ΠΑΝΕΠΙΣΤΗΜΙΟ ΠΑΤΡΩΝ
ΤΜΗΜΑ ΜΗΧΑΝΟΛΟΓΩΝ ΚΑΙ ΑΕΡΟΝΑΥΠΗΓΩΝ ΜΗΧΑΝΙΚΩΝ
ΚΑΤΑΣΚΕΥΑΣΤΙΚΟΣ ΤΟΜΕΑΣ
ΕΡΓΑΣΤΗΡΙΟ ΣΥΣΤΗΜΑΤΩΝ ΠΑΡΑΓΩΓΗΣ ΚΑΙ ΑΥΤΟΜΑΤΙΣΜΟΥ

ΔΙΠΛΩΜΑΤΙΚΗ ΕΡΓΑΣΙΑ

Σχεδιασμός και ανάπτυξη υβριδικού κελιού, βασισμένο σε συνεργασία ρομπότ και ανθρώπου, για την ανακατασκευή ελαστικών βαρέων οχημάτων.

Design and Development of a Hybrid Cell, based on Human Robot Collaboration, for the Remanufacturing of Heavy Transport Vehicle Tires

Ευάγγελος Ανδρωνάς

1080527

Επιβλέπων Καθηγητής:
Σωτήρης Μακρής
Αναπληρωτής Καθηγητής

ΠΑΤΡΑ, Σεπτέμβριος 2026

University of Patras, Mechanical Engineering and Aeronautics Department
Evangelos Andronas
© 2026 – All rights reserved

ΠΕΡΙΛΗΨΗ

Σχεδιασμός και ανάπτυξη υβριδικού κελιού, βασισμένο σε συνεργασία ρομπότ και ανθρώπου, για την ανακατασκευή ελαστικών βαρέων οχημάτων.

Ευάγγελος Ανδρωνάς

Η ανακατασκευή διατηρεί την αξία των προϊόντων στην κυκλική οικονομία, όμως η αυτοματοποίησή της περιορίζεται από τη μεταβλητή κατάσταση των επιστρεφόμενων προϊόντων, την έλλειψη πληροφοριών ανά τεμάχιο και την εξάρτηση από χειρωνακτική εργασία. Σε αντίθεση με τη συμβατική παραγωγή, η παραλλαγή και η κατάσταση κάθε τεμαχίου προσδιορίζονται κατά την εκτέλεση, ενώ οι απαιτούμενες κατεργασίες δεν εντάσσονται σε προκαθορισμένο πλάνο. Η εργασία παρουσιάζει γενική μεθοδολογία σχεδιασμού ρομποτικών κελιών ανακατασκευής, βασισμένη στην αξιολόγηση κατάστασης, την ιχνηλασιμότητα και τον καθορισμό κατεργασίας ανά τεμάχιο, με τη συνεργασία ανθρώπου-ρομπότ να διατρέχει τα στάδια. Εφαρμόζεται στην αναγόμωση ελαστικών βαρέων οχημάτων, με στόχο την αυτοματοποιημένη ανίχνευση και τον εντοπισμό ελαττωμάτων και την τοπική αφαίρεση υλικού στις εσωτερικές και εξωτερικές επιφάνειες του σκελετού, με ασφαλή παρέμβαση του χειριστή όπου απαιτείται. Το κελί χρησιμοποιεί ένα ρομπότ υψηλού ωφέλιμου φορτίου με συνεργατική αρπάγη πολλαπλών βαθμών ελευθερίας, η οποία συγκρατεί, μετρά, περιστρέφει και αναπροσανατολίζει τον σκελετό, τοποθετώντας τον στους σταθμούς αντίληψης και τοπικής αφαίρεσης χωρίς πρόσθετο σύστημα χειρισμού. Η εικονική επαλήθευση σε CATIA/DELMIA επιβεβαίωσε πλήρη επιφανειακή κάλυψη κατά τον έλεγχο, προσβασιμότητα του εργαλείου αφαίρεσης υλικού χωρίς συγκρούσεις και κάλυψη των ζωνών ασφαλείας. Η μεθοδολογία αποτελεί βάση σχεδιασμού ευέλικτων, ανθρωποκεντρικών κελιών ανακατασκευής για τεμάχια άγνωστης κατάστασης, εφαρμόσιμη και σε άλλες διεργασίες.

Λέξεις κλειδιά: Κυκλική Οικονομία; Ανακατασκευή; Σχεδιασμός Ρομποτικού Κελιού; Συνεργασία Ανθρώπου-Ρομπότ; Αυτοματοποιημένος Οπτικός Έλεγχος; Εικονική Επαλήθευση; Αναγόμωση Ελαστικών

ABSTRACT

Design and Development of a Hybrid Cell, based on Human Robot Collaboration, for the Remanufacturing of Heavy Transport Vehicle Tires

Evangelos Andronas

Remanufacturing retains product value within the circular economy, but its automation is hindered by the variable condition of returned products, limited item-specific information, and reliance on labor-intensive operations. Unlike conventional manufacturing, each incoming core must be assessed at runtime, and its processing cannot follow a fixed plan. This thesis presents a general methodology for designing robotic remanufacturing cells that combines condition assessment, item-level traceability, and per-item processing decisions, with human–robot collaboration throughout. The methodology is applied to heavy-duty transport vehicle tire retreading, targeting automated defect detection and localized material removal across the casing's inner and outer surfaces, alongside safe operator intervention where automation is insufficient. The developed cell uses a single high-payload robot equipped with a multi-DOF collaborative gripper that secures, measures, rotates, and reorients the casing, positioning it at perception and skiving stations without separate handling equipment. Virtual verification in CATIA/DELMIA confirmed full-surface inspection coverage, collision-free skiving-tool reachability, and feasibility of the designed safety zones. The methodology thus provides a transferable basis for designing flexible, human-centric cells capable of processing returned products of uncertain condition.

Keywords: Circular Economy; Remanufacturing; Robotic Cell Design; Human–Robot Collaboration; Automated Inspection; Virtual Verification; Tire Retreading

LIST OF TABLES

Table 1 Physical characteristics of the reference tire, as a new tire and as a buffed casing. 39

Table 2 Defect taxonomy for the reference scenario. 42

Table 3 Targets set for the developed system. 44

Table 4 Requirements derived from the reference scenario. 46

Table 5 Categories of data that a Re-X traceability record may hold. 55

Table 6 Operation categories available to a Re-X processing cell. 57

Table 7 Specifications for the Powered Roller Gripper. 74

Table 8 Per-operation duration estimates 105

Table 9 Estimated total processing time per tire, by defect count, under a fully automated scenario and a scenario with manual intervention 106

TABLE OF FIGURES

Fig. 1 The circular economy butterfly diagram. (Source: Ellen MacArthur Foundation [3]) 16

Fig. 2 The 9R framework of circularity strategies, ordered by decreasing value retention. 17

Fig. 3 Methodology overview diagram. 21

Fig. 4 The main five stages of the tire retreading process, with defect detection and localization and skiving marked as the two stages addressed within the scope of this thesis. 35

Fig. 5 Current industrial practice at the reference site: (a) the manual skiving workstation, (b1) an operator performing skiving on a buffed casing, (b2) the tools used for manual skiving 37

Fig. 6 3D-model buffed tire casing and its main dimensions of the widely used size 315/80R22.5. 39

Fig. 7 Representative examples of the defect on a buffed casing. 41

Fig. 8 The proposed workflow for Re-X system operation. 51

Fig. 9 The classes of condition variability a core may present, the sensing each requires, and the resulting condition profile 53

Fig. 10 The Section 4.2 Engineering activities, moving towards physical implementation. 59

Fig. 11 Tire anatomy and terminology used throughout this work: crown, sidewalls, bead, inside crown, and inside sidewalls. 69

Fig. 12 Cell layout and the main components of the designed robotic cell. 70

Fig. 13 Developed robotic cell with the material flow a) Tire transfer from input to inspection station, b) Transfer to skiving station, c1) Transfer to interaction area for manual processing, c2) Transfer back to skiving station, d) Transfer to output rack after completing process 72

Fig. 14 Securing mechanism and the main components. 76

Fig. 15 Rotation mechanism and the main components. 77

Fig. 16 Powered Roller Gripper mechanisms and the gripper mounted on the NJ130-2.6 robotic arm. 77

Fig. 17 Inspection station, ZED X mini arrangement and LED illumination. 79

Fig. 18 The two-step detection sequence: Outer-surface configuration and inner-surface configuration. 80

Fig. 19 Cylindrical coordinate representation used to localize a detected defect relative to the tire's center 81

Fig. 20 Skiving tool main components and coordinate frames of the base and the TCP. 82

Fig. 21 Skiving tool reaching outer and inner surfaces a) Crown area, b) Sidewall, c) Inner surfaces 84

Fig. 22 PGR safety enclosures and AIRSKINs covering the drivetrain components. 87

Fig. 23 Radar-monitored zone configuration. 88

Fig. 24 Operator indexing the tire rotation using the two-hand control. 90

Fig. 25 The Product–Process–Resource structure of the simulated system, the tire as product, the grouped robot and gripper activities as processes, and the robot, gripper, and stations as resources. 92

Fig. 26 Activity diagrams for the a) Loading of the tire and b) the inspection. 93

Fig. 27 DELMIA rendering of two adjacent cell layouts sharing a central operator zone, with independently monitored safety regions. 94

Fig. 28 Layout issues at the detection station, and the skiving station 97

Fig. 29 Powered Roller Gripper kinematic range. 99

Fig. 30 Angle coverage of the selected cameras for both the outer and inner surface detection. 101

Fig. 31 Skiving tool reach across inner and outer casing profile. 102

Fig. 32 Activity network (PERT) of the process sequence, showing the grouped activities and their precedence..... 103

Fig. 33 Simulated process sequence in DELMIA: 1) Pick tire from input rack, 2) Outer-surface inspection at the perception station, 3) Inner-surface inspection after repositioning of the side cameras, 4) Automated skiving at the skiving station, 5) Manual skiving by the operator in the interaction area, 6) Return to the skiving station, 7) Unload to the output rack. 108

Fig. 34 Operator working posture in a) current manual practice loading the tire & RULA score versus b) the proposed robotic solution. 109

TABLE OF CONTENTS

ΠΕΡΙΛΗΨΗ 3

ABSTRACT 5

LIST OF TABLES 7

TABLE OF FIGURES 8

TABLE OF CONTENTS 11

1.0 INTRODUCTION 14

1.1 BACKGROUND AND INDUSTRIAL CONTEXT 14

1.2 THESIS AIM AND SCOPE 20

1.3 THESIS STRUCTURE 23

2.0 STATE OF THE ART 24

2.1 AUTOMATION CHALLENGES IN REMANUFACTURING SYSTEMS 24

2.2 HUMAN-ROBOT COLLABORATIVE CELLS 25

2.3 ROBOTIC HANDLING AND PROCESSING OF LARGE OR VARIABLE-GEOMETRY COMPONENTS 28

2.4 TIRE REMANUFACTURING TECHNOLOGIES 29

3.0 INDUSTRIAL PROBLEM AND MOTIVATION 32

3.1 CHALLENGES IN RE-X PROCESSING 32

3.2 REFERENCE MANUFACTURING SCENARIO 35

3.3 PRODUCT SPECIFICATIONS AND DERIVED REQUIREMENTS 38

3.3.1 Product specifications and physical characteristics 38

3.3.2 Defect taxonomy 41

3.3.3 Performance specifications 43

3.3.4 Requirements derived for the cell design 45

4.0 APPROACH 50

4.1 REMANUFACTURING-DRIVEN DESIGN REQUIREMENTS FOR ROBOTIC CELLS 50

4.1.1 Condition assessment 52

4.1.2 Traceability layer and digital passport 54

4.1.3 Processing strategy 56

4.2 REMANUFACTURING WORKCELL ENGINEERING 58

4.2.1 Remanufacturing cell workflow definition 59

4.2.2 Process and resource definition, and planning 60

4.2.3 Workcell layout and workstation arrangement 62

4.2.4 Tooling and workstation design principles 64

4.2.5 Digital evaluation and virtual commissioning 65

4.2.6 Human-Robot collaboration strategy 65

5.0 IMPLEMENTATION 67

5.1 ROBOTIC CELL LAYOUTING 69

5.2 END-EFFECTOR 72

5.2.1 Design rationale: why an active end-effector 73

5.2.2 Gripper specifications 74

5.2.3 Main mechanisms 75

5.3 INSPECTION STATION 78

5.4 SKIVING STATION 81

5.5 HUMAN-SYSTEM INTERACTION, OPERATOR SUPPORT, AND SAFETY STRATEGY 85

5.6 VIRTUAL VERIFICATION 91

5.7 SCALABILITY OF THE PROPOSED WORKCELL 93

6.0 RESULTS 96

6.1 CELL LAYOUT AND REACHABILITY VALIDATION 96

6.2 SUBSYSTEM VALIDATION 98

6.2.1 Powered Roller Gripper 98

6.2.2 Inspection Station 100

6.2.3 Skiving Station 101

6.3 PROCESS SEQUENCE SIMULATION 102

7.0 CONCLUSIONS AND FUTURE WORK 110

7.1 CONCLUSIONS 110

7.2 FUTURE WORK 112

REFERENCES 114

1.0 INTRODUCTION

1.1 BACKGROUND AND INDUSTRIAL CONTEXT

Manufacturing systems have always evolved in response to the market and technological pressures of their time, and automation is not what distinguishes the present from the past. The mass production systems of the twentieth century were already highly automated, but that automation rested on a stable premise: a standardized product manufactured in high volumes, for which a dedicated production line could be set up and then run for years with minor changes. What has changed are the conditions under which automation must operate. Growing product variety, shorter product life cycles, and rising demand for customization have made the ability to accommodate change, rather than to repeat a fixed sequence at scale, the defining requirement of a modern manufacturing system. As framed in [1], designing and operating such a system is fundamentally a decision-making process. It reconciles the processes, equipment, and organizational structure of the system with the objectives it must serve. These objectives are commonly expressed through four fundamental manufacturing attributes: cost, quality, time, and flexibility. The shift described above amounts to a reweighting of these attributes, with flexibility gaining ground on the cost-per-unit logic that governed the mass-production era.

One response to this pressure has been the development of reconfigurable and flexible manufacturing paradigms, capable of adjusting their functionality and capacity in response to changing requirements rather than remaining fixed to a single, predetermined product or process [2]. This body of work has been developed primarily for conventional, forward manufacturing, where products are assembled from new components. A parallel and increasingly significant

pressure, however, has emerged from the opposite end of the product life cycle: the need to recover value from products once they reach the end of their initial service life, rather than treating them as waste.

This pressure is articulated most clearly through the circular economy, an economic model in which products, components, and materials are kept in circulation at their highest utility and value for as long as possible, in contrast to the linear take-make-dispose trajectory of conventional industrial practice. It is commonly represented through the butterfly diagram Fig. 1, which separates biological cycles, where biodegradable materials are returned to natural systems, from technical cycles, where products and components are kept in circulation through successive loops of reuse, repair, refurbishment, remanufacturing and, as a last resort, recycling [3]. The defining property of the technical side is the ordering of these loops: the inner ones preserve the geometry, the material, and the manufacturing effort already embodied in the product, while each loop outward destroys a greater share of that embedded value, until recycling reduces the product to material stock.

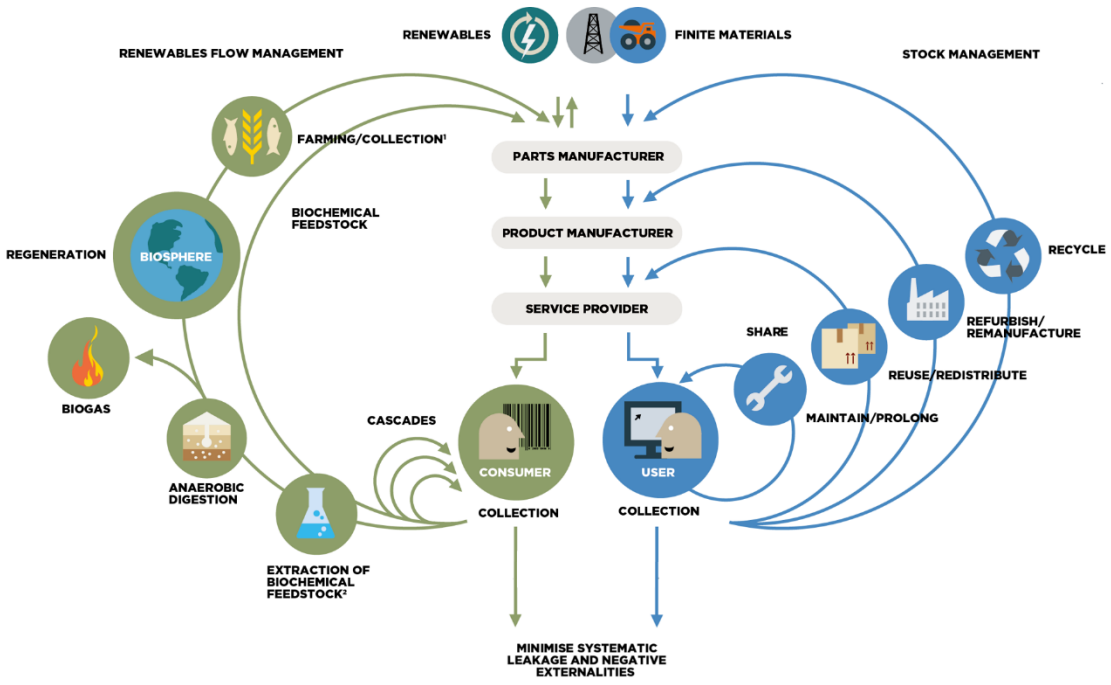


Fig. 1 The circular economy butterfly diagram. (Source: Ellen MacArthur Foundation [3])

This ordering is formalized in the 9R framework Fig. 2, which ranks circularity strategies from R0 (Refuse) and R1 (Rethink) through R4 (Repair), R5 (Refurbish), R6 (Remanufacture) and R7 (Repurpose), down to R8 (Recycle) and R9 (Recover) [4]. The strategies occupying the middle of this hierarchy, referred to collectively in this work as Re-X processes, are those of direct interest from a manufacturing systems perspective, since they are the ones that require physical processing operations to be carried out on a returned product.

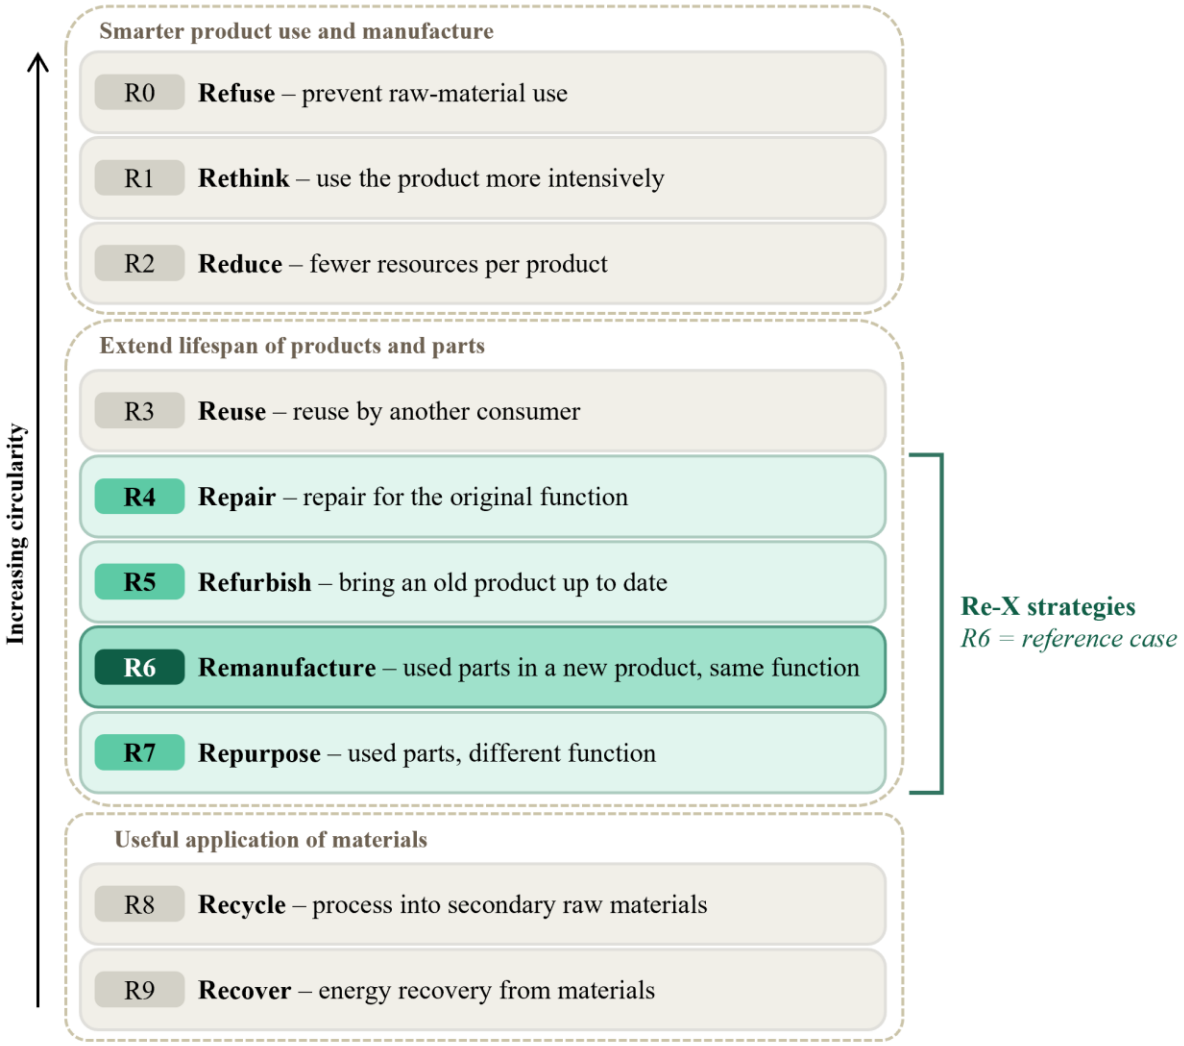


Fig. 2 The 9R framework of circularity strategies, ordered by decreasing value retention.

The shift toward these strategies is not driven by industrial considerations alone. The European Green Deal, the European Union's overarching strategy for reaching climate neutrality by 2050, established circularity as a structural objective of European industrial policy, and the subsequent Circular Economy Action Plan translated that objective into concrete measures spanning product design, waste reduction, and reuse. These measures have since been operationalized through binding instruments. The Ecodesign for Sustainable Products Regulation

extends ecodesign requirements beyond energy efficiency to cover product durability, reparability, and recyclability from the design stage, while the Right to Repair Directive obliges manufacturers to keep repair viable by ensuring the availability of spare parts and repair information, so that products are repaired rather than discarded. For tires specifically, the prohibition on landfilling whole and, subsequently, shredded tires removed disposal as a legal option and redirected end-of-life tires toward recovery and value-retention routes such as retreading [5]. The societal implications extend beyond material savings, since value retention processes are considerably more labor-intensive than the production of equivalent new goods and are therefore associated with the creation and retention of skilled, typically local employment, while reducing exposure to volatile raw-material supply chains. This dimension is explicitly recognized in the Industry 5.0 agenda, which positions human-centricity, sustainability, and resilience as the guiding objectives of the next phase of industrial development [6].

Realizing these benefits at industrial scale nonetheless remains difficult. Reviews of the field highlight that Re-X processes, and remanufacturing in particular as the most demanding among them, continue to face substantial barriers around economic feasibility, process standardization, and industrial adoption, distinguishing them from the comparatively mature and predictable environment of conventional manufacturing [5]. The root of the difficulty lies in the nature of the input. A returned product, commonly referred to as a core, is not a standardized item, its properties are not known in advance and must be established on each individual unit. Its variability has two distinct origins. The first is variant-related: even within a single product family, cores differ in size, configuration, and design generation, so the system must establish what it is dealing with before it can decide what to do with it. The second, and the more consequential, is condition-related: the wear, damage, deformation, and degradation accumulated over an

uncontrolled service life of unknown duration and severity, which cannot be inferred from a nominal specification and must instead be measured directly on each individual item. This thesis takes remanufacturing as its reference case, being the most demanding of the Re-X strategies in both respects, but the requirements that follow from these two origins of variability apply equally to repair, refurbishment, and repurposing.

Addressing this at the level of the individual processing cell is where the pace of the transition is ultimately decided. End-of-life products accumulate at industrial rates, while their processing remains largely manual and is therefore bounded by the availability and endurance of skilled operators, creating a structural mismatch between the rate at which waste is generated and the rate at which it can be processed and returned to use. Conventional fixed-automation systems cannot close this gap, lacking both the dexterity to manipulate and process items whose geometry deviates unpredictably from nominal, and the cognitive capability to assess condition and select a processing strategy without a predetermined program to follow. The established response to comparable limitations in conventional manufacturing has been to distribute the task across resources that cooperate rather than operate in isolation, combining multiple robots and human operators within a shared workspace so that flexibility emerges from the composition of the system rather than from the capability of any individual machine [7]. Supported by ICT technologies such as AI-based perception, digital twins, and item-level traceability, the same principle offers a credible route to automating or semi-automating Re-X processes: automation absorbs the repeatable, physically demanding, and geometrically well-defined portion of the work, while human judgment is retained precisely where automated decision-making reaches its limits, an area of growing research interest particularly regarding its safety and ergonomic dimensions [8]. These

two threads, the operational challenges inherent to Re-X processing and the enabling potential of human-robot collaborative systems, form the background against which this thesis is positioned.

