WHITE PAPER

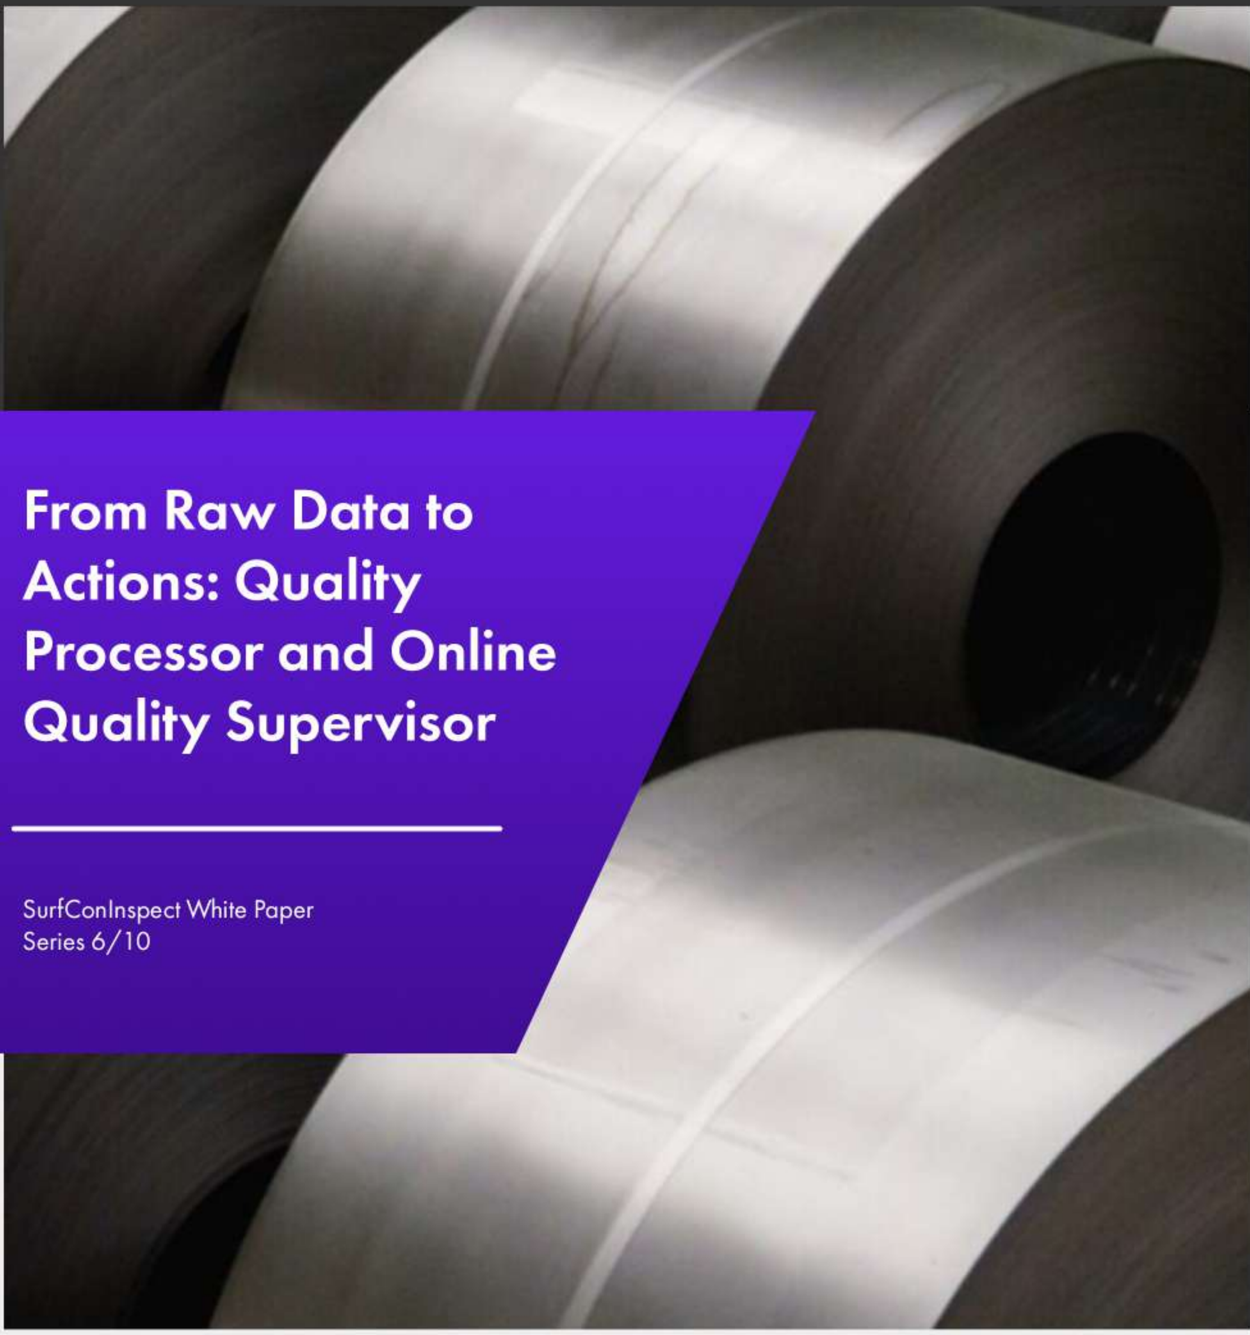

From Raw Data to Actions: Quality Processor and Online Quality Supervisor

SurfConInspect White Paper
Series 6/10

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

Inspection systems in flat steel production have evolved considerably over the past two decades. Line-scan cameras, high-resolution sensors, and—more recently—three-dimensional topographic measurement and hyperspectral imaging have extended the range of surface phenomena that can be detected and characterised at production speed. The real bottleneck in modern surface quality control is no longer detection alone. Sensors have become faster, richer, and more specialised, but production value still depends on what happens after detection. A line-scan camera, a 3D topographic channel, or a spectral detector can all produce large volumes of high-quality information at industrial speed. None of that information becomes operationally meaningful, however, until it is organised around a decision logic that reflects the physical behaviour of the strip and the practical constraints of the line. That is the problem SurfConInspect addresses directly. The project does not stop at improving what the plant can see. It builds the missing layer that turns measured quality information into structured events, actionable priorities, and responses that can be executed in time.

This paper sits at the architectural heart of the SurfConInspect concept. Earlier white papers explain how defects are measured and why geometry matters. WP6 explains what happens next: how raw inspection outputs are clustered, interpreted, weighted, and elevated into the kind of quality information a line can actually use. In the project architecture, this bridge is realised through two coordinated components. The quality processor is responsible for turning incoming inspection and contextual data into quality indices linked to a specific quality aspect. The online quality supervisor receives these outputs from multiple processors, orchestrates them in real time, and converts them into operator alerts or direct process interventions. This is the point at which inspection becomes control.

