Available online at www.sciencedirect.com

ScienceDirect

Procedia CIRP 142 (2026) 85–90

www.elsevier.com/locate/procedia

36th CIRP Design Conference (CIRP Design 2026)

A Practical Case Study on Homomorphic Encryption for Secure and Sustainable Product Development in Sheet Metal Design

Simon Jess[a,b,*], Lukas Ihrig[b], Artur Krause[a,b], Claudius Messerschmidt[b], Felix Förster[a], Nikola Bursac[a]

[a]*ISEM, Hamburg University of Technology, 21073 Hamburg, Germany*
[b]*TRUMPF Werkzeugmaschinen SE & Co. KG, 71254 Ditzingen, Germany*

* Corresponding author. *E-mail address:* simon.jess@isem-tuhh.de

Abstract

Modern product development increasingly relies on data-driven approaches to make informed decisions. Especially within the framework of System Generation Engineering (SGE), where new products are based on reference systems, and thereby a valuable knowledge base is generated. Preserving the sovereignty of this data, particularly in the context of product designs, is crucial. Simultaneously, there is a strong desire to leverage this data for data-driven services that provide valuable predictions, such as production costs, processing times, production risk, or the Product Carbon Footprint (PCF). To resolve this conflict of interest, encryption techniques offer a solution. This study investigates the practical applicability of homomorphic encryption (HE): Using PCF prediction as a case study, it demonstrates that machine learning models can forecast such values. Subsequently, it examines whether this is also feasible during inference with homomorphically encrypted data. To this end, a neural network trained on real-world industrial data is compared with a model that performs computations on encrypted data. The results demonstrate that prediction quality decreases only slightly when using HE, confirming its suitability for preserving information quality. However, this advantage comes with a significant performance overhead: computation times for the encrypted calculations increase substantially, posing a challenge for practical application, especially in time-critical scenarios. The paper concludes that HE is a promising approach to reconcile data sovereignty with data-driven innovation. However, its broad industrial application crucially depends on future advancements in technology's computational efficiency and scalability.

Keywords: Machine Learning, Data-Driven Design, Product Carbon Footprint, Homomorphic Encryption, Data Sovereignty

1. Introduction

In modern product development, the ability to leverage data is becoming a decisive competitive factor. Yet the stakeholders who need analytics and those who own the most valuable data are often not the same. Large machine tool manufacturers and service providers have the resources to build sophisticated AI solutions, but they frequently lack the large-scale, domain-specific datasets of their customers required to train reliable models. On the one hand, small and medium-sized enterprises such as contract manufacturers generate rich, high-quality data from daily operations, but they often lack the expertise and infrastructure to apply advanced machine learning methods.

This imbalance creates a paradox in engineering design: the actors with the strongest capabilities in developing data-driven services depend on the willingness of others to share sensitive design and production data. For the data owners, however, disclosing such information would mean revealing core know-how and competitive advantages. The result is a deadlock: sustainable design services such as Product Carbon Footprint (PCF) prediction tools cannot unfold their full potential, even though both sides would benefit from their application.

System Generation Engineering (SGE) emphasizes the value of knowledge carried forward from one product

10.1016/j.procir.2026.05.228

generation to the next. If this knowledge could be combined with privacy-preserving data analytics, new opportunities would arise for reconciling innovation with confidentiality. Homomorphic Encryption (HE) offers precisely this promise. By enabling computations directly on encrypted data, HE allows machine learning models to deliver predictions without exposing the underlying datasets.
This paper explores the feasibility of applying HE to the prediction of PCF values in sheet metal design. It examines whether data sovereignty can be preserved while still providing designers and engineers with reliable, data-driven decision support for sustainability in early product development.

In current product development, systems are developed in generations according to the model of the SGE. Systems are improved from generation to generation based on the knowledge and experience gained from usage. Knowledge and data play an increasingly crucial role in this process, particularly with regard to the use of data-driven tools to assist in the early stages of product development with data-driven decisions [1]. One such tool is the PCF calculator for sheet metal parts, which was developed by Krause et al. and is being investigated in the context of sheet metal design [2].

As Krause et al. demonstrate, using a sheet metal part based on geometry can reduce the required PCF value from generation to generation. This reduces both the part in production and the CO_2e required by the material and thus contributes to sustainability in sheet metal construction [2].