1.2 THESIS AIM AND SCOPE

This thesis aims to propose an end-to-end methodology for the design of robotic cells for Re-X processes, structured around the specific challenges that distinguish the processing of returned products from conventional manufacturing automation. In doing so, it contributes to the twin green and digital transition. Its green dimension lies in industrializing value-retention processes and, with them, reducing waste and demand for primary materials. Its digital dimension lies in making the proposed cells adapt to each item through perception, item-level information, and digitally validated design. The methodology is also human-centric in the sense promoted by the Industry 5.0 agenda, treating the operator not as a cost to be eliminated but as a source of judgment to be supported, reducing physical strain and exposure to hazardous operations while retaining and building on existing craft expertise. The methodology is deliberately driven by these three aspects, green, digital, and human-centric, and aims to combine them within a single design approach. It is validated on a scenario drawn from the heavy-duty transport vehicle tire retreading industry.

The methodology is end-to-end, it follows a core through the complete process, from its arrival as an unknown item to its departure as a processed one. The capabilities required along this path are treated as a single connected chain rather than as separate engineering problems, as summarized in Fig. 3.

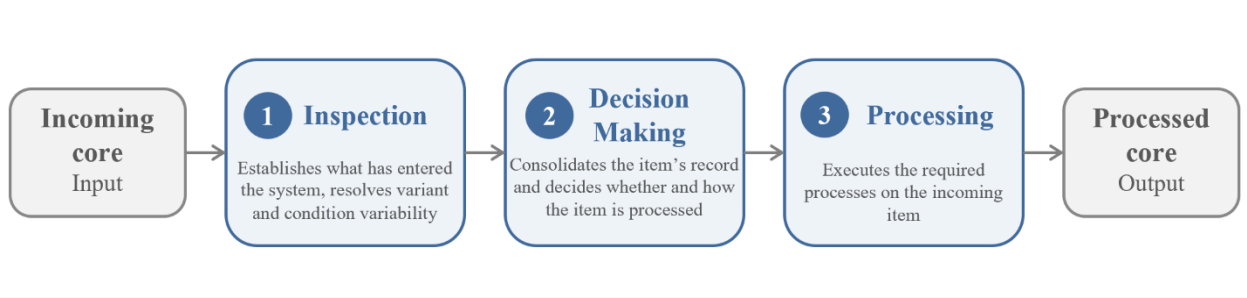


Fig. 3 Methodology overview diagram.

It is organized around three pillars:

- **Inspection** establishes what has entered the cell, resolving both the variant-related and the condition-related variability of the incoming core and generating at runtime the information that, in conventional manufacturing, pre-exists as engineering data.
- **Decision making** consolidates that information into a coherent, item-specific record and uses it to determine whether and how the item should be processed, replacing the fixed process plan of a conventional cell with one derived per item.
- **Processing** executes the resulting operations on a workpiece whose geometry and condition were not known at design time.

A human–robot collaboration strategy runs across all three pillars. It defines where along this chain the operator steps in, and which decisions or operations are handed to human judgment rather than left to automation.

The scope of this claim should be stated explicitly. The methodology covers the processing cell, and the information flows immediately around it. It does not address the reverse supply chain that collects, grades, and routes cores to the cell, the scheduling of work across multiple cells, or the business models behind Re-X operations. Nor does it develop individual process technologies from first principles; it treats the choice of a process as a design decision guided by the item's

condition. Its focus is the reasoning by which these elements are combined into a single, working cell.

The scenario selected for validation is the retreading of heavy-duty transport vehicle tires. It exhibits the defining difficulty of Re-X processing, a variable, non-standardized input, in an unusually pronounced form, while representing an industrial practice of substantial economic and environmental relevance. Within it, the work addresses two stages of the retreading process: defect detection and localization, and skiving, the localized material-removal step that prepares a damaged region of the casing for retreading.

This pairing is not arbitrary: together the two stages exercise all three pillars in sequence. Detection and localization instantiate inspection, since the condition of a buffed casing is neither documented nor predictable and must be established directly on the item. The step between them instantiates decision making, as the detected defects are characterized, recorded against the specific casing, and translated into whether it remains viable and, if so, what must be removed and where. Skiving instantiates processing, its path and extent set by the condition of the individual casing rather than by a fixed program. The scenario also exercises the collaboration strategy, since the operator confirms or overrides the proposed processing and performs it manually where automation defers. Other stages, full-surface buffing, tread application, and curing, fall outside the scope of the cell developed here, whose layout, subsystems, and operational sequence are verified in a digital environment ahead of physical realization.

The contribution of this work is accordingly a design and methodological one. It provides a structured route to designing Re-X processing cells, validated on a demanding industrial scenario, whose value is realized as improved repeatability, consistency, and operator ergonomics and safety.

1.3 THESIS STRUCTURE

This thesis is organized into seven sections, including the present introductory one. Section 2.0 reviews the relevant state of the art, covering automation challenges in remanufacturing systems, human-robot collaborative cells, robotic handling of large or non-standard components, and existing tire remanufacturing technologies. Section 3.0 presents the industrial problem and motivation underlying this work, discussing the challenges of Re-X processing within a circular economy context before introducing the specific reference scenario of heavy-duty tire retreading. Section 4.0 develops the proposed methodology for the design of Re-X cells around the three pillars introduced above, independently of any specific product. Section 5.0 applies this methodology to the tire retreading scenario, detailing the design of the cell's layout, the Powered Roller Gripper, the perception and skiving stations, the digital validation environment used, the human-system interaction and safety strategy, and a conceptual discussion of scalability. Section 6.0 presents the results obtained from the simulation and validation of the developed cell. Finally, Section 7.0 concludes the thesis and discusses directions for future work.

2.0 STATE OF THE ART

2.1 AUTOMATION CHALLENGES IN REMANUFACTURING SYSTEMS

The uncertainty inherent to remanufacturing, discussed in general terms in Section 1.1, becomes a concrete technical obstacle at the level of individual automated processes. Reviews of Industry 4.0 and robotic technologies applied to remanufacturing note that, despite growing interest in applying automation across the remanufacturing lifecycle, from disassembly to reassembly, the field continues to lag behind conventional manufacturing in automation maturity, and that remanufacturing operations remain predominantly manual precisely because core condition cannot be assumed or standardized in the way a newly manufactured part can [8]. This gap is particularly pronounced at the inspection stage, where the condition of a returning core must first be established before any subsequent decision can be made, and where the absence of a fixed reference geometry against which deviations can be measured makes generalizing an automated inspection system across a varied core population especially difficult [9].

This inspection-specific difficulty is representative of a broader pattern across other process stages. Robotic and digital technologies are increasingly proposed and piloted across disassembly, cleaning, sorting, and reassembly, but their practical deployment remains constrained by the same underlying issue, the absence of a standardized, predictable input around which a fixed automated process can be built [8]. Complexity drivers identified specifically for inspection tasks, such as wide variability in defect types and the impracticality of exhaustively characterizing every possible core condition in advance, echoing this same constraint at a narrower, more technical level [9].

This combination of factors, uncertainty that spans the full remanufacturing lifecycle and is particularly acute at the point of condition assessment, helps explain why fully automated approaches to remanufacturing continue to be the exception rather than the norm, and why hybrid, human-involved operation remains, at present, a structural characteristic of the field rather than merely a transitional stage on the way to full automation.

2.2 HUMAN-ROBOT COLLABORATIVE CELLS

Human-robot collaboration has developed into an active area of both academic and industrial interest as a means of combining the repeatability of automation with the adaptability that manual labor still offers in tasks involving variability [10]. Its industrial uptake has concentrated on precisely those processes that resist conventional automation because they involve complexity, variability, or a requirement for flexibility that fixed programming cannot accommodate. Reported deployments span automotive assembly, where collaborative operation was introduced to absorb the variety of a mixed-model line [11]; the handling of non-rigid components in the white goods sector, where part geometry cannot be predetermined and the cell must combine perception with a compliant end-effector [12]; and high-payload operations, in which the robot carries loads beyond human capability while the operator retains control of the tasks [13].

More recent activity indicates that the same capability has the potential to support processing at the other end of the product life cycle. Reviews of human-robot collaborative disassembly argue that fully automating the recovery of end-of-use products is not economically viable given their complexity and variation, and that collaboration offers a middle path, pairing

human flexibility and problem-solving with robotic precision and the handling of unsafe tasks [14]. A related review of collaborative remanufacturing reaches a similar conclusion, linking the move away from manual, labor-intensive disassembly to the growing volume of end-of-life products entering recovery [15]. Beyond these reviews, recent work has begun to demonstrate this in practice, applying human-robot collaboration to remanufacturing workstations through a methodology validated on a household-appliance use case, on an explicitly human-centric basis [16]. The case for collaboration in Re-X processing therefore rests on the same argument that motivated it in conventional manufacturing but applied to a setting where the variability that justifies it is considerably more pronounced.

Three dimensions recur across this literature, and together they define what a collaborative cell must resolve. The first is safety. Design guidance for collaborative workplaces addresses the layout, resource allocation, and protective measures that make a shared workspace safe from the outset rather than retrofitted after the fact [17]. At the control level, adaptive speed and separation monitoring switches between safety zones, letting the robot adjust its behavior to the operator's actual position [18], instead of enforcing a single conservative limit at all times. A systematic review of the field confirms that occupational health and safety criteria remain of central importance to collaborative implementations, and that the design of safe and ergonomic collaborative workcells is itself an open research field rather than a solved engineering problem [10].

The second is interaction: how the operator understands what the system is doing, and how intent is communicated back to it. Augmented reality has been applied to deliver task information within the operator's own field of view, both for guidance during execution and for training on hybrid processes [19], while action recognition frameworks work in the opposite direction,

inferring the operator's current activity from observed motion so that the system can adapt its own behavior accordingly [20], [21]. The shared objective is to reduce the interpretive burden placed on the operator, so that the state and intent of the cell are legible without recourse to a separate interface.

The third is collaboration in the strict sense, the highest level of the common human-robot interaction taxonomy (coexistence → cooperation → collaboration) [7] , in which human and robot act at the same time on the same task or workpiece, rather than sharing a workspace while working separately. Where the workpiece is deformable and its geometry cannot be assumed, model-based motion planning for co-manipulation has been proposed, in which the object's distortion is simulated in real time so that robot actions remain consistent with the operator's handling inputs. At the level of task rather than motion, dynamic allocation strategies distribute work according to real-time indicators rather than a fixed assignment: an industrial case involving a heavy, large part processed under the pace constraint of an upstream injection molding process allocated tasks according to buffer level, perceived physical fatigue, and operator hand preference, reporting reductions in total cycle time alongside improved workload balance and reduced perceived fatigue [22].

Across all three dimensions, a consistent distinction emerges between a cell that is merely technically safe for a human to enter and one that is actually designed around collaborative operation. The former treats human presence as a hazard to be contained; the latter treats the operator as an active, monitored participant in the cell's operating logic, with a role that adapts to changing conditions. For Re-X processing, where operator judgment is not a fallback but a functional requirement of the process, this distinction is not a refinement but a starting condition.

2.3 ROBOTIC HANDLING AND PROCESSING OF LARGE OR VARIABLE-GEOMETRY COMPONENTS

Robotic handling of components that are large, non-rigid, or otherwise variable in geometry departs significantly from the assumptions underlying conventional gripper design, which typically presumes a known, fixed part geometry. A comprehensive review of deformable-object grasping in industrial applications classifies the objects, grippers, and gripping strategies in use, and draws a distinction between model-based and model-free control [23]. Model-based approaches need prior knowledge of the object's rigidity and shape, so they tend to fail when those properties are not known in advance. Model-free approaches need no such model, only a relationship between input and response, so the gripper reacts to the object's behavior at the point of contact instead of requiring it to be characterized beforehand, which makes them far better suited to varied, unpredictable objects. Reviews of automated assembly with non-rigid objects reach the same conclusion, despite academic progress, industrial adoption remains limited, because handling objects whose shape cannot be fixed in advance is mechanically and control-wise complex [24].

Beyond grasping strategy alone, reliable handling of variable-geometry components has been shown to depend on tightly coupling perception and mechanical design, rather than treating gripper design and object sensing as separate concerns. This is demonstrated in the development of a robotic cell for the white goods industry, centered on the perception and handling of deformable objects whose exact geometry cannot be predetermined, which combines vision-based perception with an end-effector designed to accommodate that variability [12], and is further supported by co-manipulation work relying on model-based motion planning to compensate for

the absence of a fixed object geometry during handling [25]. Taken together, this literature points to a broader principle relevant to any system handling large or variable-geometry components: mechanical adaptability at the gripper level and the ability to sense or estimate the object's actual condition at the point of handling are not independent design choices but need to be developed jointly.

2.4 TIRE REMANUFACTURING TECHNOLOGIES

The preceding sections establish the general state of the art. It is equally necessary to examine what has already been established for the specific process selected to validate this work. The extent to which prior work has addressed automated inspection and localized processing in that process determines both what the validation case can demonstrate and where the contribution of this thesis lies.

Tire remanufacturing through retreading is a long-established industrial practice, and its role within circular economy strategies has been reaffirmed by recent literature reviews of the field. One such review surveys the broader landscape of tire remanufacturing research [26]. It notes that retreading remains one of the most mature and widely practiced forms of remanufacturing, owing to the substantial cost and material savings it offers over producing a new tire. It also observes that individual retreading operators have historically developed their own methods and quality practices, in the absence of a broadly standardized process.

Within this broader landscape, automated defect identification and localized processing has received comparatively limited attention. The closest peer-reviewed work addresses a related but distinct problem: end-of-production quality control on newly manufactured tires, where a two-

stage multimodal deep learning approach detects surface defects and classifies their severity, supported by a dedicated acquisition machine whose cameras reposition automatically to accommodate the range of tire dimensions [27]. The detection task itself is comparable to the one addressed here, and the work establishes that automated surface defect detection on a tire is feasible at industrial resolution. The conditions under which it operates, however, differ in the two respects that matter most for remanufacturing. The inspected tire is new and unbuffed, so a nominal reference geometry exists against which deviations can be measured, and the outcome is a quality verdict rather than an input to a subsequent material-removal operation. Neither condition holds for a buffed casing entering retreading, where the reference does not exist and the detection result must drive a physical process.

The step from detection to processing is addressed chiefly in the patent literature. One directly comparable example is a patented process and machine for identifying and working defects on used tires [28]. The scraped tire is rotated through at least a full revolution on a centering support, and its working surface is scanned to identify defect morphology. Each detected defect is compared against a virtual reference library to select both the operation and the tool. The selected tool is then loaded onto an anthropomorphic robot, which performs the working to form a crater suitable for subsequent coating. The same detect-classify-machine logic has been patented for generic workpieces, where an optical inspection system localizes a defect, determines its three-dimensional topography, assigns it to a defect category, selects a machining process accordingly, and generates the corresponding robot program [27], indicating that the pipeline is not specific to tires. Robotized skiving itself appears in earlier patented work, in which a robot-mounted rasp executes a cutting pattern on a casing while an optical sensor scans the cut region to detect exposure of the reinforcing wire [29].

Detection by non-visual means has also been explored. One line of patented work identifies delamination and related faults by impacting the tire with an actuator and analyzing the resulting sound wave, using a neural network trained on discriminator quantities computed from the waveform to separate genuine defects from benign features such as surface lettering and wall thickness variation [30]. A later development combines impact force with impact sound under an artificial neural network and is aimed explicitly at characterizing whether a casing is suitable for retreading rather than at localizing individual defects [31]. Both are oriented towards internal condition rather than the localized surface and near-surface defects that concern the present work, and the second states the distinction directly in seeking an overall characterization rather than a defect location.

Taken together, detection and processing have largely been pursued separately, and the peer-reviewed contribution concentrates on tires whose nominal condition is known. Automated defect detection has been demonstrated for end-of-production quality control, while the coupling of detection to an automated, localized material-removal operation on a used, buffed casing appears principally in patents. Automated, full-surface defect detection and localization integrated with the subsequent skiving operation, within a cell that also accommodates operator judgment where automated decision-making reaches its limits, therefore remains underexplored in the published literature.

3.0 INDUSTRIAL PROBLEM AND MOTIVATION

3.1 CHALLENGES IN RE-X PROCESSING

Section 1.1 identified the two origins of variability in a returned core, its variant and its condition, and the mismatch between the rate at which end-of-life products arise and the rate at which they can currently be processed. This section examines what that variability does once it enters a process chain. The difficulty is not only that the incoming core is uncertain, but that this uncertainty must be resolved step by step as the core passes through the operations. An automated Re-X system must therefore be built around managing this uncertainty across the operation chain, rather than around its presence at the input alone. Four distinct consequences follow, discussed in turn below.

The first is an inversion of the normal relationship between a product and the data describing it. In conventional manufacturing, the specification precedes the workpiece: geometry, tolerances, and material properties are fixed at design time, and every subsequent engineering decision is derived from it: the fixtures that locate the part, the toolpaths that machine it, and the tolerances against which it is later inspected are all defined in advance of the first unit being produced. In Re-X processing, the workpiece precedes the data. The item arrives first, and the information needed to process it must be produced from the item itself. Every process decision is therefore downstream of a measurement rather than downstream of a specification, and the quality of that measurement bounds the quality of everything that follows it. The pattern holds across sectors: an electric motor whose winding insulation has aged unevenly, a household appliance carrying corrosion and impact damage in unpredictable places, and a structural assembly whose

service loading is known only in aggregate all present the same problem, differing in the sensing required rather than in the nature of the difficulty.

The second is that these sources of variability multiply rather than add. A single core is described not by one unknown but by several that vary independently: which variant it is, the condition of each of its subassemblies, where on the item each instance of damage falls, and how those instances combine with one another. Because these dimensions are independent, the number of distinct overall states a core can present is obtained by multiplying the number of possibilities along each dimension, not by adding them. A modest count on each axis therefore yields an enormous number of combinations: ten variants, ten condition grades per subassembly across a handful of subassemblies, and any of several damage locations already produce more distinct situations than could be individually foreseen. The practical implication concerns how such a system can be programmed. Conventional automation is efficient because a finite list of situations can be drawn up in advance, each with a defined response; at runtime the machine identifies which of these listed situations it is facing and carries out the response prepared for it. This approach becomes untenable when the number of possible situations grows multiplicatively, since no realistic effort can pre-author a response for every combination, and any combination left unlisted is one the system cannot handle. The difficulty is thus not that the programming task is laborious but that it is the wrong strategy: a fixed catalogue of situation-response pairs cannot span a space generated by the free combination of several independent sources of variation. What is required instead is a system that determines its response from the measured condition of the individual core at runtime, rather than retrieving it from a list prepared beforehand.

The third concerns the order in which information and irreversible action occur. In conventional manufacturing, knowledge precedes action throughout: the specification exists

before the first operation, and every subsequent step is planned against it before any material is touched. In Re-X processing, the two are interleaved, because part of the information needed to decide what to do can only be obtained by first doing something that cannot be undone. A core must often be disassembled before the condition of its subassemblies can be assessed, a coating stripped before the substrate beneath it can be examined, or a worn layer removed before what remains of the underlying structure can be measured. The operation that reveals the condition is therefore also the operation that forecloses some of the options. This holds regardless of the direction in which material moves. Re-X spans both removal, disassembly, stripping, machining and addition, cladding, weld repair, patching, recoating. Both directions commit the item: material deposited over an unassessed defect is as irreversible as material incorrectly removed. The consequence for design is an ordering problem. Operations must be sequenced so that the least reversible are deferred until the information justifying them has been obtained, and the inspection performed at each such point must be sufficient at that point, rather than treated as something later stages can correct.

The fourth appears at the level of the system rather than the individual operation. Because the extent of processing required is determined per item, and because a proportion of items is rejected at a point that cannot be predicted in advance, neither the work content nor the yield of a Re-X line is deterministic. Cycle time varies from one item to the next, some items leave the line before reaching its end, and buffer levels fluctuate as a function of the incoming population rather than of the line's own behavior. Such a system cannot be balanced against a fixed takt in the way a conventional line can, and its capacity is properly expressed as a distribution rather than as a rate.

Taken together, these four consequences define what an automated Re-X system must do that a conventional one need not: generate the missing information at runtime rather than inherit

it, select a response from a space too large to enumerate rather than execute a fixed program, commit to irreversible actions under asymmetric risk, and absorb variation in work content without a fixed takt to fall back on. Each is a requirement on the system rather than a description of the domain, and each recurs in the specific process examined in the sections that follow.

3.2 REFERENCE MANUFACTURING SCENARIO

The challenges discussed in Section 3.1 are general to Re-X processing as a domain. This thesis grounds them in a specific reference scenario: the remanufacturing of heavy-duty transport vehicle tires through retreading. The scenario was selected because it exhibits the difficulty of a variable, non-standardized input in a particularly pronounced form, while representing an industrial practice of substantial economic and environmental relevance given the scale at which heavy transport tires are used and replaced at end-of-life.

Retreading proceeds through five principal stages, shown in Fig. 4. A worn tire first has its remaining tread material removed through a buffing operation, exposing the underlying casing in preparation for the application of new tread.

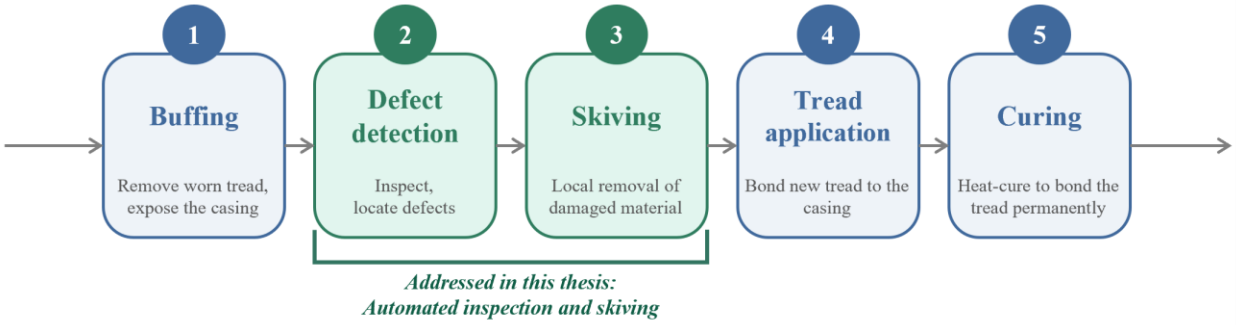


Fig. 4 The main five stages of the tire retreading process, with defect detection and localization and skiving marked as the two stages addressed within the scope of this thesis.

It is at this stage, once the tire has been buffed, that its individual condition becomes fully apparent and consequential. The buffing process itself removes material unevenly, and the exposed casing may reveal defects, whether pre-existing damage from the tire's service life or imperfections surfaced by buffing, that must be identified and addressed before the casing can be judged suitable for retreading to continue. Unlike the tire's original manufacture, where geometry and material condition are controlled by design, the buffed casing is in effect a unique object each time, requiring individual attention rather than a predetermined sequence of operations.

Addressing this requires two related capabilities. First, the casing must be inspected to detect and localize any defects present on its surfaces, since without this information no informed decision can be made about how, or whether, the casing should proceed. Second, once a defect has been located, the affected area must be prepared through a localized material-removal step, referred to as skiving, which removes damaged or unsuitable material from the identified region so that the casing surface is brought to a condition suitable for the process to continue.

In current industrial practice, both capabilities rely heavily on operator experience. Inspection is performed visually by trained personnel, and skiving is carried out manually, following strict repair specifications applied through the operator's own judgment of the defect's extent and the appropriate amount of material to remove. Reaching this level of judgment takes years of expertise, and operators work unsupervised only once that experience has been built up. Fig. 5 shows the workstation at which these operations are currently performed and an operator carrying out skiving by hand.

