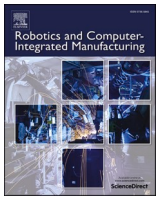
Robotics and Computer–Integrated Manufacturing 102 (2026) 103326



Contents lists available at ScienceDirect

Robotics and Computer-Integrated Manufacturing

journal homepage: www.elsevier.com/locate/rcim



A domain-informed learning framework for seam extraction in robotic welding: Generalizing to unseen seam topologies from unstructured workpiece types

Xianzhong Zhao [a,b], Haotian Liu [a], Zhaoqi Huang [a,*], Yuzhen He [a]

[a] College of Civil Engineering, Tongji University, Shanghai 200092, China
[b] Shanghai Qi Zhi Institute, Shanghai 200030, China

ARTICLE INFO

Keywords:
Robotic welding
Seam extraction
Domain-informed
Unstructured
Generalized
Real-world uncertainties
Path planning

ABSTRACT

Robust extraction of weld seams is critical for robotic welding in industries reliant on custom-fabricated or one-off components, where workpiece type and geometry are often unpredictable. This paper proposes a domain-informed generalized learning framework for seam extraction in robotic welding. The study highlights a domain-informed model generalization enhancement methodology that embeds welding expertise and operational logic of welders into the seam extraction learning approach. The domain-informed enhancement guides the model to prioritize welding positions and seam-relevant regions over specific seam pattern, and thereby enables the framework to be trained on only simple straight-seam patterns, which significantly reduces data preparation complexity and eliminates the need for model retraining or data reconstruction when faced with unfamiliar seam topologies. The study establishes a simulation-based data acquisition pipeline that incorporates critical environmental, system-related, and process-related real-world uncertainties for vision-based robotic welding systems, and enhances adaptability of the proposed seam extraction framework for industrial deployment. An adaptive weld path planning approach is developed for complex seams with multiple segments and significant curvature variations. A comprehensive experimental program was executed across both controlled laboratory conditions and real-world operational settings on four familiar cases and eight unseen scenarios, using two distinct robotic platforms. Test results reveal that the proposed domain-informed framework exhibits robust generalization capability on unseen workpiece types and seam patterns, adapts to diverse industrial settings, and maintains real-time performance for industrial deployment.

1. Introduction

Robotic welding has significantly enhanced productivity, quality, and safety in industries characterized by standardized high-volume production [1–2]. However, the adoption of robotic welding technology has been significantly slower in sectors reliant on custom-fabricated or one-off components, such as construction, and shipbuilding industries. In these fields, robotic welding has largely failed to replace manual techniques due to a lack of autonomous solutions capable of adapting to the highly variable geometries and unique project demands [3].

Conventional contact-based seam finding techniques, such as electrical or tactile touch sensing, effectively detect isolated straight seams but struggle with complex seam geometries and significantly increase cycle times for custom components [2,4]. As another option, welding robots could be programmed either offline [5] or via a teaching-playback method [6] to follow predetermined welding paths with fixed parameters. These robotic systems typically operate rigidly, and cannot adjust to unexpected changes or deviations [7–8]. By way of alternative, visual sensing-based seam extraction presents a fundamental solution by directly detecting seam geometry and patterns from the intended workpiece, based on which the robots further generate valid welding path and trajectory that accommodate real-world variations. It has been well recognized by both the scientific and technical communities that the seam extraction approach provides the critical flexibility required for widespread industrial applications [8–9].

Extensive research has attempted to extract weld seams following rule-based frameworks. Geng *et al* [8] designated the intersection lines

* Corresponding author.
E-mail address: huangzhaoqi@tongji.edu.cn (Z. Huang).

https://doi.org/10.1016/j.rcim.2026.103326
Received 14 October 2025; Received in revised form 27 March 2026; Accepted 2 May 2026
Available online 7 May 2026
0736-5845/© 2026 Elsevier Ltd. All rights are reserved, including those for text and data mining, AI training, and similar technologies.

between adjacent fitted planes as the target weld seams. An enhanced RANSAC algorithm was developed for reliable multi-plane fitting of 3D point clouds. Building upon the planar intersection approach, Yang *et al* [10] enhanced seam extraction reliability by incorporating an analysis of the distance between the fitted planes in the target workpieces. Zhang et al [11] proposed that weld seams generated localized variations in point distribution and developed a method to extract straight seams by identifying these density variations across scanned surfaces. Zhou et al [12] established unique geometry-based seam extraction rules for each investigated weld type (including V-butt, square-butt, and fillet welds in corner and pipe joints). Patil *et al* [13] proposed a centroid-based approach and extracted three-dimensional seams from noisy background. It is generally stated that the rule-based solutions rely rigidly on predefined logical rules designed by human experts. Even though they could be effective and efficient in specific cases, the rule-based approaches could hardly adapt to unexpected variations or handle complex scenarios.

Recent advances leverage learning-based approaches (i.e., using deep neural networks) for complex seam extraction tasks. Several approaches have focused on processing 2D image data. Liu et al [14] employed an L-CNN architecture to extract all observable straight seams from the input image, thereby forming the set of candidate seams that required additional processing to identify the intended seam. Cheng et al [15] developed deep learning-based solutions and extracted weld seams from real-time images in construction sites. Zhang et al [16] leveraged additional depth information (compared with 2D images) by applying learning-based approaches directly to RGB-D data, and effectively addressed the challenge of irregular geometry in seam extraction. More recently, researchers have begun to directly process 3D point cloud data for a more intrinsic geometric understanding. Ma et al [17] achieved seam extraction by developing optimized deep neural networks to map seam-relevant features from registered geometric CAD models onto their corresponding real-world point clouds. Song et al [18] developed a lightweight neural network and achieved efficient seam segmentation from welded joint point clouds.

These recent studies [14–18] have demonstrated that learning-based approaches offer superior flexibility over rule-based solutions. The learning-based approaches could easily handle geometric variations in weld seams, provided the workpiece types and seam patterns are familiar to the trained models. However, their seam extraction performance on unstructured workpieces or topologically distinct seam configurations outside the training data remains unestablished. In practice, this limitation forces the engineers to construct new qualifying training datasets and retrain the model for each new workpiece type or novel seam pattern, which is undoubtedly labor-intensive, time-consuming, and incompatible with real-time application. As such, a generalized seam extraction framework that enhances conventional learning-based solutions is critically needed for industrial robotic welding deployments.

Additionally, the majority of seam extraction research has primarily focused on well-controlled laboratory conditions. Real-world production variables critical for robotic welding deployment, including environmental variations (e.g., dynamic lighting and shadows), hardware limitations (e.g., sensor noise), as well as process-related uncertainties from robotic motion (e.g., occlusions caused by varying detection angles and orientations), have not been adequately addressed [8,10–18]. Consequently, a data acquisition pipeline that incorporates the effects of these real-world uncertainties is essential to enhance the adaptability and robustness of seam extraction for actual industrial deployment.

The first part of this paper proposes a generalized learning framework for extraction of unseen weld seams from unstructured workpiece types (Section 2). A simulation-based data acquisition pipeline that integrates critical real-world uncertainties for robotic welding deployment is developed. The resulting dataset enables comprehensive model training and enhances the adaptability of the framework to diverse industrial settings (Section 3). A domain-informed model generalization enhancement methodology is highlighted, which strategically embeds welding expertise and operational logic of welders into the seam extraction learning framework. The proposed robotic welding seam extraction framework thereby requires only minimal training data (comprising basic straight-seam patterns) and exhibits generalized performance on unseen scenarios without model retraining (Section 4). Additionally, an adaptive weld path planning approach is established for complex seams with multiple segments and significant curvature variations (Section 5). The second part of this paper accesses the performance of the proposed seam extraction framework through a comprehensive experimental program (Section 6). The generalization ability, adaptability, accuracy, and computational efficiency of the proposed seam extraction framework are thoroughly discussed (Section 7). The paper concludes with main contributions and future work of this study (Section 8).

2. The generalized seam extraction framework

Fig. 1 presents the proposed domain-informed learning framework for robotic welding seam extraction. The three key components of this study include (i) a domain-informed model enhancement methodology, (ii) a simulation-based data acquisition pipeline incorporating real-world uncertainties, and (iii) an adaptive weld path planning approach for complex seam patterns. Each component is designed to address distinct challenges in real-world robotic seam extraction. It is noted that Component (i) distinguishes the proposed learning framework from the conventional baseline method. Whereas, Components (ii) and (iii) are also employed to reinforce the conventional approach, thereby establishing a strong baseline framework for a fair comparison.

The domain-informed model enhancement methodology (Component (i)) reinforces the generalization capability of the seam extraction framework, and constitutes the core of the proposed framework. The study analyzed operational logic of welders and embedded welding expertise in robotic seam extraction by applying a domain-guided sparsity technique to the raw data. Comparative tests suggested that the domain-informed framework generalized to unseen workpiece types and novel seam topologies, whereas the baseline framework failed in unfamiliar scenarios. It must be pointed out that, the domain-informed enhancement enabled the framework to be trained on only simple straight-seam patterns, which significantly reduced data preparation complexity and eliminated the need for model retraining or data reconstruction when faced with unfamiliar scenarios.

Ablation studies demonstrated that the generalization enhancement is attributable to the proposed domain-informed methodology itself, rather than to the optimized PointNeXt model developed in this study for industrial deployment. In addition to the developed optimized PointNeXt model, the domain-informed enhancement consistently improved generalization performance across several other model architectures, indicating the broad applicability of the domain-informed approach for enhancing generalization in seam extraction tasks.

Both the simulation-based data acquisition pipeline (Component (ii)) and the adaptive weld path planning approach (Component (iii)) enhances adaptability of the seam extraction framework for industrial deployment. The simulation-based data acquisition pipeline (Component (ii)) explicitly incorporated real-world uncertainties critical for vision-based robotic welding systems, and thereby enhanced the adaptability for industrial deployment. The pipeline combined environmental uncertainties (dynamic lighting, shadow effects, and etc.), system-related uncertainties (sensor noise and resolutions), and process-related uncertainties (robotic motion and variations in measurement positions and orientations). The pipeline also enabled automated and efficient data generation.

The path planning approach ensured operational adaptability through employment of a local-curvature-variation based control and a principal-component-analysis based hybrid point sequence determination method. The established path planning approach generated optimal robotic welding paths for complex seam patterns, including those with

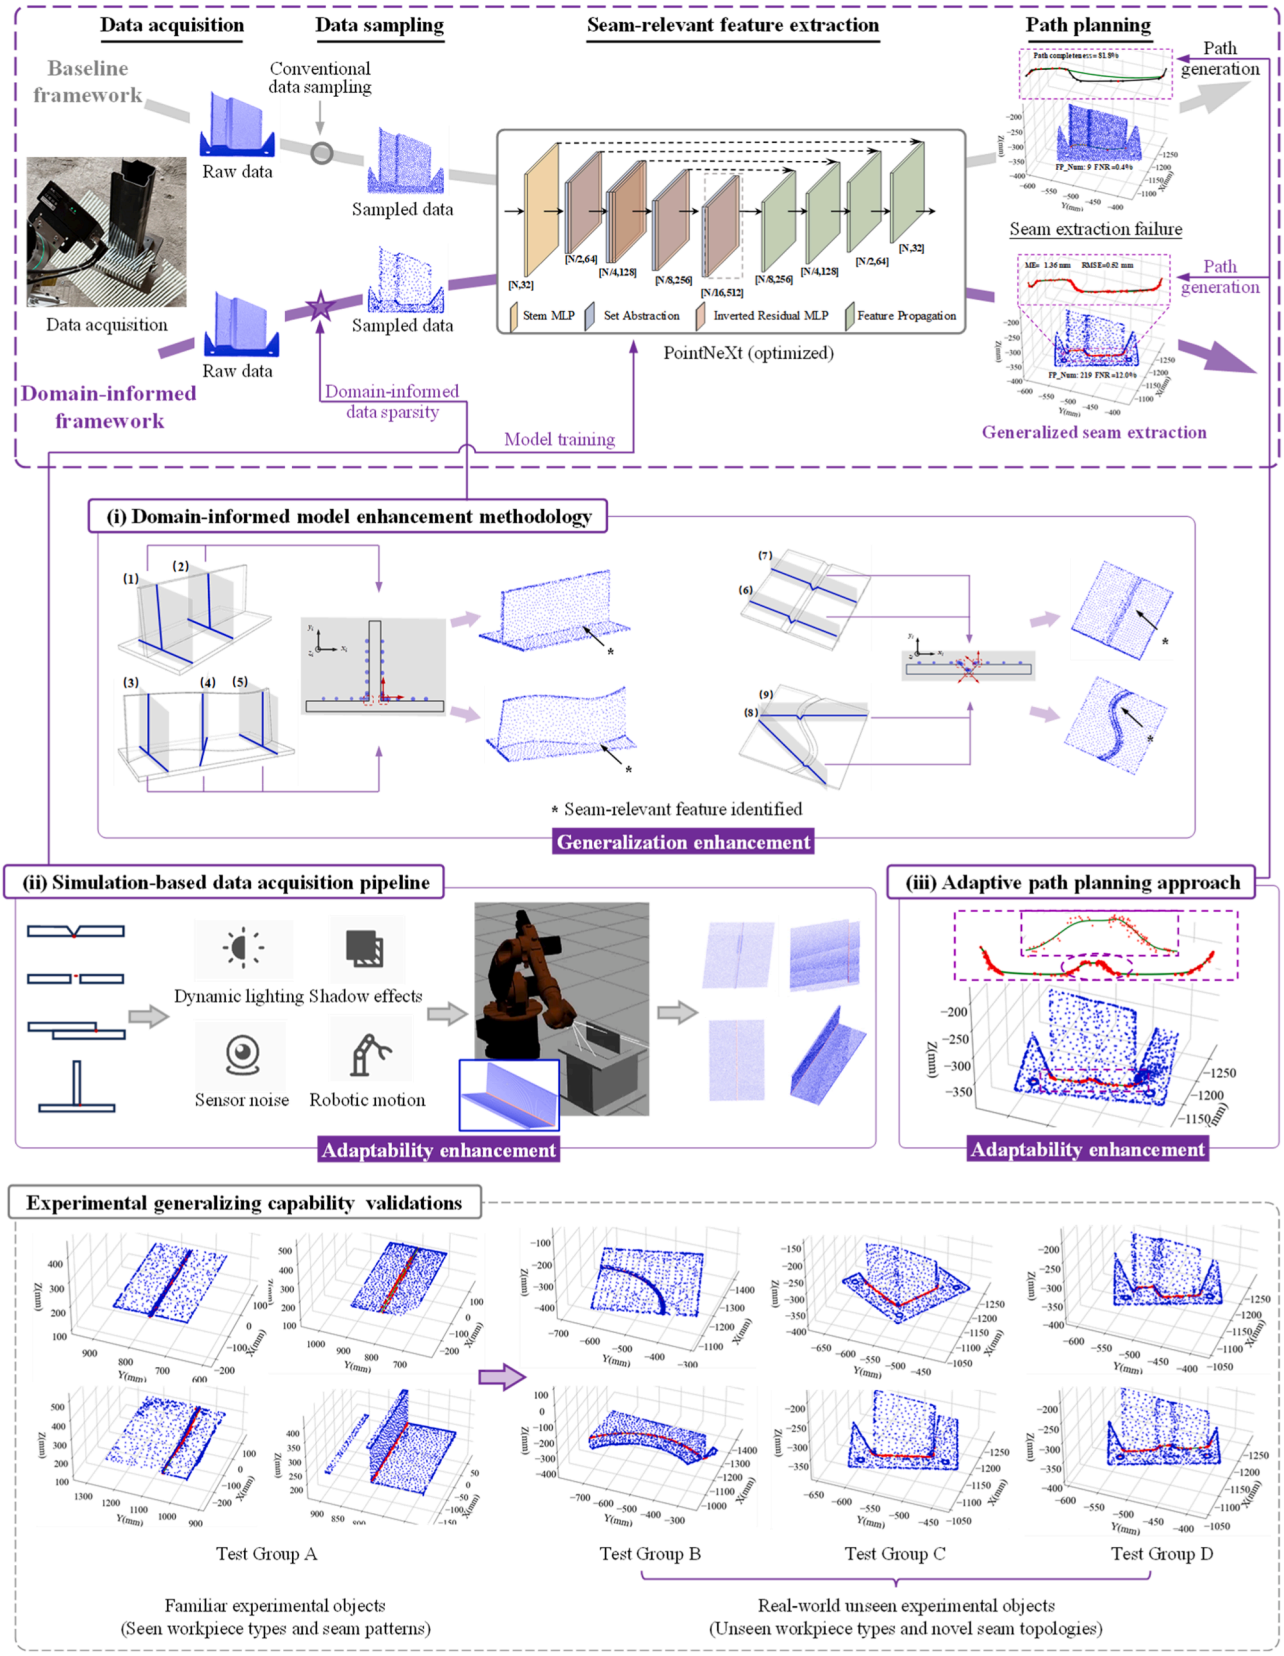


Fig. 1. Domain-informed seam extraction framework for robotic welding.

small intermediate segments and sharp corner geometries.