There are various ways to develop and provide such data-driven solutions. One option is to use machine learning, which is particularly useful for handling complexity when there are many influencing factors and parameters to consider. However, training machine learning models requires a large amount of training data. In the case of the PCF value of sheet metal parts, this training data comprises two-dimensional sheet metal parts intended for processing. Since this data is often unavailable or not shared with experts such as machine manufacturers or data scientists, because it contains the developer's knowledge and expertise, alternative methods are increasingly necessary [3]. In the context of machine learning, a distinction is usually made between feature extraction, reducing the information content to non-critical data, and the application of encryption to the data.

2. Theoretical Background

2.1. Machine Learning in Product Development

This study uses a supervised learning approach to predict the PCF value. Supervised learning is often divided into two types: cluster prediction and regression, which is the prediction of a specific value. In the study, the PCF is predicted using regression. Regressions are based on predicting a target variable based on input data and the provided target value, the so-called label [4]. By sufficiently abstracting the relationships between the input values, the function can be adjusted to predict the target variable represented by the input data and the label. The result is a function shown in (1) that represents the relationship between the input data (e.g., a vector x) and the probability distribution for the target (e.g., a value or vector y) [4]:

$$f(x) = \mathbb{E}[y \mid x] \tag{1}$$

In neural networks, this function is represented by so-called layers. These consist of nodes and weights of the nodes as well as the connections between the nodes via edges [4]. These weights and edges are adjusted using backpropagation [5]. In this process, the actual expected value, the label, and the prediction are compared after the target value has been predicted. If there is a deviation between the prediction and the label, the adjustment is carried out backwards in the neural network [6].

2.2. Fully homomorphic encryption for privacy preserving machine learning

The basis of fully homomorphic encryption (FHE) for privacy-preserving machine learning is based on the classic encryption concept. This concept involves encrypting the raw data into what is known as ciphertext. The goal is to convert the raw data into unreadable, encrypted data using an encryption algorithm. To ensure that the ciphertext can be decrypted, the symmetric encryption method can be used. This allows an encrypted character string to be decrypted back into a readable character string using the key used for encryption [7].

Traditional encryption schemes, such as AES or RSA, protect data only during storage or transmission (data at rest and data in transit). To perform computations, such as the weighted sums and activation functions in a neural network, data encrypted with these standard schemes must be decrypted first. This exposes sensitive information in the plaintext domain during processing, creating a security vulnerability. In contrast, Homomorphic Encryption allows mathematical operations (specifically addition and multiplication) to be performed directly on the ciphertext. This ensures that the data remains encrypted throughout the entire processing pipeline, preserving data sovereignty even during active computation.

The FHE approach is currently implemented primarily using two different algorithms. One approach is the Brakerski-Fan-Vercauteren scheme (BFV), another is the Cheon-Kim-Kim-Song scheme (CKKS). The CKKS scheme offers the advantage of also enabling full addition, subtraction and multiplication with floating point numbers in the calculations [8]. Since most machine learning algorithms are based on floating point operations, FHE is suitable for use in machine learning because it enables full addition, subtraction and multiplication despite encryption [9]. Homomorphic encryption, specifically FHE, in the context of privacy-preserving machine learning [8]. To train a model in the context of FHE, a model based on the ciphertext is trained. This means that prediction and training are performed with encrypted data, thus mapping the function for the relationship between the input as ciphertext and the output as ciphertext. This can then be decrypted back into a readable and interpretable character string using the appropriate key and the symmetric encryption used [3,7].

2.3. SGE – System Generation Engineering

The concept of SGE reflects the prevailing practice of product development in engineering domains. It characterizes the development of technical systems as a generational progression, in which each new system iteration is derived from at least one reference system—typically the immediate predecessor, but potentially also including competitor products or various technical solutions [10]. Within the SGE, product development is understood as a process of systematic variation and adaptation. Proven structural and functional elements from reference systems are selectively retained, modified, or replaced to realize the subsequent system generation. These modifications are classified into three categories: carryover variation (CV), where elements are adopted without change; attribute variation (AV), involving parametric modifications; and principle variation (PV), where fundamental design principles are redefined [11]. The reliance on reference systems enables the integration of empirical data generated during the operation of previous system generations. Wagenmann et al., machine data collected from tools operating in real customer environments can be systematically analyzed to gain a deeper understanding of system performance under actual usage conditions. This real-world evidence allows development decisions to be based on verified application behavior rather than assumptions or over-engineered specifications. As a result, unnecessary system attributes can be eliminated, and product portfolios streamlined without compromising functional value [12].

