

Computer Vision for Welding Applications

Active Contour – Concept: Edge-detecting curve that evolves to fit object boundaries in an image. Related terms: Snakes, level set, segmentation. Explanation: The algorithm initializes a contour near the feature of interest; internal forces maintain smoothness while external forces derived from image gradients attract the contour to edges. Example: In weld bead detection, an active contour can outline the bead profile despite varying illumination. Practical application: Real-time monitoring of bead geometry to adjust welding parameters. Challenges: Sensitive to initialization, may converge to local minima in noisy weld images.

Adaptive Thresholding – Concept: Dynamic determination of pixel intensity cut-off based on local image statistics. Related terms: Otsu's method, global threshold, binarization. Explanation: The image is divided into windows; each window computes a threshold (e.G., Mean or Gaussian weighted) and pixels are classified accordingly. Example: Separating spatter from the weld pool in infrared weld images where illumination varies across the field. Practical application: Enables robust detection of defects such as porosity under uneven lighting. Challenges: Window size selection affects performance; computational load increases with high-resolution frames.

Annotation – Concept: Manual or semi-automatic labeling of images with ground-truth information. Related terms: Labeling, dataset, supervised learning. Explanation: Annotators draw bounding boxes, polygons, or keypoints to indicate weld features (e.G., Bead edges, cracks). Example: A dataset of TIG welds annotated with crack locations for training a convolutional neural network (CNN). Practical application: Provides the necessary supervision for defect classification models. Challenges: Time-consuming, inter-annotator variability, requires domain expertise.

Artificial Neural Network (ANN) – Concept: Computational model inspired by biological neurons, consisting of layers of interconnected nodes. Related terms: Deep learning, perceptron, backpropagation. Explanation: Input images are transformed through weighted connections and activation functions to produce predictions. Example: An ANN predicts weld penetration depth from grayscale images captured by a high-speed camera. Practical application: Real-time quality assessment without explicit feature engineering. Challenges: Requires large labeled datasets; risk of overfitting on specific welding setups.

Background Subtraction – Concept: Technique to isolate foreground objects by removing a static background model. Related terms: Frame differencing, motion detection, foreground extraction. Explanation: A background image (or model) is built from a series of frames; each new frame is compared, and significant differences are classified as foreground. Example: Isolating the moving arc and droplet in a MIG welding video to focus analysis on the weld pool. Practical application: Reduces processing load by ignoring static surroundings. Challenges: Lighting changes, camera jitter, and welding sparks can corrupt the background model.

Batch Normalization – Concept: Normalization layer that stabilizes training by scaling and shifting activations. Related terms: Layer normalization, dropout, learning rate. Explanation: For each mini-batch,

mean and variance of activations are computed and used to normalize the output, followed by learned scale and shift parameters. Example: In a deep CNN for weld defect detection, batch normalization accelerates convergence and improves accuracy. Practical application: Enables deeper networks for subtle defect patterns. Challenges: Adds memory overhead; must be correctly handled during inference.

Bezier Curve – Concept: Parametric curve defined by control points, used for smooth shape representation. Related terms: Spline, NURBS, curve fitting. Explanation: The curve is a weighted sum of control points, providing an intuitive way to model continuous weld bead contours. Example: Fitting a Bezier curve to the extracted edge of a weld bead to compute curvature. Practical application: Quantitative assessment of bead uniformity. Challenges: Determining the optimal number of control points for complex bead shapes.

Binary Classification – Concept: Predicting one of two possible classes for each input sample. Related terms: Logistic regression, softmax, SVM. Explanation: The model outputs a probability that an image contains a defect (e.g., Crack) versus being defect-free. Example: A CNN classifies each frame as “spatter present” or “no spatter.” Practical application: Simple decision logic for automated welding robots. Challenges: Imbalanced datasets where defective samples are scarce.

Blob Detection – Concept: Identifying regions in an image that differ in properties such as brightness or texture. Related terms: Laplacian of Gaussian, Difference of Gaussians, region of interest. Explanation: Algorithms locate connected components that satisfy size and intensity criteria. Example: Detecting metal droplets in a laser welding video as bright blobs against a darker background. Practical application: Monitoring droplet formation frequency to infer process stability. Challenges: Distinguishing true droplets from noise caused by arc fluctuations.