3. A simulation-based data acquisition pipeline for real-world uncertainties

This section develops a synthetic point cloud generation pipeline specifically designed for learning-based seam extraction in robotic welding applications (Fig. 2). The pipeline incorporated critical real-world uncertainties that affect data acquisition through a simulation-based approach, and facilitated automated and efficient large-scale training dataset generation.

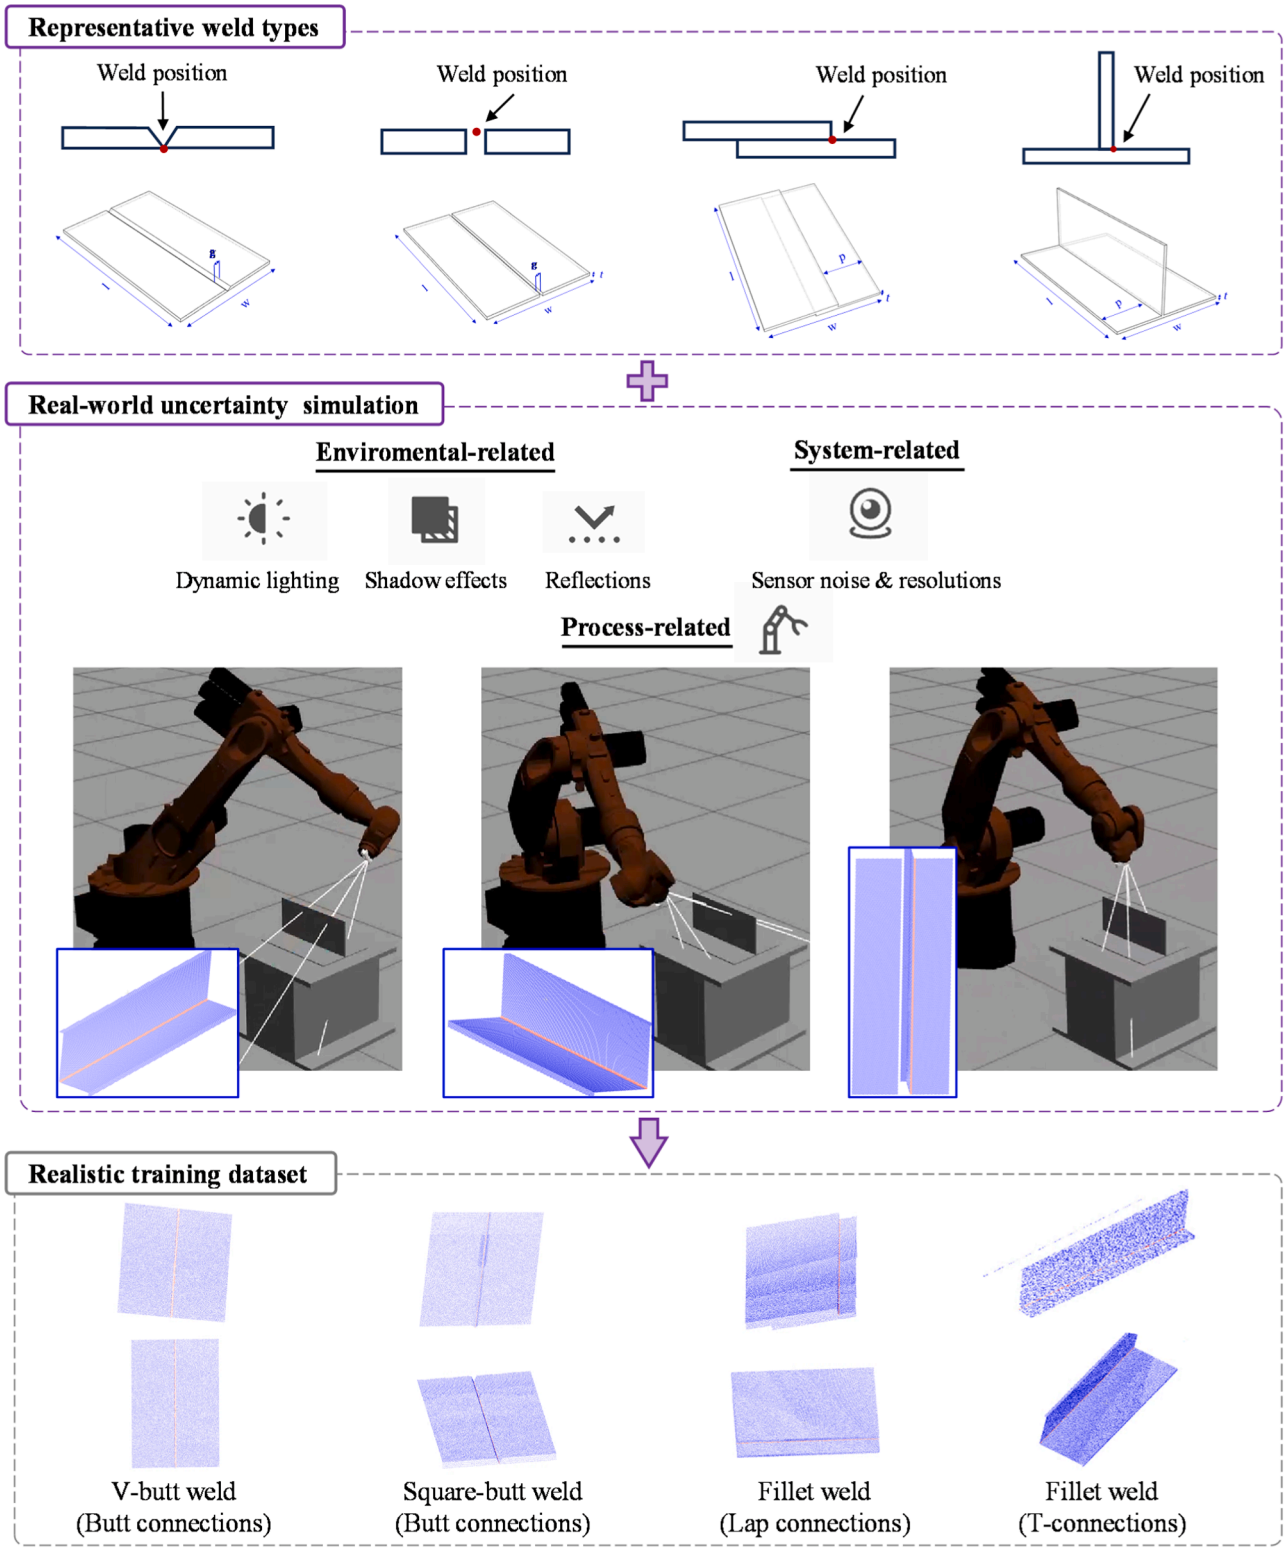


Fig. 2. Simulation-based data acquisition pipeline for real-world uncertainties and representative instances in the training dataset.

3.1. Representative weld types and geometric variations

The study considered the most prevalent weld types (i.e., butt and fillet welds) along with their common connection configurations (i.e., butt, lap, and T-connections). Fig. 2 presents the cross-sectional schematics and corresponding 3D diagrams for each weld configuration. Key geometric variables include connection length (l) and width (w), material thickness (t), butt weld groove width (g), and fillet weld location (p). The allowable variation ranges for these parameters are quantitatively defined in Table 1.

It is noticed that the training dataset consists exclusively of simple straight-seam patterns, and requires no highly-varied seam modeling, which significantly reduces the complexity of dataset preparation. Engineers with basic experience in 3D modeling and scripting can easily incorporate additional weld types (e.g., flare groove welds and etc.) into the pipeline. This requires only an extension of the parametric modeling component in Rhinoceros®, without modification to the physics-based simulation (Gazebo) or robotic control (MoveIt) components of the pipeline.

3.2. Critical real-world uncertainties under consideration

Considering common robotic welding and their vision system settings, this investigation incorporated critical real-world uncertainties that affect point cloud acquisition performance, including dynamic environmental conditions, system limitations, and process variability. For environmental uncertainties, the study considered dynamic lighting conditions (both natural and artificial light changes) and their secondary effects (shadows and reflections) since these factors directly interfere with the coded light patterns projected by the structured light sensor for point cloud acquisition. For system-related uncertainties, the study modeled sensor noises and sensor resolution constraints to account for their impact on point cloud reconstruction accuracy. Multiple noise levels and resolution parameters were systematically evaluated to simulate the performance range of commercial-grade structured light sensors under various operational conditions. For process-related uncertainties, the study intentionally simulated realistic data acquisition scenarios and generate training samples with controlled occlusion patterns for robust model training. The uncertainties modeled during synthetic point cloud generation, along with their influencing factors and real-world manifestations, are summarized in Table 2.

Table 2
Real-world uncertainties accounted in data acquisition.

Uncertainty type	Influencing factor	Manifestation
Environmental	Dynamic lighting conditions (Gazebo)	Natural light changes
		Artificial light variability
	Light interference (Gazebo)	Shadow effects
		Reflections
System-related	Sensor limitations (Gazebo)	Sensor noises
		Sensor resolutions
	Workpiece variations (Rhinoceros)	Weld types
		Geometric variations
Process-related	Sensor viewpoint shift (MoveIt)	Partial occlusions

3.3. The data acquisition pipeline and uncertainty modelling details

The data acquisition pipeline integrated Rhinoceros® for batch modeling of workpiece geometries, Gazebo [19] for physics-based simulation, and MoveIt [20] for robotic manipulation, and thereby achieved automated and efficient generation of large-scale training datasets for robotic welding seam extraction.

The scripting capabilities of Rhinoceros® enabled efficient batch generation of multiple geometric models, which were based on the workpiece diagrams displayed in Fig. 2 and incorporated the variables from Table 1.

Gazebo (the open-source 3D robotics simulator [19]) was employed to create interactive industrial environments that reproduced the intended real-world application scenario (see Fig. 2). The setup consisted of a 6-degree-of-freedom robotic arm with a depth-based sensor mounted in an Eye-in-Hand configuration. The sensor was simulated using an Intel RealSense D435 model and generated depth-based point clouds at 1280×720 resolution to match the physical sensor configuration. In order to simulate real-world data imperfections in sensor data, a zero-mean Gaussian noise ($\varepsilon \sim N(0, \sigma^2)$) with a variable standard deviation was applied. The standard deviation (σ) of the noise was randomly assigned between 0.1 mm and 0.2 mm for each simulation instance. This range was selected to model the typical magnitude of stochastic measurement errors encountered in industrial depth sensing.

Natural light variations were replicated by implementing three distinct intensity levels, including low, medium, and high lighting intensity. For low intensity, the diffuse (RGB) and ambient (RGB) parameters were set to (0.3, 0.3, 0.3) and (0.2, 0.2, 0.2). Medium intensity used values of (0.8, 0.8, 0.8) for diffuse and (0.4, 0.4, 0.4) for ambient. High intensity was defined by diffuse (1.0, 1.0, 1.0) and ambient (0.6, 0.6, 0.6). Artificial lighting was simulated using a directional light model originating from either the top (direction = (0, 0, −1)) or the side (e.g., direction = (−0.9, 0.1, −0.3)) in the world coordinate system. Shadows and reflections were naturally generated by built-in illumination models based on the configured light sources and material properties.

MoveIt (the open-source robotics manipulation framework [20]) was used to simulate both the point cloud acquisition process and corresponding robotic motion trajectories, and thereby replicated real-world operational conditions. Occlusions were introduced by systematically varying the sensor viewpoint via the MoveIt motion planner. For each workpiece, point clouds were acquired from multiple poses (e.g., top and side views) with small random perturbations. This approach intentionally occludes portions of the weld seam and surrounding geometry in a series of scans. Representative viewpoint variations and resulting occlusion patterns are illustrated in Fig. 2.

Table 1
Statistics of the training dataset.

	V-butt weld (Butt connection)	Square-butt weld (Butt connection)	Fillet weld (Lap connection)	Fillet weld (T-connection)
Schematic diagram				
Length l (mm)	300–400	300–400	300–400	300–400
Width w (mm)	200–400	200–400	200–400	200–400
Element thickness t (mm)	6–10	2–6	2–8	2–8
Groove width g (mm)	–	10–18	2–8	–
Seam location p (mm)	60–280	–	–	60–280
no. of data items	600	560	750	580
FNR*	0.8~2.4 %	0.7~4.0 %	1.2~3.8 %	0.6~5.9 %

* Feature-point-to-none-feature-point ratio.

3.4. Training dataset

A large-scale realistic training dataset was created utilizing the developed automated data generation platform. As summarized in Table 1, totally 2490 instances are involved in the dataset. Each weld configuration subclass contains between 560 and 750 unique samples. Fig. 2 shows a pair of representative samples from the dataset for each weld type under consideration. It is noted that each instance in the training set has implicitly incorporated the effects of various real-world uncertainties encountered in actual deployments. Even though not explicitly shown here, quantitative analysis has been conducted to characterize the discrepancy between synthetic and real-world sensor data. Analysis demonstrates that while a reasonable sim-to-real gap exists, our simulation pipeline accurately reproduces the order of magnitude of geometric noise found in industrial settings, and preserves the critical local geometric properties essential for effective model training required in this study. The experimental results presented in the following sections also demonstrate that the training dataset enhances seam extraction adaptability to uncontrolled industrial settings.

4. Domain-informed extraction for novel seam topologies from unstructured workpiece types

Conventional learning-based seam extraction methods are highly dependent on their training data. They commonly yields satisfactory performance on familiar workpiece types and seam topologies and fails to generalize to, or exhibits significant degradation on, unseen workpiece types or novel seam topologies. This reliance on data is incompatible with real-time industrial applications. The necessity to reconstruct the training dataset and retrain the model for every unfamiliar scenario is prohibitively labor-intensive and time-consuming.

