Proceedings of the ASME 2024
International Mechanical Engineering Congress and Exposition
IMECE2024
November 17-21, 2024, Portland, Oregon

IMECE2024-147406

PRIVACY-PRESERVING IN-SITU MONITORING IN ADDITIVE MANUFACTURING THROUGH HYPERDIMENSIONAL COMPUTING

Fardin Jalil Piran[1], Prathyush P. Poduval[2], Hamza Errahmouni Barkam[2], Mohsen Imani[2], Farhad Imani[1,*]

[1]School of Mechanical, Aerospace, and Manufacturing Engineering, University of Connecticut, Storrs, CT
[2]Department of Computer Science, University of California Irvine, Irvine, CA

ABSTRACT

Machine learning (ML) models are used for in-situ monitoring in additive manufacturing (AM) for defect detection. However, sensitive information stored in ML models, such as part designs, is at risk of data leakage due to unauthorized access. To address this, differential privacy (DP) introduces noise into ML, outperforming cryptography, which is slow, and data anonymization, which does not guarantee privacy. While DP enhances privacy, it reduces the precision of defect detection. This paper proposes combining DP with Hyperdimensional Computing (HDC), a brain-inspired model that memorizes training sample information in a large hyperspace, to optimize real-time monitoring in AM while protecting privacy. Adding DP noise to the HDC model protects sensitive information without compromising defect detection accuracy. Our studies demonstrate the effectiveness of this approach in monitoring anomalies, such as overhangs, using high-speed melt pool data analysis. With a privacy budget set at 1, our model achieved an F-score of 94.30%, surpassing traditional models like ResNet50, DenseNet201, EfficientNet B2, and AlexNet, which have performance up to 66%. Thus, the intersection of DP and HDC promises accurate defect detection and protection of sensitive information in AM. The proposed method can also be extended to other AM processes, such as fused filament fabrication.

Keywords: Privacy-Preserving Model, Real-time Monitoring, Additive Manufacturing, Hyperdimensional Computing

1. INTRODUCTION

Additive Manufacturing (AM) fabricates three-dimensional objects by layering materials as per digital blueprints, which allows for the creation of intricate designs that traditional manufacturing cannot easily achieve [1]. Despite these benefits, AM encounters significant obstacles, such as slow production speeds, higher costs, and issues with part quality that lead to defects causing approximately 10% of part failures [2, 3]. Typical defects include porosity, layer misalignment, foreign particle inclusion, delamination, and incomplete fusion, all of which negatively impact the dimensional accuracy and mechanical strength of the parts produced [4]. Consequently, many AM-manufactured parts require additional finishing to meet surface quality standards. Moreover, AM processes are often lengthy, taking anywhere from hours to days, and the cost of equipment and materials is high compared to traditional methods. As a result, in-process defect detection is required to improve economic and time efficiencies by enabling immediate adjustments to process parameters or halting defective production runs [5]. Various sensing technologies are utilized for defect detection, including ultrasound [6], acoustic emission [7], laser scanning [8], electromagnetic [9], radiographic [10], and thermographic methods [11]. Among these, vision-based strategies are particularly valued for their high degree of automation and effectiveness in detecting and evaluating defects, thereby playing a crucial role in the in-situ quality control systems in AM [12, 13].

While various strategies for identifying defects in AM processes exist, many are not suitable for in-situ applications due to limitations such as the inability to detect fine defects or the lack of validated accuracy metrics. Machine learning (ML) algorithms offer a promising enhancement for defect detection precision. These algorithms process and analyze multiple data features beyond simple point distances, providing more detailed analysis [14, 15]. The integration leverages the capabilities of advanced sensing technologies with the robust analytical power of ML, offering deeper insights into the AM process [16, 17].

The deployment of in-process monitoring systems in AM requires data analysts to access proprietary organizational data, which introduces the risk of insider threats. This concern arises when individuals misuse their authorized access to confidential information for unauthorized purposes, such as data breaches or personal gain, as demonstrated by Theoharidou et al. [18]. The risk is further intensified by insiders' thorough understanding of the in-situ sensing data management and operational proto-

*Corresponding author: farhad.imani@uconn.edu

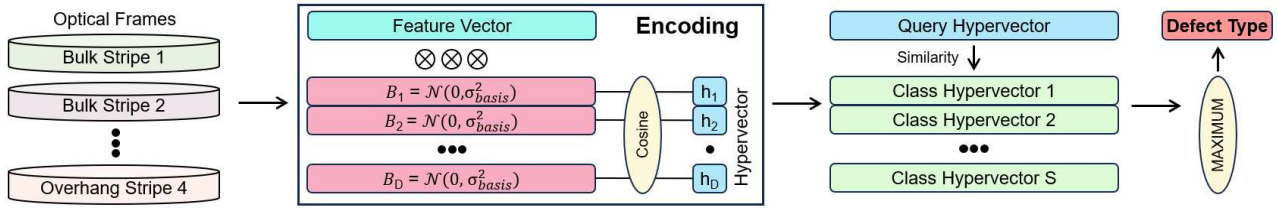

FIGURE 1: HYPERDIMENSIONAL COMPUTING FRAMEWORK: ENCODING, TRAINING, INFERENCE, AND RETRAINING.

cols. Such knowledge could lead to the unauthorized disclosure of sensitive information by exploiting predictive analytics and context-specific data [19].

To mitigate privacy concerns, several privacy-preserving strategies have been devised, including cryptographic techniques, data anonymization, and differential privacy (DP). Cryptographic methods are essential for protecting identities and controlling access to systems, thereby providing the confidentiality and integrity of information [20]. Despite their effectiveness, cryptographic methods require considerable computational resources, potentially slowing down data processing. As an alternative, data anonymization involves removing sensitive information before data analysis, reducing the risk of data breaches, especially in complex, high-dimensional datasets [21]. However, anonymization has its limitations, such as the permanent loss of original data and the inability to provide complete privacy protection [22]. Even with anonymized datasets, re-identification of individuals is still possible. For instance, an analysis of anonymized Netflix data released for a competition demonstrated that individual users could be identified by correlating their reviews with data from other websites like IMDb, revealing their viewing habits [23].

