La Metallurgia Italiana
International Journal of the Italian Association for Metallurgy

DOI 10.36146/2026_01_25

Digital Twin Framework and Federated Learning for Multi-plant Knowledge Sharing in Decision Support for Electric Steelmaking and Beyond

P. Kannisto, B. Kleimt, A. Chandgude, N. García, I. Guardiola, R. Lazcano, C. Perea, S. Dettori, A. Zaccara, L. Vannini, V. Colla, D. Blazquez, B. Fernandez, I. Garcia

The processes of electric steelmaking are complex and difficult to control to achieve sustainable production. To strive towards competitiveness and green transformation, steelmakers apply the Electric Arc Furnace (EAF) to circulate scrap into new products. This saves resources compared to iron-ore-based production, enabling both circular economy and green energy sources. However, efficient EAF operation faces difficulties in state monitoring and control decisions.

Fortunately, the control can be facilitated with optimization based on Artificial Intelligence (AI) and Digital Twins (DT). Still, DT accuracy can suffer from input data fluctuation or coverage limitations in development and validation datasets. The fluctuation stems from the environment-related variation, especially the scattering and not-exactly-known chemical composition of secondary raw materials.

For decision support, this work suggests a DT framework with Federated Learning (FL) for multi-plant schemes, focusing on electric steelmaking. The framework can deliver both historical data and message-oriented online data to the DTs. It builds upon a container orchestration system (Kubernetes) for software lifecycle management and resource scaling. Importantly, the framework implements FL to exploit network-wide knowledge. That is, the DTs share knowledge with a centralized server that aggregates a global model distributed to the participants, broadening data diversity. Still, all data remain local, which preserves privacy.

The framework applies FL for two types of process DTs, EAF and the subsequent Ladle Furnace (LF). FL can optimize EAF parameters although EAF is not AI but a dynamic model. Conversely, the LF model is composed of a set of neural networks. The results from a prototype system with actual data prove the concept. Firstly, the DTs accurately estimate process variables online, such as the chemical composition and temperature. Secondly, FL experiments indicate potential for model parameter optimization and enhanced performance. Besides, the framework concept is applicable for even more DTs and across industries.

KEYWORDS: DATA-DRIVEN SYSTEMS, ARTIFICIAL INTELLIGENCE (AI), PROCESS INDUSTRY;

INTRODUCTION

The processes of electric steelmaking are difficult to control in an optimal way [1], which makes it beneficial to share control-related knowledge between multiple plants within the same enterprise. The optimization can occur with Digital Twins (DT) and optimization models, based on Artificial Intelligence (AI) or other techniques. Because the coverage of measurement data has inevitable limitations, it is advantageous to share the related knowledge between plants. This can occur possibly within the same enterprise or co-operating partners if their production processes have similarities.

Petri Kannisto, Bernd Kleimt, Akhilesh Chandgude
BFI (VDEh-Betriebsforschungsinstitut), Düsseldorf, Germany

Nicolás García, Irene Guardiola, Raquel Lazcano, Carmen Perea
EVIDEN, Madrid, Spain

Stefano Dettori, Antonella Zaccara, Lorenzo Vannini, Valentina Colla
Scuola Superiore Sant'Anna, Pisa, Italy

David Blazquez, Bárbara Fernandez, Irene Garcia
Celsa Group, Barcelona, Spain

To enable knowledge sharing between production plants, this article presents a DT framework with Federated Learning (FL) to let the local optimization models of each plant to learn from the others. In FL, the participating parties share their local model with a centralized service that creates global weights, effectively isolating the local datasets and models yet enabling network-wide benefits [2]. The framework has received its motivation from electric steelmaking, but the concept is generalizable across process industries and other domains, such as manufacturing. The research has been carried out as a collaboration between multiple partners within the research project ALCHIMIA.

The research question of this article is: What kind of framework can operate Digital Twins and optimization models, including AI, for data-driven purposes to enable Federated Learning for network-wide benefits between industrial plants when multiple organizations develop the software components?

Earlier, the same topic has been discussed in a conference presentation [3]. This article is an extension to the conference abstract and presentation.

The structure of this document is based upon the method Design Science Research (DSR), which aims to build solutions based on tangible, relevant requirements, contributing to the scientific rigour [4]. The next section surveys the state of art, setting the scientific background. Then come the requirements for the framework, followed by the design and the related proof of concept. Finally, the results are discussed and concluded.

RELATED WORK

Earlier projects have aimed to increase the intelligence level of industrial process control with AI and related technologies. For instance, DTs can be a part of a system that forms cognitive capabilities for production optimization, as shown in CAPRI [5] and COGNITWIN [6]. On the other hand, AI systems can develop autonomous Self-X functionality with the help external supporting entities, such as the Autonomic Manager as shown in s-X-AIPI [7]. In similar cases, FL would enable sharing the knowledge of the optimization models.

Regarding FL, most earlier manufacturing- or process-industry-related works have focused on the mathematical aspects instead of information-system-level questions with production plants and DTs. As an exception, a system design, not only a structure but also operational workflows, has been proposed related to industrial asset data [8]. Additionally, another system design has been proposed for algorithms in condition monitoring [9]. Across domains, the recently identified FL challenges include scalability and resource constraints, privacy preservation, heterogeneity as well as FL applications in new sectors [10]. Generally, the FL framework suggested in this work, concretely integrated with production plants and DT, is situated in a research gap.