This section proposes a domain-informed model enhancement methodology that realizes generalized seam extraction on arbitrary seam topologies and unfamiliar workpiece types, without requiring their explicit inclusion in the training data (Section 4.1). A specialized point-wise model was developed that balanced the seam extraction accuracy with computational efficiency (Section 4.2), so as to meet the demands of real-time industrial deployment.

4.1. The domain-informed generalization enhancement methodology

This study analyzed welding expertise and recognized that regulations in production-related documents (e.g., EN 12345 [21]) and design-related specifications (e.g., EN 1993-1-8 [22]), and symbolic representation on technical drawings for welded components (e.g., ISO 2553 [23]), uniformly prioritize weld types and welding positions over seam topology or weld paths. Interviews with certified welders confirmed this prioritization aligns with their operational logic, which

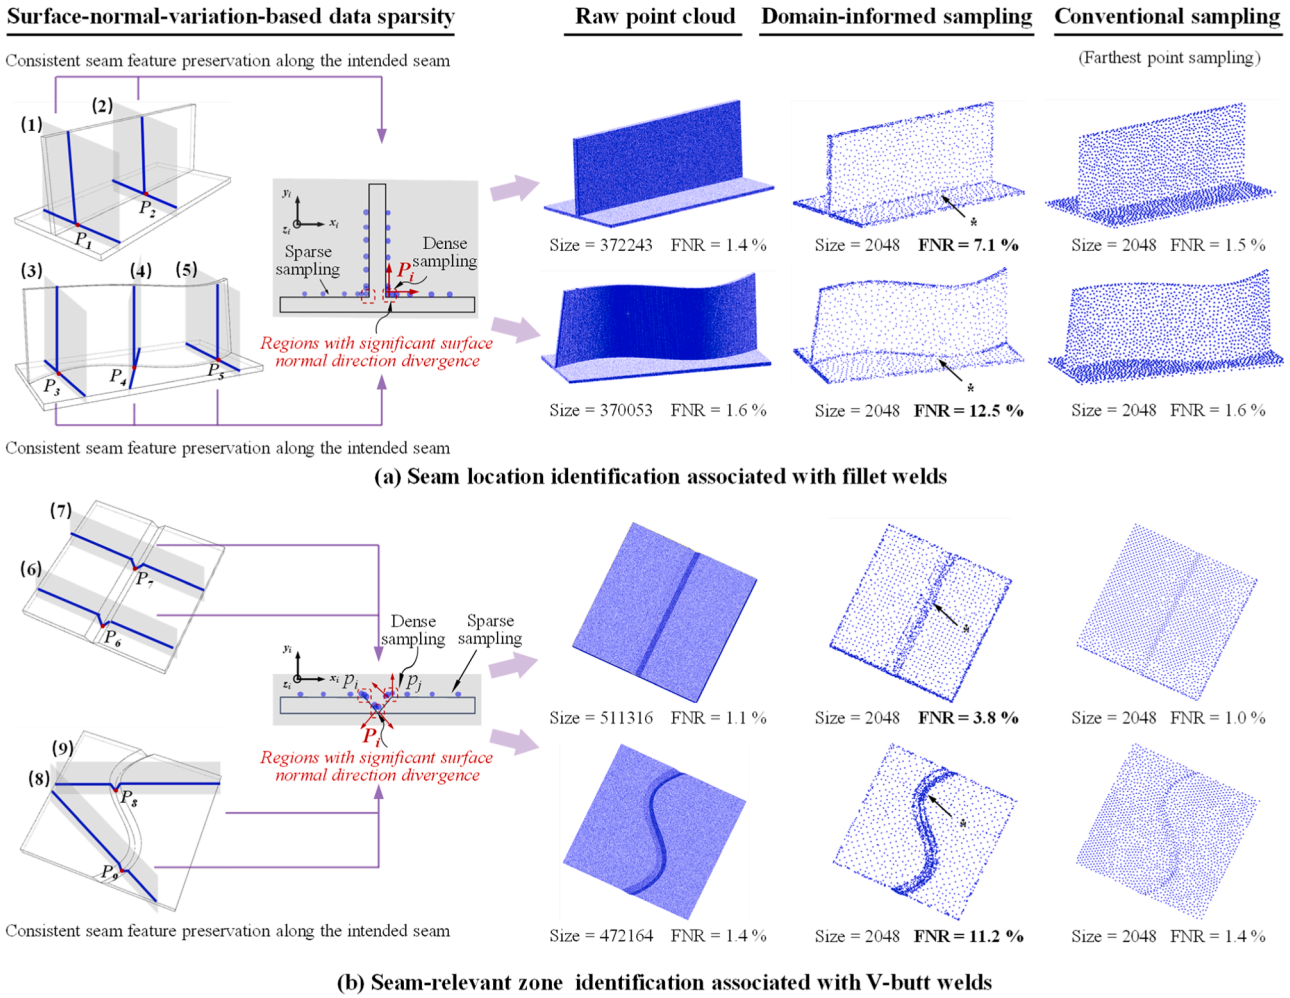


Fig. 3. Domain-informed model enhancement methodology for seam extraction.

treats any arbitrary seam path as a sequence of infinitesimal straight segments with smooth transitions once the weld type and position are fixed. These observations motivated a generalization enhancement methodology that embeds this operational logic into the robotic welding seam extraction learning framework.

The basis for the model enhancement was established by analyzing a broad range of welding sketches from technical documents [21,23], which revealed that weld positions are typically characterized by a high directional divergence of surface normals. As Fig. 3 illustrates for fillet welds (a) and V-butt welds (b), this occurs because a weld seam is usually defined by the intersection of two surfaces. At any point P_i along the intended seam, the paired surface normals inherently exhibit significant directional divergence, irrespective of the global seam topology or curvature.

The model enhancement methodology leveraged this local geometric divergence as the foundation for a domain-informed sparsity technique (see Fig. 3). The methodology analyzed the graded changes in surface normal directions, and applied dense sampling in regions with significant directional divergence so as to locate weld positions and identify seam-relevant features. Whereas the sampling density was deliberately reduced in areas of marginal divergence (e.g., flat or uniformly curved surfaces) to prevent the over-representation of non-seam regions and enhance computational efficiency. It must be mentioned that, the methodology detects a fundamental geometric property (i.e., rate of directional changes in surface normals) rather than specific seam shapes, and thereby remains equally effective for both straight and arbitrarily curved seam patterns.

Practically, the surface normal was estimated at each point by applying principal component analysis (PCA) to its k-nearest neighbors $\{p_1, \cdots, p_k\}$ ($k = 15$ in this study). Provided that $\mu = \frac{1}{k}\sum_{i=1}^{k} p_i$ denotes the centroid and $C = \frac{1}{k}\sum_{i=1}^{k}(p_i - \mu)(p_i - \mu)^T$ denotes a 3×3 covariance matrix, the surface normal for a query point p_q was established by solving

$$Cv_j = \lambda_j v_j \tag{1}$$

for $j \in \{0,1,2\}$, with eigenvalues sorted as $\lambda_0 \le \lambda_1 \le \lambda_2$. The eigenvector v_0 associated with the smallest eigenvalue λ_0 was taken as the estimated surface normal at the query point. The corresponding unit normal is given by (n_x, n_y, n_z).

The estimated surface normals were then mapped to a spherical coordinate system for the subsequent identification and sampling of seam-relevant points. Each unit normal was converted to spherical coordinates as:

$$\theta = arccos(n_z) \in [0, \pi], \qquad \Phi = atan2(n_y, n_x) \in (-\pi, \pi] \tag{2}$$

The continuous (θ, ϕ) space was further partitioned into $m \times m$ spherical units (note that a total of m^2 spherical units were obtained), with each unit containing a set of points with similar normal directions.

The local curvature measure at query point p_q was approximated using the eigenvalues ($\lambda_0 \le \lambda_1 \le \lambda_2$) from the PCA of it local neighborhood as $\sigma_q = \lambda_0/(\lambda_0+\lambda_1+\lambda_2)$. The curvature intensity for the set of points (S_j) within the j-th spherical unit of interest (U_j) was then established as the average of their local curvature measures:

$$\mu_j = \frac{1}{|S_j|}\sum_{p_i \in S_j} \sigma_i \tag{3}$$

To prioritize seam-relevant regions (characterized by high curvature), an exponential weighting scheme was applied which assigned higher sampling weights (i.e., dense sampling) to spherical units with greater curvature intensity. The sampling weight ω_j for a unit U_j was defined as:

$$\omega_j = exp\left(\mu_j\right) \Big/ \sum_{k=1}^{m^2} exp(\mu_k) \tag{4}$$

Given a total sampling budget of N_s points, the number of sampling points N_j allocated to the j-th spherical unit U_j is proportional to its sampling weight as $N_j = \max(1, [\omega_j N_s])$.

Fig. 3 compares the workpiece point clouds processed by the proposed domain-informed sparsity technique and a conventional farthest point sampling method (which is a popular technique in hierarchical point cloud learning [25] due to its ability to ensure excellent spatial coverage while reducing point density). As a general statement, the domain-informed sparsity directly identifies welding positions across all scenarios, regardless of weld type, workpiece, or seam pattern. Consistent feature preservation is achieved along the entire target seam. The seam extraction performance is quantified through the feature-to-non-feature-point ratio (FNR) defined as:

$$\text{FNR} = \frac{N_f}{N - N_f} \times 100\% \tag{5}$$

where N_f denotes the count of seam-relevant feature points labeled in the workpiece point clound of interest, and N is the total number of points in the cloud.

It is evident that a substantial increase in the FNR is achieved, rising from 1.4 ~ 1.6 % in the raw point cloud to 7.1 ~ 12.5 % in the sampled point cloud for the T-connection fillet weld (Fig. 3(a)), and from 1.1 ~ 1.4 % to 3.8 ~ 11.2 % for the V-butt weld (Fig. 3(b)). In contrast, the conventional farthest point sampling technique resulted in only marginal changes in the FNR. Although not explicitly shown here, other conventional sampling methods, including ramdom sampling and uniform sampling, yielded similar results and failed to produce a noticeable increase in the FNR.

It is noteworthy that for specific weld types like V-butt welds, the methodology identified not only the primary welding position (i.e., the root of the groove, P_i) but also other seam-relevant locations (e.g., the top edges of the groove, p_i and p_j) due to their inherent surface normal divergence (see Fig. 3(b)). This potential ambiguity is resolved during the training phase, wherein the framework learns to differentiate the true weld position from disturbing seam-relevant regions through weld-type-specific labels provided in the training data.

Ablation studies were conducted to evaluate the impact of the proposed domain-informed methodology on model performance. Instead of only considering a single specific model, comprehensive analysis across multiple well-known point-wise architectures (i.e., PointNet [24], PointNet++ [25], DeepGCN [26], DGCNN [27], and PointNeXt [28]) were performed to assess whether the proposed approach led to performance gains across diverse models.

Table 3 compares the model performance (Ins.mIoU) of the domain-informed approaches with their baselines. It is evident that the proposed domain-informed methodology consistently improves performance across all evaluated models, with Ins.mIoU gains ranging from 12.8 % (PointNeXt) to 87.9 % (PointNet). The ablation study confirms the broad applicability of our domain-informed approach for enhancing generalization in seam extraction tasks.

Table 3
Performance comparison (Ins.mIoU) of domain-informed approaches and their baselines.

Network	Baseline	Domain-informed	Performance enhancement
PointNet (2017) [24]	47.87	89.97	+ 87.9 %
DGCNN (2019) [27]	58.42	82.13	+ 40.6 %
DeepGCN (2019) [26]	63.02	87.93	+ 39.5 %
PointNet++ (2017) [25]	78.37	88.60	+ 13.0 %
PointNeXt (2022) [28]	80.15	90.44	+ 12.8 %
PointNeXt (optimized *)	81.48	90.82	+ 11.5 %

* The PointNeXt architecture was selected based on the ablation study results. An optimized version was subsequently developed and employed in this study.

4.2. Optimized model development

This section develops a high-performance robotic welding seam extraction model for industrial deployment. According to the ablation test results (Table 3), PointNeXt [28] is considered in this study as it outperforms all the other models. Optimization is implemented through (i) model scaling and (ii) loss function refinement.

Fig. 4 presents the PointNeXt architecture along with the specific scaling strategies employed in this study. In general, scaling to the width and depth of the model are implemented simultaneously so as to achieve an optimal balance between seam extraction accuracy and computational efficiency considering the industrial requirements for real-time applications. Width scaling is accomplished by changing the channel size of the StemMLP layer (C), while depth scaling is achieved through changing to the number of InvResMLP blocks (B). Table 4 quantifies the trade-offs induced by different scaling strategies, including Class-wise mean Intersection over Union (Cls.mIoU), Instance mean Intersection over Union (Ins.mIoU), Throughout, Floating-point operations (FLOPs), and model Parameters. The configuration with $C = 32$ and B = [1,1,2,1] achieved the highest Cls.mIoU of 90.67 % and the second-highest Ins.mIoU of 90.82 % utilizing only 4.3 MB of parameters. Whereas, another candidate configuration with $C = 160$ and $B = 0$ yielded a negligible gain of only 0.04 % in Ins.mIoU at a substantial cost of approximately 10.7 times more FLOPs and 5.2 times more parameters. Considering the requirements for real-time robotic welding deployment, the configuration with $C = 32$ and B = [1,1,2,1] was adopted as the optimal balanced choice for industrial applications.

Instead of conventional Cross-Entropy Loss and Focal Loss, this study adopts the Poly Loss formulation from Leng et al [29] as it allows flexible control over how the loss penalizes different confidence levels (Eq. (6)). The hyperparameters ε and γ control such behavior, with p_t representing the predicted probability for the true class. A grid search on the finalized PointNeXt architecture ($C = 32$, B = [1,1,2,1]) determined the optimal configuration to be $\varepsilon = 0$ and $\gamma = 1$, which yielded the highest performance as shown in Fig. 5.

$$L_{poly} = (1+\varepsilon)(1-p_t)^{1+\gamma} + \sum_{j=2}^{\infty}\frac{1}{j}(1-p_t)^{j+\gamma} \tag{6}$$

It is pointed out that, besides PointNeXt [28], other point-wise models (including but not limited to, PointNet [24], PointNet++ [25], DeepGCN [26], and DGCNN [27]) are likely to achieve comparable seam extraction performance when enhanced by the proposed domain-informed methodology as suggested by the ablation test results (Table 3). Given that the performance gains lare argely attributable to the proposed domain-informed methodology rather than the specific model architecture, researchers and engineers are also encouraged to adopt future state-of-the-art models for robotic welding seam extraction tasks.

5. An adaptive weld path planning approach for complex seam patterns

This investigation develops an adaptive weld path planning approach for real-world complex seams featuring (i) small intermediate segments, and / or (ii) sharp corners near endpoint regions. Non-Uniform Rational B-Spline (NURBS) curves were employed for weld path generation. Using the set of feature points $\{P_1, \cdots, P_n\}$ extracted by the seam extraction framework as control points, a weld path can be generated as a cubic NURBS curve:

$$C(u) = \frac{\sum_{i=1}^{n}\omega_i P_i N_{i,3}(u)}{\sum_{i=1}^{n}\omega_i N_{i,3}(u)} \tag{7}$$

where $N_{i,3}(u)$ are the cubic B-spline basis functions, and ω_i is the NURBS weight corresponding to the i-th control point (P_i). Conventional weld path generation typically uses fixed control point locations with uniform weights.