2.4. Data-driven Sustainability in Product Development

The environmental impact of a product is essentially determined during the development process. Appropriate decisions, similar to design decisions regarding individual machine functionality, must be made at an early stage of product development. Aiming to enhance the sustainability of mechatronic systems during operation, Krause et. al. identified several influencing attributes that can be utilized to e.g. guide sheet-metal design practice by focusing on modification of the attributes of a part geometry. Modifying e.g. the existing contour of a part geometry, shapes can be merged and re-designed to ensure technical functionalities are preserved but the necessary resource consumption for its manufacturing is reduced [2]. This can be achieved by minimizing e.g. the total contour length of the part geometry, resulting in shorter tool paths that must be machined and reduced duration of the manufacturing processes [13]. Although the analysis of operational data and part geometries can provide significant potential for optimization, gathering and analyzing such data remains challenging due to confidentiality concerns, as manufacturing processes and part geometries are often proprietary and confidential. Therefore, new approaches are required that enable the secure analysis of sensitive data—ensuring both customer confidentiality and data protection—while still providing the expected added value.

3. Objectives and Research Methodology

The development of data-driven services that take data protection and the protection of sensitive data into account requires the use of appropriate technical solutions. This study examines an approach to implementing privacy-preserving machine learning for predicting the PCF of a sheet metal part. The present case study shows that data sovereignty for sensitive data can be ensured even when considering data-driven solutions for data-driven decision-making in the SGE.

The methodological approach for the paper is based on the design science research process model according to Hevner et al. [14]. This provides systematic support in an iterative process which, in combination with the CRISP-DM framework, is used in the iterative development and investigation of artifacts such as the trained models. Three research questions (RQ) are used as guiding questions in the investigation.

RQ1: To what extent is the use of geometry-related data sufficient for making an accurate prediction of the PCF using neural networks?

RQ2: What are the conditions to train neural networks for PCF prediction on encrypted data, and how does the use of FHE affect model accuracy and training effort?

RQ3: What are the technical, algorithmic and performance-related limitations of combining fully homomorphic encryption with neural networks for PCF predictions?

At the heart of Hevner's DSR is the design cycle, which is linked to the CRISP-DM approach. This results in the following steps: business understanding, data understanding and preparation, followed modeling, evaluation and deployment [15]. The business understanding phase illustrates the relevance cycle of DSR, i.e. the necessity arising from the environment. Research questions are employed to identify problems, develop solutions and validate them within the iterative design cycle. The first step is to understand the data, which involves analyzing the available data to determine which geometric information is necessary for predicting the PCF. This can then be used as a benchmark. Solution development, model creation and training then use an unencrypted neural network and readable data, as well as a model in which the data is encrypted using FHE prior to training. To evaluate the results, the differences between the models and training processes are listed, and the results are measured against these to identify areas for improvement in future iteration cycles. The artefact developed and presented in this case study contributes to the rigor cycle for the knowledge base through the knowledge gained in the form of limitations and an initial practical approach to implementing privacy-preserving machine learning.

4. Parameters with influence on the prediction of a PCF and need of machine learning model

Few parameters are required to determine the PCF value of a geometric sheet metal part [2]. However, if this geometry-related data is unavailable, cannot be made available for data protection reasons or would compromise the owner's data

sovereignty, it is almost impossible to determine the PCF of a sheet metal part. For this reason, this study investigates the approach of a machine learning model to identify the necessary parameters and predict the PCF value of a sheet metal part based on its geometric properties.

The following approach was investigated for the study. A machine learning model with a supervised learning approach was developed. This is based on labelled data, where the data label is the PCF value. For this purpose, a dataset of 9,968 sheet metal panels with different geometries was used, also known as nesting. The PCF was then calculated for each individual part to have it available as a target value for the supervised learning approach, and the value for the entire sheet metal panel was summed based on this.