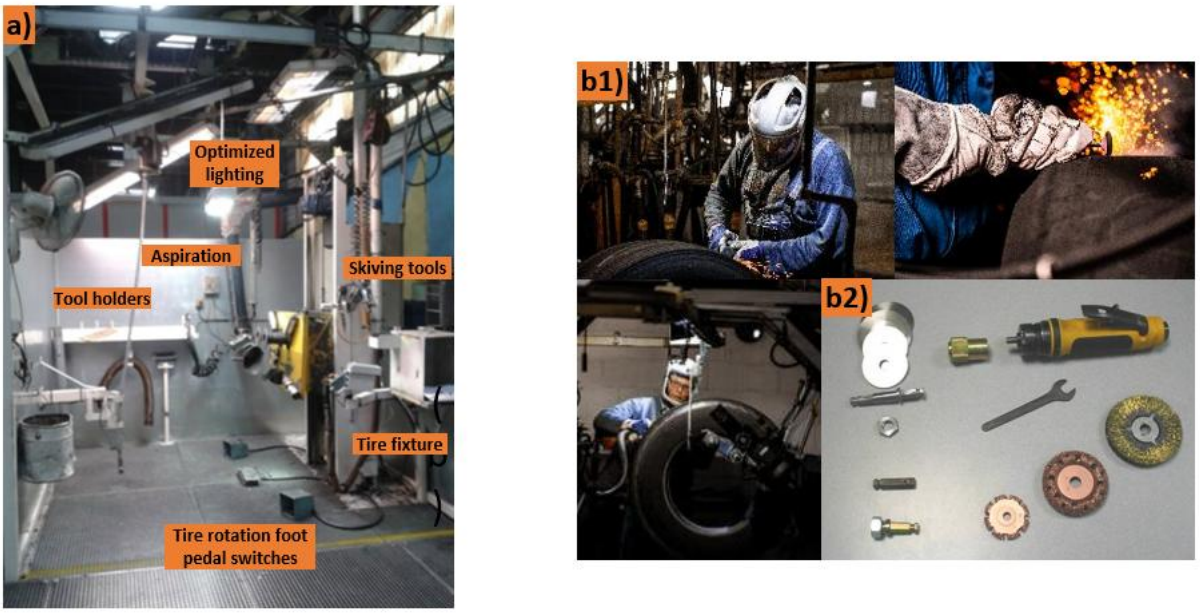


Fig. 5 Current industrial practice at the reference site: (a) the manual skiving workstation, (b1) an operator performing skiving on a buffed casing, (b2) the tools used for manual skiving

While operator expertise is a genuine asset, reliance on manual execution introduces variability of its own, since outcomes depend on individual skill and day-to-day consistency, and it exposes operators to a physically demanding, repetitive task carried out in close proximity to rotating equipment. The motivation is therefore not to remove the operator, whose judgment remains necessary under the input variability described in Section 3.1, but to support that judgment with automated defect localization and consistent, repeatable handling of the casing itself.

These considerations establish the motivation for the cell but not yet a specification. To design against it, the scenario must be expressed in quantities a designer can build to: the physical characteristics of the casing that any handling solution must accommodate, the defects the perception system must detect and the processing must address, and the terms on which the

resulting cell can be assessed. The following section establishes these and derives from them the concrete requirements that constrain the design developed in Section 5.0 .

3.3 PRODUCT SPECIFICATIONS AND DERIVED REQUIREMENTS

3.3.1 Product specifications and physical characteristics

The reference scenario is built around the 315/80R22.5 casing. This is the single most widely used tire size on trucks and heavy transport vehicles, and it is the workhorse size of the line-haul whose tires enter retreading in the greatest numbers, which makes it both the most representative case for the work and the one of greatest practical consequence. The size designation is itself informative: the 315 gives the nominal section width in millimeters, the 80 the aspect ratio as a percentage of that width, the R radial (steel-belted) construction, and the 22.5 the rim diameter in inches. These fixed geometric facts, encoded in the designation, already establish much of the physical envelope that any handling and inspection solution must accommodate, and are collected with the diametric and mass figures in Table 1.

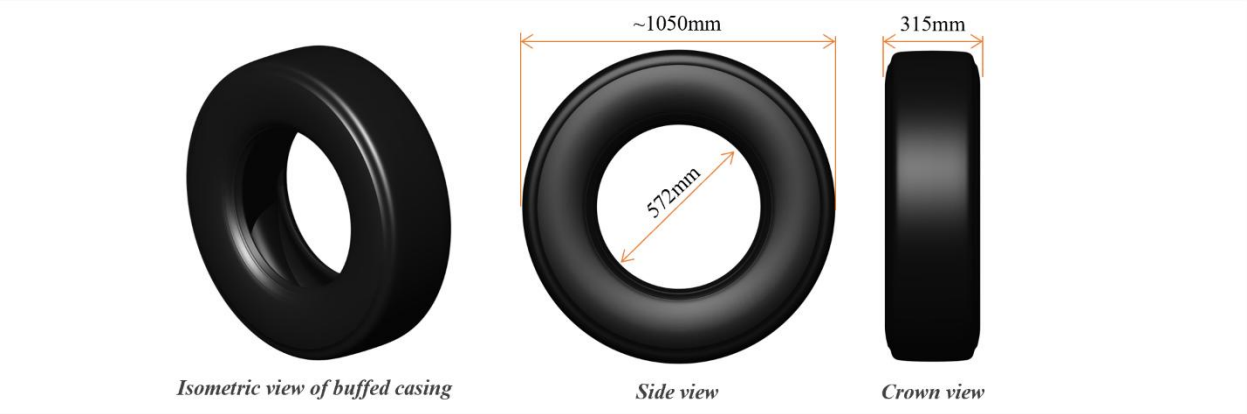


Fig. 6 3D-model buffed tire casing and its main dimensions of the widely used size 315/80R22.5.

The new-tire figures correspond to the nominal outer diameter of approximately 1074 mm and a mass in the region of 68–70 kg for this size. The buffed-casing figures are estimates rather than nominal values, because the buffed casing is not a standardized product. Buffing removes the worn tread down to a thin tread left above the belt package, so that the diameter is reduced by roughly the combined depth of the removed tread on both sides of the crown, and the mass falls by the mass of the rubber removed. The ranges should be treated as working values for dimensioning the cell and replaced with measured figures from the project product data where available.

Table 1 Physical characteristics of the reference tire, as a new tire and as a buffed casing.

Property	New tire	Buffed casing
Size designation	315/80R22.5	—
Section width	315 mm (nominal)	≈ 315 mm

Property	New tire	Buffed casing
Aspect ratio	80 %	—
Rim / bead diameter	22.5 in (572 mm)	22.5 in (572 mm)
Outer diameter	≈ 1074 mm	≈ 1040–1055 mm (est.)
Mass	≈ 68–70 kg	≈ 58–63 kg (est.)
Construction	Radial, steel-belted	Radial, steel-belted

Two characteristics of the buffed casing matter for the cell beyond these nominal figures. First, the casing is handled without a rim and without inflation pressure, so it is not dimensionally stable, it deforms under its own weight and takes up a shape that depends on how and where it is supported, rather than holding the toroidal form it has when mounted and inflated. Second, and as a direct consequence of the buffing process, the outer diameter of any individual casing is not known in advance. Buffing removes material unevenly and to a depth set per casing by its condition, so the actual diameter varies from one casing to the next and cannot be read from the size designation. It must instead be established on each individual item, a measurement the cell performs at pick-up as described in Section 5.2.

3.3.2 Defect taxonomy

The defects that must be detected on the buffed casing, together with their location, the response each requires, and the criterion beyond which the casing is rejected rather than repaired, are summarized in Table 2 and illustrated in Fig. 7. This taxonomy defines the defect types the perception system must detect, and the processing stage subsequently acts upon. The detailed classification and severity assessment of individual defects lies outside the scope of this thesis; what matters for the design of the cell is the range of defect types it must accommodate, their distribution across the casing surfaces, and the fact that the response in every repairable case reduces to a localized material-removal operation, the skiving step introduced in Section 3.2.

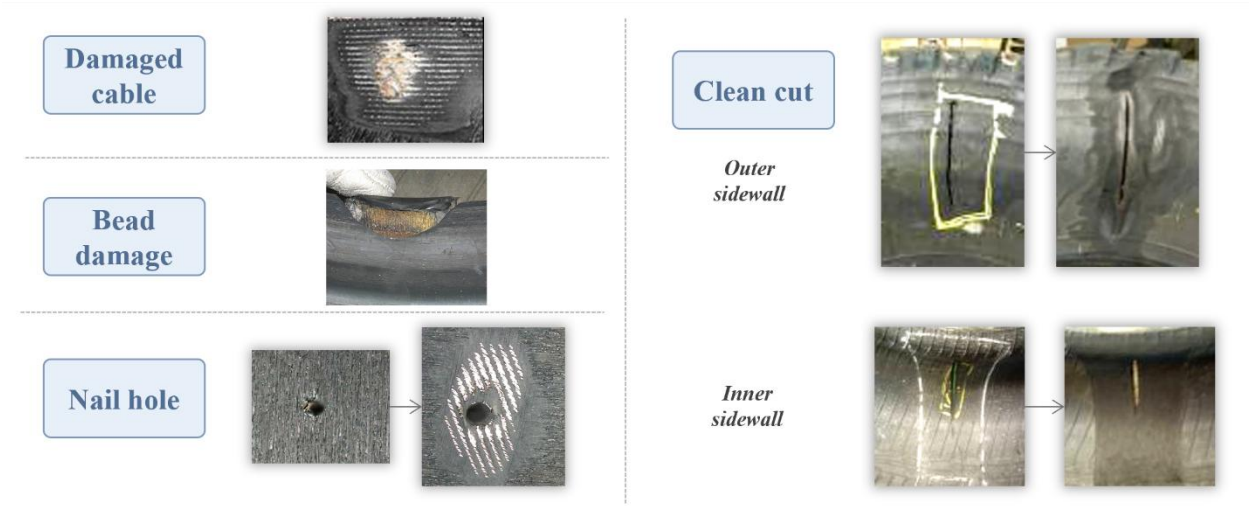


Fig. 7 Representative examples of the defect on a buffed casing.

Table 2 Defect taxonomy for the reference scenario.

Defect type	Typical extent	Location	Required response	Rejection criterion
Damaged cable (steel cord)	Exposed or broken steel cords over a local area of the belt package	Crown	Remove damaged rubber and steel around the injury	Broken cords over an area larger than ≈ 130 × 20 mm on the crown or ≈ 40 mm from the shoulder edge, or damage reaching the bead wire bundle; more than the allowed number of repairs per casing
Nail hole	Puncture up to ≈ 10 mm in diameter; larger penetrations addressed as a larger localized repair	Crown	Remove damaged rubber and steel around the hole	Located outside the crown, or exceeding the repairable size for that region
Bead damage	Damage confined to the bead rubber or chafer, or reaching the bead wire bundle	Bead	Remove and rebuild the damaged bead rubber where the wire bundle is intact	Bead wire bundle bent, broken, or distorted → casing scrapped
Cracks	Shallow surface cracking, or deeper cracks extending toward the casing cords	Sidewall, crown, inner surfaces	Remove cracked material down to sound rubber	Cracks reaching the bead wire bundle, or extending through the casing structure
Clean cut	Cut up to ≈ 130 × 20 mm on the crown or ≈ 127 × 41 mm on the sidewall, provided the cords are not fully severed	Crown, sidewall	Remove damaged material to a clean crater	Cut exceeding the repairable size for its region, severing the casing cords, or reaching the bead wire bundle

The extent to which a given defect can be repaired depends on where it lies on the casing, and the repairable size differs by region: approximately 130 × 20 mm on the crown, up to approximately 40 mm from the shoulder edge, and approximately 127 × 41 mm on the sidewall. Damage confined to the rubber, including at the bead, is recoverable by removing the affected

material and rebuilding it, whereas damage that reaches the bead wire bundle or otherwise severs the load-bearing structure of the casing cannot be restored by material removal and renders the casing unfit for retreading. These bounds define the defect population the cell is designed to process: the smallest defect that must be reliably detected sets the resolution required of the perception system, while the tolerance on the resulting skived region sets the accuracy to which a defect must be localized and the tool positioned, both of which Section 5.0 depends on.

3.3.3 Performance specifications

This subsection sets out the performance targets the developed system is designed to meet. They define what the system must achieve as designed, expressed so that each can be verified against the design and its digital validation. Three targets are set, spanning automated inspection, automated processing, and operator safety, and the extent to which each is met is reported in Section 6.0 .

The first target is **full-surface automated inspection:** the system must be capable of automated defect detection and localization across the full surface of the casing, spanning both its outer and inner surfaces, with no region left uninspected. For each defect found, the inspection must yield its location on the casing, its type, and the parameters describing its extent, since these are the quantities the subsequent processing decision and skiving operation act upon. Because a defect may lie anywhere on a buffed casing, partial coverage would leave the condition assessment incomplete and undermine every processing decision that follows from it.

The second target is **full-surface automated processing**: the system must be capable of performing the material-removal (skiving) operation across the full surface of the casing, again

spanning both the outer and inner surfaces, so that any localized defect the inspection stage identifies can subsequently be reached and processed. This target is stated separately from the first because inspection reach and processing reach are satisfied by different means and are not equally straightforward to achieve, as the requirements that follow make clear.

The third target concerns **operator safety.** The complexity and variability of the process mean that the operator must be able to intervene directly in the process, so a fenceless human–robot collaboration is required. Because this places the operator in close proximity to a high-payload manipulator handling a heavy casing, the target is that every safety function protecting the operator is built throughout from components and measures rated to a Category 3 / Performance Level d baseline under ISO 13849-1, with the corresponding SIL 2 under IEC 62061, so that the collaboration is realized at the safety integrity expected for an application of this class.

Taken together, these targets frame the developed system as successful if it can inspect and process the full surface of the casing automatically while keeping the operator safe by design. All three are collected in Table 3, and the extent to which each is met is reported in Section 6.0 .

Table 3 Targets set for the developed system.

#	Target	Definition
T1	Full-surface automated inspection	Automated defect detection and localization across the full casing surface, outer and inner, with no region left uninspected
T2	Full-surface automated processing	Automated material-removal (skiving) across the full casing surface, outer and inner, at any localized defect identified by inspection

#	Target	Definition
T3	Operator safety under HRC	Every operator-protecting safety function designed to a Category 3 / PL d baseline (ISO 13849-1; SIL 2 under IEC 62061)

3.3.4 Requirements derived for the cell design

The physical characteristics of the casing (Section 3.3.1) impose the first group of requirements. Because the casing must be carried and held for the full duration of the cycle at the mass and diameter of Table 1, the cell must command the corresponding payload and working volume. Because the casing is handled without a rim or inflation and is therefore not dimensionally stable, it must be actively held in shape rather than merely gripped at a point, since a casing left to deform under its own weight cannot present a repeatable surface to either station. And because a defect may lie anywhere on the casing, on the crown, either shoulder, either sidewall, or the inner surfaces, every point of both surfaces must be brought within reach of the perception sensors, and, separately, within reach of the skiving tool. These two reachability demands are stated separately, for perception and for processing, because they are satisfied by different means and, as shown below, are not equally easy to meet.

The defect taxonomy (Section 3.3.2) imposes the second group. The smallest defect that must be caught sets the resolution the perception system has to achieve, and the tolerance on the resulting skived region sets the accuracy to which a defect must be localized and the tool positioned. The processing itself must then produce, per defect class, the specific prepared geometry that retreading requires. The interleaving of inspection and processing established in

Section 3.1 adds a requirement of a different kind. Inspection and skiving are distinct operations, and a defect once located during inspection must still be found again when it is time to process it. The requirement is therefore that a defect's location, once established, remains associated with the casing and recoverable at the point of processing, however that association is maintained.

Finally, the input variability of Section 3.1 means that not every defect can be expected to be handled automatically. This leads to a requirement that applies no matter how much of the process is automated: whenever an automated decision is uncertain, or a defect falls outside what the automation can handle, a human must be able to step in, either to make the decision or to carry out the operation directly. The cell must therefore be designed to accommodate human intervention rather than assume full autonomy.

Table 4 Requirements derived from the reference scenario.

#	Requirement	Derived from	Constrains
R1	Handle and hold the buffed casing throughout the cycle at the mass and diameter of Table 1	3.3.1	Robot payload class, gripper actuation, cell footprint
R2	Maintain control of casing shape during handling, not merely grip it	Absence of rim and inflation, 3.3.1	Gripper architecture
R3	Present every point of both casing surfaces, outer and inner, to the perception system	Defects occur on crown, shoulders, and sidewalls, 3.3.1	Casing rotation mechanism, scan strategy
R4	Present every point of both casing surfaces, outer and inner, to the skiving tool	A defect requiring skiving may lie anywhere R3 can detect one	Casing rotation and re-orientation, tool degrees of freedom
R5	Detect defects down to the minimum extent in Table 2	3.3.2	Sensor resolution, standoff distance, scan density
R6	Localize defects to within the skiving tolerance	3.3.2	Calibration chain, robot repeatability

#	Requirement	Derived from	Constrains
R7	Execute skiving to the crater geometry required per defect class	3.3.2	Tool selection, force control, path generation
R8	Retain the defect location established at inspection and reproduce it at the skiving tool without re-detection	Interleaving of inspection and processing, 3.1	Traceability layer, angular encoding
R9	Allow the operator to confirm or override the proposed processing path, and to execute skiving manually where automation defers	3.3.1, 3.3.2	Interaction design, safety strategy, cell layout

Most of these requirements constrain a single subsystem and are addressed directly in Section 5.0 . One interaction among them, however, shapes the entire cell, and is worth drawing out here because it forecloses the arrangement that first suggests itself.

The interaction is between R1 and the two reachability requirements, R3 for perception and R4 for skiving. Together they demand that every part of the casing surface be brought before the perception sensors and, later, before the skiving tool, since a defect may lie anywhere and must be reachable both to be seen and to be processed. The natural way to achieve this is to set the casing down on a powered turntable that spins it in place, with a fixed perception station on one side and a skiving robot on the other reaching in to process the defects the perception station has found. For perception, having a sensor array around a rotating casing can be arranged to view the whole of it. It breaks down at skiving, rotation about a single vertical axis presents the crown and the outer circumference to a robot reaching in from the side, but it does not present the inner sidewall, the bead region, or the underside of the shoulder, which face inward or downward in every way the casing is turned. Reaching a defect on those surfaces would require the skiving robot to carry many additional degrees of freedom to work around and inside the casing, or a further

means of flipping and re-seating the casing on the turntable between passes. Either way, a first robot is still needed simply to load the casing onto the turntable, since a buffed casing of the order of 60 kg (Table 1) is an order of magnitude beyond the payload of the collaborative arms usual in human-robot cells. The result is a cell built from a loading robot, a powered turntable heavy enough to hold and spin that casing, and a skiving robot, in which reaching the whole surface still is not guaranteed without additional motion or re-seating, and in which the loading robot stands idle through the processing it enabled.

The alternative is to have a single robot both carry the casing and hold it in a gripper that can itself rotate it, presenting the surface to stationary stations rather than reaching in to a turned casing. Reach and re-orientation then come from the same resource: the robot's own motion places any region of the casing at a fixed perception station and a fixed skiving station, and the gripper's rotation sweeps the full circumference, while a change of the robot's pose brings the inner sidewall and shoulder into view and into reach that a turntable could not. The turntable is no longer needed, and the robot that would otherwise only have loaded it is no longer idle, since that same robot now performs part of the processing. This is the route taken in this thesis, and it turns a rotating, holding gripper from a design preference into a requirement: R1, R2, R3, and R4 are satisfiable together only if handling, rotation, and reach for both stations are consolidated into one active end-effector on one manipulator. Separating perception reachability (R3) from skiving reachability (R4) is what exposes this, since a fixed sensor array can be arranged to see a rotating casing, but no fixed arrangement reaches an arbitrary inner-sidewall defect on one.

This establishes the central requirement carried into Section 5.0 : a robot able to handle and manipulate the casing at the mass and reach of Table 1 that being a high-payload manipulator, rather than one in the payload class typical of collaborative cells. This choice has a direct

consequence for safety, which Section 5.0 addresses in full: a manipulator moving a heavy casing carries enough energy that operator protection cannot rest on power and force limiting alone, and must instead be built on maintaining separation between the operator and the moving casing, with monitored stopping when that separation is not assured. Section 5.0 discusses the gripper, the stations, and the safety case that follow from this choice.

4.0 APPROACH

4.1 REMANUFACTURING-DRIVEN DESIGN REQUIREMENTS FOR ROBOTIC CELLS

A conventional robotic cell is typically engineered around a part that is known in advance. This is a product-oriented design approach, the cell is built to a specified product or product family, as captured in design frameworks such as Product Line Engineering (PLE). In most cases its geometry, tolerances, and material properties are available from CAD models or datasheets before the cell is built, and the cell's programming, tooling, and process logic can be fixed accordingly. Much of established robotic cell design depends on the availability of this prior description of the part: the robot's motions can be programmed offline against the known geometry, and the tooling and fixtures can be shaped to fit it, all before the first item is handled. A cell for Re-X processing generally cannot rely on it to the same degree. The incoming item, commonly termed the core in the remanufacturing literature [15] has already completed one or more service lives before it arrives. Its original specification may still be known, but its present state is not: its geometry may have drifted from nominal, it may carry damage that is not externally visible, and its service history is usually unavailable when it enters the cell. The item is therefore largely an unknown that must be characterized before any processing decision can be made, rather than a known input the processing can be planned around beforehand.

This shift, from a largely known input to a largely unknown one, propagates into the design decisions that follow. Where a cell built around a single well-defined part can rely on fixed programming, dedicated fixtures, and a predetermined process sequence, a Re-X system must be

designed on the expectation that incoming items vary from one to the next and that the processing path cannot be fully fixed at design phase. It must instead be able to build the missing knowledge at runtime, for each item, and act on it. This section sets out the three capabilities the system must build for itself in order to do this:

- **Condition assessment:** establishing what has actually entered the cell, since that information does not pre-exist and must be generated on the item itself.
- **Traceability:** retaining and organizing that information into a per-item record, and connecting it to the item's wider digital passport.
- **Processing strategy:** using that information to determine how the individual item is processed.

These are shown as a staged workflow in Fig. 8, and each is developed in the subsections that follow.

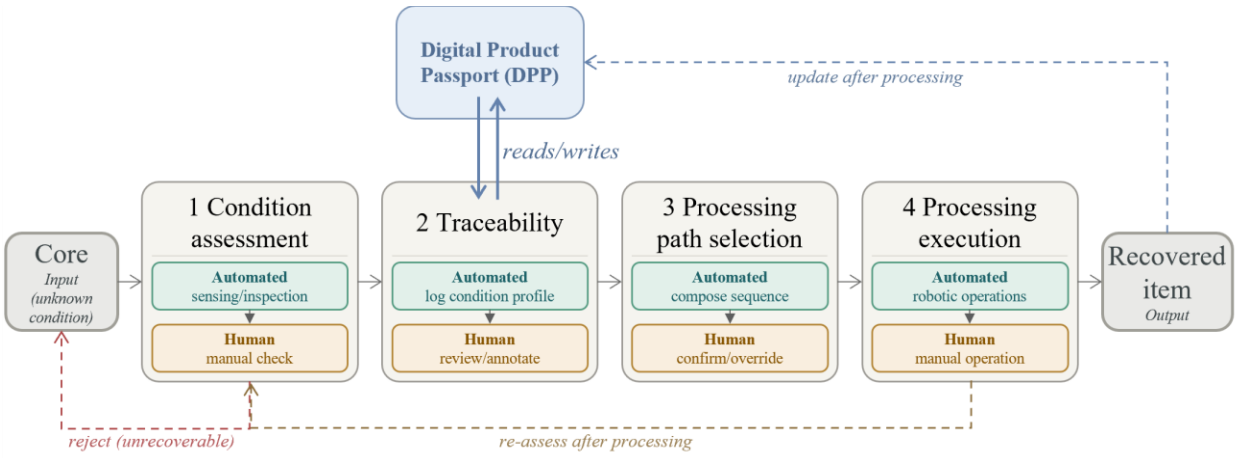


Fig. 8 The proposed workflow for Re-X system operation.

4.1.1 Condition assessment

Before any operation can be planned, the system has to establish what has actually entered it. In a conventional line this information would be available in advance, from engineering drawings and product specifications; for a returned core it does not pre-exist and must be generated on the item itself. The core carries two distinct kinds of variability, and the assessment stage exists to resolve both.

The first is identity variability. Even within a single product family, cores may differ in size, configuration, or design generation, and the system must determine which variant it is handling before anything else can proceed. The second, and the more demanding, is condition variability: the wear, damage, deformation, and degradation accumulated over an unknown service life, which cannot be read from the nominal specification and must be measured directly.

Condition variability is not a single quantity but a range of distinct phenomena, differing in their physical nature, their cause, and the sensing needed to detect them. It is useful to separate them by type. Dimensional and geometric deviation is the extent to which the item's shape has drifted from nominal, through wear, creep, or plastic deformation, detected by comparing measured geometry against a reference. Surface defects are localized features on the item's surfaces, including cracks, cuts, scratches, dents, pitting, and corrosion, which may indicate damage that compromises function or merely cosmetic wear that does not. Contamination covers foreign matter on the surface, such as dust, dirt, grease, oxidation, or residues, which may obscure the surface beneath and often has to be distinguished from, or removed before, genuine defects can be assessed. Subsurface and structural defects are conditions below the visible surface, such as delamination, internal cracking, fatigue, or voids, which surface imaging alone cannot reveal and

which require volumetric or non-contact internal sensing. Part of the task is deciding which findings actually matter: telling a true defect from a harmless surface mark, and damage severe enough to reject the item from ordinary wear that can be tolerated. The two ways of getting this wrong are not equal. Flagging a harmless mark as a defect wastes processing effort, whereas missing a real structural defect lets an unsound item continue through the process.

