industrial partner is key, as multiple iterations are required to have a product ready to market.

8 INEOS1 UC Specification: Reactor stability at Geel plant

8.1 Direct Impact / lessons learned

8.1.1 Use of AI

The objective was to develop a reactor model / digital twin allowing to support the operator in reducing reactor instability in the real live plant.

8.1.2 End user

The use case did not find a solution, no service was deployed.

8.1.3 Ethics

Overview: the majority categories for ethical recommendations in Ineos-1, were General AI/Operator(s) interaction (GAI), Error Handling (ERRH), and Facilitate interaction/engagement with the AI system (IN). The least successful recommendations in the UC were aimed at clarifying responsibility for deciding when the AI tool is reliable and acknowledging that operator knowledge cannot practically supplement the *complexity* aspect of the AI analysis with regard to reliability. The most successful recommendation was that of simplifying the use case by focusing on exploratory use of the AI at the level digital twins alone. Overall, this UC had the lowest rate of implementation of ethical recommendations for the project.

Lessons learned: 1) further research on *AI complexity overload* situations, i.e., situations where the analysis of a process by the AI system is visible – not black box – but prohibitively complex in terms of the operator or engineer being able to relate the AI suggestions to their own experience and human capacities, needs to be carried out. The implications for human oversight of AI systems are most important here, for example: are there alternatives to oversight?

Best practices: 1) ethics teams can help to reconsider design solution stages by re-focusing on the human aspect of these stages and the problems that planned solution stages might bring to the operators and engineers (end users).

8.2 Indirect Impact / lessons learned

8.2.1 Comments on the requirements Developers Perspective

Requirements, whether they were met or not, which were the most problematic and why :

The main requirement was to find a method to ensure the stability of the reactor by optimizing the control loop and by giving advices to the console operator. To achieve this, the causes of the oscillations needed to be identified.

Data was available from a long period allowing using different AI methods with the aim to find reasons for oscillations. Correlations between certain process conditions and oscillations could be found but these were related to certain products and not controllable properties as such making the result unusable for process improvements. Discussions with process engineers were invaluable when interpreting the data analysis results.

In a reactor, measurement possibilities are very limited, and for that reason physical and hybrid modelling were concluded to be the only alternatives to identify the causes and through that means to try to develop remedies. Thus, the goal was to develop a dynamic hybrid model that could be used to analyze process transients especially in conditions in which problems have been observed. The initial process description didn't fully reveal the complexity of the process and especially of the process control. Thus, the work started based on limited understanding of the process. As the work progressed, more questions arose and several discussions with process engineers, detailed drawings and descriptions of the control system were finally required to understand all the details that need to be included in a predictive process model. In this process discussions with process engineers were invaluable.

The selected physics-based modelling approach included modelling in two scales: 1) a dynamic 1D process model based on descriptions of all interconnected unit processes, a 1D description of the reactor, and detailed description of process control and 2) a dynamic 3D reactor model based on computational fluid dynamic (CFD) description of the phenomena inside the reactor. Information on mass and heat transfer produced by the 3D model is used to set the mixing description in the 1D model. The purpose was that the 1D model could be used as a digital twin to assist operators.

The complexity of the process that became clear during the work lead to a failure in reaching the implementation stage in the project. The 1D model, which would require detailed description of all control loops and the large numbers of sub-processes, proved too slow and complicated to be used as a digital twin and finally the work focused on producing a 1D dynamic model for the reactor and its controls. However, the computational speed could be improved in the future with improved algorithms if the work continues in some other project which could facilitate a digital twin of the full process.

To overcome computational limitations, CFD modelling was done with two methods: resolved modelling of hydrodynamics in fine computational meshes and time-averaged filtered modelling of all relevant phenomena in a coarse mesh. The resolved simulations were used to derive filtering closures for the filtered modelling approach. The filtered 3D CFD model describes mixing in the reactor as planned but requires further testing and integration with all sub phenomena. Slow computational speed limits the usability of the resolved simulation method in modelling of the entire reactor during required long time periods while the resolved model is designed for that type of simulations.

