SURF CON INSPECT

WHITE PAPER



Open data, synthetic data, and domain adaptation: why we need better data

SurfConInspect White Paper
Series 9/10

About Us

WHO WE ARE

SurfConInspect (SCI) is a European project funded by the Research Fund for Coal and Steel (RFCS). Our consortium brings together leading steel producers, research institutes, technology developers and innovation experts.

The partners combine industrial expertise, academic knowledge and advanced digital solutions to transform the way steel quality is controlled.



OUR MISSION

The steel industry faces increasing pressure to become more sustainable, efficient and competitive—key goals of the European Green Deal.
Our mission is to enable zero-defect manufacturing in flat steel production by:

- Detecting surface defects earlier and more accurately.
- Reducing waste and improving resource efficiency.
- Supporting operators with advanced visualization tools.
- Providing real-time control actions directly in production lines.

MAIN ACTIVITIES

- Industrial trials in tin-plate and automotive steel production.
- Development of AI-based inspection methods.
- Application of AR for real-time operator support.
- Creation of open datasets to enhance research and industrial collaboration.
- Dissemination of results to industry, research and society through events, publications and digital platforms.

EXPECTED RESULTS

- Zero-defect production through optimized inspection and control.
- Reduced waste and energy consumption, contributing to the
  European Green Deal.
- Higher competitiveness for European steel producers.
- Transferable technologies with potential applications beyond the steel sector.
- Improved working conditions, with operators supported by innovative visualization and decision-making tools.

Executive Summary

Classification performance in automatic surface inspection for flat steel production is constrained, in practice, far more by data quality than by classification architecture. Over the past decade, the research community has invested substantially in algorithmic development, convolutional neural network architectures, attention mechanisms, transfer learning schemes, and the improvements in accuracy on the standard published benchmarks have been real. What these benchmarks do not capture is the behaviour of the same classifiers when deployed on production lines whose imaging conditions, material grades, and defect frequency distributions differ from those of the datasets on which the classifiers were trained. That divergence is where the practical performance deficit originates, and it is a data problem rather than an architectural one.

Three distinct sources drive this deficit. Critical defect categories, dents, shell defects, severe abrasions, are rare in well-controlled production environments and therefore underrepresented in training datasets assembled from routine inspection records. The resulting class imbalance biases classifiers toward non-critical assignments for the events that carry the highest downstream risk. A separate problem arises when a classifier trained at one facility or on one inspection system is deployed at another: systematic differences in sensor characteristics, illumination geometry, and surface texture shift the image statistics of each defect category in ways that degrade classification performance without any change in the physical defects themselves. And even within a single facility, the absence of shared, production-representative benchmark datasets means that the research community has no common reference against which to evaluate whether claimed algorithmic advances translate to genuine improvements under industrial conditions.

SurfConInspect addresses all three of these data challenges through a coordinated set of strategies whose development is documented in this white paper. Synthetic data generation through generative deep learning models produces training examples for underrepresented defect categories under the wet-surface and variable-illumination conditions that characterise pickling-line inspection, the conditions, described in WP2, where the gap between training and production distributions is widest. Domain adaptation methods transfer classification knowledge between the thyssenkrupp Rasselstein (RAS) and Salzgitter Flachstahl (SZFG) inspection installations, reducing the relabelling cost of cross-facility deployment. And the open data initiative assembles a labelled dataset of at least 30,000 surface defect images from the consortium's industrial partners and makes it available to the research community as a production-representative benchmark. The connection between these data strategies and the hybrid classification architecture described in WP7 is direct: the architecture is designed to extract maximum classification improvement from the datasets that these strategies can realistically produce.

1. The data dependency of image-based classification
2. The scarcity problem: critical defects in well-controlled production
3. The domain gap: classifiers across systems and facilities
4. Synthetic data generation: building the missing training distribution
5. Domain adaptation in practice: bridging RAS and SZFG
6. The open data initiative
7. Data quality and the design of the hybrid classification architecture
8. Connecting data quality to the quality control architecture
9. The SurfConInspect perspective
10. Conclusions