In order to adaptively modify the control point locations, this study applied a one-dimensional Gaussian filter to the original control point set $\{P_1, \cdots, P_n\}$:

$$P'_i = \frac{\sum_{j=i-w}^{i+w} P_j \exp\left(-\frac{(j-i)^2}{2\sigma^2}\right)}{\sum_{j=i-w}^{i+w}\exp\left(-\frac{(j-i)^2}{2\sigma^2}\right)} \tag{8}$$

Here, $P_i^{'}$ is the modified coordinate of the i-th control point (P_i)

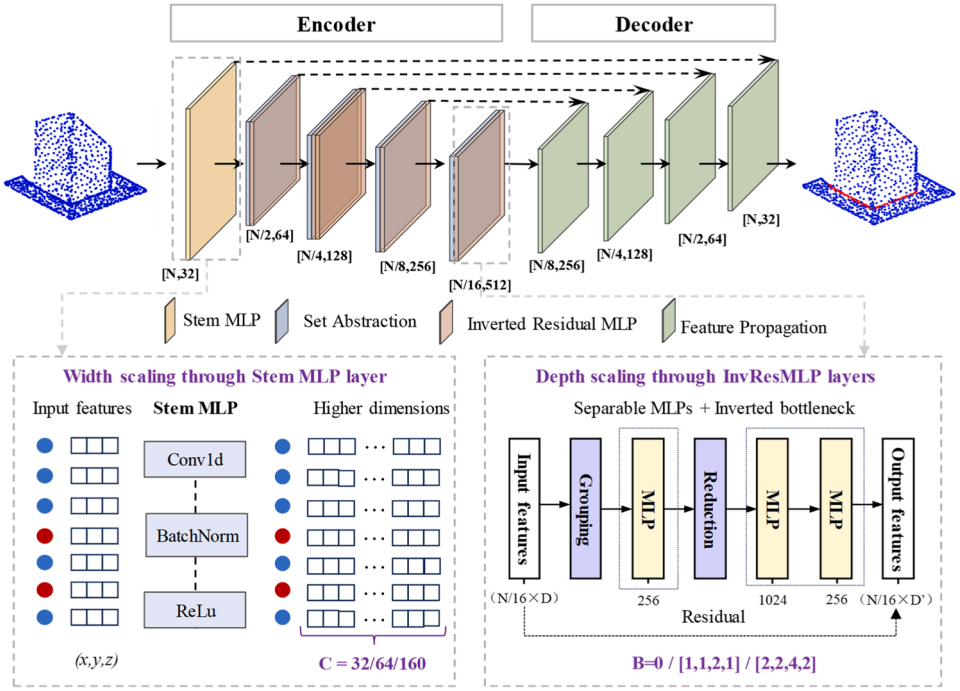
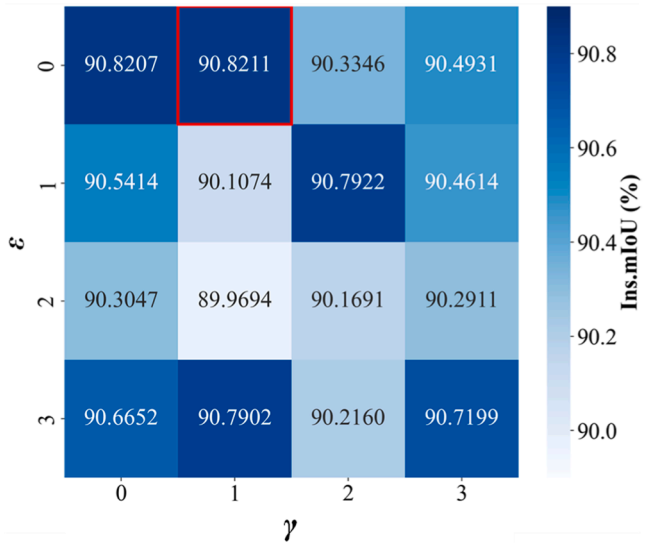


Fig. 4. PointNeXt architecture and model scaling for seam extraction.

Table 4
PointNeXt network performance across diverse model scaling strategies.

Depth scaling	Width scaling	Ins.mIoU (%)	Cls.mIoU (%)	Throughput (ins./sec.)	FLOPs (G)	Params (MB)
B = 0	$C = 32$	90.44	90.24	463.5	4.5	1.0
	$C = 64$	90.45	90.22	249.0	17.8	3.7
	$C = 160$	90.86	90.62	76.8	110.1	22.5
B = [1,1,2,1]	$C = 32$	90.82	90.67	254.0	10.3	4.3
	$C = 64$	90.55	90.32	123.9	40.8	16.9
	$C = 160$	90.77	90.56	34.0	253.2	104.6
B = [2,4,2,2]	$C = 32$	90.37	90.11	171.4	16.1	7.6
	$C = 64$	90.51	90.28	80.7	63.8	30.0
	$C = 160$	90.57	90.39	21.9	396.3	186.7



$$L_{poly} = (1+\varepsilon)(1-p_t)^{1+\gamma} + \sum_{j=2}^{\infty}\frac{1}{j}(1-p_t)^{j+\gamma}$$

Fig. 5. Loss function refinement.

obtained as a weighted average of the original point and its neighbors within a specified window. The parameter σ is the standard deviation of the Gaussian kernel that controls the degree of smoothing. The neighborhood range was defined by a half-window width of $w = 4\sigma$. Iterations were performed over the index range $[i\text{-}w,\ i + w]$.

The Gaussian filter scale σ was determined based on the overall control point density of the entire path. For the i-th control point P_i, the average Euclidean distance to its k nearest neighboring control points (with $k = 7$) was calculated as:

$$\overline{d}_i = \frac{1}{k}\sum_{j=1}^{k} \| P_i - P_j^{kNN} \| \tag{9}$$

A global density-based length scale was then established as the average nearest-neighbor distance of all control points:

$$\sigma_{len} = \frac{1}{n}\sum_{i=1}^{n}\overline{d}_i \tag{10}$$

The average sequential spacing in the spatial domain was established as:

$$\Delta s = \frac{1}{n-1}\sum_{i=1}^{n-1} \| P_{i+1} - P_i \| \tag{11}$$

The final Gaussian filter scale (in sample units) was obtained by scaling the density-based length scale by the average sequential spacing:

$$\sigma = \alpha \cdot \sigma_{len} / \Delta s \tag{12}$$

where α is a global smoothing scale adjustment coefficient (taken between 1.0–1.6 in this study).

The adaptability of the weld path planning approach was further enhanced by assigning a dynamic weight to each control point based on the local curvature of the weld path. For internal modified control points P_i' (where $i = 2, \ldots, n\text{-}1$), a local curvature measure was established using the triple of consecutive points $P_{i\text{-}1}'$, P_i', and P_{i+1}':

$$\kappa_i = \frac{2\times \| (P_i - P_{i-1}) \times (P_{i+1} - P_i) \|}{\| P_i - P_{i-1} \| \cdot \| P_{i+1} - P_i \| \cdot \| P_{i+1} - P_{i-1} \|} \tag{13}$$

After computing the raw curvature κ_i for all internal points, the values were normalized to a standard range:

$$\widetilde{\kappa}_i = \frac{\kappa_i - \min(\kappa)}{\max(\kappa) - \min(\kappa)} \in [0,1] \tag{14}$$

The adaptive weight ω_i for the i-th control point ($i = 2, \ldots, n\text{-}1$) was then defined as:

$$\omega_i = 1 + \beta \cdot \widetilde{\kappa}_i \tag{15}$$

Here, a curvature sensitivity coefficient $\beta = 2$ was employed in this study to amplify the weight increment in high-curvature regions, and thereby pulls the NURBS path closer to the control points at corners and along small intermediate segments.

Fig. 6 illustrates the path planning results for an intended seam with small intermediate segments and critical geometric features. It can be noticed that conventional path generation approaches using fixed control points with uniform weights struggle with such geometries. They either over-smooth the path and erase critical small-scale details when using a large σ (see Fig. 6(b)); or overfit local noise and lead to curvature discontinuities with a small σ (see Fig. 6(c)). The adaptive solution dynamically adjusts both the positions of control points and their corresponding weights based on local curvature variations, and thereby achieves an optimal balance between path smoothness and geometric detail preservation (see Fig. 6(d)).

It should be noted that valid weld path generation is dependent upon establishing the correct sequential order of the control points. In complex scenarios with sharp corners near weld seam end points (see a representative case in Fig. 7(a)), the path generation algorithm may incorrectly connect distant points and produce an invalid welding path for small seam segments due to significant curvature variations (see invalid path generation results in Fig. 7(b)). To address such challenges, a principal-component-analysis (PCA) based hybrid point sequence determination strategy was developed. As schematically illustrated in Fig. 7(c), the dominant direction of the entire feature point set was first established through global PCA. Critical regions were then segmented into local windows, where PCA was performed independently. The curvature continuity and point distances were subsequently checked. The point sequence would be adaptively re-ordered if necessary (for instance, when distort curvature was detected). As shown in Fig. 7(b), reliable weld paths are generated when using the adaptive path planning

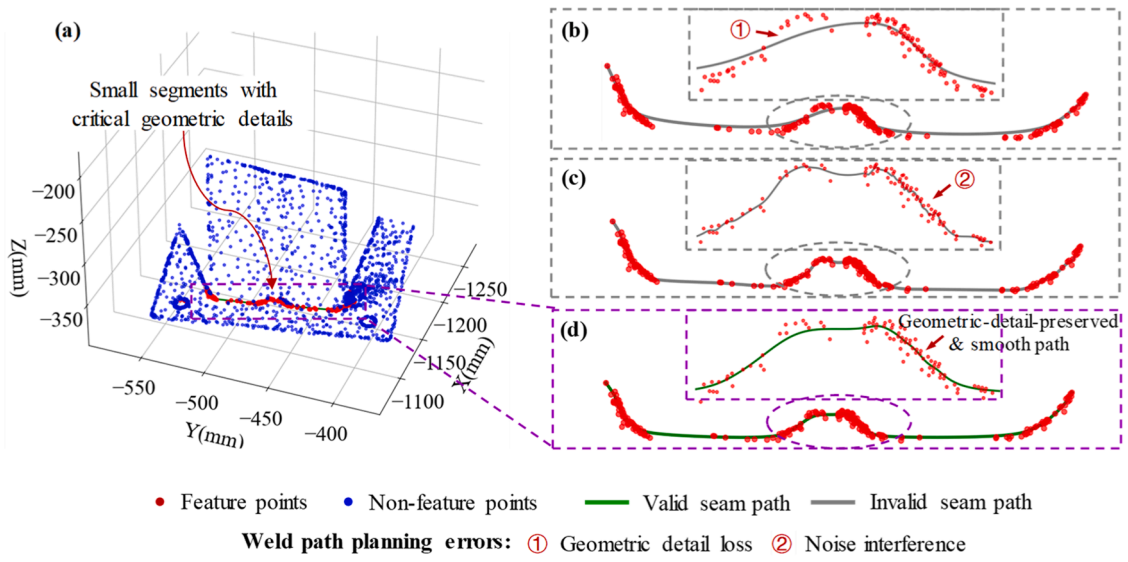


Fig. 6. A curvature-based weld path planning approach. (a) A representative seam with complex geometry. (b) An invalid weld path where geometric detail is lost. (c) An invalid weld path resulting from overfitting to local noise. (d) A valid weld path that balances between path smoothness and geometric detail preservation.

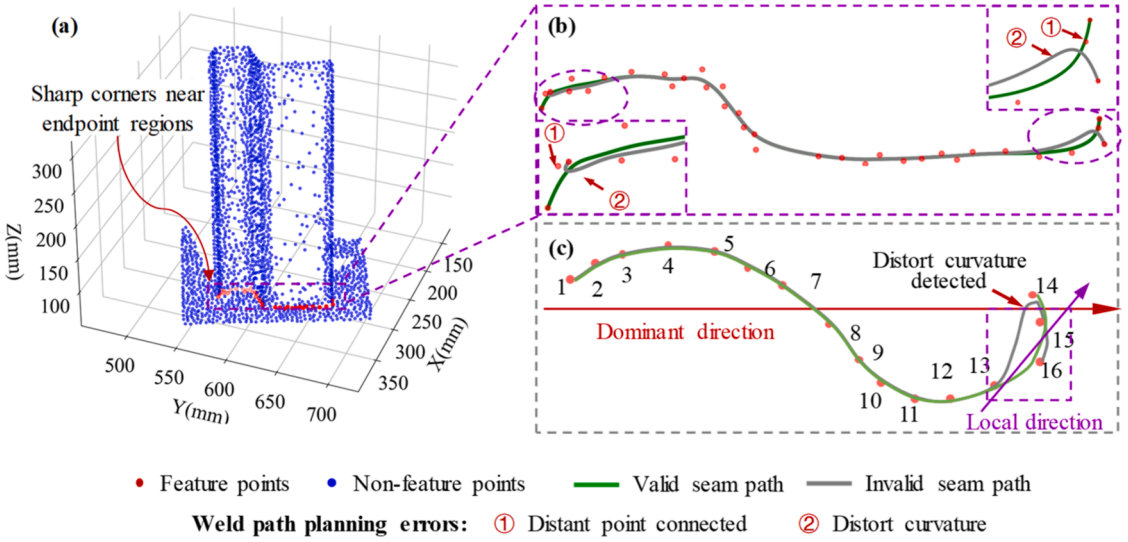


Fig. 7. A principal-component-analysis (PCA) based hybrid point sequence determination approach for weld path planning. (a) A representative case. (b) Comparison of a valid path with correctly ordered control points and an invalid distorted path resulting from the incorrect connection of endpoints. (c) Schematic diagram for the hybrid point sequence determination approach.

approach in conjunction with the hybrid point sequence determination method.

6. Experiments and results

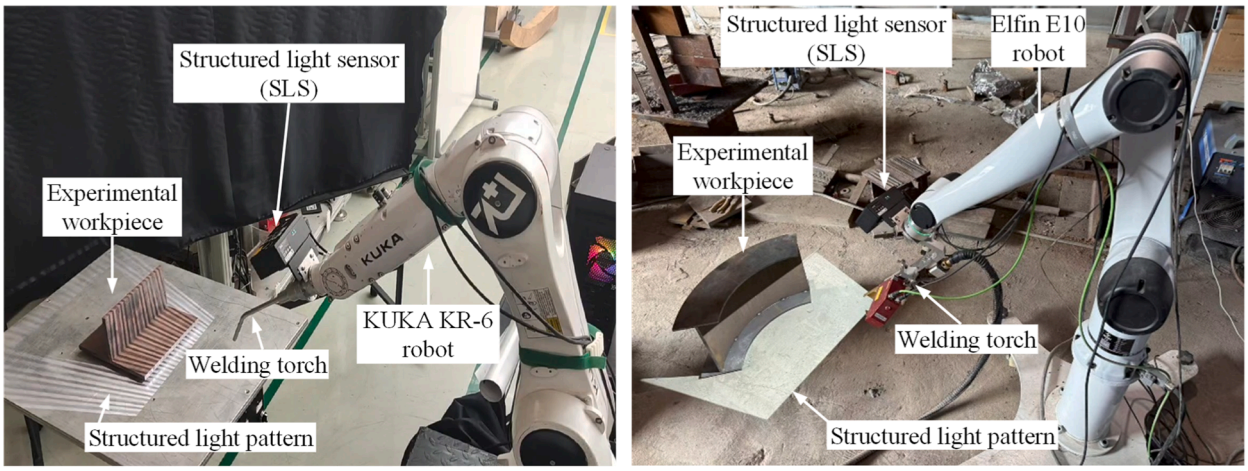
6.1. Experimental settings

Experiments were conducted on two robotic systems under distinct environmental conditions to assess real-world deployment robustness of the proposed seam extraction framework. Fig. 8(a) depicts the controlled laboratory setup in which a KUKA KR-6 welding robot served as the experimental platform. The environment had uniform indoor LED lighting and avoided uneven illumination or glare that could affect the measurements. The experimental workpieces were placed on a table under carefully optimized conditions to eliminate observable shadows from surrounding objects, such that potential disruptions to the vision-based sensing were prevented.