Camera Calibration – Concept: Determining intrinsic and extrinsic parameters of a camera to correct geometric distortions. Related terms: Lens distortion, reprojection error, extrinsic matrix. Explanation: Using a known pattern (e.g., Checkerboard), the calibration routine computes focal length, principal point, and distortion coefficients. Example: Calibrating a high-speed camera mounted at a 45° angle to the weld joint. Practical application: Accurate measurement of weld dimensions in pixel space. Challenges: Heat-induced lens deformation and vibration during welding.

Convolution – Concept: Linear operation that applies a kernel (filter) across an image to extract local features. Related terms: Filter, stride, padding. Explanation: Each output pixel is the weighted sum of neighboring input pixels defined by the kernel. Example: Edge detection kernels highlight bead edges in grayscale images. Practical application: Foundation of CNN layers for feature extraction. Challenges: Choosing appropriate kernel sizes to capture both fine and coarse features.

Convolutional Neural Network (CNN) – Concept: Deep learning architecture that processes grid-like data using convolutional layers. Related terms: Feature map, pooling, fully connected layer. Explanation: CNNs learn hierarchical features—from edges to shapes—by stacking convolutions, activations, and down-sampling. Example: A CNN trained on thousands of weld images classifies defects such as porosity, cracks, and undercut. Practical application: End-to-end automated weld inspection. Challenges: Requires extensive labeled data, high computational resources for training, and careful architecture design to avoid over-parameterization.

Cross-Entropy Loss – Concept: Loss function measuring the difference between predicted probability distribution and true labels. Related terms: Log loss, classification loss, softmax. Explanation: For binary or multi-class tasks, the loss penalizes confident but incorrect predictions heavily. Example: During training of a weld defect detector, cross-entropy guides the network to assign high probability to the correct defect class. Practical application: Provides a smooth gradient for backpropagation. Challenges: Sensitive to label noise; class imbalance may require weighting.

Data Augmentation – Concept: Synthetic expansion of training data by applying transformations. Related terms: Rotation, scaling, flipping, noise injection. Explanation: Images are randomly rotated, shifted, or have contrast adjusted to mimic real-world variability. Example: Augmenting a limited set of arc welding images with random brightness changes to improve robustness. Practical application: Mitigates overfitting in deep models when acquiring large weld datasets is costly. Challenges: Over-augmentation can create unrealistic samples that degrade performance.

Depth Map – Concept: Image where each pixel encodes distance from the camera to the scene. Related terms: Stereo vision, LiDAR, disparity. Explanation: Using two cameras or structured light, disparity between corresponding points is converted to depth values. Example: Generating a depth map of the weld pool to assess melt-pool geometry. Practical application: 3-D reconstruction of weld seams for precise robot path planning. Challenges: Reflective metal surfaces cause speckle, leading to noisy disparity estimates.

Edge Detection – Concept: Identifying pixel locations where intensity changes sharply, indicating object boundaries. Related terms: Canny, Sobel, gradient magnitude. Explanation: Operators compute gradients and apply non-maximum suppression and hysteresis thresholds to produce thin edges. Example: Detecting the outer contour of a weld bead in a high-contrast image. Practical application: Provides inputs for contour-based measurements of bead width. Challenges: Noise from welding arc can generate spurious edges; parameter tuning is essential.

Feature Extraction – Concept: Process of deriving informative representations from raw image data. Related terms: Handcrafted features, descriptors, dimensionality reduction. Explanation: Techniques such as Histogram of Oriented Gradients (HOG) or Local Binary Patterns (LBP) capture texture and shape cues. Example: HOG descriptors encode the orientation of gradients along the weld seam to differentiate good versus defective welds. Practical application: Used in traditional machine-learning pipelines where deep learning is infeasible. Challenges: Handcrafted features may not generalize across welding processes.

Gaussian Blur – Concept: Smoothing filter that reduces high-frequency noise while preserving edges modestly. Related terms: Low-pass filter, kernel, sigma. Explanation: Convolution with a Gaussian kernel averages pixel values weighted by distance from the center. Example: Pre-processing weld images to suppress arc flicker before edge detection. Practical application: Improves stability of subsequent segmentation steps. Challenges: Excessive blurring can erase fine defect details.

Gradient Descent – Concept: Optimization algorithm that iteratively updates model parameters to minimize a loss function. Related terms: Learning rate, momentum, stochastic gradient descent. Explanation: Parameters move opposite to the gradient of the loss with respect to those parameters. Example: Training a CNN for weld defect detection using stochastic gradient descent with momentum. Practical application:

Core method for fitting deep models to welding image data. Challenges: Choosing appropriate learning rates; risk of getting trapped in local minima.

