
Professional Certificate in Artificial Intelligence for Welding Processes

Computer Vision for Welding Applications

Computer vision for welding applications is a multidisciplinary field that blends image processing, machine learning, and domain-specific knowledge of welding processes. Mastery of the terminology is essential for engineers, researchers, and technicians who wish to develop robust inspection systems, automate quality control, or integrate vision-based guidance into robotic welders. The following glossary-style explanation presents the most important terms, organized thematically, and illustrates each concept with practical examples and typical challenges encountered in real-world welding environments.

Image acquisition refers to the capture of visual data from the welding zone using cameras or sensors. The quality of the acquired image directly influences every downstream algorithm. Key sub-concepts include:

Sensor type – cameras may be visible-light, infrared (IR), or hyperspectral. Infrared imaging is valuable for monitoring weld pool temperature, while hyperspectral sensors can reveal chemical composition of spatter. **Resolution** – the number of pixels in the horizontal and vertical dimensions. Higher resolution enables finer detection of defects such as micro-cracks, but also increases data volume and processing latency. **Frame rate** – the number of images captured per second (fps). Real-time weld monitoring typically requires 30–120 fps to track rapid changes in the molten pool. **Exposure time** – the duration the sensor collects light for each frame. Short exposure reduces motion blur in high-speed welding but may decrease signal-to-noise ratio (SNR) in low-light conditions.

Practical example: A robotic GMAW (gas metal arc welding) cell uses a 4-megapixel, 60-fps RGB camera positioned at a 45° angle to the weld joint. The camera's exposure is set to 5 μs to freeze the bright arc, and an IR camera simultaneously records temperature maps for process control.

Pixel is the smallest addressable element in a digital image. Each pixel stores intensity values that represent the scene's radiance. In welding, pixel intensity often correlates with brightness of the arc, temperature of the molten metal, or reflectivity of the surface. Understanding pixel representation is the foundation for all subsequent processing steps.

Bit depth – the number of bits used to encode each pixel's intensity. An 8-bit image can represent 256 gray levels, whereas a 12-bit image provides 4096 levels, allowing finer discrimination of subtle temperature gradients in the weld pool. **Color space** – the mathematical model that defines how color information is stored. Common spaces include RGB, HSV, and CIE LAB. Converting from RGB to HSV (Hue-Saturation-Value) can simplify segmentation of the bright arc (high Value) from the darker background.

Example: When detecting spatter, converting the RGB image to the HSV space isolates the high-Saturation region where metallic particles appear as bright, saturated spots against a low-Saturation background.

Histogram is a graphical representation of the distribution of pixel intensities. In welding inspection,

histograms help assess exposure quality and guide threshold selection for binary segmentation.

Global thresholding – a single intensity value that separates foreground (e.G., Weld pool) from background. Otsu's method automatically computes an optimal threshold by maximizing inter-class variance. Adaptive thresholding – thresholds vary across the image based on local statistics, useful when illumination is non-uniform due to arc flicker or shadows.

Example: A GMAW process exhibits bright arcs that cause local overexposure. Adaptive thresholding partitions the image into small blocks, computes a local Otsu threshold for each block, and successfully isolates the weld pool despite varying brightness.

Edge detection identifies locations where intensity changes sharply, corresponding to physical boundaries such as the weld bead edge, heat-affected zone (HAZ), or crack tips.

Sobel operator – computes gradient magnitude in horizontal and vertical directions using convolution kernels. Canny detector – a multi-stage algorithm that includes noise reduction, gradient calculation, non-maximum suppression, and hysteresis thresholding for robust edge maps.

Challenge: The intense arc produces strong gradients that can overwhelm edge detectors, leading to false edges. Pre-processing with a Gaussian blur reduces noise, while setting high hysteresis thresholds in the Canny detector filters out spurious arcs.

Segmentation partitions an image into meaningful regions. In welding, segmentation separates the weld bead, spatter, HAZ, and potential defects.

Region-based segmentation – techniques such as region growing start from seed points and expand until a homogeneity criterion is met (e.G., Similar intensity). Clustering – algorithms like k-means group pixels by color or texture features, useful for distinguishing metal from background. Deep segmentation – convolutional neural networks (CNNs) such as U-Net or Mask RCNN learn to predict pixel-wise class labels from annotated training data.