The paper explains these concepts without assuming prior knowledge of control systems engineering or data architecture. It is written for the same broad audience addressed throughout the SurfConInspect white paper series—plant managers, quality engineers, line operators, and production supervisors—who need to understand not only what advanced inspection systems can measure, but how the information they generate can be structured and used to make the production process more stable, more efficient, and more reliably aligned with quality requirements. The descriptions are grounded in the four industrial use cases at the heart of the project: automated cutting control for heavy defects at the RAS pickling line, process speed optimisation based on scale affection, online quality grading at the RAS tinning line, and online quality grading at the SZFG electrolytic galvanising line.

The gap between measurement and action

Every surface inspection system installed on a flat steel production line produces, at some level of abstraction, a record of what is on the strip surface. In the simplest case, this record is a set of images captured by line-scan cameras, from which detection algorithms extract regions that differ from the normal background texture and classify them as events of one type or another. In more advanced systems of the kind described in the preceding papers of this series, the record includes height maps from three-dimensional sensors, spectral response data from hyperspectral cameras, and a variety of positional and contextual metadata that locates each detected event precisely on the coil. As inspection systems have become more sophisticated, this record has become richer, more detailed, and more multi-dimensional. But richness of data alone does not guarantee quality improvement. Between the moment at which a defect is detected and classified and the moment at which an appropriate response is initiated—whether by an automated system or by a human operator—there is a gap that is filled neither by sensing technology nor by classification algorithms, but by the logic of decision-making.