The first step in identifying the necessary model parameters is to use the SHapley Additive exPlanation (SHAP) methodology to determine the relevant input parameters. Figure 1 shows the SHAP values for the features of a simple neural network used for predicting the PCF with a supervised learning approach.

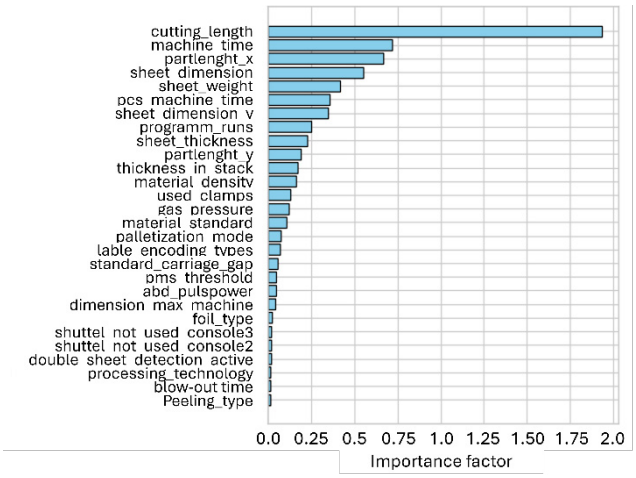

Fig. 1. SHAP-Scores for PCF prediction with supervised learning

The SHAP score represents a metric for the relevance of a parameter to the prediction and clearly shows that parameters such as 'cuttingLength', 'machineTime', 'xLength', 'sheetDimensionX', 'sheetWeight', 'sheetDimensionY' and 'thinkness' play a relevant role in mapping the relationship between the PCF and the input data.

This confirms the parameters used to calculate the PCF and shows that a correlation can be established between the relevant parameters and the calculated target of the PCF. This is visible in the SHAP value for each parameter shown in Figure 1. The higher the SHAP value, the more relevant the parameter is in relation to the nodes and weights that contribute to the decision. Figure 1 also shows other parameters that have no influence on the calculation but show a slight correlation with the prediction.

5. PCF prediction with machine learning and fully homomorphic encrypted machine learning

The training design is the same for both neural networks, with and without encrypted data, and is as follows: Both models are trained with the same data for 100 epochs each with early stopping, if the loss has reached a plateau. The architecture of both models is also structured identically for comparability. It is a neural network with an input layer of 256, the first hidden layer with 128 neurons, the second hidden layer with 64 neurons, a third hidden layer with 32 neurons, and an output layer with 1 neuron for a linear output for the regression prediction of the PCF. Cross-validation is not used to validate or test the models. The advantage of this is that neither the test nor the validation data are used at any point during training. The chosen split is 70% training data, 15% test data and 15% validation data. With the 9,968 pieces of geometry data used, this equates to 6,978 pieces of geometry data for training, and 1,495 pieces each for testing and validation.

In this study, accuracy refers to the percentage of predictions falling within a defined tolerance threshold of the true value. To provide a rigorous performance evaluation for the regression task, standard metrics were calculated alongside this accuracy. The unencrypted baseline model achieved a Mean Squared Error (MSE) of 0.035 and a Mean Absolute Error (MAE) of 0.149, resulting in a threshold accuracy of 80.3%. The model trained on homomorphically encrypted data achieved comparable results, recording an MSE of 0.042 and an MAE of 0.162, with an accuracy of 78.1%. These results demonstrate that the transition to the encrypted domain resulted in only a marginal loss of predictive power. An early stopping within 100 epochs was not reached.

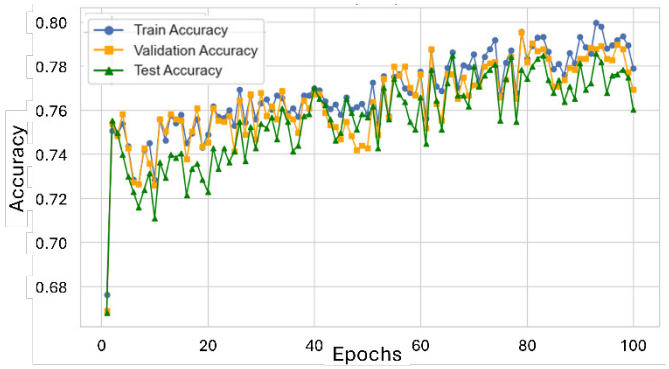

Fig. 2. Accuracy of the training over training time of a neural network for PCF prediction without FHE