A second difficulty is spatial. An item must in general be assessed over the whole of its surface, and often within it, not only on the faces that happen to be presented. Many products have regions that are intrinsically hard to reach, such as internal geometries, cavities, bores, undersides, and concave or hidden surfaces. Defects are just as likely to occur there, yet these regions are much harder to bring within view of a sensor. Achieving complete coverage is therefore partly a sensing problem and partly a manipulation one: the item, the sensor, or both must be moved so that every region of interest is reachable, a requirement that couples the assessment stage directly to the handling capability of the system and, in turn, to the layout and tooling decisions of Section 4.2.

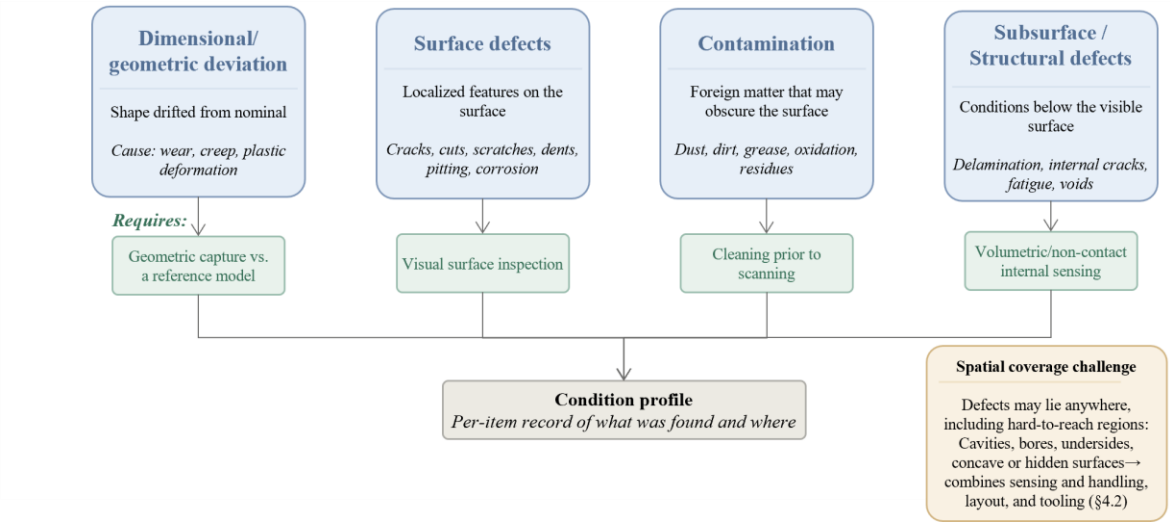


Fig. 9 The classes of condition variability a core may present, the sensing each requires, and the resulting condition profile

Resolving both kinds of variability across the whole of the item requires an inspection capability that behaves less like a conventional presence or quality check and more like an automated diagnostic process. It generally combines several sensing modalities, chosen to cover the defect types above, and turns their raw output into a structured, quantified description of the item. This is a different sensing philosophy from that of a conventional cell, where sensors mostly confirm the presence or position of an already-known part; here the sensing is what defines the item in the first place. Its output is not a pass or fail signal but a condition profile, a per-item record of what was found and where, which becomes the input to every decision the system subsequently makes about that item.

4.1.2 Traceability layer and digital passport

The output of condition assessment is only useful if it is retained, organized, and made available to the rest of the system rather than consumed once and discarded. Immediately after inspection, an item is described by a scattered set of readings, often produced by several independent sensing sources, each capturing a different facet of its condition: dimensional data, defect types and locations, contamination, wear classification, and any variant identification. Left unconsolidated, this has limited value, because no single downstream station holds a complete picture of the item it is about to act on.

Consolidating these readings is the role of the traceability layer, the cell's information backbone. It gathers the readings produced during condition assessment, merges them into a single record for that item, and stores the record for any later stage to retrieve. This per-item record is a runtime-generated counterpart to the engineering data that a known part would carry in a

conventional line; the difference is that it is built on the fly for each item, rather than fixed once at the design phase. Because every processing decision can be traced back to the record that justified it, the layer also makes the process auditable.

The record can hold considerably more than the current cycle strictly requires, and the categories it may span are summarized in Table 5: the item's identity and provenance, the condition profile measured this cycle, the processing history generated within the system, and, where available, the item's history from earlier life cycles.

Table 5 Categories of data that a Re-X traceability record may hold.

Category	Example contents	Origin
Identity and provenance	Product variant, unique item identifier, original specification, material composition	Assessment (variant ID) and external records
Measured condition	Dimensional deviation, defect type, location, and severity, contamination, wear classification	Condition assessment, this cycle
Processing record	Operations performed, parameters used, material removed or added, accept/reject outcome	Processing stages, this cycle
Life-cycle history	Condition and processing records from previous cycles, number of prior cycles	Persistent record across cycles

The last of these categories is what makes the record more than a within-cycle convenience. If an item's earlier condition and processing persist between cycles, assessment in the current cycle

can draw on them, comparing present condition against past rather than starting from nothing, which is directly useful for tracking how a recurring item is degrading over successive lives.

This is also the point at which the system connects to a development beyond it. Under the EU Ecodesign for Sustainable Products Regulation, a Digital Product Passport (DPP) is to become a machine-readable record accompanying each product, holding its material composition, durability and performance data, and end-of-life information. A traceability layer of the kind described here is, functionally, the mechanism by which such a record would be read and written at the point of processing: it can be initialized from an item's existing passport where one is present, and it produces exactly the condition and processing data the passport is intended to carry. Seen this way, the traceability layer is not only an internal convenience but the point at which a Re-X system participates in the item's life-long digital record.

4.1.3 Processing strategy

With a consolidated characterization of the item available from the traceability layer, the system can determine how that specific item should be processed. This is where the variability discussed across this section becomes a concrete consequence for the processing itself rather than an abstract property. Because each item's condition differs, the operations it needs, and their order, cannot be fixed ahead of time in the way hard automation fixes a single repeated sequence. The processing sequence is instead composed at runtime from the item's condition profile, so that each item receives only the operations its own state calls for.

The detailed engineering of any individual process technology is outside the scope of this thesis. It is nonetheless useful to set out the categories of operation a Re-X system can draw on,

since their breadth is itself part of what sets Re-X processing apart. Table 6 organizes these operations into two groups: handling operations, which change the item's position or orientation without altering it, and processing operations, which act on the item itself. The processing operations span the full remanufacturing cycle, from disassembly and cleaning, through inspection and the material-changing operations that recondition the item, to reassembly and testing. It also includes the outright replacement of components, an option available to remanufacturing that has no equivalent in the processing of a single new part.

Table 6 Operation categories available to a Re-X processing cell.

Class	Category	Purpose	Representative operations
Handling	Handling and transport	Move, orient, or hold the item without altering it	Grasping, rotation, indexing, fixturing, transfer between stations
Processing	Disassembly and separation	Break the item down into components, or separate joined parts, for individual processing	Unfastening, cutting, pressing, de-bonding
Processing	Cleaning	Remove contamination, residues, or corrosion ahead of inspection or processing	Washing, degreasing, blasting, ultrasonic and chemical cleaning
Processing	Inspection and sorting	Establish condition and route the item or component accordingly, including rejection	Metrology, defect detection, condition grading, accept/reject sorting
Processing	Subtractive	Remove material to correct geometry or eliminate localized damage	Turning, milling, drilling, grinding, skiving, machining
Processing	Additive and restorative	Add or rebuild material to restore lost geometry	Cladding, weld repair, thermal spray, patching, recoating
Processing	Forming	Restore or adjust shape without removing material	Cold or hot deformation, straightening, re-forming
Processing	Surface finishing	Restore surface texture, tolerance, or cosmetic condition	Vibratory finishing, honing, polishing, blasting

Class	Category	Purpose	Representative operations
Processing	Component replacement	Replace a subcomponent that cannot be restored with a new or serviceable one	Removal and refitting of worn or failed parts
Processing	Reassembly	Rebuild the item from reconditioned and replaced components	Fastening, welding, adhesive bonding, fitting
Processing	Testing and verification	Confirm the processed item meets the required performance or quality standard	Functional testing, dimensional verification, quality validation

Composing a processing sequence, in this sense, is the activity of selecting, ordering, and parameterizing the appropriate combination of operations from Table 6 for the item's condition profile, separately from the resources that carry those operations out. It is this per-item composition, rather than the execution of any one operation, that characterizes a Re-X system: the processing plan is built for each item from its measured condition rather than issued once for a known part. Whether that sequence is composed automatically or with the human in the loop depends on the system's implementation and its operating scenarios (Section 4.2.2). Accordingly, this thesis treats the composition of the processing sequence as a methodological activity the design must support, rather than as something it claims to automate in full.

4.2 REMANUFACTURING WORKCELL ENGINEERING

Section 4.1 established what makes a Re-X system different at the level of input and decision-making. This section addresses how that understanding is turned into an engineered cell. The engineering activities, defining the cell's function, resourcing it, arranging it in space, designing its tooling, and validating it digitally, are presented in sequence, in an order that runs

from function toward physical realization. They are not carried out as a single forward pass: a decision taken in later steps often exposes a constraint on an earlier one, requiring that earlier decision to be revisited.

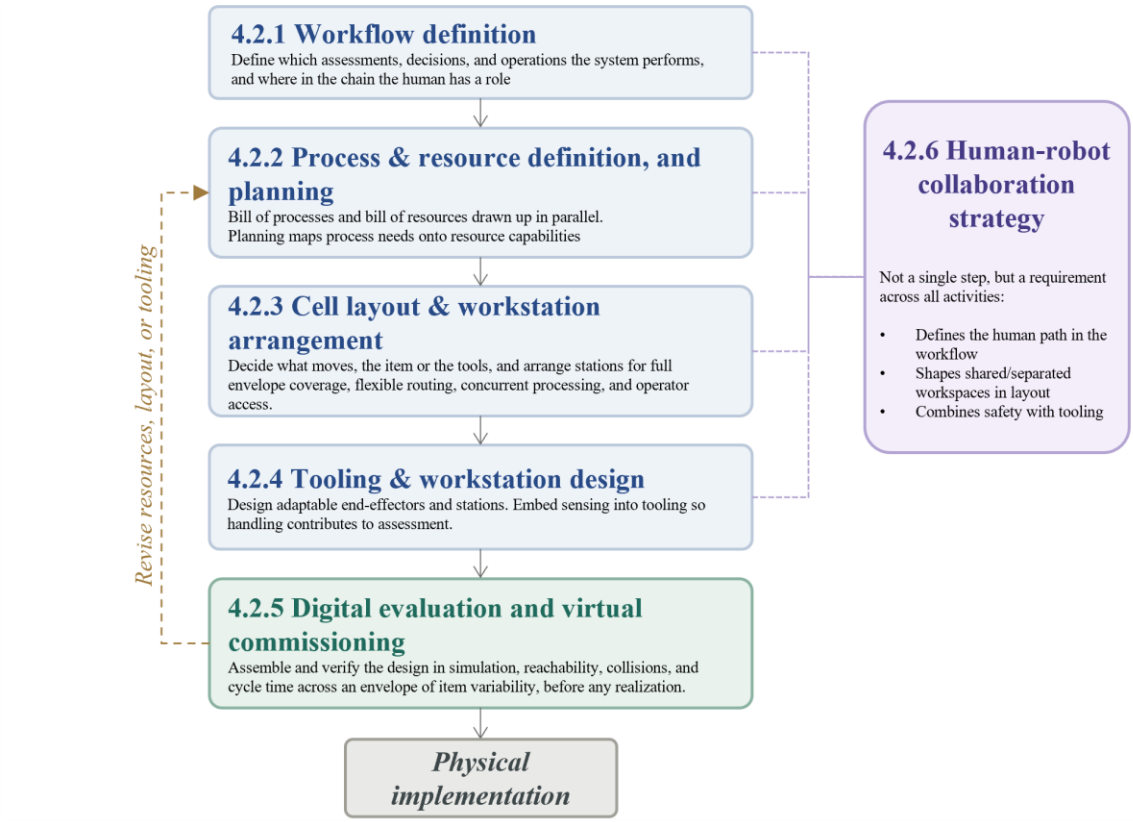


Fig. 10 The Section 4.2 Engineering activities, moving towards physical implementation.

4.2.1 Remanufacturing cell workflow definition

The first activity is to translate the general operational logic into a concrete workflow for the particular system under design. Every station and decision point serves one link of the assessment-to-processing chain, and workflow definition is the exercise of making that chain concrete: which assessments are performed, what decisions follow from them, which operations are available, and where the human enters the loop.

This chain cannot be a strictly linear path. Because each item's condition determines which operations it requires, the workflow takes the form of a branching structure: some items bypass whole groups of operations after assessment, others need repeated passes through a stage before proceeding, and some are rejected outright. What stays constant across all cases is the order of the chain: assessment before decision, decision before execution, with the human in the loop at each. What varies from one item to the next is only which operations are carried out at each step.

The workflow is defined first because everything that follows is derived from it. The operations the system must perform, the resources it needs, and the scenarios it must support are all read off this common structure, which keeps them consistent with one another instead of being specified in isolation and reconciled afterwards. Fixing the intended operation of the system before committing to any hardware is what allows each later design decision to be traced back to a role the workflow already defines.

4.2.2 Process and resource definition, and planning

With the workflow defined, two descriptions are drawn up alongside each other, and reconciled only once both are in place. The first is the bill of processes: the full set of operations the cell must be able to perform, mapped against the range of conditions the incoming items are expected to present, and drawn from the operation categories of Table 6. The second is the bill of resources: the candidate set of resources able to perform such operations, human operators and automated equipment alike, each described by what it is capable of, its working envelope, its accuracy, and the operations it can carry out.

Planning is then the step of mapping one onto the other: matching the specifications of each required process against the capabilities of the available resources, and on that basis allocating each operation to the resource type best suited to it. This allocation is where the division between automated and manual execution is actually decided, according to whether an operation can be reliably performed by automated equipment or is better assigned to an operator given its complexity or the variability it must manage. Presenting it as a mapping between process requirements and resource capabilities, rather than as an ad hoc assignment, keeps the allocation traceable: each decision to automate or not can be referred back to a specific mismatch or match between what an operation demands and what a resource offers. The outcome fixes the number, type, and specification of the robots, tools, and machinery the cell requires, sized to the operations they have been assigned, together with the operations retained for the operator.

Defined alongside the resource allocation are the operating scenarios the cell must support. A scenario is a mode in which the cell can run. This work considers two: a fully automated scenario, in which the cell executes the assessment-to-processing chain without operator involvement, and a human-involved scenario, in which the operator supervises or performs the operations the automated path cannot resolve.

A scenario is kept distinct from the degree of human–robot interaction it involves, because the two describe different things. The active scenario defines who does what. The interaction defines how closely the human and robot are coupled when they act, whether they share a task and workspace, share a task across separated workspaces, or work independently, and is treated as part of the collaboration strategy (Section 4.2.6). Keeping them separate is necessary because one scenario can be realized through several interaction patterns, each with its own layout and safety implications.

4.2.3 Workcell layout and workstation arrangement

Layouting translates the resource list and the operating scenarios into a spatial arrangement: where the assessment, processing, and human stations sit relative to one another, how the item moves between them, and where people and material enter, move, and leave. In a Re-X cell, two properties established in the preceding sections make this more than a conventional layout exercise.

The first is the branching workflow. Because not every item follows the same sequence of stations, the layout must support flexible routing between stations rather than a fixed linear flow, while not becoming so freely connected that material flow turns unmanageable. The arrangement therefore balances the range of paths the workflow requires against the tractability of the resulting flow.

The second concerns what moves. This is a question a Re-X layout must answer explicitly, whereas a conventional single-station typology leaves it implicit. In classical manufacturing-systems terms, a station may be organized so that the workpiece is fixed while resources act on it, or so that the workpiece moves between fixed resources, and cells may be laid out as fixed-position, functional, or flow arrangements accordingly [1]. Re-X processing perturbs this in two ways. First, because the item must often be inspected and processed over its whole surface, including hard-to-reach regions, the layout must resolve whether that coverage is achieved by moving the item relative to fixed sensors and tools, by moving the sensors and tools relative to a held item, or by some combination, a decision with direct consequences for which resource carries the item and how the stations are placed around it. Second, because the same item variability motivates human involvement, the layout must place human and automated stations so that the

interaction patterns of the active scenarios are physically supported, keeping the operator's habitual working area clear of automated stations during normal operation while still allowing intervention where the workflow calls for it.

These two considerations interact. A layout that consolidates handling into a single resource that carries the item to successive fixed stations answers the "what moves" question one way and tends toward a compact, fixed-position-like arrangement; one that distributes the item across a line of stations answers it differently and tends toward a flow arrangement with transport between stations. Which is preferable depends on the item, its size and mass, the coverage its assessment demands, and the degree of parallelism the throughput target requires, since stations that can operate on different items concurrently relieve a bottleneck that a strictly sequential arrangement would impose. The layout activity is the point at which these are weighed together, and its output is the spatial configuration that the tooling and validation activities then work within. A further question sits underneath these, which the classical typology was not framed to answer: whether the circular character of the process, rather than merely the presence of a robot, perturbs the station types themselves, for instance by requiring an arrangement in which a single resource carries the item through inspection and processing rather than fixing it at a station or transferring it along a line. The cell of Section 5.0 is one instance of such an arrangement, and can illustrate what that perturbation looks like in a concrete case; whether it generalizes across circular processes is beyond what a single case can establish, and is returned to as an open direction in Section 7.0 .

4.2.4 Tooling and workstation design principles

Tool design, which here means the design of the robot end-effectors and any dedicated handling equipment, follows from the resource list and the variability the cell must absorb. Because no single item geometry or condition can be assumed, tooling cannot be optimized around one fixed case as it would be in a conventional cell; it must work reliably across a defined envelope of dimensional and condition variation. This favors adaptable or compliant handling mechanisms over rigid single-purpose fixtures.

It also motivates a feature particular to this setting: embedding sensing into the tooling itself, so that handling equipment contributes to condition awareness rather than acting only as a passive interface. This is where tooling design and the circular character of the process meet most directly. In a Re-X cell the item is unknown, so a handling device that measures as it handles it, performs part of the condition assessment of Section 4.1.1 before the item even reaches a dedicated sensing station. A gripper that registers an item's dimension as it secures it, for instance, turns an unavoidable handling action into a first condition measurement. Integrating sensing into tooling is thus not incidental in a circular context but a way of exploiting handling actions the cell must perform anyway to reduce the burden on dedicated assessment, and it is pursued directly in the design of Section 5.0 .

Workstation design follows the same logic. Dedicated stations are built around the operations allocated to them during planning, but remain adjustable enough to accommodate the range of cases the workflow routes to them, rather than being locked to the narrow tolerances a fixed, known part would allow.

4.2.5 Digital evaluation and virtual commissioning

Before any physical resource is committed, the proposed layout, tooling, and workflow are assembled and evaluated in a digital environment, in which the system is built and its operation verified in simulation ahead of physical realization. This digital-twin based evaluation follows an established practice of designing, validating, and virtually commissioning robotic and human–robot collaborative cells in simulation before committing to hardware [32], [33] . The evaluation checks what is hard to establish on paper, the reachability across the full range of item positions, absence of collisions between resources, and cycle-time behavior across the operating scenarios and process paths the cell must support. In a Re-X context, this carries particular weight, because it is the full envelope of item variability, rather than a single nominal case, that tends to expose design gaps, and simulation is where that envelope can be exercised at low cost.

If the evaluation shows that the configuration cannot accommodate that envelope, the relevant earlier decisions, in resource allocation, layout, or tooling, are revisited before proceeding. This is the iterative return noted at the start of Section 4.2: the activities are presented in sequence for clarity, but it is digital evaluation that closes the loop back onto them. Only once the design is validated in this environment does it move on to the construction of physical tools and stations, connecting the engineering of this section to the physical realization in Section 5.0 .

4.2.6 Human-Robot collaboration strategy

Given the variability of the incoming items and the breadth of what the cell must handle, pursuing full automation of every scenario would carry disproportionate cost and complexity, and

would in places exceed what current automated assessment and decision-making can reliably support. The operator therefore remains part of the cell by design, a human-in-the-loop approach [34] considered from the workflow and scenario stage onward rather than added afterwards as a contingency.

The operator's involvement takes different forms depending on the active scenario: resolving cases the automated path cannot settle, confirming or overriding a proposed processing sequence before it is executed, and carrying out operations that remain impractical to automate across the range of conditions involved. The extent of this involvement scales with the operating scenarios of Section 4.2.2, from occasional oversight in a largely automated run to direct execution when the automated path defers.

Where the operator shares a workspace or a task with the robot, the safety strategy follows the standard classification of human-robot interaction by whether the two share a common task and workspace, share a common task across separated workspaces, or work independently within the same cell [7]. Each configuration implies a different combination of safety measures, from safety-rated monitored stops where workspaces are cleanly separated, through speed and separation monitoring, to power and force limiting where robot and operator work in close and potentially continuous proximity. These choices are inseparable from the layout and tooling decisions of Section 4.2, a configuration that overlaps human and robot working areas without providing for it at the layout stage introduces safety-driven inefficiencies that are much harder to resolve afterwards than if the interaction had been designed for from the outset. Because the safety strategy must be determined jointly with layout, tooling, and workflow rather than after them, the collaboration strategy is drawn together here, once the individual activities it touches are in place, rather than being confined to any one of them.

5.0 IMPLEMENTATION

The methodology of Section 4.0 was defined independently of any specific product, structured around the logic that governs a Re-X system: variable input, condition assessment, information retention, processing, and human involvement. This section applies that methodology to a concrete industrial case, the retreading of heavy transport vehicle tires, in order to validate it. Where Section 4.0 sets out the methodology that is the central contribution of this work, this section materializes it: it shows how the general activities are realized as a specific cell design, and in doing so demonstrates that the methodology yields a coherent, workable design when instantiated on a demanding real case.

Within the broader retreading process, this work addresses two of its constituent stages: defect detection and localization, corresponding to the condition assessment activity of the methodology, and skiving, the material-removal operation that prepares a damaged region of the tire for retreading, corresponding to the processing activity. The remaining stages of a complete retreading process, such as buffing of the full tire surface, application of new tread material, and curing, fall outside the scope of this cell.

The cell developed here is therefore a targeted instantiation of the methodology: a single robot, the Comau NJ130-2.6, equipped with a custom Powered Roller Gripper, moves the tire between a perception station, where defects are identified and localized, and a skiving station, where the identified defect is processed. Human involvement is retained by design rather than as a fallback. As established in Section 4.2.6, full automation of every case is neither presently achievable nor economically warranted at this level of process maturity: certain defect presentations cannot yet be reliably processed automatically, and the proposed material removal

cannot always be committed to without human review. The cell is accordingly designed so that an operator can confirm or override the proposed processing, and carry out skiving manually where the automated path defers, with the safety strategy this entails developed in Section 5.5. The sections that follow describe how each of these elements was realized for this specific case, in terms of the concrete hardware, sensing, and human interaction the cell relies on.

Because those descriptions refer repeatedly to specific regions of the tire, both to state where defects may occur and to reason about which surfaces the cell must reach, inspect, and process, the terminology used throughout this section is fixed here. A tire casing is described in terms of its crown, the outer band that carries the tread; its two sidewalls, the flexible walls running from crown to bead; and its two beads, the stiff inner edges that would seat against the wheel rim. Equally important for this work is the distinction between the tire's outer and inner surfaces: defects, and therefore the reach requirements they impose, arise not only on the externally visible crown and sidewalls but also on their inner counterparts, the inside crown and inside sidewalls, which face into the tire's cavity and are considerably harder to bring before a sensor or a tool. These regions are illustrated in Fig. 11 and the terms are used with these meanings from Section 5.1 onward.

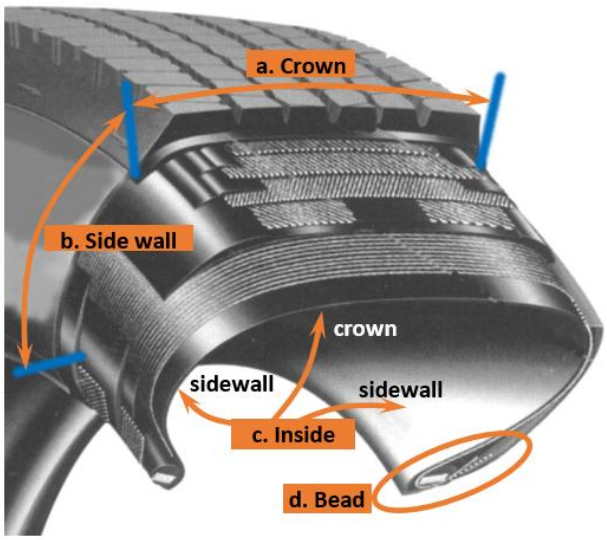


Fig. 11 Tire anatomy and terminology used throughout this work: crown, sidewalls, bead, inside crown, and inside sidewalls.

5.1 ROBOTIC CELL LAYOUTING

The methodology requires a workflow structured around condition assessment, information retention, and processing, executed within a layout that supports both automated and human-involved operation. For the tire case, this becomes a cell built around two automated operations, perception and skiving, with the operator able to intervene where either cannot be resolved automatically. Perception depends on stable, repeatable positioning relative to a fixed sensor arrangement; skiving depends on accurate execution against the specific geometry and location of each defect, using the tool described in Section 5.4. The layout described here supports both operations, together with the exception path for human intervention, within a single compact cell built around one robot.

