POLITECNICO DI TORINO

Collegio di Ingegneria Informatica, del Cinema e Meccatronica



Master's Degree Thesis in Mechatronic Engineering

Comparative Study of Scanning Techniques in a Reverse Engineering Process for the Acquisition and Topology Optimisation of Robot End-Effector

Supervisors

Prof. Paolo Minetola

Politecnico di Torino

Prof. Santiago Carlos Gutiérrez Rubert

Universitat Politècnica de València

Candidate

Alessandro Luigi Zin

Academic Year 2025 – 2026

Table of Contents

Abstract 1
Acknowledgements 3
List of Figures 5
List of Tables 15
Acronyms and Abbreviations 19
1. Introduction 21
1.1 Objectives and Scope 21
1.2 Workflow Overview and Thesis Structure 22
1.3 Limitations 24
1.4 Alignment with UN SDGs 25
2. Component Selection and Characterisation 27
2.1 Application Context 27
2.2 Selection Criteria and Component Description 29
2.3 Geometric, Functional and Material Characteristics 31
2.4 Printing and Assembly 33
2.5 Normal Working Conditions 37
3. 3D Scanning 41
3.1 Overview and Technology Selection 41
3.2 Handheld Structured-Light Scanning 46
3.3 Fixed-Head Structured-Light Scanning 52
3.4 X-ray Computed Tomography 58
3.5 Comparison and Issues 69
4. Reverse Engineering of the Infill 73
4.1 Setup for the Reverse Engineering 73
4.2 CAD Reconstruction in SolidWorks from the CT Data 73
4.3 Dimensional Comparison and Considerations 78
5. FEM Structural Analysis 83
5.1 FEM Setup: Material Model, Constraints and Load Cases 83
5.2 Results on the original CAD 92
5.3 Results on RE CT and Discussion 97
6. Topology Optimisation 103
6.1 Fundamentals of Topology Optimisation 103

6.2 Design-Space Setup 104
6.3 Raw Optimisation Results 107
7. Redesign and Manufacturing 111
7.1 Redesign Based on Topology Optimisation 111
7.2 Structural Verification of Redesigned Parts 113
7.3 Mass Reduction Assessment 117
8. Cost and Time Analysis 119
8.1 Scope and Methodology of Cost Estimation 119
8.2 Time Breakdown by Phase 120
8.3 Cost Summary Table 121
9. Conclusions 123
9.1 Encountered Issues and Limitations 123
9.2 Summary of Results 124
Bibliography 127
Appendices 129
A - Technical drawings of the components 130
B - EinScan Pro 2X Inspected Results 136
C - GOM ATOS Core 300 Inspected Results 139
D - Phoenix v|tome|x s Results 142
E - Outer-Shell Extraction Algorithm (HPR) 161
F - HSV and Filters Algorithm 169
G - FEM Servo Stall forces 173
H - FEM SF & Displacement Results 179
I - Technical Specifications Document 203

Abstract

This thesis studies how much mass can be removed from the structural parts of a robotic end-effector without loss of structural integrity and compares the scanning techniques that make such a redesign possible. The component is a servo-driven parallel gripper, reproduced as an acrylonitrile styrene acrylate (ASA) specimen by fused deposition modelling (FDM) from the publicly available computer-aided design (CAD) data, with a target mass reduction of 10–30 [%] on its main structural parts. The specimen was digitised with three non-contact techniques: handheld structured-light scanning (EinScan Pro 2X), fixed-head structured-light scanning (GOM ATOS Core 300) and X-ray computed tomography (CT). Measured against the nominal model, the three agreed on a mean deviation of −0,040 to −0,070 [mm], set by the shrinkage of the printed parts rather than by the instruments, with standard deviations of 0,147, 0,117 and 0,096 [mm] respectively. The CT was taken as the master for the reverse engineering in SolidWorks not for its accuracy but because it alone resolved the internal infill. The reconstructed model was analysed by finite elements in Altair Inspire under the reference load cases and then lightened by topology optimisation, after which the frame and claws were remodelled to be “additive-manufacturing friendly” and reprinted. The redesign reduced the printed frame by 40,4 [%] and each claw by 53,8 [%], with governing safety factors of 3,4 and 6,0 against a target of 2,0; combined with the change of manufacturing route, the reworked specimen weighs 89,5 [g] against the 133 [g] catalogue value, an overall reduction of 32,7 [%], 23,8 [%] of which comes from in-house fabrication.

Keywords: reverse engineering; 3D scanning; computed tomography; topology optimisation; additive manufacturing; robot end-effector; lightweighting.

Acknowledgements

Ultima cosa che ho scritto di questa tesi ma sicuramente la più importante, la parte che ho rimandato più a lungo per mancanza di parole. Durante questi cinque anni ho avuto momenti alti e bassi, situazioni di stress e momenti nei quali onestamente avrei voluto mollare o voltare pagina. Nonostante ciò, grazie anche alle persone che mi sono state accanto durante questo tempo ho stretto i denti e ho continuato ad andare avanti (ma non senza lamentele... anzi!).

Non so se questa laurea mi servirà per i miei obiettivi di vita, ma sicuramente se siamo qui oggi a leggere queste parole, lo devo a chi ha avuto la fortuna (o meno) di avere a che fare con me.

Parto ringraziando chi l'odissea l'ha condivisa con me: i miei compagni di viaggio di Torino e di Vicenza, con i quali ho condiviso momenti di tempesta in sessione ed estati folli al mare. Con voi ho passato dei bellissimi momenti: giocare a frisbee invece di andare a lezione, videochiamate disperate prima degli esami, laboratori randomici nei quali contava più il divertimento che vivevamo del contenuto che imparavamo e le vacanze assieme a base di silicio e soluzioni a base di etanolo. Assieme abbiamo fatto tante cose e con il cuore in mano spero di poterne fare tante altre: grazie per tutte queste avventure, vi voglio bene.

Ai miei ex docenti e colleghi del Chilesotti, che mi hanno sempre spinto a dare il meglio di me. Nonostante qualche visione divergente, ci siamo sempre trovati su strade convergenti. È anche sicuramente grazie a loro se oggi su questa tesi c'è scritto "ingegneria".

Un grazie particolare va poi a queste persone che sono state più che centrali durante questi anni:

- a mia Madre, che nonostante i "non respirare, non mangiare, non guardare, non parlare che poi le tende si sporcano" durante questi anni ha sempre sostenuto con convinzione questo percorso, perché lo completassi e lo portassi a casa.
- a mio Padre, che grazie al portarmi a fiere di modellismo e di scienza è riuscito a passarmi l'amore per le materie tecniche che poi si è sviluppato in un interesse profondo per la tecnologia. Se non mi avesse aiutato con i teoremi di Euclide non so se saremmo qui oggi.
- a mio Fratello, che nonostante un carattere "complesso" e nonostante litigassimo molto da piccoli, è una persona che stimo moltissimo e che nel tempo ha saputo regalarmi sicuramente molte soddisfazioni e del quale vado fiero dicendo "quello è mio fratello".
- ai miei compagni di avventura a Valencia, con cui ho condiviso momenti splendidi e avventure indimenticabili: a Matteo, con cui ogni discorso finiva immancabilmente in "ragazzi, io vado al bar"; a Pietro, con cui ho affrontato l'ultimo anno di corsi assieme e che mi è stato d'aiuto durante i momenti di difficoltà con la tesi; e a Elisa, con cui ho costruito un rapporto di odio-amore, in questo preciso ordine cronologico, e alla quale voglio davvero un mondo di bene.

Menzione speciale per Ugo, il mio gatto d'infanzia, che è sempre stato un supporto emotivo e al quale ho voluto sempre un bene infinito.

Ma il grazie più grande va a chi queste pagine non potrà leggerle: a te, nonna Vanna, con cui ho passato la mia infanzia, tra le patatine fritte a base d'olio e bigoli co l'arna. A te che hai sempre creduto in me e che te ne sei andata proprio il giorno in cui avremmo dovuto festeggiare il primo grande traguardo. Sappi che oggi abbiamo raggiunto lo scalino successivo e ce ne saranno ancora molti altri. Ti ricorderò per sempre.

PS: un grazie anche a me stesso, che nonostante tutte le difficoltà non ha mai mollato. Se oggi siamo qui, è grazie a quella tenacia, a quella testardaggine e a quella convinzione.

List of Figures

Figure 1: Servo-Driven Parallel Gripper Kit by servocity.com as presented in the assembly file. 21

Figure 2: Thesis workflow diagram illustrating the full design cycle: from component selection and 3D printing, through 3D scanning and FEM analysis, to topology optimisation, redesign, and final reprint. 22

Figure 3: KUKA industrial robotic arms operating within an automated manufacturing cell. 27

Figure 4: ABB FlexPicker delta robot designed for high-speed industrial pick-and-place operations. 28

Figure 5: Original end-effector CAD visualised in SolidWorks. 30

Figure 6: Ultimaker Cura slicing preview of the frame component, showing the 20 [%] grid infill pattern and solid perimeter shells adopted for all prints in this work. 35

Figure 7: First-layer deposition of the two gear arm components on the Ultimaker S3 build plate, illustrating the adhesion quality achieved with the 95 [°C] bed temperature and adhesion-spray surface preparation. 35

Figure 8: Still frame extracted from a video recording of the frame print in progress. The 20 [%] grid infill pattern is clearly visible in the interior; minor surface irregularities attributable to nozzle movement are also apparent but do not affect the structural performance of the component. 36

Figure 9: Reference frame adopted for the structural analysis and load-case definition. 37

Figure 10: End-effector in the worst-case geometric configuration adopted for the structural analysis: fully extended kinematic chain. This configuration maximises the bending moment arm at the frame mounting points and is therefore used as the reference for all FEM and topology optimisation runs. 38

Figure 11: Position of the centre of mass of the representative 150 [g] incompressible payload relative to the two claw tips. For the worst-case analysis, the contact reaction force from the object is assumed to act at the foremost point of the finger profile. 38

Figure 12: LK Altera Type S 10.7.6 five-axis coordinate measuring machine (CMM) used for precision dimensional inspection, shown here as a representative example of contact-based measurement methods. Dimensional Metrology Laboratory, UPV. 42

Figure 13: Application of photogrammetry techniques for the digital acquisition of architectural heritage. 43

Figure 14: Roland PICZA LPX-60 laser-triangulation 3D scanner that captures object geometry by projecting a laser beam onto the surface and triangulating its reflected position, generating a precise point cloud. 43

Figure 15: Shining3D EinScan Pro 2X 2020 handheld structured-light 3D scanner that projects fringe patterns onto the object surface, computing geometry through phase-shift analysis and SLAM to generate a high-accuracy point cloud. 44

Figure 16: GE Phoenix v|tome|x s: an industrial X-ray computed tomography (CT) system that reconstructs both internal and external geometry of components through volumetric radiographic acquisition, without any contact with the part. 45

Figure 17: Capturing high-density point clouds using the EinScan Pro 2X 2020 handheld 3D scanner in hybrid scanning mode. 46

Figure 18: Mounting stand with markers in place and gripper held open, with stain-remover spray (quitamanchas) applied to enable acquisition of the dark surfaces of the servo motor. . 47

Figure 19: Assembly flipped on the stand with the base markers removed, allowing acquisition of the claw tips that had previously been covered by the modelling clay used as a temporary fixture. 48

Figure 20: EXScan Pro preview of the acquired point cloud, prior to the removal of out-of-interest regions. The preview provides an effective qualitative assessment of the reconstruction quality before proceeding to the meshing stage. 48

Figure 21: First mesh obtained after removal of the redundant points. Hole-filling is then performed, on the gaps left by the markers. 49

Figure 22: STL mesh of the complete end-effector assembly imported into Autodesk Meshmixer. Geometric defects are visible across the assembly, concentrated particularly on the Frame component. 49

Figure 23: Isolated and repaired Frame mesh obtained by applying Boolean and trimming operations to the raw EinScan point cloud, used as the input for the GOM Inspect dimensional comparison. 50

Figure 24: Surface deviation map of the EinScan-derived Frame mesh with respect to the nominal CAD model, obtained after pre-alignment and global best-fit registration in GOM Inspect. Colour scale represents signed dimensional deviation in millimetres. 51

Figure 25: GOM ATOS Core 300 fixed-head SLS system with the gripper component under structured-light projector before starting the acquisition process. 52

Figure 26: Right claw with retro-reflective markers applied for ATOS acquisition. Compared with the EinScan campaign (see Figure 19), the markers used here are substantially smaller and do not interfere with the component's edges. 53

Figure 27: Link as seen by the ATOS cameras during marker placement. Green crosses indicate markers detected by the system; not all markers are identified in the first preview frame. 54

Figure 28: Working-distance indicator showing detected markers relative to the optimal focal plane. A leftward bias means the component is too close to the scanner; a rightward bias means it is too far. 54

Figure 29: Result after the first Link’s acquisition. The preview shows the detected points already meshed and the markers numbered with unique identifiers. 55

Figure 30: Link after approximately 20 acquisitions. Where green points appear without an underlying hole, markers were removed during the procedure to acquire the surface beneath; this introduces a local reconstruction error, visible for example at the green spot near the right-hand hole. 56

Figure 31: Final mesh obtained after background-point removal and meshing in the GOM software. 56

Figure 32: ATOS deviation map of the Link against the nominal CAD model. The map visualises the local distance between the acquired mesh and the reference geometry, expressed in millimetres. 57

Figure 33: Views of the Phoenix v|tome|x s system. The left image shows the external cabinet; the right image shows the interior, with the X-ray emitter (yellow body) on the right, the

rotating specimen stage at the centre, and the detector positioned to the left, just outside the FOV (field of view). 58

Figure 34: Driven Gear Arm specimen wrapped in a low-density soft polymeric foam, used to hold it in position and to keep it raised above the metal stage so that the base is excluded from the reconstructed volume. 59

Figure 35: Radiograph of the Driven Gear Arm embedded in the low-density foam support, acquired with the settings listed in Table 5. 60

Figure 36: Grey-value histogram of the reconstructed CT volume. Three principal peaks are visible, corresponding to the background air, the low-density foam support, and the ASA bulk polymer, in order of increasing grey value. 62

Figure 37: Full view of the reconstructed Gear Arm after cleaning the grey-value histogram of the background and low-density material. 63

Figure 38: Sectioned view of the cleaned mesh, exposing the internal infill structure. Note the voids within the boundary walls and the imperfections and irregularities affecting the infill. 63

Figure 39: CT deviation map of the original full Gear Arm against the nominal CAD model. The map visualises the local distance between the acquired mesh and the reference geometry, expressed in millimetres. 64

Figure 40: Section view of the original full Gear Arm. The CT acquisition allows the internal infill structure of the part to be observed in detail. 65

Figure 41: Cavities within the material, caused by the non-ideally homogeneous layered structure of the 3D-printed part. With a search distance of 1,5 [mm] these cavities are considered by the comparison and bias the mean and standard deviation values. 65

Figure 42: Detail of the region near the teeth. Not only are internal cavities (in pink) detected, but also some of the faces belonging to the boundary between internal void and material. These contribute to biasing the data. 66

Figure 43: (a) CT deviation map of the emptied and repaired Gear Arm against the nominal CAD model; (b) A section of the emptied and repaired Gear Arm (observe the difference with the Figure 40). 67

Figure 44: Representative cross-section extracted with GOM Inspect from the CT volume, used as the geometric reference for the infill reconstruction. 74

Figure 45: Representation of the HSV colour space: hue as the angular coordinate, saturation as the radial distance from the central axis, and value as the height. 75

Figure 46: Infill walls isolated by HSV thresholding with the bounds of Table 10. 75

Figure 47: Segmented mask after the morphological closing performed with a 3 × 3 structuring element. 76

Figure 48: Processed cross-section applied as a reference image onto the original Frame CAD solid in SolidWorks. 76

Figure 49: Reconstructed Frame, section view showing the rebuilt infill. 77

Figure 50: Reconstructed Claw, section view showing the rebuilt infill. 77

Figure 51: Frame STEP model exported from SolidWorks with reversed surface normals, revealed by the inverted front and back face shading. 78

Figure 52: Erroneous surface comparison caused by the reversed normals, with the reconstruction matched against the wrong faces. 79

Figure 53: Frame STEP model after inverting the normals in GOM Inspect, with shading consistent with the intended infill geometry. 79

Figure 54: Surface comparison of the reconstructed Frame infill against the CT mesh (global best-fit alignment, 1,5 [mm] maximum search distance). 80

Figure 55: Section comparison of the Frame CT mesh against the reconstructed infill. 80

Figure 56: Spurious thin walls in the Frame CT data, produced by nozzle travel with residual attached material and identified as print defects not present in the reconstruction. 81

Figure 57: Surface comparison of the reconstructed Claw infill against the CT mesh. 81

Figure 58: Section comparison of the Claw CT mesh against the reconstructed infill. 82

Figure 59: Print defects identified in the Claw CT data. 82

Figure 60: (a) global coordinate reference system adopted for the frame component; (b) the 8 fixed boundary supports applied to the mounting holes, each restraining all six degrees of freedom. 84

Figure 61: Front (left) and side (right) views of the point-mass positions relative to the Frame: (1) Frame mass; (2) augmented payload mass (170 [g]); (3) servomotor mass. 85

Figure 62: The four connectors and their attachments: (1) the 170 [g] payload mass tied to the four servo mounting holes, the two Link-to-Frame holes and the slave-gear hole; (2) the driven-gear axis to the four servo mounting holes; (3) the slave-gear axis to the slave-gear hole; (4) the servo centroid to the four servo mounting holes. 86

Figure 63: The 9,81 [m/s²] gravitational acceleration acting on the three masses of the system along the −Z direction. 86

Figure 64: The 30 [m/s²] acceleration acting on the three masses of the system along the +X direction. 87

Figure 65: Servo-stall forces applied to the Frame, computed in Appendix G. 87

Figure 66: Inspire view of the Claw, showing the two revolute pairs at A and B and the reference frame centred at A. 88

Figure 67: Application on the Claw pad of the gravitational force due to half the object weight. 89

Figure 68: Application of the acceleration load along the X axis on the Claw pad. 90

Figure 69: Application on the Claw of the object's reaction to the gripping (closing) force, computed in Appendix G. 90

Figure 70: Reverse-engineered version of the Frame (left) and Claw (right), with the same loads and constraints applied as in the nominal models. 91

Figure 71: Overall envelope of the 3 cases on the solid Frame (CAD) in terms of SF. Min and Max SF are identified in the figure. 92

Figure 72: Overall envelope of the 3 cases on the solid Frame (CAD) in terms of displacement from the original position. Min and Max displacement are identified in the figure. 93

Figure 73: Overall envelope of the 3 cases on the solid Claw (CAD) in terms of SF. Min and Max SF are identified in the figure. 94

Figure 74: Overall envelope of the 3 cases on the solid Claw (CAD) in terms of displacement. Min and Max displacement are identified in the figure. 95

Figure 75: Overall envelope of the 3 cases on the reconstructed Frame (RE, with infill) in terms of SF. Min and Max SF are identified in the figure. 97

Figure 76: Overall envelope of the 3 cases on the reconstructed Frame (RE, with infill) in terms of displacement. Min and Max displacement are identified in the figure. 98

Figure 77: Overall envelope of the 3 cases on the reconstructed Claw (RE, with infill) in terms of SF. Min and Max SF are identified in the figure. 99

Figure 78: Overall envelope of the 3 cases on the reconstructed Claw (RE, with infill) in terms of displacement. Min and Max displacement are identified in the figure. 100

Figure 79: (left) original Claw part; (right) the same part with all non-functional holes filled, where the algorithm might remove material to topologically optimise the part. 104

Figure 80: Identification of the Non-Design Space regions (orange) to be excluded from the optimisation, at the functional attachment holes and the gripping surface, whose original shapes are preserved. 105

Figure 81: (left) Design Space (burgundy) and Non-Design Space (grey) of the Claw; (right) Design Space (burgundy) and Non-Design Space (grey) of the Frame. 106

Figure 82: Bonded contacts (blue) defined on the interfaces between Design Space and Non-Design Space; (left) Claw, (right) Frame, represented by hundreds of blue dots. 106

Figure 83: Raw topology-optimisation result of the Claw with the extrusion constraint. 107

Figure 84: Raw topology-optimisation result of the Frame with the extrusion constraint. ... 108

Figure 85: PolyNURBS fit on the optimised surface (grey) with the connected Non-Design Space parts (orange), Claw..... 108

Figure 86: The PolyNURBS part compared with the previously defined design space for the Claw. 109

Figure 87: PolyNURBS fit on the optimised surface (grey) with the connected Non-Design Space parts (orange), Frame. 109

Figure 88: The PolyNURBS part compared with the previously defined design space for the Frame. 110

Figure 89: Redesign of the Claw following the PolyNURBS prototype, ensuring full contact between the optimised volume and the remaining components. 111

Figure 90: Redesign of the Frame following the PolyNURBS prototype, ensuring full contact between the optimised volume and the remaining components. 112

Figure 91: Result envelope of the gravity, X-axis acceleration and servo-stall cases on the redesigned Claw. The minimum SF, 6,0, is reached under linear acceleration along the X axis. 113

Figure 92: Result envelope of the gravity, X-axis acceleration and servo-stall cases on the redesigned Frame. The minimum SF, 3,4, is reached under linear acceleration along the X axis, at a sharp unfilleted corner of the servo slot. 114

Figure 93: Result envelope of the gravity, X-axis acceleration and servo-stall cases on the redesigned Claw. The maximum displacement, 1,457 [mm], is reached under linear acceleration along the X axis. This displacement is measured relative to the anchoring holes

with the link and the gear, not with respect to the global reference system placed at the Frame–robot attachment. 115

Figure 94: Result envelope of the gravity, X-axis acceleration and servo-stall cases on the redesigned Frame. The maximum displacement, 1,112 [mm], is reached under linear acceleration along the X axis. 116

Figure 95: Comparison between the original and redesigned Frame, both printed in ASA. . 117

Figure 96: Comparison between the original and redesigned Claw, both printed in ASA. ... 117

Figure 97: The manipulator fully assembled with the redesigned and optimised parts. 125

Figure 98: Surface deviation map of the EinScan-derived Gear Arm mesh with respect to the nominal CAD model, obtained after pre-alignment and global best-fit registration in GOM Inspect. 136

Figure 99: Surface deviation map of the EinScan-derived Link mesh with respect to the nominal CAD model, obtained after pre-alignment and global best-fit registration in GOM Inspect. 137

Figure 100: Surface deviation map of the EinScan-derived Claw mesh with respect to the nominal CAD model, obtained after pre-alignment and global best-fit registration in GOM Inspect. 138

Figure 101: ATOS deviation map of the Gear Arm against the nominal CAD model. The map visualises the local distance between the acquired mesh and the reference geometry, expressed in millimetres..... 139

Figure 102: ATOS deviation map of the Frame against the nominal CAD model. The map visualises the local distance between the acquired mesh and the reference geometry, expressed in millimetres..... 140

Figure 103: ATOS deviation map of the Claw against the nominal CAD model. The map visualises the local distance between the acquired mesh and the reference geometry, expressed in millimetres..... 141

Figure 104: (left) one of the 1500 radiographs acquired for the Link and servo support; (right) the corresponding three-peak grey-value histogram. The three peaks correspond to the background (air), the low-density infill, and the solid material. 143

Figure 105: (a) full view of the Link (green) and servo support (blue) reconstructed after cleaning the grey-value histogram of the background and the low-density material; (b) view of the Link (green) and servo support (blue) mesh cleaned in VG Studio and sectioned to reveal the infill structure. 144

Figure 106: (a) CT deviation map of the original full Link against the nominal CAD model; (b) section view of the original full Link; (c) internal surfaces wrongly recognised and considered by the comparison. 146

Figure 107: (a) CT deviation map of the original full servo support against the nominal CAD model; (b) support section with some cavities marked, wrongly considered by the comparison. 147

Figure 108: (a) CT deviation map of the emptied and repaired Link against the nominal CAD model; (b) section of the emptied and repaired Link (note the difference with Figure 106-b). 148

Figure 109: (a) CT deviation map of the emptied and repaired servo support against the nominal CAD model; (b) section of the emptied and repaired servo support (note the difference with Figure 107-b). 149

Figure 110: (left) one of the 1500 radiographs acquired for the Claw; (right) the corresponding three-peak grey-value histogram. 150

Figure 111: (a) full view of the Claw reconstructed after cleaning the grey-value histogram of the background and the low-density material; (b) view of the Claw cleaned in VG Studio and sectioned to reveal the infill structure. 151

Figure 112: (a) CT deviation map of the original full Claw against the nominal CAD model; (b) section view of the original full Claw; (c) internal surfaces wrongly recognised and considered by the comparison. 153

Figure 113: (a) CT deviation map of the emptied and repaired Claw against the nominal CAD model; (b) section of the emptied and repaired Claw (note the difference with Figure 112-b). 154

Figure 114: (top) one of the 1500 radiographs acquired for the Frame; (bottom) the corresponding three-peak grey-value histogram. 155

Figure 115: (a) full view of the Frame reconstructed after cleaning the grey-value histogram of the background and the low-density material; (b) view of the Frame cleaned in VG Studio and sectioned to reveal the infill structure. 156

Figure 116: (a) CT deviation map of the original full Frame against the nominal CAD model; (b) section view of the original full Frame; (c) internal surfaces wrongly recognised and considered by the comparison. 158

Figure 117: (a) CT deviation map of the emptied and repaired Frame against the nominal CAD model; (b) section of the emptied and repaired Frame (note the difference with Figure 116-b). 159

Figure 118: Free-body diagram of the Claw in the planar reference system. F is the gripping force transmitted by the object; A and B are the revolute pairs connecting the Claw to the gear arm and to the Link; l, m and n are the lever arms used in the moment equilibrium about A. 174

Figure 119: Free-body diagram of the Link. The two revolute pairs B′ (to the Claw) and C (to the Frame) are aligned along the planar y axis at a distance z; with no other applied load the Link behaves as a two-force member. 175

Figure 120: Free-body diagram of the slave gear arm. The mesh forces Ft and Fr act at the pitch contact point P at a distance gr from the axis D; the Claw reaction acts at A′ at a distance gl from D. 176

Figure 121: Free-body diagram of the driven gear arm. The servo torque Ms acts about the axis E together with the mesh reaction at the pitch point and the mirrored Claw reaction at the arm hole. 177

Figure 122: Factor-of-safety distribution on the nominal CAD Frame under the gravity load case. 179

Figure 123: Displacement-magnitude contour of the nominal CAD Frame under the gravity load case; peak displacement $8{,}655 \times 10^{-3}$ [mm]. 180

Figure 124: Factor-of-safety distribution on the nominal CAD Frame under the acceleration-along-X load case. 181

Figure 125: Displacement-magnitude contour of the nominal CAD Frame under the acceleration-along-X load case; peak displacement 0,678 [mm], at the free end of the bracket. 182

Figure 126: Factor-of-safety distribution on the nominal CAD Frame under the servo-stall load case. 183

Figure 127: Displacement-magnitude contour of the nominal CAD Frame under the servo-stall load case; peak displacement 0,213 [mm]. 184

Figure 128: Factor-of-safety distribution on the nominal CAD Claw under the gravity load case. 185

Figure 129: Displacement-magnitude contour of the nominal CAD Claw under the gravity load case; peak displacement $2{,}928 \times 10^{-3}$ [mm]. 186

Figure 130: Factor-of-safety distribution on the nominal CAD Claw under the acceleration-along-X load case. 187

Figure 131: Displacement-magnitude contour of the nominal CAD Claw under the acceleration-along-X load case; peak displacement 0,197 [mm]. 188

Figure 132: Factor-of-safety distribution on the nominal CAD Claw under the servo-stall (gripping) load case. 189

Figure 133: Displacement-magnitude contour of the nominal CAD Claw under the servo-stall (gripping) load case; peak displacement 0,057 [mm]. 190

Figure 134: Factor-of-safety distribution on the reconstructed (CT) Frame under the gravity load case. 191

Figure 135: Displacement-magnitude contour of the reconstructed (CT) Frame under the gravity load case; peak displacement $1{,}057 \times 10^{-2}$ [mm]. 192

Figure 136: Factor-of-safety distribution on the reconstructed Frame under the acceleration-along-X load case. 193

Figure 137: Displacement-magnitude contour of the reconstructed Frame under the acceleration-along-X load case; peak displacement 0,821 [mm]. 194

Figure 138: Factor-of-safety distribution on the reconstructed Frame under the servo-stall load case. 195

Figure 139: Displacement-magnitude contour of the reconstructed Frame under the servo-stall load case; peak displacement 0,256 [mm]. 196

Figure 140: Factor-of-safety distribution on the reconstructed Claw under the gravity load case. 197

Figure 141: Displacement-magnitude contour of the reconstructed (CT) Claw under the gravity load case; peak displacement $3{,}766 \times 10^{-3}$ [mm]. 198

Figure 142: Factor-of-safety distribution on the reconstructed Claw under the acceleration-along-X load case. 199

Figure 143: Displacement-magnitude contour of the reconstructed Claw under the acceleration-along-X load case; peak displacement 0,233 [mm]. 200

Figure 144: Factor-of-safety distribution on the reconstructed Claw under the servo-stall (gripping) load case. 201

Figure 145: Displacement-magnitude contour of the reconstructed Claw under the servo-stall (gripping) load case; peak displacement 0,077 [mm]. 202

List of Tables

Table 1: Mass breakdown of the FDM-fabricated specimen, by subassembly. The 133 [g] reference value is the manufacturer's catalogue mass for the original injection-moulded gripper. 32

Table 2: Comparison of nominal mechanical and thermal properties of 3DJAKE ASA and 3DJAKE ABS filaments, as declared by the respective manufacturer datasheets. 33

Table 3: Maximum, minimum, mean and standard deviation of the surface-distance values shown in Figure 24, for the Frame mesh reconstructed in Meshmixer. The values were extracted after surface comparison in GOM Inspect with a maximum search distance of 1,5 [mm]. 51

Table 4: Maximum, minimum, mean and standard deviation of the surface-distance values shown in Figure 32, for the Link mesh reconstructed with the ATOS scanner and GOM Scan. 57

Table 5: Scan parameters for the Driven Gear Arm (Figure 35) acquisition on the Phoenix v|tome|x s. 61

Table 6: Comparison of the values for the original Driven Gear Arm and for the same part after applying the HPR algorithm and the GOM Inspect hole-closing routine. 68

Table 7: Global error indicators (mean deviation and pooled standard deviation) of the three scanning technologies across the four common components, on the outer surface. (*) The EinScan was acquired on the fully assembled end-effector, under conditions not homogeneous with the other two techniques, and its figures are a qualitative baseline only. 69

Table 8: Summary of the three digitisation techniques across the criteria relevant to the choice of dataset. 72

Table 9: Comparison of the two candidate reconstruction strategies. 73

Table 10: HSV threshold bounds applied to the cross-sectional image. 74

Table 11: ASA material card combining supplier and literature data. The supplier values were measured on standard specimens according to ASTM D790 and therefore describe the bulk polymer rather than the FDM printed material. 83

Table 12: Servo stall load set applied to the Frame, computed in Appendix G. The sign of the servo reaction torque TX was verified against the deformation pattern of a preliminary run. 88

Table 13: Actions transmitted by the gripped object and applied at the Claw pad through the surface connector. The Claw own mass is applied separately as a gravity and an acceleration body load. The gripping force is normal to the pad; the gravity and acceleration shares are tangential and are carried by friction. 90

Table 14: SF and peak displacement of the nominal CAD Frame under the three load cases. The minimum SF governs the verification, and the acceleration along X is the sizing case. .. 93

Table 15: SF and peak displacement of the nominal CAD Claw under the three load cases. The servo-stall (gripping) case yields the lowest SF and drives the topology optimisation of Chapter 6; under acceleration the true limit is the pad–object friction rather than the strength of the finger. 95

Table 16: CAD versus reconstructed (CT) comparison for the Frame, minimum SF and peak displacement under the three load cases, with the percentage deviation of the reconstructed

model relative to the CAD baseline. Displacement deviations are approximate, owing to the coarse read-out of the contour plots. 98

Table 17: CAD versus reconstructed (CT) comparison for the Claw, minimum SF and peak displacement under the three load cases, with the percentage deviation relative to the CAD baseline. The Claw is essentially insensitive to the infill, with deviations within a few per cent except under gravity. 100

Table 18: Governing minimum SF and maximum displacement of the Frame and Claw across the CAD, reverse-engineered (RE) and topologically optimised (TO) models. 116

Table 19: Mass of the Frame and Claw for the printed original and the TO printed part, with the percentage mass reduction with respect to the printed original. 118

Table 20: End-effector mass at each stage of the workflow, from the original commercial unit to the topology-optimised specimen. 118

Table 21: Engineering time by phase. 120

Table 22: Cost summary of the project under the company (A) and academic (B) scenarios. 121

Table 23: Maximum, minimum, mean and standard deviation of the surface-distance values shown in Figure 98, for the Gear Arm mesh. 137

Table 24: Maximum, minimum, mean and standard deviation of the surface-distance values shown in Figure 99 for the Link mesh. 137

Table 25: Maximum, minimum, mean and standard deviation of the surface-distance values shown in Figure 100 for the Claw mesh. 138

Table 26: Aggregated EinScan handheld-SLS error indicators (mean and std.) across the four components. See the introductory paragraphs of this appendix for the underlying assumptions. 138

Table 27: Maximum, minimum, mean and standard deviation of the surface-distance values shown in Figure 101, for the Gear Arm mesh. 139

Table 28: Maximum, minimum, mean and standard deviation of the surface-distance values shown in Figure 102 for the Frame mesh. 140

Table 29: Maximum, minimum, mean and standard deviation of the surface-distance values shown in Figure 103 for the Claw mesh. 141

Table 30: Aggregated GOM ATOS Core 300 fixed-head SLS error indicators (mean and std.) across the four components. See the introductory paragraphs of this appendix for the underlying assumptions. 141

Table 31: Values set on the Phoenix v|tome|x s machine for the acquisition of the Link and servo support. 143

Table 32: Comparison between the original Link and servo support and their state after application of the HPR algorithm and hole closing in GOM Inspect. 149

Table 33: Values set on the Phoenix v|tome|x s machine for the acquisition of the Claw. 150

Table 34: Comparison between the original Claw and its state after application of the HPR algorithm and hole closing in GOM Inspect. 154

Table 35: Values set on the Phoenix v|tome|x s machine for the acquisition of the Frame. 155

Table 36: Comparison between the original Frame and its state after application of the HPR algorithm and hole closing in GOM Inspect. 159

Table 37: Global CT indicators computed on the original full meshes and on the post-HPR, hole-treated meshes in equal weighting. 160

Table 38: Input data for the numerical evaluation of the servo stall load set. Lengths measured from the CAD model in the gripping configuration. 177

Table 39: Servo stall load set transmitted to the Frame, expressed in the global reference system of Section 5.1. The set coincides with Table 12 of Chapter 5. 178

Acronyms and Abbreviations

Acronym	Expansion
Geometry & CAD	
CAD	Computer-Aided Design
STL	Standard Tessellation Language
STEP	Standard for the Exchange of Product Data
PolyNURBS	Polygonal Non-Uniform Rational B-Spline
Scanning & Reverse Engineering	
RE	Reverse Engineering
CT	Computed Tomography
CMM	Coordinate Measuring Machine
SLS	Structured-Light Scanning
FOV	Field of View
HPR	Hidden Point Removal
HSV	Hue, Saturation, Value
Structural Analysis	
FEM	Finite Element Method
SF	Safety Factor
TO	Topology Optimisation
Manufacturing	
AM	Additive Manufacturing
FDM	Fused Deposition Modelling
CNC	Computer Numerical Control
DfAM	Design for Additive Manufacturing
ABS	Acrylonitrile Butadiene Styrene
ASA	Acrylonitrile Styrene Acrylate
Robotics	
DOF	Degree of Freedom
EE	End-Effector
General	
2D	Two-Dimensional
3D	Three-Dimensional
AI	Artificial Intelligence
SDG	Sustainable Development Goal

1. Introduction

1.1 Objectives and Scope

The continuous evolution of industrial automation and high-performance robotics has brought the optimisation of mechanical components to the forefront of modern engineering research. In this context, the reduction of the mass and rotational inertia of moving elements, and in particular of end-effectors (EE) mounted at the distal end of robotic manipulators, represents a primary objective for improving dynamic performance, reducing settling times, and enabling higher operational speeds without compromising structural integrity: even a modest reduction in EE mass yields measurable improvements in peak acceleration, cycle time, actuator torque demand, and energy consumption.