Given the loss curve in Figure 3 it does not reach a sufficient plateau for early stopping. Nevertheless, the loss curve shows that it is approaching a plateau and not converging towards zero. This ensures that, despite the significant fluctuations in accuracy shown in Figure 2, there is no overfitting and the model generalizes sufficiently. The stability of the training process is also evident in that, although there are fluctuations in accuracy, these are not the result of varying loss, but can be assumed to be due to the limited availability of training data. This once again clearly shows that the assumption from Chapter 4 can be covered by analyzing the relevant parameters using XAI methodologies for predicting the PCF value with a neural network. Based on this, it is assumed that the results achieved with the unencrypted training and model are sufficient for an initial comparison with a model trained on encrypted data.

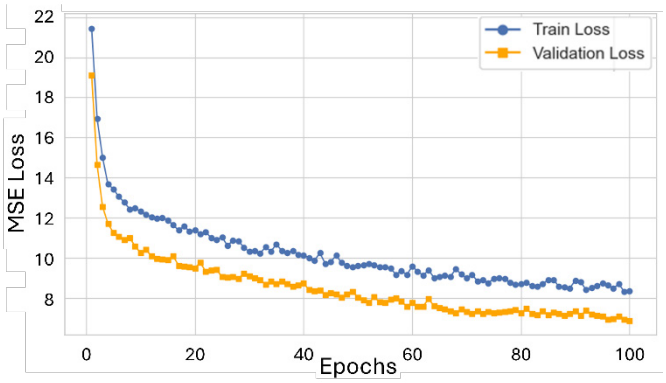

Fig. 3. Training and validation loss over training time of the neural network for PCF prediction without FHE

As mentioned at the beginning, the same structure and training procedure are used to implement the FHE approach for privacy-preserving machine learning. The TenSEAL library is used to implement CKKS. This means that the model itself remains the same, only the input data differs in that it is no longer raw data but encrypted ciphertext.

Figure 4 shows the training accuracy curve for the model trained with encrypted data as input. After 81 epochs the early stopping was triggered, since the loss curve has stagnated. Within 81 epochs, the training and thus the model achieved an accuracy of 78.96%.

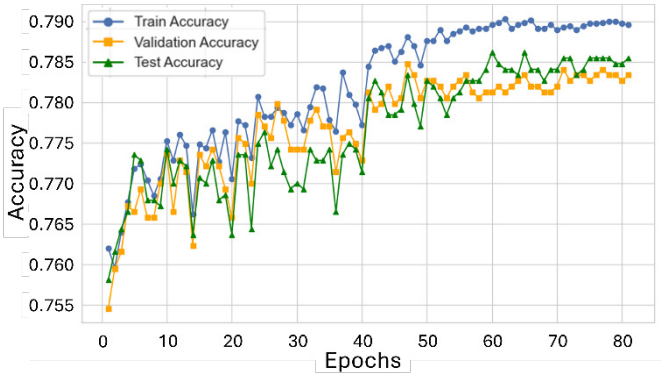

Fig. 4. Accuracy of the training over training time of a neural network for PCF prediction with FHE

Furthermore, compared to training with unencrypted data, it becomes clear that not only does the loss curve (see Figure 5) converge so strongly that premature termination is triggered after 81 epochs, but also that the accuracy of the prediction for the encrypted data is not as unstable and scattered. Both loss curves show a continuously decreasing loss, suggesting that both the model with unencrypted data and the model with encrypted data are able to predict the PCF value based on the given relevant information about the geometry using regression.

The challenge of implementing training based on encrypted data lies in backpropagation within neural networks to account for the discrepancy between the predicted PCF value and the label when adjusting the neural network. CKKS allows full addition, subtraction and multiplication, but not division. For this reason, the model is also trained using encrypted input data in this case study. After prediction, backpropagation (i.e. adjustment of the weights with unencrypted data) is performed using the key known for the encryption. This is due to the activation function used in the neurons, which was implemented with a ReLU function and requires the function to be differentiated during backpropagation.