Practical application: A Mask RCNN model trained on annotated TIG weld images can simultaneously detect the weld bead contour and classify spatter as separate instances, enabling automated counting of spatter particles per pass.

Morphological operations manipulate binary masks to refine segmentation results. They are particularly helpful for removing small noise and bridging gaps in weld bead contours.

Erosion – shrinks foreground regions, eliminating isolated pixels. Dilation – expands foreground, filling small holes. Opening – erosion followed by dilation, useful for removing spatter noise while preserving bead shape. Closing – dilation followed by erosion, helps fill gaps in the bead contour caused by missing pixels.

Example: After thresholding the weld pool, a small amount of spatter remains attached to the bead mask. Applying an opening operation with a 3-pixel structuring element removes the spatter while keeping the bead intact.

Contour is a curve joining all continuous points along a boundary with the same intensity. Contour extraction allows measurement of geometric weld attributes.

Chain code – a compact representation of a contour using directional symbols. Polygonal approximation – reduces contour points to a set of line segments while preserving overall shape, facilitating calculation of bead width and reinforcement.

Case study: By extracting the bead contour and approximating it with a polygon, the system computes the maximum bead width and compares it to specification limits, flagging over-fill or under-fill conditions.

Feature extraction converts raw pixel data into informative descriptors that capture shape, texture, or intensity patterns. Features are the building blocks for machine-learning classifiers.

Hand-crafted features – include Histogram of Oriented Gradients (HOG), Local Binary Patterns (LBP), and Gabor filters. Deep features – are automatically learned by convolutional layers of a neural network during training.

Practical scenario: A support-vector-machine (SVM) classifier uses HOG descriptors extracted from the weld pool region to differentiate between stable and unstable arc conditions, providing early warnings for process deviations.

Descriptors are numerical vectors that summarize local image information. In welding, descriptors often capture texture of the HAZ or the shape of a crack.

SIFT (Scale-Invariant Feature Transform) – detects keypoints at multiple scales and computes orientation-normalized descriptors, robust to changes in view angle and illumination. SURF (Speeded-Up Robust Features) – a faster alternative to SIFT, employing approximated Gaussian derivatives. ORB (Oriented FAST and Rotated BRIEF) – combines FAST keypoint detection with BRIEF descriptors, offering high speed for real-time applications.

Challenge: The intense and fluctuating illumination of the welding arc can cause keypoint detectors to produce many false positives. Applying a narrow band-pass filter to suppress high-frequency arc flicker reduces spurious keypoints.

Convolutional Neural Network (CNN) is a deep-learning architecture designed to process grid-like data such as images. CNNs have become the dominant approach for weld defect detection, bead tracking, and process monitoring.

Layer types – include convolutional layers (feature extraction), pooling layers (spatial down-sampling), and fully-connected layers (classification). Activation function – non-linear functions such as ReLU (Rectified Linear Unit) that introduce non-linearity, enabling the network to learn complex patterns. Loss function – measures the discrepancy between predicted and true labels; common choices are cross-entropy for classification and mean-squared error for regression.

Example: A CNN trained on thousands of X-ray images of welds can automatically localize porosity defects, achieving detection accuracy above 95 % when compared with manual radiographic interpretation.

Training dataset comprises labeled images used to teach a machine-learning model. For welding, datasets must capture the diversity of welding conditions, materials, joint designs, and defect types.

Annotation – the process of marking ground-truth information on images, such as bounding boxes around cracks or pixel-wise masks for bead regions. Ground truth – the accurate reference data against which model predictions are evaluated. Data split – division of the dataset into training, validation, and test subsets, typically 70%/15%/15% to avoid overfitting.

Challenge: Annotating weld images is labor-intensive because experts must identify subtle defects that may be invisible to the naked eye. Semi-automated tools that suggest candidate regions for human verification can accelerate the process.

Region of interest (ROI) is a sub-image that contains the weld joint or specific area of interest. Focusing computation on the ROI reduces processing time and improves robustness.

ROI selection – can be manual (operator draws a box) or automatic using object detection networks (e.G., YOLO, SSD) that locate the joint in the full frame. ROI scaling – resizing the ROI to a fixed size before feeding it into a CNN, ensuring consistent input dimensions.

Example: An automated GMAW system first runs a YOLOv5 detector to locate the joint, crops a 256 × 256-pixel ROI around the bead, and then passes the ROI to a defect-classification CNN.