The data dependency of image-based classification

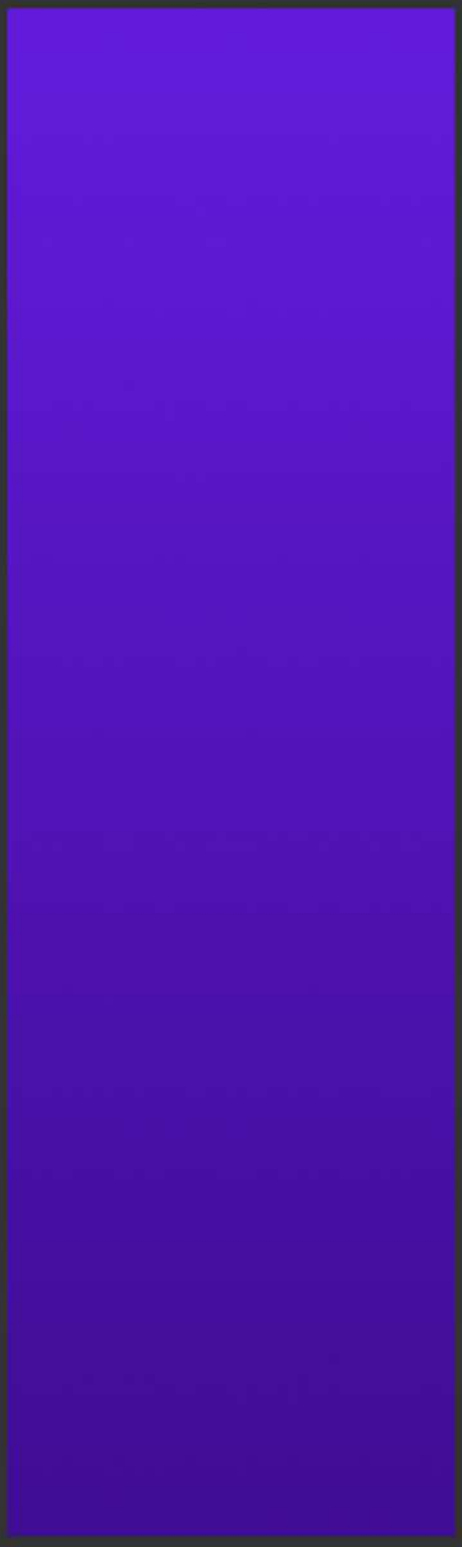
Classification systems in automatic surface inspection are trained through supervised learning: given a collection of labelled images, the classifier learns to associate visual patterns with defect categories by minimising prediction error across the training set. The quality of this association, expressed as the classifier's ability to assign the correct category to previously unseen examples, depends on how well the training collection represents the distribution of defect appearances encountered in production. When the training distribution and the production distribution diverge, the classifier's decision boundaries, which were learned to separate categories as they appear in training data, no longer correctly separate them as they appear in production. The magnitude of the resulting performance degradation is proportional to the degree of divergence.



In pickling-line inspection environments of the kind described in WP5, the sources of divergence between training and production distributions are well characterised. Wet surfaces modify the reflectance contrast of defect images relative to dry conditions, as established in WP2: even a micrometric water film alters the local intensity distribution through absorption within the film and reflection at the liquid-air interface, shifting defect image statistics in ways that move defect appearances away from the regions of feature space where the classifier has been trained. Changes in material grade alter surface texture and baseline reflectance. Gradual wear and recalibration of sensor components, illumination sources, camera gain settings, optical window fouling, introduce slow shifts in image statistics that accumulate over the operational lifetime of the system. And the deployment of a classifier trained at one facility onto a line at a different facility combines all these sources of divergence simultaneously.

The practical consequence is one that experienced inspection engineers at flat steel facilities recognise: classifiers that achieve high accuracy on validation datasets consistently underperform when measured against the downstream quality outcomes, coil break rates at tandem rolling mills, coating defect frequencies at finishing lines, customer claim rates, that constitute the true operational metric of classification quality. As documented in WP4, the defects most directly associated with these downstream outcomes are dents, shell defects, and abrasions in the near-edge zone of the strip, precisely the categories whose visual appearance is most sensitive to surface state variability and therefore most likely to diverge between training and production conditions. Closing the gap between benchmark accuracy and operational performance requires improving the match between the training distribution and the production distribution, a goal that cannot be achieved by architectural refinements to the classifier alone.