The component selected for study is the **ServoCity Servo-Driven Parallel Gripper Kit (made by ServoCity, SKU: 3219-0001-0002)** [1], see Figure 1.



Figure 1: Servo-Driven Parallel Gripper Kit by servocity.com as presented in the assembly file.

The quantitative objective of the thesis is to achieve a **mass reduction of approximately 10–30 [%]** on the main structural components of the EE, excluding the servo and fasteners, while maintaining structural integrity under all reference load cases.

To this end, an integrated physical-to-digital-to-physical workflow is adopted, combining 3D scanning and reverse engineering (RE), finite element method (FEM), topology optimisation (TO), and additive manufacturing (AM).

The effectiveness of analogous integrated workflows for lightweight robotic component redesign has been demonstrated in recent literature, with reported mass reductions in the range of 10–40 [%] for comparable gripper and robotic arm optimisation problems. [2] [3]

1.2 Workflow Overview and Thesis Structure

The present thesis follows a sequential workflow, Figure 2, each phase of which broadly corresponds to a distinct chapter. The logical progression of the work is summarised below, followed by a discussion of the principal limitations inherent to the methodology adopted.

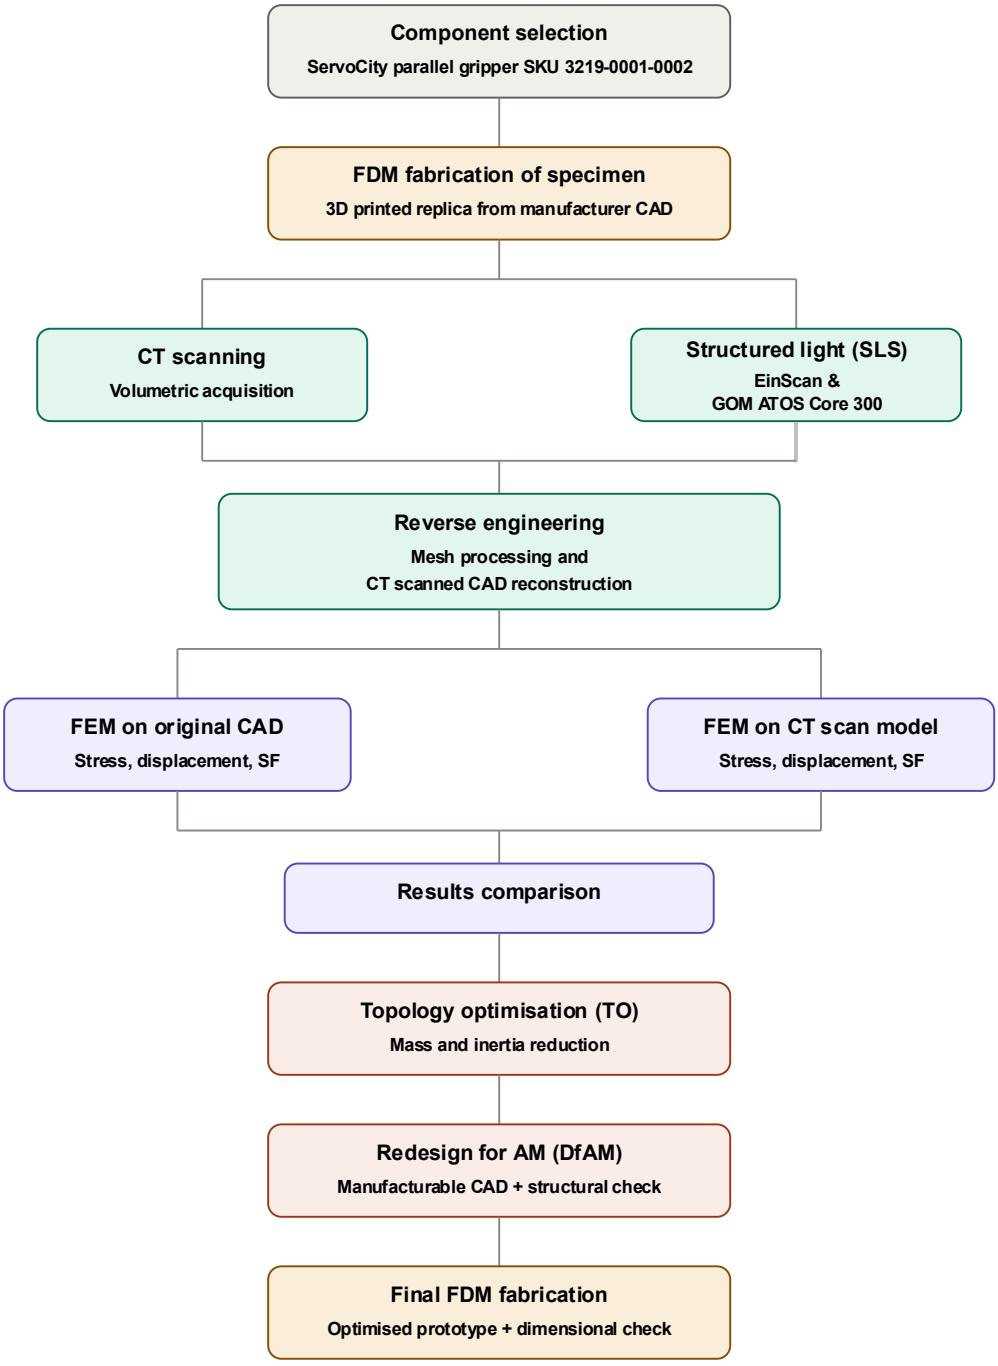

Figure 2: Thesis workflow diagram illustrating the full design cycle: from component selection and 3D printing, through 3D scanning and FEM analysis, to topology optimisation, redesign, and final reprint.

The component selected for study (introduced in Section 1.1 and fully characterised in Chapter 2) was not physically procured. Given that the original CAD model was publicly released by the manufacturer and that international shipping costs at the time of writing were

disproportionate relative to the component price, the decision was made to fabricate the specimen directly by **Fused Deposition Modelling (FDM)** using the available CAD data as the reference geometry. This approach allowed the reverse engineering workflow to be executed under controlled and reproducible conditions, while simultaneously serving as a proof of concept for the physical-to-digital-to-physical pipeline that constitutes the methodological core of this work.

Chapter 2 addresses the characterisation of the EE and all the steps between the initial idea and the 3D-printed model.

Chapter 3 addresses the **3D scanning** phase. The FDM-fabricated specimen is digitised using three complementary non-contact techniques: X-ray computed tomography (CT) on the Phoenix v|tome|x s system (Waygate Technologies / GE), structured-light scanning on the EinScan Pro 2X (Shining3D), and structured-light scanning on the GOM ATOS Core 300 (Carl Zeiss / GOM). The three techniques are applied in parallel to simulate the standard reverse engineering scenario, in which an existing physical component is acquired and digitally reconstructed, and to enable a direct methodological comparison in terms of achievable resolution, dimensional accuracy, acquisition time, and suitability for components of this scale and material composition. The data obtained from the three scanning techniques are first meshed within their respective acquisition software; the resulting meshes are then imported into GOM Inspect 2019 and Meshmixer/MeshLab, where results are corrected where necessary and, if possible, the geometry is remodelled as faithfully as possible to the original.

Chapter 4 presents the **reverse engineering (RE)** partial workflow. Only the CT-derived model is then used as the basis for the Frame and Claw CAD reconstruction in SolidWorks, since the EinScan and ATOS outputs do not resolve the internal cavities of the gripper**.** The reconstructed model is then compared dimensionally against the CT-derived pieces. The CT-derived pieces act as masters, due to the high precision of this scanning technique.

Chapter 5 presents the **FEM** structural analysis. The original CAD model is subjected to the reference load cases defined in Altair Inspire, characterising the stress distribution, displacement field, and safety factors (SF) under the nominal operating conditions. The analysis is then repeated on the reconstructed model of Chapter 4 that approximates the internal infill geometry observed in the CT scan. The comparison between the two FEM results quantifies the influence of the printed internal structure on the mechanical response and validates the use of the simplified solid model as input to the optimisation stage.

Chapter 6 describes the **topology optimisation (TO)** study, in which material is selectively redistributed within the design domain to minimise mass subject to the structural constraints established by the FEM analysis of the original model.

Chapter 7 then presents the redesign of the optimised components as manufacturable CAD models. The raw polygonal non-uniform rational B-spline (PolyNURBS) output of Altair Inspire is used as a geometric guideline: the Frame and the Claw are remodelled as single monolithic parts, removing the re-entrant fillets introduced by the automatic fit and locally increasing thicknesses, in accordance with design-for-additive-manufacturing (DfAM) principles. The chapter closes with the structural verification of the redesigned geometries, their reprinting, and the assessment of the achieved mass reduction.

Chapters 8 and 9 present, respectively, the **cost and time analysis** of the overall workflow and the **conclusions** of the work, including quantification of the performance improvement in terms of inertial load reduction under the reference scenarios and directions for future development.

1.3 Limitations

Several limitations of the present work deserve explicit acknowledgement.

First, the specimen subjected to scanning and reverse engineering is an **FDM-fabricated replica** of the commercial component rather than a unit procured directly from the manufacturer. As a consequence, the mechanical properties of the scanned part, in particular its anisotropic stiffness and layer-dependent strength characteristics, differ from those of the original, which may have been produced by injection moulding or an alternative polymer processing route. While this does not significantly affect the geometric fidelity of the RE workflow, it introduces a degree of uncertainty in the FEM analysis if the material model does not fully account for the FDM-induced anisotropy.

Second, the performance improvement assessment is conducted entirely through **numerical simulation** and does not include experimental dynamic validation on a physical robotic installation.

Third, the component selection was initially constrained by the **operational requirements of the Zeiss Xradia 620 Versa, which imposed strict restrictions on specimen size and material density. Following the switch to the Phoenix v|tome|x s system**, these constraints were partially relaxed; the selection is nonetheless retained as it reflects the conditions under which the component was originally chosen. This necessarily limited the selection pool to small, non-metallic components, excluding a broader class of industrial end-effectors that would otherwise have been more representative of high-performance robotic applications.

1.4 Alignment with UN SDGs

The 2030 Agenda for Sustainable Development, adopted by all United Nations Member States in 2015, establishes seventeen Sustainable Development Goals (SDGs) as a universal call to action to end poverty, protect the planet, and ensure prosperity for all. Within the academic and engineering domains, aligning research activities with the SDGs has become an increasingly recognised standard of responsible practice. The present thesis intersects with three SDGs, as discussed below.



SDG 9 (Industry, Innovation and Infrastructure) calls for sustainable industrialisation and the fostering of technological innovation. This thesis directly embodies these principles through the adoption of an advanced digital-to-physical workflow, integrating CT scanning, RE, FEM, TO, and AM, that represents precisely the type of innovation-driven industrial methodology SDG 9 promotes. The application of additive manufacturing as the primary fabrication route further supports SDG 9's emphasis on sustainable industrial infrastructure, enabling near-net-shape production and the realisation of geometrically complex structures that would be unachievable through conventional subtractive methods.



SDG 12 (Responsible Consumption and Production) calls for the sustainable management of natural resources and the reduction of waste throughout production cycles. Topology optimisation, as the methodological core of this work, is fundamentally an exercise in material efficiency: achieving equivalent or superior structural performance with a reduced material volume. Furthermore, additive manufacturing generates significantly less waste than computer numerical control (CNC) subtractive machining, where material is removed from a solid billet, a shift that is directly aligned with SDG 12's principles of responsible consumption.



SDG 13 (Climate Action) calls for urgent action to reduce anthropogenic greenhouse gas emissions across all sectors of the economy. The alignment of this thesis with SDG 13 is indirect but substantive: the mass and inertia reduction achieved through TO directly reduces the torque demand on the servo actuator, lowering the electrical energy consumed per operational cycle. Over the lifetime of an industrial robotic system, which may involve millions of actuation cycles, this cumulative energy saving can translate into a measurable reduction in carbon emissions, the magnitude of which depends on the energy mix of the operating context.

2. Component Selection and Characterisation

2.1 Application Context

Industrial robotics has undergone a profound transformation over the past three decades, driven by the increasing demand for faster, more precise, and more energy-efficient automated systems. Among the various robotic architectures developed for manipulation tasks, two main families stand out in high-performance applications: **serial manipulators** and **parallel manipulators**.

Serial manipulators, characterised by an open kinematic chain in which each link is connected to the next in sequence, offer large workspaces and high dexterity, but are inherently limited in dynamic performance by the cumulative inertia of their structure. Each joint must move the payload and also accelerate and decelerate all downstream links, including, critically, the EE mounted at the tip of the chain. This propagation of inertia through the kinematic structure means that even a modest mass reduction at the distal end yields disproportionate benefits in terms of reduced joint torques, improved acceleration capacity, and lower energy consumption [4] [5]. Figure 3 shows a typical industrial context for such considerations, with articulated arms operating inside an automated manufacturing cell.



Figure 3: KUKA industrial robotic arms operating within an automated manufacturing cell.

Parallel manipulators, and particularly Delta robots, were developed to address precisely this limitation. Introduced by Reymond Clavel in the 1980s [6], and subsequently protected by US patent [7], the Delta architecture connects the moving platform, to which the EE is attached, to a fixed base through multiple independent kinematic chains acting in parallel. This configuration keeps all actuators at the base, dramatically reducing the moving mass and enabling peak accelerations exceeding **100 [m/s²]** in industrial implementations [8]. Delta robots are today widely deployed in pick-and-place operations in the food, pharmaceutical, and electronics industries, where cycle times on the order of tenths of a second are standard

requirements. An example is the delta robot of Figure 4, designed precisely for high-speed pick-and-place operations.



Figure 4: ABB FlexPicker delta robot designed for high-speed industrial pick-and-place operations.

In both architectures, the EE is the terminal element of the kinematic chain and the component most directly in contact with the manipulated object. Its mass and inertial properties directly affect the peak achievable acceleration, the torque demand on the actuators, the mechanical stress on transmission elements, and the structural vibration response of the system. In high-speed applications, the maximum achievable acceleration is bounded not only by actuator torque limits but also by the structural integrity of the components under inertial loading; lighter EE designs experience lower internal stresses during rapid direction reversals, reducing fatigue risk and improving settling time, a key throughput metric in automated cells. This is the context in which the present thesis is positioned.

2.2 Selection Criteria and Component Description

The selection of the component to be studied was driven by two early constraints: the characteristics of the CT scanning equipment initially earmarked for this work and the time budget available for the experimental campaign. The CT system originally planned was a Zeiss Xradia 620 Versa (by the Universitat Politècnica de València), whose maximum admissible specimen envelope and X-ray penetration depth effectively restrict the analysable specimens to small, non-metallic parts: metallic components, particularly steel or aluminium alloys, would attenuate the beam beyond the threshold required to recover internal geometry at the resolutions needed for RE. These constraints, combined with the limited time window available for scanning, fabrication and analysis, oriented the selection from the outset towards a small, polymer-based component.

During the work the CT system actually used was switched to the Phoenix v|tome|x s (Waygate Technologies / GE, by the Politecnico di Torino), discussed in detail in Chapter 3, based on laboratory availability. The change relaxed some of the size and resolution constraints in practice, but the small, polymer-based scope of the component had already been adopted as a precondition of the workflow and was retained for consistency and simplicity of the subsequent reverse-engineering, FEM and TO pipeline.

Within this small, polymer-based scope, the component selection was further refined by three additional engineering criteria:

1. **Functional relevance**: the component must constitute a genuine EE intended for use on a robotic manipulator, so that the mass and inertia reduction objectives of the thesis are directly applicable to a real operational context.
2. **Geometric complexity**: the component must present sufficient geometric complexity and volume to make the scanning, RE, FEM, and TO workflow non-trivial and methodologically meaningful. Simple prismatic or axisymmetric geometries would not adequately exercise the digital-to-physical pipeline.
3. **CAD availability**: the availability of a public nominal CAD model from the manufacturer was considered a significant practical advantage. Given the time and resource constraints of the present work, access to the nominal geometry would have allowed the RE accuracy assessment (Chapter 4) to be performed by direct dimensional comparison against a known reference, rather than requiring full manual reconstruction from scratch.

The component selected based on these criteria is the **ServoCity Servo-Driven Parallel Gripper Kit** [1], Figure 5, a commercially available servo-actuated parallel gripper. The selection satisfies all three criteria: the gripper is a genuine robotic EE, its assembly presents a complex multi-body geometry including interlocking gear elements and four-bar linkage mechanisms, and the nominal CAD model is publicly released by the manufacturer.

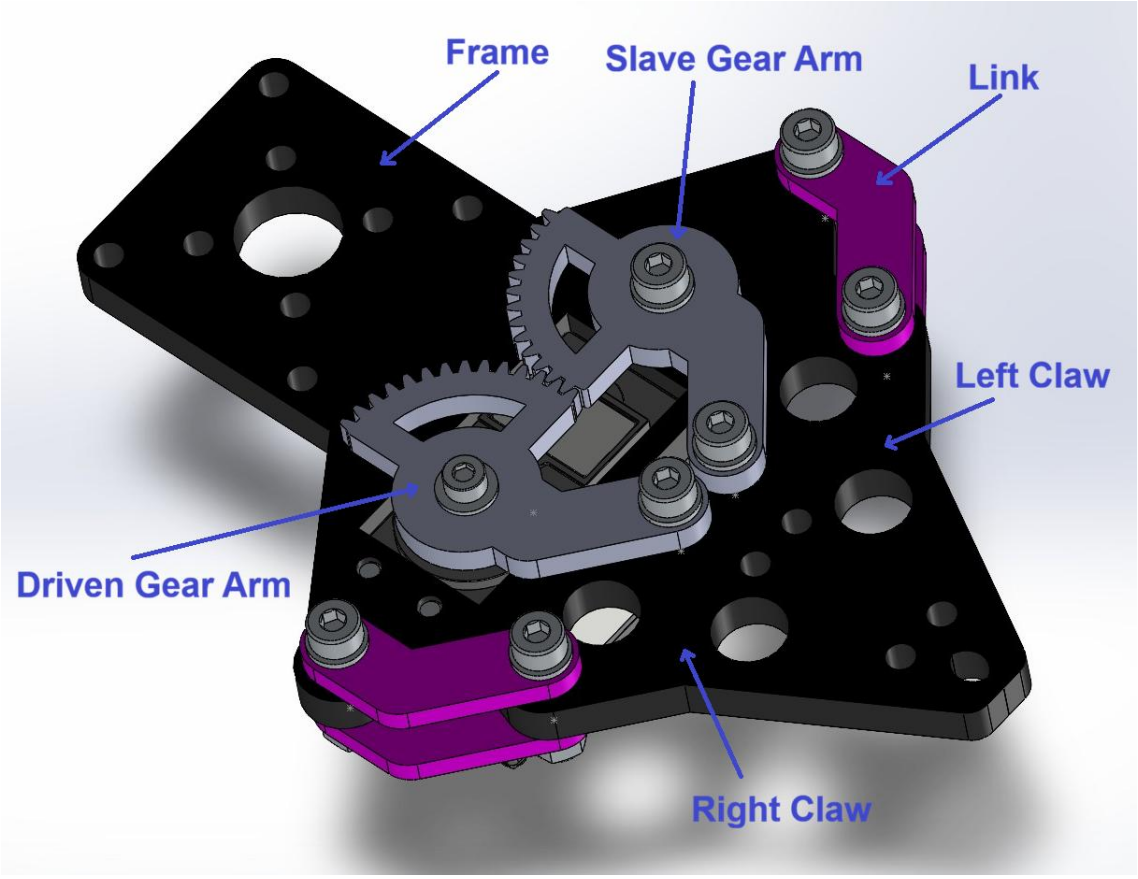

Figure 5: Original end-effector CAD visualised in SolidWorks.

2.3 Geometric, Functional and Material Characteristics

The EE is a compact parallel gripper in which the grasping motion of two opposing fingers is driven by a single rotary servo actuator. The mechanical architecture is based on a symmetric four-bar linkage arrangement: the rotation of the servo output shaft is transmitted to both gripper fingers simultaneously through a gear-and-linkage system, ensuring that the two fingers move in parallel and in equal and opposite directions, maintaining a constant orientation of the gripped object throughout the opening and closing stroke.

The assembly consists of the following primary structural components:

- **Frame:** the main structural body, providing the mounting interface for the servo and the pivot points for the linkage system.
- **Left/right Claws:** the two grasping fingers, each articulated through two linkage elements.
- **Links:** the connecting rods forming the four-bar mechanism.
- **Driven/Slave Gear Arms:** the gear elements transmitting servo rotation to finger displacement.
- **Servo Supports:** mechanical standoffs for servo mounting.
- **Servo Motor:** the rotary actuator providing the input torque to the gear-and-linkage transmission.
- **Screws and Fasteners:** M3 and M4 socket head screws, locknuts and washers that complete the assembly.

The original commercial gripper, as supplied by ServoCity, is manufactured by injection moulding in two polymer grades declared by the supplier: acrylonitrile butadiene styrene (ABS) for the main structural bodies and acetal plastic for the gear-arm and linkage elements, and is assembled with a standard servo and hardware kit. Its nominal mass, as reported in the manufacturer's catalogue, is 133 [g].

The specimen analysed in the present work is not a commercial unit, but an FDM-fabricated replica produced from the publicly available CAD data. It differs from the catalogue unit in three respects: the polymer parts are FDM-printed in ASA by 3DJAKE.com on an Ultimaker S3 rather than injection-moulded in ABS / acetal plastic; the original servo actuator (Proton Servo, SKU: 2002-0180-0002 [1]) has been substituted with a different model (Robbe Futaba S3003); and the fastener set has been adapted accordingly. The total assembled mass of the specimen therefore differs from the catalogue value. The mass of each subassembly, as measured on the FDM-fabricated specimen, is summarised in the table below.

Component	Quantity × Mass [g]	Role in optimisation
Frame	1 × 15,6	Design domain (to be optimised)
Claws	2 × 5,2	Design domain (to be optimised)
Links, Gear Arms, Servo Spacers	1 × 10,3 (total)	Non-design FDM parts (retained as printed)
Total printed polymer mass (assembled)	36,3	Sum of all FDM-printed parts
Servo motor	37,5	Non-design (excluded from TO)
Screws, locknuts, washers	27,6	Non-design (excluded from TO)
Total assembled specimen mass	101,4	Full as-built gripper

Table 1: Mass breakdown of the FDM-fabricated specimen, by subassembly. The 133 [g] reference value is the manufacturer's catalogue mass for the original injection-moulded gripper.

As can already be noted at this stage, the act of fabricating the end-effector in-house has yielded an immediate mass reduction: the total assembled specimen weighs 101,4 [g] against the 133 [g] catalogue value of the original commercial unit, a decrease of approximately 23,76 [%]. This **mass reduction** is attributable to a combination of factors: primarily the 20 [%] infill density and the generally lower density of ASA relative to the injection-moulded ABS and acetal parts specified by the manufacturer, together with the substitution of the original 55 [g] Proton Servo with the 37,5 [g] Robbe Futaba S3003 [9], dictated mainly by availability. The servo swap, however, entails a trade-off in actuator performance: the replacement unit delivers a maximum stall torque of **402 [N·mm] at 6,0 [V]**, compared to the 530 [N·mm] rated output of the original Proton servo at the same voltage, representing a reduction of approximately 24,1 [%] in available torque. This trade-off is considered acceptable within the scope of the present work, since the primary objective is to demonstrate the validity of the reverse-engineering and topology-optimisation methodology rather than to optimise the actuator selection.

2.4 Printing and Assembly

All FDM prints carried out in the present work, namely the **original specimen** used for 3D scanning and reverse engineering, and the **post-optimisation redesigns** subsequently fabricated for prototype validation, were produced on a single **Ultimaker S3** printer using the same **ASA** filament supplied **by 3DJAKE**. The use of a single printer and a single material batch across the whole experimental campaign was a deliberate methodological choice, aimed at minimising the number of independent variables when comparing the structural performance of the original geometry and of the topologically optimised redesign.

The selection of ASA over the more conventional ABS was driven by three considerations. First, the **mechanical performance of ASA** is at least equivalent to that of ABS in terms of stiffness and impact resistance, and is in fact superior in tensile strength, as summarised in Table 2; the substitution therefore does not introduce any penalty for the structural validation objectives of the prototypes and, by virtue of the higher tensile strength of ASA, marginally improves the load-bearing capacity of the printed components. Second, ASA offers a **noticeably more stable printing process**: its tendency to warp and to delaminate between layers is considerably reduced relative to ABS, which translates into better first-layer adhesion, lower part-to-part variability, and improved dimensional repeatability across successive prints, an essential requirement when geometric comparison between the original and redesigned components is part of the methodology. Third, ASA does not release **styrene fumes** during extrusion, which is a relevant practical advantage in a university laboratory environment where dedicated ventilation systems are not always available.

The nominal mechanical and thermal properties of the two filaments, as declared by the manufacturer, are summarised in Table 2 for direct comparison. All ASA values originally reported in [kgf/cm²] in the manufacturer datasheet have been converted to [MPa] to allow consistency with the ABS datasheet, which is reported in SI units. [10][11]

Property	Symbol / [Unit]	3DJAKE ASA	3DJAKE ABS	Test method
Physical properties				
Specific gravity	ρ [g/cm³]	1,10	1,06	ISO 1183 / ASTM D792
Melt flow rate (MFR)	[g/10 min]	15	21	ISO 1133 / ASTM D1238
Tensile strength	σ_t [MPa]	78,5	49	ASTM D790 / ASTM D638
Tensile modulus	E [MPa]	2 168	2 350	ASTM D790 / ASTM D638
Mould shrinkage	[%]	0,5	n/a	ISO 294-4
Thermal properties				
Printing temperature	[°C]	210 to 250	220 to 250	(manufacturer)
VICAT softening temperature (A)	[°C]	95	92	ISO 306 / ASTM D792
Recommended bed temperature	[°C]	60 to 100	100 to 110	(manufacturer)

Table 2: Comparison of nominal mechanical and thermal properties of 3DJAKE ASA and 3DJAKE ABS filaments, as declared by the respective manufacturer datasheets.

A direct inspection of Table 2 confirms the observations reported above. All values are adopted as declared in the manufacturers' technical datasheets [10], [11]; the ASA figures, originally reported in kgf/cm², have been converted to [MPa] (1 [kgf/cm²] = 0,0980665 [MPa]) for consistency with the ABS datasheet. They are carried unchanged into the material model used for the structural analyses of Chapter 5. ASA exhibits a tensile strength approximately **60 [%] higher** than that of ABS (78,5 against 49 [MPa]), while the tensile modulus is comparable (2 168 against 2 350 [MPa]); the slightly higher VICAT softening temperature of ASA (95 against 92 [°C]) further confirms its thermal robustness in the operating range relevant for the present

application. The lower MFR of ASA (15 against 21 [g/10 min]) is consistent with a more viscous melt during extrusion, which contributes to the reduced warping tendency observed in practice. Overall, the comparison supports the choice of ASA as a structurally superior alternative to ABS in tensile strength, and a broadly equivalent one in stiffness, while additionally offering practical advantages in terms of process stability and laboratory health and safety.

The main slicing parameters and process choices, identical for all prints produced in this work, are summarised below:

- **Layer height**: **0,15 [mm]**, with **8 top/bottom layers** and **3 wall perimeters**, chosen to ensure a sufficient solid shell thickness in regions subject to stress concentration.
- **Infill**: grid pattern at **20 [%] density**, selected as a compromise between print time, material consumption, and structural stiffness, and kept constant across all prints to ensure comparability of the resulting parts.
- **Nozzle temperature**: **245 [°C]**, within the upper range of the manufacturer-recommended interval (210 to 250 [°C]), to guarantee a stable and continuous material flow at the printing speeds adopted.
- **Build plate temperature**: **95 [°C]**, with the printer enclosure closed throughout the print, to ensure adequate adhesion of the first layer and to prevent corner lifting or detachment in regions of high residual thermal stress during cooling.
- **Build plate preparation**: thorough cleaning of the glass build plate with a spatula and acetone before each print, followed by **application of a thin uniform layer of adhesion spray** (lacquer), to ensure consistent first-layer bonding and to eliminate variability associated with surface contamination between successive prints.

The same printer, filament batch, slicing profile, and operator (myself) were used for every part produced in this work, so that any differences observed between the original specimen and the optimised redesigns can be attributed primarily to the geometry itself rather than to fabrication variability. The implications of infill-dependent effective stiffness on the FEM modelling are discussed in Chapter 5, while the cross-reference between the present section and the AM workflow proper is provided in Chapter 7. The process is documented in Figure 6 to Figure 8. Figure 6 shows the Ultimaker Cura slicing preview of the frame, in which the 20 [%] grid infill and the solid perimeter shells adopted for all prints can be appreciated. Figure 7 shows the first-layer deposition of the two gear arm components on the build plate, whose adhesion quality reflects the 95 [°C] build-plate temperature and the lacquer surface preparation; Figure 8 is a still frame captured during the frame print, in which the same 20 [%] grid infill is clearly visible in the interior of the part before the closure of the top layers.