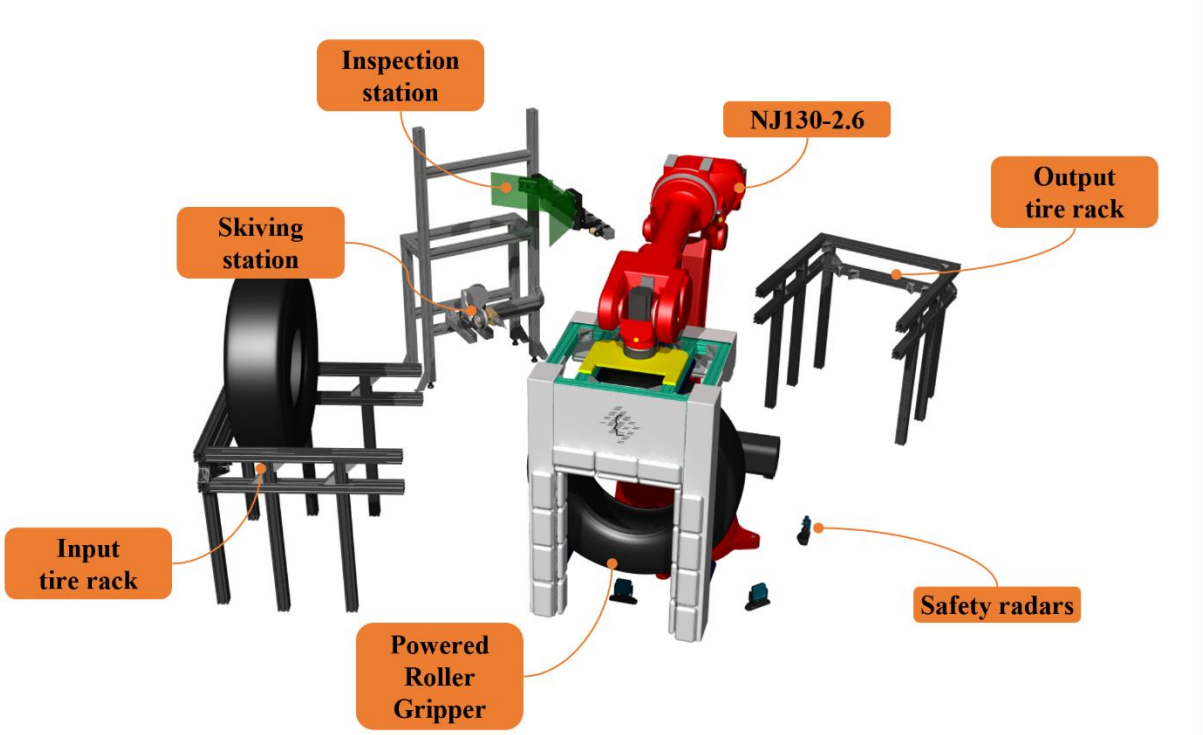


Fig. 12 Cell layout and the main components of the designed robotic cell.

The choice of robot is driven by two combined requirements. First, the manipulator must carry, for the full duration of each cycle, the combined weight of the tire and the gripper holding it. Second, as developed in Section 5.2, the decision to have the robot itself hold and continuously reorient the tire, rather than presenting it to a stationary fixture, requires sufficient reach to bring the tire to both stations from a single mounting position, and sufficient repeatability to preserve the positional accuracy that both defect localization and skiving execution depend on. Together these place the robot in a higher payload class than is typical of human-robot collaborative applications. The Comau NJ130-2.6 was selected on this basis: a 130 kg payload-class, six-axis manipulator with a maximum horizontal reach of approximately 2.98 m and a repeatability of 0.07 mm, meeting the combined mass of gripper and tire while retaining the positional accuracy required at both stations, and doing so without a second robot introduced purely to resolve reach or reorientation, an alternative examined in Section 5.2.

The robot's working envelope is shared between two vertically arranged stations: the perception station in the upper region of the workspace, and the skiving station below it. Stacking the stations vertically within one reach envelope, rather than distributing them horizontally, keeps the floor footprint compact and avoids long horizontal robot travel, or a second robot or transport axis, between the two operations.

From the point the tire enters the cell it remains held by the robot throughout, via the Powered Roller Gripper. The robot first presents the tire to the perception station, where detection and localization are carried out (Section 5.3), then transfers it directly to the skiving station below (Section 5.4). No intermediate transport mechanism, conveyor, or turntable is required, since the tire is transferred and reoriented entirely through the robot's own motion and the gripper's rotational axis.

The skiving station sits within the same reach envelope as the perception station while preserving a dedicated zone for the operator, used only when automated skiving defers to manual intervention. The two automated operations, perception and skiving, are grouped together on one side of the cell, while the operator's zone is kept apart from them, outside their footprint during normal operation. The associated safety measures and monitored zones are addressed in Section 5.5.

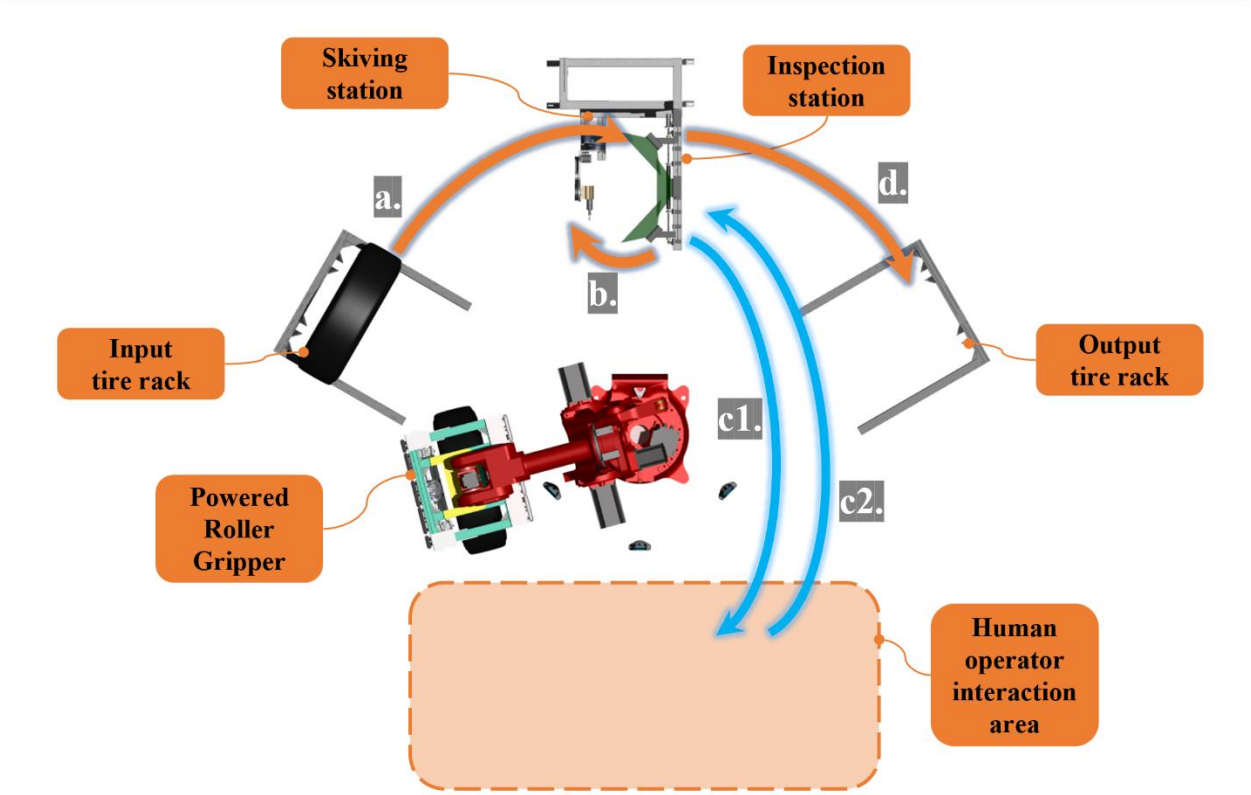


Fig. 13 Developed robotic cell with the material flow a) Tire transfer from input to inspection station, b) Transfer to skiving station, c1) Transfer to interaction area for manual processing, c2) Transfer back to skiving station, d) Transfer to output rack after completing process.

5.2 END-EFFECTOR

Handling a heavy transport vehicle tire across two functionally different stations, stable positioning for perception and precise presentation for skiving, ruled out a conventional single-purpose end-effector and shaped the design of the cell before the gripper itself was finalized. This subsection first sets out the specifications the gripper must meet, then describes the core mechanisms through which each is achieved.

5.2.1 Design rationale: why an active end-effector

An early concept placed rotation at the perception station itself: a stationary roller-based turntable that would rotate the tire in place, supported by the turntable rather than held by the robot. This met the rotation requirement for perception but exposed a reachability problem at skiving: with the tire resting on the turntable rather than held and freely reorientable by the robot, presenting an arbitrary defect location, anywhere on the outer or inner surface, to the skiving tool would demand either a very high number of degrees of freedom at the tool, or a second dedicated robot to carry and position the skiving head around the tire. Beyond reachability, a second robot would bring its own controller, safety zone, and the coordination of two robots working on one part would increase the complexity of the task.

This led to the approach taken here. Rather than leaving the tire's rotation to a fixture and the reach to every defect to a second robot, both capabilities are consolidated into a single active end-effector, the Powered Roller Gripper (PRG), carried on one high-payload manipulator able to both hold the tire and freely reorient it. Because the robot holds the tire and can continuously reorient it, any location on the tire — inner or outer surface — can be presented to either station through the robot's own six-axis motion, with no fixture at perception and no second robot at skiving.

5.2.2 Gripper specifications

The requirements the gripper must satisfy, derived from the tire, the two stations, and the operator role, are collected in Table 7. They fall into grasping, rotation, sensing, and safety groups, and each is traced in Section 5.2.3 to the mechanism that provides it.

Table 7 Specifications for the Powered Roller Gripper.

#	Group	Specification	Target / value
G1	Grasping	Secure a buffed 315/80 R22.5 casing across its diameter tolerance	Ø 950-1200 mm range
G2	Grasping	Maintain grasp for the full cycle across both stations	Fail-safe under power loss
G3	Grasping	Apply and hold a controlled clamping force without continuous motor torque	1.6 kN
G4	Grasping	Provide open approach clearance for loading around the tire	200 mm clearance
R1	Rotation	Rotate the secured tire for scanning and skiving presentation	10 rpm
R2	Rotation	Provide fine, controllable angular resolution	Effective reduction ≈ 20:1
R3	Rotation	Monitor angular position continuously throughout the cycle	Encoder on rotation drive
S1	Sensing	Measure the tire's effective diameter at pick-up	Via securing-travel encoder
SF1	Safety	Retain grip under power or drive failure	Self-locking, non-back-drivable
SF2	Safety	Keep the operator's hands clear of the roller nip during rotation	Two-hand control (Section 5.5)
SF3	Safety	Enclosed drivetrain components not requiring operator access	Fixed guards (Section 5.5)

5.2.3 Main mechanisms

The PRG is built around two functionally independent mechanisms sharing one frame: a securing mechanism that grasps and holds the tire, and a rotation mechanism that turns it once secured. The tire is held between four rollers in two pairs, two lower and two upper, replicating the grasping principle of a conventional four-finger gripper.

Securing mechanism (G1–G4, SF1, S1). The lower roller pair is passive, mounted on a moving assembly that lowers out of the way so the gripper can be positioned around the tire, then raises to bring the tire into contact with the upper rollers and complete the four-point grasp. This vertical motion is driven by a servomotor with an integrated encoder, through a chain drive to a lead screw with a trapezoidal thread. The trapezoidal thread is chosen for its self-locking behavior: because the mechanism is not back drivable under axial load, the tire stays clamped even under power or drive failure, without the motor maintaining holding torque **(SF1)**. The same servomotor's encoder measures the assembly's vertical travel directly; since its geometry relative to the upper rollers is fixed and known, the travel to secure the tire yields the tire's effective diameter at pick-up **(S1)**. This matters because a buffed casing's diameter is not the catalogue value, so the PRG performs the first measurement contributing to the item's condition profile, consistent with Section 4.1.1, before the tire reaches the perception station.

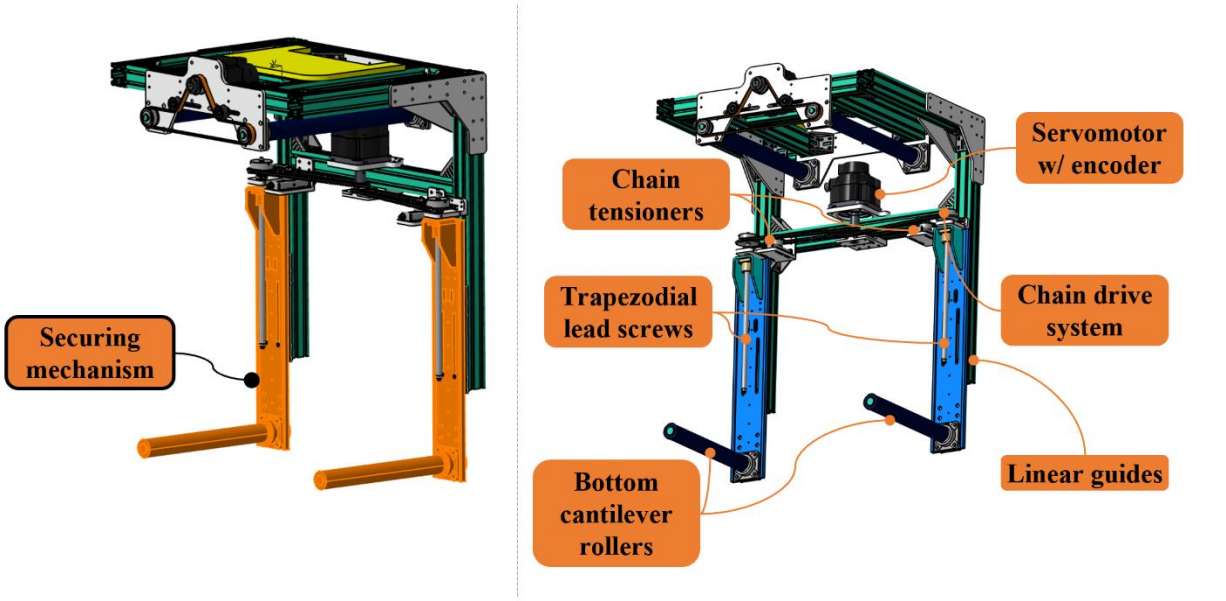


Fig. 14 Securing mechanism and the main components.

Rotation mechanism (R1–R3). Once secured, rotation is provided by the upper roller pair, driven by a second servomotor with encoder, independent of the securing drive. Torque is transmitted to the roller shaft through a 1:1 belt drive, while the effective reduction between roller and tire surface is far higher, approximately 20:1, from the roughly 50 mm rollers driving a tire on the order of 1 m in diameter **(R2)**. This gives fine, controllable rotation despite the compact rollers. Rotation is continuous rather than indexed **(R1),** so the tire turns smoothly under the fixed cameras during perception and under the skiving tool during processing. The encoder on the rotation drive gives continuous knowledge of the tire's angular position relative to the gripper frame throughout the cycle **(R3).** This angular information is what lets a defect found during perception be relocated at skiving without re-detection: the encoder reading taken at detection is kept as a reference, and the tire is rotated back to that angle for skiving. This closes the link to the traceability layer of Section 4.1.2 for this specific implementation.

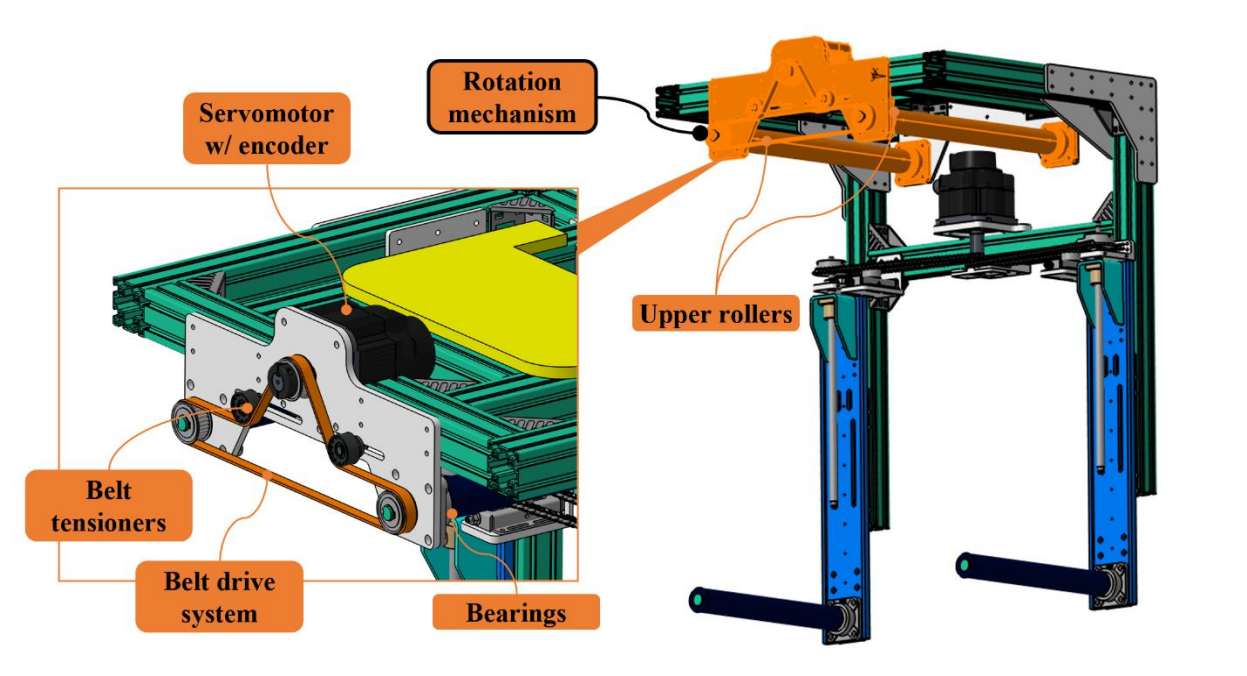


Fig. 15 Rotation mechanism and the main components.

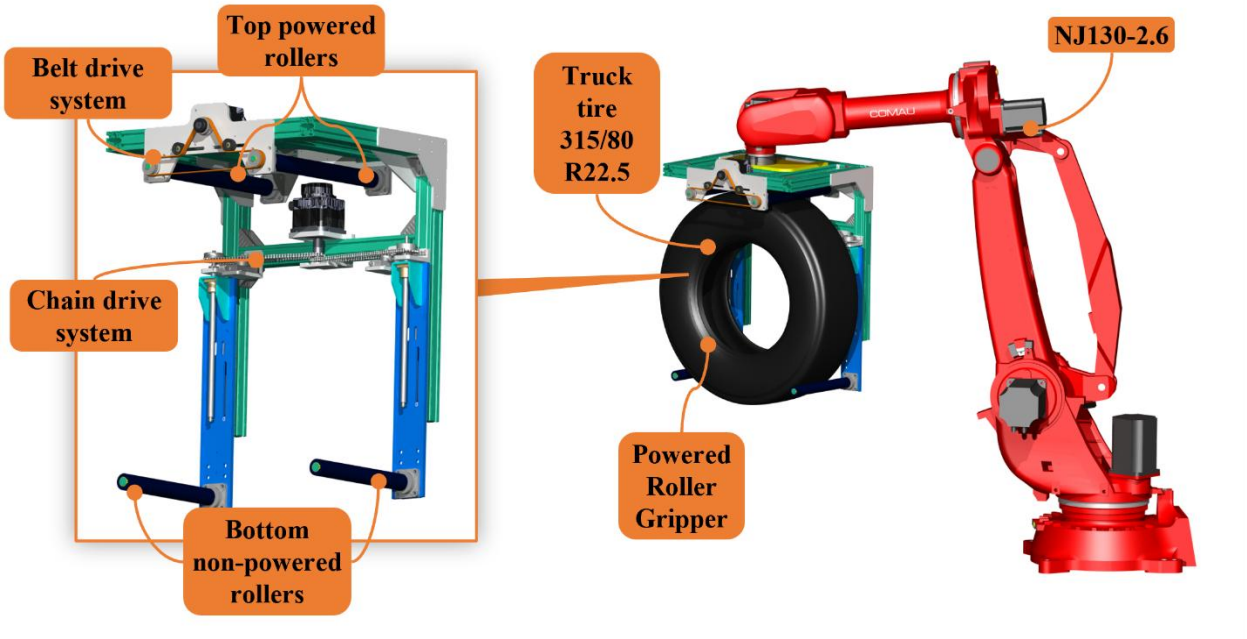


Fig. 16 Powered Roller Gripper mechanisms and the gripper mounted on the NJ130-2.6 robotic arm.

5.3 INSPECTION STATION

Selecting the sensing approach required weighing measurement accuracy against the geometric variability of the surface being inspected. A laser profilometer was considered first, offering high-accuracy surface measurement well suited in principle to detecting the localized material loss a defect represents. That accuracy, however, is most reliable over flat or gently curved surfaces; profilometer measurement becomes less straightforward across a surface that varies as sharply in curvature as a tire does from crown to sidewall, where a fixed scanning geometry would need adapting, or the sensor repositioning considerably, between regions. This limited its suitability for a station required to cover the tire's full outer and inner surface rather than the crown alone.

This motivated stereo depth cameras instead. They support image-based detection and localization, capturing a whole region of the surface in a single view rather than scanning it point by point; this lets them cover curved and variable geometry without reconfiguration between regions, at relatively low cost. The ZED X Mini units selected offer lower absolute depth accuracy than a profilometer, but this is acceptable because a defect here is a localized cavity, missing material relative to the surrounding surface. Depth information alone reveals its presence and extent without the sub-millimeter accuracy a profilometer would provide; detailed classification of a detected defect from that depth information is outside the scope of this thesis. Three cameras are arranged around the tire, as below, to extend coverage across the full surface rather than relying on a single viewpoint or a scanning motion.

The perception station uses three stationary-mounted ZED X Mini cameras: one central, facing the tire directly, and two symmetric to its left and right, angled inward toward the tire's axis

so each captures an oblique viewpoint complementary to the central one. Each has an integrated LED strip along its top and bottom edges for consistent, localized illumination regardless of ambient light.

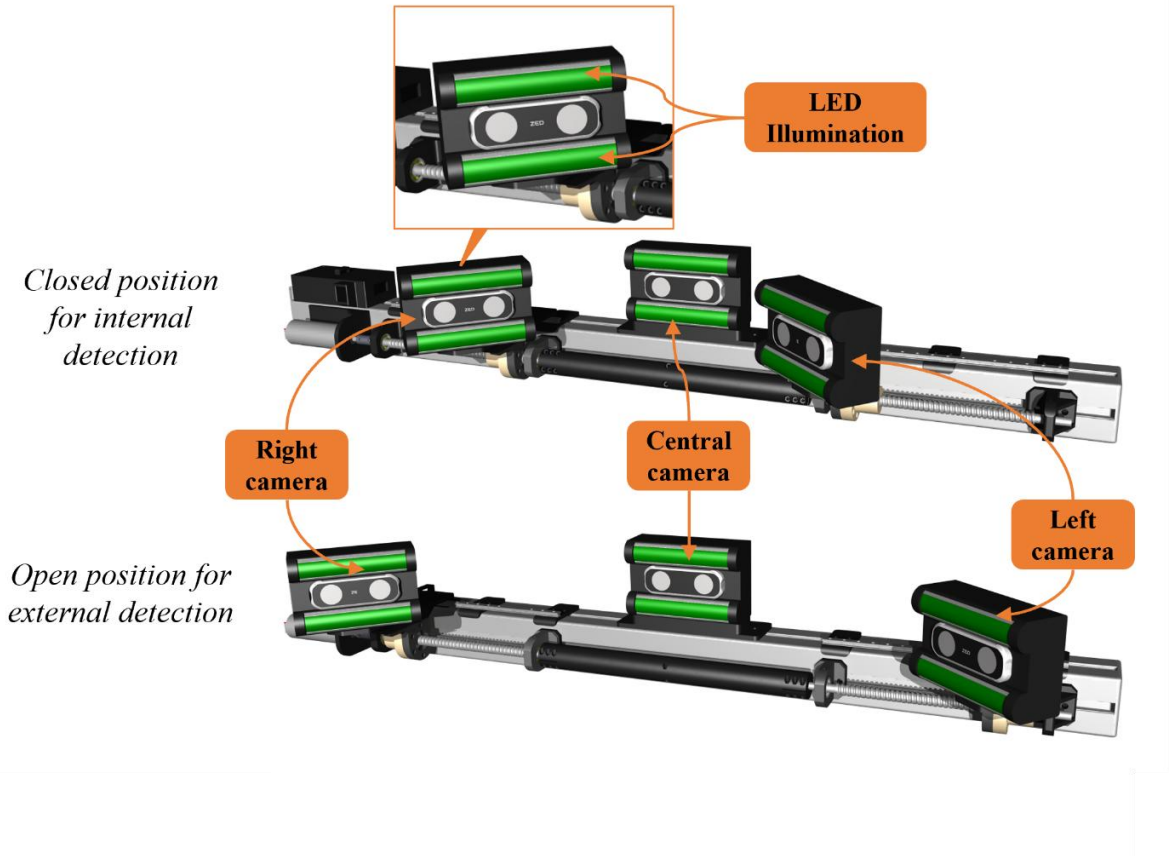


Fig. 17 Inspection station, ZED X mini arrangement and LED illumination.