Fig. 8(b) documents the experimental setup in an uncontrolled real-world factory environment, where an Elfin E10 robot was employed as the operational platform. The workpieces were placed directly on rough concrete floor. The factory relied on natural light and employed uneven artificial lighting. Observable shadows from surrounding objects were cast onto the target workpieces, which introduced potential noise for the vision-based measurements.

Both robotic systems primarily consisted of a 6-degree-of-freedom robotic arm, a structured light sensor, and an industrial computer. The robots were equipped with a welding torch as their end effector. The structured light sensor was fixed mounted to the robotic arm using the Eye-in-Hand configuration. Even though not explicitly shown here, the structured light sensor was calibrated to ensure alignment with the robot. Key specifications of the sensor are presented in Fig. 8.

Prior to each test, the robot was positioned to ensure that the target seams remained entirely within the field of view of the structured light sensor, with the sensor maintained at an optimal working distance from



(a) Controlled laboratory environment **(b) Uncontrolled real-world factory environment**

Structured light sensor (SLS) specifications

Working distance (mm)	Near field of view (mm)	Far field of view (mm)	Resolution (Pixels)	Calibration accuracy (mm)	Acquisition time (s)
300-600	220×150 @0.5	440×300 @1.0	1280×1024	0.1	0.6-1.1

Fig. 8. Robotic welding systems and experimental settings.

the target. As shown in Figs. 8(a) and (b), during the tests, the sensor projected a structured light pattern onto the experimental workpiece and generated high-precision point cloud.

6.2. Model training

The training used a workstation equipped with an AMD EPYC™ 7452 processor (32 cores) and four NVIDIA® GeForce RTX™ 3090 GPUs (24GB VRAM) using PyTorch on a Linux OS. The Adam optimizer with a learning rate of 1×10^{-3} and weight decay of 1×10^{-4} were employed. The training spanned 600 epochs, each comprising 63 mini-batches with a size of 32. It is mentioned that these settings were optimized through preliminary experiments. A batch size of 32 was selected to fully leverage GPU memory and to avoid out-of-memory errors. The 600-epoch schedule ensured complete model convergence, with additional training providing no further benefit.

The training processes of both the domain-informed and baseline frameworks are visualized through monitoring curves, including mIoU (mean Intersection over Union, Fig. 9(a)), Loss (Fig. 9(b)), Precision (Fig. 9(c)) and Recall (Fig. 9(d)). Both frameworks demonstrate robust training processes with progressive improvement in key metrics (mIoU, Precision, and Recall) alongside smooth loss convergence. However, the baseline framework exhibits less steady growth in mIoU and reveals a noticeable Precision-Recall trade-off (achieving 93 % Precision but only 85 % Recall in final evaluation). In contract, the domain-informed framework exhibits consistently superior performance across all metrics. More stable and steady growth in mIoU is observed. A balanced improvement in both Precision and Recall is achieved (95 % and 94 % respectively) with marginal trade-offs. Additionally, the zoom-in view in Fig. 9(b) suggests that the domain-informed framework exhibits fewer and smaller spikes in the training loss curve. Overall, the domain-informed framework exhibits stronger convergence behavior compared to the baseline.

6.3. Experimental objects and test scenarios

The proposed seam extraction framework was evaluated through four experimental object categories. Generally, Test Groups A & B considered simple cases with single straight or curved seam (see Figs. 10,11), whereas Test Groups C &D involved real-world structural joints with complex seam geometries (Groups C & D, Figs. 12,13). Test Group C evaluated a 100 × 100 hot-rolled square hollow section requiring continuous welds along all four edges. Test Group D assessed challenging cold-formed thin-walled open sections with all-around weld requirements. For each real-world joint in Groups C &D, the robotic welding system encountered complex seam patterns combining multiple straight segments and small-radius curves from various observation angles.

It is pointed out that, the model were only trained on the workpiece types and seam patterns in Test Group A. The framework's performance on Group A reflects its ability to handle familiar cases. For all the other test groups (Groups B, C, and D), however, the workpiece types and seam topologies were unseen during training. The performance of the framework on other test groups evaluates its generalization capability to new workpiece types and novel seam topologies, which is critical for real-world deployment where inputs are unpredictable.

In total, the study comprised twelve test scenarios across nine distinct experimental objects. For each test scenario, the ground truth seam was established using a lead-through teaching approach. The robot arm was manually guided along the intended weld seam, with the coordinates of multiple intermediate points (i.e., taught points) along the seam recorded by the robotic welding system. Care was taken to ensure precise alignment of the torch tip (tool control point) with both the start and end points of the seam. Additional taught points were recorded at small and curved portions of the seams to ensure accurate seam reconstruction. Following the teaching phase, the adaptive path planning approach developed in this study (see Section 5) was implemented to interpolate a smooth spline through the taught points, generating the final ground-truth seam.

6.4. Experimental results

Experimental results across all twelve scenarios are presented in Figs. 10–13. In the figures the environmental settings and target seam

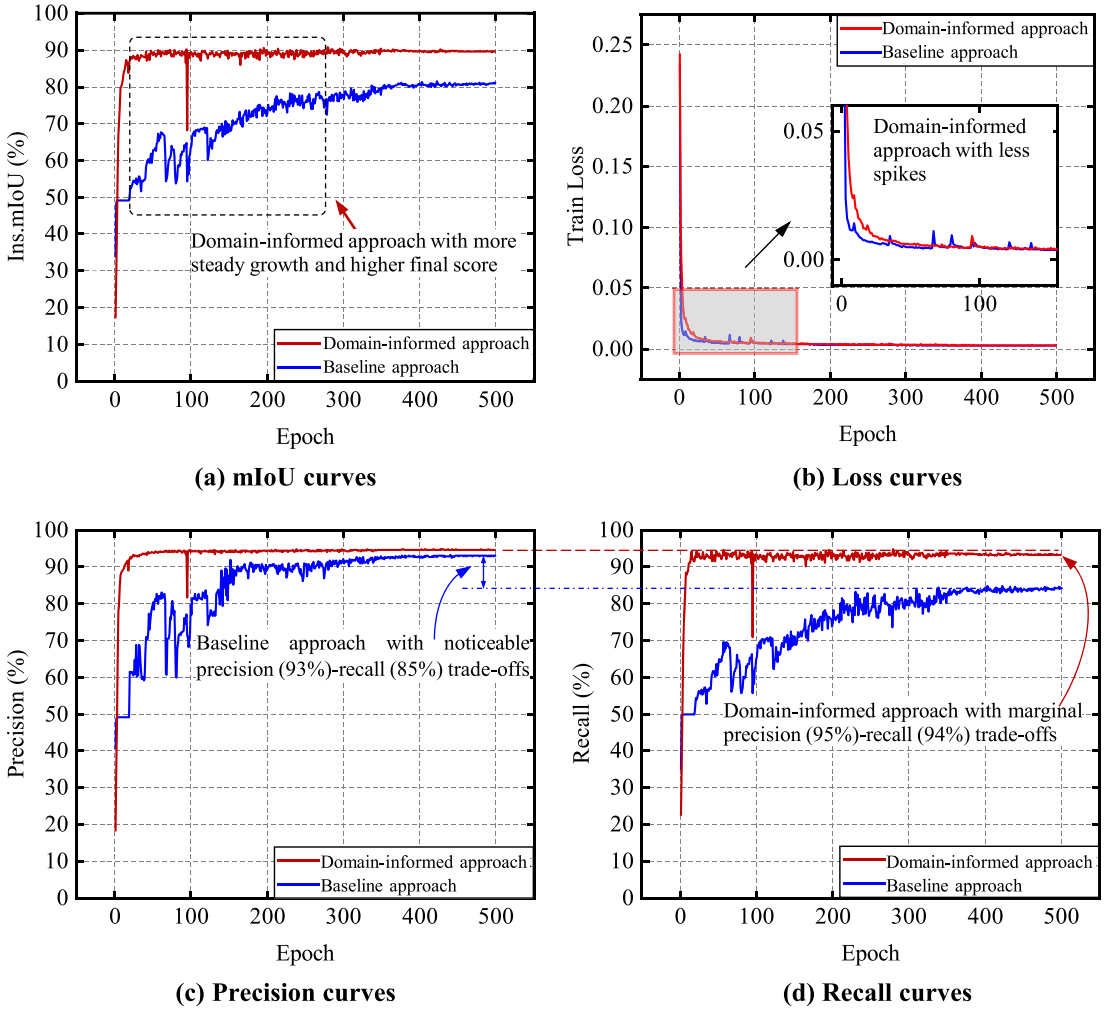


Fig. 9. Monitoring curves for network training.

details are documented, and performance of both domain-informed and baseline frameworks for each test case is explicitly compared. The zoom-in views highlight feature point extraction along the target seams and show the predicted seam paths against their ground truth. Possible seam extraction failures and inaccuracies (by the baseline framework) are identified in the zoomed views. Key point cloud characteristics, including the feature point count (FP_Num) and feature-point-to-non-feature-point ratio (FNR), are labeled in the corresponding graph.

Generally, the point distributions processed by the domain-informed and baseline frameworks exhibit substantial differences, though performance evaluations employ consistent 2048-point clouds for both frameworks for fair comparison. Critical results in Figs. 10–13 demonstrate that, compared with the baseline, the domain-informed framework identifies more feature points and considerably increases the feature-point-to-non-feature-point ratio (FNR) from a baseline range of 0.4 ~ 3.1 % to 8.6 ~ 20.3 %.

Table 5 summarizes the quantitative results associated with each test scenario. The generalizing capability of the seam extraction frameworks is quantified by path completeness, which is established as the ratio of predicted seam path length to the corresponding ground truth seam length. The seam extraction accuracy is quantified through the maximum error (ME, the largest Euclidean deviation) and root mean square error (RMSE) between the predicted and ground truth paths:

$$ME = \max_{1\le i\le N} d_i \tag{16}$$

$$RMSE = \sqrt{\frac{1}{N}\sum_{i=1}^{N} d_i^2} \tag{17}$$

Here, $d_i = \|p_i^{\text{pred}} - p_i^{\text{gt}}\|_2$ denotes the point-wise Euclidean deviation, where $\{p_i^{\text{gt}}\}_{i=1}^{N}$ and $\{p_i^{\text{pred}}\}_{i=1}^{N}$ are the sets of corresponding points along the ground-truth and predicted seam paths, respectively.

It is noted that, for test scenarios with incomplete seam path prediction (path completeness < 95 %), the result is considered a seam extraction failure (N/A). The ME and RMSE values are not reported for these cases.

7. Discussions

This section conducts an in-depth analysis of the domain-informed seam extraction framework. The generalizing ability, adaptability, accuracy, and computational efficiency of the framework are evaluated and thoroughly discussed.

7.1. Generalization capability to new workpiece types and unseen seam topologies

In general, the comprehensive experimental program demonstrates distinct generalization capabilities between the domain-informed and