REQUIREMENTS FROM INDUSTRIAL OPTIMIZATION

General functional requirements

The main purpose of the system is to provide decision support for process optimization. That is, there is a Human in the Loop to apply their final judgment instead of any direct process control by the models. The decision support relies on the measurement data available from the actual processes.

Particularly in this work, the focus is on two-unit processes: the Electric Arc Furnace (EAF) and Ladle Furnace (LF). These are a part of the common production route in electric steelmaking (see Figure 1). Within the context of the relevant research project ALCHIMIA, multiple scrap-handling-related aspects were considered as well. These include, for instance, scrap mix optimization and scrap processing, but these are out of scope. Additionally, any steps from casting onwards remain as future work.

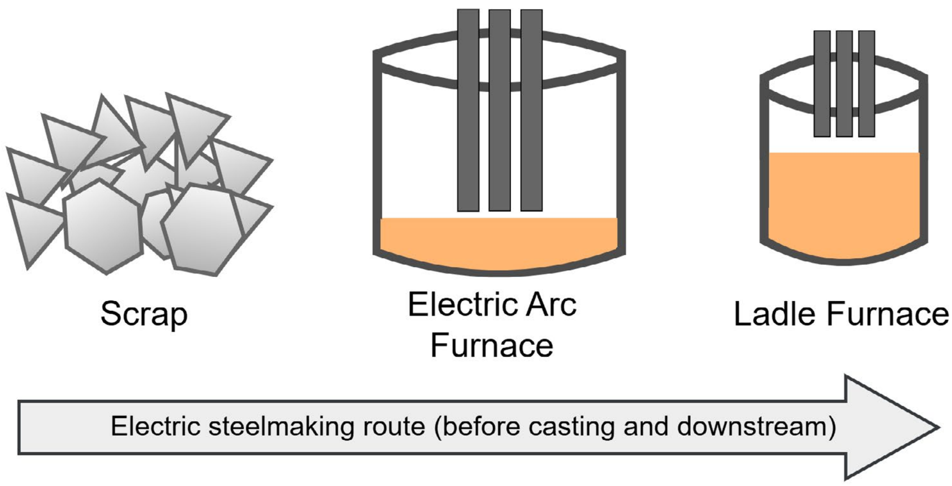

Fig.1 - Electric steelmaking route.

The system should apply FL to generate benefits within an ecosystem of participating production plants. The framework should enable this for any optimization model as long as there are similar models in the ecosystem and the principles of the model itself do not prevent FL. Concretely, FL is inapplicable if the source data is too heterogeneous between the plants or if the underlying principle of the model applies no learning of any kind. For instance, the scrap data in electric steelmaking typically differs considerably between plants. Besides, FL is typically applied to Machine Learning models, and an explicitly developed model can be inapplicable.

Development-related functional requirements

The system should follow DevOps and MLOps [11] principles to enable automatic or close-to-automatic, iterative, cyclic development with early feedback about error. These principles enable the developer team to deploy without relying on a separate operations team with possibly conflicting priorities (such as the avoidance of repeated updates). Despite the name, the MLOps cycle can validate and optimize a variety of models regardless of the technology, not only to Machine Learning.

Concretely, the systems should be deployed with the Infrastructure as Code (IaC) principle to avoid manual installation. Technologies, such as Docker and Kubernetes, enable this by letting the developer define textual manifest files. These will be interpreted by the infrastructure, which automatically sets up the system based on the developer's definition instead of a manual installation and configuration for each component. This enables a unified, cloneable environment to increase productivity.

Non-functional requirements

Security is paramount in any modern information system. A professional security approach builds upon a risk assessment, and the so-called CIA triad (confidentiality, integrity, and availability) is the basis for all information security. Additionally, the related term cyber security widens the scope to protecting the environment, people and assets instead of mere information.

Scalability is another key factor, even in industrial plants, due to the high amount of data generated as well as unforeseen technological developments. That is, the design must consider a scenario where the system grows in size, complexity and data volume. Concretely, any centralized tools, such as message brokers and software infrastructure, should support load balancing and encourage the elimination of direct point-to-point dependencies between software components.

Heterogeneity is an inevitable feature in industrial data. There is a variety of processes, equipment types and software components, and these can come from multiple manufacturers. Besides, new technologies will introduce unforeseen data-related needs. This means that the system must enable diverse platforms and technologies, preferably applying standards or other well-established methods for data integration.

FRAMEWORK DESIGN

Figure 2 illustrates the logical structure designed for the Digital Twin and Federated Learning framework. The following paragraphs will explain these aspects, starting from the overarching security, end user goals and raw data and leading to the optimization models and FL.