Object detection identifies and localizes instances of predefined classes within an image. In welding, object detection can find spatter, droplets, or the weld pool itself.

Single-stage detectors – such as YOLO (You Only Look Once) or SSD (Single Shot MultiBox Detector), which provide fast inference suitable for real-time monitoring. Two-stage detectors – like Faster R-CNN, which first proposes regions and then classifies them, offering higher accuracy at the cost of speed.

Practical use: A YOLOv3 model deployed on an edge GPU processes 60 fps, flagging any frame where spatter count exceeds a threshold, allowing the controller to adjust wire feed speed on the fly.

Inference is the phase where a trained model processes new data to generate predictions. In welding applications, inference must often meet stringent timing constraints.

Latency – the time elapsed from image capture to output generation. Low latency (Throughput – number of frames processed per second; high throughput enables monitoring of multiple welding heads simultaneously).

Challenge: Complex CNNs can exceed the computational budget of embedded controllers. Model compression techniques such as pruning, quantization, or knowledge distillation reduce model size while preserving accuracy.

GPU (Graphics Processing Unit) accelerates parallel computations required by deep-learning inference. Modern welding inspection stations often integrate a dedicated GPU or an edge AI accelerator (e.G., NVIDIA Jetson series).

CUDA cores – parallel processing units within NVIDIA GPUs that execute matrix operations efficiently.

Tensor cores – specialized hardware for mixed-precision matrix multiplication, offering up to 8× speedup for FP16 operations.

Example: Deploying a TensorRT-optimized version of a U-Net segmentation model on a Jetson Xavier yields 120 fps inference, sufficient for high-speed laser welding monitoring.

Model compression reduces the memory footprint and computational demand of deep-learning models.

Pruning – removes redundant network connections based on low weight magnitude. Quantization – converts 32-bit floating-point weights to 8-bit integers, often with negligible loss in accuracy. Knowledge distillation – trains a smaller “student” model to mimic the outputs of a larger “teacher” model, achieving comparable performance with fewer parameters.

Challenge: Aggressive quantization may degrade detection of small cracks because subtle intensity differences become indistinguishable. Careful calibration with representative welding data mitigates this risk.

Transfer learning leverages knowledge from a pre-trained model (e.g., ImageNet) to accelerate training on a welding-specific dataset. By fine-tuning only the final layers, developers can achieve high accuracy with limited annotated images.

Feature extractor – the frozen convolutional base that provides generic visual features. Fine-tuning – updating weights of deeper layers to adapt to the target domain, often necessary when the welding images differ significantly from natural images.

Example: A ResNet-50 network pre-trained on ImageNet is fine-tuned on a set of 2 000 annotated weld defect images, reducing training time from weeks to a few hours while reaching 93 % classification accuracy.

Domain adaptation addresses the shift between training and deployment environments, such as differences in lighting, camera type, or material reflectivity.

Unsupervised adaptation – aligns feature distributions without requiring labeled target data, using techniques like adversarial training. Style transfer – synthetically modifies source images to resemble target domain characteristics, augmenting the training set.

Challenge: A model trained on flat-plate weld images may perform poorly on curved-tube welds due to geometric distortion. Applying a geometric augmentation pipeline that simulates curvature improves robustness.

Data augmentation artificially expands the training set by applying transformations to existing images, helping the model generalize to unseen conditions.

Geometric augmentations – include rotation, scaling, translation, and perspective warping, useful for varying joint angles. Photometric augmentations – adjust brightness, contrast, saturation, or add Gaussian noise to simulate varying arc intensities. Elastic deformation – mimics material warping that can occur in real welds,

preserving label consistency.

Example: During training, each weld image is randomly rotated between -15° and $+15^\circ$, brightened by up to 20%, and corrupted with salt-and-pepper noise. This yields a model resilient to camera misalignment and arc flicker.

Real-time processing denotes the ability to analyze images as they are captured, providing immediate feedback for control loops.

Pipeline parallelism – splits the processing stages (acquisition, pre-processing, inference) across multiple threads or cores to reduce overall latency. Batch size – in real-time scenarios, batch size is often set to 1 to minimize waiting time, though micro-batching can improve GPU utilization if latency budgets permit.

Practical implementation: A welding robot controller receives a processed bead width measurement every 20 ms, allowing it to adjust torch speed dynamically to maintain a constant bead profile.