To address privacy concerns, as introduced by Dwork [24], integrating noise into data processing algorithms ensures that databases differing by only one record remain indistinguishable from the perspective of the same algorithm, thereby significantly enhancing privacy protection. However, this approach necessitates a careful balance between enhancing privacy and maintaining algorithmic performance. Increasing the noise level strengthens privacy protections but may also negatively impact the algorithm's accuracy and efficiency. Therefore, it is crucial to find an optimal balance that ensures robust data protection without compromising utility and effectiveness.

Building upon the goal of simultaneously enhancing privacy and maintaining performance, this paper introduces the integration of DP with Hyperdimensional Computing (HDC), an advanced machine learning framework. HDC's exceptional speed in processing queries makes it particularly suitable for scenarios requiring swift monitoring. By incorporating DP into HDC, we create a robust framework that enhances privacy while preserving performance. The unique data manipulation method of HDC, which utilizes high-dimensional vector spaces for computations, combined with the privacy guarantees of DP, ensures rapid analysis and maintains an optimal balance between privacy protection and efficiency.

2. RESEARCH METHODOLOGY

2.1 Hyperdimensional Computing

The HDC architecture is inspired by the human brain, using high-dimensional vectors to mimic cognitive functions [25, 26]. This technique enables HDC to imitate the brain's methods of information processing and memory retention across various cognitive tasks [27, 28]. As shown in Figure 1, the HDC model is divided into four phases: encoding, training, inference, and retraining.

Encoding focuses on creating a high-dimensional framework by converting feature vectors into hypervectors. In this framework, each element of a hypervector equally contributes to encoding information, ensuring that no single component disproportionately influences the overall data representation. The process involves mapping a feature vector $\vec{F} = \{f_j\}_{j=1}^{J}$, where each f_j belongs to $\mathbb{R}$, into a hypervector $\vec{H}$. This mapping is achieved by computing the dot product between $\vec{F}$ and a unique, randomly generated vector for each dimension, given by the equation $h_i = \cos(\vec{F} \cdot \vec{B}_i)$. Each $\vec{B}_i$ is a random vector drawn from a Gaussian distribution with a mean of zero and a variance of σ^2_{basis}. These vectors, $\{B_i\}_{i=1}^{D}$, serve as the random basis vectors, are determined initially, and remain fixed throughout. This approach efficiently translates continuous feature vectors into continuous-valued hypervectors, making them suitable for classification tasks in various fields, such as anomaly detection [29, 30].

When discussing the key elements of the encoding methodology, it is essential to consider the impact of the variance, σ^2_{basis}, in the Gaussian distribution used for generating the random basis vectors, $\vec{B}_i$. The selection of variance plays a crucial role, shaping the encoding strategy to be either exclusive or inclusive. By choosing a larger σ^2_{basis}, an exclusive encoding strategy is employed, which increases the spread of the basis vectors' values. This increased dispersion enhances the uniqueness of each hypervector, aiding in their distinction. Such an approach is particularly beneficial in scenarios where precise differentiation between encoded vectors is required, such as in tasks involving the accurate identification of unique patterns.

Conversely, an inclusive encoding strategy is defined by a reduced variance, σ^2_{basis}, leading to a more concentrated distribution of the basis vector values. This concentration results in hypervectors that are more similar to each other, which is advantageous for tasks requiring generalization over similar features. Adjusting the variance parameter provides nuanced control over the encoding's specificity and generalizability, showcasing the versatility of HDC in addressing a wide range of classification challenges.

Training begins with the creation of class-specific hypervectors, denoted as $\vec{C}_s$, for each class s, where s ranges from 1 to S, representing the total number of classes. Each class hypervector is formed by summing all the associated hypervectors for that class, as expressed by the following equation:

$$\vec{C}_s = \sum \vec{H}^s \tag{1}$$

where $\vec{H}^s$ represents the hypervectors associated with class s. This equation illustrates how class hypervectors aggregate the features of the training samples from their respective classes.

Combining hypervectors from the same class allows for the merging of shared characteristics and attributes within a high-dimensional framework. This process ensures that the core aspects of each class are integrated into a unified, clear representation, which is essential for accurate and efficient classification. By aggregating hypervectors, the HDC framework leverages the principle of superposition inherent to high-dimensional spaces, improving its ability to differentiate between classes by examining their collective features.

The creation of class hypervectors is a crucial step in the HDC training process, serving as the foundation for the subsequent inference and retraining phases. This step allows HDC to develop a thorough understanding of each class's distinct high-dimensional profile, which is vital for the model's accurate pattern recognition and predictive capabilities.

Inference is a critical phase where the class association of a hypervector is determined by comparing the query hypervector with the class hypervectors established during training. Cosine similarity is particularly effective for this comparison, as it measures the alignment of two vectors within the high-dimensional space. This similarity is computed as the dot product of two vectors, normalized by the product of their norms. The formula for calculating the cosine similarity between a class hypervector $\vec{C}_s$ and a query hypervector $\vec{H}_q$ is given by:

$$\mathrm{Cos}(\vec{C}_s, \vec{H}_q) = \frac{\vec{C}_s \cdot \vec{H}_q}{||\vec{C}_s|| \cdot ||\vec{H}_q||} \tag{2}$$

During the inference process, the HDC model assesses the query hypervector $\vec{H}_q$ against each class hypervector using cosine similarity metrics. The class whose hypervector $\vec{C}_s$ shows the highest similarity to $\vec{H}_q$ is identified as the most likely class affiliation for $\vec{H}_q$. This method allows the HDC model to classify new data by leveraging the learned patterns and utilizing the spatial properties unique to high-dimensional spaces to distinguish between different classes.