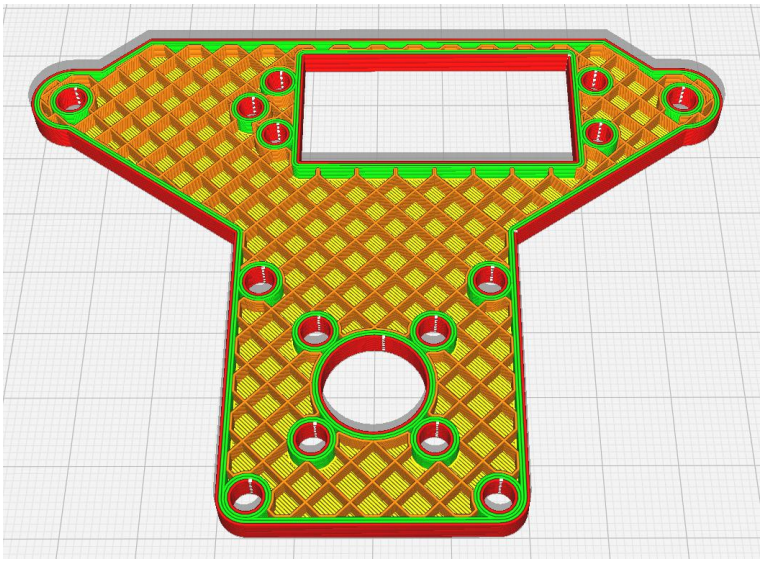

Figure 6: Ultimaker Cura slicing preview of the frame component, showing the 20 [%] grid infill pattern and solid perimeter shells adopted for all prints in this work.

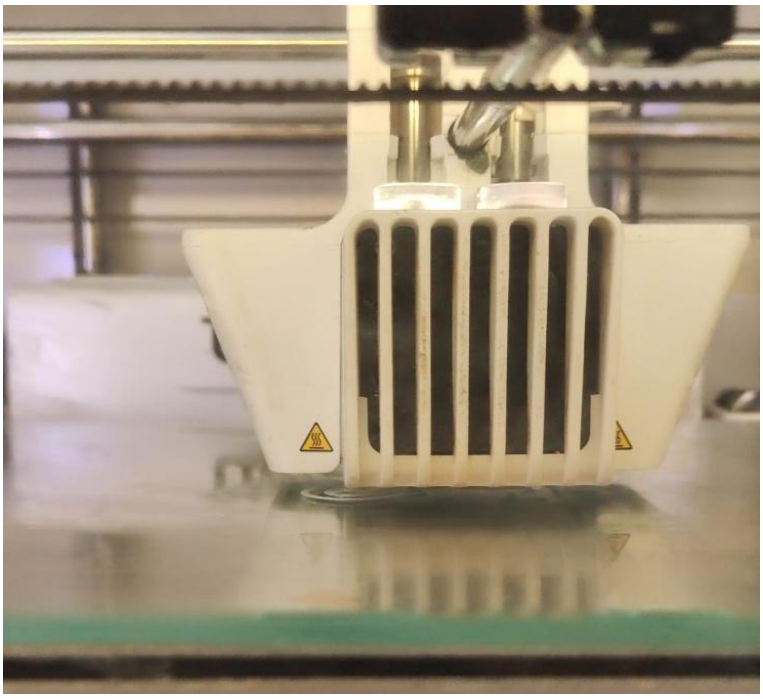

Figure 7: First-layer deposition of the two gear arm components on the Ultimaker S3 build plate, illustrating the adhesion quality achieved with the 95 [°C] bed temperature and adhesion-spray surface preparation.

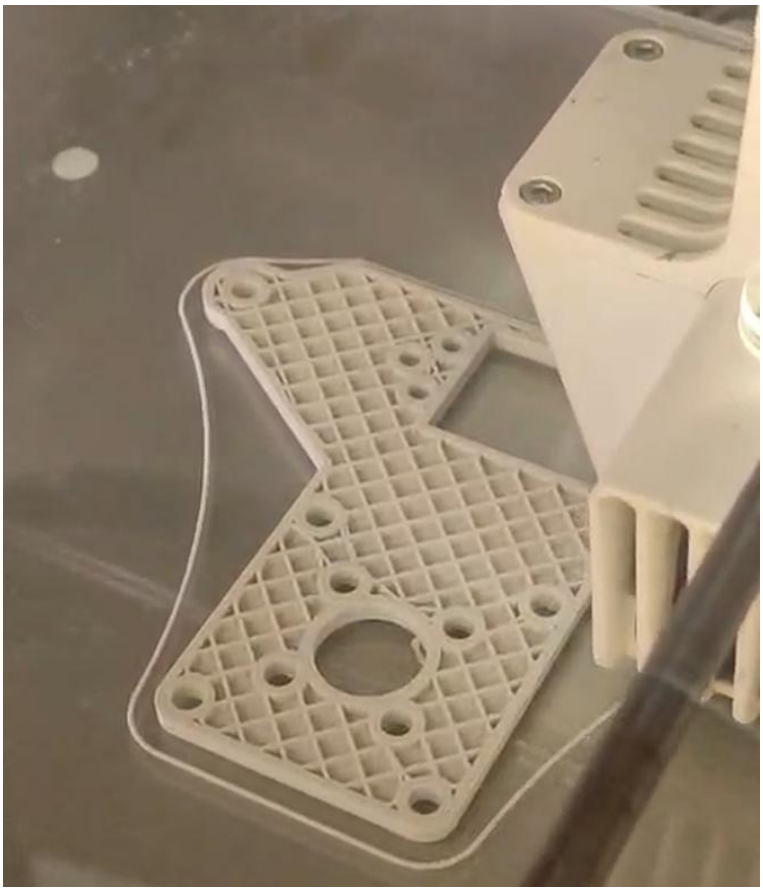

Figure 8: Still frame extracted from a video recording of the frame print in progress. The 20 [%] grid infill pattern is clearly visible in the interior; minor surface irregularities attributable to nozzle movement are also apparent but do not affect the structural performance of the component.

2.5 Normal Working Conditions

The components targeted for structural analysis and topology optimisation are the frame and the two claw fingers, identified as the main interesting elements of the assembly. To drive the optimisation toward a physically meaningful design, the loading scenarios summarised below have been defined; their detailed implementation in the FEM environment, including boundary conditions, contact definitions and mesh setup, is deferred to Chapter 5.

The end-effector is considered in the configuration represented in Figure 9, with the claw tips pointing downward and the interface to the robot wrist located at the upper face of the frame. This orientation has been chosen because it is the configuration most representative of the pick-and-place operations for which the gripper is intended, where the payload is collected from below and transferred to a target position. In the reference frame adopted throughout, the X–Y plane coincides with the horizontal working plane of the manipulator and the Z axis is aligned with the vertical travel direction of the end-effector; all forces, accelerations and displacement constraints defined in Chapter 5 are expressed in this coordinate system.

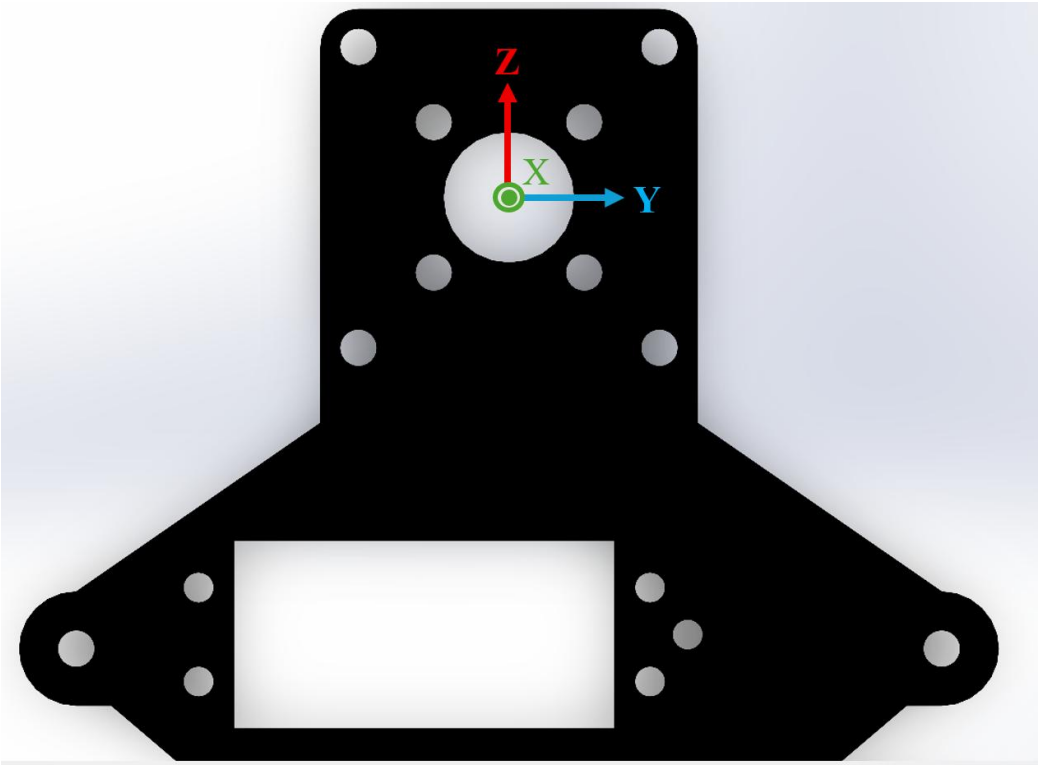

Figure 9: Reference frame adopted for the structural analysis and load-case definition.

The geometric configuration adopted for the analysis corresponds to the fully extended position of the kinematic chain, with both the guide linkages and the gear-links aligned vertically and the claws in the open state. This configuration maximises the lever arm of any external load applied at the claw tips with respect to the frame mounting points and therefore represents the worst case for both bending and torsion of the load-bearing components, Figure 10. Selecting this configuration as the reference for the FEM and TO ensures that the optimised geometry retains adequate margin throughout the operating range.

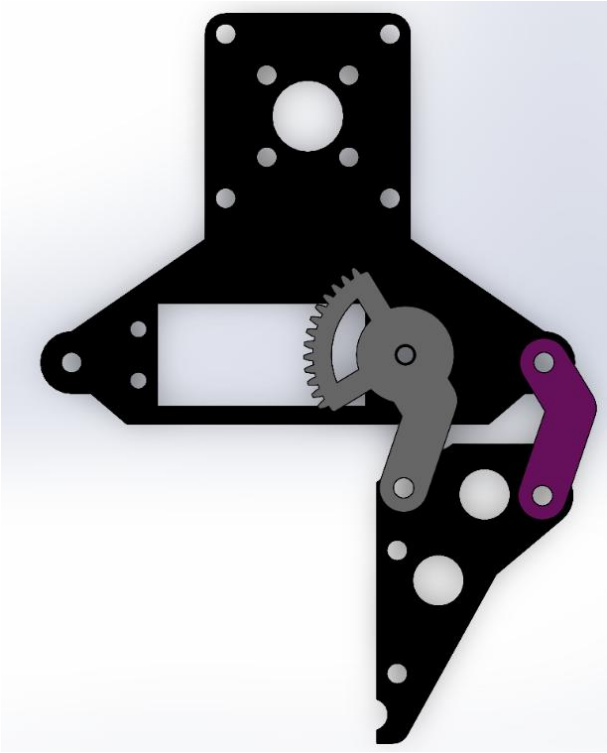

Figure 10: End-effector in the worst-case geometric configuration adopted for the structural analysis: fully extended kinematic chain. This configuration maximises the bending moment arm at the frame mounting points and is therefore used as the reference for all FEM and topology optimisation runs.

A representative payload of 150 [g], modelled as a rigid incompressible cube whose centre of mass coincides in the Z axis with the gripper tips, Figure 11, has been adopted as the reference object held by the end-effector. The value reflects the typical mass of the small components handled in high-speed sorting applications, while the rigid-cube idealisation eliminates the need to characterise contact stiffness or deformation of the gripped object, which fall outside the scope of this work.

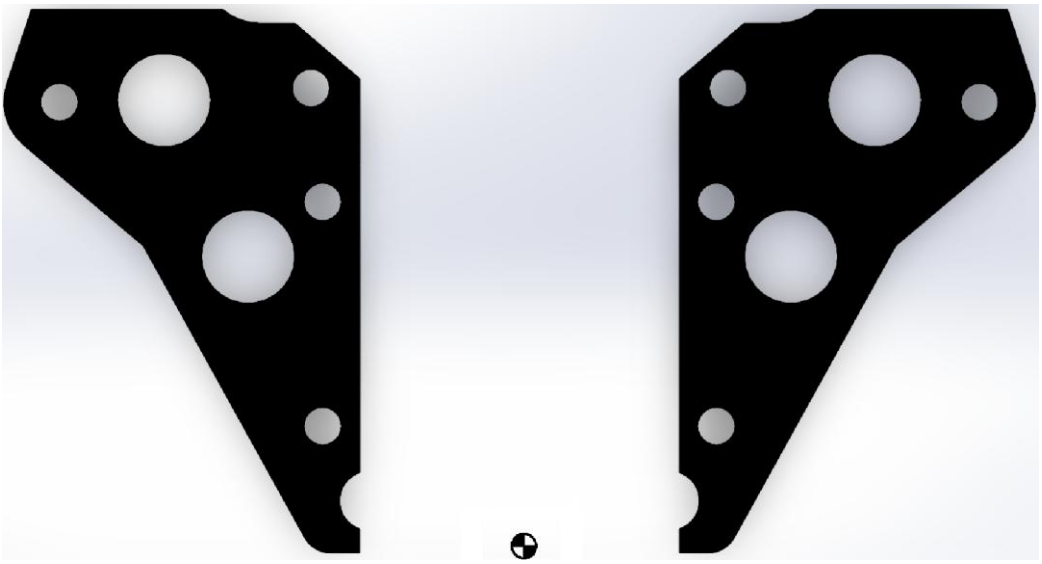

Figure 11: Position of the centre of mass of the representative 150 [g] incompressible payload relative to the two claw tips. For the worst-case analysis, the contact reaction force from the object is assumed to act at the foremost point of the finger profile.

Four load cases have been defined for the structural analysis:

1. **Gravity.** The gravitational field is applied to all masses of the system, namely the frame, the two claws, the servomotor (represented as a concentrated mass at its estimated centre of gravity) and the payload at the claw tips. This load case quantifies the stress state induced by self-weight alone and serves as a baseline against which the relative magnitude of the other loads is assessed.
2. **Servo stall.** When the gripper holds an incompressible object the actuation chain becomes kinematically blocked, and the servomotor reaches its stall torque while no rotation of the gear pair occurs. Under this condition the mechanism transmits the maximum torque the servo can produce, which is conservatively assumed to act statically on the structure. **Two distinct contributions** reach the frame and are modelled accordingly: the **reaction at the gear mesh**, decomposed into a **tangential and a radial component through the gear pressure angle**, which loads the bearing pins of the Servo and Driven Gear Arms; and the **reaction torque of the servo stator**, which is reacted by the four mounting holes of the servo on the frame.
3. **Operational acceleration.** The end-effector is mounted on a high-speed manipulator, and during typical pick-and-place trajectories the wrist is subjected to translational accelerations whose peaks can reach approximately 3 [*g*] (g denoting the gravitational acceleration). A uniform linear acceleration of **30 [m/s²]** (3,06 times the gravity) is therefore applied normal to the plane of the frame, simulating the inertial loading experienced by all the masses of the system during the acceleration phase of a working cycle. This direction has been selected because it produces the largest bending of the frame about its weakest axis and is consequently the most demanding orientation among those that the wrist commonly experiences.
4. **Linear combination.** The three previous load cases are superimposed by linear combination, under the assumption of linear-elastic material behaviour, to identify the worst-case stress state that the structure must withstand during operation. This combined scenario also makes it possible to rank the individual contributions and verify which load dominates the design, an aspect that is anticipated qualitatively in this section and confirmed quantitatively in Chapter 5.

Two simplifications have been adopted at this stage of the analysis. The mass of bolts, washers, spacers and other small fasteners has been treated in a simplified but conservative manner: the combined mass of the claws and of the fasteners belonging to the moving finger assemblies (approximately 20 [g] in total) has been lumped at the centre of mass of the payload rather than distributed at the actual positions of the individual components, while the remaining hardware, located close to the frame mounting points and therefore acting on a negligible lever arm, is not added to this lumped term. Consequently, all load cases involving acceleration or gravity acting on the payload (Gravity and operational acceleration) operate on an effective payload mass of 170 [g] rather than the nominal 150 [g]. This lumping strategy is justified on two grounds: it is the simplest possible treatment that avoids the need to characterise the individual mass and position of each fastener, and it places the lumped masses at the point of maximum lever arm with respect to the frame supports, producing bending loads more severe than those of the physically distributed masses and thereby introducing a conservative margin into the analysis rather than an unconservative approximation. The centre of mass of the servomotor has been placed at the geometric centroid of its bounding rectangular parallelepiped, coinciding with the physical centre of the servo body. The internal layout of a commercial hobby servo is not publicly documented in sufficient detail to allow a rigorous estimation of the true CoM position; the servo is therefore treated as a black box whose mass is assumed to be uniformly distributed

within its outer envelope. The centroid of the parallelepiped is consequently the best estimate available without disassembly of the unit and is expected to be a reasonable approximation of the actual CoM for the purposes of the present structural analysis.

Concerning the strength criterion, a minimum safety factor of 2,0 has been imposed throughout the topology optimisation. The value has been selected based on common engineering practice for FDM polymer parts subjected to quasi-static loads, accounting for the well-documented anisotropy of the printed material, the moderate scatter of its mechanical properties between print runs, and the absence of detailed fatigue characterisation for the ASA grade adopted. Any deviation from this threshold during the subsequent stages of the work is documented and justified in the corresponding chapter.

3. 3D Scanning

3.1 Overview and Technology Selection

The acquisition of three-dimensional geometric data from physical objects has undergone significant development over the past two decades, giving rise to a wide range of technologies that differ substantially in operating principle, achievable resolution, measurement uncertainty, cost, and field of application. In the context of RE and quality control, the selection of the most appropriate scanning technique is a critical methodological decision, as it directly conditions the accuracy of the reconstructed CAD model and the feasibility of all subsequent processing steps.

Most popular and available 3D measurement technologies can be organised according to a two-level hierarchical classification, illustrated schematically in the figure below.

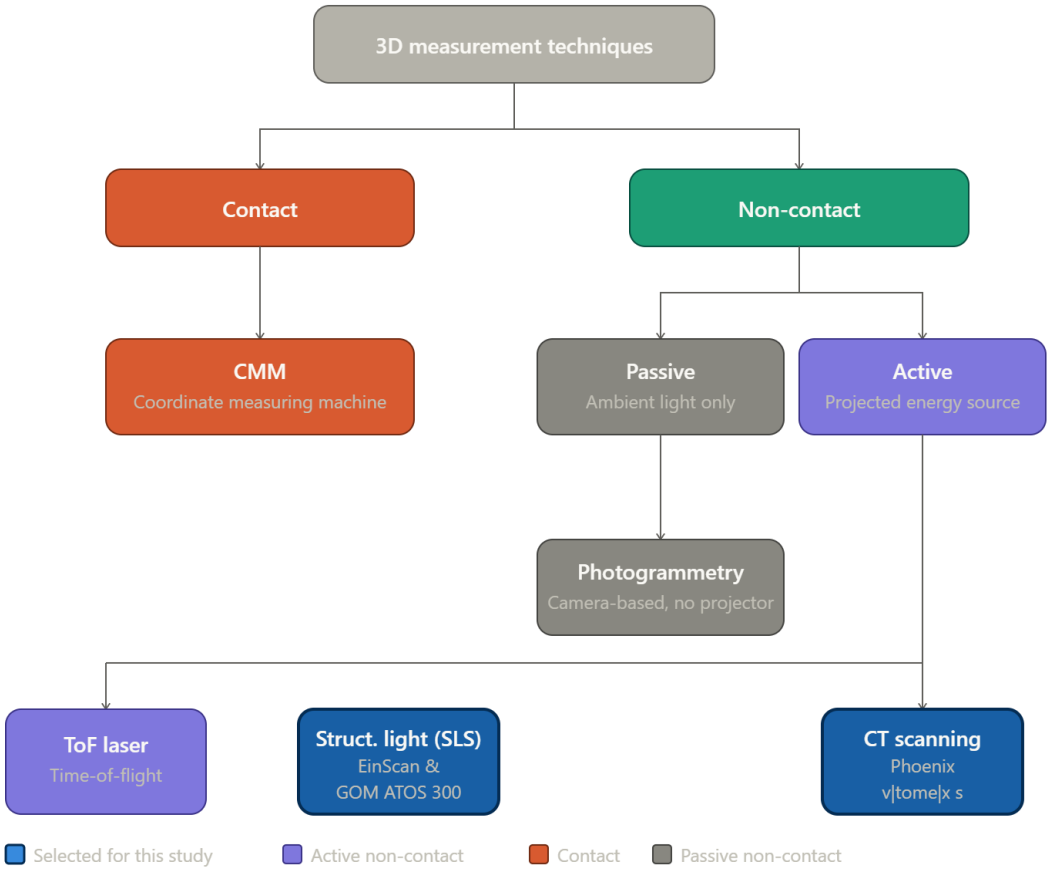

At the first level, a fundamental distinction is drawn between **contact** and **non-contact** methods.

Contact methods, of which the Coordinate Measuring Machine (CMM) is the most representative industrial example, operate by physically probing the object surface point by point using a calibrated stylus. While CMMs offer measurement uncertainties in the sub-micrometre range under controlled laboratory conditions, their applicability is inherently limited by three factors: point-by-point acquisition is time-intensive; the technique is unsuitable for compliant or geometrically complex surfaces; and it provides no information whatsoever on internal geometry. For a component such as the one under study, which presents freeform gripper-claw profiles, small-radius fillets, and a gear tooth interface, contact methods are impractical as a primary digitisation strategy and are therefore excluded from further

consideration. The five-axis CMM shown in Figure 12, available at the Dimensional Metrology Laboratory of the UPV, was not employed in this work: it is included here only as a representative example of contact-based measurement systems.



Figure 12: LK Altera Type S 10.7.6 five-axis coordinate measuring machine (CMM) used for precision dimensional inspection, shown here as a representative example of contact-based measurement methods. Dimensional Metrology Laboratory, UPV.

Non-contact methods overcome these limitations by extracting geometric information without any physical interaction with the object surface. They form the primary focus of this chapter and are further classified at the second level of the hierarchy.

At the second level, non-contact techniques are subdivided into **passive** and **active** systems, according to whether they rely on an externally projected energy source or on ambient illumination alone.

Passive systems, most notably photogrammetry and stereo vision, reconstruct object geometry by processing multiple photographs taken from different viewpoints, Figure 13, exploiting feature-matching and triangulation algorithms to infer three-dimensional coordinates from two-dimensional image correspondences. While photogrammetry can achieve millimetre-level accuracy on large, well-textured objects and requires no specialised hardware beyond a calibrated camera, its performance degrades significantly on surfaces with low texture, specular or near-specular reflectance, or fine geometric detail. Two of these characteristics are present in the FDM-fabricated specimen under study, which exhibits a uniform light-grey ASA finish with regular layer lines but limited natural texture and a slightly glossy, partially specular appearance. Passive methods are therefore considered unsuitable for the present application.



Figure 13: Application of photogrammetry techniques for the digital acquisition of architectural heritage.

Active systems address these limitations by projecting a controlled energy source, typically structured visible light, laser (by triangulation or time-of-flight), or penetrating X-ray radiation, onto the object and analysing the interaction between that source and the material to extract geometric information. Within this category, three families of techniques are most relevant to industrial RE at the scale of the present component:

- **Laser triangulation i**nfers point-to-sensor distances from the angular offset of a laser spot projected onto the surface and imaged by a CCD/CMOS detector held at a fixed baseline from the emitter. A Roland PICZA LPX-60 laser-triangulation scanner (Figure 14) was physically available at the host institution but was not adopted in this work. The unit is an older-generation, now-discontinued model whose resolution, acquisition speed, and post-processing software have been widely surpassed in current reverse-engineering practice; other techniques were therefore preferred.



Figure 14: Roland PICZA LPX-60 laser-triangulation 3D scanner that captures object geometry by projecting a laser beam onto the surface and triangulating its reflected position, generating a precise point cloud.

- **Structured-light scanning (SLS)** projects a series of known fringe or coded light patterns onto the object surface and reconstructs geometry by measuring the deformation of those patterns as seen by one or more calibrated cameras, through triangulation. SLS, as shown in Figure 15, is today one of the most widely adopted non-contact digitisation techniques in industrial reverse engineering and quality control, valued for its combination of sub-millimetre accuracy and rapid acquisition. However, SLS, like the other optical techniques already presented, is an inherently **surface technique**: it reconstructs only the external envelope of the object and is blind to any internal features, cavities, or material inhomogeneities.



Figure 15: Shining3D EinScan Pro 2X 2020 handheld structured-light 3D scanner that projects fringe patterns onto the object surface, computing geometry through phase-shift analysis and SLAM to generate a high-accuracy point cloud.

- **X-ray computed tomography (CT)** exploits the differential attenuation of X-ray radiation through matter to reconstruct a full volumetric representation of the object, including both external geometry and internal features such as wall thicknesses, voids, embedded components, and sub-surface defects, Figure 16. Unlike all optical methods, CT is not limited to the object surface: it provides a complete three-dimensional density map from which any cross-section or isosurface can be extracted. This volumetric capability is particularly relevant in the context of FDM-fabricated specimens, where internal infill geometry, layer boundaries, and potential void defects may depart from the nominal CAD model in ways that are entirely invisible to surface-based techniques.



Figure 16: GE Phoenix v|tome|x s: an industrial X-ray computed tomography (CT) system that reconstructs both internal and external geometry of components through volumetric radiographic acquisition, without any contact with the part.

Based on the discussion above, **SLS and CT are adopted as the two digitisation techniques for this study**. The two methods are deliberately complementary in their information content: SLS reconstructs the external envelope of the specimen with sub-millimetre accuracy in a few seconds per scan position, while CT provides a full volumetric representation that includes the internal infill geometry, the layer boundaries, and any sub-surface defects of the FDM-printed component, none of which are accessible to surface-based techniques. Their combined use is also consistent with the size and material composition of the component and is made possible by the availability of both instruments at the host institution. Beyond these practical considerations, the side-by-side application of SLS and CT to the same specimen offers a direct methodological comparison between the two routes, in terms of achievable resolution, dimensional fidelity, and processing complexity, which constitutes one of the original methodological contributions of this chapter.

3.2 Handheld Structured-Light Scanning

The first 3D digitisation technique adopted in this work was handheld SLS, performed using an EinScan Pro 2X (2020 edition, Shining3D) operated through the proprietary EXScan Pro software, Figure 17. The instrument couples a structured-light projector with two cameras arranged at a fixed angular baseline; the manufacturer declares a nominal scan accuracy of 0,045 [mm] and a minimum point spacing of 0,2 [mm], the finest resolution available in the software, which was adopted for all acquisitions in this work. [12]

The EinScan Pro 2X was selected as the entry point of the digitisation workflow based on two converging factors: its immediate availability at the laboratory of the institution (UPV), and the opportunity it offered to develop operator familiarity with handheld SLS acquisition before progressing to higher-accuracy instruments. The handheld acquisition was therefore conceived as a methodological demonstration rather than as a source of geometry for the downstream reverse-engineering workflow: it provided a first-order digital reconstruction of the assembly and a qualitative baseline against which the subsequent ATOS Core 300 and X-ray CT acquisitions could be evaluated.



Figure 17: Capturing high-density point clouds using the EinScan Pro 2X 2020 handheld 3D scanner in hybrid scanning mode.

The acquisition was performed on the **fully assembled EE**, mounted at the claw tips on a flat work surface using modelling clay as a temporary fixture to prevent unintended displacement during scanning, Figure 18. Two surface preparation steps were applied prior to acquisition. First, retro-reflective circular markers, consisting of black dots with concentric white rings, were affixed at geometrically distributed positions across the assembly to provide stable reference points for the SLAM-based tracking algorithm. Given the relatively small dimensions of the structural components, the physical size of the available markers proved disproportionate with respect to the local surface features, a factor that subsequently introduced reconstruction

artefacts in the final mesh, as discussed below. Second, a thin uniform coat of stain remover spray (*quitamanchas*) was selectively applied to the servo motor body and to other surfaces presenting near-zero reflectance, creating a diffuse white matte finish that improved the camera response on regions that would otherwise have returned insufficient light to the sensor. The coating was reversible and did not affect the dimensional integrity of the polymer structural components.

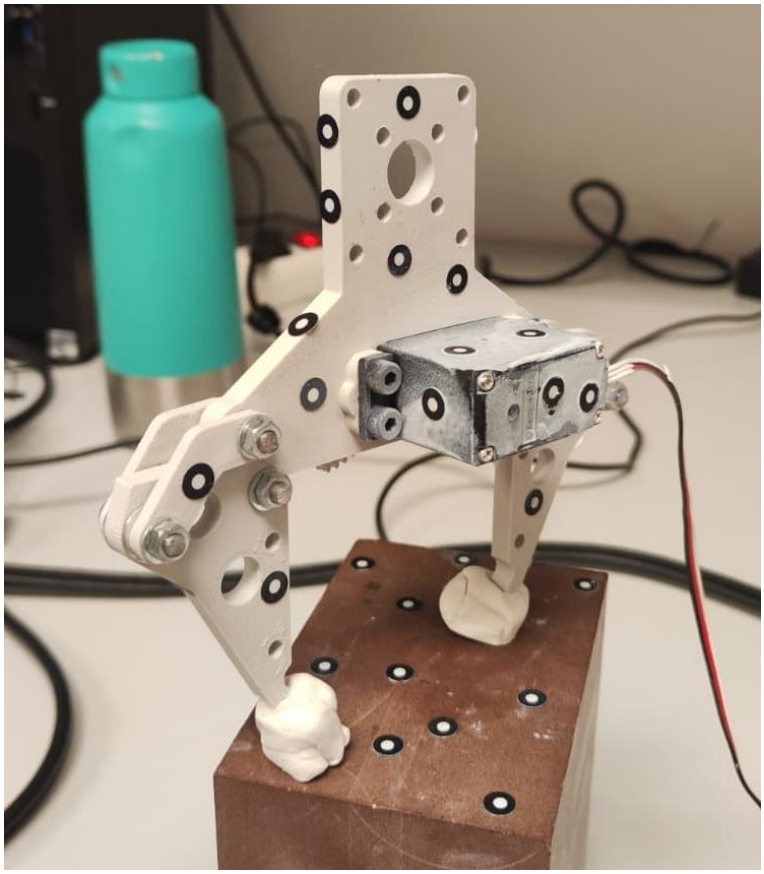

Figure 18: Mounting stand with markers in place and gripper held open, with stain-remover spray (quitamanchas) applied to enable acquisition of the dark surfaces of the servo motor.

The scanning trajectory was iterative. The scanner was held freehand and moved continuously around the stationary specimen, with real-time visual feedback provided by the EXScan Pro preview window allowing the operator to monitor point cloud density and adjust scanner distance, orientation, brightness and focus dynamically during the acquisition. Multiple passes were performed, alternating between acquisition of new points on poorly covered regions and deletion of noisy or redundant data, until the preview indicated satisfactory surface coverage. The underside of the assembly, including the claw tips, was acquired in a separate pass after physically reorienting the specimen, a step made possible by the marker-based coordinate system, which allowed the two partial scans to be registered into a single unified point cloud without loss of spatial reference. The complete acquisition session, including setup, surface preparation, four full scan attempts and a primary iterative refinement of the point cloud, required approximately **6 hours of effective scanning time** before a satisfactory dataset was obtained, as shown in Figure 19, Figure 20 and Figure 21.