The scarcity problem: critical defects in well-controlled production



The most direct manifestation of the training data problem in surface inspection is the scarcity of labelled examples for the defect categories that matter most operationally. A pickling line processing material within its specified parameters generates relatively few occurrences of heavy topographic defects per coil, dents, shell defects, severe abrasions, compared to the volume of non-critical surface events that populate the inspection record at far higher frequency: liquid droplet reflections, minor scale patches, background structures, transient illumination variations. This frequency distribution reflects the effectiveness of upstream process control at preventing critical defects, and it is a desirable operational condition. Its consequence for classifier training is a class imbalance that directly undermines the classifier's ability to learn reliable representations of the minority classes.

A classifier trained on an imbalanced dataset, where critical defect categories constitute a few percent of training examples and non-critical categories constitute the remainder, develops decision boundaries that are positioned to minimise error on the majority class. Ambiguous cases, which are precisely the ones where critical defects under difficult imaging conditions fall, get classified as non-critical, because that is the assignment that minimises the training loss on the predominantly non-critical training set. The result is an elevated false negative rate for critical categories: not because the classifier is generally inaccurate, but because the imbalance in its training data has shaped its decision boundaries away from the regions where critical defects actually occur under variable production conditions.

Collecting more data from routine production does not correct this imbalance, it reproduces it, because the imbalance reflects the actual frequency distribution of defects in production. What is required is targeted augmentation of the training set with additional examples of the underrepresented critical categories, either by mining historical inspection records for rare events, by imaging deliberately induced defects under controlled conditions, or by generating synthetic examples computationally. The last option is the most scalable and the most flexible in terms of the surface conditions that can be represented, but its effectiveness depends on the realism of the generated images, a criterion that generative models can approach but not guarantee without validation against real production data.

The scarcity problem interacts with the wet-surface imaging challenge described in WP2. A dent on a strip section covered by a thin, uniform water film produces an image that differs from the same dent imaged under dry conditions or under a thick, inhomogeneous film. A training dataset that contains only dry-condition images of dents does not represent the full range of wet-condition appearances that the same defect type produces in production. The resulting false negatives in wet-condition inspection are attributable not to the scarcity of dents in absolute terms but to the scarcity of wet-condition dent images specifically, a more precise characterisation of the data deficit that points directly to the need for surface-state-stratified training data.

The domain gap: classifiers across systems and facilities

When a classifier trained on data from one ASIS installation is deployed on a different installation, at the same facility after sensor recalibration, or at a different facility, the images it receives in production differ systematically from those on which it was trained. This systematic difference, the domain gap, degrades classification performance in proportion to its magnitude, and its sources in surface inspection are well characterised.

Differences in camera sensor characteristics, quantum efficiency, dynamic range, noise floor, produce images with different signal-to-noise ratios for defects of the same geometric severity. Differences in illumination geometry, the angle and spectral composition of the light source relative to the strip surface, change the reflectance signature of surface features in ways that shift the feature distributions of each defect category. Differences in surface texture between material grades alter the background against which defects must be detected, affecting the contrast and spatial frequency content of defect images. Process condition differences, acid concentration, the degree of surface wetness at the inspection point, introduce the same sources of image variability documented in WP2 and WP8, superimposed on the sensor-level differences between systems. The domain gap between the RAS pickling line in Andernach and the SZFG electrolytic galvanising line in Salzgitter, the two industrial inspection environments central to SurfConInspect, incorporates all of these sources simultaneously.

The standard approach to managing domain gaps in industrial inspection has been manual recalibration: after deploying a classifier from one system onto another, inspection engineers relabel a representative set of images from the new system and retrain or fine-tune the classifier on the relabelled data. This process is effective but expensive. The cumulative cost of manual recalibration across the multiple lines operated by industrial partners with several production lines each represents a significant fraction of total inspection system maintenance effort. For the SurfConInspect consortium, which includes RAS and SZFG as industrial partners each operating multiple ASIS installations, this cost is substantial in absolute terms and constitutes a practical constraint on the frequency with which classifiers can be updated to reflect current production conditions.

Domain adaptation methods reduce this cost by learning the transformation between source and target image distributions without requiring target-domain labels. The theoretical foundations of unsupervised domain adaptation, adapting a classifier trained on labelled source data to an unlabelled target domain, are well established in the statistical learning literature. Methods based on discrepancy minimisation between source and target feature distributions, including approaches based on the Sliced Wasserstein Discrepancy have demonstrated effective transfer of classification knowledge across domains with moderate divergence.