Histogram Equalization – Concept: Contrast enhancement technique that spreads intensity values across the full range. Related terms: CLAHE, intensity scaling, dynamic range. Explanation: The cumulative distribution of pixel intensities is remapped to a uniform distribution. Example: Enhancing low-contrast infrared weld images to reveal subtle temperature gradients. Practical application: Improves visibility of defects in thermal images. Challenges: May amplify noise; global equalization can over-enhance bright regions.

Homography – Concept: Projective transformation mapping points from one plane to another. Related terms: Perspective transform, planar registration, warp. Explanation: Represented by a 3×3 matrix, it aligns images taken from different viewpoints. Example: Aligning a side-view camera image with a top-view image of the same weld joint for multi-view analysis. Practical application: Enables fusion of complementary visual information. Challenges: Requires accurate correspondence points; metal surface reflectivity can hinder reliable matching.

Image Pyramid – Concept: Multi-scale representation of an image created by repeatedly down-sampling. Related terms: Gaussian pyramid, Laplacian pyramid, scale space. Explanation: Lower-resolution levels capture coarse structures; higher levels retain fine details. Example: Detecting cracks of varying widths by scanning across pyramid levels with a fixed-size filter. Practical application: Provides scale invariance for defect detection. Challenges: Increases memory usage; down-sampling may lose small defect information.

Intersection over Union (IoU) – Concept: Metric measuring overlap between predicted and ground-truth bounding boxes. Related terms: Jaccard index, localization accuracy, evaluation metric. Explanation: $\text{IoU} = (\text{area of overlap}) / (\text{area of union})$. Example: An object detector for weld spatter achieves an IoU of 0.78 on the test set. Practical application: Thresholds IoU to decide whether a detection is correct. Challenges: Sensitive to annotation precision; small objects can yield low IoU even with accurate localization.

K-Nearest Neighbors (KNN) – Concept: Instance-based classifier that assigns a label based on majority vote of nearest samples. Related terms: Distance metric, lazy learning, classification. Explanation: Feature vectors of test images are compared to training vectors; the label of the majority among the K closest is chosen. Example: Classifying weld images into “acceptable” or “defective” using HOG features and $K=5$. Practical application: Simple baseline for weld quality prediction. Challenges: Computationally intensive at inference; performance degrades with high-dimensional features.

Kernel Density Estimation (KDE) – Concept: Non-parametric method to estimate probability density of data points. Related terms: Bandwidth, smoothing, probability distribution. Explanation: Each data point contributes a kernel (e.g., Gaussian) and the sum approximates the underlying distribution. Example: Estimating the distribution of bead width measurements from vision data to detect process drift. Practical application: Provides statistical insight for process control. Challenges: Choice of bandwidth influences bias-variance trade-off; high dimensionality reduces accuracy.

Laser Triangulation – Concept: Optical ranging technique that measures distance by intersecting a laser beam with a camera’s field of view. Related terms: Structured light, depth sensing, profilometry. Explanation:

The displacement of the laser line on the sensor is proportional to surface height. Example: Capturing the cross-section profile of a weld pool in real time. Practical application: Precise measurement of bead geometry for closed-loop control. Challenges: High temperatures and plasma can scatter the laser, degrading signal quality.

Leaky ReLU – Concept: Activation function that allows a small gradient when the unit is not active. Related terms: ReLU, activation, non-linearity. Explanation: $F(x) = x$ for $x > 0$, else αx with α small (e.g., 0.01). Example: Used in deep CNNs for weld defect detection to mitigate dead neuron problems. Practical application: Improves gradient flow in deep networks. Challenges: Choice of α affects sparsity and performance.

Linear Discriminant Analysis (LDA) – Concept: Dimensionality reduction technique that maximizes class separability. Related terms: Fisher's criterion, projection, classification. Explanation: Projects data onto a lower-dimensional space where between-class variance is maximized relative to within-class variance. Example: Reducing HOG feature vectors of weld images before feeding them to a support vector machine. Practical application: Accelerates training of traditional classifiers. Challenges: Assumes Gaussian class distributions; may not capture non-linear relationships.