In conventional line practice, the operator often acts as the hidden integration layer of the quality system. The inspection system reports anomalies, but the mental work of interpreting them, comparing them with process context, and deciding whether they justify action has largely remained human. On slower or less information-dense lines, that can be workable. On modern flat steel lines, it becomes structurally fragile. The SurfConInspect documents explicitly recognise that multiple quality aspects may need to be supervised in parallel and that intervention windows can be very short, especially in the pickling use cases. Under those conditions, leaving aggregation and prioritisation entirely to the operator creates not only occasional mistakes but recurring uncertainty in the control process. The architecture described here should therefore be understood as a way of redistributing cognitive load: away from continuous raw-data interpretation, toward curated supervisory judgement.

The consequence of this overload is not simply occasional error. It is systematic uncertainty in the quality control process—uncertainty about whether a detected defect has been correctly assessed, uncertainty about whether the appropriate response has been initiated in time, and uncertainty about whether the cumulative effect of many individually minor events on a single coil has been properly considered. This uncertainty has direct operational costs: unnecessary cutting of material that could safely have been processed, failure to cut material that subsequently causes a coil break, and missed opportunities to adjust process parameters in response to quality trends that were present in the inspection data but were not extracted and communicated in a form that the operator could act on.

The quality processor and online quality supervisor architecture of SurfConInspect is designed to address this gap systematically. Rather than generating a raw stream of defect events and leaving the interpretation entirely to the operator, the architecture introduces a structured layer of computational intelligence—not in the sense of machine learning or artificial inference, but in the engineering sense of a defined logic that converts raw data into quality indices and translates quality indices into control actions according to rules established in advance through the combined knowledge of process experts, quality engineers, and inspection system engineers. The result is a system in which the operator is supported rather than replaced: the machine handles the continuous, high-speed aggregation and interpretation of inspection data, while the human operator retains authority over decisions that require contextual knowledge or judgement that has not been captured in the rule set.

The concept of the quality processor

The quality processor is the smallest meaningful decision unit in the SurfConInspect architecture, and its strength lies in its specificity. Rather than attempting to create one monolithic quality engine for all surface-related issues, the project adopts a modular concept in which each processor supervises one quality aspect independently. This principle is stated clearly in the project overview: surface quality includes many different phenomena with different physical origins, different measurement needs, and different intervention consequences. Each quality processor therefore synchronises the data relevant to its own task, applies a dedicated quality model, and produces quality indices that describe the current state of the coil with respect to that task alone. This modularity is not just sound software engineering. It is what keeps the outputs interpretable for industrial users.

It is best understood by analogy with the control loops that are already standard in other aspects of flat steel production. A thickness control system at a cold rolling mill, for example, continuously measures strip thickness through X-ray or isotope gauges, compares the measured value against the target specification, computes a control error, and adjusts the roll gap to drive the error toward zero. The entire cycle—from measurement to comparison to correction—is performed automatically, continuously, and at a speed that matches the dynamics of the process being controlled. The quality processor applies the same principle to surface quality, but for a quantity that is inherently more complex than thickness: instead of a single scalar measurement, it works with a stream of discrete, spatially distributed defect events that must be aggregated into a meaningful quality representation before comparison against any kind of threshold or target.

What makes a quality processor useful is not merely that it receives defect data, but that it transforms those data through a quality model linked to a process decision. In the SurfConInspect requirements, this transformation is described consistently across the use cases: clustered raw data from inspection and material context are turned into quality indices that describe the current state of the coil. For heavy-defect handling at the pickling line, this means weighting defect depth, size, class, and edge position in relation to downstream coil-break risk. For scale affection, it means aggregating area-based information over coil sections to support speed adaptation. For line grading applications, it means condensing distributed inspection evidence into operator-relevant assessments of current coil quality. The common structure is the same across all cases: raw events are too granular to support action directly, so the processor creates a more decision-ready quality representation.

The cutting-control case is where the quality processor becomes most visibly tied to downstream physics. The project documentation already makes clear that the relevant variables are not simply defect presence or absence, but height or depth, lateral extent, defect type, steel type, and position relative to the strip edge. In practical terms, this means that the severity score is not a generic anomaly score. It is a process-informed risk estimate. A defect that is geometrically similar to another may be assigned a very different priority depending on its edge position or downstream sensitivity. This is consistent with the broader project evidence from the topographic defect work, where edge proximity emerges as a major severity amplifier. The value of the quality processor, then, is that it formalises this logic so the system can apply it continuously and consistently instead of leaving it to ad hoc interpretation.