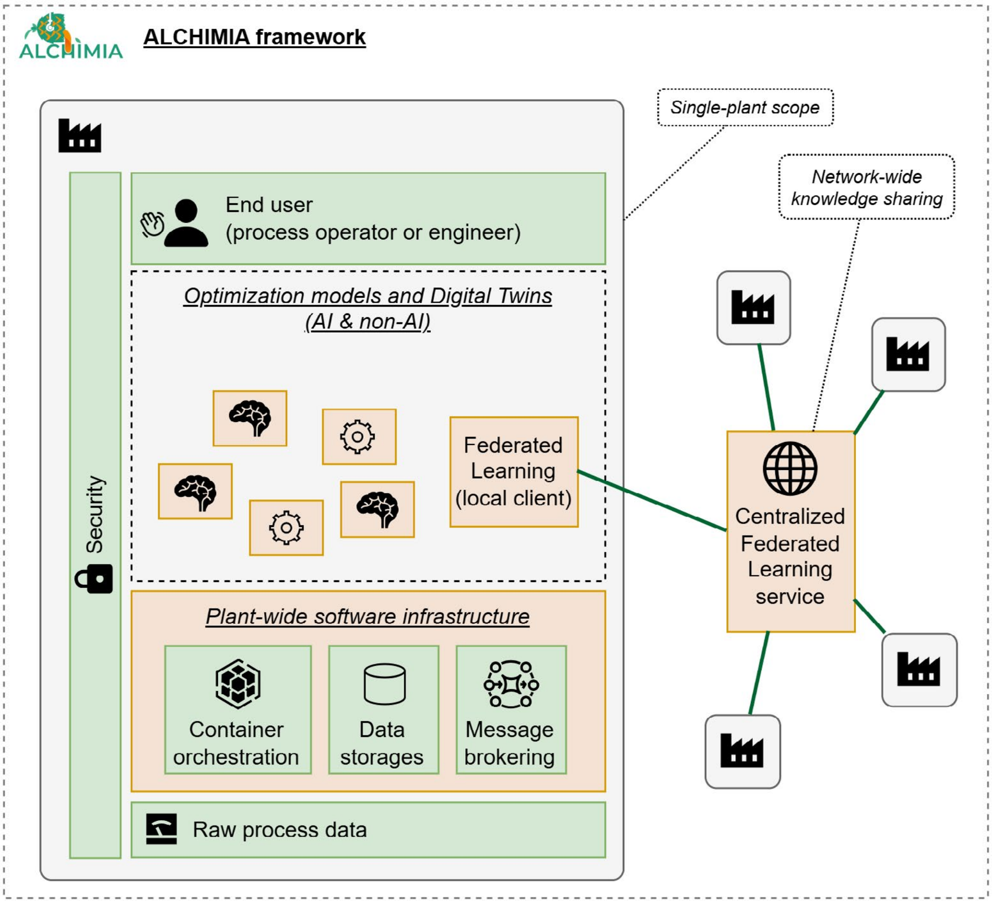

Fig.2 - ALCHIMIA framework from the viewpoint of a production plant.

Security

While the concrete security measures depend on the security policy of each participating plant, certain elements are generic. In the user and developer scope, each person is authenticated upon accessing the end user interface or the developer interfaces. In the software component level, each network endpoint enforces user authentication and access control. Many components provide built-in authentication mechanisms for state-of-art tools, such as single sign-on and Lightweight Directory Access Protocol (LDAP) to authenticate against existing credentials. Additionally, the users can have finer-grained permissions to assign administrator, developer and end user roles, possible on the level of a certain unit process or software component.

Besides user authentication, security zones and domains are in place. The defence-in-depth principle creates layers of security, and the decision support systems should be physically separated from production systems.

Additional mechanisms are applicable for availability and consistency. In error situations, data loss can be avoided with backup systems. In case of an unauthorized use or erroneous software, logging enables activity tracking. More mechanisms are to be deployed as indicated by a continuously active security assessment process.

End user

From the framework viewpoint, the "end user" aspect refers to the various decision support interfaces created for the DTs. In the steelmaking context, the end user should receive benefit thanks to stable recommendations, improved temperature predictions and reduced guesswork.

To emphasize that these are DTs rather than Digital Shadows, there can be optimization functionality to let the operator define setpoints. For instance, in relation to the EAF process, the setpoints can include the tapping temperature, the target carbon content and the target meltdown degree to trigger each scrap basket addition. Thanks to web technologies, the user interfaces can potentially be accessed from anywhere, but practically the access is likely restricted to control rooms or at least to the domain of the enterprise.

Raw process data

The measurement data from the processes in the foundation of the DTs, and the framework places no restrictions related to the physical data sources or formats. Concretely, the data sources can be databases, control systems, Manufacturing Execution Systems (MES), sensors or anything available at the plant.

Plant-wide software infrastructure

The infrastructure provides a layer of services to supply data for the DTs in a suitable format. For historical data, databases are applied. Respectively, any event-based, message-oriented data are best supplied with a message broker, such as Apache Kafka or a product compliant with Message Queueing Telemetry Transport (MQTT). For example, historical data can describe scrap additions and past calculations for the process states. On the other hand, the actual event of scrap addition can be notified with a message describing the scrap types and masses. Because any data delivery necessitates preprocessing, the infrastructure includes appropriate components, referred to as Extract-Transform-Load (ETL).