Local Binary Patterns (LBP) – Concept: Texture descriptor that encodes the relationship of a pixel with its neighbors. Related terms: Histogram, texture analysis, feature descriptor. Explanation: For each pixel, neighboring intensities are thresholded against the center value, forming a binary code. Example: LBP histograms differentiate smooth weld beads from rough, defect-prone surfaces. Practical application: Fast, illumination-invariant feature for defect classification. Challenges: Sensitive to noise; limited discriminative power for complex patterns.

Mask R-CNN – Concept: Instance segmentation network extending Faster R-CNN to predict pixel-wise masks. Related terms: Region proposal network, segmentation, object detection. Explanation: After generating bounding boxes, a small FCN predicts a binary mask for each Region of Interest. Example: Segmenting individual spatter particles in a MIG welding image. Practical application: Enables precise quantification of spatter area and count. Challenges: Requires extensive annotated masks; higher computational cost than bounding-box detectors.

Mean Shift – Concept: Mode-seeking algorithm that iteratively shifts data points toward regions of higher density. Related terms: Clustering, bandwidth, kernel density. Explanation: Each point moves to the mean of points within a kernel window until convergence. Example: Grouping detected edge points to identify continuous weld seams. Practical application: Unsupervised identification of weld trajectories. Challenges: Bandwidth selection critically influences cluster granularity; may merge separate seams if too large.

Mid-Infrared Imaging – Concept: Thermal imaging using wavelengths typically between 3–5 μm . Related terms: IR camera, emissivity, temperature mapping. Explanation: Captures emitted radiation from hot surfaces, providing temperature distribution. Example: Monitoring the cooling rate of a weld pool to detect potential cracking zones. Practical application: Non-contact temperature monitoring for process optimization. Challenges: Requires calibration for metal emissivity; atmospheric absorption can affect accuracy.

Multispectral Vision – Concept: Imaging across several spectral bands beyond visible light. Related terms: Hyperspectral, band selection, sensor fusion. Explanation: Different wavelengths reveal distinct material properties; combining them enhances defect detection. Example: Using visible and near-IR bands to distinguish oxidation from spatter. Practical application: Improves robustness of defect classification under varying lighting. Challenges: Increased data volume; need for band-specific calibration.

Neural Architecture Search (NAS) – Concept: Automated method for discovering optimal network structures. Related terms: Hyperparameter optimization, search space, reinforcement learning. Explanation: Algorithms explore combinations of layers, connections, and operations to maximize performance on a validation set. Example: NAS identifies a lightweight CNN tailored for on-board weld inspection hardware. Practical application: Generates efficient models for embedded welding robots. Challenges: Computationally expensive; risk of overfitting to specific datasets.

Optical Flow – Concept: Pixel-wise motion estimation between consecutive frames. Related terms: Lucas-Kanade, Horn-Schunck, dense flow. Explanation: Assumes brightness constancy and computes displacement vectors that best explain intensity changes. Example: Tracking the progression of the weld pool front in a high-speed video. Practical application: Deriving melt-pool velocity for adaptive control. Challenges: Strong illumination fluctuations and plasma occlusions can break brightness constancy assumption.

Over-fitting – Concept: Model learns noise and specific patterns in training data, reducing generalization. Related terms: Regularization, validation loss, model complexity. Explanation: Training error continues to decrease while validation error starts to increase. Example: A deep CNN achieves 99% accuracy on training weld images but only 70% on unseen data. Practical application: Indicates need for more data, dropout, or simpler architecture. Challenges: Detecting subtle over-fitting when validation set is small.

Patch-Based Learning – Concept: Training models on small image regions (patches) rather than whole frames. Related terms: Sliding window, local context, crop. Explanation: Each patch is labeled and fed to the network, allowing focus on fine-grained features. Example: Classifying 64 × 64 pixel patches for presence of micro-cracks. Practical application: Enables detection of small defects that might be diluted in full-image predictions. Challenges: Requires careful patch selection to avoid class imbalance; stitching predictions can produce seams.

Principal Component Analysis (PCA) – Concept: Linear dimensionality reduction method that projects data onto orthogonal axes of maximal variance. Related terms: Eigenvectors, covariance matrix, feature compression. Explanation: The first few principal components capture most variability, allowing compact representation. Example: Reducing high-dimensional texture descriptors from weld images before feeding to a classifier. Practical application: Lowers computational load and mitigates the curse of dimensionality. Challenges: Assumes linear relationships; may discard subtle defect cues.