Modularity is one of the strongest design decisions in the entire SurfConInspect architecture. The project overview explicitly states that multiple quality processors can run independently in parallel, each representing a single quality aspect of the production process. That approach matters for at least three reasons. It allows each processor to remain tied to a well-defined physical question. It makes commissioning more manageable, because thresholds and logic can be tuned by use case rather than across a monolithic system. And it preserves traceability, since each quality event can be traced back to a specific processor, model, and threshold logic. In industrial deployment, those properties are not secondary. They are what make the system maintainable and credible over time.

From quality index to quality event

A quality index, as produced by the quality model within a quality processor, is a continuously updated numerical representation of the current state of the coil with respect to a defined quality aspect. As long as the index remains within acceptable bounds—as long as the detected defect severity remains below the critical threshold, or the scale affection fraction remains within the permissible range for the current production speed—the quality processor operates in a monitoring mode, continuously updating its assessment without generating any output that requires a response. This silent monitoring is as important as the detection of critical events, because it provides the baseline record of quality behaviour from which deviations can be identified and from which process trends can be analysed retrospectively.

The transition from quality index to quality event is the moment at which monitoring becomes supervision. As long as the index remains within acceptable bounds, the processor is effectively building a structured baseline of normal or tolerable quality behaviour. Once a threshold is crossed, the system needs a more formal object of communication. That is the function of the quality event. In the SurfConInspect logic, the event is not just an alarm flag. It is the structured package that carries the identity of the affected quality aspect, the threshold crossing itself, the relevant supporting measurements, the coil position, and the timing needed to correlate the signal with plant operations and later traceability. This is why the event becomes the natural interface between processor and supervisor: it compresses meaning without losing context.

Threshold setting is where quality control ceases to be a purely technical exercise and becomes a matter of industrial judgement. The project material strongly supports the need for a commissioning phase in which indices are first observed in monitoring mode and then translated into intervention logic. Set thresholds too low and the system will flood the operator and supervisor with low-value signals, creating the same habituation problem already noted for audible alerts in the human-machine interface. Set them too high and the system becomes elegantly quiet while letting process-relevant deterioration pass downstream. SurfConInspect therefore treats thresholds as process-specific parameters rather than universal constants. They depend on production context, material, downstream route, and the costs of false positive versus false negative action. That is a mature and industrially credible position.

One of the architecture's most valuable features is that it allows quality events to be interpreted not only individually but relationally. A heavy topographic event near the strip edge and an elevated scale-affection condition are not the same kind of problem, but their simultaneous occurrence may be more informative than either signal taken alone. That is exactly the type of multi-aspect interpretation the online quality supervisor is meant to enable. The project documentation describes the supervisor as the component that receives all incoming quality events, combines their implications, and transforms them into coherent control logic. This is where a modular architecture becomes more than a collection of independent channels. It becomes a supervisory view of coil condition.

The online quality supervisor

The online quality supervisor is best described as the plant-facing intelligence layer of the architecture. Each processor remains faithful to one quality aspect. The supervisor is where those aspect-specific outputs are orchestrated into a single operational response. This orchestration may include filtering, prioritising, merging several events into one recommendation, or turning a single event into an immediate intervention. The supervisor does not replace the processors. It gives them operational coherence.

The supervisory logic that the online quality supervisor applies to incoming quality events is structured around two fundamental categories of control action. Automatic supervisory control is the right strategy when three conditions are met simultaneously: the quality event is clearly above threshold, the response is well defined, and the intervention window is too narrow for reliable human review. SurfConInspect explicitly positions both pickling-line use cases in this direction. The project overview states that automated cutting and speed control are well suited to automatic supervisory control, even though the underlying quality models are not trivial. This nuance is important. Automation is not being justified because the decision is simplistic. It is being justified because the decision, once properly modelled, is time-critical and mechanically unambiguous enough to benefit from direct machine-to-machine response. In the SurfConInspect framework, when a defect event exceeds the severity threshold associated with coil-break risk, the supervisor sends a cutting signal to the cutting station controller, which initiates the mechanical cutting procedure without waiting for operator input.

Human supervisory control is not a fallback for architectural incompleteness. It is the correct design choice whenever process context, customer implications, or borderline quality trade-offs still matter more than execution speed. The project overview states this clearly for the online grading use cases at the tinning and galvanising lines: the response to emerging quality events is not always straightforward and often depends on information not fully available to the control system, including customer-related context. In those cases, the supervisor's job is not to remove the operator from the loop but to sharpen the loop: filter the noise, structure the evidence, and present the event in a form that supports a defensible decision. That is a very different role from raw alarm display.