Additionally, the infrastructure provides an execution environment with IaC for remote deployments to facilitate software development in external organisations. For IaC, Kubernetes provides the foundation. This enables the developers to build their components, such as optimization models and the user interfaces, as microservices using textual manifest files. This eliminates the need to manually install and configure each software component as well as facilitates the setup of testing environments that resemble the actual production environment. This automation enables DevOps and MLOps to increase the automation degree of software updates. Additionally, the container orchestration system provides a physically separated network, which isolates the software from the outside world, facilitating the configuration of the security measures. Finally, Kubernetes promotes scalability now that load balancing is built-in feature in the platform, and the software component instances (referred to as "pods") can be configured for mutual load balancing. These benefits apply to each optimization model and user interface.

Optimization models and Digital Twins

The optimization models and DTs operate on the data from the ETL and generate decision support for the end user. The concrete functionality depends on the use case, but a usual approach is to create at least one DT or model for each unit process being optimized. In this work, this refers to EAF and LF.

Federated Learning

FL enables enterprise-wide or even ecosystem-wide benefits depending on the connected organizations (see figure 3). In this scheme, each participant shares its local model with the centralized service, which will generate global weights for the participants. That is, no source data is shared within the network, and only the centralized service sees the properties of the local models. Thus, in principle even competing organizations could participate in the same network. On the other hand, even if a single enterprise operates all the plants, the scheme reduces concerns related to data sharing by restricting the scope.

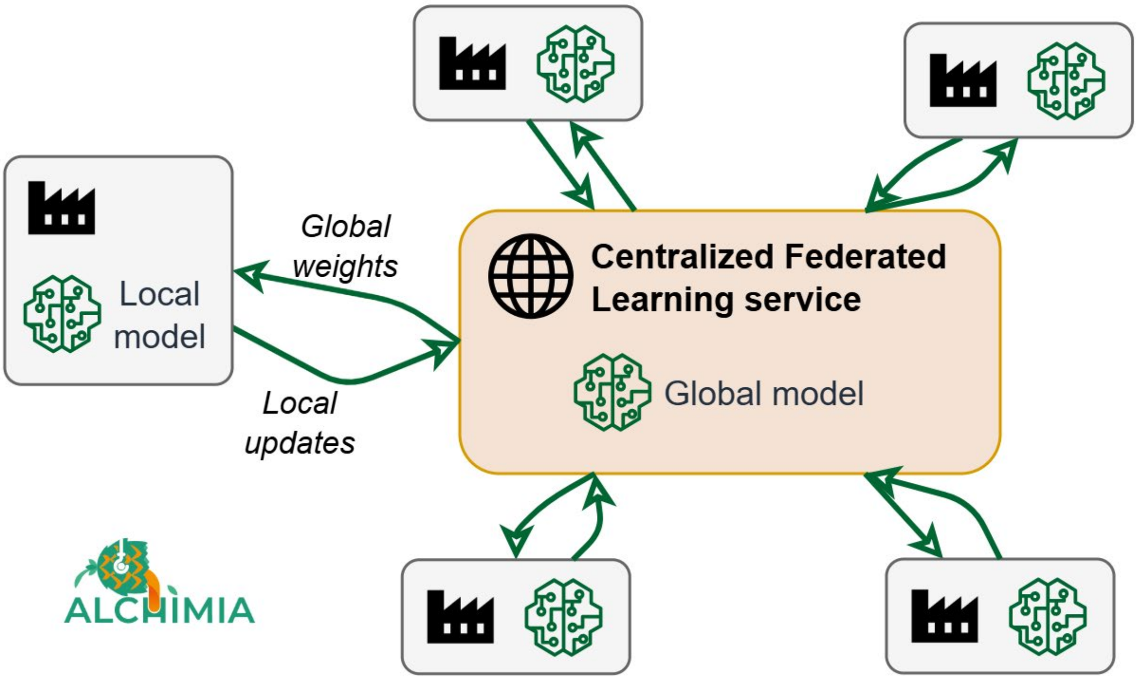

Fig.3 - In Federated Learning, the global model enables the local models to be improved.

In ALCHIMIA, each participating plant operates a client to send its local updates to the central server and receive global weights to learn from the others. Despite the asymmetries between the plants, such as the computational infrastructure, the physical processes and the scrap types, the steelmaking processes share the same properties and measures, which enable FL. ALCHIMIA applies FL to improve the performance of EAF and LF optimization models.

PROOF OF CONCEPT WITH INDUSTRIAL DATA

Software infrastructure and components

For the proof of concept, the software infrastructure was built on top of Kubernetes. In this case, it is shared between three plants within the same enterprise. With Kubernetes, the software development can occur with any language and environment as long as this can be packaged into a suitable virtualized software image. In practice, the developers applied Docker to build their images and then push these into a remote image registry. Then, the software was deployed with Kubernetes manifests that define the address of the image in the registry as well as the configuration of the environment. This includes but is not limited to network addresses, user credentials to other applications (such as databases) and volumes for persistent data.