Quadratic Discriminant Analysis (QDA) – Concept: Classification technique modeling each class with its own covariance matrix. Related terms: Gaussian classifier, decision boundary, discriminant function. Explanation: Computes posterior probabilities assuming multivariate normal distributions with class-specific covariances. Example: Distinguishing between “porosity” and “underfill” defects using statistical features extracted from

weld images. Practical application: Provides probabilistic output useful for risk-based decision making. Challenges: Requires sufficient samples per class to estimate covariances reliably.

Region of Interest (ROI) – Concept: Sub-area of an image selected for focused processing. Related terms: Cropping, masking, attention. Explanation: By limiting analysis to the weld zone, irrelevant background is ignored, improving efficiency. Example: Defining an ROI that tightly encloses the weld bead based on prior knowledge of joint geometry. Practical application: Reduces false positives from surrounding metal structures. Challenges: ROI must adapt to varying joint positions; inaccurate ROI can exclude relevant features.

ResNet – Concept: Residual network architecture that facilitates training of very deep models via shortcut connections. Related terms: Skip connection, identity mapping, deep learning. Explanation: Layers learn residual functions relative to the input, alleviating vanishing gradient problems. Example: A ResNet-34 model classifies weld defects with high accuracy while maintaining reasonable inference speed. Practical application: Enables deployment of deep models on industrial PCs. Challenges: Increased parameter count; careful tuning of learning rates required.

Reinforcement Learning (RL) – Concept: Learning paradigm where an agent interacts with an environment to maximize cumulative reward. Related terms: Policy, Q-learning, Markov decision process. Explanation: The agent selects actions (e.g., Welding speed adjustments) based on visual state observations and receives feedback. Example: An RL controller uses real-time vision of bead geometry to decide voltage changes. Practical application: Adaptive welding processes that self-optimize. Challenges: Designing appropriate reward functions; ensuring safety during exploration.

Scale-Invariant Feature Transform (SIFT) – Concept: Keypoint detector and descriptor robust to scale, rotation, and illumination changes. Related terms: Keypoint, descriptor, feature matching. Explanation: Detects extrema in Difference of Gaussian space and encodes local gradients into a 128-dimensional vector. Example: Matching SIFT features between a reference weld image and a live frame to assess alignment. Practical application: Provides reliable correspondence for homography estimation. Challenges: Computationally intensive; patent restrictions may limit commercial use.

Semantic Segmentation – Concept: Pixel-level classification assigning each pixel to a predefined class. Related terms: Fully convolutional network, encoder-decoder, class map. Explanation: The network outputs a map where each pixel label indicates "weld pool," "spatter," "background," etc. Example: A U-Net model segments the melt pool region from surrounding metal in infrared images. Practical application: Enables precise measurement of pool area and shape. Challenges: Requires dense annotations; class imbalance (few defect pixels) can hinder learning.

Shape Context – Concept: Descriptor capturing the distribution of points relative to a reference point, useful for shape matching. Related terms: Contour descriptor, point set, similarity measure. Explanation: A log-polar histogram records relative positions of neighboring points. Example: Comparing shape contexts of bead outlines to detect deviations from the desired profile. Practical application: Automated verification of weld bead geometry against design specifications. Challenges: Sensitive to noise in contour extraction; computational cost for large point sets.

SIFT – Concept: See Scale-Invariant Feature Transform entry. Related terms: Feature matching, keypoint detection. Explanation: (Refer to the earlier definition.) Example: (Refer to earlier example.) Practical application: (Refer to earlier application.) Challenges: (Refer to earlier challenges.)

Support Vector Machine (SVM) – Concept: Supervised learning model that finds the hyperplane maximizing margin between classes. Related terms: Kernel trick, soft margin, decision boundary. Explanation: By mapping data into higher dimensions via kernels, SVM can separate non-linearly separable data. Example: Using a radial basis function kernel to classify weld images into “acceptable” vs. “Defective” based on texture features. Practical application: Effective with limited training data where deep learning is infeasible. Challenges: Scaling to large datasets; selection of kernel and hyperparameters critical.

TensorFlow – Concept: Open-source machine learning framework for building and deploying models. Related terms: Keras, computational graph, GPU acceleration. Explanation: Provides high-level APIs for constructing CNNs, training loops, and inference pipelines. Example: Implementing a weld defect detector with TensorFlow’s eager execution for rapid prototyping. Practical application: Enables integration of vision models into welding robot controllers. Challenges: Version compatibility; memory management for high-resolution images.