The strongest framing of supervisory control is as a graded response spectrum rather than a fixed system split. The project materials already support this with examples of filtering, prioritising, and differentiated action ranges. A processor may compute a quality index continuously, but the supervisor decides how that index enters operations: silent monitoring, low-priority flagging, operator alert, or direct actuation. This graduated logic is essential in steelmaking because the economics of intervention are rarely linear. A marginal event should not necessarily trigger the same response as a clearly critical one, and a system that cannot express that nuance quickly becomes either overactive or distrusted. The supervisor architecture is valuable precisely because it allows that nuance to be formalised.

Data synchronisation and multi-source integration

The quality processor and online quality supervisor architecture operates on data from inspection systems, but in many of the SurfConInspect use cases, inspection data alone is not sufficient to generate a complete and reliable quality assessment. The quality of a coil is not determined solely by the surface defects detected at the pickling line inspection zone. It also depends on upstream process history—the casting conditions, the hot rolling parameters, the coil handling and storage record—and on the characteristics of the specific material being processed, including its grade, thickness, and intended downstream application. Integrating this contextual information with inspection data is a prerequisite for quality decisions that are well-calibrated to the actual requirements of each coil.

Data synchronisation is the quiet enabler of the entire SurfConInspect concept. Without reliable position alignment and source integration, even the best quality models would be acting on partial or misplaced evidence. The project documentation emphasises this in two ways. First, the SCI framework is designed to synchronise multiple measurements and in-coil positions in real time. Second, the exploitation and overview documents describe methods for synchronising data from multiple sensors and calculating indices related to specific quality aspects. This matters especially at the pickling line, where time, strip position, and intervention location are inseparable. If the event is real but the coil segment is misaligned, the control action will still fail industrially.

In a continuous pickling line, the strip moves through multiple process zones between the entry section—where the coil is unwound and welded to the preceding coil to form a continuous strip—and the exit section, where the processed strip is recoiled into finished coils. The same physical point on the strip passes through the inspection zone at one time and through the exit recoiler at a later time, and the delay between these two moments depends on the current line speed and the accumulated length of strip in the loop sections that buffer the line against speed variations. Establishing a reliable correspondence between inspection data collected at one point in time and coil position data collected at another requires a position tracking system that follows the movement of every point on the strip from entry to exit, accounting for the elongation and speed variations that occur along the way.

In the SurfConInspect architecture, this position tracking is implemented through a combination of encoder signals from the drive systems of the various line sections and a computational model that propagates position information forward through the line geometry as the strip advances. The result is a continuous position estimate for each point on the strip, expressed as a distance from the coil head, that remains synchronised with the inspection data across the entire coil length. This synchronisation is the foundation on which the quality processor's quality model operates: without it, the aggregation of defect events into a quality index cannot be reliably associated with specific coil segments, and the cutting or grading decisions derived from the quality index cannot be correctly positioned on the material.

The integration of upstream process data—from the hot rolling mill, the continuous caster, or the material tracking system—into the quality assessment adds a further dimension to this synchronisation challenge. Process data generated at an earlier production stage is timestamped and associated with slab or coil identifiers that must be mapped onto the strip position coordinates used by the inspection system. This mapping requires a data infrastructure capable of handling the asynchronous arrival of information from multiple sources and of resolving the identity correspondences between coils at different stages of processing. In the SurfConInspect project, this infrastructure is provided by the distributed SCI framework, which manages the communication between the quality processors, the online quality supervisor, the inspection system data streams, and the plant-level process data systems through a modular, message-based architecture that can be adapted to the specific IT environment of each industrial partner.

Use case illustration: automated cutting at the pickling line

The automated cutting use case is the clearest demonstration of the architecture's purpose because it connects geometry, risk modelling, and physical action in one short interval. The rule-based examples in the requirements are especially useful here because they make the logic concrete without exposing confidential thresholds: a defect may be cut directly, handled later downstream, or ignored depending on height, position from edge, type, and steel type. This gives the white paper a valuable opportunity to show that the quality processor is not an abstract analytics layer. It is a rule-bearing process actor.