The HDC's inference, which employs cosine similarity, is engineered to deliver both high accuracy and computational efficiency, making it highly adaptable for various classification tasks. Notably, the norm of the query hypervector, $||\vec{H}_q||$, remains consistent across all classes, allowing it to be excluded from detailed calculations. Similarly, the magnitude of each class hypervector, $||\vec{C}_s||$, stays constant during the inference stage since these vectors are fixed during training and only need to be calculated once. Consequently, the core of the inference process is simplified to computing dot products between the query hypervector and class hypervectors, a remarkably fast process. This swift processing makes HDC exceptionally suited for real-time monitoring in AM, where rapid and accurate data classification is crucial.

Retraining is an essential element of the HDC approach, greatly improving the model's accuracy and adaptability. This phase involves comparing hypervectors of training samples with the class hypervectors established during the initial training phase. By identifying and correcting inaccuracies, the retraining process enhances the overall performance of the model.

If a hypervector, $\vec{H}^s$, is incorrectly classified into a wrong class, $s^{'}$, instead of its correct class, s, the HDC model updates to correct this error. The process adjusts the class hypervectors according to the magnitude of the misclassification, as shown in the following equations:

$$\begin{aligned} \vec{C}_s &= \vec{C}_s + \vec{H}^s, \\ \vec{C}_{s'} &= \vec{C}_{s'} - \vec{H}^s \end{aligned} \tag{3}$$

This recalibration effectively adjusts the class hypervectors by incorporating the misclassified hypervector, $\vec{H}^s$, into its correct class, s, and removing it from the incorrect class, $s^{'}$. These modifications bring the correct class's hypervector closer to $\vec{H}^s$ and distance it from the hypervector of the incorrect class, thereby decreasing the chances of future misclassifications.

The retraining process is inherently iterative, enabling the HDC model to continuously improve its classification accuracy by correcting errors. Each update incrementally enhances the model's ability to capture nuanced variations within the data, boosting its efficiency and reliability. This ongoing optimization is particularly valuable in environments where data distributions shift or new patterns emerge.

Retraining enhances the HDC model into a progressively more effective classification tool, capable of adapting to evolving data landscapes without needing to restart the training process. This adaptability, along with the inherent efficiency of the HDC method, makes it an excellent solution for real-time data processing and decision-making across various applications.

2.2 Differential Privacy

In our study, we tackle the critical issue of preserving confidentiality in real-time monitoring, with a particular focus on the risks posed by insiders. Traditional systems often inadvertently expose sensitive information due to their extensive data accessibility requirements. To address this, we have implemented DP mechanisms. These mechanisms ensure that only outcomes modified according to DP standards are accessible, rather than the original, unmodified data. This approach significantly enhances security, greatly reducing the risk of privacy violations by insiders. It not only protects sensitive data but also strengthens the integrity of real-time monitoring against internal security threats.

DP enhances the security infrastructure of real-time monitoring by restricting access to unprocessed data and allowing only the dissemination of results that have been modified using DP techniques. DP operates by intentionally adding noise to obscure the original data, a process governed by key metrics: the privacy budget (ϵ) and the privacy loss threshold (δ). These parameters are crucial in determining the noise intensity and the level of privacy protection, significantly reducing the likelihood of inferring

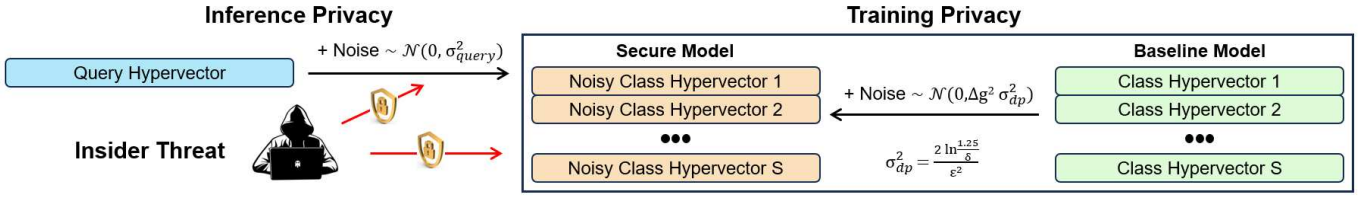

FIGURE 2: ENHANCING PRIVACY IN HDC TRAINING AND INFERENCE WITH GAUSSIAN NOISE INTEGRATION FOR DIFFERENTIAL PRIVACY.

individual data points. This approach ensures that the privacy of data subjects is preserved without compromising the utility of the data for analysis, thereby providing a robust solution to the challenge of protecting sensitive information in real-time monitoring in AM from potential insider threats.

In the context of AM, to elucidate the concept of DP, consider two datasets, I_1 and I_2, which are identical except for a single data point. A predictive model M is said to exhibit (ϵ, δ)-DP if, for any such pair of datasets I_1 and I_2, the model satisfies the following inequality:

$$\mathbb{P}[M(I_1)] \leq e^{\epsilon} \cdot \mathbb{P}[M(I_2)] + \delta \tag{4}$$

This criterion ensures that the presence or absence of a single data entry does not significantly impact the model's output, thereby protecting the confidentiality of individual data points.

In predictive modeling applications, particularly within AM, the Gaussian mechanism is a key technique for enforcing DP. This method involves strategically adding noise to the model's output to effectively mask the influence of any single data point. Consider a function $g : I \rightarrow \mathbb{R}^D$ responsible for data manipulation or transformation in the model. The Gaussian mechanism, characterized by the noise level σ^2_{dp}, is implemented as follows:

$$M(I) = g(I) + \mathcal{N}(0, \Delta g^2 \sigma^2_{dp}) \tag{5}$$

The term Δg, known as the sensitivity of g, quantifies the maximum expected difference in g's output when comparing any two datasets, I_1 and I_2, that differ by only one element:

$$\Delta g = \max_{I_1, I_2} \|g(I_1) - g(I_2)\| \tag{6}$$

To achieve (ϵ, δ)-DP, the noise level, σ^2_{dp}, must adhere to the following condition:

$$\sigma^2_{dp} > 2 \ln \frac{1.25}{\delta} \cdot \frac{1}{\epsilon^2} \tag{7}$$

In our specific case, we set δ to 10^{-4}, based on the previous studies in the DP field, such as Abadi et al. [31], which suggest that δ should be smaller than the inverse of the dataset size. We determine the optimal σ^2_{dp} using the above equation, aiming to balance the trade-off between preserving data privacy and maintaining the operational performance of our predictive model in the AM setting.