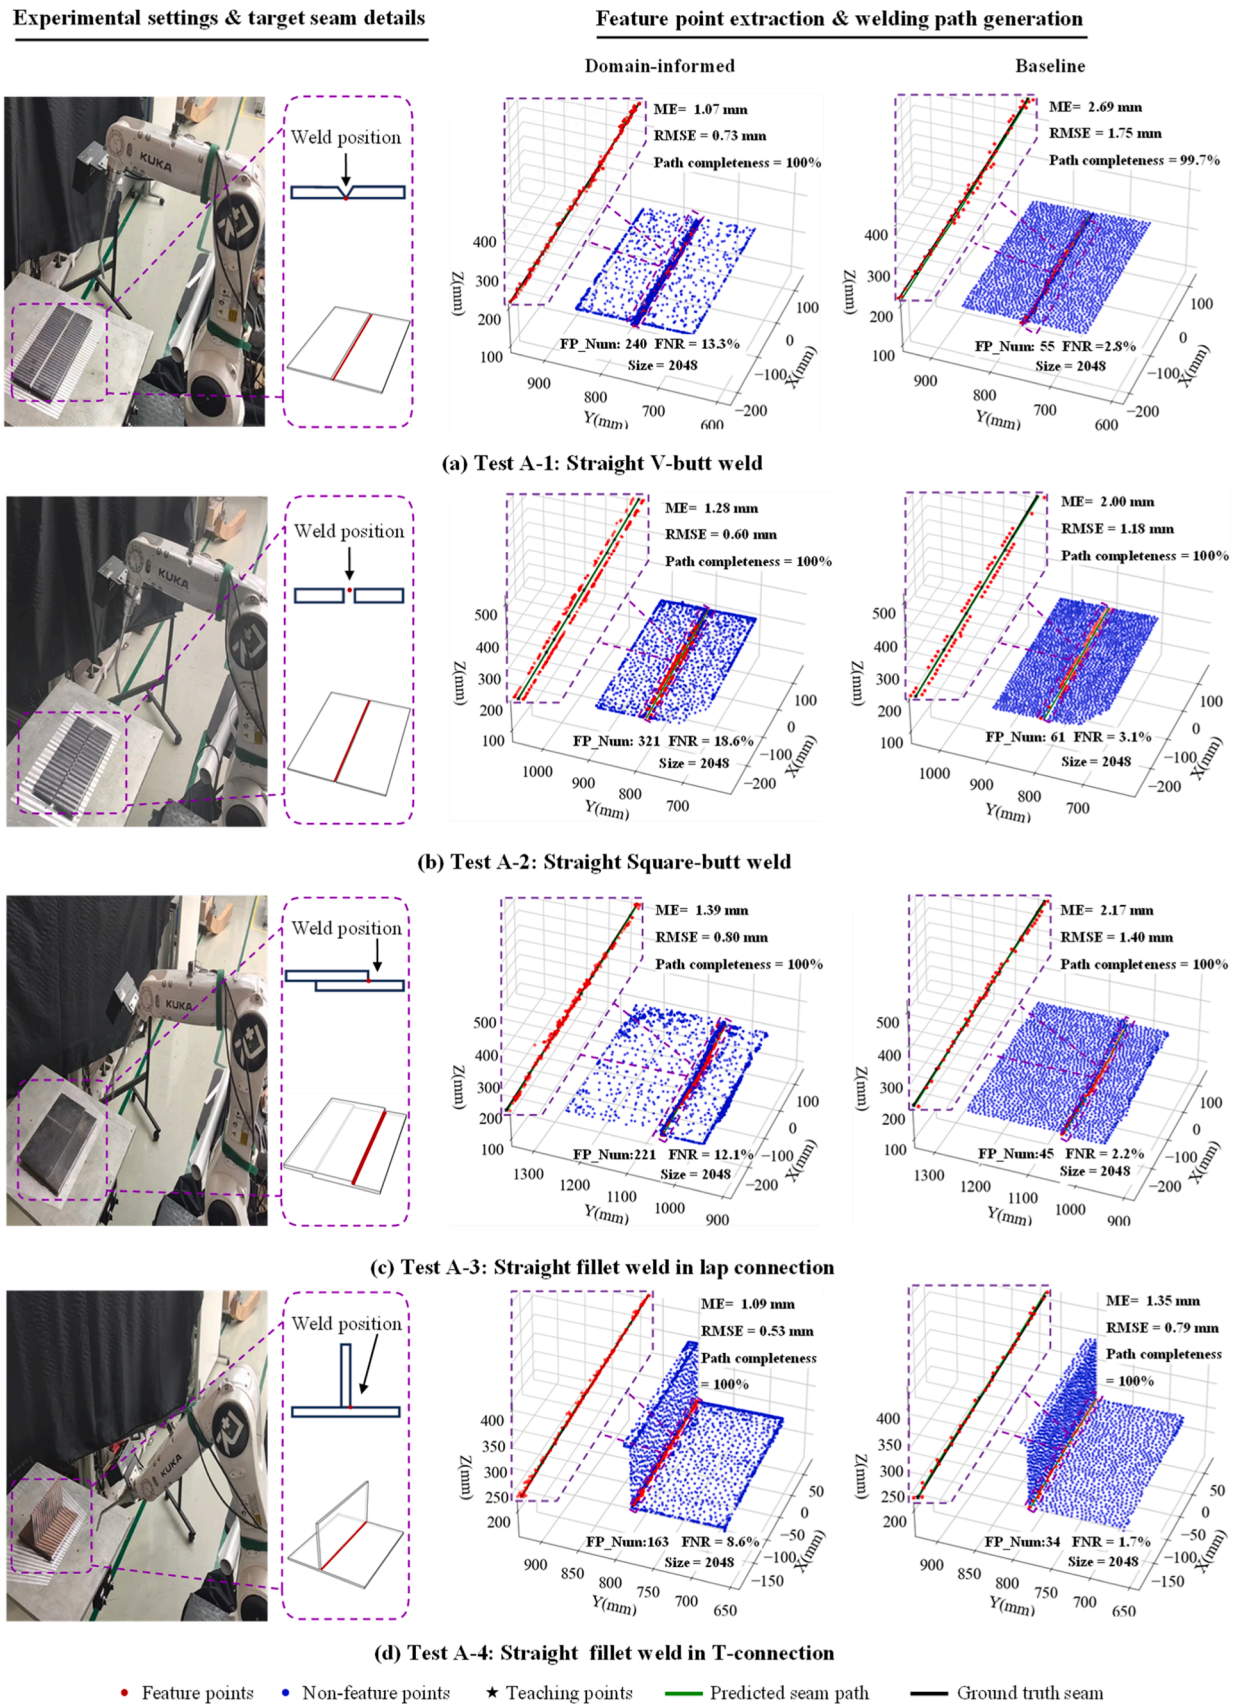


Fig. 10. Experimental straight seam extraction performance on familiar workpieces.

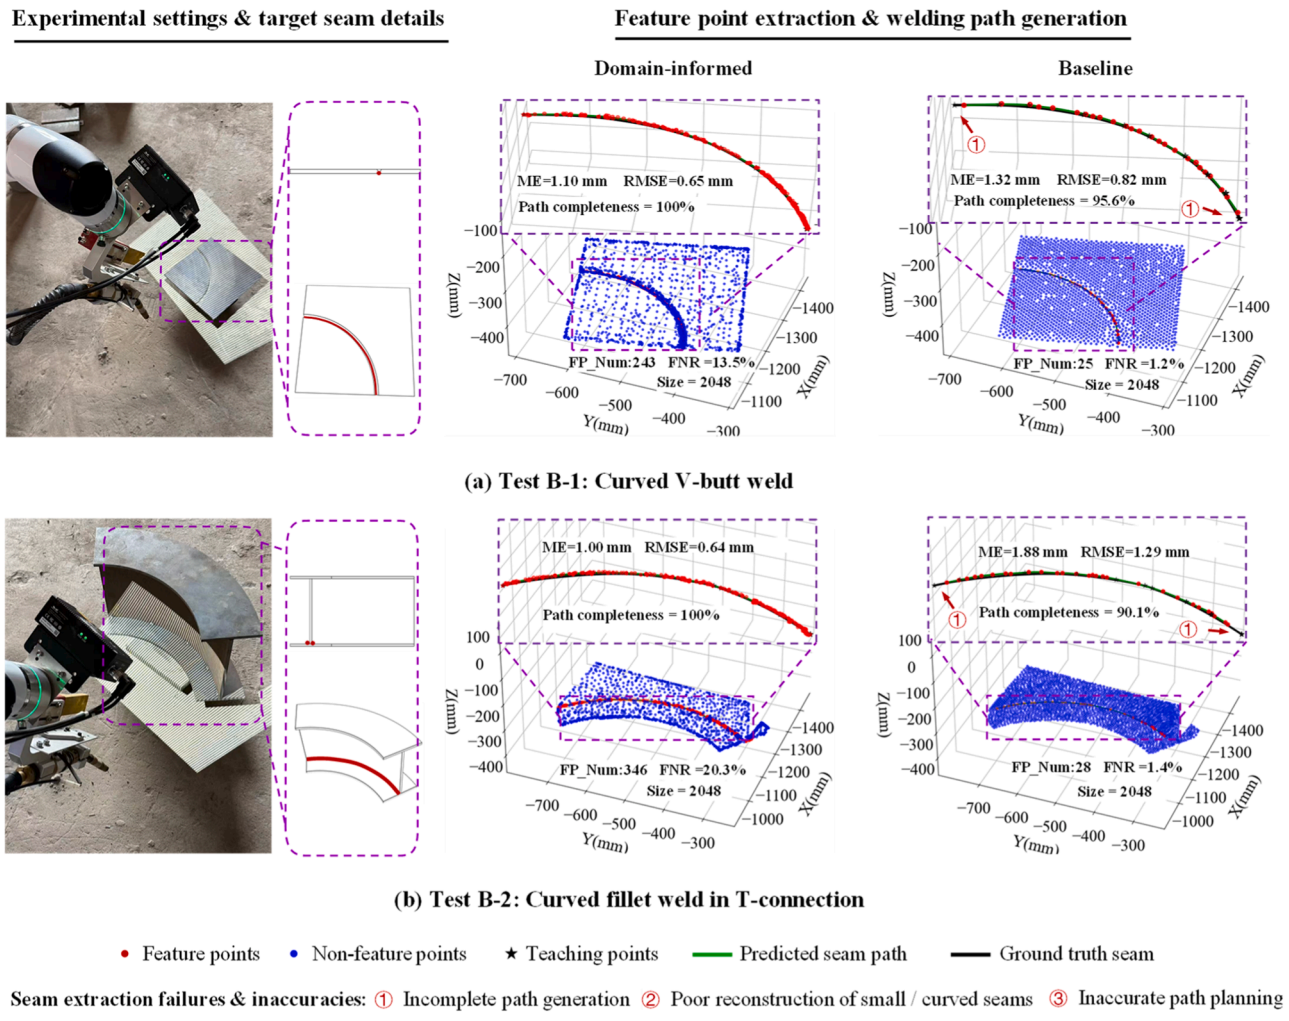


Fig. 11. Experimental curved seam extraction performance on unseen workpieces.

baseline frameworks. Even though the baseline framework successfully identifies simple straight seams from familiar workpieces (Test Group A) and achieves 99.7 % ~ 100 % predicted path completeness on the four basic test cases (Fig. 10), significant performance degradation in the baseline approach could be clearly observed when dealing with new types of workpieces or unstructured real-world structural joints. As the workpiece types and target seam topologies grow more complex, the predicted path completeness by the baseline framework drops to 90.1 % ~95.6 % for moderately challenging cases and further to a lowest of 67.1 % for real-world scenarios (see Table 5 Column (12)).

As evident from Figs. 11–13, when dealing with unseen weld seams, the baseline approach fails to detect adequate feature points at critical regions and thus results in incomplete path generation near the endpoint regions (see Seam extraction failure ①) and poor reconstruction of small / curved seam segments (see Seam extraction failure ②). Evident deviations between the predicted seam paths and their ground truth could be observed.

In contrast, the domain-informed framework identifies significantly more feature points than the baseline along the entire seam path. Critical seam regions, including the endpoint zones and the intermediate small / curved sections, are well reproduced. The resulting seam paths closely align with the corresponding ground truth seams. Quantitative analysis indicates that the domain-informed framework consistently achieves 100 % predicted path completeness from simple (Fig. 10) and moderately complex (Fig. 11) to highly challenging scenarios (Figs. 12,13). Quantitative results in Table 5 indicate only marginal variances are observed between the domain-informed path prediction and the ground truth measurements. These observations highlight the exceptional generalization capability of the domain-informed seam extraction framework, particularly when handling new types of workpieces and novel seam topologies.

Remarkably, the proposed domain-informed framework requires only minimal training data (basic straight-seam patterns) and maintains its performance on novel workpiece types without model retraining, which offers significant practical advantages for robotic welding applications where workpiece variability is common but training data reconstruction and model retraining are costly.

7.2. *Adaptability to operational settings and geometric variations*

Experimental results demonstrate that the domain-informed framework not only identifies target weld seams from controlled laboratory environments (Test Group A, Fig. 10) but also achieves robust seam extraction performance when confronted with real-world industrial conditions including variable lighting conditions, shadow interference, cluttered backgrounds, and etc. (Test Groups B, C, and D, Figs. 11–13). Such observations confirm that the proposed framework effectively adapts to diverse operational scenarios, and highlight its practical applicability in industrial settings.

Test Group D comprised a pair of complex cold-formed thin-walled sections featuring nominally similar profile but varied geometric parameters (Tests D-1 & D-2 used a 100 × 120 section, whereas Tests D-3 & D-4 used a 140 × 110 section). As evidenced in Fig. 13, for both sections the domain-informed framework effectively reproduces the complex

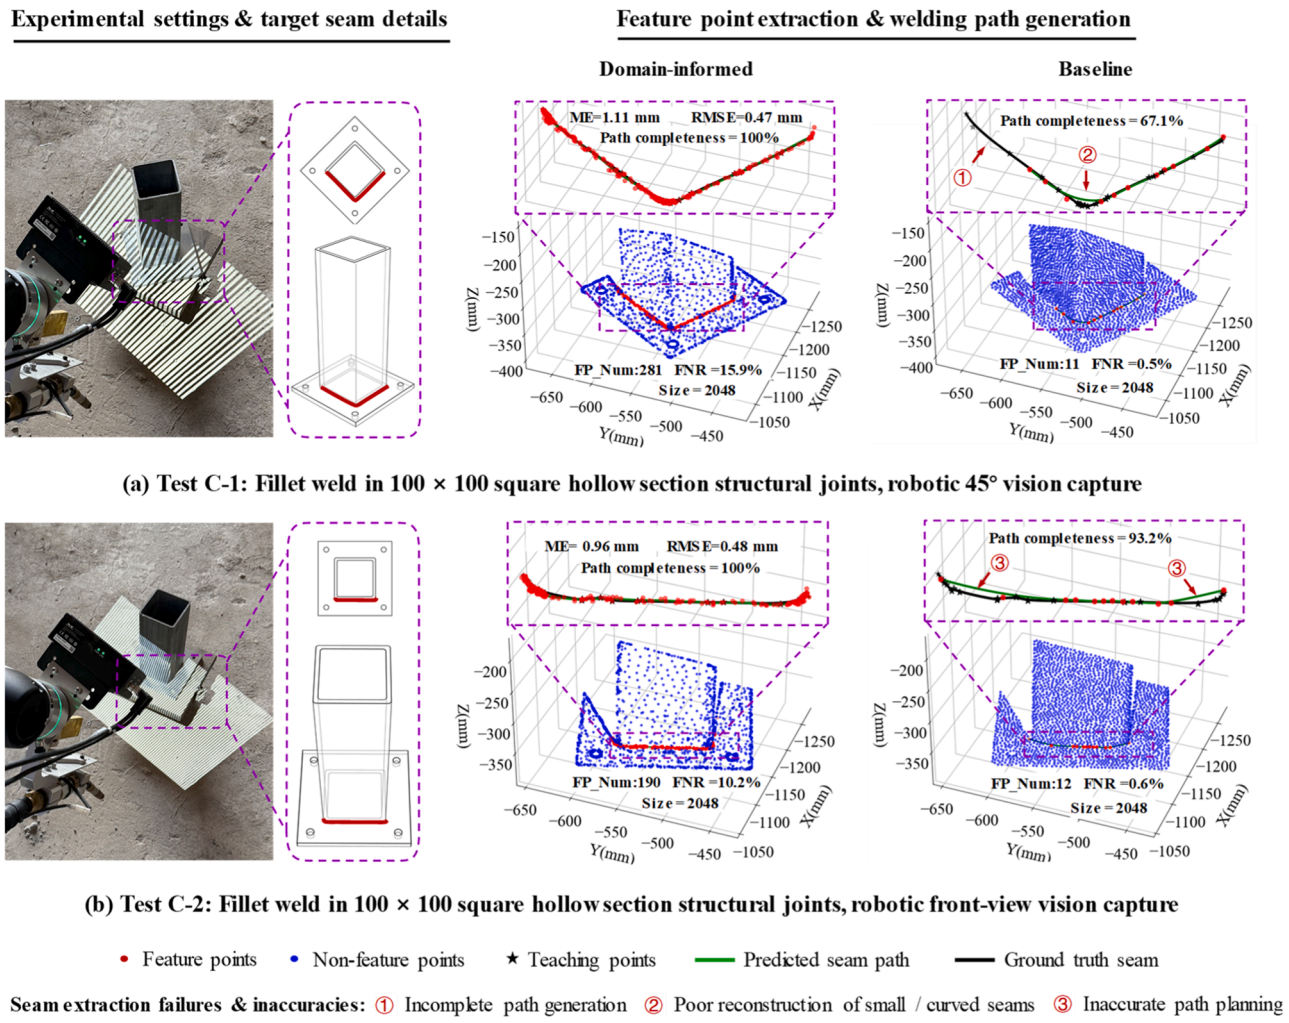


Fig. 12. Experimental seam extraction results for a real-world base joint with a 100 × 100 hot-rolled square section.

seam patterns and maintains geometric fidelity with their ground truth counterparts. These results confirm the robust adaptability of the proposed framework to geometric variations in seam patterns.