The severity index for a given defect event reflects the combination of its geometric properties—depth, profile sharpness, lateral extent—and its positional properties, specifically its distance from the nearest strip edge. As described in WP4, defects located within the near-edge zone of the strip are assigned higher severity scores than geometrically equivalent defects in the central region, because the elevated tensile stress near the strip edges during cold rolling amplifies the stress concentration effect of the defect and increases the probability of crack initiation. A dent with a depth of 0.5 mm at the strip centre might receive a severity index below the cutting threshold; the same dent located 50 mm from the strip edge might receive an index above the threshold. The model encodes this positional dependency quantitatively, so that the severity assessment reflects the actual downstream risk rather than the geometric properties of the defect in isolation.

When the severity index for a defect event exceeds the cutting threshold, the quality processor generates a quality event and transmits it to the online quality supervisor. The supervisor receives the event, confirms that it satisfies the conditions for automatic supervisory control—severity above the cutting threshold, defect in a position where cutting is mechanically feasible, no conflicting operational constraints from other active quality processors—and generates a cutting signal directed at the cutting station controller. The cutting station controller initiates the cutting procedure, positioning the cutting mechanism and executing the cut while the strip is in motion, removing the section of strip containing the defect from the coil and directing it to the scrap or evaluation path. The entire sequence, from defect detection to cut execution, must be completed within the time window defined by the distance between the inspection zone and the cutting station and the current line speed.

The industrial importance of automated cutting lies in what it removes from the operator's workload and what it prevents downstream. Before structured decision support, the operator had to monitor continuous inspection results, distinguish critical from non-critical events under time pressure, and act before the physical intervention window closed. SurfConInspect reduces that burden for the class of events that are both severe and modelled well enough to justify immediate automatic action. This is consistent with the broader project ambition of achieving a new level of online surface quality control: not by bypassing the operator everywhere, but by reserving human attention for the situations where human context still adds value.

Use case illustration: process speed control via scale affection

The speed-control use case is important because it shows that the quality processor concept is not limited to singular critical defects. Here the monitored quantity is distributed, area-based, and cumulative. Residual scale is not evaluated as an isolated event but as an evolving surface condition over coil length, and the relevant output is not a cut decision but a dynamic recommendation on process speed. This makes the use case conceptually valuable: it demonstrates that the same architecture can support both defect-centric control and condition-centric control. In both cases the processor performs the same basic industrial function—converting inspection data into a quality index aligned with a process lever—but the aggregation logic and the action type are completely different.

The degree of scale affection of a coil is operationally important because it determines how effectively the pickling acid has penetrated and dissolved the oxide layer across the full strip width and along the full coil length. Coils with high scale affection—where a significant fraction of the surface remains covered by residual oxide after passing through the pickling bath—require more aggressive treatment conditions or, more practically, a lower processing speed that increases the residence time of the strip in the acid bath and allows more complete scale removal. Coils with low scale affection can be processed at higher speeds without compromising the chemical treatment quality. The optimisation of line speed based on scale affection is therefore a direct path to productivity improvement: running the line faster on coils that can support it, while slowing down only when the scale characteristics of the material genuinely require it.

The quality processor for scale affection computes a scale affection index by aggregating the spatial coverage of scale-related detection events over sliding windows of coil length. As the strip advances through the inspection zone, each detected scale event contributes to the running average of scale coverage, which is updated continuously and smoothed to remove short-range fluctuations that reflect local variability in the pickling process rather than systematic under-treatment. When the scale affection index exceeds the threshold associated with the current production speed—indicating that the strip is accumulating more residual scale than is acceptable at that speed—the quality processor generates a quality event recommending a speed reduction. When the index falls below the threshold associated with a higher speed level, the processor generates an event recommending a speed increase. The online quality supervisor receives these events and translates them into speed adjustment signals directed at the line drive system, subject to the operational constraints—maximum ramp rate, minimum speed floor—that prevent abrupt changes incompatible with the mechanical stability of the line.

A distinctive feature of the scale affection use case is the role of the two-dimensional imaging channel in the quality index computation. Scale is a non-topographic phenomenon: it has no consistent geometric signature that would make it reliably detectable in the three-dimensional height map. The three-dimensional channel, which is highly effective for the detection and characterisation of topographic defects, is largely insensitive to scale coverage because scale-covered and clean-steel regions may be geometrically indistinguishable. The scale affection index is therefore computed from the two-dimensional imaging channel alone, while the three-dimensional channel continues to monitor for topographic defects in parallel. This complementarity reflects the integrated multi-channel design philosophy of the SurfConInspect inspection system and the recognition that different quality aspects require different sensing modalities.