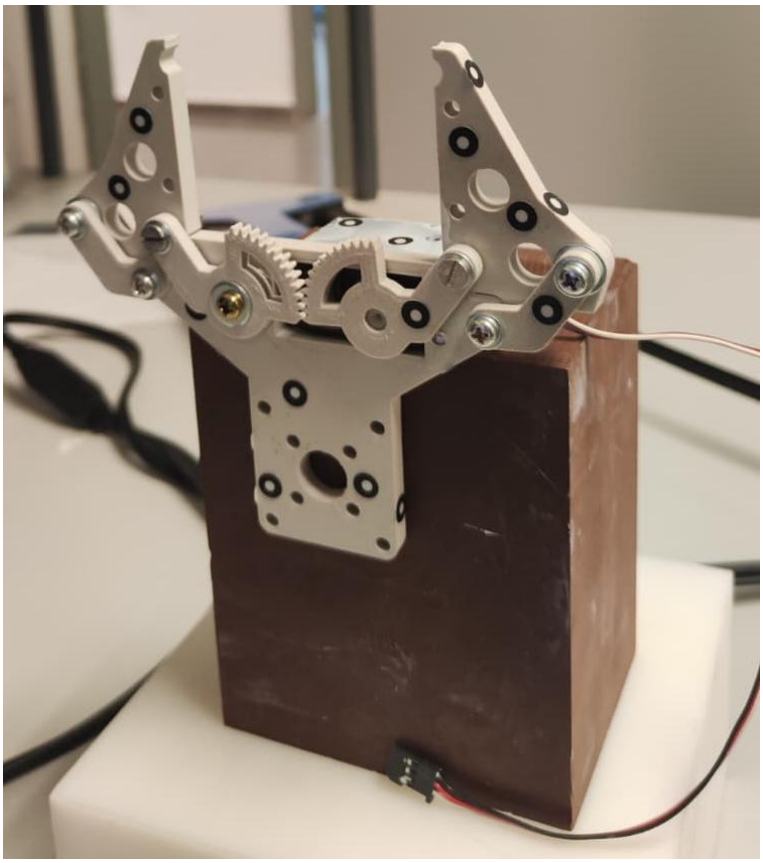

Figure 19: Assembly flipped on the stand with the base markers removed, allowing acquisition of the claw tips that had previously been covered by the modelling clay used as a temporary fixture.

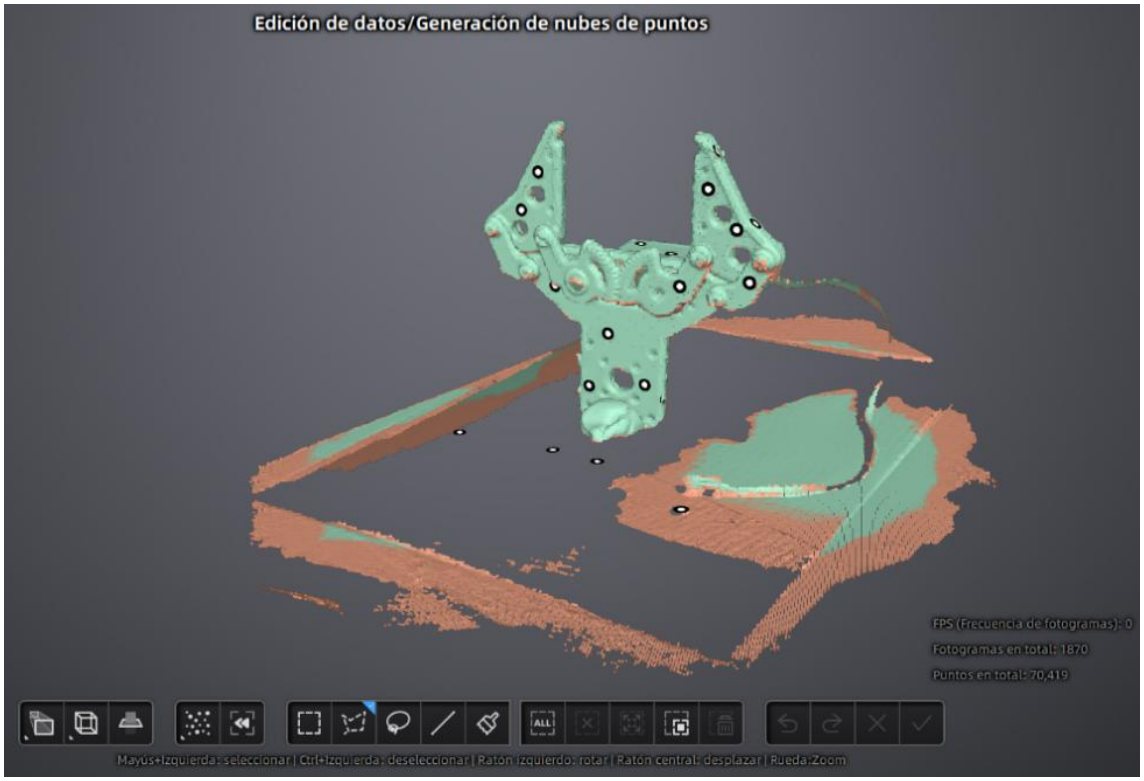

Figure 20: EXScan Pro preview of the acquired point cloud, prior to the removal of out-of-interest regions. The preview provides an effective qualitative assessment of the reconstruction quality before proceeding to the meshing stage.

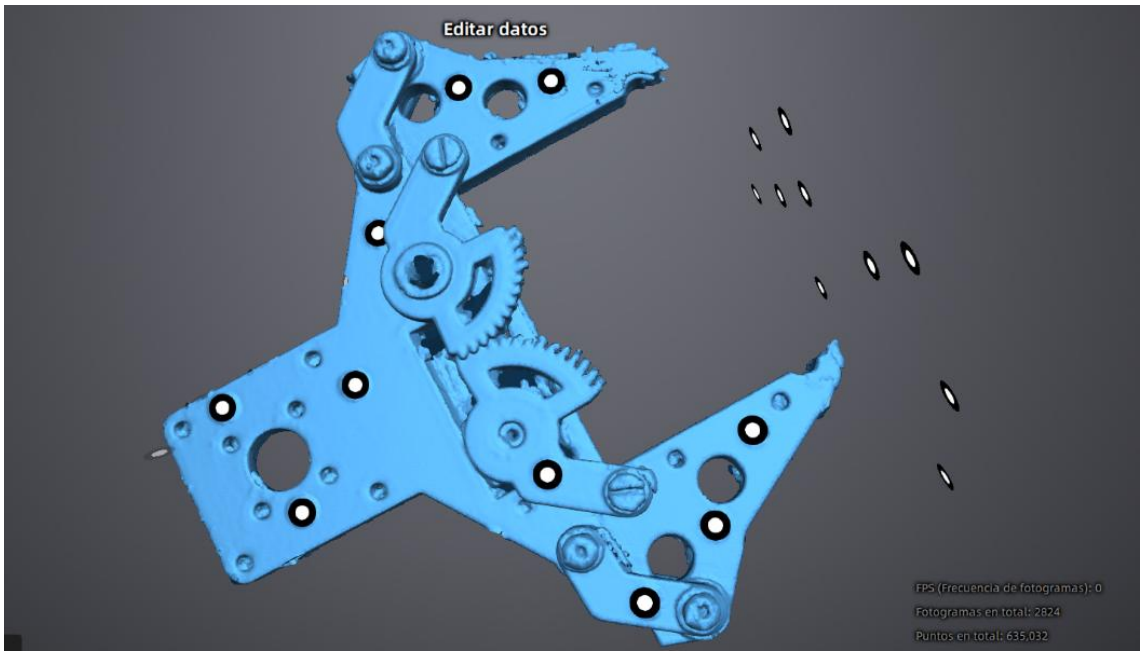

Figure 21: First mesh obtained after removal of the redundant points. Hole-filling is then performed, on the gaps left by the markers.

The consolidated point cloud was processed within EXScan Pro through automatic meshing, followed by hole-filling and mesh simplification, and exported as a Standard Tessellation Language (STL) file. The regions most affected by marker footprints were subsequently subjected to manual repair in Autodesk Meshmixer to remove the localised surface bumps, as illustrated in Figure 22 and Figure 23. For the individual structural components (Claws, Links, Driven Gear Arm), the corresponding manual repair stage required on the order of five hours of work per component before the mesh reached a quality compatible with downstream use; the Frame mesh discussed in the remainder of this section is presented in its Meshmixer-repaired form. The combination of long acquisition times and lengthy mesh-repair effort already qualifies handheld SLS, in this specific application, as a feasible but disproportionately time-consuming digitisation route.

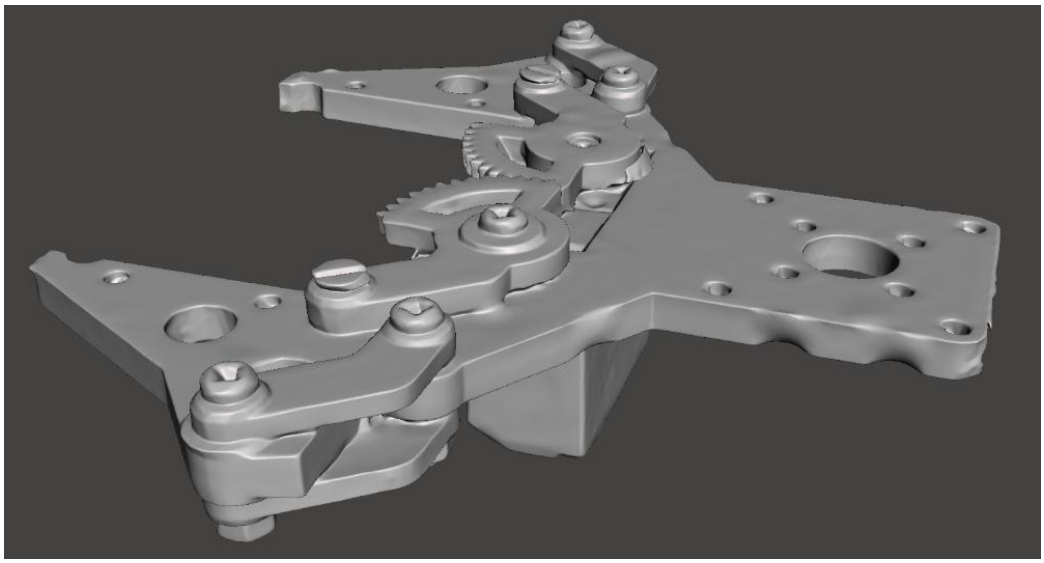

Figure 22: STL mesh of the complete end-effector assembly imported into Autodesk Meshmixer. Geometric defects are visible across the assembly, concentrated particularly on the Frame component.

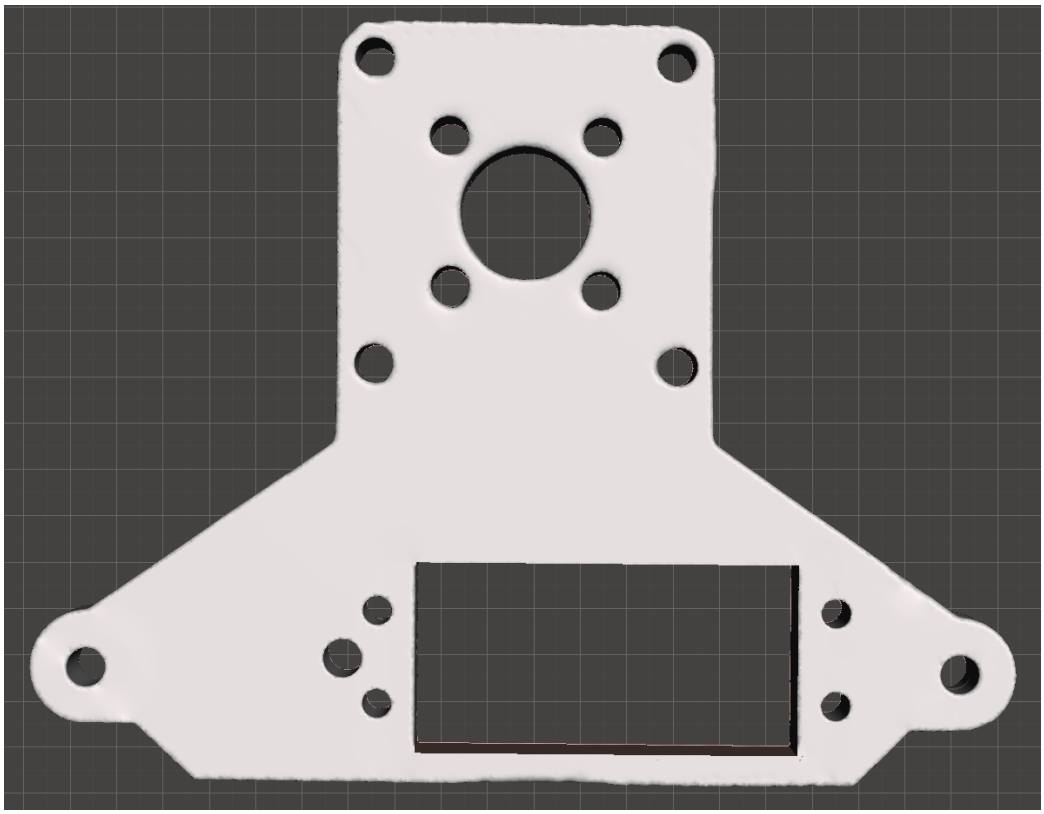

Figure 23: Isolated and repaired Frame mesh obtained by applying Boolean and trimming operations to the raw EinScan point cloud, used as the input for the GOM Inspect dimensional comparison.

The final Meshmixer-repaired mesh was imported into GOM Inspect for quantitative dimensional comparison with the nominal CAD model. Only the Frame is presented here as a representative case; a full component-by-component breakdown of the EinScan results is reported in Appendix B, and a cross-scanner comparison of mean deviation values is provided in Section 3.5. The EinScan acquisition was performed on the fully assembled end-effector, whereas the subsequent ATOS and CT campaigns described in Sections 3.3 and 3.4 were carried out on the individual disassembled components; the datasets are therefore not directly comparable without accounting for this difference in acquisition geometry, and any cross-scanner comparison must be interpreted with that caveat in mind. Within those constraints, the EinScan Frame serves as a meaningful qualitative baseline for the magnitude of deviations obtainable with handheld SLS at this scale of application.

The deviation map of the EinScan-derived Frame mesh against the nominal CAD model is shown in Figure 24. Following pre-alignment and global best-fit registration in GOM Inspect, the surface comparison was performed adopting a maximum search distance of 1,5 [mm], a value applied consistently across all three scanning technologies considered in this thesis (EinScan, ATOS, CT) to ensure methodological homogeneity of the comparison; the rationale behind this choice is detailed in Appendix B. These results are collected in Table 3, which reports the extreme values, mean and standard deviation of the measured surface distances: the mean dimensional deviation across the analysed surface was −0,02 [mm], with a **standard deviation of 0,15 [mm]**. The **maximum positive deviation** recorded was **+1,47 [mm]**, localised at a reconstructed hole whose geometry was partially compromised by marker proximity during acquisition. The **maximum negative deviation** was **−1,44 [mm]**, corresponding to a mesh artefact (a degenerate triangle attributed to a surface patch inconsistent with the nominal CAD geometry) rather than a true geometric deviation of the component. It should also be noted that the parameters used when exporting the reference CAD model to STL

(triangle size, number of triangles, chord error) influence such point-wise extremes, as discussed further in Section 3.5. For the purposes of this assessment, the mean and standard deviation are therefore of greater interest than the maximum and minimum values.

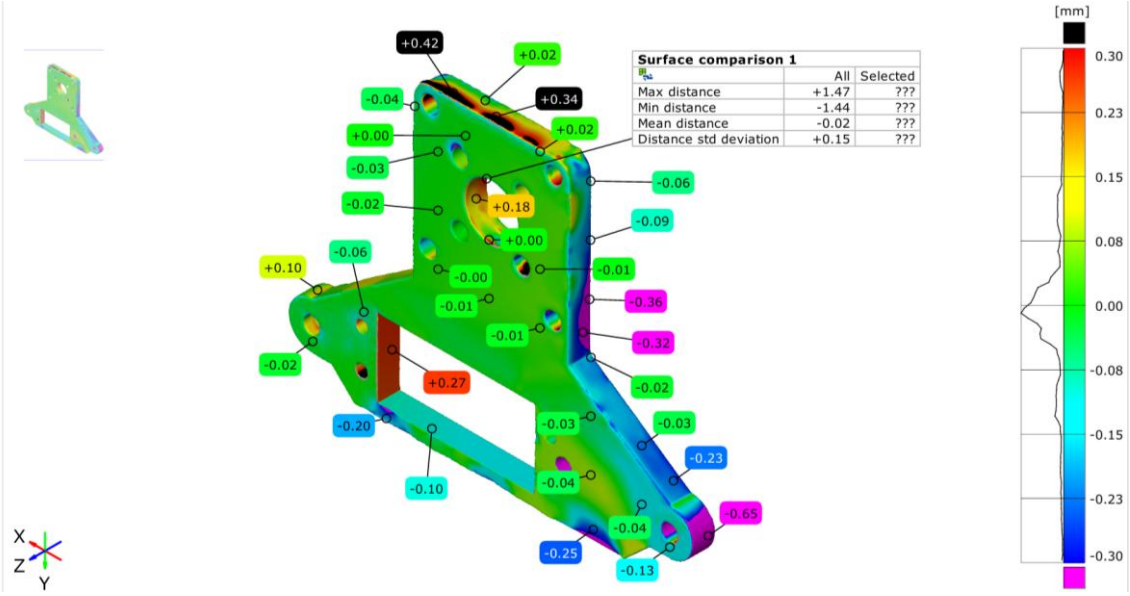

Figure 24: Surface deviation map of the EinScan-derived Frame mesh with respect to the nominal CAD model, obtained after pre-alignment and global best-fit registration in GOM Inspect. Colour scale represents signed dimensional deviation in millimetres.

Frame Surface Info	Value [mm]
Max distance	+1,47
Min distance	−1,44
Mean distance	−0,02
Distance std deviation	0,15

Table 3: Maximum, minimum, mean and standard deviation of the surface-distance values shown in Figure 24, for the Frame mesh reconstructed in Meshmixer. The values were extracted after surface comparison in GOM Inspect with a maximum search distance of 1,5 [mm].

A discussion of the EinScan Pro 2X results, including a comparison with the fixed-head SLS and CT acquisitions and the related error analysis, is presented in Section 3.5. The complete set of deviation maps, tables and per-component statistics for this scanning technology is reported in Appendix B.

3.3 Fixed-Head Structured-Light Scanning

The second 3D digitisation technique adopted in this work was fixed-head SLS, performed using a **GOM ATOS Core 300** (Carl Zeiss/GOM) at the Politecnico di Torino, during the same experimental campaign as the X-ray CT acquisition. The instrument operates on an optical principle, projecting (like the EinScan) structured fringe patterns onto the object surface and reconstructing geometry from their deformation as captured by two calibrated cameras. The fundamental architectural difference with respect to the handheld configuration lies in the **fixed optical head**: the scanner is rigidly mounted on a support arm and remains stationary during each individual acquisition, with the object manually repositioned between successive scans rather than the scanner being moved around it. The nominal point spacing of the ATOS Core 300 is 0,18 [mm], with optimal working distance of 440 [mm] and a measuring area of 320×230 [mm]. [13]

The adoption of the ATOS Core 300 was motivated by the availability of the instrument at the laboratory in Torino during the CT experimental campaign; the system is shown in Figure 25. Figure 26 shows the right claw prepared for the acquisition, with the retro-reflective markers required by the ATOS for the alignment of the individual scans.

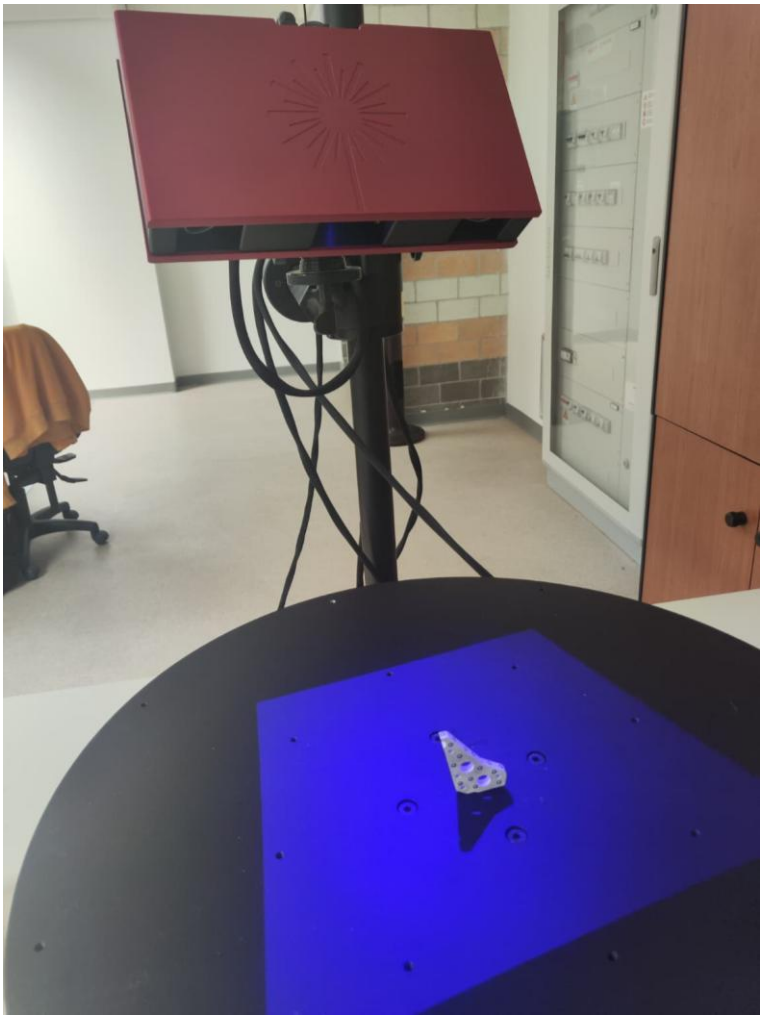

Figure 25: GOM ATOS Core 300 fixed-head SLS system with the gripper component under structured-light projector before starting the acquisition process.

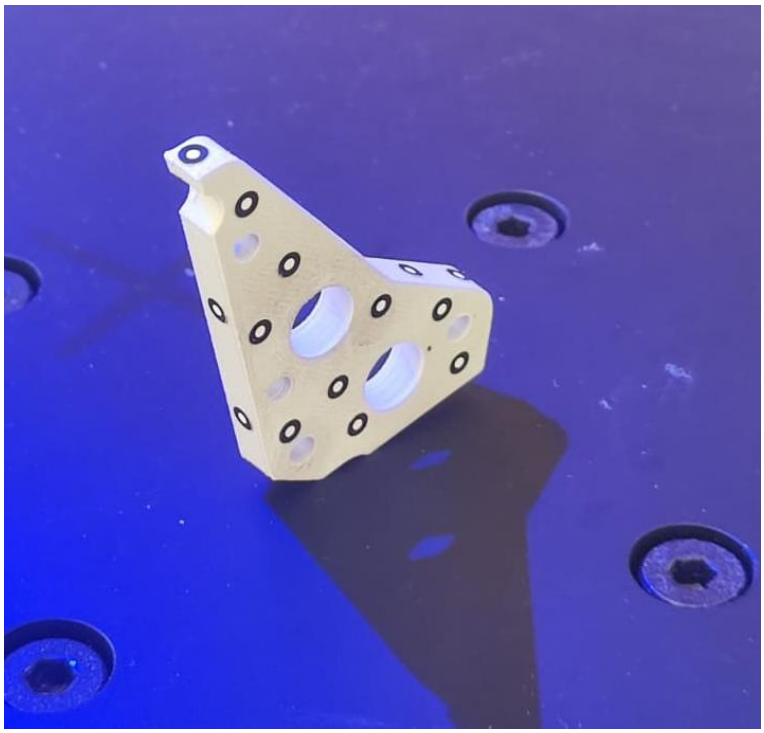

Figure 26: Right claw with retro-reflective markers applied for ATOS acquisition. Compared with the EinScan campaign (see Figure 19), the markers used here are substantially smaller and do not interfere with the component's edges.

The **polymer components were acquired individually** in the disassembled configuration. The disassembly itself was not driven by ATOS-specific considerations: the parts were already separated in preparation for the CT acquisition, based on the recommendation discussed in Section 3.4. **The ATOS scanning was therefore carried out on the same set of components while the tomographic acquisition was in progress**, which proved convenient from a workflow perspective and additionally avoided the surface occlusion and optical reflectance issues that would have been associated with the servo motor body in an assembled configuration. The components acquired were the **Frame**, the **Right Claw**, the **Right-Upper Link** and the **Driven Gear Arm**; the servo spacers could not be acquired because their geometry did not allow reliable marker placement, and their limited structural relevance did not justify alternative acquisition strategies.

The acquisition followed a structured **five-step procedure**, here illustrated through the reconstruction of the **Link** taken as a representative example. The same workflow was applied without modification to all other components (see Appendix C).

1. **Marker placement and camera calibration.** Retro-reflective circular markers of approximately 2 [mm] diameter were applied to the component using tweezers, with particular attention to placement near surface discontinuities (edges, holes, and step transitions) where tracking robustness is most critical. The reduced marker size, significantly smaller than those used in the EinScan campaign, allowed placement even on narrow surfaces without interfering with adjacent geometric features, directly addressing one of the principal limitations observed in the preceding acquisition. Camera exposure was then adjusted independently for marker detection and surface texture acquisition with the **GOM Scan** software, verifying the absence of pixel saturation (visible as red regions in the live preview) on both channels before proceeding. As shown in Figure 27, the **markers are visible in the preview as green**

crosses, indicating successful detection but not yet unique identification, which occurs only after the first acquisition is initiated.

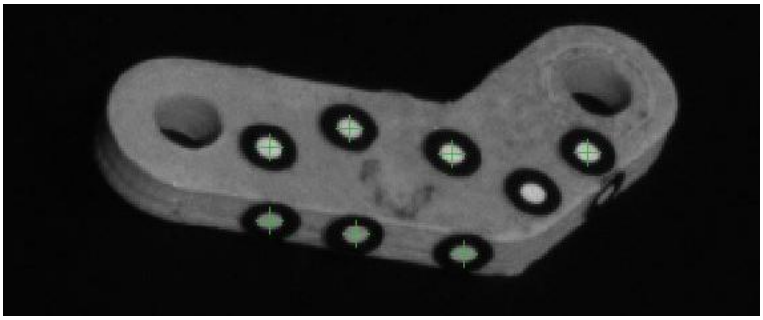

Figure 27: Link as seen by the ATOS cameras during marker placement. Green crosses indicate markers detected by the system; not all markers are identified in the first preview frame.

2. **Working distance calibration.** The scanner was positioned at the nominal working distance from the specimen (~ 440 [mm]), verified through the real-time distance indicator provided by the GOM Scan software, as shown in Figure 28. The objective of this calibration is to ensure that the markers, and consequently the reconstructed surface points, lie as close as possible to the optimal focal plane of the instrument. A scanner positioned too close to the specimen would project the acquired points in front of this plane, while an excessive working distance would project them behind it, in both cases degrading the depth measurement accuracy. The sensitivity of the system to mechanical disturbances was evident throughout this phase: vibrations transmitted through the work surface, including those induced by the operator leaning against the table, were sufficient to trigger automatic scan abort, confirming the sub-millimetre positional resolution of the tracking system.

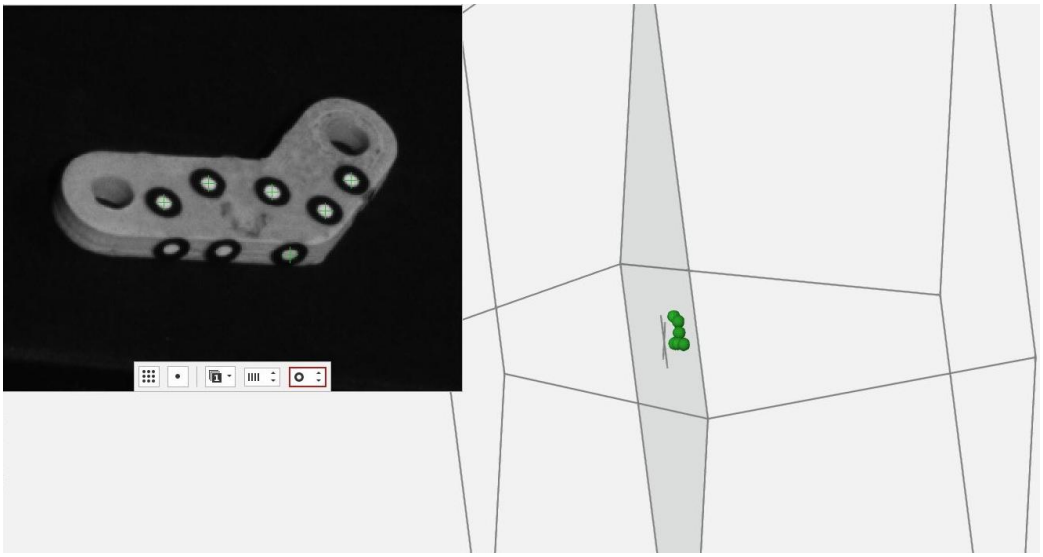

Figure 28: Working-distance indicator showing detected markers relative to the optimal focal plane. A leftward bias means the component is too close to the scanner; a rightward bias means it is too far.

3. **Initial acquisition and marker identification.** Scanning was initiated through the software interface via the spacebar shortcut, with each individual acquisition lasting approximately 3 to 5 seconds. Following the first scan, as shown in Figure 29, a significant portion of the visible surface was already reconstructed, and all detected markers were automatically numbered and assigned unique identifiers. These identifiers

serve as the geometric reference frame for all subsequent acquisitions, allowing the software to register each new scan into the same coordinate system even after the component has been physically repositioned. Markers not detected in the initial frame, due to suboptimal local exposure or unfavourable viewing angle, were recovered in subsequent passes as the scanner orientation was varied.

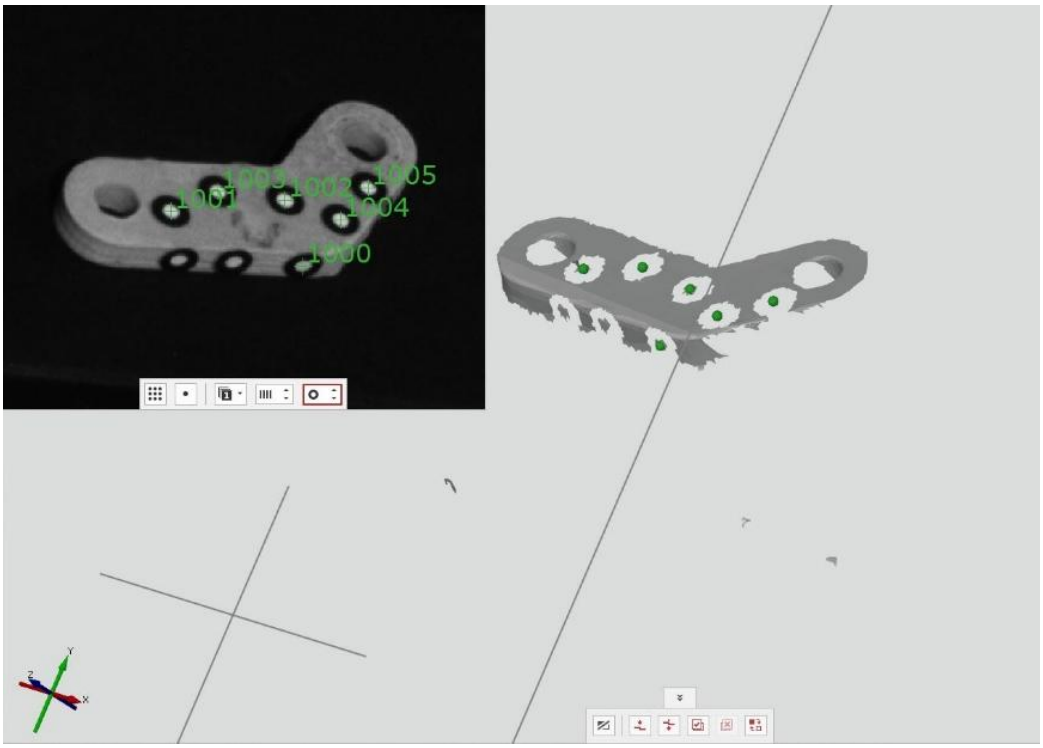

Figure 29: Result after the first Link's acquisition. The preview shows the detected points already meshed and the markers numbered with unique identifiers.