7.3. Accuracy

Table 5 summarizes experimental deviations between the predicted seam path and their ground truth. Generally, the domain-informed framework achieves a maximum error (ME) in Euclidean distance ranging from 0.89 mm (Test D-2) to 1.39 mm (Test A-3), whereas, the baseline framework exhibits higher MEs of 1.32 mm (Test B-1) to 2.69 mm (Test A-1). Furthermore, the domain-informed framework exhibits considerably lower root mean square error (RMSE) values (0.41 ~ 0.80 mm) compared to the baseline framework (0.79~1.75 mm) (see Table 5 Columns (4)-(5) and (13–14)). These results demonstrate that the proposed domain-informed framework achieves substantially improved precision and consistency in seam extraction compared to the baseline framework. It is also pointed out that both the ME and RMSE values of predicted weld seam by the domain-informed framework remain well below the designed throat thickness for fillet welds or the weld width for butt welds, indicating that the proposed framework achieves the necessary seam extraction accuracy for robust industrial implementation.

In addition to the Euclidean distance measurements, Columns (6)-(11) and (15)-(20) of Table 5 also presents maximum seam path deviations along the *x*-, *y*-, and *z*-axes. These results provide additional experimental data and can be qualitatively compared with those reported in related works examining similar workpiece types and comparable seam geometries (e.g., Yang et al [10], Fang and Tian [30]). However, we avoid direct quantitative comparisons in this paper as differences in workpiece geometric details and experimental settings preclude strictly fair evaluation. Nevertheless, it is stated that the proposed domain-informed framework exhibits competitive seam extraction performance compared to currently available approaches in terms of both precision and consistency.

7.4. Computational efficiency

Table 6 quantifies the processing efficiency of the proposed framework for representative test scenarios on real-world structural joints (i.e., Tests C-1 and D-2). The total processing time and its division between seam extraction and path planning stages are presented (see Table 6 Columns (13)-(15)). The results demonstrate that the total processing time for both complex cases remained below 0.3 s, undoubtedly meeting the requirement for real-time robotic welding deployment. The seam extraction stage alone required less than 0.04 s (inference time), which underscores the high computational efficiency of the optimized PointNeXt model for this task.

This section extends the computational efficiency evaluation with a point-cloud-size sensitivity analysis using experimentally acquired raw point clouds from Tests C-1 and D-2. Fig. 14 compares the seam extraction performance across three resolutions (2048, 1024, and 512

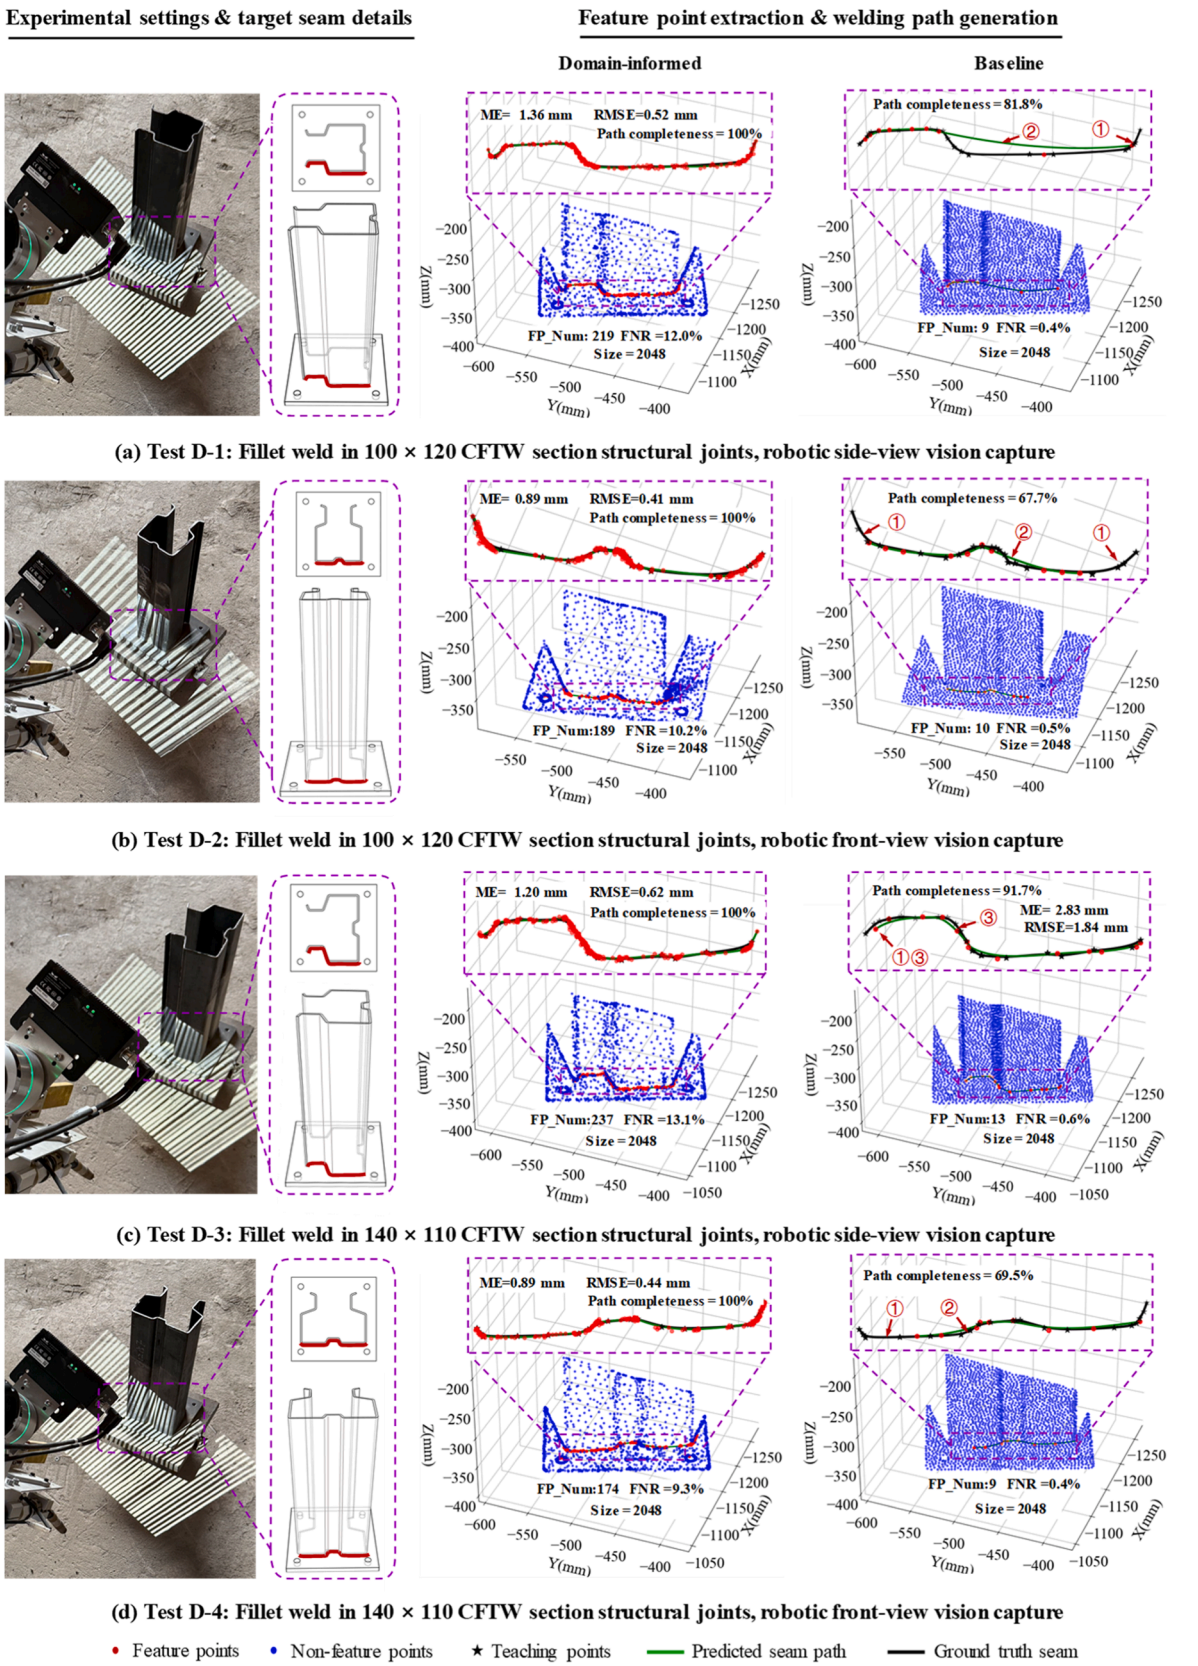


Fig. 13. Experimental seam extraction results for real-world base joints with complex cold-formed thin-walled sections.

Table 5
Experimental results.

TestID	Weld type	Workpiece/joint type	Domain-informed framework								Baseline framework							
			Path compl.	Path deviations Eucl. dist.		x-axes		y-axes		z-axes		Path compl.	Path deviations Eucl. dist.		x-axes		y-axes		z-axes	
				ME	RMSE	ME	RMSE	ME	RMSE	ME	RMSE		ME	RMSE	ME	RMSE	ME	RMSE	ME	RMSE
(1)	(2)	(3)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)	(19)	(20)
A-1	V-butt	Butt	100 %	1.07	0.73	0.16	0.11	0.95	0.61	0.75	0.38	99.7 %	2.69	1.75	0.69	0.24	2.66	1.70	0.63	0.38
A-2	Sq.-butt	Butt	100 %	1.28	0.60	0.14	0.08	1.26	0.47	0.89	0.35	100 %	2.00	1.18	0.39	0.19	1.94	1.12	0.58	0.31
A-3	Fillet	Lap	100 %	1.39	0.80	0.24	0.13	1.29	0.68	0.67	0.40	100 %	2.17	1.40	0.27	0.14	1.78	1.08	1.48	0.88
A-4	Fillet	T-shaped	100 %	1.09	0.53	0.09	0.05	1.07	0.40	0.63	0.34	100 %	1.35	0.79	0.14	0.08	1.34	0.72	0.60	0.31
B-1	V-butt	Curved, Butt	100 %	1.10	0.65	0.85	0.31	0.57	0.24	0.90	0.52	95.6 %	1.32	0.82	0.84	0.48	1.05	0.45	0.80	0.48
B-2	Fillet	Curved, T-	100 %	1.00	0.64	0.91	0.51	0.31	0.19	0.67	0.34	90.1 %	N/A							
C-1	Fillet	SHS100×100	100 %	1.11	0.47	0.88	0.30	0.47	0.20	0.86	0.30	67.1 %	N/A							
C-2	Fillet	SHS100×100	100 %	0.96	0.48	0.77	0.31	0.82	0.18	0.61	0.32	71.5 %	N/A							
D-1	Fillet	CFTW100×120	100 %	1.36	0.52	0.85	0.36	1.06	0.30	0.82	0.23	81.8 %	N/A							
D-2	Fillet	CFTW100×120	100 %	0.89	0.41	0.61	0.29	0.42	0.16	0.67	0.25	67.7 %	N/A							
D-3	Fillet	CFTW140×110	100 %	1.20	0.62	0.76	0.39	0.96	0.29	0.73	0.37	91.7 %	N/A							
D-4	Fillet	CFTW140×110	100 %	0.89	0.44	0.72	0.26	0.48	0.16	0.83	0.31	69.5 %	N/A							
Max.				**1.39**	**0.80**								**2.69**	**1.75**						
Min.				**0.89**	**0.41**								**1.32**	**0.79**						

Notes: Test scenarios with incomplete seam path prediction (path completeness < 95 %) are deemed seam extraction failures. For these cases, the Maximum Error (ME) and Root Mean Square Error (RMSE) values are not reported.

points), where the 2048-point size served as the default in previous experiments. The corresponding quantitative performance metrics for each resolution are also summarized in Table 6.

As shown in Fig. 14, the domain-informed framework maintains satisfactory performance with reduced point densities. Complete seam extraction in 100×100 square hollow section base-joint (100 % path completeness) and near-complete extraction in 100×120 cold-formed thin-walled section base-joint (95.1 % path completeness) are achieved when using 1024-point clouds. Analysis of Table 6 indicates that, compared to 2048-point clouds, processing with 1024-point clouds achieves substantial efficiency improvements. The total processing time decreases by 24.9 ∼ 27.5 %, seam extraction stage becomes 50.0 % faster, and path planning efficiency improves by 20.8 ∼ 24.0 %. It is pointed out that, though these gains entail a marginal compromise in seam extraction accuracy, such improvements in computational efficiency are crucial for tasks involving large-scale workpieces with lengthy seams, where processing speed is paramount for real-time robotic welding deployment.

For 512-point clouds, however, a sharp degradation in predicted path completeness is observed. This indicates that, for the examined welding applications, such point size has fallen below the minimum viable threshold for reliable seam extraction. Practically, engineers are advised to perform the sensitive analysis, and optimize the trade-off between extraction accuracy and computational efficiency by selecting appropriate point cloud resolutions.

8. Conclusions and limitations

This paper proposed a generalized learning framework for robust seam extraction in robotic welding applications. The proposed framework integrates the three primary components developed in this study, including a domain-informed model enhancement methodology, a simulation-based data acquisition pipeline, and an adaptive weld path planning approach. The domain-informed model enhancement methodology forms the core of the seam extraction framework, directly strengthening its generalization capability for extracting unstructured seam topologies from unseen workpiece types. Meanwhile, the data acquisition pipeline and the adaptive path planning approach enhance the operational adaptability of the framework. The generalization capability and adaptability of the proposed framework have been thoroughly validated through comprehensive experiments. The major conclusions of this study are as follows:

(i) The proposed domain-informed model generalization enhancement methodology strategically embeds welding expertise and operational logic of welders into the robotic welding seam extraction learning framework by applying a surface normal divergence based sparsity technique to raw workpiece point clouds. This enhancement consistently yields a significant increase in the feature-to-non-feature-point-ratio (FNR) in the sampled data and thereby guides the model to prioritize welding positions and seam-relevant regions over specific seam pattern. Comprehensive experimental validation confirmed that the generalization enhancement is directly attributed to the proposed domain-informed methodology rather than other components of in the framework. Ablation studies demonstrated that the domain-informed methodology provides universal improvements to generalizing capability across diverse models, which underscores the broad applicability of the methodology.

An optimized robotic welding seam extraction model was developed using PointNeXt as the backbone architecture. The optimized model balances the seam extraction accuracy with computational efficiency for industrial real-time applications. Quantitative evaluation indicates that the optimized model achieves a mean seam extraction deviation of only 1.1 mm with a

Table 6
Processing efficiency of the proposed seam extraction framework and a point cloud size sensitive analysis for representative scenarios on real-world structural joints.

Test ID	Weld type	Workpiece /joint type	Point size	Domain-informed framework										
				Path compl.	Path deviations (mm)								Processing time consumed (ms)		
					Eucl. dist.		x-axes		y-axes		z-axes		Total	Seam extraction	Path planning
					ME	RMSE	ME	RMSE	ME	RMSE	ME	RMSE			
(1)	(2)	(3)	(4)	(5)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
C-1	Fillet	SHS100×100	2048	100 %	1.11	0.47	0.88	0.30	0.47	0.20	0.86	0.30	28.84	3.92	24.92
			1024	100 %	1.21	0.50	0.92	0.26	0.68	0.25	1.06	0.33	20.90 (27.5 % ↓)	1.96 (50.0 % ↓)	18.94 (24.0 % ↓)
			512	87.9 %	N/A								17.06	1.11	15.95
D-2	Fillet	CFTW100×120	2048	100 %	0.89	0.41	0.61	0.29	0.42	0.16	0.67	0.25	27.84	3.92	23.92
			1024	95.1 %	1.50	0.66	1.40	0.55	0.53	0.19	0.69	0.32	20.90 (24.9 % ↓)	1.96 (50.0 % ↓)	18.94 (20.8 % ↓)
			512	86.2 %	N/A								17.56	1.10	16.46

Notes: Test scenarios with incomplete seam path prediction (path completeness < 95 %) are deemed seam extraction failures. For these cases, the Maximum Error (ME) and Root Mean Square Error (RMSE) values are not reported.