Detection is carried out in two steps, one per surface. The first covers the outer surface, the crown and both sidewalls; the second covers the inner surface, the bead, inside crown, and inside sidewalls. Because the two surfaces need substantially different viewpoints, the two side cameras are mounted on an opposing lead-screw mechanism, driven by a single servomotor with integrated encoder, that repositions both simultaneously between two fixed positions, one for outer-surface detection and one for inner. The central camera stays fixed, while the tire is rotated continuously

by the PRG (Section 5.2), so each configuration progressively covers the full circumference of its surface before the side cameras move for the second step.

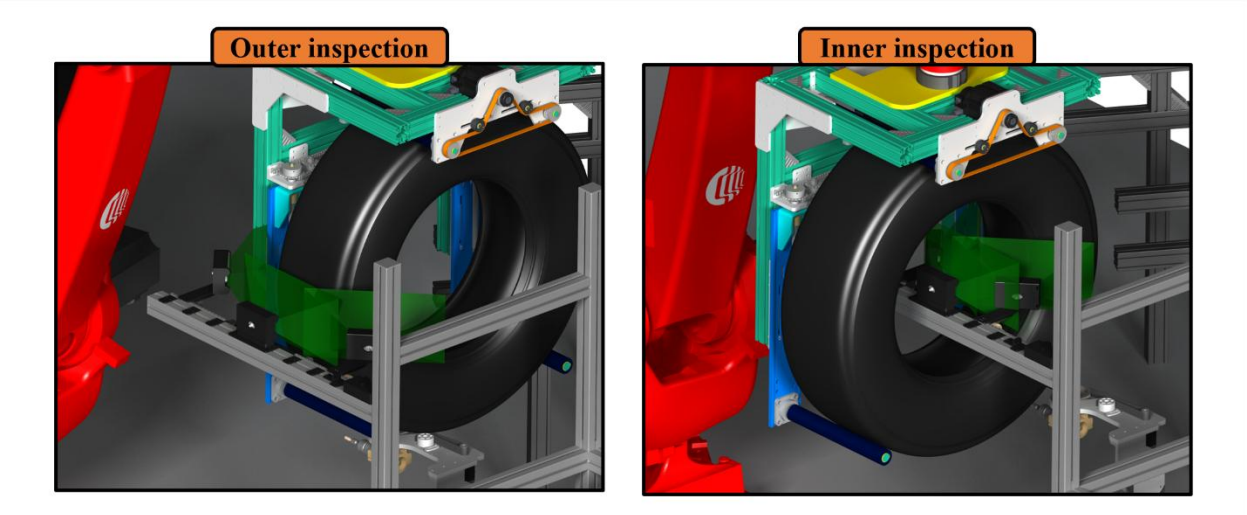


Fig. 18 The two-step detection sequence: Outer-surface configuration and inner-surface configuration.

Establishing a detected defect's location combines several independently known quantities. The cameras are extrinsically calibrated to the robot, so each camera's position relative to the robot frame is known, together with the robot's joint state at detection. This is combined with the tire's angular position, from the gripper's rotation encoder, and its diameter, measured at pick-up (Section 5.2). Together, these express a detected defect in cylindrical coordinates relative to the tire's own center: a radius, an angular position referenced to the encoder-tracked rotation at detection, and an axial position along the tire's width, locating the defect fully within the tire's geometry rather than only around its circumference.

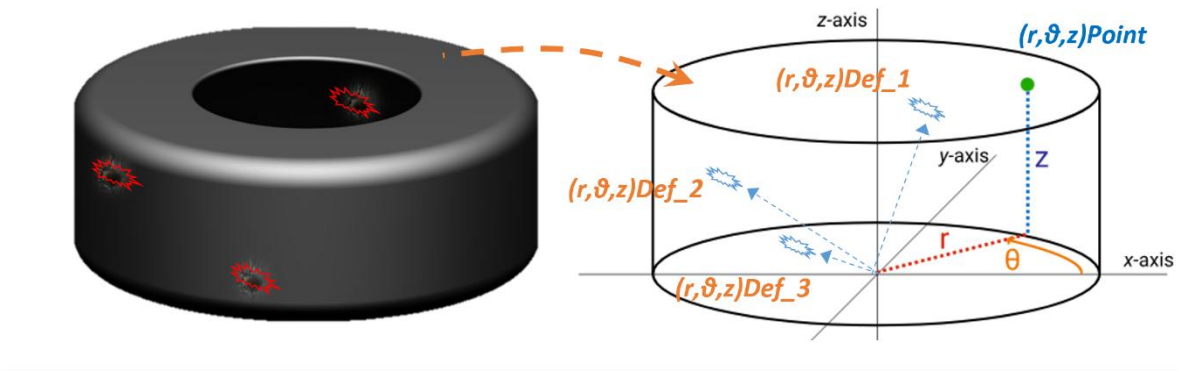


Fig. 19 Cylindrical coordinate representation used to localize a detected defect relative to the tire's center.

5.4 SKIVING STATION

Both stations share the same principle within the cell: neither the cameras nor the skiving tool moves to the tire; both stay stationary while the robot, through the PRG, brings the tire to each in turn. Once a defect's cylindrical position is established at perception, the robot presents the tire to the skiving station within the same reach envelope, and the tire is rotated back to the encoder-referenced angle at which the defect was detected, so that skiving acts on the located defect without re-detection.

The skiving tool is a custom two-axis assembly carrying a rotary cutting spindle. The two axes are revolute, each with a full 360° range, and together they orient the spindle so that it can approach a localized defect on any region of the tire, the crown, either sidewall, or the corresponding inner surfaces, rather than from a single fixed direction. The cutting rotation of the spindle itself is not counted among these degrees of freedom, being the process motion rather than a positioning one. The spindle is an off-the-shelf item: a high-speed pneumatic rotary spindle of the kind used for automated surface finishing and material removal, fitted with a spherical (ball-

nose) tungsten-carbide burr, and able to maintain constant contact pressure against a varying surface through radial deflection.

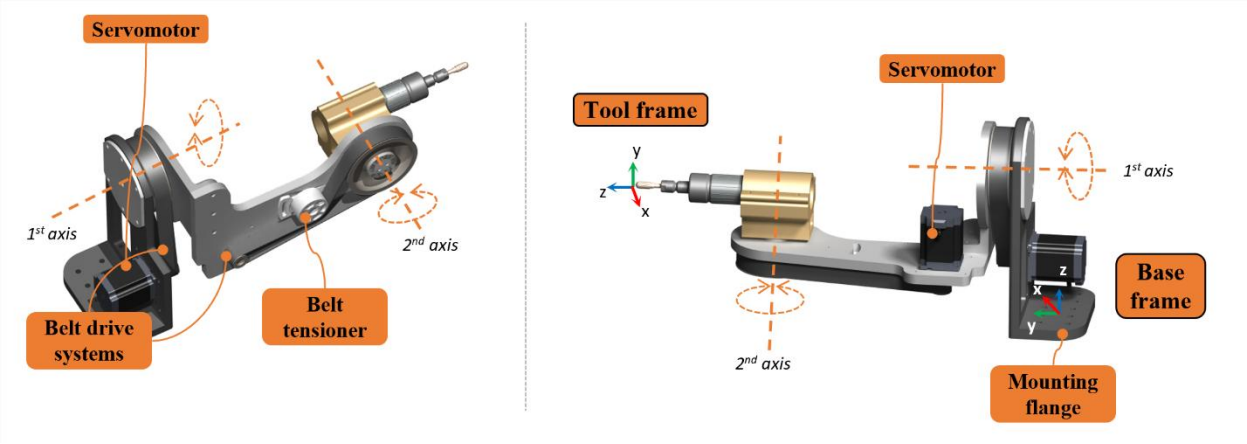


Fig. 20 Skiving tool main components and coordinate frames of the base and the TCP.

The tool is given two controllable axes of its own, rather than being a fixed spindle that the robot presents the tire to, for two reasons. The first is reachability within a constrained envelope. Once the tire and the Powered Roller Gripper are attached to the robot, the volume they jointly occupy, together with the robot's own joint limits, leaves regions of the tire surface that a single fixed tool cannot reach. A spindle mounted at a fixed pose would have to sit within this already-restricted envelope, and from there could not present itself to every defect location across both the outer and inner surfaces. Two controllable tool axes remove this restriction, letting the spindle reorient to the surface at each defect regardless of where in the envelope the defect is held.

The second concerns the dynamics of the cutting motion itself. Skiving follows a fine trajectory over a small region, and the accuracy of that trajectory depends on how precisely the moving part can be controlled through it. Were the tool fixed and the tire moved against it, that trajectory would have to be produced by articulating the robot while it carries the full combined mass of the tire and the gripper. This means a large inertia, held at a substantial distance from the robot's axes, being asked to trace a small, exacting path. Moving so large a mass along a fine

skiving trajectory is considerably harder to do accurately than moving the tool. Placing the fine, trajectory-generating motion on the light two-axis tool instead, while the heavy tire is held at its indexed orientation, keeps the most demanding part of the motion on the low-inertia element. The robot is then left to handle the coarser positioning of the spindle around the tire. This is reflected in how the degrees of freedom are used during skiving, described below. The gripper's rotation is stopped, and the cut is executed by the robot and tool axes together, with the tool carrying the exacting part of the path.

During skiving the degrees of freedom are used differently from during perception. At the perception station the gripper rotates the tire continuously in view of the fixed cameras. At the skiving station the gripper instead indexes the tire to the angular position at which the defect was localized, using the encoder reference captured at detection, and then that rotation axis is held stopped, so the defect is fixed in space for the duration of the cut. The skiving motion is then executed by the robot's six axes and the tool's two axes, actively jogged together to bring the spindle to the defect and work it, while the tire remains stationary. The active kinematic chain during cutting is therefore the robot's six axes plus the tool's two, with the tire held at a fixed orientation, in contrast to the continuously rotating presentation used for inspection.

Regarding its kinematics, because the tool carries two actuated joints, what they resolve is not a full pose but the orientation of the cutting axis, the robot supplying the remaining position and the roll about it. With the base frame of Fig. 20 attached to the mounting flange and its z axis along the first tool axis, a Denavit–Hartenberg description of the two-link chain gives the cutting axis as a unit vector whose polar angle from z is θ_2 and whose azimuth about it is θ_1, so that the two axes together constitute a spherical pointing mechanism. The inverse kinematics is accordingly posed on that direction rather than on a pose. The direction required at the defect is

the inward surface normal of the casing, which perception establishes in the tire's own frame. Describing that same direction in the coordinates of the base gives the components $n_x\ n_y\ n_z$, from which the two joint angles follow in closed form as Eq. (5.1). The sign selects between two admissible solution branches, both reachable given the full 360° range of each joint and available to collision avoidance; the solution contains none of the link offsets, and it degenerates only at $\theta_2 = 0$ and $\theta_2 = \pi$, where the spindle becomes collinear with the first axis and the first joint ceases to influence the approach direction.

$$\theta_2 = \text{atan2}\left(\pm\sqrt{{n_x}^2 + {n_y}^2}, n_z\right), \qquad \theta_1 = \text{atan2}\left(n_y, n_x\right) \tag{5.1}$$

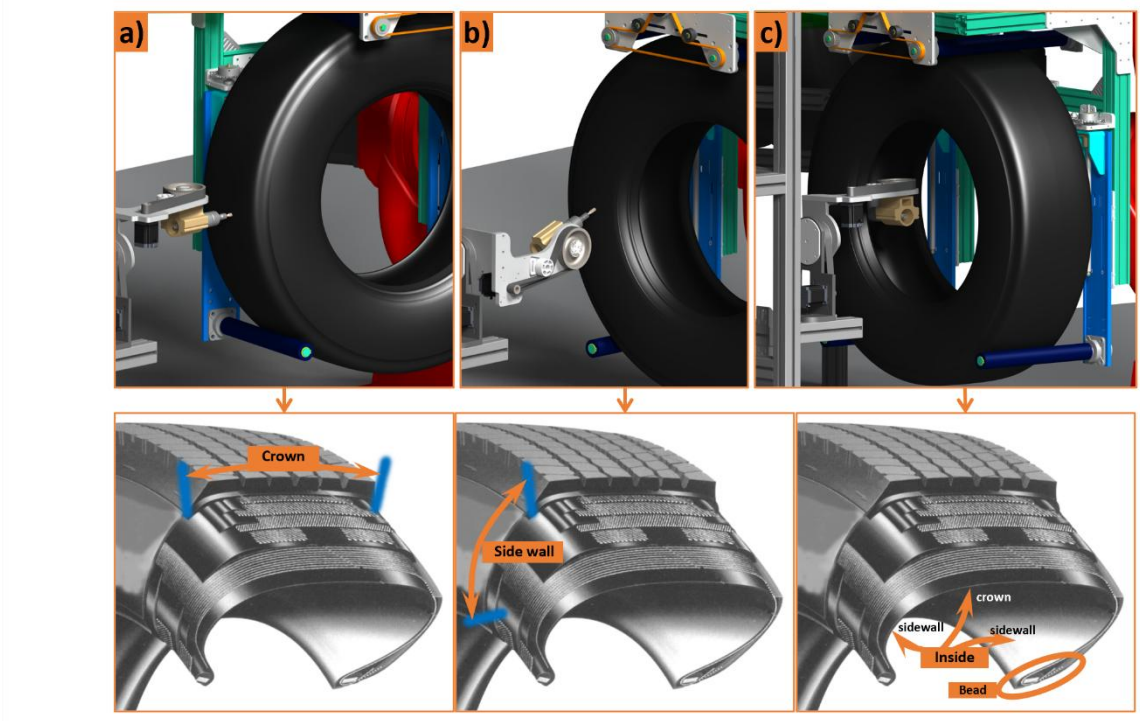


Fig. 21 Skiving tool reaching outer and inner surfaces a) Crown area, b) Sidewall, c) Inner surfaces

One reach limitation follows from the physical size of the mechanism. For defects on the inner sidewall or inner bead located on the same side as the tool, the bodies of the tool's links are

too bulky relative to the tire's internal dimensions to bring the spindle to the defect directly. This is resolved operationally by rotating the tire 180° about the robot's sixth axis, presenting the same defect on the opposite side, where the tool can reach it without interference. A mechanically slimmer link geometry that would remove the need for this reorientation is identified as a target for the next design iteration.

The scope of this work is the mechanical and kinematic design of the tool and the demonstration that, together with the robot and gripper, it can reach an arbitrary localized defect across the tire's surface. The generation of the skiving path and the control of the cutting process are outside this scope. The tool is designed to be presented with a reachable, localized defect; this section establishes the reachability of that defect, rather than the removal operation performed on it.

5.5 HUMAN-SYSTEM INTERACTION, OPERATOR SUPPORT, AND SAFETY STRATEGY

The operator works in close physical proximity to the Powered Roller Gripper during manual skiving, an end-effector that actively rotates a heavy transport vehicle tire through powered rollers while continuing to bear its full weight. Accordingly, the safety assessment for this cell has to consider the complete application, robot, gripper, rollers, tire, drive components, and skiving operation, rather than relying on the collaborative certification of the robot arm alone. This follows the scope of ISO 10218-2, which addresses the integration of industrial robots with end-effectors and other equipment into complete applications and cells, together with the general risk-assessment and risk-reduction methodology of ISO 12100.

The robot selected for this cell, the Comau NJ130-2.6 (Section 5.1), is a standard industrial manipulator, not certified as collaborative. Its safety strategy therefore cannot rely on inherent power- or force-limiting behavior at the joint or drive level. Contact between the operator and the robot while it is in motion is accordingly not treated as an outcome to be mitigated after the fact, but as a condition to be prevented entirely, through zone separation and a monitored, non-moving robot state whenever the operator is present. The strategy is structured as a layered architecture around this principle: mechanical hazard elimination and enclosure, a monitored robot state during close manual work, safety-rated presence and separation sensing, and a supplementary tactile contact layer addressing the PRG's exposed structure rather than the robot arm itself.

At the mechanical level, all drivetrain components of the PRG that do not require direct operator access, the securing mechanism's chain drive and trapezoidal lead screw, the rotation drive's belt transmission, the roller shafts, and associated motor outputs, are enclosed within fixed guards, consistent with ISO 14120 and the safety-distance requirements of ISO 13857. Only the minimum roller surface necessary to contact the tire remains exposed. This distinction matters at the roller-tire interface: ISO 13854, governing minimum gaps to avoid crushing, explicitly does not address drawing-in or shearing hazards, so the roller nip requires direct physical guarding rather than clearance alone.



Fig. 22 PGR safety enclosures and AIRSKINs covering the drivetrain components.

The lower roller pair is an exception. Supported in a cantilevered arrangement from the moving securing assembly, it cannot be fully enclosed without obstructing the open approach path required for loading; only the base and bearings of each lower roller are guarded, and the roller body remains exposed. This is addressed through the two-hand control below rather than additional enclosure, consistent with the intended use of ISO 13851 for hazard zones that cannot be adequately protected by a fixed guard.

Presence and separation monitoring is provided by an Inxpect safety radar system, integrated as a volumetric 3D detection layer around the skiving station's interaction area. Radar was chosen over an optical safety scanner for its immunity to airborne dust, debris, and particulate, conditions the skiving process itself generates, which would compromise an optical system here.

The system is certified to PL d under ISO 13849-1 and SIL 2 under IEC 62061. It implements a two-zone concept: an outer zone, in which operator approach transitions the cell toward a reduced, safe operating state, and an inner zone, in which operator entry brings the robot's main axes to a monitored standstill, the term used in ISO 10218:2025 for a supervised, non-hazardous stationary state distinct from an unsupervised stop, from which resumption is only permitted once the zone is confirmed clear.

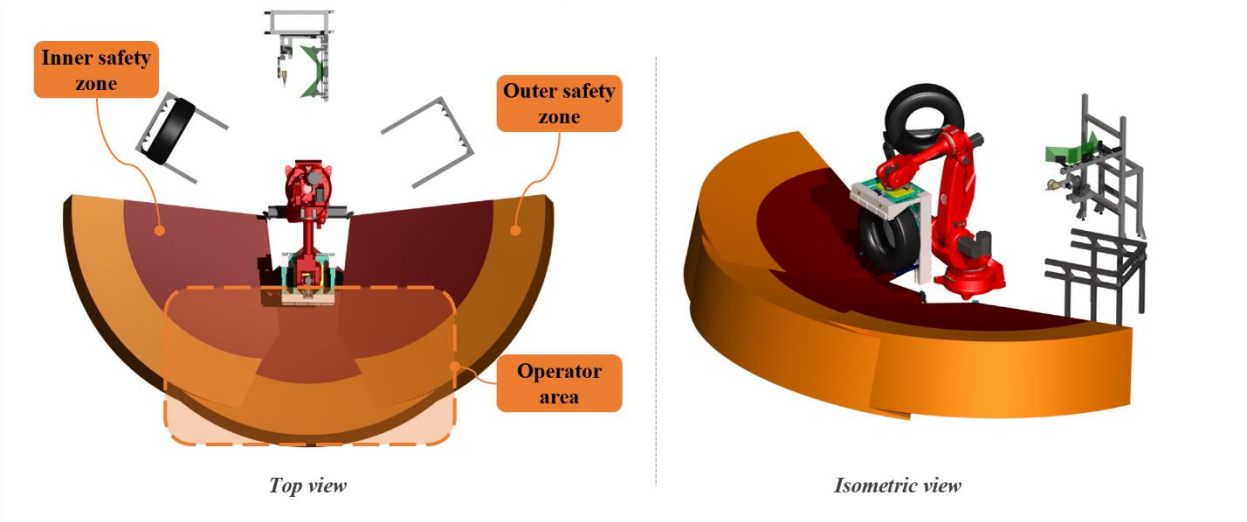


Fig. 23 Radar-monitored zone configuration.

Within this framework, two distinct forms of tire motion are permitted while the operator is present, and are treated as separate safety functions. The first is indexing rotation, used to bring the next defect into an orientation optimal for the operator to see and manually process, rather than relative to any automated tool. This is carried out through the PRG's rotation drive alone, with the main robot axes in monitored standstill, commanded exclusively through two dedicated buttons on the gripper, positioned to keep both of the operator's hands clear of the roller-tire interface for the duration of actuation. Requiring simultaneous, sustained pressure on both buttons, consistent with ISO 13851 for two-hand control devices, ensures rotation is possible only while neither hand

can reach the moving interface, and releasing either button halts rotation immediately. This same control governs the exposed lower rollers: rotation, the condition under which their exposed surface presents a drawing-in hazard, is possible only while both buttons are actuated, so the operator's hands cannot be at the roller interface during the one state in which it is hazardous. Outside active rotation the exposed rollers present a static rather than a dynamic surface. Because this measure protects the operator engaging the controls, the radar-monitored inner zone is specified to permit roller rotation only when no person other than that operator is present, so the two-hand control is not relied on to protect anyone besides its own operator.

The second permitted motion is repositioning, used when the tire must be moved by the robot itself into a more ergonomically favorable orientation, rather than indexed by the rollers. Because this is robot motion, it is not permitted while the operator remains in the inner zone: the operator must withdraw to the outer zone, or leave the cell, before the radar-supervised system allows the repositioning, after which the robot returns to monitored standstill once the operator re-enters.

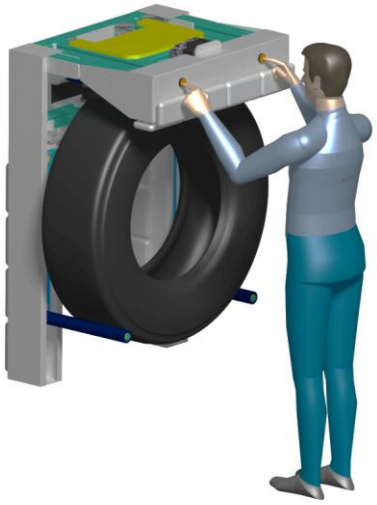


Fig. 24 Operator indexing the tire rotation using the two-hand control.

Grip retention during all these states is addressed inherently by the PRG's mechanical design rather than a separate safety function. As in Section 5.2, the securing mechanism's trapezoidal lead screw is self-locking and not back-drivable under axial load, so the tire stays clamped even under power or drive failure. This satisfies ISO 14118's requirement to prevent hazardous movement from loss of electrical or stored energy, particularly relevant since the tire remains suspended for the full cycle.

As a supplementary contact layer, a certified tactile skin such as AIRSKIN, rated Category 3 / PL e under ISO 13849-1, is considered for the PRG's exposed external frame and structural surfaces the operator may incidentally contact during skiving. It is not applied over the roller-tire contact area, since a pressure-sensitive surface there would be compromised by the grip force the mechanism applies; that interface remains addressed through mechanical guarding and two-hand control.

The specific performance level required for each safety function, radar-based zone monitoring, monitored-standstill transitions, two-hand roller control, and the tactile layer, would

need to be established through a formal risk assessment under ISO 12100 and the safety-function methodology of ISO 13849-1, rather than assigned from component datasheets. The architecture here is accordingly proposed as a conservative design baseline, built around Category 3 / PL d-capable components throughout, consistent with the level typically required for comparable heavy-payload industrial applications.

5.6 VIRTUAL VERIFICATION

Following the digital-evaluation activity of Section 4.2.5, before any component of the cell is committed to physical construction, the complete concept, robot, PRG, perception station, and skiving station is assembled and evaluated within a digital environment, using CATIA V5 for component design and DELMIA for cell-level simulation and validation.

Each component is first validated individually within CATIA before integration into the wider cell simulation. For the perception station, this involves designing the camera frustums for the three ZED X Mini units, so their fields of view can be visualized directly against the tire geometry, and both the camera positions and the required side-camera stroke can be determined and optimized for full coverage of the outer and inner surfaces (Section 5.3). For the PRG, the same approach is applied to its kinematics, so the securing and rotation mechanisms can be tested, validated, and iteratively redesigned where the initial configuration did not perform as intended.

With the components validated, the full cell is modelled in DELMIA to evaluate the arrangement and interaction of all components together within a shared digital environment. This allows the complete operational sequence to be simulated end-to-end: the robot picking up the tire,

transferring it to perception, presenting it to skiving, and returning it, before any equivalent sequence exists physically. The end-to-end simulation confirms reachability across the shared envelope, verifies that no collisions occur between robot, gripper, and station components throughout the motion sequence, and validates the spatial arrangement of Section 5.1.

This iterative use of simulation, at component level in CATIA and cell level in DELMIA, allows design issues to be identified and corrected digitally, consistent with the closed-loop relationship between digital evaluation and earlier design decisions in Section 4.2.5.