4. **Iterative acquisition and marker management.** Approximately 20 individual scans were required per component to achieve satisfactory surface coverage. As shown in Figure 30, several observations can be made on the intermediate reconstruction. First, the preview mesh is already substantially complete, with the only missing regions corresponding to the marker footprints, which are subsequently closed automatically by the software through tangent surface interpolation. Second, the number of markers visible on the component at this stage is lower than at the start of the acquisition: a subset of markers was progressively removed during the procedure to extend acquisition coverage to the underlying surface, a critical step for achieving complete reconstruction of lateral surfaces and regions near holes. Third, the raw point cloud includes background noise originating from the dark work surface; this noise, analogous to that observed in the EinScan acquisition, is easily removed through a selection of the component followed by inverted-selection deletion.

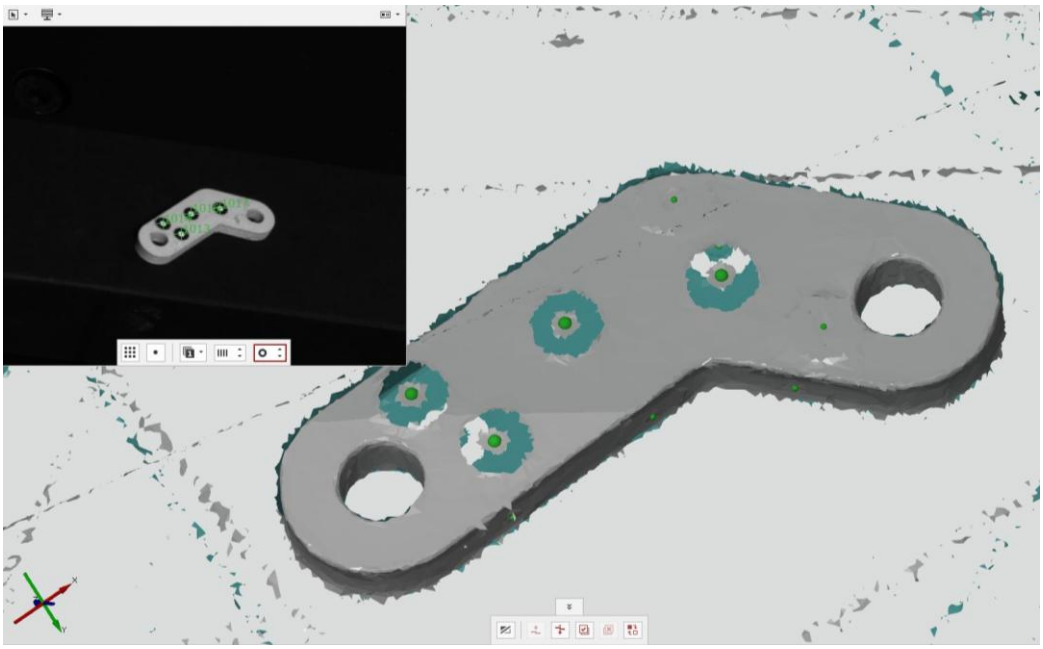

Figure 30: Link after approximately 20 acquisitions. Where green points appear without an underlying hole, markers were removed during the procedure to acquire the surface beneath; this introduces a local reconstruction error, visible for example at the green spot near the right-hand hole.

5. **Meshing and post-processing.** The consolidated point cloud was meshed using the GOM software, which offers three quality settings (coarse, medium, and precise). The **precise** option was selected for all acquisitions to maximise geometric fidelity. After noise removal, the resulting mesh exhibited a high level of surface detail with only minor residual imperfections, primarily small, unclosed regions near screw holes and narrow internal radii. These residual holes were closed using the tangent surface interpolation tools in GOM Scan, producing a closed watertight mesh for each component, as shown in Figure 31.

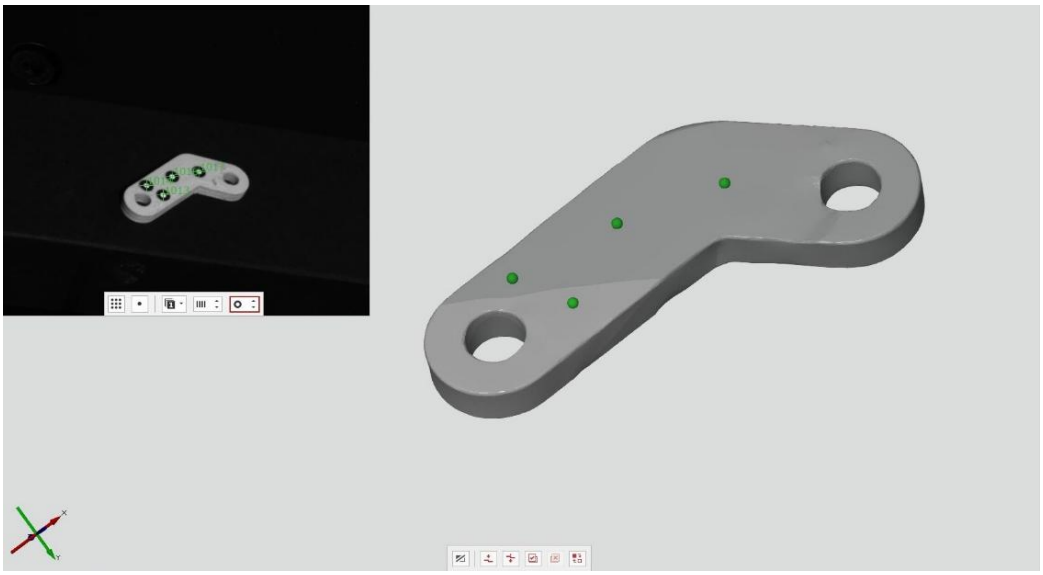

Figure 31: Final mesh obtained after background-point removal and meshing in the GOM software.

The final meshes of all components were imported into GOM Inspect for quantitative dimensional comparison with the nominal CAD model, like for the EinScan. Again, only the Link is presented here as a representative case; a full component-by-component breakdown of the ATOS results is reported in Appendix C. The ATOS acquisition was performed on the

individually disassembled components, consistently with the CT campaign; this acquisition geometry differs from that of the EinScan campaign, which was conducted on the fully assembled end-effector, and any cross-scanner comparison must be interpreted in light of this difference.

The deviation map of the ATOS-derived Link mesh against the nominal CAD model is shown in Figure 32. Following pre-alignment and global best-fit registration in GOM Inspect, the surface comparison was performed adopting a maximum search distance of 1,5 [mm], a value applied consistently across all three scanning technologies considered in this thesis (EinScan, ATOS, CT) to ensure methodological homogeneity of the comparison. The **mean dimensional deviation** across the analysed surface was **−0,07 [mm]**, with a **standard deviation** of **0,06 [mm]**. The **maximum positive deviation** recorded was **+0,14 [mm]**; the **maximum negative deviation** was **−0,51 [mm]**, localised at a narrow-undercut region where optical line-of-sight occlusion limited complete surface reconstruction, see Table 4.

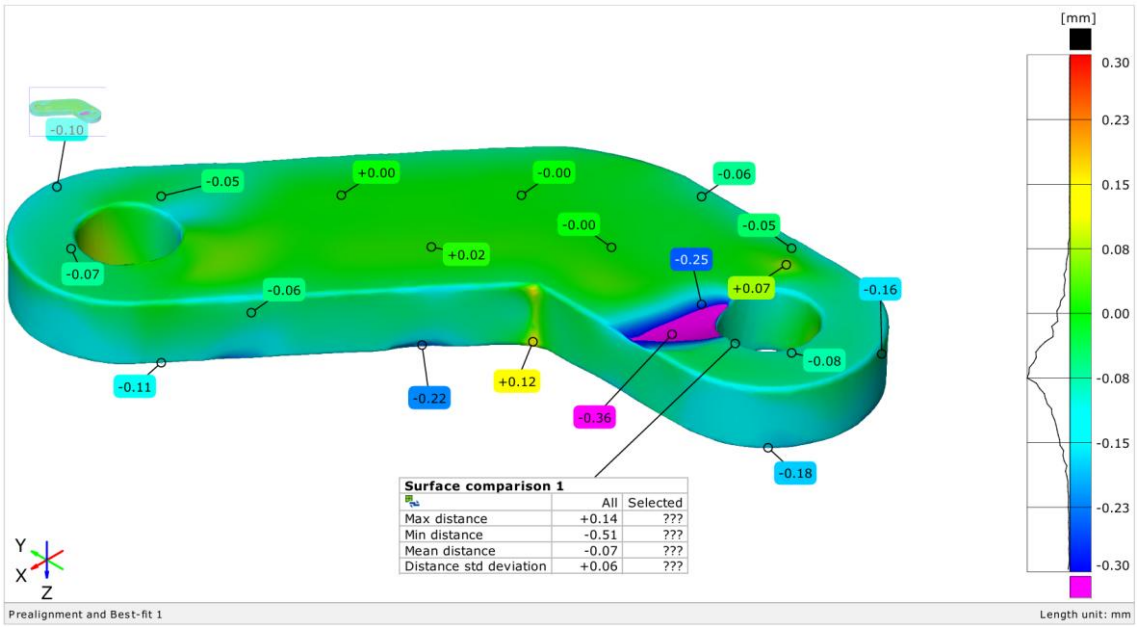

Figure 32: ATOS deviation map of the Link against the nominal CAD model. The map visualises the local distance between the acquired mesh and the reference geometry, expressed in millimetres.

Link Surface Info	Value [mm]
Max distance	+0,14
Min distance	−0,51
Mean distance	−0,07
Distance std deviation	0,06

Table 4: Maximum, minimum, mean and standard deviation of the surface-distance values shown in Figure 32, for the Link mesh reconstructed with the ATOS scanner and GOM Scan.

The mesh obtained required minimal post-processing: aside from automatic hole-filling of marker footprints and background-point removal, no further manual repair was necessary. This contrasts markedly with the EinScan acquisition, which required several hours of manual work per component in Meshmixer.

A discussion of the GOM ATOS Core 300 results, including a comparison with the EinScan and CT acquisitions and the related error analysis, is presented in Section 3.5.

3.4 X-ray Computed Tomography

The third 3D digitisation technique adopted in this work was X-ray CT, performed on a **Phoenix v|tome|x s** system (Waygate Technologies) at the Politecnico di Torino. As anticipated in Chapter 1, the CT acquisition was originally planned on the **Zeiss Xradia 620 Versa** available at the Universitat Politècnica de València; however, limitations in instrument suitability for the present specimens, in particular the restricted working volume relative to the Frame dimensions, the limited time slots available within the experimental window of this thesis, and the instrument usage fees applied at the UPV facility (15 €/h), led to the relocation of the campaign to the Politecnico di Torino during a return visit to Italy, where the Phoenix system was accessible together with the GOM ATOS Core 300 described above. The two acquisition campaigns were therefore carried out within the same experimental window, on the same set of disassembled components.

The Phoenix v|tome|x s is a high-resolution industrial microfocus CT system [14], equipped with a 240 [kV] / 320 [W] open directional microfocus X-ray tube, a high-dynamic flat-panel digital detector, and a precision rotation stage. It supports dual-detector configurations and vertical stage translation, allowing the working volume to be adapted to specimens of varying size without compromising the achievable voxel resolution, as shown in Figure 33.

The adoption of CT was motivated primarily by the opportunity to access the internal structure of the components, in particular the FDM infill pattern and inter-layer interfaces that remain entirely inaccessible to optical methods, and by the higher absolute geometric accuracy achievable with the Phoenix instrument relative to the two SLS routes previously applied.

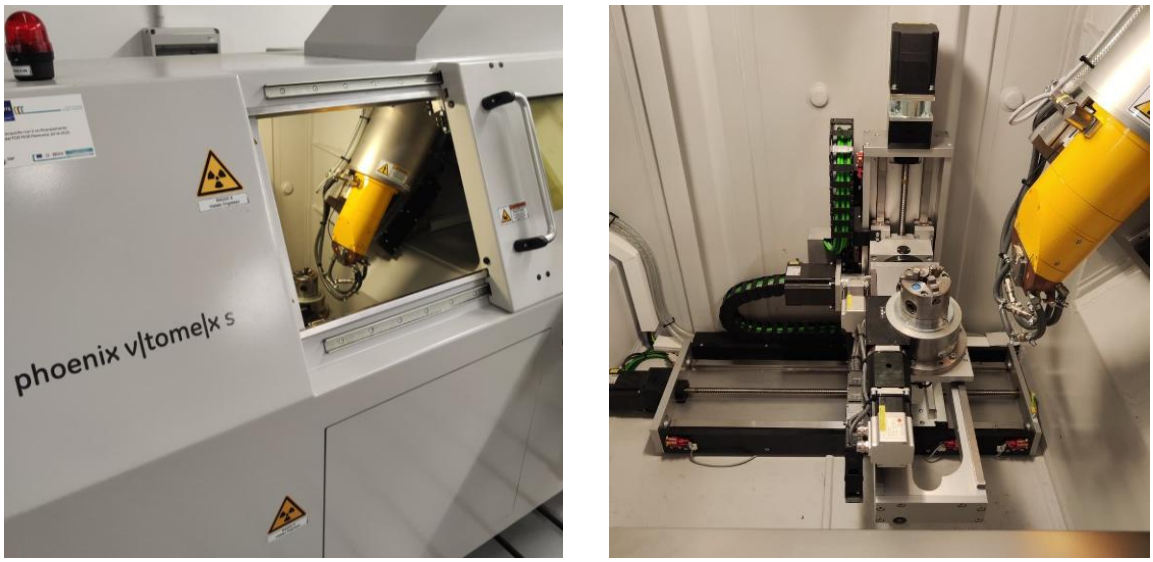

Figure 33: Views of the Phoenix v|tome|x s system. The left image shows the external cabinet; the right image shows the interior, with the X-ray emitter (yellow body) on the right, the rotating specimen stage at the centre, and the detector positioned to the left, just outside the FOV (field of view).

The longer acquisition time per component and the dependence on access to a specialised facility motivated the per-component planning strategy described below; the acquisition workflow also required dedicated post-processing in Volume Graphics VG Studio MAX to extract usable surface meshes from the reconstructed voxel volumes.

Of the polymer components of the EE, as for the ATOS Core 300 campaign, **four configurations** were acquired: the **Frame** individually, the **Right Claw** individually, the **Driven Gear Arm** individually, and the **Link together with the Servo Support** in a combined acquisition. The decision to scan only one of the two Claws, rather than both, was taken to limit

the use of the CT facility, given the fully symmetric geometry of the assembly: the geometry of the second finger is recovered by mirroring of the first. The entire CT campaign, including specimen mounting, parameter optimisation, sequential acquisition of all four configurations, dataset reconstruction, and preliminary visual inspection, occupied a **full working day** at the host facility, from 09:00 to approximately 19:00.

The acquisition procedure adopted for each component followed a structured sequence, described in the following paragraphs.

Specimen mounting. Each polymer component was embedded within low-density polymeric foam supports, selected for their negligible X-ray attenuation relative to the specimen material. Owing to their very low density, the foam supports are essentially transparent to the X-ray beam at the operating energies adopted and therefore introduce no spurious contrast in the reconstructed volume; their function is purely mechanical, maintaining the specimen in a stable position throughout the rotation sequence without inducing detectable artefacts in the data. The foam also elevates the specimen above the metallic rotation stage, ensuring that the high-attenuation metal remains outside both the X-ray beam path and the reconstructed volume. Had the metal base been intercepted by the cone-beam rays contributing to the lower specimen slices, its strong absorption and scattering would have introduced streak artefacts, dark banding and photon-starvation effects at the polymer surfaces nearest the stage. Such artefacts would have biased the local surface determination during threshold-based segmentation, degrading the dimensional fidelity of those regions. The adopted solution is shown in Figure 34: the specimen wrapped in a low-density polymeric foam that holds it in position while remaining nearly transparent to the X-rays.

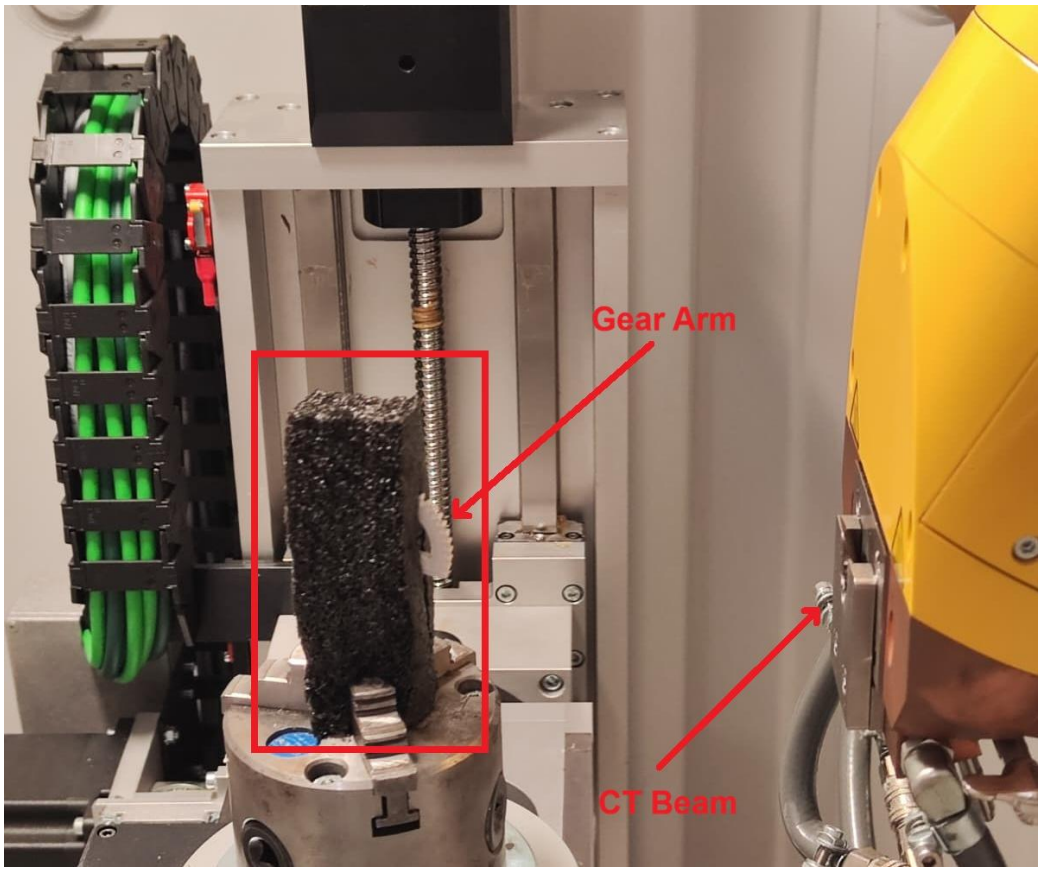

Figure 34: Driven Gear Arm specimen wrapped in a low-density soft polymeric foam, used to hold it in position and to keep it raised above the metal stage so that the base is excluded from the reconstructed volume.

Geometric calibration and beam filtering. With each specimen mounted on the rotation stage, the source-to-object distance was adjusted to achieve a target **voxel resolution** (the edge length of the cubic volumetric pixel that constitutes the smallest reconstructible unit of the CT volume) between **50 and 55 [µm]** across all acquisitions, ensuring comparable spatial resolution among the different components despite their differing sizes. A **0,1 [mm] copper pre-filter** was placed at the source output to attenuate the lower-energy portion of the X-ray spectrum (a process commonly referred to as *beam hardening pre-correction*), thereby reducing the beam-hardening cupping artefacts [15] that would otherwise arise from the polychromatic nature of the X-ray beam, particularly along the longer absorption paths through the thicker sections of the Frame.

Tube parameter optimisation. The X-ray source was activated and the acquisition parameters tuned iteratively on trial radiographs, varying the tube voltage and current and visually assessing the resulting image contrast and signal-to-noise ratio. The final acquisitions were performed at a tube voltage of **160 [kV]** with tube currents in the range **120 to 160 [µA]**, the specific value selected on a per-component basis depending on the specimen thickness and absorption profile. Higher currents were adopted for the thicker components (Frame), to increase the photon flux through the longest absorption paths, while the smaller components (Driven Gear Arm, Link with Spacer) were acquired at the lower end of the range. A representative trial radiograph of the Driven Gear Arm, acquired with these settings, is shown in Figure 35; the complete set of scan parameters adopted for that acquisition is collected in Table 5. In the radiograph the outer wall and the internal infill of the part stand out clearly against the near-transparent foam support, confirming that the selected voltage and filtering provide enough contrast to resolve the internal geometry. Table 5 groups these parameters by category: the reconstruction geometry (magnification, voxel size and focus-to-object distance), the acquisition settings (number and resolution of the projections), the detector configuration, and the X-ray source parameters (voltage, current, copper filter and focal-spot size).

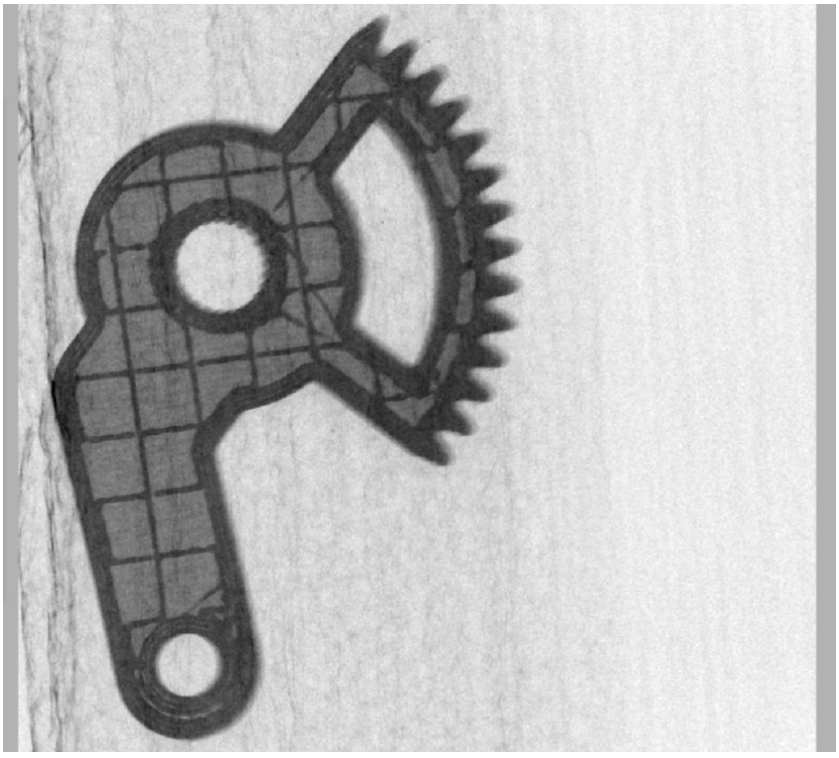

Figure 35: Radiograph of the Driven Gear Arm embedded in the low-density foam support, acquired with the settings listed in Table 5.

Scan information	Value / Unit
Geometry	
Magnification	3,874563
Voxel size	51,618725 [µm]
FOD	210 [mm]
Acquisition	
Number of images	1500
Resolution Images	990 × 1000 pixels
Detector	
Type	DXR-250 RT
Timing	200 [ms]
Skip frames	4
X-Ray	
Voltage	160 [kV]
Current	120 [µA]
Filter	0,1 Cu [mm]
Focal spot size	19,2 [µm]

Table 5: Scan parameters for the Driven Gear Arm (Figure 35) acquisition on the Phoenix v|tome|x s.

Acquisition sequence and skip parameter. For each component, a complete tomographic acquisition consisted of a full **360° rotation** about the vertical Z-axis, with a fixed angular increment of **360° / 1500 = 0,24°** between successive radiographs. A **skip parameter of 4** was applied throughout: after each angular increment, the system waited a time interval equal to four times the integration time of a single radiograph (4×T) before triggering the next acquisition. This dwell period allows the residual mechanical oscillations induced by the rotation stage motion to decay before the projection is recorded, suppressing motion blur and improving the geometric consistency of successive projections. Even small positional perturbations of the specimen can translate into noticeable reconstruction artefacts.

Per-component acquisition strategy. The geometry of the components, in combination with the limited FOV of the detector at the chosen voxel resolution, required different acquisition strategies for each part:

- The **Driven Gear Arm** and the **Link with Servo Support** (the two smallest acquisitions) were completed with a single-detector, single Z-position acquisition of 1500 projections each, with a total scan time of approximately **35 minutes per component**.
- The **Right Claw** required two consecutive vertical Z-axis acquisitions, totalling **3000 projections**, to cover its full longitudinal extent while preserving the voxel resolution below 55 [µm]. The total scan time was approximately **1 hour 10 minutes**.
- The **Frame**, owing to its substantially larger footprint, required both a dual-detector configuration and three sequential vertical Z-positions, for a total of **2 detectors × 3 positions × 1500 projections = 9000 projections**. The total scan time was approximately **3 hours 30 minutes**.

Reconstruction and dataset assembly. Upon completion of each acquisition, the integrity of the data was first verified by visual inspection of the radiographs to confirm that no positional shift of the specimen had occurred during the rotation sequence; no such displacement was observed in any of the acquisitions, and no re-scan was required. The raw projection datasets were then reconstructed using the proprietary acquisition and reconstruction software **Phoenix datos|x**, which implements an optimised filtered back-projection algorithm. For the multi-position acquisitions (Frame and Claw), the **partial volumes** corresponding to each Z-position **were merged in datos|x into a unified reconstructed volume** covering the full geometry of the component. As shown in Figure 36, once again for the Gear Arm, the grey-value histogram

of each reconstructed volume typically exhibits three distinct peaks: a **background peak** corresponding to air, a **low-density material peak** whose location partially overlaps with the background (negligible in the present case, given the monomaterial nature of the polymer components), and a **high-density peak** corresponding to the bulk polymer of the specimen.

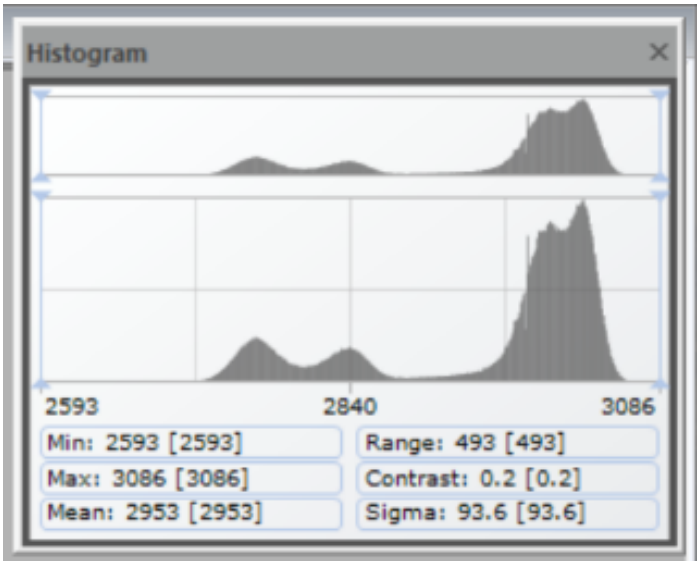

Figure 36: Grey-value histogram of the reconstructed CT volume. Three principal peaks are visible, corresponding to the background air, the low-density foam support, and the ASA bulk polymer, in order of increasing grey value.

Threshold selection, ROI definition and meshing in VG Studio. The reconstructed volumes were imported into **Volume Graphics VG Studio MAX** for **surface extraction and mesh generation**. A **grey-value threshold** was first applied to isolate the material peak from the background, retaining only the voxels assigned to the polymer. A piecewise-linear **grey-value linearisation** was then applied to **enhance the local contrast** between solid and air regions across the volume, improving the geometric fidelity of the subsequent surface extraction step. A reference coordinate system was defined as consistently as possible across all components to facilitate the dimensional comparisons, and a region of interest (ROI) was defined around each individual part within the multi-component reconstructions. Surface meshing was then performed.

The meshing step was executed **twice** for each component. The first iteration, performed with the standard meshing settings of **VG Studio**, produced surface meshes in which several fine geometric features were poorly resolved or lost (in particular, small fillet radii and gear tooth profiles). The second iteration was therefore performed using the **precise** meshing option of VG Studio, which produces a **higher triangle density and more accurate boundary localisation at the cost of substantially larger file sizes**. With this setting, all geometric features observed in the original volumes were correctly transferred into the surface meshes. The most demanding case was the Frame mesh, with approximately **13 million triangles** and a file size of approximately **634 MB**. Figure 37 and Figure 38 show the refined mesh of the Gear Arm after all the filters have been applied: the full reconstructed volume, and a sectioned view exposing the internal infill together with its voids. Representative reconstructed volumes and sections of all four acquisitions are presented in Appendix D.

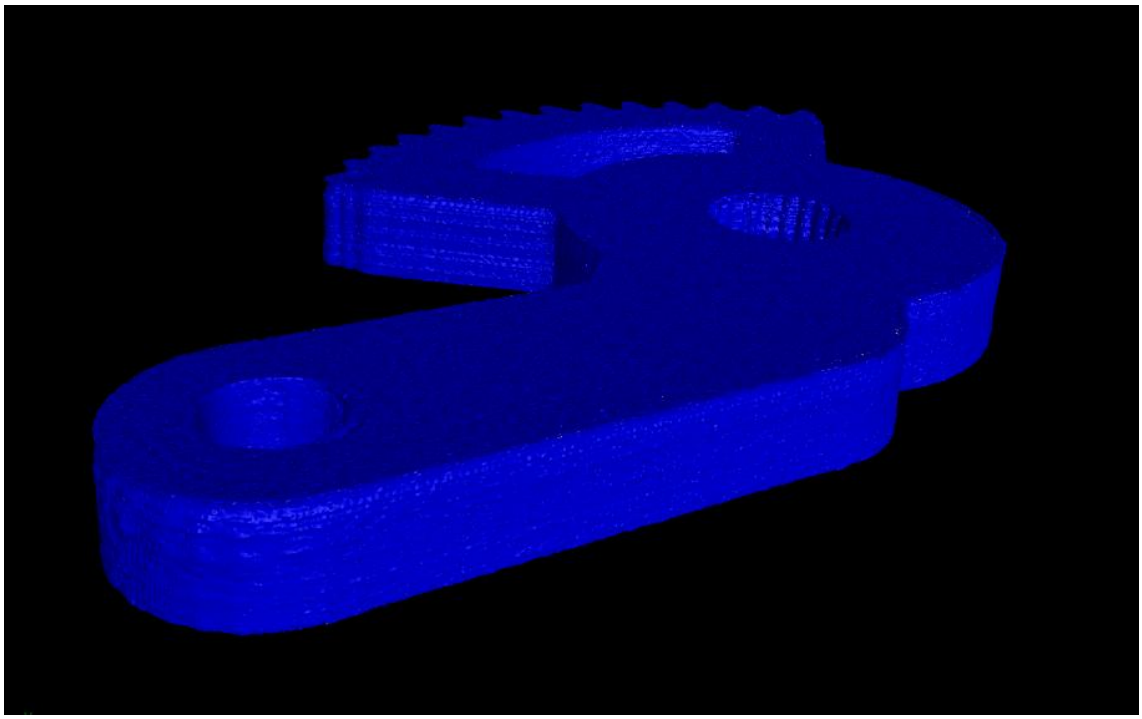

Figure 37: Full view of the reconstructed Gear Arm after cleaning the grey-value histogram of the background and low-density material.



Figure 38: Sectioned view of the cleaned mesh, exposing the internal infill structure. Note the voids within the boundary walls and the imperfections and irregularities affecting the infill.

The resulting meshes were imported into GOM Inspect for quantitative dimensional comparison against the nominal CAD model, following the same procedure adopted for the other two technologies: pre-alignment, global best-fit registration, and surface comparison with a maximum search distance of 1,5 [mm]. A full component-by-component breakdown of the CT results, including the Frame, the Right Claw and the Link with Servo Support, is reported in Appendix D, while the cross-scanner comparison of the mean deviation values is provided in Section 3.5.

The central question for the CT route is whether, retaining the same GOM Inspect parameters used for the other two techniques, a meaningful result can be obtained. The answer is no, and the reason is illustrated in the figures below. At first glance, looking at the deviation map and the overall reconstruction (Figure 39 and Figure 40), the result appears acceptable. However, on closer inspection the problem becomes evident: with a maximum search distance of 1,5 [mm], the GOM Inspect nearest-point algorithm also captures part of the interior of the part, namely the internal skeleton formed by the infill and by the wall separating the infill voids from the outer shell, as well as some of the cavities present within what should be solid material, visible in Figure 41. Retaining these data for the comparison is therefore meaningless, as they do not belong to the outer surface being assessed.

It should further be noted that the CT mesh is neither clean nor single-bodied: it presents a considerable number of holes (801), which cannot all be removed reliably within a single GOM Inspect hole-repair iteration. A detail of the region near the teeth, where internal cavities and surface artefacts are both visible, is shown in Figure 42.