Synthetic data generation: building the missing training distribution

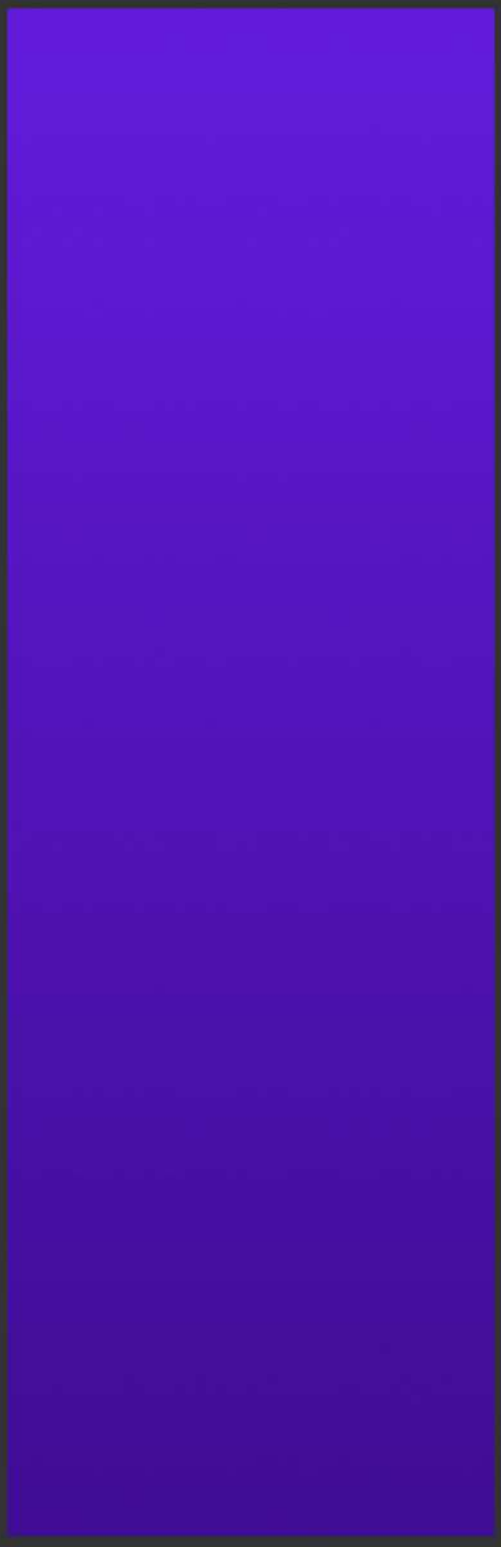
Synthetic data generation produces training examples computationally for defect categories and surface conditions that are underrepresented in the real inspection record. For surface inspection in pickling-line environments, the target of synthetic generation within SurfConInspect is specifically the combination of critical defect types, heavy dents, shell defects, with the wet-surface imaging conditions that characterise the RAS inspection zone. This combination is the one that produces the most severe training data deficit: critical defect images acquired under wet conditions are rare in the real training set because critical defects are rare in production, and because wet-condition acquisitions represent only a fraction of the imaging conditions encountered during the operational year.



The generation of synthetic defect images for ASIS training has been developed through several approaches in the literature. The accuracy of a simulation-based method in which defect images are rendered by superimposing a physically modelled defect profile onto a measured background texture depends on the fidelity of the physical model, a requirement that aligns well with the three-dimensional measurement data available from the combined inspection system described in WP5, where the depth and profile of each defect are directly measured and can serve as geometric inputs to the rendering model. The limitation is that the rendering requires detailed knowledge of the illumination geometry and sensor characteristics of the target system, which must be independently measured and may change over the operational lifetime of the system.

Generative approaches based on Variational Auto-Encoders (VAE) and Generative Adversarial Networks (GAN) learn the statistical structure of the training image distribution directly from data and generate new examples by sampling from the learned distribution. Within SurfConInspect, VAE and GAN architectures are trained on the available real examples of heavy dents and shell defects from the RAS and SZFG inspection systems and used to generate additional training examples representing the same defect types under a wider range of surface states, including wet-surface variants for which the real training set is most deficient. Generated images are evaluated against real production images using the Fréchet Inception Distance between real and synthetic distributions, and are reviewed by inspection engineers familiar with the appearance of these defect types in production before inclusion in the training set. This combination of quantitative and qualitative validation guards against the inclusion of synthetic examples that are statistically plausible but physically unrealistic.