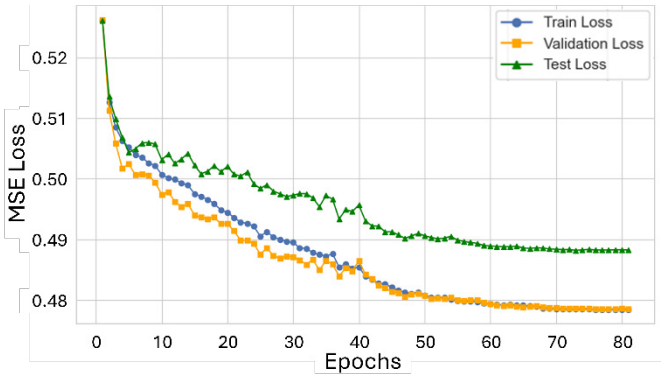

Fig. 5. Training and validation loss over training time of the neural network for PCF prediction with FHE

Regarding computational overhead, the experiments showed that training time increased by a factor of two, rising from approximately 5 minutes for the plaintext model to 10 minutes for the encrypted model. For the practical application in a design tool, the inference latency is critical. The study measured an increase in inference time per sample from 0.05 seconds (plaintext) to 0.11 seconds (encrypted). Although this represents a statistical doubling, an inference time of 0.11 seconds remains imperceptible to the human user and is fully acceptable for real-time quotation and design checks in an industrial workflow. Training is carried out on the same hardware, with the time required to encrypt and decrypt the data not considered.

Training shows that both models can predict the PCF value. While the models are not yet ready for application at this stage of development, they already demonstrate their potential, and the challenges involved. Notably, the minimal deviation indicates that a model operating with encrypted data is equally promising as one operating with plain text geometry information.

6. Discussion

The study demonstrates that FHE enables privacy-preserving prediction of PCF with only minor reductions in accuracy compared to conventional models. This is highly relevant for engineering design, where early-phase decisions shape sustainability, cost, and manufacturability. Designers and engineers could integrate sustainability considerations into their work without exposing proprietary geometric or process data, thereby reconciling confidentiality with innovation.

FHE opens the door to a wide range of design-support services. Beyond PCF prediction, secure quotation and cost-estimation services could provide reliable forecasts on encrypted data without revealing manufacturing know-how. Manufacturability checks on encrypted CAD models could

deliver feedback on tooling feasibility or process times while preserving confidentiality. Risk assessment services could be built from aggregated encrypted datasets, highlighting probabilistic failure modes or quality deviations across small and medium-sized enterprises. Lifecycle-oriented services, such as predictive maintenance planning or recycling recommendations, could be enabled by encrypted usage data, extending privacy-preserving analytics into later product lifecycle phases.

Within the framework of SGE, these services strengthen the systematic reuse of knowledge from earlier product generations. Encrypted operational data can refine design rules, sustainability targets, and performance benchmarks without forcing SMEs to disclose sensitive information. This reinforces generational improvement and expands the scope of data-driven support across design, manufacturing, and operation.

Despite these promising implications, several constraints remain. First, the predictive accuracy of around 80% was limited by dataset size and model simplification for explainability. Second, a critical methodological constraint identified in this study concerns the scope of privacy preservation. While the inference phase (application of the model) can be performed fully on encrypted data, allowing a client to receive PCF predictions without revealing their geometry, the training phase currently faces limitations. Due to the complexity of computing non-linear derivatives required for backpropagation on ciphertext, the weight updates in this case study were performed using decrypted gradients. Consequently, the current approach assumes a trusted environment for the training process, whereas the subsequent application of the service is trustless.

For the design community, these findings suggest that FHE should be viewed as an enabling technology rather than a standalone solution. Hybrid approaches that combine FHE with federated learning or secure multi-party computation should be investigated to balance privacy and performance. Integration into established design environments, CAD, PLM, or generative design tools, will be crucial for adoption. Further empirical studies with larger, more diverse datasets are needed to validate generalizability and industrial robustness.

7. Conclusion and Outlook

This study illustrates that FHE can become a key enabler of secure and sustainable product development. By enabling computations on encrypted data, it allows design teams to access advanced analytics while maintaining strict data sovereignty. This creates opportunities for sustainability, cost prediction, manufacturability checks, and risk assessment to be embedded in early design phases.