The role of the operator within the supervisor architecture

SurfConInspect does not make the operator disappear from quality control. It changes where the operator adds value. Continuous aggregation, threshold tracking, and multi-stream vigilance are shifted into the processor-supervisor architecture because they are repetitive, time-critical, and structurally difficult to sustain by human attention alone. The operator remains essential where judgement, context, and accountability matter most. That is especially visible in the human supervisory control use cases, where the system presents a structured assessment of what is critical, marginal, or stable, but the final grading decision remains human. This is a more realistic and industrially credible vision of advanced automation than claims of full autonomy.

This repositioning of the operator's role is most clearly illustrated in the human supervisory control use cases—the quality grading applications at the RAS tinning line and the SZFG electrolytic galvanising line. In these use cases, the online quality supervisor does not generate automatic control actions. Instead, it presents the operator with a quality assessment of the current coil—a summary of which quality aspects are within specification, which are marginal, and which have generated events that require a response—and provides, for each quality event, the supporting data the operator needs to make a decision: the severity of the event, its position on the coil, its relationship to other events on the same coil, and, where relevant, its implications for specific downstream processes or customer specifications. The operator retains full authority over the final decision but makes that decision with the benefit of a systematically organised, quantitatively grounded quality assessment that would not have been accessible without the supervisory architecture.

The design of the operator interface for the human supervisory control use cases reflects the lessons of human factors engineering applied to quality control environments. Information must be presented at the right level of aggregation: too much detail overwhelms, too little abstraction conceals the pattern. Priority must be communicated visually and immediately, so that the operator can identify the most urgent quality event without searching. Contextual information must be accessible but not intrusive. The design challenges of achieving this balance are addressed within the broader context of operator decision support described in WP3 of this series, and the SurfConInspect AR demonstrator represents one approach to making quality information more directly accessible within the operator's working environment.

The training implications of this new operator role are also significant. Operators who are accustomed to interpreting raw inspection images and making quality judgements based on visual pattern recognition must develop a complementary understanding of quality indices and quality events—what they represent, how they are computed, what assumptions underlie the severity thresholds, and when the automated assessment can be trusted and when it should be questioned. This understanding cannot be assumed; it must be actively developed through structured training, through transparency about the logic of the quality models, and through the accumulation of operational experience in which the predictions of the supervisory system are systematically compared against subsequent downstream outcomes. Building this experience base is itself one of the most valuable long-term outputs of the SurfConInspect project.

Traceability and data-driven process improvement

One of the less visible but more durable outputs of the architecture is traceability. Because processors compute indices continuously and the supervisor logs events and responses, the system naturally builds a historical record of surface quality behaviour across coils, positions, and time. That record is strategically important. It supports incident reconstruction, threshold refinement, model tuning, and the gradual accumulation of plant-specific knowledge about which quality signals truly matter downstream. This point becomes stronger when linked to the project's wider ambition around data-driven improvement, benchmarking, and open-data concepts. SurfConInspect is not only building a real-time control architecture. It is also laying down a structured quality memory for future optimisation.

The most direct application of this traceability record is the investigation of downstream quality incidents. When a coil break occurs at a cold rolling mill, or when a customer complaint is received about a surface defect in a delivered product, the first question is always where the defect originated and why it was not detected or acted upon at the pickling line. With a comprehensive traceability record, this question can be answered quantitatively: the relevant coil segment can be identified in the database, the inspection history of that segment can be retrieved, and the question of whether the defect was detected, how it was classified, and what response it triggered can be answered from the data rather than from reconstructed recollection. This capability transforms the investigation of quality incidents from a forensic exercise that depends on operator memory into a systematic data analysis that yields actionable conclusions.

The traceability record also supports a longer-term process improvement function. Over time, as the database accumulates quality histories for a large number of coils, it becomes possible to identify systematic patterns in the relationship between pickling-line quality events and downstream outcomes: which types of defect events at which positions on the strip are most strongly correlated with subsequent coil breaks; which combinations of scale affection index and processing speed are associated with surface quality problems at the tinning or galvanising line; which upstream process conditions—hot rolling parameters, slab casting records—are predictive of the quality events observed at the pickling stage. These patterns, once identified, provide the basis for targeted improvements in process parameters, inspection threshold settings, and intervention rules that progressively reduce the frequency and severity of quality incidents. The SurfConInspect architecture is designed with this continuous improvement function in mind, and the open data framework described in WP9 of this series represents one approach to making the accumulated quality data available for research and benchmarking purposes beyond the boundaries of individual production facilities.