During the training phase of the HDC model, DP is incorporated by making the class hypervectors differentially private through the addition of Gaussian noise, as shown in Figure 2. This approach effectively reduces the likelihood that an adversary can infer specific details about the original feature vectors from the class hypervectors.

The selection of the appropriate noise level, which is critical for enhancing the model's security, depends on the sensitivity of the HDC model. This sensitivity is determined by identifying the maximum norm among all encoded training samples, as illustrated in Eq. (8). Higher sensitivity requires the addition of more noise, thereby strengthening the model's security. Consequently, the highest sensitivity observed across all training samples is used as a collective sensitivity indicator, ensuring robust privacy measures throughout the dataset.

$$\Delta g = \max \|\vec{H}\| \tag{8}$$

In AM contexts, ensuring secure data exchange during the inference phase is crucial. Query hypervectors are enhanced with an additional layer of security by appending Gaussian noise before their dissemination, as shown in Figure 2. This is achieved by applying Gaussian noise with a distribution of $\mathcal{N}(0, \sigma^2_{\text{query}})$ to the query hypervectors, thereby maintaining the privacy of the contained information.

3. EXPERIMENTAL DESIGN AND RESULTS

Addressing privacy concerns in the dissemination and analysis of sensor data from AM processes is a paramount consideration. The need to protect sensitive manufacturing data, especially in an era where intellectual property and competitive advantage are closely guarded, necessitates the exploration of privacy-preserving methodologies. This study, therefore, not only focuses on the technical aspects of in-situ sensing and defect detection but also the implementation of DP measures. These measures ensure that the valuable insights gained from high-speed camera data do not compromise the confidentiality of the manufacturing process or the design specifications of the parts being produced. This introduction sets the stage for a comprehensive investigation into the experimental designs and methodologies employed in this research. The subsequent sections will delve into the specifics of sensor integration, data acquisition, and the analytical techniques used to achieve the dual objectives of enhancing manufacturing precision and ensuring data privacy in the realm of AM.

3.1 Experimental Setup

This section outlines the instrumentation of sensor equipment on Laser Powder Bed Fusion (LPBF) machine, specifically the EOS M270 model used at the National Institute of Standards and Technology (NIST). The equipment includes a visible high-speed camera aimed at observing the melt pool dynamics and the dispersion of particles around it. The camera monitors the formation and behavior of a melt pool created in a sample made

from nickel alloy 625 (Inconel 625, UNS N06625) for an unsupported overhang with a 40.5-degree angle. The investigation focuses on sensor data collected from three different heights of the build process: 6.06 mm, 7.90 mm, and 9.70 mm, each incorporating structures with overhangs. The objective is to uncover patterns in thermal activity during the process, particularly in the last segments of scanning before or right after the formation of an edge, excluding the contour scan, as it involves the final scanning vectors. Apart from these, all scans are included in the analysis. The scanning procedure involves four passes over each layer within a 4 mm height increment, rotating the direction of each layer's scans by 90 degrees to achieve a vertical striping effect observable under the thermal camera [32].

Despite its simplicity, the piece's design used in this study highlights the technology's limitations in capturing complex geometrical features. To facilitate the understanding of thermal dynamics in overhang constructions, NIST employed a test piece with a manageable size and shape, avoiding the issues with the infrared camera's focus on larger objects. The study indicates that non-coaxial sensor alignment could lead to less precise thermal readings due to potential image blurring when observing larger pieces. Additionally, visible high-speed camera footage is presented in a 256 × 256 pixel frame, which shot at a rate of 1000 frames per second.

3.1.1 High-Speed Camera Integration and Data Acquisition. In the pursuit of advancing in-situ monitoring within the LPBF process, our experimental framework integrates a high-speed visible spectrum video camera, characterized by a frame rate of 1000 frames per second and a 256 × 256 pixel frame. Positioned strategically within the build chamber, the camera captures detailed thermal patterns as the laser fuses powder layers, creating a dynamic record of the melting process. The specificity of the camera's placement and settings allows for an unprecedented view into the thermal dynamics at play, particularly in the formation of overhang features and the solidification of bulk areas.

The data acquisition phase is meticulously designed to ensure the capture of high-fidelity images, which are then windowed to a manageable size for processing. This optimization balances the need for detailed thermal data against the practical considerations of data storage and processing speed. The focus on high-speed camera data, among other sensor inputs, stems from its direct correlation with the thermal phenomena under investigation, offering a rich dataset for subsequent analysis.

3.1.2 Classification Framework for Enhanced Precision. Building on the foundational data captured, the study introduces a nuanced classification framework that delineates eight classes based on specific overhang and bulk characteristics across four distinct strips. This classification is not merely academic; it has profound implications for understanding the thermal behavior in AM and the structural integrity and surface quality of the printed parts. Each class represents a unique thermal signature, corresponding to different combinations of overhang features and bulk regions, thereby enabling a detailed analysis of potential defects and irregularities.

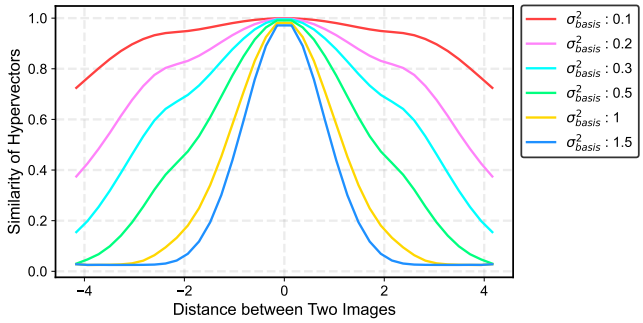

FIGURE 3: COMPARATIVE OUTCOMES OF EXCLUSIVE AND INCLUSIVE ENCODING IN OPTICAL IMAGING.