Domain adaptation in practice: bridging RAS and SZFG



The domain gap between the RAS pickling line at Andernach and the SZFG electrolytic galvanising line at Salzgitter is substantial in every dimension that affects ASIS image statistics. The surface texture of hot-rolled tin-plate substrate differs from that of cold-rolled automotive-grade coil in its roughness profile and spatial frequency content. The illumination configurations of the two inspection systems are adapted to the specific defect populations that each line generates, introducing systematic differences in the contrast and spatial frequency structure of the images they produce. A classifier trained exclusively on RAS data will not perform reliably at SZFG without adaptation, and the reverse holds equally.

Domain adaptation for the RAS-to-SZFG transfer targets the defect categories that occur at both facilities: dents, scratches, and shell defects arising from the common hot rolling and pickling process history shared by all flat steel coils before they diverge toward tin-plate or automotive-grade finishing. For these shared categories, the domain gap is a matter of image statistics, the appearance of the same physical defect type under different surface and illumination conditions, rather than of fundamentally different defect morphologies. Domain adaptation methods can bridge this statistical gap without requiring target-domain labels by learning a feature representation in which the image distributions of source and target domains overlap sufficiently for the source-trained classifier's category knowledge to generalise.

The Sliced Wasserstein Discrepancy provides a computationally tractable measure of this distributional overlap that scales to the high-dimensional feature spaces generated by convolutional neural network encoders. Minimising the SWD between source (RAS) and target (SZFG) feature distributions, while maintaining the discriminative structure of the source classifier's category boundaries, yields a classifier that retains the category associations learned from RAS labelled data while adapting its internal representations to the statistical properties of SZFG images. The adaptation is unsupervised with respect to SZFG defect labels: only unlabelled SZFG images are required for the adaptation procedure.

The practical scope of domain adaptation is bounded by the degree to which source and target domains share the same defect categories. For the non-topographic defect types specific to the SZFG galvanising environment, surface conditions related to zinc coating adhesion and electrolytic process residues, addressed in WP2, domain adaptation from RAS data is inapplicable. RAS pickling data contains no examples of these categories, and no amount of distributional alignment can transfer category knowledge that was never present in the source training set. For these SZFG-specific categories, labelled training data must be acquired directly from the SZFG inspection system, and the domain adaptation component of the SurfConInspect data strategy is complemented by targeted data collection at SZFG for the categories that lie outside the shared defect taxonomy.

The open data initiative

The research literature on surface defect classification in steel production has long been constrained by the limited scale and representativeness of publicly available datasets. The NEU surface defect dataset contains 1,800 images across six defect classes; the SEVERSTAL Kaggle dataset contains approximately 12,500 images across four classes. These datasets have enabled meaningful algorithmic research and established common evaluation protocols, but they are orders of magnitude smaller than the training sets used by commercial ASIS vendors, and they cover a fraction of the defect classes present in productive systems, which routinely classify more than fifty distinct defect types. Classifier accuracy figures reported against these benchmarks provide limited guidance on the performance that the same classifiers will achieve in production, where the class structure, the surface conditions, and the frequency distributions of events all differ substantially from those of the published benchmarks.

The open data initiative within SurfConInspect assembles a labelled dataset of surface defect images from the inspection systems of the consortium's industrial partners, RAS and SZFG, and makes it publicly available for research and benchmarking. The target size of at least 30,000 labelled images is calibrated to a scale at which deep learning classifiers trained on the dataset are expected to reach performance levels relevant to industrial deployment. The defect category structure follows the standardised surface defect catalogues for hot-rolled and cold-rolled flat steel, providing a class taxonomy recognised across the European flat steel industry and enabling cross-facility and cross-system comparison of classifier performance.

Surface inspection images from production facilities carry commercial sensitivity because the distribution of defect types and frequencies in a large image collection reflects the process capability of the facility that generated them. A competitor with access to a sufficiently large and labelled collection of surface images from a specific facility could characterise that facility's defect distribution and draw inferences about its process stability. The data governance framework for the SurfConInspect open dataset addresses this risk through image normalisation, metadata stripping, and selective exclusion of images that carry facility-identifying information beyond the surface feature itself. The framework draws on the GAIA-X initiative for European data infrastructure and the FAIR principles, Findable, Accessible, Interoperable, Reusable, adapted to the specific requirements of industrial surface inspection data, including the treatment of process-correlated metadata and the definition of access conditions that permit research use while preventing competitive exploitation.