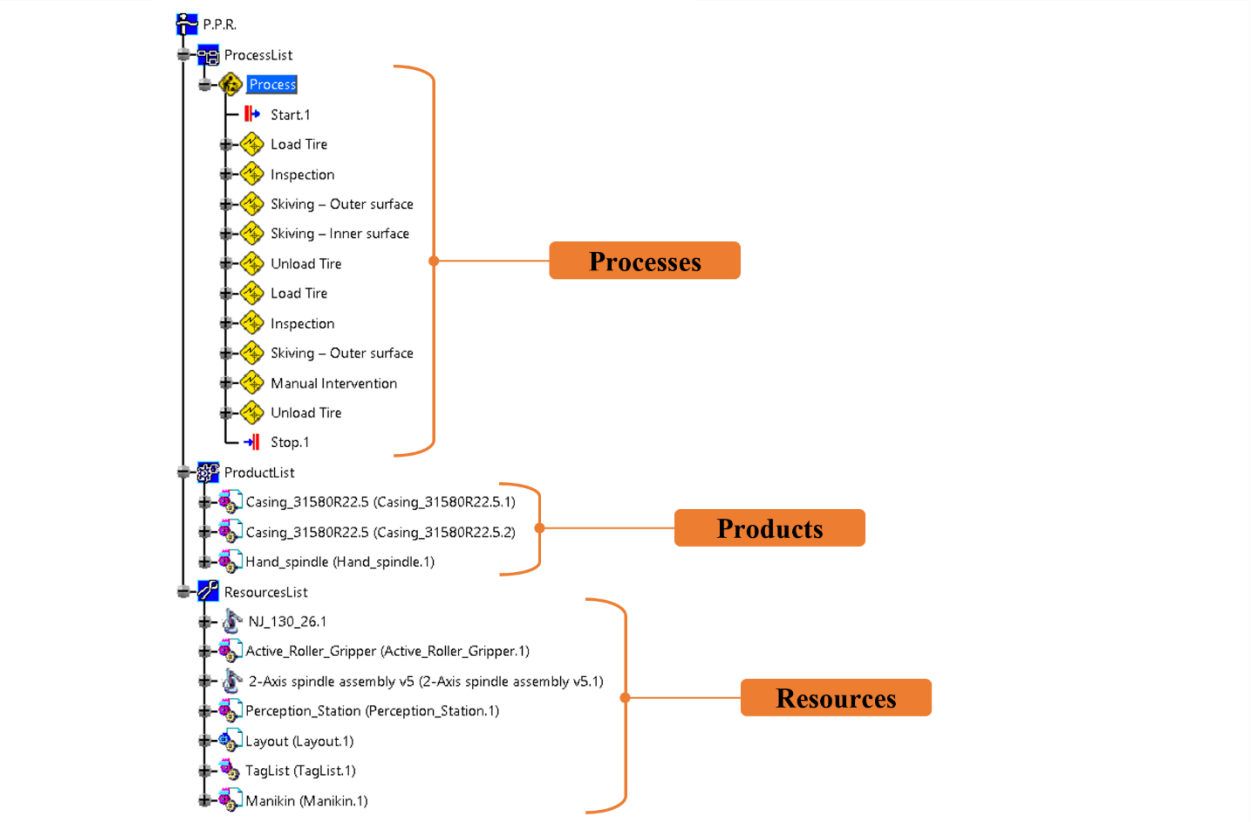


Fig. 25 The Product–Process–Resource structure of the simulated system, the tire as product, the grouped robot and gripper activities as processes, and the robot, gripper, and stations as resources.

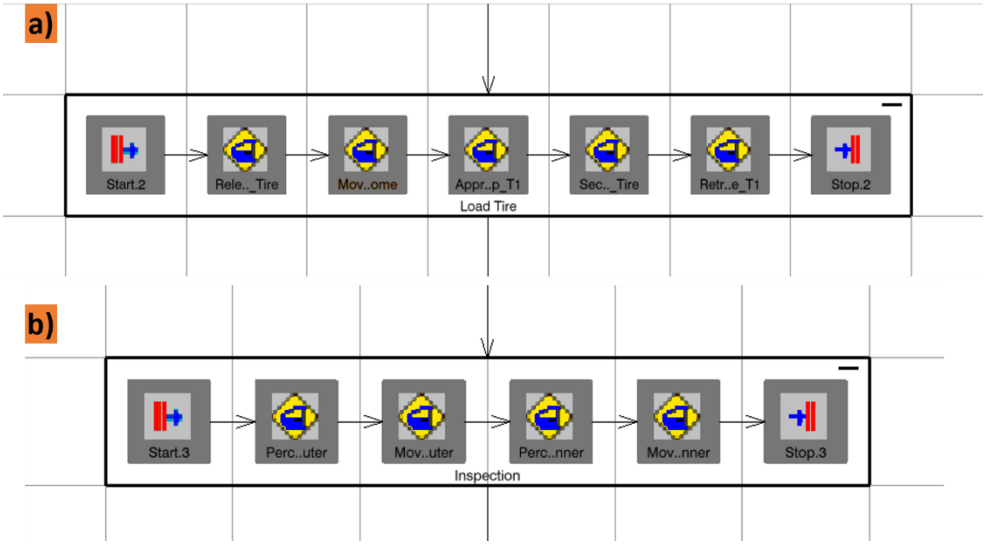


Fig. 26 Activity diagrams for the a) Loading of the tire and b) the inspection.

5.7 SCALABILITY OF THE PROPOSED WORKCELL

The single-cell design validated above is treated as the basic unit for scaling along two dimensions, throughput and size-variant coverage, each addressed by replicating validated units rather than redesigning the cell.

Throughput scaling follows from skiving being automated by default while the operator intervenes only as an exception, which makes the human a shared, comparatively scarce resource rather than one dedicated to a single cell. This suggests two identical cells side by side, served by a single operator who moves between them as needed, rather than doubling the robot or station count within one cell. Each cell retains its own independently monitored safety zone; entry into one cell's interaction zone is muted only when that cell requires intervention, while the neighboring cell continues automated operation at full speed.

Sharing one operator across two or more cells, however, turns throughput scaling into a line-balancing problem. The operator becomes a shared resource whose time must be distributed across cells that each generate intervention events unpredictably, according to the condition of the tires they happen to be processing. If two cells call for manual skiving at the same moment, one must wait, and the idle time this introduces, on the waiting cell's robot, or on the operator when no cell needs attention, is what determines whether adding cells actually yields proportional throughput. Balancing this is not a matter of layout but of scheduling: a dynamic task planner would be needed to allocate the operator, and any other shared resource, across the cells in a way that keeps resource idle time low and resolves competing demands for intervention as they arise. The design and evaluation of such a planner lies outside the scope of this thesis.

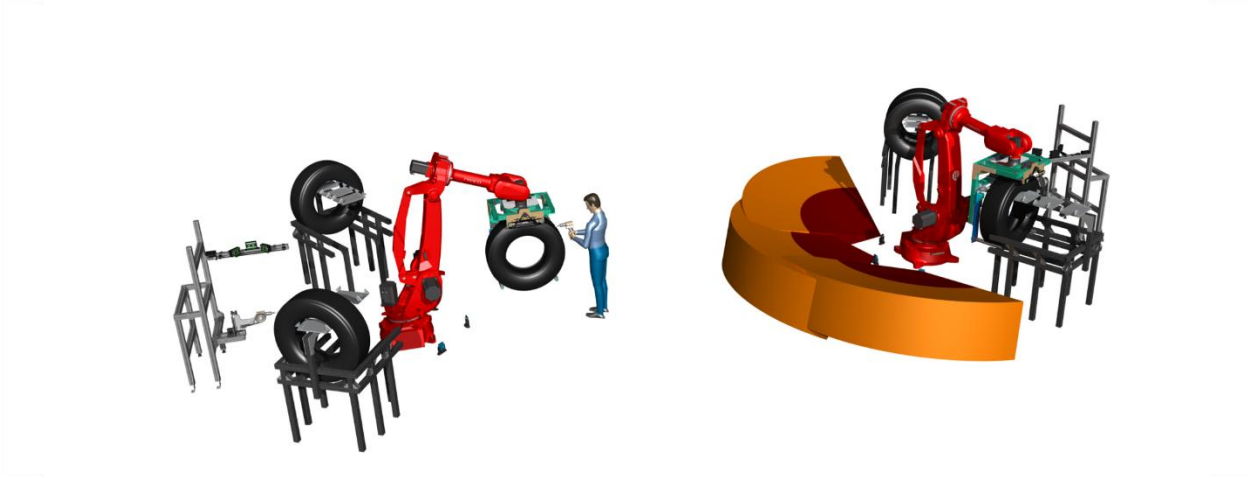


Fig. 27 DELMIA rendering of two adjacent cell layouts sharing a central operator zone, with independently monitored safety regions.

Coverage of tire size variants beyond the 315/80 R22.5 range follows the same replication logic rather than a reconfigurable design. As in Section 5.2, a single PRG is scoped to a defined size range, and extending coverage to meaningfully larger or smaller tires is more practically

achieved through separate, appropriately sized cells, each serving its own product range, than through one cell attempting the full spectrum. This mirrors the throughput approach in treating the validated single-cell unit as the building block for industrial-scale deployment, with capacity and product coverage both extended through replication rather than added complexity within any one cell.

6.0 RESULTS

6.1 CELL LAYOUT AND REACHABILITY VALIDATION

The layout described in Section 5.1 was validated iteratively within the CATIA/DELMIA environment of Section 5.6, through repeated adjustment of component positions until full reachability was confirmed across the complete working envelope shared by the perception and skiving stations, covering the entire outer and inner surface of the tire required by both. This process surfaced two issues that shaped the final layout. The first was a collision between the Powered Roller Gripper and the robot during the detection step, caused by the perception station's initial position sitting too close to the skiving tool mounted on the same structure; this was resolved by repositioning the perception station down and to the side to create adequate clearance. The second was a comparable reach limitation at the skiving station, where the tool's position was adjusted until it could reliably access the complete tire surface exposed for processing at each defect location. In both cases, the layout presented in Section 5.1 reflects the outcome of this correction process rather than the initial configuration, refined over several iterations before reaching its final form.

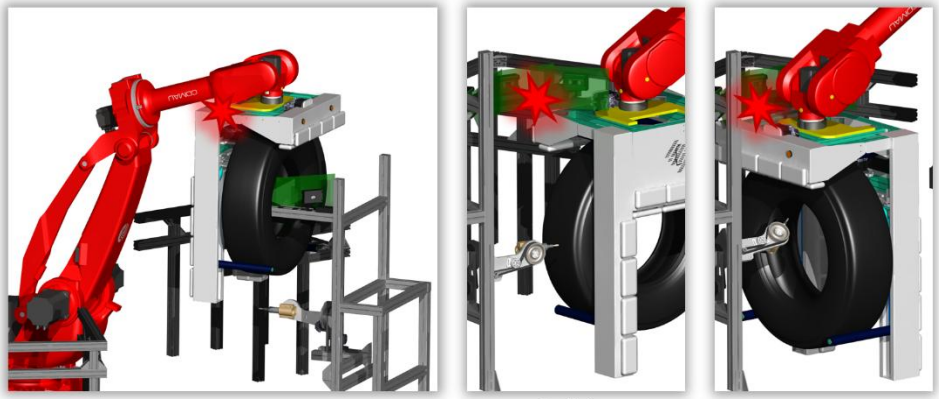


Fig. 28 Layout issues at the detection station, and the skiving station

As part of the same validation activity, the safety concept for the operator interaction zone at the skiving station was tested for spatial feasibility using the monitored regions of the safety radar introduced in Section 5.5. The concept establishes two nested zones around the interaction area: an outer zone in which robot and tool speeds are reduced upon operator entry, and an inner zone in which the robot and tool are brought to a monitored standstill.

The two zones were sized using the ISO 13855 minimum protective separation distance formula, $S = (K \times T) + C$, applied around the skiving station's operator interaction area rather than the robot's base, since the robot does not operate at full extension at this station. Using an approach speed of $K = 1600$ mm/s (whole-body detection, applicable to a radar-based volumetric zone), a radar response time of 100 ms (Inxpect S201A series), a conservative placeholder robot stopping time of 500 ms, and a generic detection-capability allowance of $C = 850$ mm pending confirmation from Inxpect, the resulting protective separation distance is $S \approx 1.81$ m. Combined with the bounding box of the Powered Roller Gripper holding a maximum-size tire at the interaction area (≈ 750 mm from center), this yields an illustrative inner (stop) zone radius of approximately 2.6 m and an outer (reduced-speed) zone radius of approximately 3.2 m from the interaction-area center.

These figures are deliberately conservative and retained as such rather than optimized down, for two reasons specific to this cell's operating logic. First, during automated operation the robot works exclusively at the rear of the cell, at the perception and skiving stations, and does not approach the operator interaction area at all, so the calculated zone need not account for continuous shared operation. Second, the only motion that does bring the robot near the interaction area, tire repositioning during manual intervention, is itself fully interlocked with the radar zones: repositioning executes only once the zone is confirmed clear, and the operator may re-enter only once the robot has completed its motion and returned to monitored standstill. The ISO 13855 margin is retained not to protect against ongoing proximity but as the safety-rated backstop that makes this interlock trustworthy, ensuring that even an unexpected re-entry attempt during the brief repositioning window would be detected with sufficient distance for the robot to stop before contact. On this basis, the conservative sizing is treated as a deliberate safety margin rather than an inefficiency, pending confirmation of the robot's certified stopping time and the radar's certified detection-capability rating through a formal risk assessment. This zoning logic was verified against the cell layout to confirm that the two regions could be defined consistently with the physical arrangement of the skiving station and the robot's working envelope.

6.2 SUBSYSTEM VALIDATION

6.2.1 Powered Roller Gripper

Validation of the Powered Roller Gripper was carried out kinematically within CATIA, exercising both of its mechanisms across the full size envelope established in Section 5.2. The

securing mechanism, the lower roller pair lowering to admit the tire and raising to complete the four-roller grasp, was confirmed to open wide enough to approach and position the gripper around the largest tire in range and to close fully onto the smallest, so that a stable four-point grasp is maintained across the whole envelope. The gripper accommodates a maximum buffed tire diameter of 1200 mm with a corresponding width of 370 mm without modification, comfortably above the nominal 315/80 R22.5 buffed casing it is sized for and leaving margin for the diameter spread that buffing introduces. The rotation mechanism, the upper roller pair, was likewise verified to rotate the held tire continuously through a full revolution, the motion on which both the perception sweep and the presentation of a localized defect to the skiving tool depend. Finally, the kinematic relationship underlying the pick-up diameter measurement, whereby the vertical travel of the securing assembly maps to the tire's effective diameter, was confirmed geometrically across the range, establishing that the measurement the gripper performs at grasp remains valid for any casing within the accommodated envelope.

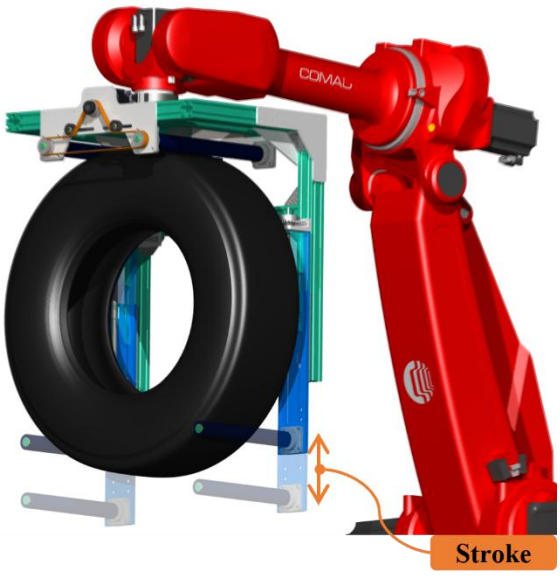


Fig. 29 Powered Roller Gripper kinematic range.

6.2.2 Inspection Station

The inspection station was validated by visualizing each of the three ZED X Mini cameras' fields of view against the tire geometry within CATIA, in both the outer- and inner-surface detection configurations and at the extremes of the accommodated diameter range. Each camera snapshot was confirmed to cover approximately 26° of the tire's circumference on the outer surface and approximately 50° on the inner surface, corresponding to 14 and 8 rotation steps respectively for full 360° coverage of each surface, with sufficient overlap between successive snapshots that no angular region of either surface is left unobserved as the gripper rotates the tire between them. The larger coverage obtained on the inner surface follows from the closer working distance and viewing geometry required to image the inner circumference, and was accommodated without adjusting the arrangement. The side-camera lead-screw assembly was verified to provide a stroke of 145 mm, sufficient to reposition the two side cameras between the outer- and inner-surface configurations across the full tire dimension range accommodated by the gripper, without any further adjustment to the camera mounts.

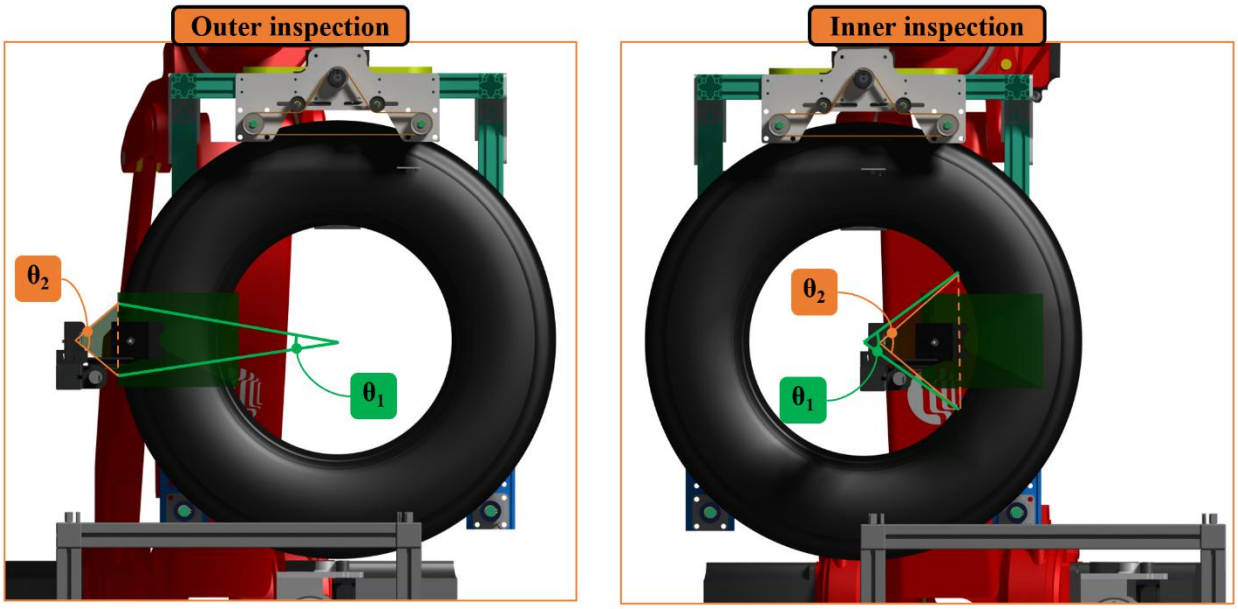


Fig. 30 Angle coverage of the selected cameras for both the outer and inner surface detection.

6.2.3 Skiving Station

The skiving tool was validated by driving its two-axis assembly through its full range within CATIA and checking, at each pose, that the tool remained correctly oriented against the tire's local surface. Both revolute joints were confirmed to provide a full 360° range of motion, giving the tool the additional dexterity a single robot arm could not supply on its own and allowing it to reach and process both the outer and inner surfaces at a previously localized defect across the crown, shoulders, and sidewalls. One reach limitation was identified: for defects on the inner sidewall or inner bead lying on the same side of the tire as the tool, the tool's own structure partially obstructs the approach angle needed to reach the defect directly. This is resolved operationally, by rotating the tire 180° about the robot's sixth axis so that the defect is presented from the opposite,

unobstructed side; while effective, this adds a repositioning step whenever it applies, and a mechanically slimmer tool structure that removes the obstruction is identified as a target for the next design iteration, consistent with Section 5.4.

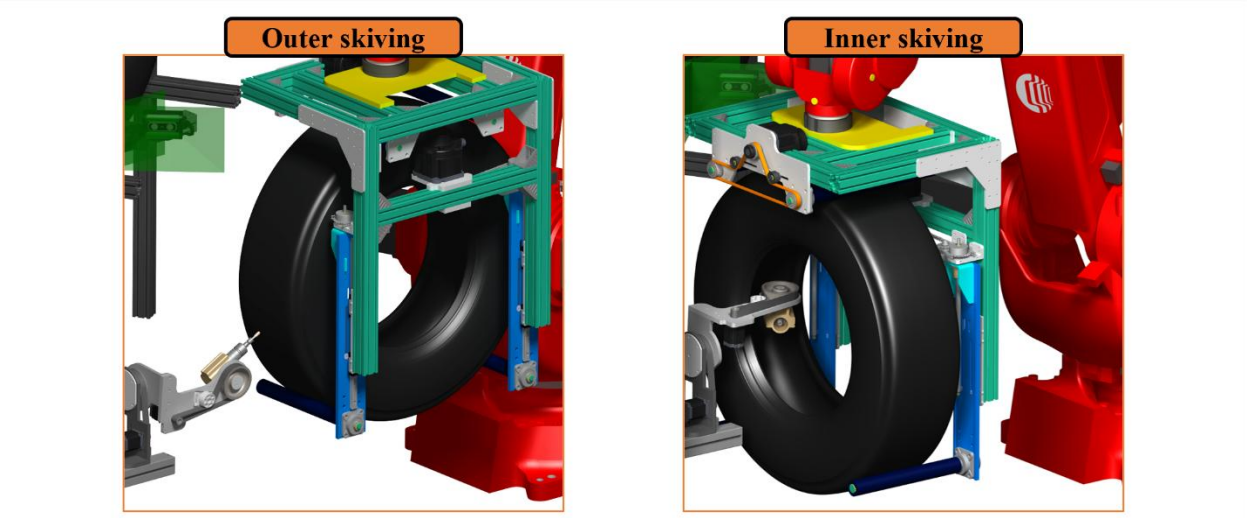


Fig. 31 Skiving tool reach across inner and outer casing profile.

6.3 PROCESS SEQUENCE SIMULATION

A representative process sequence, covering the robot picking up the tire, transferring it to the perception station, presenting it to the skiving station, and returning it, was simulated within the environment of Section 5.6. The simulation provides the duration of the robot motions that connect the operations, pick-up, transfer between stations, repositioning, and the switch between the automated and manual paths, so that a per-tire processing time is built up from repeatable, individually determined components rather than asserted as a single figure.

The sequence was modelled in the Product–Process–Resource (PPR) structure of the simulation environment, with the tire as the product, the individual robot and gripper motions as

processes, and the robot, gripper, and stations as the resources that carry them out. The low-level motions are grouped under higher-level activities, loading, inspection, transfer, skiving of each surface, and unloading, so that the sequence can be read at either level. Their precedence and the branch between the automated and manual paths are expressed as an activity network, from which the ordering below and the per-operation durations are taken.

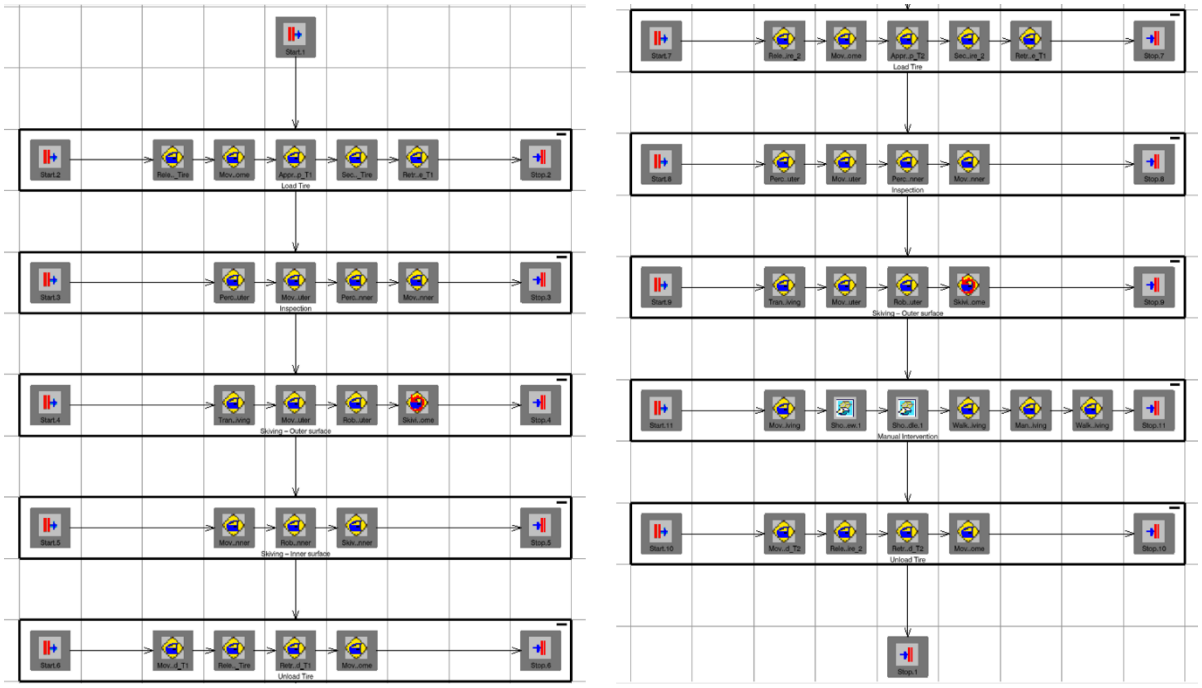


Fig. 32 Activity network (PERT) of the process sequence, showing the grouped activities and their precedence.