3.1.3 Importance of Privacy in Additive Manufacturing. As we navigate through the complexities of in-situ sensing and data analysis, the imperative of preserving privacy emerges as a critical concern. The integration of DP techniques in the processing of sensor data addresses this concern head-on, ensuring that while the data provides invaluable insights into the LPBF process, it does not expose sensitive information related to part designs or proprietary manufacturing techniques. This dual focus on enhancing manufacturing precision while safeguarding data privacy forms the cornerstone of our approach, setting a new standard for research and practice in the field of AM.

3.2 Experimental Results

3.2.1 Training Privacy. Within the HDC framework, optical frames are processed using specialized real-time monitoring framework to ensure privacy. This procedure involves converting the training dataset's feature vectors into hypervectors within a high-dimensional space. Subsequently, these hypervectors are combined to create class hypervectors, which form the foundation of the HDC model. Retraining is conducted to refine these models further. To augment the model's robustness, noise is added to both the class hypervectors and the encoded queries.

Determining the optimal variance, denoted as $\sigma^{2*}_{\text{basis}}$, is critical for encoding efficiency. Figure 3 demonstrates the impact of variance on the encoding quality for optical images, highlighting the trade-offs between exclusive and inclusive encoding strategies. A lower σ^2_{basis} results in high similarity across hypervectors regardless of the actual distance between the corresponding images. For instance, with $\sigma^2_{\text{basis}} = 0.1$, hypervectors exhibit a similarity exceeding 0.7, even for images at maximal separation. Conversely, increasing σ^2_{basis} gradually aligns hypervector similarity with the physical distance between images. Notably, at a σ^2_{basis} of 1.5, images spaced more than a certain distance apart yield virtually zero similarity in their hypervectors. This observation underscores the necessity of selecting a σ^2_{basis} that neither overly compresses nor dilutes the feature space, ensuring hypervectors accurately reflect the original feature vectors.

The primary aim of our research is to strike an optimal balance between safeguarding privacy and maintaining high system performance. The encoding variance, represented as σ^2_{basis}, is pivotal not only in the encoding process but also in achieving this balance between privacy enhancement and performance op-

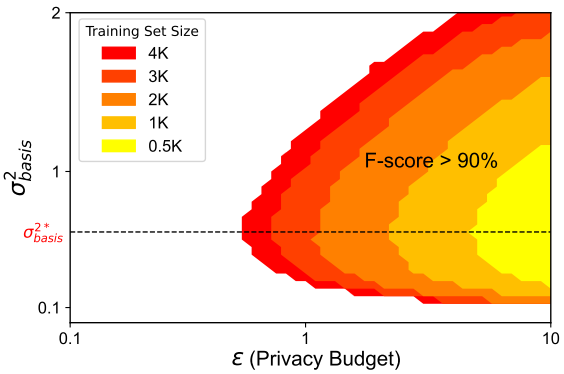

FIGURE 4: PRIVACY-PERFORMANCE TRADE-OFF IN HYPERDIMENSIONAL COMPUTING: IMPACT OF ENCODING VARIANCE ON F-SCORE AND PRIVACY BUDGET.

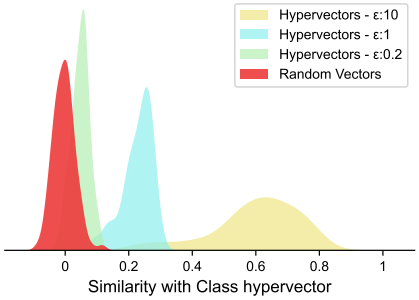

FIGURE 5: EFFECT OF DIFFERENTIAL PRIVACY ON HYPERDIMENSIONAL COMPUTING MEMORIZATION.

timization. We utilize the F-score to assess our model's performance. The statistical accuracy is measured in terms of the statistical F-score, which combines both the type I (false alarm) and type II (failing to detect) statistical errors. Figure 4 depicts the HDC model's effectiveness across varying ϵ levels for both exclusive and inclusive encoding methods. Our approach involved experimenting with various sizes of training datasets to delineate regions where the F-score surpasses the 90% threshold. Based on these findings, a select subset of the training data was used to develop the HDC model. Additional samples were then employed for validation to ascertain the optimal $\sigma^{2*}_{\text{basis}}$ value.

Our search was geared towards identifying a $\sigma^{2*}_{\text{basis}}$ that yields a high F-score, indicative of enhanced performance, alongside a low ϵ, signifying robust privacy measures. It became evident that the ideal $\sigma^{2*}_{\text{basis}}$ does not depend on the size of the training dataset. For instance, at a low $\sigma^{2}_{\text{basis}}$ such as 0.1, the F-score invariably falls below 90%. As $\sigma^{2}_{\text{basis}}$ increases, the F-score climbs above 90%, coinciding with smaller ϵ values. Nonetheless, elevating $\sigma^{2}_{\text{basis}}$ beyond the optimal point results in the need for a higher ϵ to maintain an F-score above 90%. Particularly, when $\sigma^{2}_{\text{basis}}$ reaches 2, achieving an F-score above 90% requires an impractically high ϵ or may not even be feasible for certain training dataset sizes. Conversely, an intermediate $\sigma^{2}_{\text{basis}}$, neither too low nor too high, facilitates achieving an F-score above 90% at more acceptable ϵ levels, such as 0.6 in our evaluation.