Latency budget is the maximum allowable delay from sensor capture to actuation command. In welding, latency budgets are dictated by the dynamics of the melt pool and the speed of the welding head.

Deterministic latency – predictable and bounded, essential for safety-critical applications such as autonomous welding in aerospace manufacturing. Jitter – variability in latency, which can cause unstable control if not mitigated.

Challenge: Networked camera systems introduce variable transmission delays. Using a dedicated Ethernet link with real-time protocols (e.g., TSN – Time-Sensitive Networking) reduces jitter and ensures compliance with a 10 ms latency budget.

Heat-affected zone (HAZ) is the region of base material whose microstructure is altered by welding heat, but which does not melt. Vision systems can estimate HAZ size by analyzing temperature gradients in IR images.

Thermal gradient – the rate of temperature change across the material; steeper gradients often indicate a narrower HAZ. Isotherm mapping – contour lines of equal temperature derived from calibrated IR data, enabling quantitative HAZ measurement.

Example: An IR camera captures a sequence of temperature frames during a TIG pass. By extracting isotherms at 400°C and 600°C , the system computes the HAZ width and alerts the operator if it exceeds design limits.

Weld pool is the localized region of molten metal formed by the welding arc. Monitoring the pool shape and dynamics provides insight into process stability.

Pool geometry – includes width, length, and depth; deviations can cause under-fill or over-fill. Arc oscillation – periodic movement of the arc, which can be inferred from high-frequency variations in pool brightness.

Practical observation: In a robotic MIG welding trial, a sudden increase in pool width correlates with a drop

in welding voltage, suggesting a transition to a more conductive mode. The vision system flags the event for corrective action.

Bead reinforcement refers to the extra metal deposited above the base surface, forming a raised profile. Excessive reinforcement can lead to stress concentrations, while insufficient reinforcement may reduce joint strength.

Profile extraction – using laser line scanners or structured-light projection to capture the 3-D shape of the bead. Cross-sectional analysis – slicing the 3-D point cloud to measure reinforcement height at multiple locations.

Example: A laser triangulation sensor mounted beside the torch scans the solidified bead. The resulting point cloud is fitted with a spline, from which the maximum reinforcement of 1.2 Mm is compared against a specification of 1.0 ± 0.2 Mm.

Porosity is a common welding defect characterized by gas-filled cavities within the weld metal. Detecting porosity early can prevent costly re-work.

X-ray imaging – provides high-resolution volumetric data, but requires shielding and safety precautions. Ultrasonic testing (UT) – uses high-frequency sound waves; combined with image processing, UT data can be visualized as B-scans for defect detection.

Vision-based approach: A high-resolution X-ray detector captures radiographs of a weld joint. A CNN trained on labeled porosity patterns classifies each image as “acceptable” or “defective,” achieving a false-negative rate below 2%.

Crack detection is critical for safety-critical industries (e.g., Pressure vessels, aerospace). Cracks often appear as thin, high-contrast lines in optical or X-ray images.

Line detector – algorithms such as the Hough Transform identify linear features by voting in a parameter space. Frangi filter – enhances tubular structures, useful for emphasizing crack-like patterns before segmentation.

Challenge: Surface contaminants and spatter can produce line-like artifacts that confuse crack detectors. Combining morphology-based filtering with a confidence score derived from a secondary classifier reduces false positives.

Undercut is a groove formed at the weld toe, weakening the joint. Vision systems can measure undercut depth by analyzing the edge profile of the bead.

Depth map – obtained from stereo vision or structured-light techniques, providing per-pixel height information. Thresholded depth – regions where depth exceeds a preset limit are flagged as undercut.

Example: A stereo camera pair mounted on a welding robot captures the bead after each pass. The disparity map is converted to a depth map, and any area where the depth exceeds 0.3 Mm is highlighted for operator review.

Spatter consists of molten droplets expelled from the arc that solidify on surrounding surfaces. Excessive spatter can cause defects and require post-process cleaning.

Particle counting – segment spatter blobs and count them to assess process cleanliness. Size distribution – compute area or equivalent diameter for each spatter particle, providing insight into arc stability.

Practical workflow: After each weld, a vision system applies a color-threshold in HSV space to isolate bright metallic particles, then uses connected-component labeling to count and size them. If the average particle size exceeds a predefined threshold, the welding parameters are adjusted.