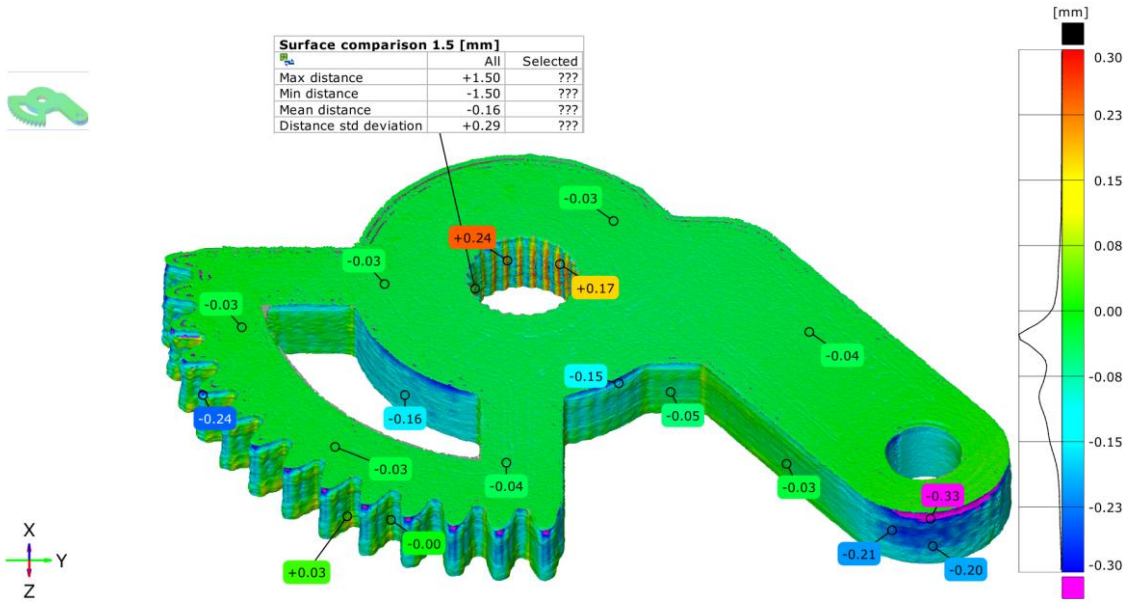

Figure 39: CT deviation map of the original full Gear Arm against the nominal CAD model. The map visualises the local distance between the acquired mesh and the reference geometry, expressed in millimetres.

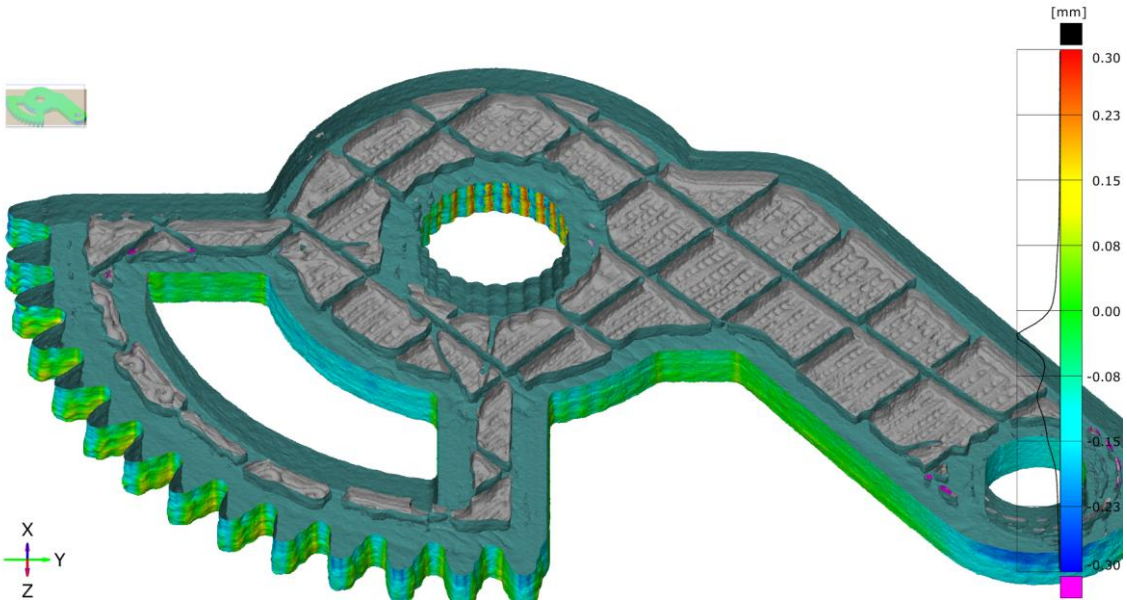

Figure 40: Section view of the original full Gear Arm. The CT acquisition allows the internal infill structure of the part to be observed in detail.

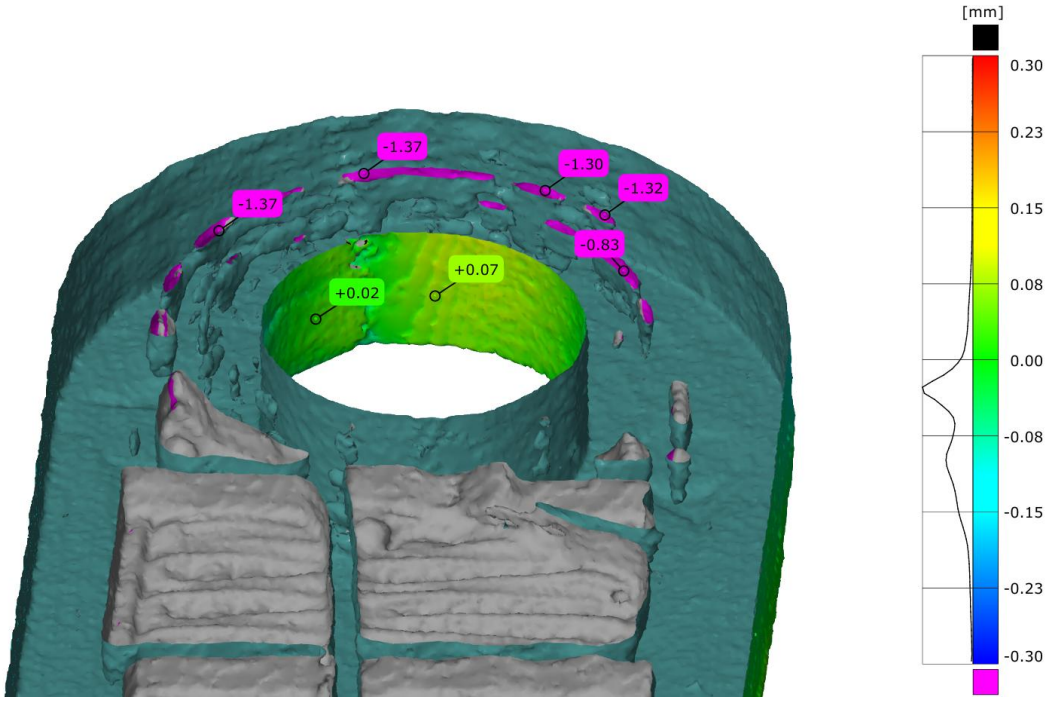

Figure 41: Cavities within the material, caused by the non-ideally homogeneous layered structure of the 3D-printed part. With a search distance of 1,5 [mm] these cavities are considered by the comparison and bias the mean and standard deviation values.

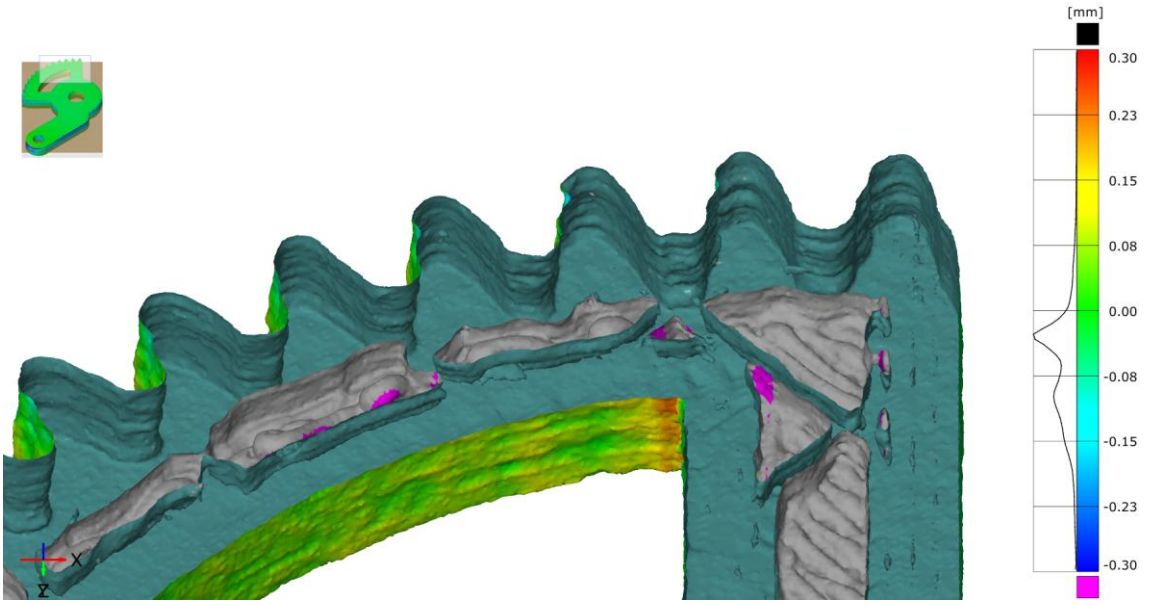

Figure 42: Detail of the region near the teeth. Not only are internal cavities (in pink) detected, but also some of the faces belonging to the boundary between internal void and material. These contribute to biasing the data.

Having established that the presence of the infill is unnecessary for the dimensional assessment, and indeed detrimental to the evaluation of mean and standard deviation as carried out for the other technologies, two possible strategies were identified:

1. The first is to **reduce the maximum search distance**. This easily resolves the partial inclusion of infill data and of the solid-void boundary, but it does not solve the problem of the cavities: these lie within a distance range compatible with the genuine deviations of the part (as shown in the figures above) and therefore remain the most evident and most frequently detected source of error. Moreover, this approach leaves the mesh affected by the hole problem, which proves difficult to fill reliably.
2. The second strategy is to **address the problem at the software level, by applying a dedicated algorithm to remove all internal geometry prior to the comparison.** The method adopted is **Hidden Point Removal (HPR),** introduced by Katz, Tal and Basri [16]. The underlying principle is that a **surface element belongs to the outer shell if and only if it is visible from at least one viewpoint located outside the part**. The algorithm distributes many virtual viewpoints over a sphere surrounding the component and, for each of them, determines which points of the mesh are visible by means of a spherical-flip transformation followed by a convex-hull operation. A point that is never visible from any external viewpoint is, by definition, interior geometry (infill, inner wall, sealed cavity) and is discarded. Crucially, the operator works directly on the point positions rather than on rays traced through the mesh, so it does not leak through the small holes of the fragmented CT mesh, a limitation that had made simpler ray-casting and voxel-based approaches unsuitable for this data. The criterion is purely geometric and intrinsic to the scan: no reference to the nominal CAD is required at this stage, so the measured point positions are never altered, and the subsequent comparison remains independent of the reconstruction.

Given the mixed engineering background of this work and the fact that the second strategy extracts more meaningful data, this route was chosen. The supporting code, developed with the assistance of an artificial intelligence (AI) coding tool (Claude, Anthropic) and subsequently reviewed, tested and fully commented, is presented in Appendix E.

In synthesis, the algorithm proceeds as follows: it loads the CT mesh and consolidates duplicate vertices; it distributes the viewpoints over a bounding sphere; it applies the HPR operator from each viewpoint and takes the union of all visible points; it retains a triangle when a sufficient number of its vertices are visible; and it exports the resulting outer-shell mesh in the original measured positions, with no point displaced or reconstructed.

Once this algorithm has been applied and the remaining holes have been closed automatically and iteratively within GOM Inspect, the following results are obtained, Figure 43. The mean and standard deviation are now substantially changed, and the model is completely hollow internally, free of any interior geometry.

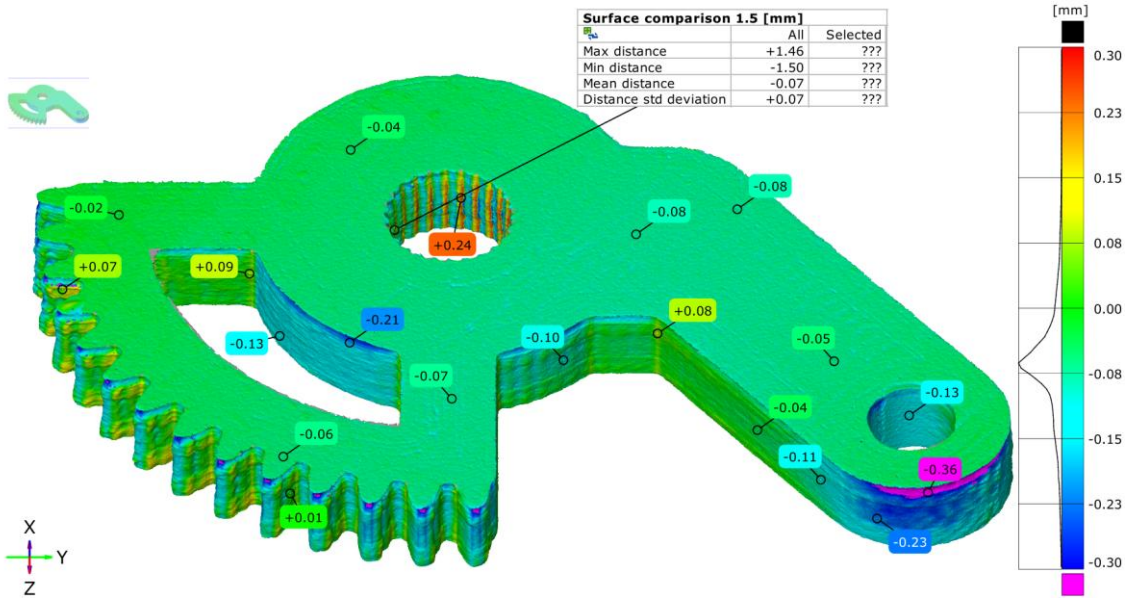

(a)

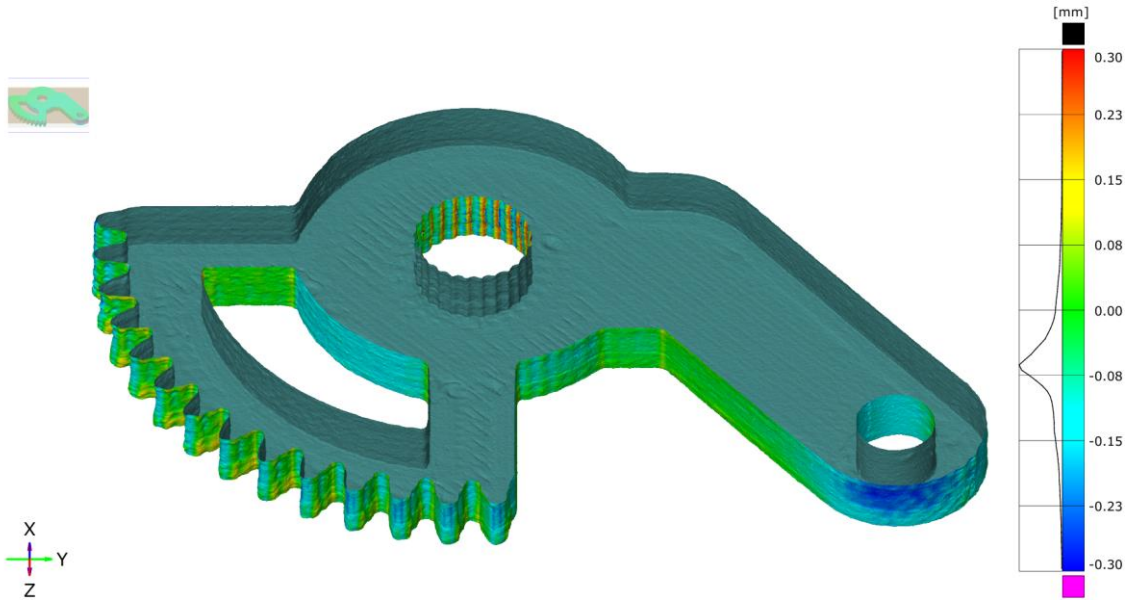

(b)

Figure 43: (a) CT deviation map of the emptied and repaired Gear Arm against the nominal CAD model; (b) A section of the emptied and repaired Gear Arm (observe the difference with the Figure 40).

The effect is also visible in the deviation histograms. The distribution obtained after applying the HPR algorithm (Figure 43-a) exhibits a single peak centred around the mean, whereas the

distribution of the original full mesh (Figure 39) is clearly multimodal, with two distinct peaks corresponding respectively to the outer shell and to the internal geometry (inner wall and infill). It should also be noted that the combination of the HPR algorithm and the GOM Inspect hole-closing and mesh error removal routines reduced the number of holes from approximately 800 to 0.

The deviation map of the CT-derived Gear mesh, processed with the HPR algorithm and with the hole-closing routine of GOM Inspect, against the nominal CAD model is shown in Figure 43. Following pre-alignment, global best-fit registration and surface comparison with a maximum search distance of 1,5 [mm] (like the others), the **mean dimensional deviation** across the analysed surface was **−0,07 [mm]**, with a **standard deviation** of **0,07 [mm]**. The **maximum positive deviation** recorded was **+1,46 [mm]**, and the maximum negative deviation was **−1,50 [mm]**.

Component	Max [mm]	Min [mm]	Mean [mm]	St.D. [mm]	# of Holes
Gear Arm					
Original	+1,50	−1,50	−0,16	0,29	801
After HPR and Holes treatments	+1,46	−1,50	−0,07	0,07	0

Table 6: Comparison of the values for the original Driven Gear Arm and for the same part after applying the HPR algorithm and the GOM Inspect hole-closing routine.

As summarised in Table 6, both indicators dropped substantially once the internal geometry was removed: the mean deviation fell from −0,16 [mm] to −0,07 [mm] and the standard deviation from 0,29 [mm] to 0,07 [mm]. This confirms that the infill and inner-wall geometry, rather than genuine surface deviations, were the dominant source of bias in the original CT comparison, and that the HPR-based outer-shell extraction restores a result consistent with the two optical scanning routes.

A discussion of the Phoenix v|tome|x s CT scanner results, including a comparison with the EinScan and GOM ATOS Core 300 acquisitions and the related error analysis, is presented in Section 3.5. The complete set of deviation maps, tables and per-component statistics for this scanning technology is reported in Appendix D.

The fully commented Python implementation of the HPR algorithm is provided in Appendix E.

3.5 Comparison and Issues

The three techniques were applied to the same physical parts, so it makes sense to start from the numbers they returned, then to ask where those numbers come from. The order matters, because the deviation each technique reports against the nominal CAD is not, for the most part, the error of the scanner. It is the end of a long chain that begins at the CAD model and runs through every file conversion, through the printing itself, and only at the very end through the scanning and the reconstruction.

The two numbers, and which one matters. Each comparison is summarised by two figures: the **mean deviation** from the nominal CAD and the **standard deviation** of that deviation over the surface. Both are useful, but they say different things. The mean is a systematic offset, whether the surface lies on average inside or outside the nominal. The standard deviation is the one that describes the dispersion, how much the local deviation varies from point to point across the part, and it is the better single measure of how faithfully the surface tracks the nominal. It is the figure leaned on most in what follows. Table 7 collects both, aggregated over the four parts.

Technique	Global mean deviation [mm]	Global standard deviation [mm]
EinScan Pro 2X (handheld SLS) *	−0,040	0,147
GOM ATOS Core 300 (fixed-head SLS)	−0,045	0,117
Phoenix v|tome|x s (X-ray CT, post-HPR)	−0,070	0,096

Table 7: Global error indicators (mean deviation and pooled standard deviation) of the three scanning technologies across the four common components, on the outer surface. (*) The EinScan was acquired on the fully assembled end-effector, under conditions not homogeneous with the other two techniques, and its figures are a qualitative baseline only.

How the figures were obtained. Each part was compared in GOM Inspect after pre-alignment and global best-fit registration, with the same 1,5 [mm] maximum search distance for all three techniques, enough to catch the real deviations without matching points on the opposite wall. The per-part results were then combined into the global ones with equal weighting: the global mean is the plain average of the per-part means, and the global standard deviation is the pooled value

$$\sigma_{pooled} = \sqrt{(\text{mean}(\sigma_i^2) + \text{var}(\mu_i))}$$

which is the **dispersion of a set of distributions** taken together when each counts equally. The **first term is the average scatter within the parts**, the **second adds the spread between their means**. Averaging the per-part standard deviations directly would have been wrong, since **variances add and standard deviations do not**, and it would have dropped the between-part spread altogether. **Weighting each part by its triangle count was rejected on purpose**: the aim is to describe the technique on a typical part of this family, not the total surface of one assembly, and the triangle count is fixed by part size and meshing density, both arbitrary, so weighting on them would have let the Frame, at about 13 million triangles against tens of thousands for the Link, drown out the smaller parts and would have folded a process choice into the result. The same scheme is used for all three techniques, so any bias it carries is shared. One honest limit: four parts is a small sample, and the between-part term rests on four numbers, so these globals are a summary of this campaign rather than a robust estimate of each technique in general, and no confidence interval is attached to them.

What the numbers say. The means agree closely, all between about −0,04 and −0,07 [mm], and all slightly negative: on average the parts come out a touch smaller than nominal. The

standard deviations, 0,147, 0,117 and 0,096 [mm], fall in the same sub-tenth-of-a-millimetre band. CT is the lowest, but the gap to ATOS, around 0,02 [mm] on both figures, is no larger than the resolution of the instruments themselves, the 50 to 55 [µm] CT voxel and the optical point spacing, so the two cannot honestly be told apart: at this scale they are equivalent. The EinScan is a little noisier, and because it was acquired on the fully assembled end-effector, with each part then segmented out of that single scan rather than digitised individually, its figures are not homogeneous with the other two. It is therefore retained as a rough reference rather than a true peer.

Where the small differences come from. Because all three measured the same physical parts, everything that happened before the measurement is common to them. The only links that differ are the last two, **the scanning and the reconstruction**, and that is where the small spread between the standard deviations arises: **optical noise**, **marker footprints** and **line-of-sight gaps** on one side, **voxel size**, **grey-value thresholding** and the **emptying-and-patching step on the other**. The agreement on the mean is not a coincidence either: three instruments are looking at the same shrunk parts, and the offset they agree on is built into the parts, not into the scanners. To see why, it helps to follow the whole chain that produced the deviation.

The chain of errors, from CAD to mesh. The deviation measured against the nominal CAD is the sum of everything that happens between the design model and the final mesh. Taken in order, each step adds its own contribution:

- **From CAD to STEP.** The native SolidWorks model is exported to a neutral Standard for the Exchange of Product Data (STEP) file. This is a boundary-representation to boundary-representation step and is essentially lossless for the analytic surfaces involved; beyond sub-micron healing tolerances it adds nothing measurable. It is named only for completeness, as the first link.
- **From STEP to STL.** To be printed, the model is tessellated into an STL, which replaces every curved surface with flat triangles. Planar faces survive exactly; curved ones, the fillets, the bores, the gear flanks, are approximated within the chordal tolerance set at export, with the facets cutting slightly inside convex curves. The effect is small, at most a few tens of microns, and systematic on rounded features.
- **From STL to G-code.** The slicer (**Cura v5.10.2**) cuts the STL into layers of 0,15 [mm] and turns each into toolpaths. Two things enter here. The layering itself approximates any non-vertical surface as a stack of steps, the staircase effect, largest on shallow slopes. And the outer contour is traced by a deposited bead of finite width, about 0,4 [mm]; where that bead is placed relative to the nominal edge, together with any horizontal-expansion compensation, shifts the printed outer surface inward or outward. Both contributions are partly systematic.
- **The print itself.** This is the first major contributor, and the one that explains the negative mean. As the **ASA cools from extrusion temperature it shrinks**, by a few tenths of a percent, which on parts tens of millimetres across amounts to tenths of a millimetre, and **it shrinks toward being undersized**: that is the systematic offset all three techniques pick up. On top of it sit warping from uneven cooling, worst on the large flat Frame, over or under extrusion from flow calibration, the slight widening of the first layers known as elephant foot, and the ordinary positioning error of the machine. These add both a systematic part, which lands on the mean, and a scattered local part, which lands on the standard deviation.

- **Post-processing of the print.** Removing imperfections and deburring the parts leaves local marks where material is taken off or left behind. On these components the effect is minor and local, but it is real, and it sits on the measured surface.
- **Scanning.** Only here does the instrument enter. The optical scanners add their point spacing and noise, the line-of-sight gaps in undercuts, and, for the EinScan, the bumps left by oversized markers. The CT adds its voxel size and, more importantly, the grey-value threshold that decides where the surface lies, a few levels of which move the extracted surface by microns to tens of microns; the copper pre-filter keeps beam hardening from making this worse.
- **Reconstruction.** Finally, the raw data becomes a mesh: surface extraction and meshing density, the interpolation that fills unmeasured holes, the Meshmixer post-process treatments to obtain the EinScan's pieces, the HPR emptying for the CT, and the best-fit alignment, which by minimising the overall residual quietly absorbs part of a uniform offset and spreads the rest. The mean and standard deviation reported are the figures left after all of this.

Put together, the chain separates cleanly into two kinds of error:

1. The **systematic links**, shrinkage above all, with tessellation and perimeter placement behind it, set the mean, which is why the mean is slightly negative and why the three techniques agree on it: they are looking at the same shrunk parts.
2. The **random and local links**, local print defects and scanning noise, set the standard deviation, and that is the only place the techniques differ at all, because only their last two links differ.

The measurement, in other words, is one source of deviation, but most probably not the main one. The manufacturing, and for the EinScan also the reconstructions, are.

Two caveats on the comparison. First, the comparison itself has a reference-side contribution: the deviations reported by GOM Inspect depend not only on the acquired mesh but also on the parameters used to tessellate the nominal model for the comparison, the chordal tolerance and the size and number of triangles. On curved features a coarse tessellation adds a systematic contribution of the order of the chordal tolerance itself, while on planar faces, which the tessellation reproduces exactly, the triangle density still fixes the number of points available to the comparison and to the best-fit registration; too few points weaken the alignment and locally distort the reported deviations. The export tolerances used here were kept well below the deviations being measured, so this contribution is small, but it is not zero and it belongs to the chain above. Second, the near-parity between the CT and the ATOS in Table 7 should not be read as equal intrinsic accuracy. The pooled standard deviation describes this campaign, on these parts, after the HPR and hole-closing treatment; the CT chain carries error sources of its own, the finite focal spot, the 50 to 55 [µm] voxel, the grey-value threshold that decides where the surface lies, residual beam hardening and the patching of unmeasured regions, which explain why, in general use, the technique is not credited with accuracies below about ±0,3 [mm] despite the favourable dispersion measured here. On this class of parts the ATOS therefore remains the reference for absolute surface accuracy.

Table 8 sets the three techniques side by side across the criteria that bear on the choice of which dataset to carry forward.

Criterion	EinScan Pro 2X	ATOS Core 300	Phoenix v|tome|x s
Nominal resolution	0,2 [mm] point spacing	0,18 [mm] point spacing	50 to 55 [µm] voxel
Mean deviation [mm]	−0,040	−0,045	−0,070
Standard deviation [mm]	0,147	0,117	0,096
External surface coverage	Partial (assembled piece)	Near complete (some marks/edge)	Complete (100 [%]), but fragmented
Internal geometry access	No	No	Yes (full volumetric)
Acquisition configuration	Fully assembled	Disassembled	Disassembled
Acquisition time / part	≈ 6 [h] (assembly)	≈ 20 scans/part (≈ 15 [min])	35 [min] to 3 [h] 30
Post-processing burden	Heavy (Meshmixer)	Light (GOM)	Medium-high (VG + HPR + GOM)
Input for FEM material model	No	No	Yes (infill resolved)

Table 8: Summary of the three digitisation techniques across the criteria relevant to the choice of dataset.

What this means for the dataset, and the hand-off to Chapter 4. Two conclusions follow. First, on dimensional grounds the choice between the datasets is close to indifferent: the three agree, ATOS and CT are indistinguishable within instrument resolution, and what deviation there is comes mostly from the print (and, for the EinScan, from the Meshmixer post-processing) rather than from the scan. Accuracy, then, does not decide which dataset to take forward. Second, what does decide it is the one thing only CT provides, the inside of the parts. The infill, the wall thicknesses and the internal defects cannot be recovered from any outer-surface scan, and they are exactly what the work needs next.

The CT dataset is therefore the one carried into Chapter 4, where it is used to evaluate the effect of the infill on the part. The finer phenomena that the CT also brings to light, the layer structure and its irregularities, the partial inter-layer fusion and the print defects visible in the volumes, are taken up at the opening of that chapter, where they properly belong.

4. Reverse Engineering of the Infill

4.1 Setup for the Reverse Engineering

Once the scans had been acquired, the work was organised to quantify the mechanical difference between a part produced by FDM with internal infill and the same part treated as fully solid. The two configurations differ in several aspects, and the literature shows that infill density is a primary driver of the mechanical response of FDM parts, while the infill pattern plays a secondary role [17]. The question addressed here is narrower: whether the 20 [%] infill selected for fabrication behaves differently under load from the fully solid original geometry, and whether any such difference is large enough to matter for the structural assessment. Answering it requires the internal geometry of the Claw and of the Frame to be recovered, so that both can be analysed in Chapter 5. The CT acquisition, combined with the original CAD model, provided the information needed to reconstruct that internal geometry by reverse engineering.

4.2 CAD Reconstruction in SolidWorks from the CT Data

Direct surfacing of the CT mesh inside SolidWorks was not practical. The reconstructed volume contained approximately 13 million triangles, well beyond what the CAD environment handles efficiently, so two viable routes were considered. They are summarised in Table 9.

Approach	Working principle	Main advantage	Main limitation
A. Mesh decimation and ScanTo3D	Decimate the CT mesh in MeshLab or with a Python script, import the reduced mesh, and surface it with the SolidWorks ScanTo3D module.	Allows the wall thickness and the top and bottom layers to be measured and reproduced as faithfully as the data permit.	Requires an external decimation step whose result is not guaranteed, and the surfacing quality depends heavily on the decimated mesh.
B. Image guided rebuild on the original CAD	Extract a cross-sectional image of the infill, import it into SolidWorks as a sketch reference or as a texture, scale it to the part, and rebuild the internal structure on the original shell.	Reuses the validated original shell and keeps the model light and fully parametric.	Hybrid and approximate by construction, although the deviations reported in Appendix D between the CT-scanned and the original parts indicate that the approximation is acceptable.

Table 9: Comparison of the two candidate reconstruction strategies.

The second approach was adopted. The reconstruction started from cross-sectional images of the infill such as the one in Figure 44.