The data integration occurred with two main tools depending on the needs. Firstly, any historical data was stored in Structured Query Language (SQL) databases that operated in the Kubernetes environment. Secondly, any event-based, message-oriented data were brokered with Apache Kafka. Kafka enables topic-based publish-subscribe communication, designed to be scalable even when the number of data consumers is high and the amount of data traffic varies. On the other hand, message-oriented communication makes it straightforward to re-use existing, possibly standardized, data models. ALCHIMIA created three so called Smart Data Models for the FIWARE ecosystem, grouped into the subject IndustrialProcess[1] in the domain Smart Manufacturing. This covers three message structures generic within process industry and especially batch processes: MaterialAddition, ProcessChemicalAnalysis and ProcessEvent (such as the end of a heat). Related to data, all processing followed the privacy and confidentiality policies of each plant.

For the models referred to in this document, the multiple Kubernetes components (or pods) were necessary. These are summarized in table 1. Both EAF and LF necessitate

[1] https://github.com/smart-data-models/dataModel.IndustrialProcess

multiple components now that the models and user interfaces are separate. Additionally, the EAF model stores its results into a database. For both EAF and LF, the components were installed for the three participating plants.

Tab.1 - Kubernetes pods for Digital Twins for each participating steel plant.

Pod	API	Persistent data
EAF web application	HTTP (user interface)	Configuration
EAF model	None (only client towards others)	None
EAF model results database	SQL	Calculated EAF state; model logs
LF web application	HTTP (user interface)	Configuration; historical data
LF model	HTTP	Configuration
Life Cycle Assessment (LCA) web service	HTTP	Configuration

Decision support for Electric Arc Furnace

The EAF decision support system includes three main views, two of which display results from the DT during online operation. First, a diagram displays the evolution of the calculated steel temperature and meltdown degree as well as the measured steel temperature (see figure 4). The same view shows suggestions when to add scrap baskets and when to end burner operation as well as the remaining electrical energy and oxygen inputs. Second, another view shows the calculated carbon and oxygen content along with the actual measured values. Third, there is a view to browse historical heat information, calculated by the DT, to observe past measures and performance. Additional screenshots were presented in [3].

Fig.4 - One of the views in EAF decision support, currently in endpoint control stage (some data hidden due to the corporate privacy and confidentiality policy).

Additionally, the operator can set control targets for the meltdown degree for the scrap basket addition and burner shutdown, the tapping temperature as well as the target carbon and oxygen content. This makes the DT an active, adaptable optimization tool for the Human in the Loop.
Technically, the DT has been developed in Python. Model consists of multiple submodels that can be parametrized (see [1] for more information). The web application technology is Flask.

Decision support for Ladle Furnace

The LF decision support system comprises four main views for the operators. First, the dashboard shows the current chemical composition, steel weight and tapping temperature. The second shows the initial and final chemical composition (see figure 5). The third view shows the initial and final steel temperature and an overview of the added materials. Finally, the fourth view summarizes a Life Cycle Assessment (LCA) evaluation of the operation. Additional screenshots were presented in [3].