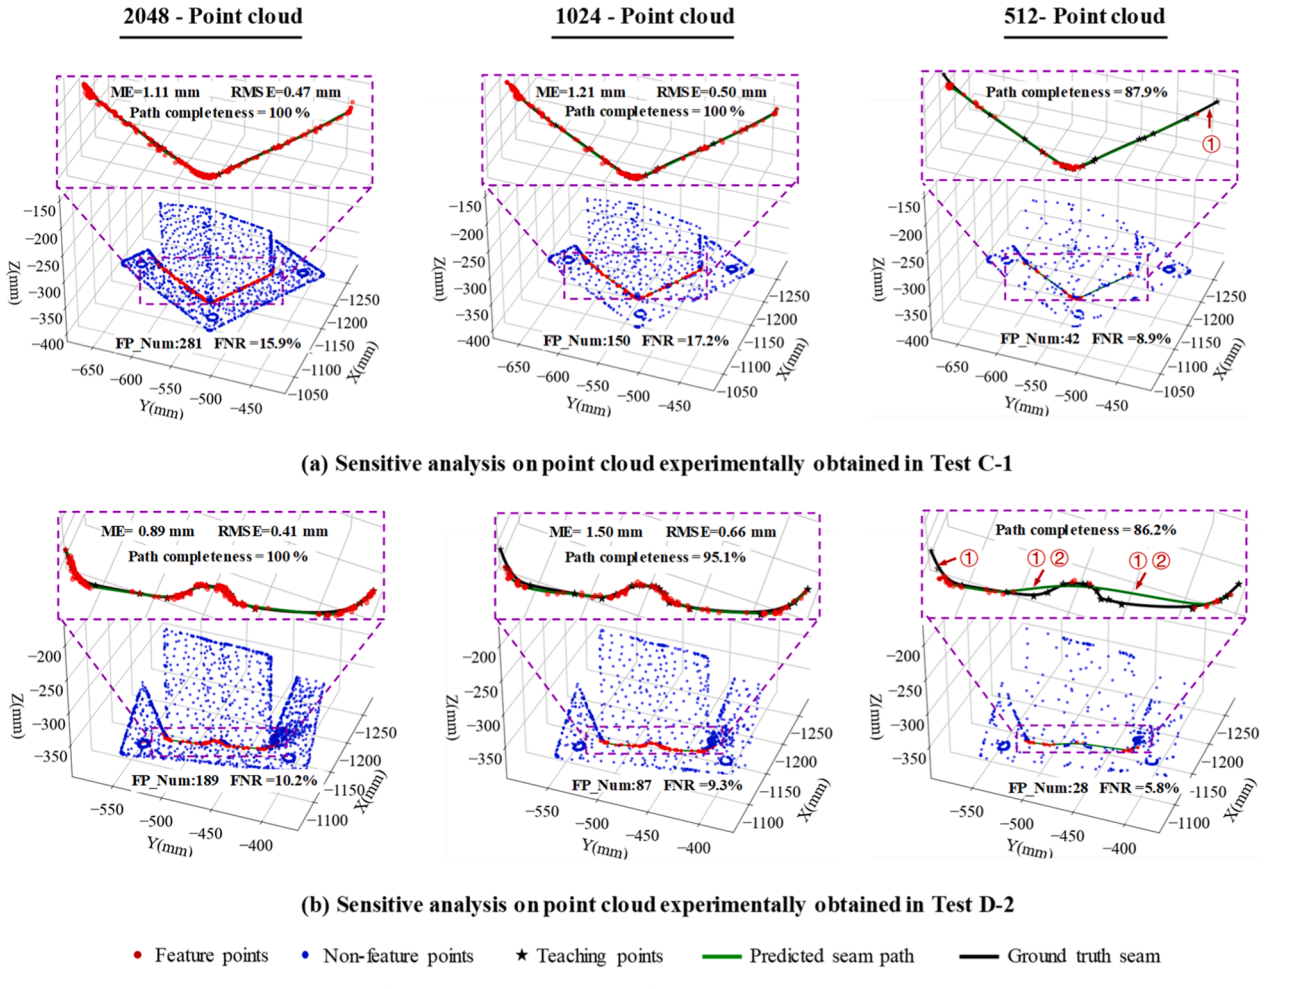


Fig. 14. Point cloud processing sensitive analysis.

maximum inference time of 0.04 s, which fully meets the requirements for real-world robotic welding deployment.

(ii) The established simulation-based data acquisition pipeline incorporated critical real-world variables that directly impact vision-based measurements, including dynamic lighting conditions (both natural and artificial light variations), their secondary effects (shadows and surface reflections), and sensor noise. Occlusion patterns caused by robotic motion were intentionally

simulated to replicate authentic operational scenarios. The pipeline enabled the automated generation of large-scale realistic workpiece point cloud datasets. Experimental validation confirmed that this comprehensive data generation pipeline ensures robust training coverage and facilitates consistent model adaptation across diverse operational settings and robotic platforms.

(iii) The adaptive weld path planning approach implemented a curvature-based solution to achieve an optimal balance between path smoothness and geometric detail preservation. A principal-component-analysis (PCA) based hybrid point sequencing algorithm was developed to avoid invalid welding path generation. In conjunction these two methods guarantee robust robotic welding path generation for complex real-world seams with small intermediate segments and sharp corners.

(iv) A comprehensive experimental validation was performed to evaluate the proposed domain-informed robotic welding seam extraction framework. A total of twelve experimental tests were conducted across both controlled laboratory settings and unstructured real-world factory environments using distinct robotic platforms. The evaluation encompassed four distinct test groups: one with four basic cases of familiar workpieces, and three groups containing eight challenging scenarios with new workpiece types and novel seam topologies. In every test scenario, the performance of the proposed domain-informed framework was compared against the baseline framework.

Experimental results reveal that the proposed domain-informed framework exhibits robust generalization capability, achieving complete and accurate seam extraction on unseen workpiece types and seam patterns. The proposed framework also adapts to diverse environmental conditions, and maintains real-time performance for industrial deployment. Test results confirm that the domain-informed framework requires only minimal training data (comprising basic straight-seam patterns) and maintains robust performance on unseen scenarios without model retraining. In contrast, the baseline framework failed all the unseen test cases, with its predicted path completeness deteriorating sharply as the seam pattern complexity increases.

It is pointed out that, while this research has demonstrated the robust generalization capability of the proposed domain-informed seam extraction framework, its current application is primarily focused on isolated weld seams in structural steel fabrication. The performance of the proposed seam extraction framework may be less satisfactory in more complex scenarios, such as those in shipbuilding industry, where multiple stiffeners and obstacles can interfere with seam extraction. Future work should extend the framework to handle environments with a high density of closely spaced welds and obstructions, as found in shipbuilding and similar heavy industries.

CRediT authorship contribution statement

Xianzhong Zhao: Writing – review & editing, Resources, Methodology, Funding acquisition. **Haotian Liu:** Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation. **Zhaoqi Huang:** Writing – original draft, Visualization, Validation, Methodology, Investigation, Formal analysis, Conceptualization. **Yuzhen He:** Visualization, Validation.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This work is supported by Shanghai Qi Zhi Institute (Grant No. SQZ202309) and the China Postdoctoral Science Foundation under Grant Number 2024M752413.

Data availability

Data will be made available on request.

References

[1] B. Wang, S.J. Hu, L. Sun, T. Freiheit, Intelligent welding system technologies: state-of-the-art review and perspectives, J. Manuf. Syst. 56 (2020) 373–391, https://doi.org/10.1016/j.jmsy.2020.06.020.
[2] A. Rout, B.B.V.L. Deepak, B.B. Biswal, Advances in weld seam tracking techniques for robotic welding: a review, Robot. C.-Int. Manuf. 56 (2019) 12–37, https://doi.org/10.1016/j.rcim.2018.08.003.
[3] Y. He, Z. Huang, H. Liu, J. Ye, Y. Lu, X. Zhao, Self-adaptive seam detection framework for unmanned structural steel welding robots in unstructured environments, Autom. Constr. 175 (2025) 106221, https://doi.org/10.1016/j.autcon.2025.106221.
[4] K. Weman, G. Lindén, MIG Welding Guide, CRC Press, New York, 2006.
[5] C. Zheng, Y. An, Z. Wang, H. Wu, X. Qin, B. Eynard, Y. Zhang, Hybrid offline programming method for robotic welding systems, Robot. C.-Int. Manuf. 73 (2022) 102238, https://doi.org/10.1016/j.rcim.2021.102238.
[6] G. Du, M. Chen, C. Liu, B. Zhang, P. Zhang, Online robot teaching with natural human–robot interaction, IEEE Trans. Ind. Electron. 65 (2018) 9571–9581, https://doi.org/10.1109/TIE.2018.2823667.
[7] L. Yang, Y. Liu, J. Peng, Advances techniques of the structured light sensing in intelligent welding robots: a review, Int. J. Adv. Manuf. Technol. 110 (2020) 1027–1046, https://doi.org/10.1007/s00170-020-05524-2.
[8] Y. Geng, M. Lai, X. Tian, X. Xu, Y. Jiang, Y. Zhang, A novel seam extraction and path planning method for robotic welding of medium-thickness plate structural parts based on 3D vision, Robot. C.-Int. Manuf. 79 (2023) 102433, https://doi.org/10.1016/j.rcim.2022.102433.
[9] Q. Guo, Z. Yang, J. Xu, Y. Jiang, W. Wang, Z. Liu, W. Zhao, Y. Sun, Progress, challenges and trends on vision sensing technologies in automatic/intelligent robotic welding: state-of-the-art review, Robot. Comput.-Integr. Manuf. 89 (2024) 102767, https://doi.org/10.1016/j.rcim.2024.102767.
[10] L. Yang, Y. Liu, J. Peng, Z. Liang, A novel system for off-line 3D seam extraction and path planning based on point cloud segmentation for arc welding robot, Robot. C.-Int. Manuf. 64 (2020) 101929, https://doi.org/10.1016/j.rcim.2019.101929.
[11] Y. Zhang, Y. Jiang, X. Tian, X. Xu, Y. Geng, L. Min, A point cloud-based welding trajectory planning method for plane welds, Int. J. Adv. Manuf. Technol. 125 (2023) 1645–1659, https://doi.org/10.1007/s00170-022-10699-x.
[12] B. Zhou, Y. Liu, Y. Xiao, R. Zhou, Y. Gan, F. Fang, Intelligent guidance programming of welding robot for 3D curved welding seam, IEEE Access 9 (2021) 42345–42357, https://doi.org/10.1109/ACCESS.2021.3065956.
[13] V. Patil, I. Patil, V. Kalaichelvi, R. Karthikeyan, Extraction of weld seam in 3D point clouds for real time welding using 5 DOF robotic arm, in: 2019 5th International Conference on Control, Automation and Robotics (ICCAR), 2019, pp. 727–733, https://doi.org/10.1109/ICCAR.2019.8813703.
[14] J. Liu, T. Jiao, S. Li, Z. Wu, Y.F. Chen, Automatic seam detection of welding robots using deep learning, Autom. Constr. 143 (2022) 104582, https://doi.org/10.1016/j.autcon.2022.104582.
[15] J. Cheng, H. Jin, X. Qian, Real-time weld seam feature extraction in construction sites, Autom. Constr. 160 (2024) 105330, https://doi.org/10.1016/j.autcon.2024.105330.
[16] Y. Zhang, Y. Geng, X. Tian, Y. Sun, X. Xu, A novel weld seam extraction method with semantic segmentation and point cloud feature for irregular structure workpieces, Robot. Comput.-Integr. Manuf. 95 (2025) 102987, https://doi.org/10.1016/j.rcim.2025.102987.
[17] Y. Ma, J. Fan, S. Zhao, F. Jing, S. Wang, M. Tan, From model to reality: a robust framework for automatic generation of welding paths, IEEE Trans. Ind. Electron. 71 (2024) 16153–16164, https://doi.org/10.1109/TIE.2024.3395792.
[18] L. Song, H. Wang, Y. Zhang, Z. Qiao, F. Han, LWSNet: a lightweight network for automated welding point cloud segmentation, Measurement 243 (2025) 116290, https://doi.org/10.1016/j.measurement.2024.116290.
[19] C.K. Liu, D. Negrut, The role of physics-based simulators in robotics, annual review of control, Robot. Auton. Syst. 4 (2021) 35–58, https://doi.org/10.1146/annurev-control-072220-093055.
[20] S. Chitta, MoveIt!: an introduction, , in: A. Koubaa (Ed.), Robot Operating System (ROS): the Complete Reference, Robot Operating System (ROS): the Complete Reference, 1, , Springer International Publishing, Cham, 2016, pp. 3–27, https://doi.org/10.1007/978-3-319-26054-9_1.
[21] CEN, Welding — multilingual terms for welded joints with illustrations, EN 12345, European Committee for Standardization (CEN), 1998.
[22] CEN, Eurocode 3—design of steel structures. Design of joints, EN 1993-1-8, European Committee for Standardization (CEN), 2005.

[23] ISO, Welding and allied processes — symbolic representation on drawings — welded joints, ISO 2553, International Organization for Standardization (ISO), 2019.
[24] R.Q. Charles, H. Su, M. Kaichun, L.J. Guibas, PointNet: deep learning on point sets for 3D classification and segmentation, in: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), IEEE, Honolulu, HI, 2017, pp. 77–85, https://doi.org/10.1109/CVPR.2017.16.
[25] C.R. Qi, L. Yi, H. Su, L.J. Guibas, PointNet++: deep hierarchical feature learning on point sets in a metric space, in: I. Guyon, U.V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, R. Garnett (Eds.), Advances in Neural Information Processing Systems, Curran Associates, Inc., 2017.
[26] G. Li, M. Muller, A. Thabet, B. Ghanem, DeepGCNs: can GCNs go as deep as CNNs?, in: 2019: pp. 9267–9276. https://doi.org/10.1109/ICCV.2019.00936.
[27] DGCNN: a convolutional neural network over large-scale labeled graphs, Neural Netw. 108 (2018) 533–543, https://doi.org/10.1016/j.neunet.2018.09.001.
[28] G. Qian, Y. Li, H. Peng, J. Mai, H. Hammoud, M. Elhoseiny, B. Ghanem, PointNeXt: revisiting PointNet++ with improved training and scaling strategies, Adv. Neural Inf. Process. Syst. 35 (2022) 23192–23204.
[29] Z. Leng, M. Tan, C. Liu, E.D. Cubuk, X. Shi, S. Cheng, D. Anguelov, PolyLoss: a polynomial expansion perspective of classification loss functions, (2022). https://doi.org/10.48550/arXiv.2204.12511.
[30] W. Fang, X. Tian, A novel model-based welding trajectory planning method for identical structural workpieces, Robot. Comput.-Integr. Manuf. 89 (2024) 102772, https://doi.org/10.1016/j.rcim.2024.102772.