At the same time, significant computational and methodological challenges remain. Until performance improves and full end-to-end encryption becomes feasible, applications will be limited to contexts where time is not critical. Nonetheless, the results demonstrate the feasibility of privacy-preserving services in engineering design and highlight the potential of FHE to bridge the gap between data-rich small and medium-sized enterprises and data-driven service providers.

For the Design Community, the implications are clear: privacy-preserving technologies are not only a cryptographic challenge but a design research frontier. They have the potential to reshape how product development ecosystems collaborate, ensuring that innovation, sustainability, and confidentiality can coexist.

References

[1] Wagenmann S, Bursac N, Rapp S, et al. Success Factors for the Validation of Requirements for New Product Generations – A Case Study on Using Field Gathered Data. Proc Des Soc 2022;2:1805–14. https://doi.org/10.1017/pds.2022.183.

[2] Krause A, Dielhenn J, Jess S, et al. Sustainable Sheet-Metal Design: Employing the Product Carbon Footprint as Support for Engineers in Developing New Product Generations. Procedia CIRP 2025;136:510–5. https://doi.org/10.1016/j.procir.2025.08.088.

[3] Benaissa A, Retiat B, Cebere B, et al. TenSEAL: A Library for Encrypted Tensor Operations Using Homomorphic Encryption 2021. https://doi.org/10.48550/arXiv.2104.03152.

[4] Goodfellow I, Bengio Y, Courville A. Deep Learning. MIT Press; 2016.

[5] Lecun Y, Bottou L, Bengio Y, et al. Gradient-based learning applied to document recognition. Proc IEEE 1998;86:2278–324. https://doi.org/10.1109/5.726791.

[6] Rumelhart DE, Hintont GE, Williams RJ. Learning representations by back-propagating errors. Nature 1986;323:533–6. https://doi.org/10.1038/323533a0.

[7] Gentry C. Fully homomorphic encryption using ideal lattices. Proc. Forty-First Annu. ACM Symp. Theory Comput., Bethesda MD USA: ACM; 2009, p. 169–78. https://doi.org/10.1145/1536414.1536440.

[8] Cheon JH, Kim A, Kim M, et al. Homomorphic Encryption for Arithmetic of Approximate Numbers. In: Takagi T, Peyrin T, editors. Adv. Cryptol. – ASIACRYPT 2017, vol. 10624, Cham: Springer International Publishing; 2017, p. 409–37. https://doi.org/10.1007/978-3-319-70694-8_15.

[9] Acar A, Aksu H, Uluagac AS, et al. A Survey on Homomorphic Encryption Schemes: Theory and Implementation. ACM Comput Surv 2019;51:1–35. https://doi.org/10.1145/3214303.

[10] Albers A, Rapp S, Spadinger M, et al. The reference system in the model of PGE: proposing a generalized description of reference products and their interrelations. Proc 22nd Int Conf Eng Des ICED19 2019;1:1693–702. https://doi.org/10.1017/dsi.2019.175.

[11] Pfaff, Felix, Götz, Gregor Theodor, Rapp, Simon, et al. Evolutionary perspective on system generation engineering by the example of the iPhone. Proc Int Conf Eng Des ICED23 2023. https://doi.org/10.1017/pds.2023.172.

[12] Wagenmann, A. Krause, S. Rapp, et al. Process Model for the Data-driven Identification of Machine Function Usage for the Reduction of Machine Variants. 2022 IEEE Int. Conf. Ind. Eng. Eng. Manag. IEEM, 2022, p. 0444–51. https://doi.org/10.1109/IEEM55944.2022.9989909.

[13] Krause A, Dannerbauer T, Wagenmann S, et al. Enhancing efficiency and environmental performance of laser-cutting machine tools: An explainable machine learning approach. 14.05.2024, Portugal: 2024. https://doi.org/10.1016/j.procir.2024.10.299.

[14] Hevner AR, March ST, Park J, et al. Design Science in Information Systems Research. MIS Q 2004;28:75–105. https://doi.org/10.2307/25148625.

[15] Chapman P, Clinton J, Kerber R, et al. CRISP-DM 1.0: Step-by-step data mining guide. Springer 1999.