The intellectual property status of inspection images is a further practical constraint. Images captured by an ASIS installed at a production facility may be owned by the facility operator or the system vendor depending on the terms of the supply contract. The IP assessment for each contributing partner is conducted before data collection begins, with the consortium's IP management structure advising on images whose ownership is contested. The labelling of the assembled dataset is performed by quality engineers at the contributing facilities, with inter-annotator agreement measured on a cross-labelling subset to validate label reliability. Images with low inter-annotator agreement are relabelled by consensus or excluded, ensuring that the published benchmark reflects reliable labels rather than the genuinely ambiguous cases that are present in any large production dataset.

Data quality and the design of the hybrid classification architecture

The choice of the hybrid classification architecture described in WP7, combining rule-based logic for well-characterised defect categories with image-based analysis for borderline cases, was made partly in awareness of the data constraints that characterise the SurfConInspect industrial context. A fully image-based classifier covering the complete defect taxonomy of the RAS and SZFG inspection systems would require training data of a volume and balance that the real inspection records of both facilities cannot currently provide. By restricting the image-based pathway to the subset of events that the rule-based system cannot resolve with sufficient confidence, the hybrid architecture reduces the volume and class balance requirements on the training dataset to a level that the combination of real, synthetic, and domain-adapted data can realistically meet.