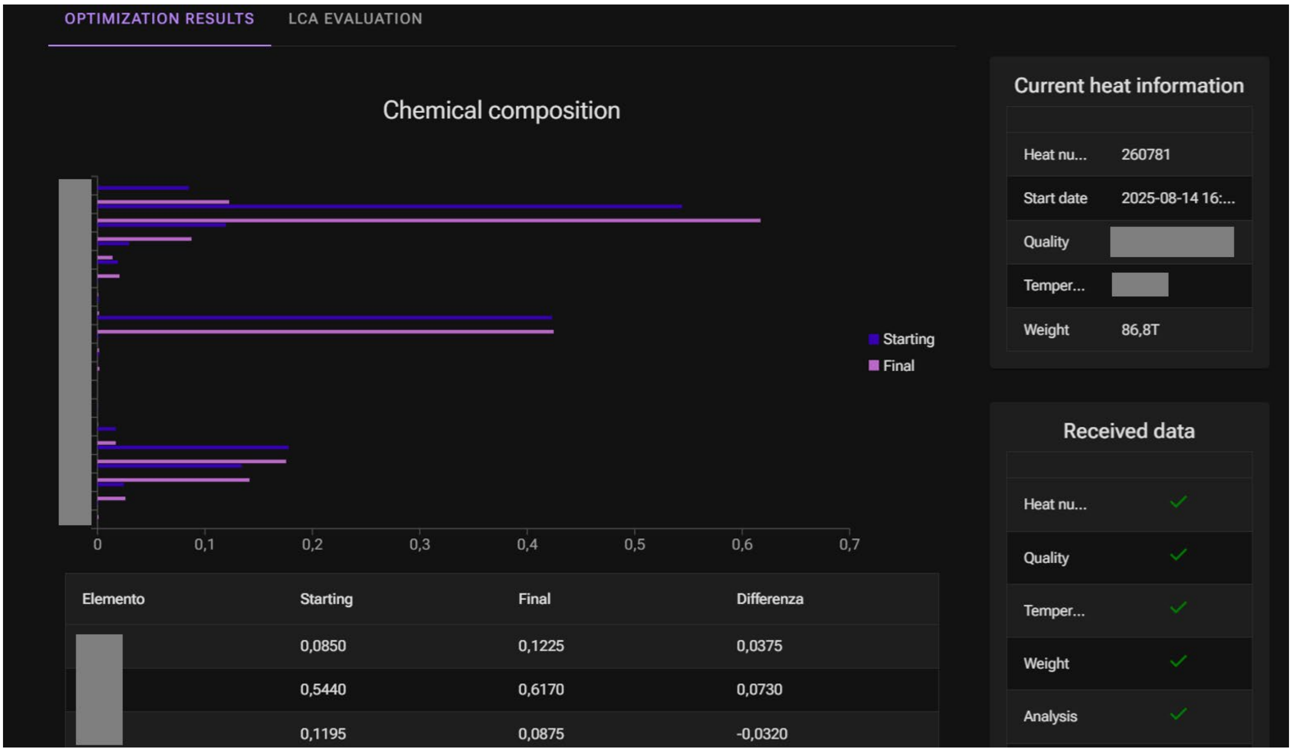

Fig.5 - A part of LF optimization view, depicting the initial and final chemical composition (some data hidden due to the corporate privacy and confidentiality policy).

Besides, the DT enables optimization by letting the operator to define prices for the added materials as well as electrical energy. The material additions include, for instance, carbon, magnesium oxide, calcium oxide and silicon manganese. Besides, the operator can define ferroalloy additions. With these inputs, the DT generates recommendations.
The DT has been developed in Python. The underlying model is a feedforward neural network, the backend technology is FastAPI, and the user interface builds upon Dash.

Federated Learning results

The FL functionality, while integrated with online data for LF, was proven with offline data for both LF and EAF. The FL subject was the temperature model for both as well as the chemical composition model for LF. For each model, FL was applied between multiple plants within the same enterprise. Because the LF model is a neural network, an FL application was straightforward due to the origin in Machine Learning applications. Conversely, the application on EAF was less conventional because this model is physical or analytical (i.e., based on equations), which means that the model lacks an adaptability similar to neural networks. However, FL proved applicable when the plants share knowledge related to model parametrization, which is a novel opening and promising for future research.

In our EAF implementation, the FL architecture comprised two plants and a server that aggregated locally optimized parameters each round using weighted FedAvg. In this approach, the server aggregates the updates from the clients with client-specific weights, so some clients have a greater influence on the global model than others. The parameters were obtained by running a non-linear least squares algorithm. Promising results were reached with two sufficiently high-quality datasets, representing two separate model instances and plants. Due to the asymmetric coverage in a third dataset, it was necessary to exclude this to prevent the Mean Absolute Error (MAE) of one model from growing. This suggests that the FL principle works but depends on data quality, coverage, and symmetry or homogeneity. Concretely, problematic heterogeneity can result from differences in measurement availability, sensor outputs and the data quality.

The LF study involved both temperature and chemical composition models, both including a model instance from three separate plants and a server that aggregated locally trained weights using FedAvg. In both cases, the results showed an improvement in two models instances, the third indicating a similar or a lowered performance. Table 2 shows the results for the temperature model, indicating that plants 1 and 3 improved, whereas plant 2 maintained its R2 score but degraded in Mean Squared Error (MSE). Plant 1 had the smallest dataset, while plant 3 had the lowest data quality and contained more than twice the samples of the other two. Therefore, improvements for these plants were expected, as they benefited from the other plants' data. In contrast, plant 2, which had the strongest training dataset, did not improve due to the imbalance in the datasets caused by plant 3's larger and lower-quality data. Again, it appears that FL performance can suffer from asymmetric or low-quality datasets, but a clear improvement is still possible. However, the federated LF temperature model meaningful enhances the performance for plants 1 and 3 and does not significantly worsen performance for plant 2.

Tab.2 - FL improved two of the three LF temperature models: plant 1 and 3.

Model	Features	Plant 1		Plant 2		Plant 3	
		R2 Score	MSE	R2 Score	MSE	R2 Score	MSE
Local	FinalTemperature	0,278	137,33	0,580	36,65	0,158	59,10
Federated	FinalTemperature	0,446	98,88	0,579	45,56	0,299	32,77
Federated	FinalTemperature 95% confidence interval	(0.363, 0.492)	(85.44, 111.49)	(0.465, 0.655)	(35.21, 58.62)	(0.232, 0.361)	(25.92, 42.04)

Lessons learned

For future efforts, multiple remarks were made during the work. The results are promising, but the proposed framework operated in a complex environment with multiple organizations, which does not come without challenges.

Kubernetes can have a steep learning curve, and if managed remotely, the overhead increases. Some of its mechanisms can surprise due to primitiveness compared to a plain container virtualization environment (such as Docker). This is especially related to the persistent volumes to keep data when a component restarts or is re-instantiated. On the other hand, the importance of security policies increases when the environment is accessed remotely, which adds overhead to both developers and administrators. Still, now that most components ran in one environment, it was straightforward to enforce security policies compared to a network of distributed, heterogeneous infrastructures. In the experiment, the sole external component was the central FL service.

When models receive data in an event-driven fashion, the importance of workflows and functional interoperability increases. Clearly, the models cannot operate properly – if at all – if an essential message arrives too late or drop completely. In any case, in an asynchronous, event-driven world, the models must be robust to operate even in exceptional situations. Additionally, any implications on the model and application state must be considered in case of inconsistencies in earlier data. For instance, when a heat starts and the earlier data have been incomplete, the objective should still be to generate clean results for the new heat.