Thermal Camera Calibration – Concept: Process of aligning temperature readings with known reference temperatures. Related terms: Emissivity correction, blackbody, radiometric calibration. Explanation: A calibrated blackbody source provides reference temperatures; the camera’s response curve is adjusted accordingly. Example: Calibrating a 3- μm infrared camera using a heated steel plate to ensure accurate weld pool temperature measurements. Practical application: Guarantees reliable temperature data for process control. Challenges: Emissivity of molten metal changes rapidly; environmental reflections can introduce errors.

Transfer Learning – Concept: Reusing a model pretrained on a large dataset for a related task with limited data. Related terms: Fine-tuning, feature reuse, domain adaptation. Explanation: Early layers capture generic visual features; only later layers are retrained on weld-specific data. Example: Adapting a ResNet pretrained on ImageNet to classify weld defects using a small annotated dataset. Practical application: Reduces data collection effort while achieving high performance. Challenges: Domain shift between natural images and welding imagery may limit transferability.

U-Net – Concept: Encoder-decoder CNN architecture designed for biomedical segmentation, adapted for weld image segmentation. Related terms: Skip connections, deconvolution, segmentation mask. Explanation: The contracting path captures context, while the expanding path restores spatial resolution, with skip connections preserving fine details. Example: Segmenting the weld pool and spatter in high-speed visible-light videos. Practical application: Provides accurate pixel-wise masks for quantitative analysis. Challenges: Requires extensive annotated masks; memory consumption grows with image size.

Unsupervised Anomaly Detection – Concept: Identifying out-of-distribution samples without explicit defect labels. Related terms: Autoencoder, one-class SVM, reconstruction error. Explanation: Models learn the distribution of normal weld images; high reconstruction error signals anomalies. Example: An autoencoder trained on defect-free weld images flags frames with unexpected crack patterns. Practical application: Early

warning system for unseen defect types. Challenges: Defining a robust threshold; false alarms caused by process variations.

Vanishing Gradient – Concept: Phenomenon where gradients become extremely small as they propagate through many layers, hindering learning. Related terms: Deep networks, activation function, gradient flow. Explanation: In deep networks with saturating activations, back-propagated error diminishes, preventing weight updates. Example: Early CNNs for weld inspection struggled to train beyond 5 layers due to vanishing gradients. Practical application: Motivates use of residual connections or ReLU variants. Challenges: Requires architectural tricks (e.g., Batch normalization) to mitigate.

Video Frame Interpolation – Concept: Generating intermediate frames to increase temporal resolution of video sequences. Related terms: Optical flow, motion compensation, super-resolution. Explanation: Algorithms estimate motion vectors and synthesize new frames that smoothly transition between existing ones. Example: Doubling the frame rate of a 500 fps welding video to 1000 fps for finer analysis of droplet dynamics. Practical application: Enables detailed study of rapid phenomena without expensive hardware. Challenges: Interpolation artifacts can mislead defect detection algorithms.

Visual Servoing – Concept: Closed-loop control method that uses visual feedback to guide robot motion. Related terms: Image-based servoing, pose estimation, feedback control. Explanation: The robot adjusts its trajectory based on real-time image measurements of the weld joint. Example: Aligning a welding torch to a joint edge detected by computer vision while the robot moves. Practical application: Increases precision of automated welding on complex geometries. Challenges: Latency in image processing; robustness to lighting changes and occlusions.

YOLO (You Only Look Once) – Concept: Real-time object detection framework that predicts bounding boxes and class probabilities in a single forward pass. Related terms: Single-shot detector, anchor boxes, non-maximum suppression. Explanation: The image is divided into a grid; each cell predicts multiple bounding boxes with associated confidence scores. Example: Detecting spatter, cracks, and undercut in a single frame of a MIG welding video at 60 fps. Practical application: Enables on-line quality monitoring with minimal latency. Challenges: Requires careful anchor box design for small weld defects; trade-off between speed and detection accuracy.

Z-Score Normalization – Concept: Statistical scaling that centers data around zero mean and unit variance. Related terms: Standardization, feature scaling, outlier detection. Explanation: For each feature, subtract the mean and divide by the standard deviation. Example: Normalizing intensity histograms of weld images before feeding them to a classifier. Practical application: Improves convergence of gradient-based learning algorithms. Challenges: Sensitive to outliers; mean and variance must be computed on training data only.