Process monitoring encompasses the continuous observation of welding parameters (current, voltage, speed) using visual cues.

Fusion of sensor data – combines visual information with traditional process signals to improve reliability. Anomaly detection – unsupervised models such as autoencoders learn normal visual patterns; deviations indicate potential faults.

Example: An autoencoder trained on normal arc images reconstructs each incoming frame. A high reconstruction error coincides with an unexpected voltage dip, prompting the controller to pause the weld and alert the operator.

Robotic welding integration refers to the seamless coupling of vision algorithms with robot motion planning and control.

Closed-loop control – the robot adjusts its trajectory based on real-time feedback from the vision system (e.g., To maintain constant bead width). Path planning – vision informs the optimal torch path, accounting for joint geometry and previously detected defects.

Scenario: During a multi-pass pipe welding operation, the vision system detects a small crack in the first pass. The robot replans the subsequent passes to avoid the defect zone, while a repair pass is scheduled later.

Quality assurance (QA) metrics derived from vision data provide quantitative evidence of weld conformity.

Defect density – number of detected defects per unit length of weld. Coverage ratio – proportion of the weld length that meets geometric specifications (e.g., Bead width within tolerance).

Implementation: After a production run, the system aggregates defect counts and computes a defect density of 0.02 Defects per meter, well below the industry threshold of 0.05, And automatically generates a QA report.

Calibration ensures that the visual measurements correspond to real-world dimensions and physical quantities.

Geometric calibration – determines camera intrinsic parameters (focal length, principal point) and extrinsic parameters (position, orientation) relative to the welding torch. Radiometric calibration – maps pixel

intensity to temperature for IR cameras, often using a black-body reference source.

Challenge: In a harsh welding environment, vibration and heat can cause camera mounts to shift, degrading calibration. Periodic automated recalibration using a known checkerboard pattern mounted on the workpiece mitigates drift.

Noise in welding images arises from multiple sources: Electrical interference from the welding power supply, photon shot noise in low-light conditions, and speckle from laser-based sensors.

Gaussian noise – statistical variation with a bell-shaped distribution, often modeled during algorithm design. Salt-and-pepper noise – impulsive disturbances that appear as isolated bright or dark pixels, typical of transmission errors.

Mitigation strategies: Applying a median filter removes salt-and-pepper noise, while a Wiener filter adapts to local variance and reduces Gaussian noise without overly blurring edges.

Illumination variance is a major challenge because the welding arc emits intense, flickering light that can saturate sensors.

High-dynamic-range (HDR) imaging – captures multiple exposures and merges them to retain detail in both bright and dark regions. Polarizing filters – reduce glare from the arc, improving contrast for downstream processing.

Example: An HDR camera captures three exposures (short, medium, long) during a weld pass. The merged HDR image preserves the bright arc core while revealing the surrounding bead geometry, enabling simultaneous monitoring of both.

Shadowing occurs when the torch or workpiece blocks illumination, creating dark regions that can be mistaken for defects.

Shadow compensation – uses illumination models to estimate and correct for shadowed areas. Multiple view fusion – combines images from different camera angles to fill in occluded regions.

Practical solution: Two cameras positioned on opposite sides of the torch provide overlapping fields of view. When a shadow is detected in one view, the system fills the missing information using the complementary view.

Texture analysis extracts statistical patterns from the surface appearance of welds, useful for classifying HAZ microstructures.

Gray-level co-occurrence matrix (GLCM) – quantifies how often pairs of pixel intensities occur at a given offset, yielding features such as contrast, homogeneity, and entropy. Local binary patterns (LBP) – encodes the relationship of a pixel to its neighbors, providing a compact texture descriptor.

Application: A classifier trained on GLCM features distinguishes between properly normalized HAZ and over-tempered zones, supporting process optimization.

3-D reconstruction creates spatial models of the weld bead and surrounding geometry.

Stereo vision – uses two cameras with known baseline to triangulate 3-D points. Structured light – projects a known pattern onto the surface; deformation of the pattern reveals depth information.

Challenge: The highly reflective metal surface can distort projected patterns, causing reconstruction errors. Applying a matte spray coating or using wavelength-specific illumination reduces specular reflections.

Machine vision pipeline outlines the sequential steps from raw image capture to final decision.

Pre-processing – includes denoising, illumination correction, and geometric rectification. Feature extraction – either hand-crafted or deep-learned descriptors. Classification / segmentation – assigns labels to pixels or regions. Post-processing – refines results with morphological operations or temporal smoothing.