Upon identifying the optimal variance for the random basis vectors, the next step involves evaluating the effects of noise addition within the HDC model. After aggregating multiple hypervectors to construct a class hypervector, Figure 5 displays a histogram that elucidates the similarity between individual hypervectors and the aggregated class hypervector, underlining the influence of incorporating different degrees of noise into the class hypervector. Based on Eq. (7) and Eq. (5), a lower ϵ indicates the addition of more noise. The similarity between the class hypervector and any random vectors remains negligible, as the class hypervector does not retain information from these unrelated vectors. With an ϵ of 10, indicating minimal noise introduction, the similarity scores between hypervectors and the class hypervector remain high, suggesting that the class hypervector preserves the information from the original hypervectors despite the noise. However, increasing the noise level within the class hypervector leads to a predominance of noise information over the original hypervector data, thereby reducing the similarity scores. For instance, at an ϵ of 0.1, signifying the addition of substantial noise, the similarity drops to levels comparable to those with random vectors, indicating that the class hypervector no longer effectively encapsulates the information from the original hypervectors.

TABLE 1: MODEL F-SCORE PERFORMANCE ACROSS DIFFERENT EPSILON VALUES.

Model	$\epsilon = 0.5$	$\epsilon = 0.7$	$\epsilon = 1$
HDC	**87.71%**	**91.77%**	**94.30%**
ResNet50	17.05%	28.90%	42.59%
AlexNet	20.27%	43.56%	48.09%
DenseNet201	14.61%	43.29%	65.41%
EfficientNet B2	20.27%	22.48%	28.11%

To evaluate the robustness of the HDC model against traditional models under varying levels of privacy, we benchmarked the HDC model against established models in the image classification field, including ResNet50, DenseNet201, EfficientNet B2, and AlexNet. For implementing DP, we utilized Opacus, an open-source library designed for PyTorch that facilitates training deep learning models with DP mechanisms [33]. Opacus is designed for ease of use, adaptability, and efficiency, ensuring seamless integration into existing machine learning workflows. It employs Differential Privacy Stochastic Gradient Descent (DP-SGD), a variation of the standard SGD algorithm. DP-SGD enhances privacy by introducing precisely calibrated noise to the gradients during the training process. This approach minimizes the impact of individual data entries by clipping gradients and adding noise, thus masking their specific contributions to the final model parameters. The features of DP-SGD provide measurable privacy assurances, which are crucial in fields where data confidentiality is essential.

Table 1 illustrates the F-score for the image classification task after applying the DP mechanism. The introduction of DP leads to varying impacts on the F-score, showing that HDC experiences a minor performance reduction compared to significant drops in other models. At a stringent privacy level, represented by an ϵ value of 0.5, HDC maintains an F-score of 87.71%, while other models fall below 21%, highlighting HDC's robustness against

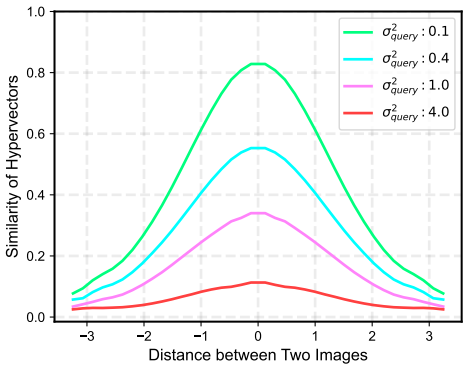

FIGURE 6: NOISE INFLUENCE ON HYPERVECTOR SIMILARITY ACROSS IMAGE DISTANCES IN HDC.

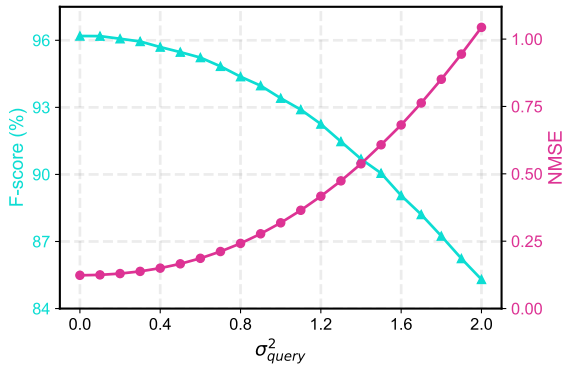

FIGURE 7: TRADE OFF BETWEEN INFERENCE DATA PRIVACY AND HDC PERFORMANCE.

high noise levels as it distributes memory evenly across all dimensions of the hyperspace. As ϵ increases to 0.7, reducing the noise, HDC's F-score rises to 91.77%, significantly surpassing other models, which remain below 44%. When ϵ reaches 1, HDC achieves an F-score of 94.30%, whereas other models—42.59% for ResNet50, 48.09% for AlexNet, and 28.1% for EfficientNet B2—show a substantial disparity, with DenseNet201 being the closest yet still 28.89% lower than HDC. This demonstrates HDC's superior capability to balance privacy and performance compared to other prevalent ML models in the field.

3.2.2 Inference Privacy. Introducing noise to encoded queries enhances their privacy, yet it tends to shift the focus of the hypervector elements towards retaining noise rather than the information from the feature vector queries. Figure 6 illustrates the relationship between hypervector similarity and the distance between query images, under the influence of various noise levels added to the hypervector. When the variance σ^2_{query} is minimal, there's a high similarity among hypervectors, which correlates with the distance between images, indicating that the hypervectors preserve the image information effectively. For instance, at $\sigma^2_{\text{query}} = 0.1$, the similarity for closely related images exceeds 0.8. Conversely, escalating the noise level diminishes this similarity, and at a significantly high σ^2_{query}, such as 4, the similarity drops to zero, irrespective of the distance between images, as the hypervectors predominantly reflect noise data. Thus, while augmenting the noise in encoded queries serves to privatize them, it inadvertently compels the hypervectors to overlook the essential query information.