It was observed that the success of FL directly depends on the availability of models and data from the clients. The data must be sufficient in quality and volume, and there must be multiple plants integrated. If this condition is not met, delays and problems will occur. On the other hand, asymmetry in data coverage can lead to no benefit from FL in a plant. These factors and risks must be considered when planning for the investment. During operation, it is advisable to monitor FL performance continuously in each participating plant.

In summary, two main factors were identified. First, data quality and availability are paramount. Second, the stages of any event-driven workflows must be harmonized early.

SOCIETAL, ECONOMIC AND ENVIRONMENTAL IMPACT

The platform presumably has an effect on the society, economy of the industrial plants, and the environment. These influences are elaborated in the following paragraphs.

Considering the society, industrial enterprises are core actors by providing workplaces, resources and revenue, but AI introduces another ethical dimension. This has been captured within the AI Act [12], considering possible AI-related factors, such as physical human interaction, general-purpose applicability, medical aid and the processing of personal data. Considering the AI Act, ALCHIMIA framework appears to pose no special requirements beyond ensuring the AI literacy of the end users. Still, this can potentially change if ALCHIMIA's is extended in the future.

Furthermore, ALCHIMIA can contribute to the economic performance of industrial enterprises. First, the DTs included help enterprises directly in optimizing their production activities. Second, FL improves the performance of the included DTs. This applies both to the related electric steelmaking use case and ALCHIMIA framework in general, providing re-usable methods and structures for industrial use cases. We expect similar to be reachable in any use case with DTs, especially if multiple production plants are involved with the same FL network.

Finally, ALCHIMIA framework helps heavy industries in reducing their environmental footprint. Although economic efforts often align with factors, such as energy and resource savings, LCA can be integrated for explicit environmental objectives. This is already available for the LF model, and a similar tool could be integrated into the EAF model. With more extensive LCA efforts, ALCHIMIA could host an entire environmental toolkit, providing online decision support.

DISCUSSION

The suggested DT framework brings novelty especially related to FL. It is among the first ones to suggest an FL framework for industrial plants. As far as is known, there have been no earlier research to either create an FL framework in metallurgy or to parametrize a physical or analytical optimization model.

The framework is significant as it applies AI for the competitiveness of European steel industry, helping to maintain jobs and the domestic production of critical resources. It improves the applicability of AI technologies in the industrial scope even when the data sources are heterogeneous yet share the same principles. Additionally, this work developed AI applications towards the green transformation, a key overall goal in modern industry.

Certain limitations exist. Only a few plants were included in the experiment, and this numerically evaluates FL improvements only with offline data (although the FL system was integrated with online data and models). Besides, it is limited to upstream processes in electric steelmaking. There could be a wider study regarding the benefits of the DT and FL framework, including more enterprises and production processes.

For the future work, multiple opportunities remain. The advantages of FL could be explored further, and there could be more focus on Continual Learning to react to any degradations in model performance. The concept could be expanded to larger enterprise networks, and more of steel production processes could be included. The advantages of the concept could be explored in other metallurgical industries as well as process industries in general (such as chemical or pulp and paper) and manufacturing. On the other hand, as the software solutions come from multiple developers and even separate organizations, certain challenges arise. Thus, the advantages of the Data Mesh [13] could be researched now that these can enforce common interoperability and security policies as well as reduce the danger of bottlenecks in data integration tasks. In larger multi-actor business ecosystems, there could be a coordinating non-profit entity to decide on the rules, as earlier suggested for industrial partner networks [14]. Finally, as the current system focuses on models to optimize single unit processes, there could be FL for the plant-wide optimization problems [15].

CONCLUSIONS

This document introduces a DT framework that operates optimization models, including AI, and exploits enterprise-wide knowledge with FL. The design of the framework is based on industrial requirements, and it has been proven with actual data. The results suggest that the concept is applicable in industrial environments.

The results have indicated that the DT platform, along with FL, can improve the performance of production processes in electric steelmaking. FL has proven to be effective even in the industrial context as demonstrated for both EAF and LF. Although the reported experiments are restricted to a certain steelmaking route, we can expect similar benefits in other industrial FL applications.

A clear challenge in industrial FL is the data heterogeneity between plants. Even the plants that operate similar processes can vary considerably in the availability of the measurements. This not only adds difficulty but also effectively prevents certain application areas, such as scrap characterization now that scrap suppliers and qualities are fundamentally different between plants. Additionally, as noticed with EAF, an inferior data quality or coverage can hamper benefits at least in a single plant.

Considering the research question presented in the introduction, we can summarize that DTs and FL in data-driven industrial use cases necessitate special attention on the local data infrastructure in plants. A scalable environment should include services for data storage, data delivery and messaging, and container orchestration to be used in co-operation between the component developers. It is essential to preprocess data to fit for the network-wide requirements, or otherwise the plant cannot benefit from FL. In any case, only the local infrastructure can guarantee FL suitability for the plant.