A typical pipeline for defect detection might: (1) Acquire a high-speed RGB frame; (2) apply a median filter; (3) convert to HSV and threshold the Value channel; (4) run a small CNN to classify each candidate region; (5) apply a morphological opening to eliminate false positives; and (6) output a defect map to the welding controller.

Temporal smoothing leverages the continuity of video streams to reduce false alarms.

Running average – maintains a moving average of pixel intensities, helping to filter out transient arc flicker. Kalman filter – predicts the next state of a measured variable (e.g., Bead width) and corrects it with the new observation, providing a statistically optimal estimate.

Example: A Kalman filter tracks bead width over successive frames, smoothing out occasional spikes caused by sensor noise while still responding promptly to genuine width changes.

Edge-aware filters preserve important structural information while reducing noise.

Bilateral filter – smooths images while respecting intensity edges, ideal for preserving the sharp boundary of the weld pool. Guided filter – uses a guidance image (often the original) to direct smoothing, offering linear-time performance.

Use case: Applying a bilateral filter to an IR image reduces high-frequency sensor noise without blurring the temperature gradient at the HAZ boundary.

Statistical process control (SPC) integrates vision-derived metrics into traditional control charts.

Control limits – upper and lower thresholds derived from process variability; measurements outside these limits indicate an out-of-control condition. Process capability index (Cpk) – quantifies how well the process meets specification limits, using visual measurements such as bead width.

Implementation: The system logs bead width for each weld, updates an X-bar chart in real time, and triggers an alarm when the width exceeds the $\pm 3\sigma$ control limits.

Explainability (XAI) is increasingly important when AI systems make safety-critical decisions in welding.

Saliency maps – highlight image regions that contributed most to a CNN’s classification, allowing engineers to verify that the model focuses on the weld pool rather than background. Rule-based post-processing – combines AI outputs with deterministic logic (e.G., “If defect size > 2 mm then reject”) to provide transparent decision criteria.

Challenge: Deep models can be opaque; integrating saliency visualizations into the operator interface builds trust and facilitates troubleshooting.

Edge computing moves processing closer to the data source, reducing latency and bandwidth usage.

On-device inference – runs the AI model directly on the camera or an attached microcontroller, avoiding the need to stream raw frames to a central server. Model offloading – dynamically decides whether to process locally or send to a more powerful edge server based on current load and latency constraints.

Scenario: In a remote shipyard, a ruggedized vision module performs on-device spatter detection. When network connectivity is available, it offloads heavy segmentation tasks to a nearby edge server for detailed defect analysis.

Safety considerations are paramount because welding environments involve high temperatures, intense radiation, and moving equipment.

Protective enclosures – shield cameras from spatter and UV exposure, extending sensor lifespan. Fail-safe design – ensures that loss of vision data defaults the welding system to a safe state (e.G., Stop arc).

Example: A vision system mounted on a robotic arm includes a transparent sapphire window with anti-reflective coating, protecting the sensor while preserving optical clarity.

Regulatory compliance may dictate documentation of inspection methods and traceability of results.

Traceability matrix – links each measured parameter (e.G., Bead width) to relevant standards such as ISO 3834 or ASME IX. Audit logs – record timestamps, operator actions, and system decisions for later review.

Implementation: The vision software automatically generates an audit log entry each time a defect is detected, storing the original image, the processed mask, and the classification outcome, thereby satisfying audit requirements.

Future trends in computer vision for welding are shaped by emerging hardware, algorithmic advances, and integration with broader Industry 4.0 Ecosystems.

Neuromorphic sensors – event-based cameras that output asynchronous brightness changes, offering microsecond latency and low data rates, ideal for capturing rapid arc dynamics. Self-supervised learning – models that learn representations from unlabeled welding video, reducing the reliance on costly annotation. Digital twins – virtual replicas of welding cells that ingest vision data to predict outcomes and optimize parameters in simulation before deployment.

These developments promise more adaptive, efficient, and intelligent welding systems, where vision not only detects defects but also proactively guides process parameters to achieve optimal joint quality.

By mastering the terminology outlined above, learners will be equipped to navigate the complex landscape of computer-vision-enabled welding, design robust inspection pipelines, and contribute to the next generation of intelligent manufacturing.