The aim of ensuring privacy during the inference phase is to strike a balance between query confidentiality and HDC model efficiency, as depicted in Figure 7. Different levels of noise are applied to encoded queries, and the F-score is evaluated to assess HDC performance. Additionally, a neural network, simulating an adversary, attempts to reconstruct the original query images from the hypervectors. The effectiveness of this reconstruction is quantified using the Normalized Mean Square Error (NMSE), which gauges the neural network's success in decoding the hypervectors. Without noise addition, the F-score surpasses 96%, indicating high HDC performance; however, the queries remain unsecured, as evidenced by a low NMSE of 0.12, indicating that the neural network can effortlessly decode the hypervectors to retrieve the images. Increasing the noise level enhances query privacy, as reflected by a rise in NMSE, but concurrently, it compromises HDC performance by shifting the hypervectors' focus from image data to noise. At a σ^2_{query} of 2, the NMSE exceeds 1, suggesting the neural network's inability to extract any meaningful information from the hypervectors, yet this comes at the cost of reduced HDC effectiveness, with an F-score of 85%. Assuming an F-score of 90% is an acceptable threshold for HDC performance, such a level is attainable with a noise variance of $\sigma^2_{\text{query}} = 1.5$, which corresponds to an NMSE of 0.6. This indicates a feasible noise threshold that maintains queries' privacy without significantly impairing HDC performance.

4. CONCLUSION

This study introduces a novel method to improve privacy and efficiency in additive manufacturing (AM) by integrating in-process monitoring with differential privacy (DP) and hyperdimensional computing (HDC). Experimental results on in-process monitoring of the overhang anomaly via high-speed melt pool data emphasize the success of this approach, demonstrating its capacity to balance operational efficiency, prediction precision, and data privacy within AM's critical framework. The proposed scalable solution opens new avenues for AM, setting a benchmark for privacy-preserving, efficient in-process monitoring. Future directions include broadening this method's application across manufacturing processes, underscoring its significance in advancing manufacturing technologies while ensuring data security.

Future work should focus on applying the proposed method to other additive manufacturing processes, such as fused filament fabrication and direct metal deposition, to explore its versatility and robustness. Additionally, while this paper studied different encoding strategies in hyperdimensional computing (HDC), it would be valuable to develop a method that can automatically determine the best encoding strategy for HDC based on the dataset's distribution, enhancing the efficiency and effectiveness of the proposed approach. Moreover, investigating the scalability of our framework to handle larger datasets and more complex geometries will be crucial. Ensuring that the method performs well under these conditions will validate its practicality for industrial applications. These future directions underscore the potential of

our method to advance manufacturing technologies while maintaining data security.

ACKNOWLEDGMENTS

This work was supported by the National Science Foundation [grant numbers 2127780, 2312517]; the Semiconductor Research Corporation (SRC); the Office of Naval Research [grant numbers N00014-21-1-2225, N00014-22-1-2067]; the Air Force Office of Scientific Research [grant number FA9550-22-1-0253]; and generous gifts from Xilinx and Cisco. The authors gratefully acknowledged the valuable contributions from the National Institute of Standards and Technology (NIST) in particular Dr. Brandon Lane for providing the data for this research.

REFERENCES

[1] Bhatia, Akash and Sehgal, Anuj Kumar. "Additive manufacturing materials, methods and applications: A review." *Materials Today: Proceedings* Vol. 81 (2023): pp. 1060–1067.

[2] Imani, Farhad and Khanzadeh, Mojtaba. "Image-guided multi-response modeling and characterization of design defects in metal additive manufacturing." *ASME International Mechanical Engineering Congress and Exposition*, Vol. 85567: p. V02BT02A046. 2021. American Society of Mechanical Engineers.

[3] Leach, Richard. *Optical measurement of surface topography*. Vol. 8. Springer (2011).

[4] Mandache, Catalin. "Overview of non-destructive evaluation techniques for metal-based additive manufacturing." *Materials Science and Technology* Vol. 35 No. 9 (2019): pp. 1007–1015.

[5] Thomas, Douglas S, Gilbert, Stanley W et al. "Costs and cost effectiveness of additive manufacturing." *NIST special publication* Vol. 1176 (2014): p. 12.

[6] Lv, Gaolong, Yao, Zhijun, Chen, Dan, Li, Yehai, Cao, Huanqing, Yin, Anmin, Liu, Yanjun and Guo, Shifeng. "Fast and high-resolution laser-ultrasonic imaging for visualizing subsurface defects in additive manufacturing components." *Materials & Design* Vol. 225 (2023): p. 111454.

[7] Kononenko, Denys Y, Nikonova, Viktoriia, Seleznev, Mikhail, van den Brink, Jeroen and Chernyavsky, Dmitry. "An in situ crack detection approach in additive manufacturing based on acoustic emission and machine learning." *Additive manufacturing letters* Vol. 5 (2023): p. 100130.

[8] Wang, Kang. "Contrastive learning-based semantic segmentation for In-situ stratified defect detection in additive manufacturing." *Journal of Manufacturing Systems* Vol. 68 (2023): pp. 465–476.

[9] Ramírez, Isaac Segovia, Márquez, Fausto Pedro García and Papaelias, Mayorkinos. "Review on additive manufacturing and non-destructive testing." *Journal of Manufacturing Systems* Vol. 66 (2023): pp. 260–286.

[10] Hu, Chuanfei and Wang, Yongxiong. "An efficient convolutional neural network model based on object-level attention mechanism for casting defect detection on radiography images." *IEEE Transactions on Industrial Electronics* Vol. 67 No. 12 (2020): pp. 10922–10930.

[11] Oster, Simon, Breese, Philipp P, Ulbricht, Alexander, Mohr, Gunther and Altenburg, Simon J. "A deep learning framework for defect prediction based on thermographic in-situ monitoring in laser powder bed fusion." *Journal of Intelligent Manufacturing* Vol. 35 No. 4 (2024): pp. 1687–1706.

[12] Chen, Ruimin, Imani, Mohsen and Imani, Farhad. "Joint active search and neuromorphic computing for efficient data exploitation and monitoring in additive manufacturing." *Journal of manufacturing processes* Vol. 71 (2021): pp. 743–752.

[13] Imani, Farhad and Chen, Ruimin. "Latent Representation and Characterization of Scanning Strategy on Laser Powder Bed Fusion Additive Manufacturing." *ASME 2022 International Mechanical Engineering Congress and Exposition*: p. 2022. 2022. American Society of Mechanical Engineers Digital Collection.