Conclusions

The conclusion works best when it closes the loop between sensing, architecture, and industrial outcome. The lasting contribution of SurfConInspect is not simply that it detects more defects, or that it applies more advanced software to inspection data. It is that it introduces a systematic, modular, and operationally intelligible way of turning inspection outputs into actions. Quality processors formalise specific quality questions. The online quality supervisor turns their outputs into plant decisions. Operators receive better-filtered and better-timed information. Process systems receive direct intervention signals where automation is justified. And the plant gains a traceable quality history that can be used to improve thresholds and models over time. That is what gives the architecture strategic relevance beyond the project itself.

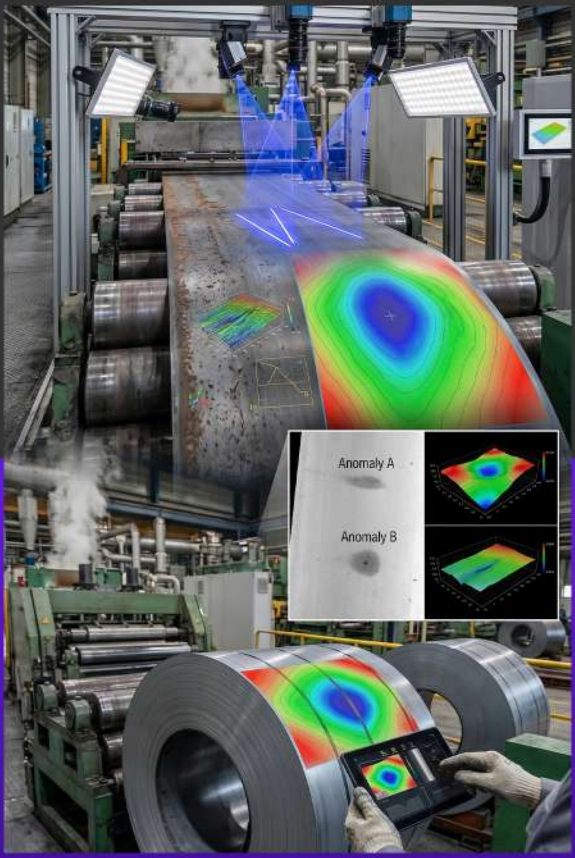

The quality processor concept—in which a dedicated computational module monitors a single quality aspect, applies a quality model to the stream of inspection events, computes a continuously updated quality index, and generates structured quality events when defined thresholds are exceeded—provides a modular, transparent, and maintainable framework for implementing multiple quality objectives simultaneously on the same production line. The online quality supervisor—which receives quality events from all active quality processors, evaluates their combined implications, and translates the result into automatic control actions or human supervisory alerts—provides the orchestrating logic that connects individual quality assessments into a coherent production control function.

The four SurfConInspect use cases—automated cutting control, process speed optimisation, tinning line quality grading, and galvanising line quality grading—illustrate how this architecture can be applied to qualitatively different quality management objectives, each requiring a different quality model, a different threshold structure, and a different balance between automatic and human supervisory control. Across all four, the common thread is the same: inspection data acquires operational value not through its volume or technical sophistication alone, but through the systematic logic that transforms it into decisions grounded in the physics of the material, calibrated to the requirements of the downstream process, and communicated to the operator and the process control system in a form that is both timely and actionable.

The traceability record generated by the architecture as a byproduct of its operational function adds a further dimension of value: a comprehensive, structured database of surface quality behaviour that supports the investigation of downstream incidents, the continuous improvement of process parameters and inspection thresholds, and the accumulation of the operational knowledge base on which future enhancements to the system can be built. In this sense, the quality processor and online quality supervisor are not only tools for managing today's production. They are the infrastructure through which the knowledge needed to improve tomorrow's production is systematically collected and preserved—and the foundation on which SurfConInspect's broader ambition of data-driven, benchmarked, open-quality management in flat steel production can be built.

FURTHER READING

- SurfConInspect Project Website
- Introductory resources on 3D surface inspection
- Related white papers in the SurfConInspect series

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 A Member of the Salzgitter Group BFI Excellence in Applied Research Sant'Anna Scuola Universitaria Superiore Pisa cetic Your Connection to ICT Research CLECIM FRANCE YSedu yourscienceEDU