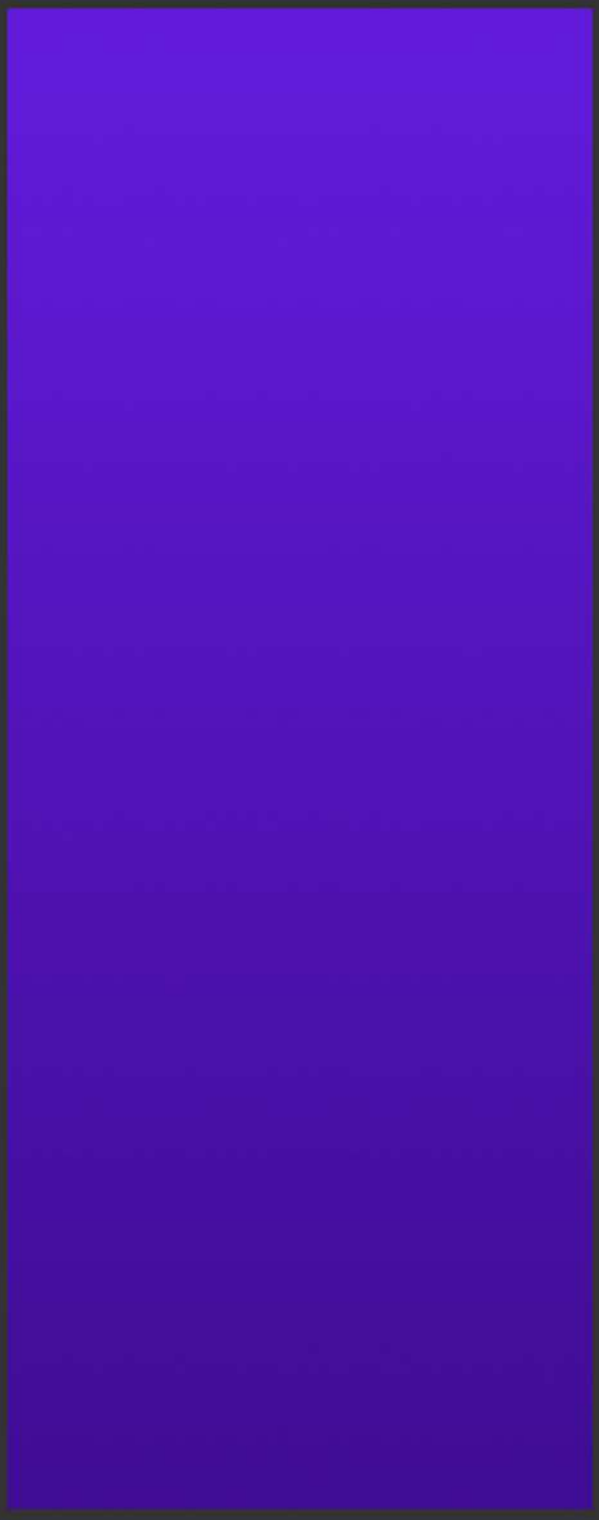


This relationship between data strategy and classification architecture is bidirectional. The data strategy is calibrated to the requirements of the classification system: synthetic augmentation targets the defect categories and surface conditions for which the hybrid system's image-based pathway is invoked most frequently, the ambiguous cases where rule-based confidence is low, because these are the cases where additional training data has the greatest impact on classification accuracy. Domain adaptation targets the transfer between RAS and SZFG for the shared defect categories where the image-based pathway is expected to provide the largest improvement over rule-based performance, because cross-facility deployment is where the domain gap most severely affects classification quality. The open dataset targets 30,000 images because that is the scale at which classifiers trained for the hybrid system's ambiguous-case function are expected to achieve industrially relevant performance levels, based on the scaling relationships observed in comparable image classification tasks.

A corollary of this architecture-data interaction is that improvements in data quality and improvements in classifier architecture are not independent variables in the SurfConInspect performance improvement framework. Adding synthetic wet-condition dent images to the training set improves the hybrid system's false negative rate for dents in wet pickling environments, but only if the image-based pathway of the hybrid system has sufficient capacity to learn from those additional examples. Adding domain-adapted representations of SZFG defect images improves classification at SZFG, but only if the rule-based component of the hybrid system has been configured with the defect attributes characteristic of SZFG surface conditions. The data strategies and the classification architecture must be developed in concert, each informed by the current state of the other, rather than treated as separable modules that can be optimised independently.

The data strategies described in this paper represent the project's systematic response to that constraint in the specific context of pickling-line and galvanising-line inspection, where the combination of wet surfaces, critical defect rarity, and multi-facility deployment creates a data environment that is more demanding than the controlled settings in which most surface inspection classification research has been conducted.

Connecting data quality to the quality control architecture



The data strategies described in this paper are not standalone research outputs; they are upstream components of the quality control architecture whose operational layers are described in WP6, WP7, and WP8. Classification reliability, the rate at which critical defects are correctly identified rather than misclassified as non-critical, determines the accuracy of the quality indices that the quality processor computes from inspection results. The cutting control processor at the RAS pickling line, described in WP6, computes a severity index for each detected event based on its classification and geometric attributes. If the classification is wrong, if a heavy dent near the strip edge is classified as a water reflection, the event does not receive a severity score, does not contribute to the threshold evaluation, and does not trigger a cutting signal. The defect proceeds to the downstream rolling mill with its full coil-break risk intact, as though the inspection system had not detected it.

The 20 percent reduction in coil break frequency that the SurfConInspect cutting control use case targets at downstream tandem rolling mills is therefore directly dependent on the false negative rate of the classification system for the dent and shell defect categories in the near-edge zone of the strip. Industrial data, as summarised in WP4, shows that approximately 90 percent of coil breaks at downstream facilities trace to defects in this near-edge zone. A false negative rate reduction of even a few percentage points for these specific categories, achievable through the combination of wet-condition synthetic augmentation and domain-adapted classifiers, translates directly into measurable reduction in the downstream failure rate and in the associated production losses.

The connection runs in the other direction as well. The traceability record generated by the quality processor and online quality supervisor, which logs every classification result alongside the subsequent downstream quality outcome for the corresponding coil, provides the longitudinal validation data needed to measure whether the data strategy improvements actually deliver the expected false negative rate reductions in production. Without this feedback from the operational quality control architecture back to the data and classification components, the performance claims of the data strategies would rest on validation set accuracy figures that may not translate to production conditions. The integration of data quality work within the broader SurfConInspect framework, rather than as a standalone research module, is what makes this operational validation possible.

The SurfConInspect perspective

The data challenges addressed in this paper, class imbalance from critical defect rarity, domain gaps between facilities, the absence of production-representative public benchmarks, are not problems that SurfConInspect discovered. They are widely recognised in the ASIS research and development community, acknowledged in the literature reviewed in WP7, and reflected in the gap between the classification performance reported in academic publications and the classification performance observed in productive inspection systems at steel facilities. What the project contributes is a coordinated engineering response to these challenges in the specific physical and operational context of pickling-line and galvanising-line inspection in European flat steel production.

The physical specificity matters. Synthetic generation of wet-condition defect images requires models of how water films modify the reflectance of steel surfaces, models that draw on the same physical understanding developed in WP2 for hyperspectral inspection in wet environments. Domain adaptation between RAS and SZFG requires knowledge of the physical sources of the domain gap, the difference in process liquid composition, surface texture, illumination geometry, to evaluate whether the gap is amenable to statistical adaptation or requires targeted data collection.