[14] Xames, Md Doulotuzzaman, Torsha, Fariha Kabir and Sarwar, Ferdous. "A systematic literature review on recent trends of machine learning applications in additive manufacturing." *Journal of Intelligent Manufacturing* Vol. 34 No. 6 (2023): pp. 2529–2555.

[15] Rescsanski, Sean, Yadollahi, Aref, Khanzadeh, Mojtaba and Imani, Farhad. "Anomaly Detection of Laser-Based Metal Additive Manufacturing Using Neural-Variational Auto-Encoder." *International Manufacturing Science and Engineering Conference*, Vol. 87233: p. V001T01A026. 2023. American Society of Mechanical Engineers.

[16] Wang, Chengcheng, Tan, Xipeng P, Tor, Shu Beng and Lim, CS. "Machine learning in additive manufacturing: State-of-the-art and perspectives." *Additive Manufacturing* Vol. 36 (2020): p. 101538.

[17] Zhao, Xiyue, Imandoust, Aidin, Khanzadeh, Mojtaba, Imani, Farhad and Bian, Linkan. "Automated anomaly detection of laser-based additive manufacturing using melt pool sparse representation and unsupervised learning." (2021).

[18] Theoharidou, Marianthi, Kokolakis, Spyros, Karyda, Maria and Kiountouzis, Evangelos. "The insider threat to information systems and the effectiveness of ISO17799." *Computers & Security* Vol. 24 No. 6 (2005): pp. 472–484.

[19] Xu, Runhua, Baracaldo, Nathalie and Joshi, James. "Privacy-preserving machine learning: Methods, challenges and directions." *arXiv preprint arXiv:2108.04417* (2021).

[20] Shi, Zhangyue, Oskolkov, Boris, Tian, Wenmeng, Kan, Chen and Liu, Chenang. "Sensor data protection through integration of blockchain and camouflaged encryption in cyber-physical manufacturing systems." *Journal of Computing and Information Science in Engineering* Vol. 24 No. 7 (2024).

[21] Neves, Flávio, Souza, Rafael, Sousa, Juliana, Bonfim, Michel and Garcia, Vinicius. "Data privacy in the Internet of Things based on anonymization: A review." *Journal of Computer Security* Vol. 31 No. 3 (2023): pp. 261–291.

[22] Sweeney, Latanya. "Matching known patients to health records in Washington State data." *arXiv preprint arXiv:1307.1370* (2013).

[23] Narayanan, Arvind and Shmatikov, Vitaly. "Robust de-anonymization of large sparse datasets." *2008 IEEE Symposium¸on Security and Privacy (sp 2008)*: pp. 111–125. 2008. IEEE.

[24] Dwork, Cynthia, Roth, Aaron et al. "The algorithmic foundations of differential privacy." *Foundations and Trends® in Theoretical Computer Science* Vol. 9 No. 3–4 (2014): pp. 211–407.

[25] Chen, Ruimin, Sodhi, Manbir, Imani, Mohsen, Khanzadeh, Mojtaba, Yadollahi, Aref and Imani, Farhad. "Brain-inspired computing for in-process melt pool characterization in additive manufacturing." *CIRP Journal of Manufacturing Science and Technology* Vol. 41 (2023): pp. 380–390.

[26] Hoang, Danny, Errahmouni, Hamza, Chen, Hanning, Rachuri, Sriniket, Mannan, Nasir, ElKharboutly, Ruby, Imani, Mohsen, Chen, Ruimin and Imani, Farhad. "Hierarchical representation and interpretable learning for accelerated quality monitoring in machining process." *CIRP Journal of Manufacturing Science and Technology* Vol. 50 (2024): pp. 198–212.

[27] Piran, Fardin Jalil, Poduval, Prathyush P, Barkam, Hamza Errahmouni, Imani, Mohsen and Imani, Farhad. "Explainable Hyperdimensional Computing for Balancing Privacy and Transparency in Additive Manufacturing Monitoring." *arXiv preprint arXiv:2407.07066* (2024).

[28] Poduval, Prathyush, Alimohamadi, Haleh, Zakeri, Ali, Imani, Farhad, Najafi, M Hassan, Givargis, Tony and Imani, Mohsen. "Graphd: Graph-based hyperdimensional memorization for brain-like cognitive learning." *Frontiers in Neuroscience* Vol. 16 (2022): p. 757125.

[29] Jalil Piran, Fardin, Barkam, Hamza Errahmouni, Imani, Mohsen and Imani, Farhad. "Hyperdimensional Cognitive Computing for Lightweight Cyberattack Detection in Industrial Internet of Things." *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, Vol. 87356: p. V007T07A013. 2023. American Society of Mechanical Engineers.

[30] Hoang, Danny, Mannan, Nasir, ElKharboutly, Ruby, Chen, Ruimin and Imani, Farhad. "Edge Cognitive Data Fusion: From In-Situ Sensing to Quality Characterization in Hybrid Manufacturing Process." *International Manufacturing Science and Engineering Conference*, Vol. 87240: p. V002T06A029. 2023. American Society of Mechanical Engineers.

[31] Abadi, Martin, Chu, Andy, Goodfellow, Ian, McMahan, H Brendan, Mironov, Ilya, Talwar, Kunal and Zhang, Li. "Deep learning with differential privacy." *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*: pp. 308–318. 2016.

[32] Lane, Brandon, Whitenton, Eric and Moylan, Shawn. "Multiple sensor detection of process phenomena in laser powder bed fusion." *Thermosense: Thermal Infrared Applications XXXVIII*, Vol. 9861: pp. 20–28. 2016. SPIE.

[33] Yousefpour, Ashkan, Shilov, Igor, Sablayrolles, Alexandre, Testuggine, Davide, Prasad, Karthik, Malek, Mani, Nguyen, John, Ghosh, Sayan, Bharadwaj, Akash, Zhao, Jessica et al. "Opacus: User-friendly differential privacy library in PyTorch." *arXiv preprint arXiv:2109.12298* (2021).