The project produced new modelling tools for polymerization reactors. Concerning the final result, the goals of the use case were not reached. The modelling tools developed in the project need further testing, and demonstration simulations, in worst case also further development, to use the models for identifying and finding corrections to the conditions leading to observed unwanted process behavior. Since the implementation stage was not reached, the ethical considerations identified in the beginning were not addressed.

8.2.2 Comments on the requirements Final User Perspective

Strong partners even with limited experience within the petrochemical industry can provide valuable contributions related to data analysis and development of digital twins / reactor models.

The use case however was so challenging that a mature solution ready for deployment could not yet be finalized.

The Covid period has proven to be an extra barrier reducing opportunity for success on this use case. Many discussions can be held virtually, but within an international consortium of partners not used to work together, frequent face to face visits of the production assets in real live do make a difference. This was not possible and is potentially a contributor to not having a service ready for deployment.

9 INEOS3 UC Specification: Rheology drift at Cologne plant

9.1 Direct Impact / lessons learned

9.1.1 Use of AI

Developed an AI service which could predict and avoid drift of product quality in the plant. This has not been successful.

9.1.2 End User

No deployment of services

9.1.3 Ethics

No ethics review on this use case as not relevant

9.2 Indirect Impact / lessons learned

9.2.1 Comments on the requirements Developers Perspective

One of the most challenging aspects of the work on INEOS3 use case stemmed from the available historical data set. While thorough historical data analysis was conducted with a primary focus on detecting the most probable degradation causes, limitations due to data availability were encountered. Rheological measurements were limited due to the nature of their sampling, despite the efforts to extend data collection. This limitation emphasized the importance of data completeness and accessibility, serving as a reminder of the critical role data plays in the success of such projects. Additionally, it was found that the understanding of the production process of the corresponding developers and was greatly enhanced by having several meetings with plant representatives and conducting two plant visits. This close collaboration was invaluable in gaining insight into the nuances of the production environment and ensuring that the conclusions were both relevant and practical.

One of the key takeaways from working on INEOS3 use case was the successful identification of the most probable causes of quality drift, a critical use case requirement. Through rigorous data analysis based on various AI techniques and collaboration with plant representatives, the potential quality drift culprits were pinpointed. However, it is worth noting that the process parameter settings optimization and the development of the operator recommendation system did not reach the same level of fruition as we had initially envisioned. Due to practical constraints of data limitation and time and resource constraints, it was decided to focus the use case on offline data analysis. This experience highlighted the importance of adaptability in project development, where sometimes, adjustments are necessary to align with real-world constraints.

Overall, INEOS3 use case of AI-PROFICIENT project has reinforced the importance of adaptability, data quality, and close collaboration with domain experts in tackling complex real-world challenges resulting in valuable offline conclusions which could improve process line in the following period.

9.2.2 Comments on the requirements Final User Perspective

Strong partners even with limited experience within the petrochemical industry can provide valuable contributions related solving complex use cases within the industry.

A learning here for future projects would be to perform a detailed evaluation and assessment of the quality and availability of production/sensor data. For this use case multiple parallel activities have been performed with data analysis being one of these. In hindsight it would have been better to focus all efforts on quality of sensor data prior to launching other activities. During the projects, an agreement was made to increase the amount of sensor data and sampling significantly to allow better analysis of the data a find correlations. The time window of these additional data sets was limited.

Recommendation for future projects would be to clearly identify and select use cases where sensor data is readily available, both in terms of frequency (time interval between to data points of the same sensor) and quantity (how many online sensor do we have, instead of having to take manual samples which are to be analyzed in a lab).

10 Conclusion

The deliverable contains valuable information for future projects related to deployment of AI services in an industrial environment.

Direct learning relate directly to the AI-PROFICIENT specific use cases.

Indirect learnings are more generic and are applicable to and relevant for a broad range of projects where technical and industrial partners team up to design and deploy AI services. The dedicated sections on ethics provide relevant insights useful to other projects and could be incorporated in an early stage to these projects in an effort to have high ethical standards embedded in these projects as from the design phase.

11 Acknowledgements

This project has received funding from the European Union’s Horizon 2020 research and innovation program under grant agreement No 957391.