The relationship between WP9 and the other papers in this series reflects this physical grounding. WP2 establishes why wet surfaces create classification ambiguity; WP9 addresses how to train classifiers that are robust to it. WP4 identifies which defect categories carry the highest downstream risk; WP9 explains why those same categories are the most underrepresented in training data and what can be done about it. WP7 specifies the classification architecture that translates data quality improvements into operational performance gains; WP9 specifies the data strategies that provide the training material that architecture requires. These are not independent research contributions that happen to coexist in the same project; they are components of an integrated quality control framework in which each layer depends on the performance of the layers beneath it.

The open dataset, available at "https://ckan-surfconinspect.cetic.be", will outlast the formal duration of the SurfConInspect project. By providing the research community with a production-representative benchmark at the scale and class coverage needed for industrially relevant classifier development, it creates a reference point against which future algorithmic advances in ASIS classification can be evaluated on data that reflects the actual conditions of flat steel production rather than the controlled conditions of laboratory experiments. Establishing that reference point is a contribution to the field that complements, and in the long term may exceed in impact, the specific classifier performance improvements achieved within the project's own industrial demonstrations.

Conclusions

Classification performance in automatic surface inspection is bounded by the representativeness of the training data relative to the production conditions in which the classifier operates. The gap between benchmark accuracy and operational performance, a gap that practitioners in industrial ASIS consistently observe, is a direct consequence of the divergence between training and production distributions. Algorithmic improvements reduce the gap only insofar as they improve a classifier's tolerance to distributional shift; they cannot compensate for training data that fundamentally misrepresents the conditions under which the classifier will be used.



Three sources of distributional mismatch affect classification in the SurfConInspect industrial context. Critical defect categories are rare in well-controlled production, producing class imbalance that biases classifier decision boundaries away from the regions where critical defects occur under variable surface conditions. The domain gap between the RAS and SZFG inspection installations, arising from differences in material grade, process liquid composition, and sensor configuration, degrades classifier performance when systems are deployed outside the conditions under which they were trained. And the absence of production-representative public benchmarks prevents the research community from measuring whether claimed algorithmic advances deliver genuine improvements under industrial conditions.

The SurfConInspect data strategies address each source with a method appropriate to its physical character. Synthetic data generation through VAE and GAN architectures trained on real production images extends the training distribution toward the wet-surface, critical-defect regime that is most underrepresented in routine inspection records. Domain adaptation through Sliced Wasserstein Discrepancy minimisation transfers classification knowledge from labelled source domains to unlabelled target domains, reducing the relabelling cost of cross-facility deployment for the shared defect categories in the RAS-to-SZFG transfer. The open data initiative provides a labelled benchmark of at least 30,000 production-representative images, governed by FAIR-aligned data principles, that establishes a common evaluation standard for surface defect classification research.

These strategies are designed to feed the hybrid classification architecture of WP7, whose performance in turn determines the accuracy of the quality indices that drive the cutting control, speed optimisation, and online grading use cases described in WP6. The improvement pathway from data quality to classification reliability to quality index accuracy to downstream industrial outcomes, reduced coil break frequency, fewer coating defects, lower customer claim rates, is the operational justification for the data quality work as a component of the SurfConInspect quality control framework, not merely as a standalone research contribution.

SURF CON INSPECT

The Consortium

SurfConInspect unites a diverse group of partners:

- **Steel producers:** thyssenkrupp Rasselstein (RAS), Salzgitter Flachstahl (SZFG)
- **Research & Academia:** VDEh-Betriebsforschungsinstitut (BFI), Scuola Superiore Sant'Anna (SSSA), CETIC
- **Technology developers:** Clecim France (CLECIM)
- **Communication & dissemination experts:** yourscienceEDU (YSE)



STAY CONNECTED

Follow our updates, publications and events on:
🌐 Project Website: **surfconinspect.eu**
💼 LinkedIn: **SurfConInspect**

Together, we are shaping the future of **sustainable steel production**.



thyssenkrupp · SALZGITTER FLACHSTAHL · BFI Excellence in Applied Research · Sant'Anna Scuola Universitaria Superiore Pisa · cetic Your Connection to ICT Research · CLECIM · YSedu yourscienceEDU