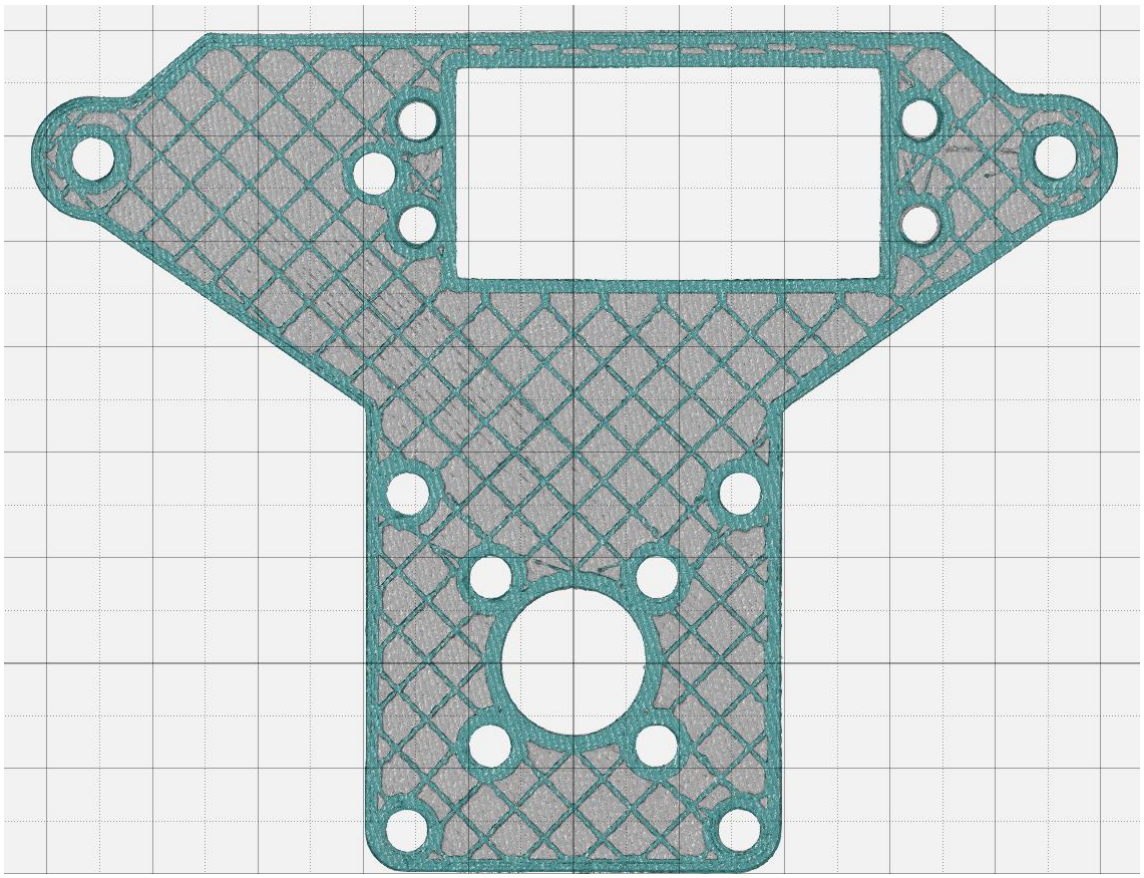

Figure 44: Representative cross-section extracted with GOM Inspect from the CT volume, used as the geometric reference for the infill reconstruction.

The cross-section is usable inside SolidWorks as it stands, but its contrast does not separate the infill walls cleanly from the surrounding material. The image was therefore processed with standard computer-vision operations. A Python script written some months earlier for an unrelated project was reused for this purpose: it applies a colour threshold in the hue, saturation, value (HSV) space, Figure 45, and then performs morphological operations (erosion and dilation) on the regions of interest. The full, commented code is reported in Appendix F.

The processing was carried out in two stages. First, an HSV threshold isolated the infill walls using the bounds listed in Table 10; the result is shown in Figure 46. The hue channel follows the 0 to 179 scale of the imaging library, while the saturation and value channels range up to 255.

Channel	Min	Max
Hue	78	101
Saturation	30	145
Value	0	255

Table 10: HSV threshold bounds applied to the cross-sectional image.

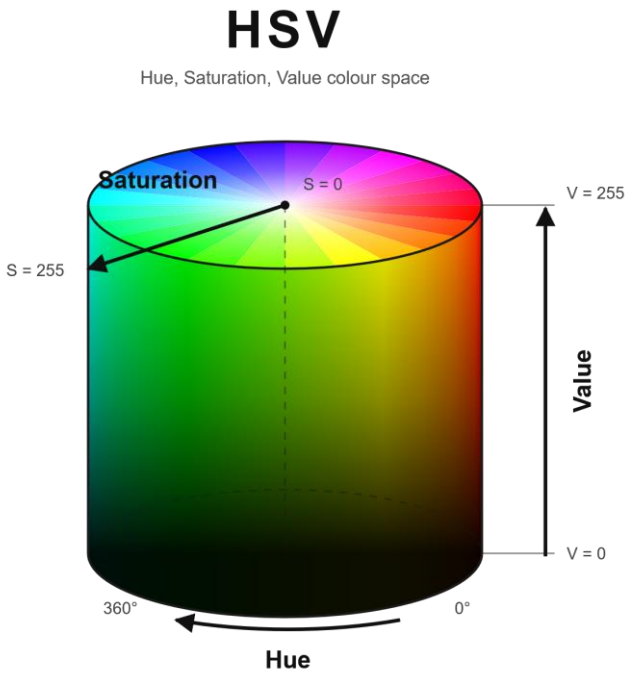

Figure 45: Representation of the HSV colour space: hue as the angular coordinate, saturation as the radial distance from the central axis, and value as the height.

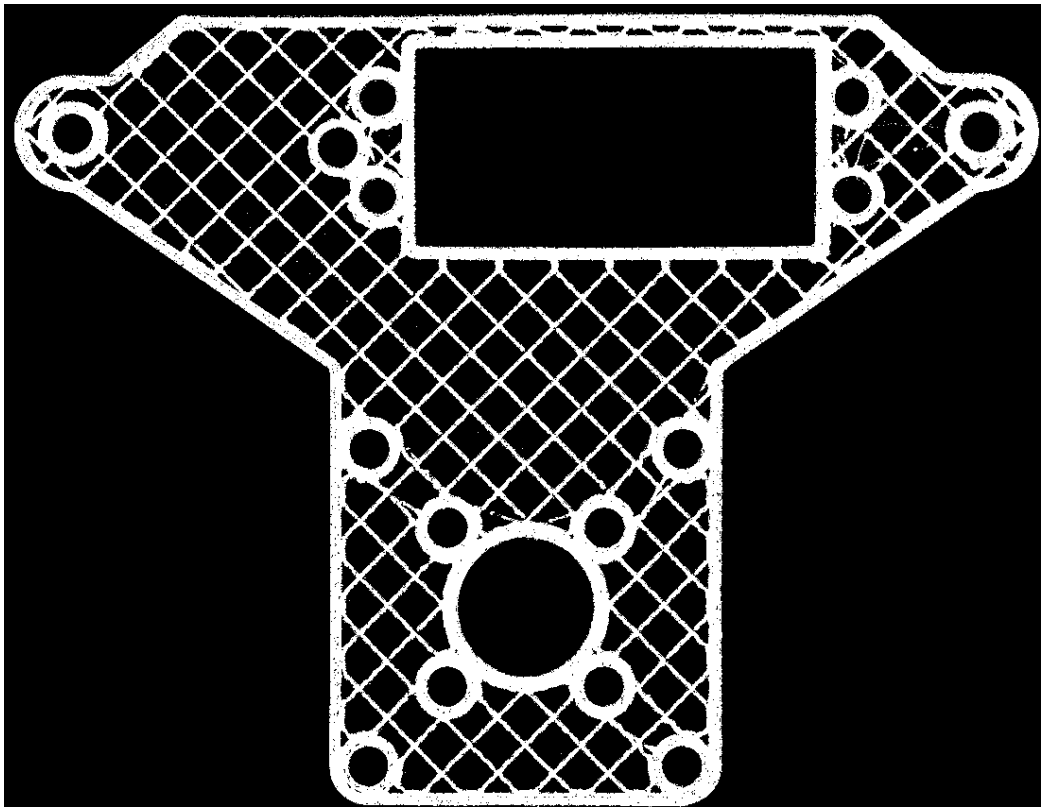

Figure 46: Infill walls isolated by HSV thresholding with the bounds of Table 10.

Second, a morphological closing (a dilation followed by an erosion) was applied with a 3 × 3 structuring element, Figure 47. This consolidated the segmented mask and bridged the small gaps left by the threshold, so that the square lattice could be traced continuously.

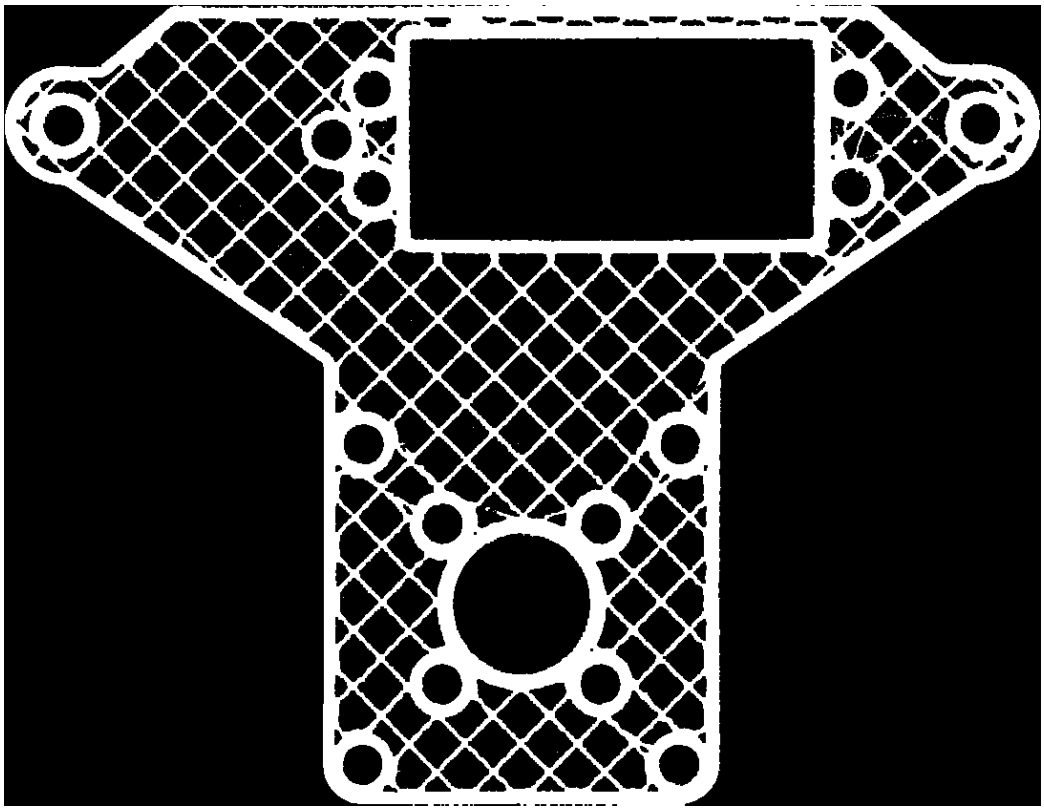

Figure 47: Segmented mask after the morphological closing performed with a 3 × 3 structuring element.

The processed section, Figure 48, was then applied as a reference image onto the original solid, which allowed the infill to be characterised directly on the validated geometry.

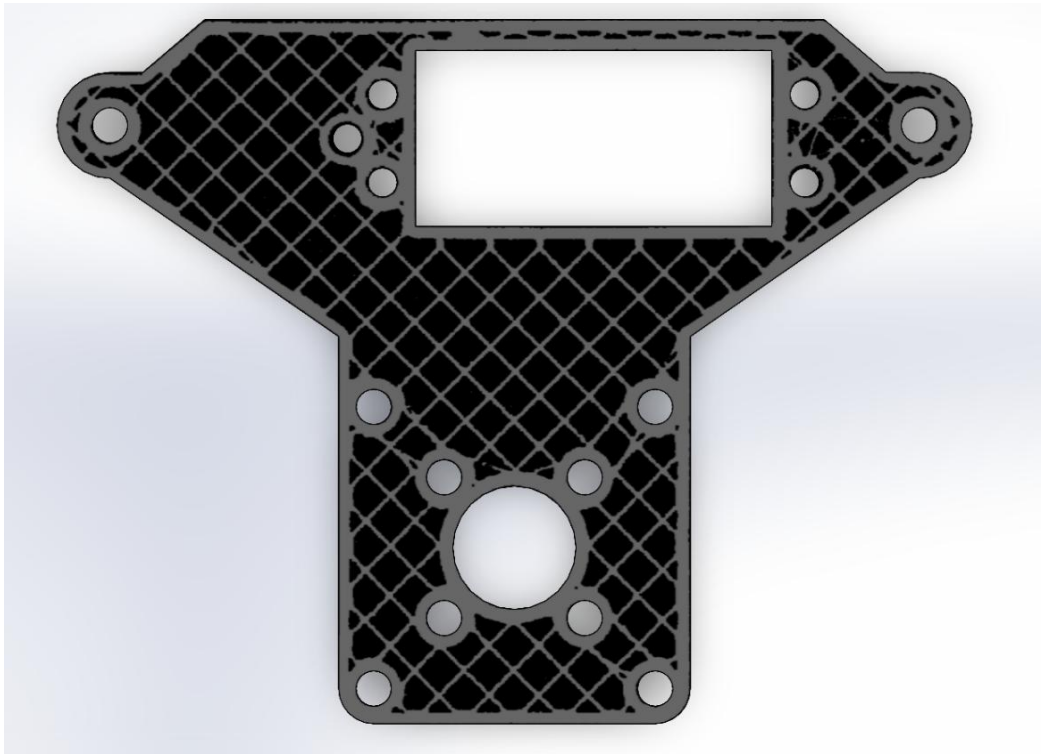

Figure 48: Processed cross-section applied as a reference image onto the original Frame CAD solid in SolidWorks.

Repeating the procedure for the Frame and for the Claw, the square pattern was reproduced and the infill rebuilt for both parts. The resulting section views are shown in Figure 49 and Figure 50.

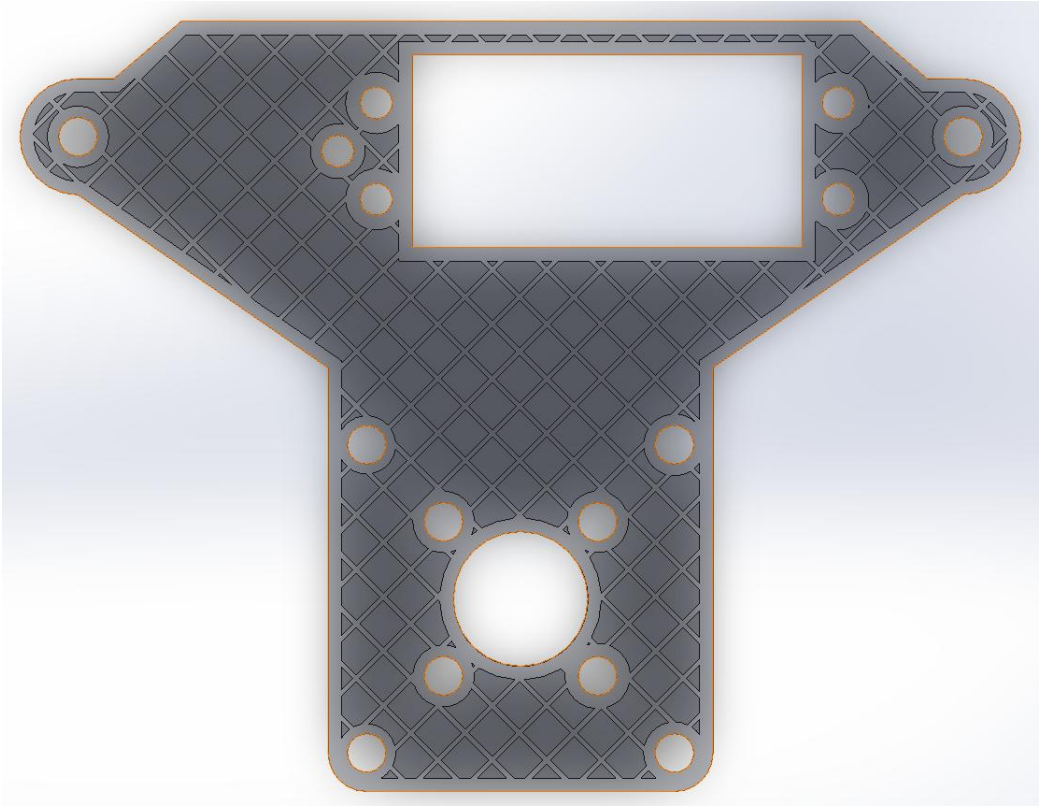

Figure 49: Reconstructed Frame, section view showing the rebuilt infill.

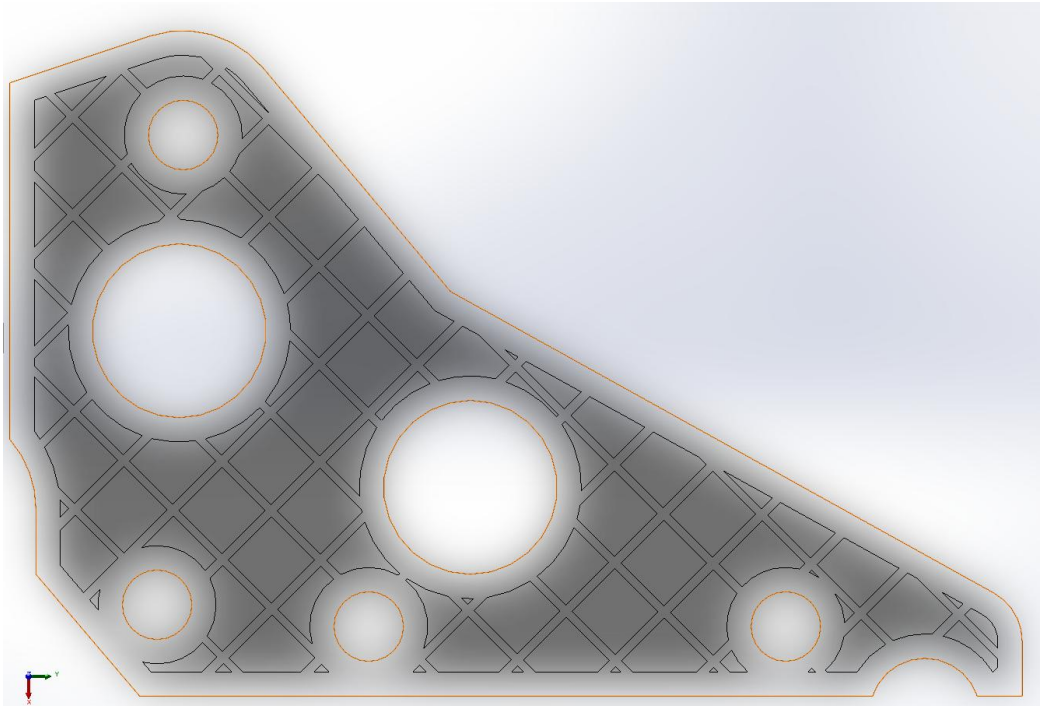

Figure 50: Reconstructed Claw, section view showing the rebuilt infill.

The reconstruction is therefore hybrid: the original CAD supplies the outer shell, while the infill is derived from the cross-sectional image processed through HSV thresholding and morphological closing.

4.3 Dimensional Comparison and Considerations

A critical issue appeared at this stage: GOM Inspect could not align either reconstructed part, the Frame or the Claw, with the corresponding CT mesh. After several hours the cause was identified. The STEP file exported by SolidWorks carried reversed surface normals, so that even a manual pre-alignment left the comparison algorithm matching the wrong faces. Figure 51 shows the reversed surface normals of the infill.

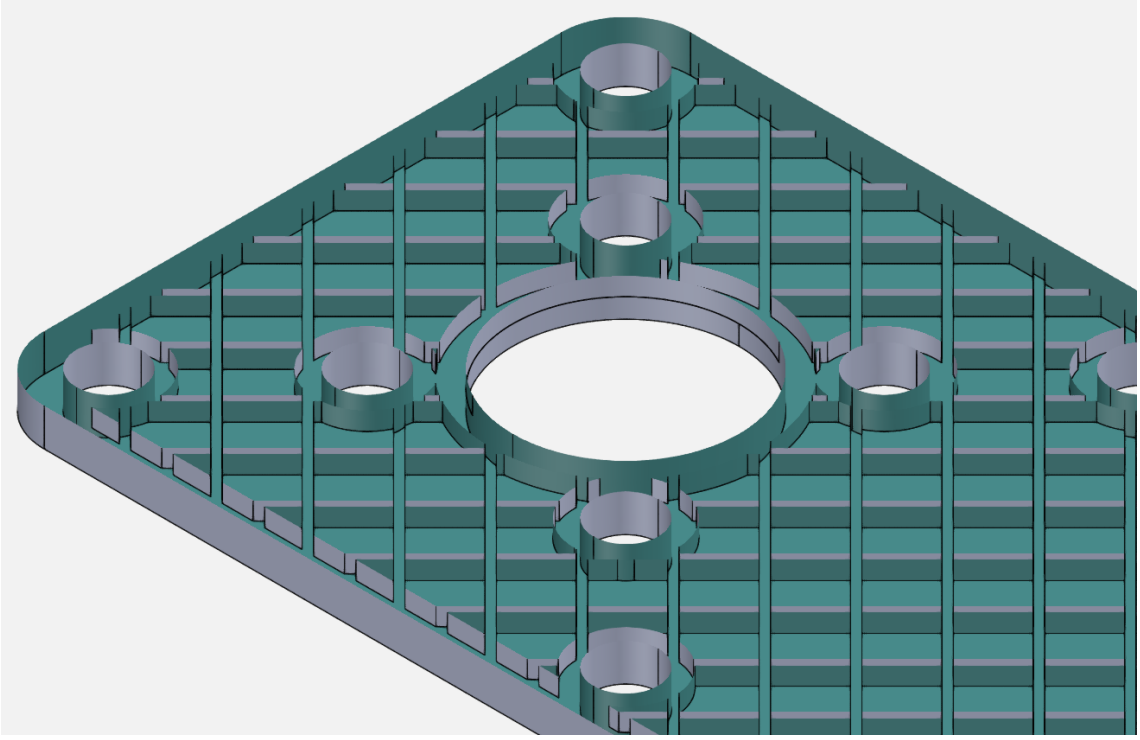

Figure 51: Frame STEP model exported from SolidWorks with reversed surface normals, revealed by the inverted front and back face shading.

The viewer shades the two sides of each surface differently: grey marks the external, outward-facing surfaces, and the blue to green shading marks the internal surfaces that face into the void squares. On the infill the two are swapped, so the external surfaces appear blue to green and the internal surfaces appear grey, which is the visible signature of the inverted normals. With the normals in this state the global best-fit alignment failed, and every other alignment matched the reconstruction against the wrong faces, as in Figure 52.

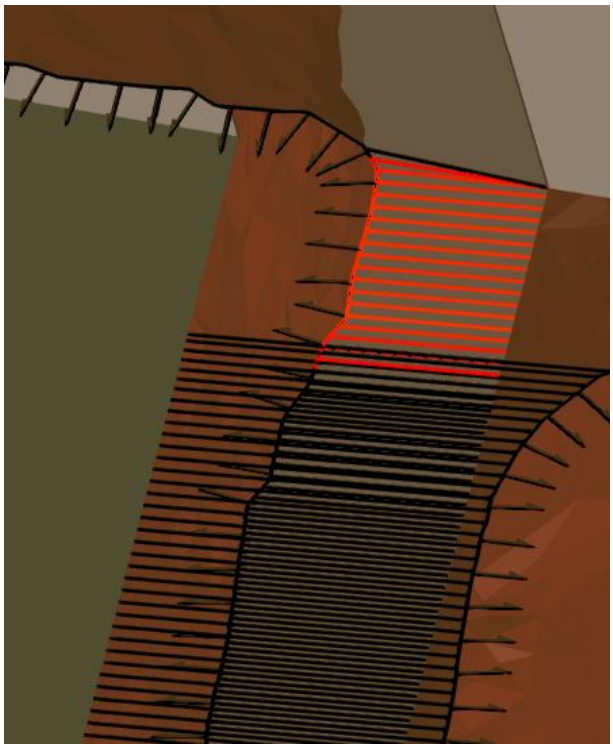

Figure 52: Erroneous surface comparison caused by the reversed normals, with the reconstruction matched against the wrong faces.

The problem was solved by inverting the normals of the STEP model inside GOM Inspect, Figure 53. After the inversion the part showed normals consistent with the intended infill geometry.

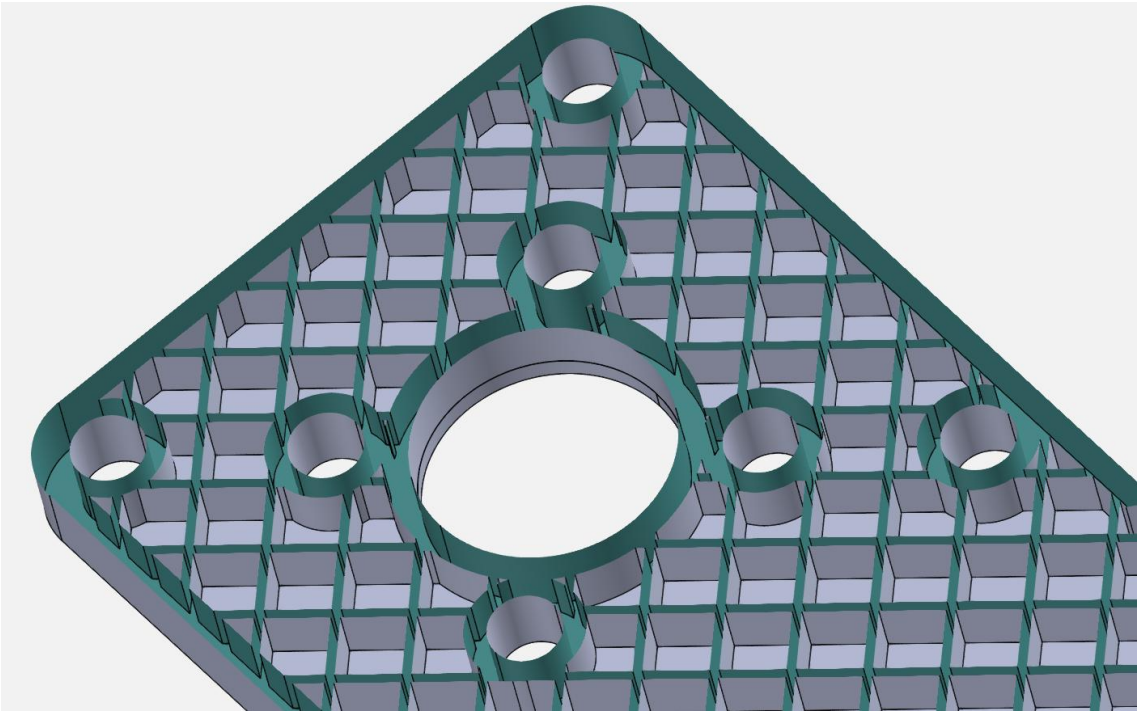

Figure 53: Frame STEP model after inverting the normals in GOM Inspect, with shading consistent with the intended infill geometry.

The standard GOM Inspect procedure was then applied: a global best-fit alignment followed by a surface comparison with a maximum search distance of 1,5 [mm], Figure 54. The mean and the standard deviation are not the focus here; the comparison serves to confirm visually that the reconstruction is compatible with the internal infill.

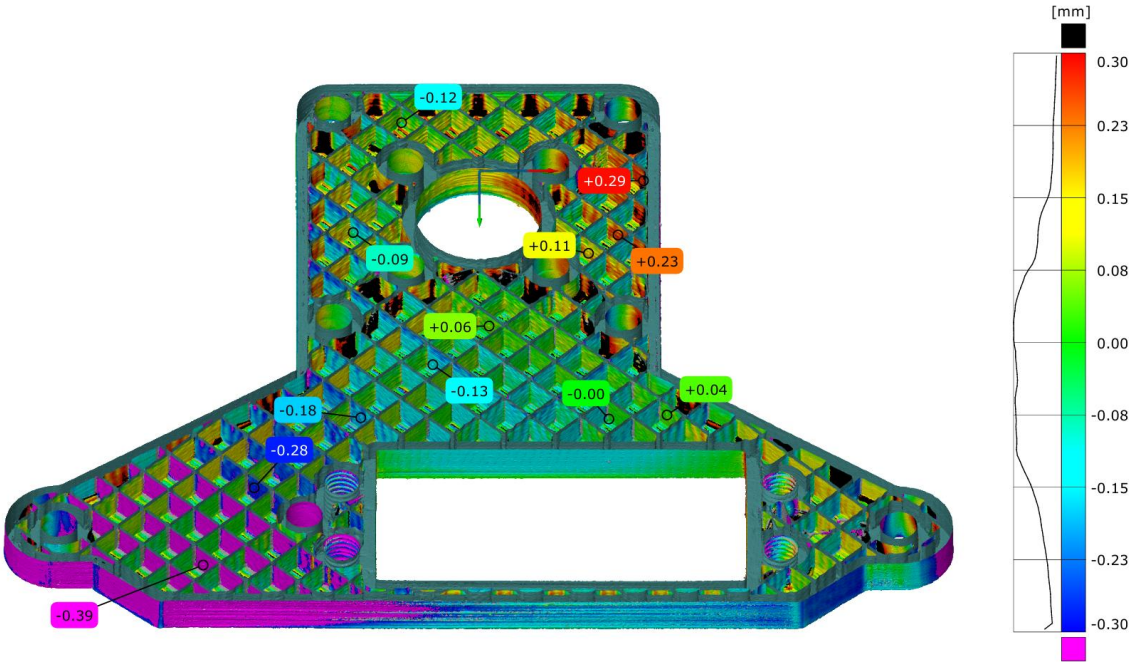

Figure 54: Surface comparison of the reconstructed Frame infill against the CT mesh (global best-fit alignment, 1,5 [mm] maximum search distance).

The structure built in SolidWorks follows the real infill closely, Figure 55, with a local deviation of at most about half a millimetre.

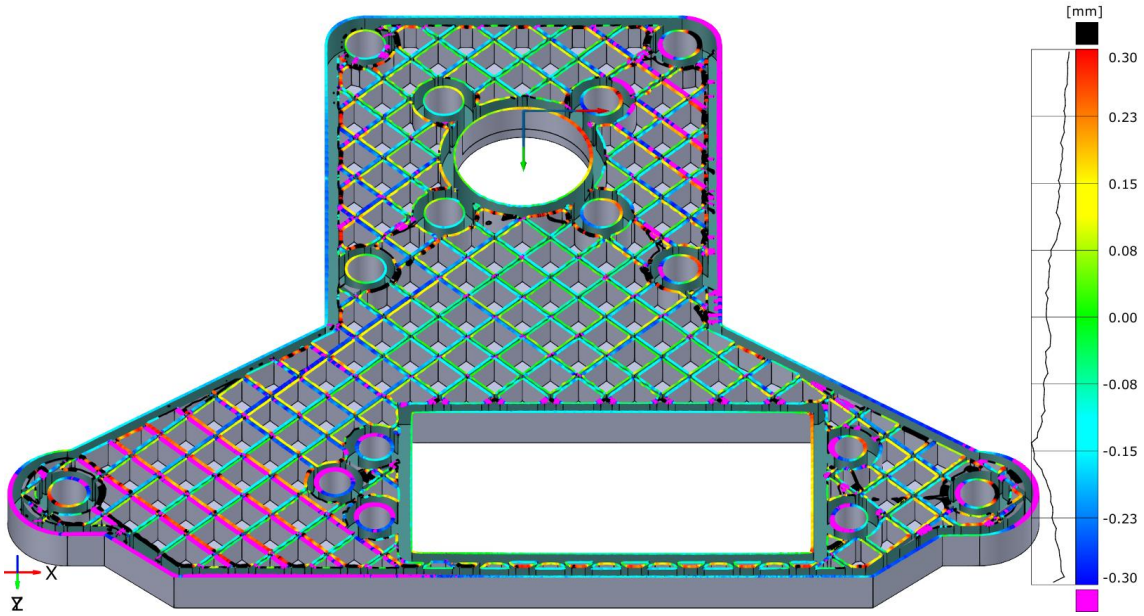

Figure 55: Section comparison of the Frame CT mesh against the reconstructed infill.

The comparison also makes clear what was deliberately left out of the reconstruction. The CT data contain small walls and ridges produced by the travel of the printer nozzle while residual

material was still attached. These features do not belong to the nominal infill; they are print defects, visible in Figure 56, and they were not reproduced in the CAD model.