Detection presents the full surface to the cameras by rotating the tire through one revolution per surface; at the 10 rpm the gripper's rotation drive provides, one revolution takes 6 s, so the outer and inner sweeps take 12 s together (the 14 and 8 snapshot steps of Section 6.2.2 set the capture frequency, not the total rotation). Adding the side-camera repositioning between the two steps, 10 s for the 145 mm stroke, gives a detection time of approximately 22 s. Skiving time, by contrast, depends on the specific defect being handled, both for the automated cut and for a manual intervention, and so remains a case-dependent parameter to be supplied from the responsible team

or a time study. In the model below, T_detect therefore takes the value derived here, while $T_process$ and T_manual remain parameters.

The processing time for a tire is expressed as a sum of per-tire fixed operations and per-defect operations. The durations below are estimates intended to feed the scaled-up planner rather than certified cycle figures; they are drawn from the DELMIA simulation with the skiving operations left as parameters:

- Per-tire (fixed): loading, presentation of the tire at the inspection station, detection, the first transfer of the tire to the skiving point after detection, and unloading.
- Per-defect: the processing operation itself and, between successive defects, a short repositioning of the tire to the next defect. Under automated processing, the operation is the automated skiving; under manual intervention, it is instead the transfer to the manual area, the manual skiving, and the return to automated skiving.

Total processing time for a tire with n defects is then:

$$T(n) = T_load + T_inspect + T_detect + T_transfer + T_unload + n \cdot T_process + (n-1) * T_reposition$$

with the manual path substituting the transfer to the manual area, the manual skiving, and the return to automated skiving, $T_switch + T_manual + T_return$, for the automated skiving $T_process$ on any defect requiring intervention.

Table 8 Per-operation duration estimates.

Operation	Type	Symbol	Duration (s)
Loading (tire pick-up)	Per-tire	T_load	55
Presentation at inspection station	Per-tire	T_inspect	60
Detection	Per-tire	T_detect	22
First transfer to the skiving point	Per-tire	T_transfer	20
Unloading (tire return)	Per-tire	T_unload	55
Repositioning to the next defect	Between defects	T_reposition	5
Automated skiving	Per-defect	T_process	parameter
Transfer to manual area	Per intervention	T_switch	30
Manual skiving	Per intervention	T_manual	parameter
Return to automated skiving	Per intervention	T_return	30

With these values, the fixed per-tire contribution is $T_load + T_inspect + T_detect + T_transfer + T_unload = 55 + 60 + 22 + 20 + 55 = 212\,s$, incurred once per tire. Each defect then adds the skiving cut $T_process$, and every defect after the first adds a

repositioning of $T_reposition \approx 5\,s$ to bring the next defect to the skiving point. A defect handled manually instead of automatically replaces its $T_process$ with the transfer to the manual area, the manual skiving, and the return to automated skiving, $T_switch + T_manual + T_return = 30 + T_manual + 30$. The skiving durations $T_process$ and T_manual remain case-dependent parameters. Substituting into the model gives the totals in Table 9. The motion values reflect the current motion profiles and would be revised if the profiles are retuned.

Table 9 Estimated total processing time per tire, by defect count, under a fully automated scenario and a scenario with manual intervention.

Defects (*n*)	Fully automated (s)	With one manual intervention (s)
2	217 + 2·T_process	277 + T_process + T_manual
5	232 + 5·T_process	292 + 4·T_process + T_manual
10	357 + 10·T_process	317 + 9·T_process + T_manual
15	382 + 15·T_process	342 + 14·T_process + T_manual

The two scenarios bound the range, the fully automated column processes every defect by automated skiving, while the manual column reroutes one defect through the operator, replacing that defect's automated skiving $T_process$ with the transfer to the manual area, the manual skiving, and the return to automated skiving, $T_switch + T_manual + T_return$, and so adds a fixed $(60 + T_manual - T_process)$ relative to the fully automated total.

The simulated sequence also incorporated a scenario in which one or two defects are handled through manual intervention rather than automated skiving, replicating the exception path of Section 5.5 for cases where the automated process cannot be relied upon. This confirmed that the sequence remains coherent when the human-intervention path is triggered, not only under the fully automated case.

Beyond validating the sequence, these per-operation timings are the quantities a multi-cell deployment would build on. As set out in Section 5.7, scaling the system shares one operator across several cells, which turns throughput into a scheduling problem, a planner balancing operator and robot tasks across cells requires exactly these normalized per-operation and per-defect durations, together with the automated-versus-manual split, to estimate each cell's demand on the shared operator and to keep resource idle time low.

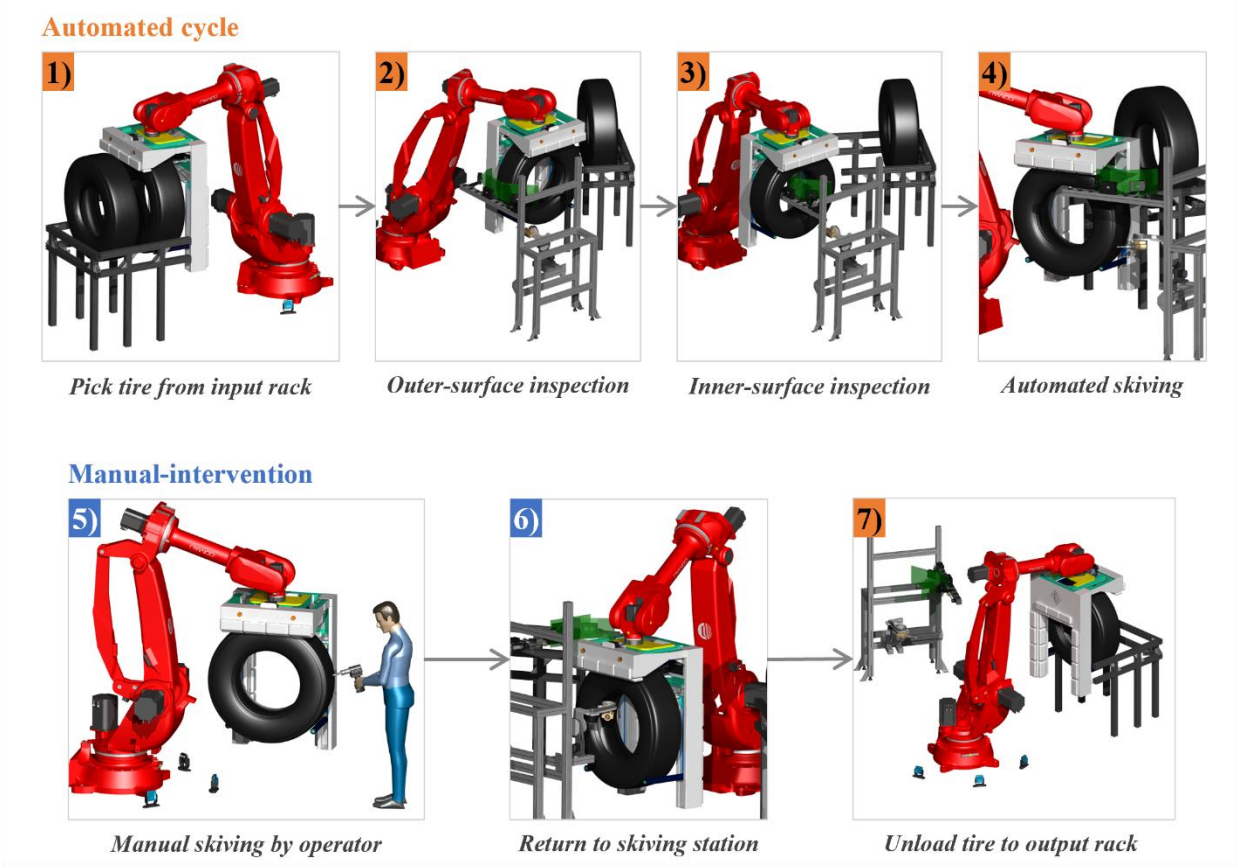


Fig. 33 Simulated process sequence in DELMIA: 1) Pick tire from input rack, 2) Outer-surface inspection at the perception station, 3) Inner-surface inspection after repositioning of the side cameras, 4) Automated skiving at the skiving station, 5) Manual skiving by the operator in the interaction area, 6) Return to the skiving station, 7) Unload to the output rack.

A further result concerns the operator's physical workload, which the cell improves in three ways relative to current practice. In the manual workstation of Section 3.2, the operator picks up and transfers the heavy casing by hand and then skives every defect manually, in postures dictated by how the tire is placed. In the proposed cell, the robot carries and holds the casing throughout, so the operator no longer lifts or transfers it. The gripper's indexing rotation (Section 5.5) brings each defect to an orientation, and the robot's movements allow for an optimal working height chosen for the operator. The manual skiving is required only for the exception defects the

automated path defers, reducing the number of defects an operator processes by hand per tire. Together these shift the operator from continuous heavy handling in awkward postures toward selective, supervised intervention in a controlled stance. A before-and-after comparison of the operator's working posture is shown in Fig. 34. The need for posture investigation is raised based on the RULA scores calculated in CATIA, regarding the initial and final lifting postures of the human operator during manual tire loading.

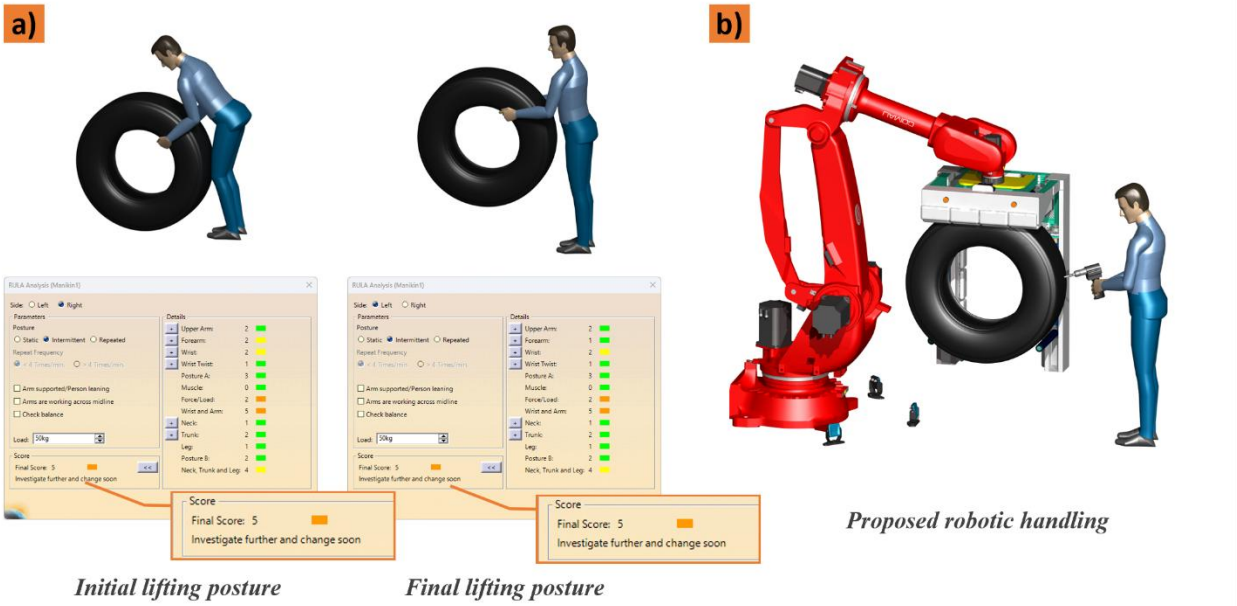


Fig. 34 Operator working posture in a) current manual practice loading the tire & RULA score versus b) the proposed robotic solution.

7.0 CONCLUSIONS AND FUTURE WORK

7.1 CONCLUSIONS

This thesis sets out to address a difficulty that distinguishes remanufacturing, and the wider family of Re-X processes, from conventional manufacturing. The incoming item is not a known part but a returned core whose variant and condition are established only after it arrives, so the information needed to process it must be generated at runtime rather than inherited from a specification. Fixed automation is therefore poorly suited to Re-X, and much of the work remains manual and operator-bound even as end-of-life volumes grow. Against this background, the thesis asked how a robotic cell can be designed to assess, decide on, and process such a variable item while keeping the operator in the loop wherever automation cannot yet reach.

The approach is a general, end-to-end methodology for the design of Re-X cells, organized around three pillars, condition assessment, a traceability layer, and a processing strategy derived per item, with a human–robot collaboration strategy running across all three. On this basis, the methodology sets out the cell-engineering activities that turn the logic into hardware, from workflow definition and process-and-resource allocation through layout and tooling to digital evaluation, worked through iteratively rather than in a single pass. The methodology is deliberately product-independent, so that the same activities apply to other Re-X processes beyond the one used to validate it.

The contribution is accordingly twofold. The first part is a transferable design methodology that moves from condition assessment through processing decision to a safety-aware collaborative layout, and the second is its applied realization as a fully specified tire-remanufacturing cell built

around a Comau NJ130-2.6 robot, a Powered Roller Gripper, a three-camera perception station, and a two-axis skiving tool.

The cell was validated digitally in CATIA V5 and DELMIA, and the three performance targets set in Section 3.3.3 were each addressed. Full-surface automated inspection (T1) was confirmed, with the perception station covering both the outer and inner casing surfaces across the full accommodated size range and no blind spots. Full-surface automated processing (T2) was confirmed with one caveat, since the skiving tool reaches any localized defect on both surfaces while a single inner-sidewall and bead edge case is resolved operationally by a 180° tire rotation. Operator safety under human–robot collaboration (T3) was met at the design-baseline level, with every operator-protecting function designed to a Category 3 / Performance Level d safety baseline and the protective separation distances sized accordingly, and formal per-function risk assessment left to future work. Alongside these, the working envelope was shown to be fully reachable and collision-free after resolving two interference issues, the gripper and station kinematics were validated across the tire size range, and a per-operation timing model was built to express cycle time as a function of defect count under both fully automated and manual-intervention scenarios.

The work is intended to be of value to several groups. For the tire-retreading industry, and for Re-X operators more broadly, it offers a concrete, safety-aware design for automating the inspection and localized processing of a variable casing while retaining the operator's judgment and relieving the physical burden of manual tire handling. For researchers and system integrators, it provides a reusable methodology and a worked reference case for designing collaborative cells around uncertain input, a problem that recurs across remanufacturing. And for the broader circular-economy agenda, it is a step toward industrializing the value-retention processes that remain largely manual today, helping to close the gap between the rate at which end-of-life products

accumulate and the rate at which they can be returned to use. In this sense the thesis argues that the obstacle to automating Re-X is not any single technology, but the absence of a coherent way to combine assessment, decision, processing, and human involvement within one cell, which is what the proposed methodology sets out to provide.

7.2 FUTURE WORK

The work opens directions ranging from immediate engineering steps to broader research. In the near term, the designed cell would move to physical implementation. Building and commissioning it would replace the simulated timings with measured cycle times, verify the safety functions through a formal risk assessment, characterize the spread of cycle times across many varied runs rather than the single representative sequence reported here.

The more significant direction lies in making the cell more adaptive, more interactive, and more accountable. A first direction is adaptive orchestration. The current workflow, with its branching per-defect paths and its exception route to manual intervention, is well suited to being expressed as a behavior-tree architecture. Expressing it this way would make the sequence reconfigurable and reactive at runtime, with the reassignment of work between operator and automation being dynamically performed as conditions change rather than following a predetermined program. This would turn the fixed logic validated here into a flexible orchestration layer that can absorb new operations and new scenarios without redesign.

A second direction is natural, language-based interaction. State-of-the-art language and multimodal models open the possibility of an operator querying, confirming, or redirecting the

cell's work in plain language rather than through dedicated controls. Such an interface would lower the expertise barrier to operating the cell and raise productivity, aligning the system more closely with the human-centric intent of Industry 5.0.

A further direction looks past the technical, to the work's wider impact. A comparison of its environmental footprint against both manual retreading and new-tire production would quantify the sustainability case that motivates the work. The concept also scales beyond a single cell, through the shared-operator, multi-cell configuration of Section 5.7 and the dynamic planner such a configuration would require. Together these directions point toward an intelligent, human-centric, and environmentally accountable approach to Re-X automation, for which this thesis lays the foundation.

REFERENCES

[1] George Chryssolouris, "Manufacturing Systems: Theory and Practice," *Manufacturing Systems: Theory and Practice*, 2006, doi: 10.1007/0-387-28431-1.

[2] Y. Koren, X. Gu, and W. Guo, "Reconfigurable manufacturing systems: Principles, design, and future trends," *Frontiers of Mechanical Engineering*, vol. 13, no. 2, pp. 121–136, Jun. 2018, doi: 10.1007/S11465-018-0483-0/METRICS.

[3] "Butterfly Diagram | Circular Economy Infographic | Ellen MacArthur Foundation." Accessed: Sep. 20, 2026. [Online]. Available: https://www.ellenmacarthurfoundation.org/circular-economy-diagram

[4] J. Potting, M. P. Hekkert, E. Worrell, and A. Hanemaaijer, "Circular Economy: Measuring Innovation in the Product Chain," 2017, *PBL Publishers*. Accessed: Sep. 20, 2026. [Online]. Available: https://dspace.library.uu.nl/handle/1874/358310

[5] M. Matsumoto, S. Yang, K. Martinsen, and Y. Kainuma, "Trends and research challenges in remanufacturing," *International Journal of Precision Engineering and Manufacturing-Green Technology 2016 3:1*, vol. 3, no. 1, pp. 129–142, Jan. 2016, doi: 10.1007/S40684-016-0016-4.

[6] J. Leng *et al.*, "Industry 5.0: Prospect and retrospect," *J. Manuf. Syst.*, vol. 65, pp. 279–295, Oct. 2022, doi: 10.1016/J.JMSY.2022.09.017.

[7] S. Makris, "Cooperating Robots for Flexible Manufacturing," *Springer Series in Advanced Manufacturing*, vol. 2021, pp. 1–409, 2021, doi: 10.1007/978-3-030-51591-1/SAVE-RESEARCH.

[8] M. Kerin and D. T. Pham, "A review of emerging industry 4.0 technologies in remanufacturing," *J. Clean. Prod.*, vol. 237, p. 117805, Nov. 2019, doi: 10.1016/J.JCLEPRO.2019.117805.

[9] J. P. Kaiser, S. Lang, M. Wurster, and G. Lanza, "A Concept for Autonomous Quality Control for Core Inspection in Remanufacturing," *Procedia CIRP*, vol. 105, pp. 374–379, Jan. 2022, doi: 10.1016/J.PROCIR.2022.02.062.

[10] L. Gualtieri, E. Rauch, and R. Vidoni, "Emerging research fields in safety and ergonomics in industrial collaborative robotics: A systematic literature review," *Robot. Comput. Integr. Manuf.*, vol. 67, p. 101998, Feb. 2021, doi: 10.1016/J.RCIM.2020.101998.

[11] G. Michalos *et al.*, "Seamless human robot collaborative assembly – An automotive case study," *Mechatronics*, vol. 55, pp. 194–211, Nov. 2018, doi: 10.1016/J.MECHATRONICS.2018.08.006.

[12] D. Andronas *et al.*, "On the perception and handling of deformable objects – A robotic cell for white goods industry," *Robot. Comput. Integr. Manuf.*, vol. 77, p. 102358, Oct. 2022, doi: 10.1016/J.RCIM.2022.102358.

[13] K. Katsampiris-Salgado *et al.*, "Collision detection for collaborative assembly operations on high-payload robots," *Robot. Comput. Integr. Manuf.*, vol. 87, p. 102708, Jun. 2024, doi: 10.1016/J.RCIM.2023.102708.

[14] M. L. Lee, X. Liang, B. Hu, G. Onel, S. Behdad, and M. Zheng, "A Review of Prospects and Opportunities in Disassembly With Human–Robot Collaboration," *J. Manuf. Sci. Eng.*, vol. 146, no. 2, Feb. 2024, doi: 10.1115/1.4063992/1170009.

[15] J. Xiao and K. Huang, "A comprehensive review on human–robot collaboration remanufacturing towards uncertain and dynamic disassembly," *Manuf. Rev. (Les. Ulis).*, vol. 11, p. 17, 2024, doi: 10.1051/MFREVIEW/2024015.

[16] D. Vallianatos, D. Kaliakatsos-Georgopoulos, D. Dimosthenopoulos, C. Gkournelos, G. Michalos, and S. Makris, "A Methodology for Applying HRC on Remanufacturing Workstations: A Use Case from Household Appliances," *Lecture Notes in Mechanical Engineering*, pp. 815–825, 2026, doi: 10.1007/978-3-032-16889-4_77/FIGURES/4.

[17] V. Villani, F. Pini, F. Leali, and C. Secchi, "Survey on human–robot collaboration in industrial settings: Safety, intuitive interfaces and applications," *Mechatronics*, vol. 55, pp. 248–266, Nov. 2018, doi: 10.1016/J.MECHATRONICS.2018.02.009.

[18] J. A. Marvel and R. Norcross, "Implementing speed and separation monitoring in collaborative robot workcells," *Robot. Comput. Integr. Manuf.*, vol. 44, pp. 144–155, Apr. 2017, doi: 10.1016/J.RCIM.2016.08.001.

[19] M. Dianatfar, J. Latokartano, and M. Lanz, "Review on existing VR/AR solutions in human–robot collaboration," *Procedia CIRP*, vol. 97, pp. 407–411, Jan. 2021, doi: 10.1016/J.PROCIR.2020.05.259.

[20] C. Gkournelos, C. Konstantinou, P. Angelakis, E. Tzavara, and S. Makris, "Praxis: a framework for AI-driven human action recognition in assembly," *Journal of Intelligent Manufacturing 2023 35:8*, vol. 35, no. 8, pp. 3697–3711, Nov. 2023, doi: 10.1007/S10845-023-02228-8.

[21] E. Coupeté, F. Moutarde, and S. Manitsaris, "Multi-users online recognition of technical gestures for natural human–robot collaboration in manufacturing," *Autonomous Robots 2018 43:6*, vol. 43, no. 6, pp. 1309–1325, Feb. 2018, doi: 10.1007/S10514-018-9704-Y.

[22] E. Montini *et al.*, "An industrial human-robot collaboration case study for workers' well-being," *Procedia CIRP*, vol. 130, no. 27, pp. 924–929, Jan. 2024, doi: 10.1016/J.PROCIR.2024.10.186.

[23] W. Yuanyang and M. N. Mahyuddin, "Grasping Deformable Objects in Industry Application: A Comprehensive Review of Robotic Manipulation," *IEEE Access*, vol. 13, pp. 33403–33423, 2025, doi: 10.1109/ACCESS.2025.3543737.

[24] S. Makris, F. Dietrich, K. Kellens, and S. J. Hu, "Automated assembly of non-rigid objects," *CIRP Annals*, vol. 72, no. 2, pp. 513–539, Jan. 2023, doi: 10.1016/J.CIRP.2023.05.003.

[25] S. Makris, E. Kampourakis, and D. Andronas, "On deformable object handling: Model-based motion planning for human-robot co-manipulation," *CIRP Annals*, vol. 71, no. 1, pp. 29–32, Jan. 2022, doi: 10.1016/J.CIRP.2022.04.048.

[26] J. Gaidhane, I. Ullah, and A. Khalatkar, "Tyre remanufacturing: A brief review," *Mater. Today Proc.*, vol. 60, pp. 2257–2261, Jan. 2022, doi: 10.1016/J.MATPR.2022.04.142.

[27] "Automatic detection and robot-assisted machining of surface defects," May 2018, doi: 10.2312/EGST.20141040.

[28] U. Patent FIG, "Process and machine for identification and working of defects on used tyres," Mar. 2010.

[29] "Skiving Machine for Retreaded Tires," May 1991.

[30] F. J. Heirtzler, "Tire defect detection system and method," Jun. 2000.

[31] Shen, "Method and apparatus for non-destructive detection of tire anomalies," Nov. 2013.

[32] N. Kousi *et al.*, "Digital Twin for Designing and Reconfiguring Human–Robot Collaborative Assembly Lines," *Applied Sciences 2021, Vol. 11, Page 4620*, vol. 11, no. 10, p. 4620, May 2021, doi: 10.3390/APP11104620.

[33] G. Wöhlke and E. Schiller, "Digital Planning Validation in automotive industry," *Comput. Ind.*, vol. 56, no. 4, pp. 393–405, May 2005, doi: 10.1016/J.COMPIND.2005.01.010.

[34] P. Tsarouchi, S. Makris, and G. Chryssolouris, "Human–robot interaction review and challenges on task planning and programming," *Int. J. Comput. Integr. Manuf.*, vol. 29, no. 8, pp. 916–931, Aug. 2016, doi: 10.1080/0951192X.2015.1130251.