ACKNOWLEDGEMENTS

This work has been supported by the project "Data and decentralized Artificial intelligence for a competitive and green European metallurgy industry" (ALCHIMIA), which has received funding from the European Union's Horizon Europe research and innovation programme under grant agreement No. 101070046. Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union. Neither the European Union nor the granting authority can be held responsible for them.

Additionally, the authors want to thank all project partners for their cooperation, particularly mentioning EVIDEN colleagues Diego Esteban for project coordination, Vera Peiss from BFI for various software development contributions, and EXUS' colleagues Giorgos Nomikos, James Lloyd, and Marcos Varveris for setting up components for data integration and preprocessing.

REFERENCES

[1] B. Kleimt, W. Krieger, D. Mier Vasallo, A. Arteaga Ayarza, I. Unamuno Iriondo, "Model-Based Decision Support System for Electric Arc Furnace (EAF) Online Monitoring and Control," Metals, vol. 13, no. 8, p. 1332, 2023. https://doi.org/10.3390/met13081332

[2] J. Wen, Z. Zhang, Y. Lan, Z. Cui, J. Cai, W. Zhang, "A survey on federated learning: challenges and applications," International Journal of Machine Learning and Cybernetics, Bd. 14, pp. 513-535, 2023. https://doi.org/10.1007/s13042-022-01647-y

[3] P. Kannisto, B. Kleimt, A. Chandgude, N. García Crespo, I. Guardiola Luna, S. Dettori, V. Colla, B. Fernandez, "Framework with Digital Twins and Federated Learning for Decision Support in Multi-plant Schemes for Electric Steelmaking and Beyond," in DTMI 2025 - International Workshop Carlo Longaretti on Digital Tools for the Metallurgical Industry, Bergamo, Italy, 2025.

[4] A. R. Hevner, "A Three Cycle View of Design Science Research," Scandinavian Journal of Information Systems, vol. 19, no. 2, pp. 87-92, 2007. https://aisel.aisnet.org/sjis/vol19/iss2/4

[5] A. Salis, A. Marguglio, G. De Luca, S. Razzetti, W. Quadrini and S. Gusmeroli, "An Edge-Cloud based Reference Architecture to support cognitive solutions in Process Industry," Procedia Computer Science, vol. 217, pp. 20-30, 2023. https://doi.org/10.1016/j.procs.2022.12.198.

[6] S. T. Johansen, P. Unal, Ö. Albayrak, E. Ikonen, K. J. Linnestad, J. Sudi, A. K. Srivastava, B. T. Løvfall, "Hybrid and cognitive digital twins for the process industry," Open Engineering, vol. 13, no. 1, p. 2022-0418, 2023. doi: 10.1515/eng-2022-0418

[7] P. Kannisto, Z. Kargar, G. Alvarez, B. Kleimt, A. Arteaga, "Resilient, Adaptive Industrial Self-X AI Pipeline with External AI Services: A Case Study on Electric Steelmaking," Processes, vol. 12, no. 12, p. 2877, 2024. https://doi.org/10.3390/pr12122877

[8] T. Hiessl, D. Schall, J. Kemnitz and S. Schulte, "Industrial Federated Learning - Requirements and System Design," in Highlights in Practical Applications of Agents, Multi-Agent Systems, and Trust-worthiness. The PAAMS Collection, 2020. https://doi.org/10.1007/978-3-030-51999-5_4

[9] R. Kanagavelu, Z. Li, J. Samsudin, S. Hussain, F. Yang, Y. Yang, R. S. M. Goh, M. Cheah, "Federated Learning for Advanced Manufacturing Based on Industrial IoT Data Analytics," in Implementing Industry 4.0: The Model Factory as the Key Enabler for the Future of Manufacturing, Springer International Publishing, vol. 202, pp. 143-176, 2021. https://doi.org/10.1007/978-3-030-67270-6_6

[10] T. Z. Sana, S. Abdulla, A. Nag, A. Das, M. M. Hassan, Z. Z. Fiza, A. Karim, S. R. R. Kabir, "Advancing Federated Learning: A Systematic Literature Review of Methods, Challenges, and Applications," IEEE Access, vol. 13, pp. 153817-153844, 2025. doi: 10.1109/ACCESS.2025.3605165

[11] S. Moreschini, F. Lomio, D. Hästbacka, D. Taibi, "MLOps for evolvable AI intensive software systems," in 2022 IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER), Honolulu, HI, USA, pp. 1293-1294, 2022. doi: 10.1109/SANER53432.2022.00155

[12] "Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations", 2024.

[13] Z. Dehghani, Data Mesh: Delivering Data-Driven Value at Scale, O'Reilly Media, Inc., 2022.

[14] P. Kannisto, D. Hästbacka, A. Marttinen, "Information Exchange Architecture for Collaborative Industrial Ecosystem," Information Systems Frontiers, vol. 22, no. 3, pp. 655-670, 2020. https://doi.org/10.1007/s10796-018-9877-0

[15] P. Kannisto, D. Hästbacka, T. Gutiérrez, O. Suominen, M. Vilkko, P. Craamer, "Plant-wide interoperability and decoupled, data-driven process control with message bus communication," Journal of Industrial Information Integration, vol. 26, p. 100253, 2022. https://doi.org/10.1016/j.jii.2021.100253