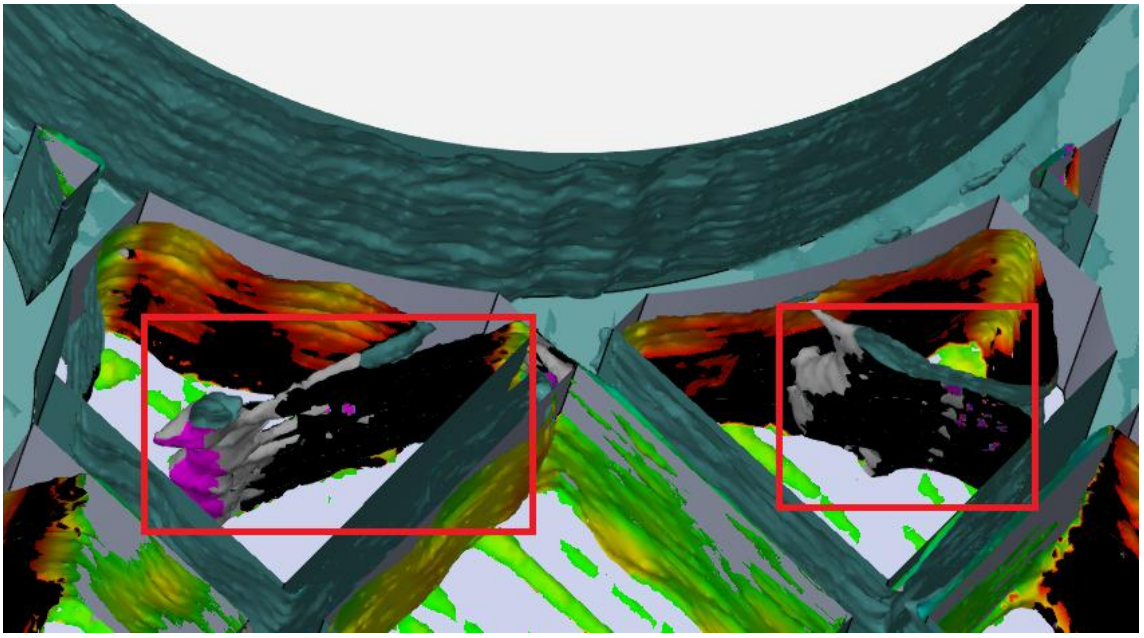

Figure 56: Spurious thin walls in the Frame CT data, produced by nozzle travel with residual attached material and identified as print defects not present in the reconstruction.

For completeness, the Claw was examined in the same way, which confirmed the reconstruction of its infill. The corresponding views are shown in Figure 57 to Figure 59.

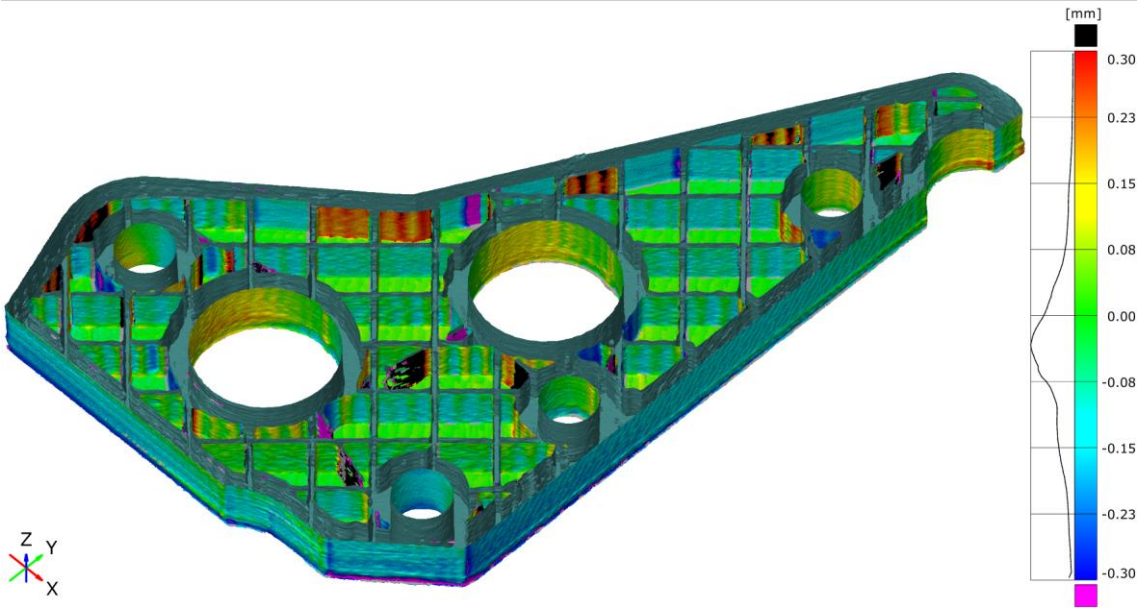

Figure 57: Surface comparison of the reconstructed Claw infill against the CT mesh.

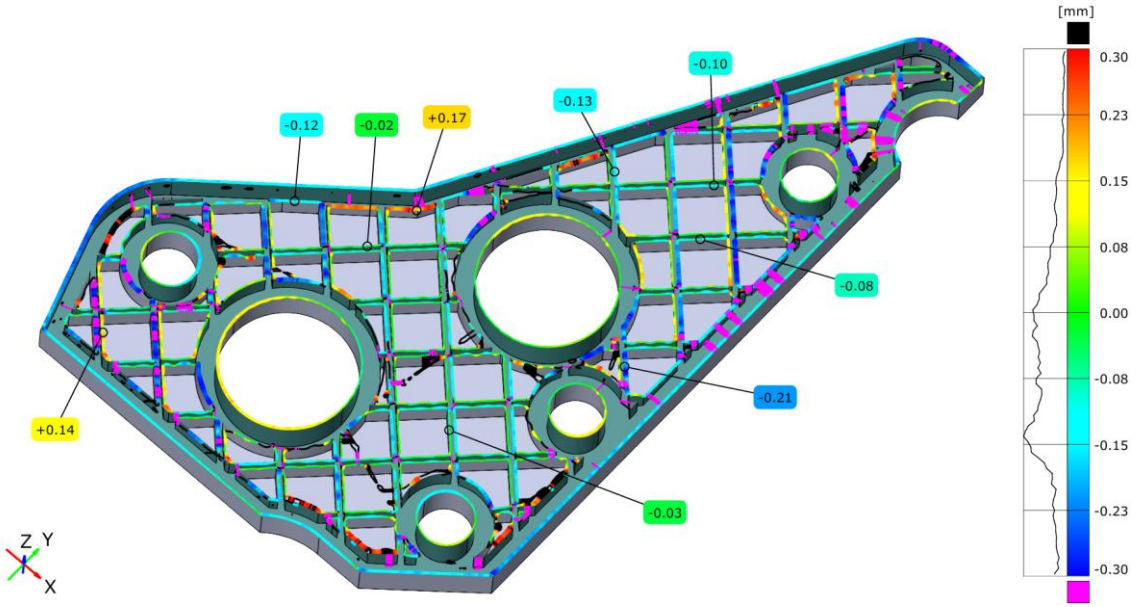

Figure 58: Section comparison of the Claw CT mesh against the reconstructed infill.

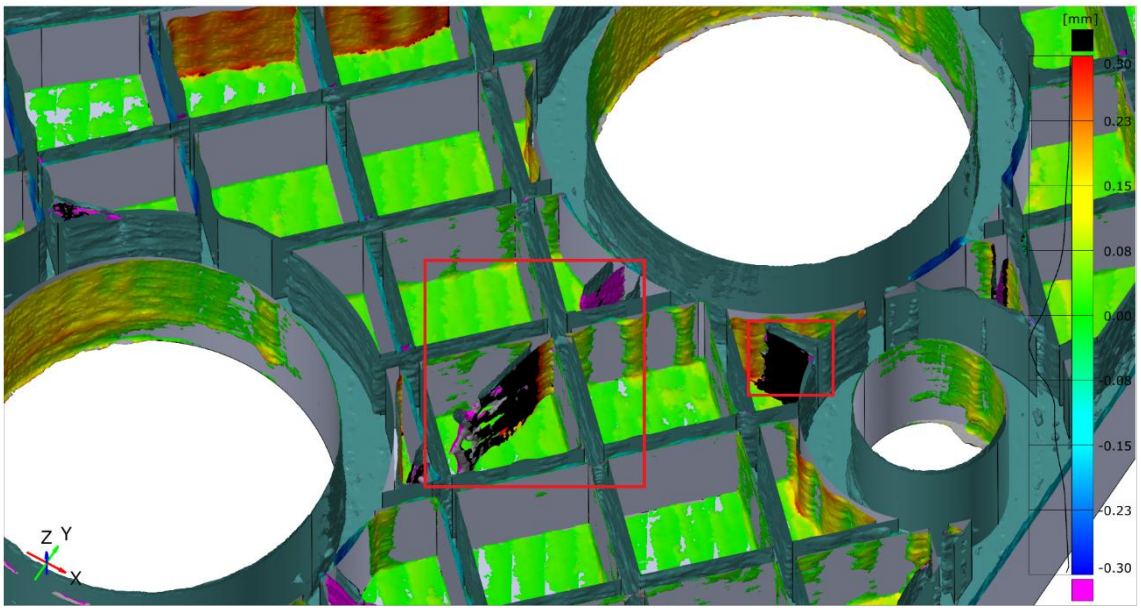

Figure 59: Print defects identified in the Claw CT data.

5. FEM Structural Analysis

5.1 FEM Setup: Material Model, Constraints and Load Cases

The FEM analysis was carried out in Altair Inspire, using the OptiStruct solver. Inspire works on the CAD geometry, and when an OptiStruct run is launched it generates a tetrahedral finite element mesh from that geometry; the average element size is either kept at the automatic value, computed from the smallest features of the part, or set explicitly, a finer size giving a more accurate but slower solution [19]: preliminary convergence tests showed that a good compromise between solution time and accuracy is an average element size of 0,5 [mm] (this value was obtained by progressively reducing the average element size until the SF stopped changing, then increasing it again to the largest size that left the SF unchanged while shortening the solution time). Three reasons motivated the choice of OptiStruct over the meshless SimSolid solver also available in Inspire. First, OptiStruct is the solver that performs the topology optimisation of Chapter 6, so the verification and the optimisation share the same solver and the same mesh-based stress field, and the results of the two phases are directly comparable rather than read on two different discretisations. Second, the Claw load case requires a partial degree of freedom support at hole B, where the action of the Link, a two-force member (Appendix G), is reproduced by restraining only the translation along the Link axis, while leaving free the rotation about the pin and the translation perpendicular to it; OptiStruct accepts this partial support, whereas the meshless solver does not handle partial degree of freedom supports on curved faces robustly. Third, the mesh-based formulation gives explicit control over the element size and the local refinement, which helps to resolve the stress field around the functional features. The reverse engineered models, although geometrically complex because they reproduce the internal infill as solid CAD geometry, remain clean solids and are meshed without difficulty.

Based on the supplier technical data sheet [10] and standard literature values for bulk ASA [18], the isotropic linear elastic properties in Table 11 were assigned. E and σy were converted from the data sheet through 1 [kgf/cm²] = 0,0980665 [MPa]; ν and the thermal properties, not reported by the supplier, were taken from literature. The thermal properties play no role in the present static, isothermal analyses and are included only for completeness of the material model.

Property	Symbol	Value / Unit	Source
Mechanical Properties			
Young's modulus	E	2168 [MPa]	ASA 3DJAKE
Poisson's ratio	ν	0,350	literature
Density	ρ	$1{,}100 \times 10^{-6}$ [kg/mm³]	ASA 3DJAKE
Yield Stress	σ_y	78,5 [MPa]	ASA 3DJAKE
Thermal Properties			
Thermal expansion	α	$9{,}000 \times 10^{-5}$ [K⁻¹]	literature
Thermal conductivity	λ	$1{,}700 \times 10^{-4}$ [W/(mm·K)]	literature
Specific heat	cp	$1{,}300 \times 10^{3}$ [J/(kg·K)]	literature

Table 11: ASA material card combining supplier and literature data. The supplier values were measured on standard specimens according to ASTM D790 and therefore describe the bulk polymer rather than the FDM printed material.

The load definition concerns the two structural components of the mechanism, the Frame and the Claw, which are characterised in turn below.

Frame: reference system and constraints. The reference system is defined as in Figure 60-a, with the gear axes parallel to the global X axis and the mechanism operating in the YZ plane, Z being the vertical axis and Y the clamping direction. The Frame is constrained through 8 supports applied to its 8 mounting holes (Figure 60-b), under the working assumption that the robot side fixture uses all of them, which is the stiffest and therefore most conservative mounting condition for the stress field.

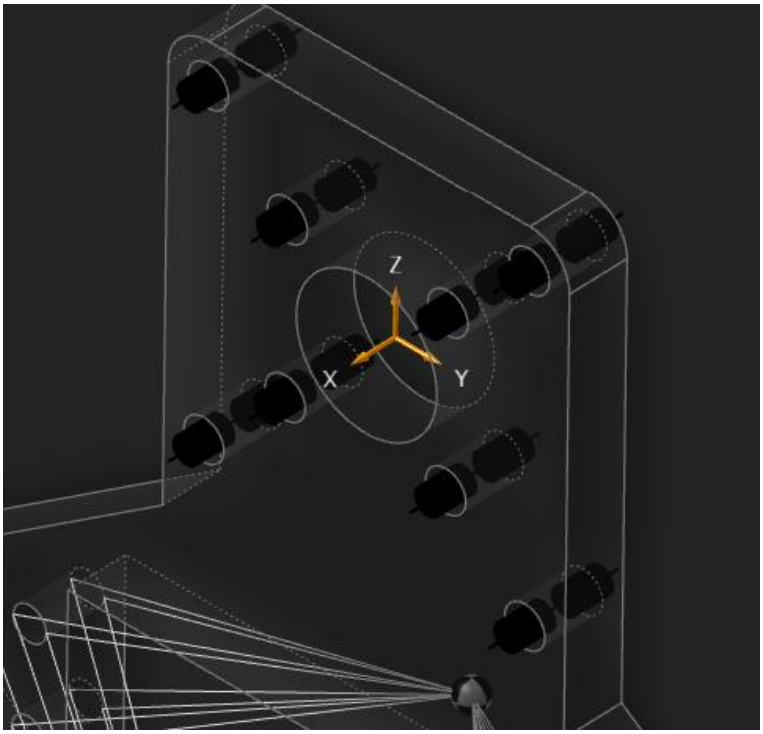

(a)

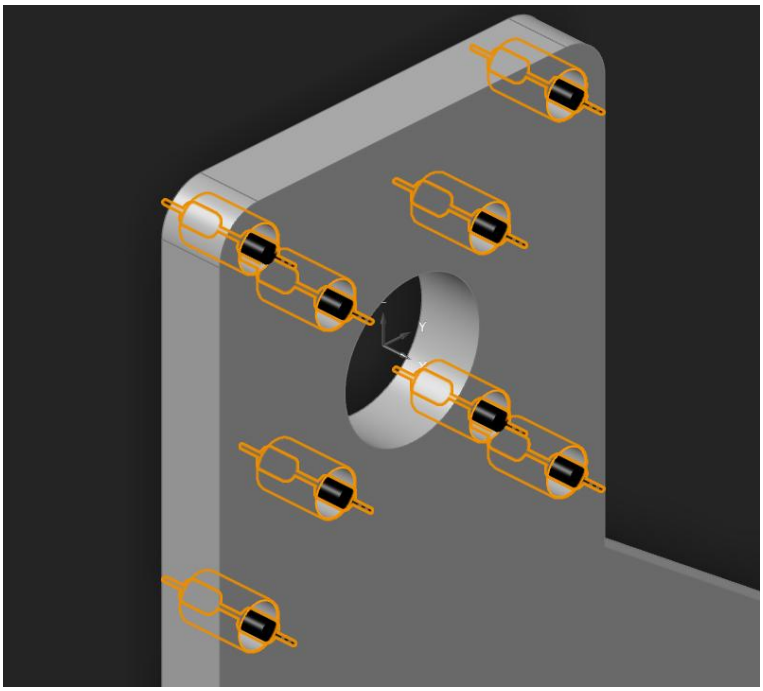

(b)

Figure 60: (a) global coordinate reference system adopted for the frame component; (b) the 8 fixed boundary supports applied to the mounting holes, each restraining all six degrees of freedom.

Frame: masses. Three masses act on the model as in Figure 61: **(1)** the distributed mass of the Frame itself, derived from the material density; **(2)** a point mass of 170 [g] (the 150 [g] payload plus 20 [g] of residual hardware) located at the maximum extension point of the gripper, which maximises the lever arm and is therefore the worst case position; **(3)** the servo mass, placed at its centroid, which under uniform density coincides with its centre of mass.

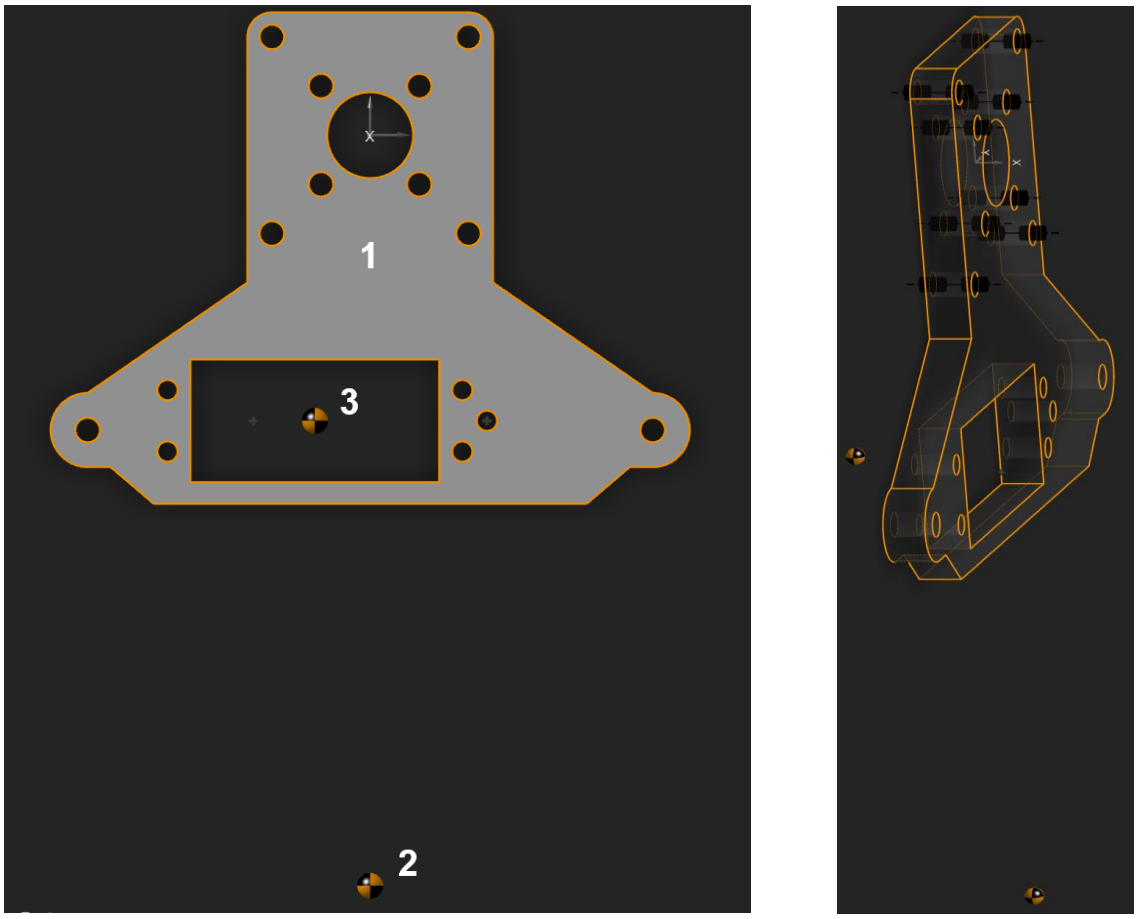

Figure 61: Front (left) and side (right) views of the point-mass positions relative to the Frame: (1) Frame mass; (2) augmented payload mass (170 [g]); (3) servomotor mass.

Frame: connectors. The point masses and the load application points are tied to the Frame through connectors. Four connectors are defined, as in Figure 62.

- **(1)** the 170 [g] point mass is connected to the 4 servo mounting holes, the 2 holes joining the Links to the Frame and the slave gear hole, representing the load path from the payload to the Frame through the intermediate components, modelled as rigid.
- **(2, 3)** since the gears are 3 [mm] thick, the forces on the two gear axes (servo driven gear and slave gear) are applied at 1,5 [mm] above the Frame face, that is at the mid thickness of the gear, as a mean application point. These two points are connected respectively to the 4 servo mounting holes (driven gear) and to the slave gear hole. No connector is needed at the Link to Frame attachments: the Link engages both faces of the Frame, so the forces act symmetrically on the same fastener, no uncompensated torsional moment arises about that axis, and only the net force along Z remains.
- **(4)** the servo centroid is connected to the 4 servo mounting holes, representing the physical attachment of the servomotor to the Frame.

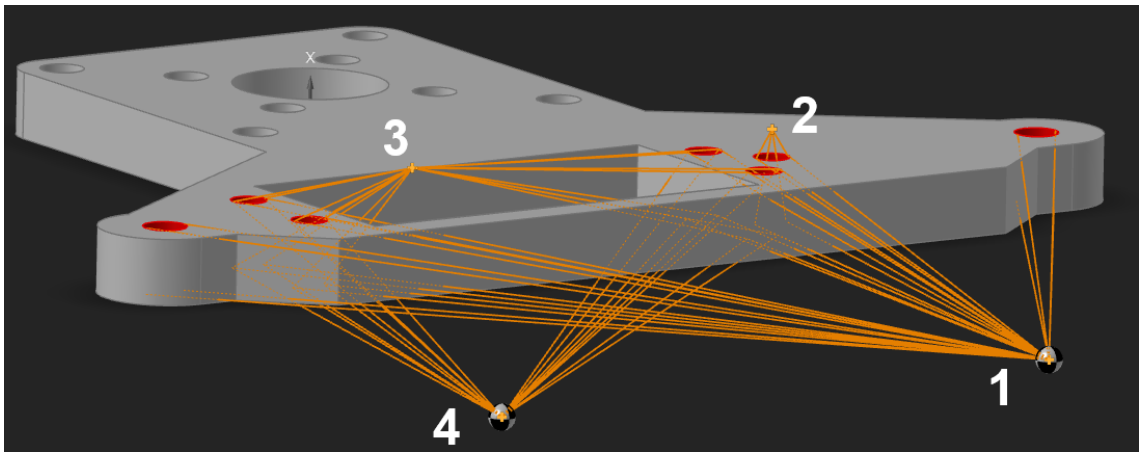

Figure 62: The four connectors and their attachments: (1) the 170 [g] payload mass tied to the four servo mounting holes, the two Link-to-Frame holes and the slave-gear hole; (2) the driven-gear axis to the four servo mounting holes; (3) the slave-gear axis to the slave-gear hole; (4) the servo centroid to the four servo mounting holes.

Frame: load cases. Three load cases are analysed:

- **Gravity.** An acceleration of 9,81 [m/s²] acts on the three masses along −Z, as in Figure 63.

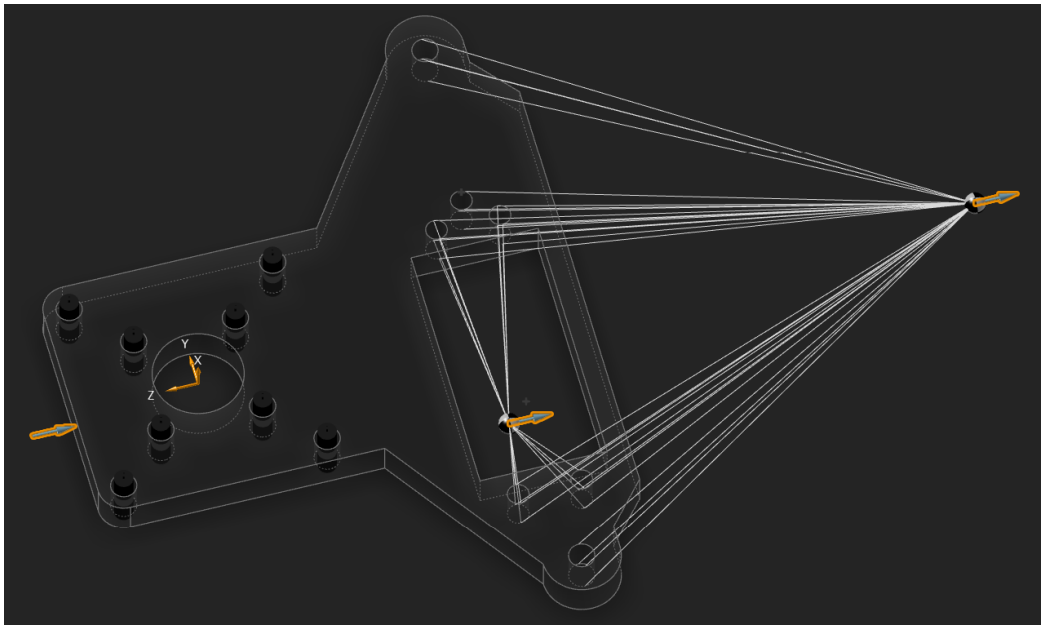

Figure 63: The 9,81 [m/s²] gravitational acceleration acting on the three masses of the system along the −Z direction.

- **Acceleration along X.** With the Frame held fixed at its 8 supports, an acceleration of 30 [m/s²] is applied along X to the three masses, Figure 64.

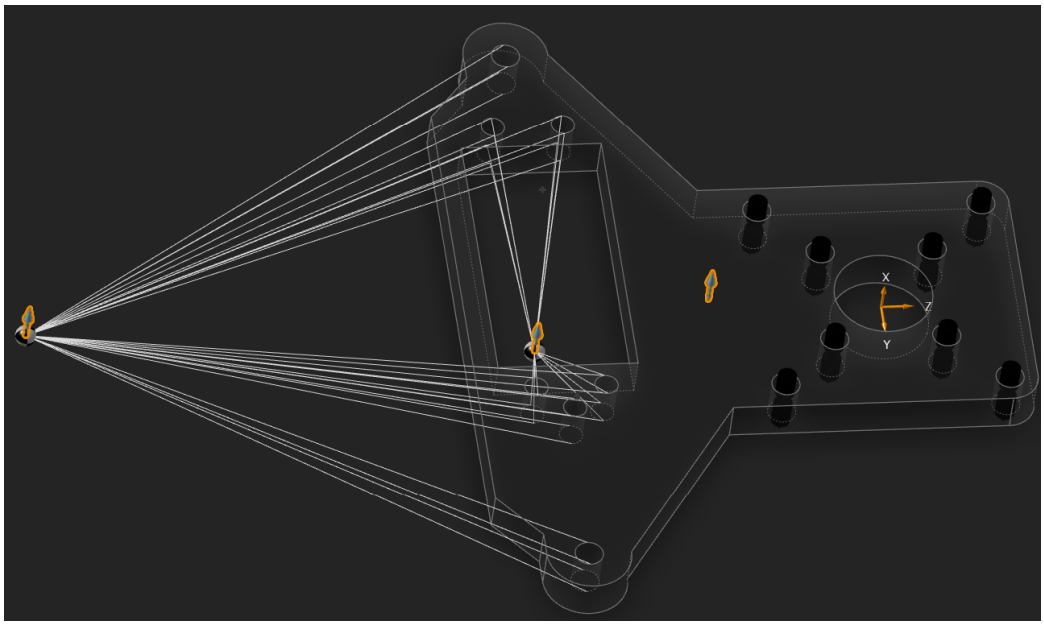

Figure 64: The 30 [m/s²] acceleration acting on the three masses of the system along the +X direction.

- **Servo stall.** The forces derived in Appendix G from the static equilibrium of the mechanism at stall act on the Link and gear arm holes, together with the servo reaction torque on its 4 mounting holes, as in Figure 65.

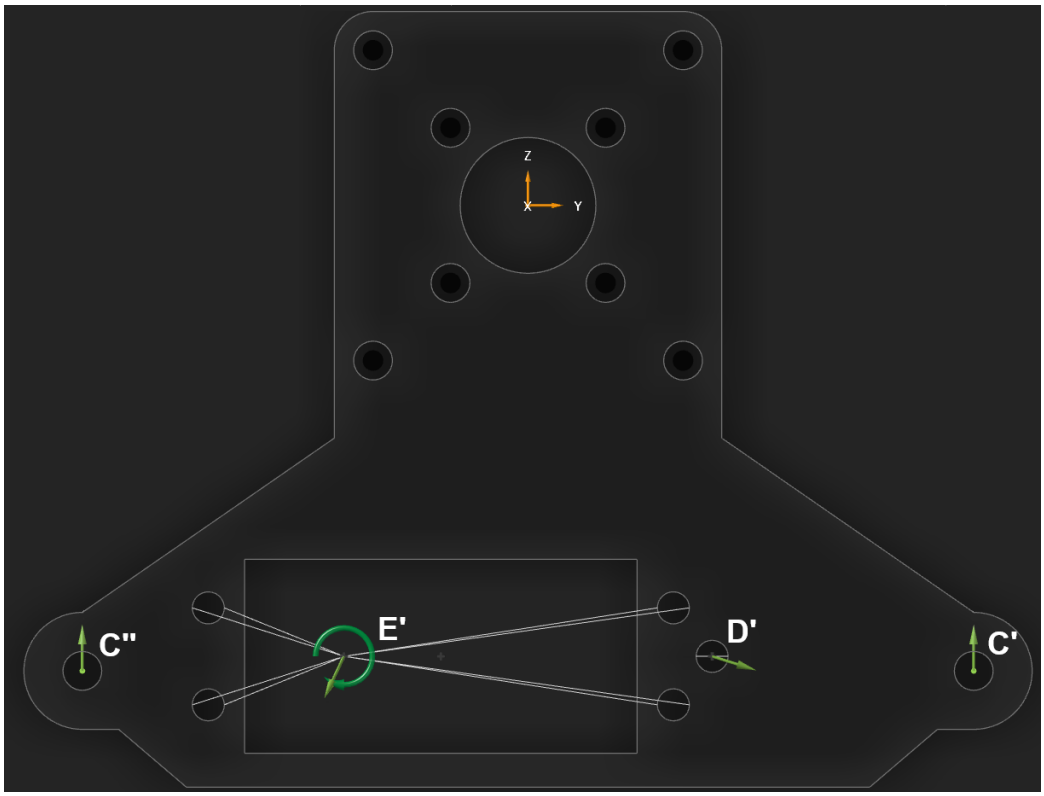

Figure 65: Servo-stall forces applied to the Frame, computed in Appendix G.

Table 12 collects these actions, referred to the points labelled in Figure 65 and resolved along the axes of the reference system. As a consistency check the force system is self-equilibrated: the Z components sum to $2 \times 14{,}298 - 3{,}686 - 24{,}910 = 0$ [N] and the Y components cancel pairwise, confirming that the applied loads represent an internal equilibrium state of the mechanism and that the supports react only to the residual moment.

Point on the Frame	Component	Value / Unit
C′ = C″ (Link holes, one per side)	FZ	+14,298 [N]
D′ (slave gear axis)	FY	+11,618 [N]
	FZ	−3,686 [N]
E′ (servo driven gear axis)	FY	−11,618 [N]
	FZ	−24,910 [N]
	TX	−402,0 [N·mm]

Table 12: Servo stall load set applied to the Frame, computed in Appendix G. The sign of the servo reaction torque TX was verified against the deformation pattern of a preliminary run.

Claw: reference system and constraints. The reference system of the Claw is centred at hole A, as in Figure 66. The finger is pinned at A and at B. The pair at A is a revolute: it leaves free only the rotation about the pin axis and locks the other two rotations and the three translations. At B the support reproduces the action of the Link which, as a two-force member, restrains the finger only along the Link axis (Z), the direction that loads the Link in tension or compression: the rotation about the pin axis and, for small displacements, the translation along Y remain free. The Claw is therefore grounded at the two pins that physically connect it to the gear arm (A) and to the Link (B).

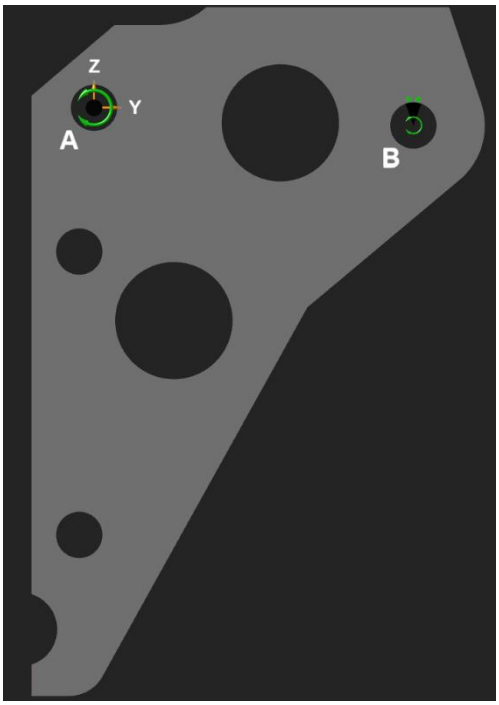

Figure 66: Inspire view of the Claw, showing the two revolute pairs at A and B and the reference frame centred at A.

Claw: masses and forces. Two masses act on the Claw: its own distributed mass, from the material density, and the share of the gripped object that the finger reacts at the pad. The object share cannot be modelled, as on the Frame, by a point mass placed at a distance from the finger, because that representation would be misleading: an offset point mass under acceleration would introduce a bending moment that in reality is cancelled by the symmetric action of the opposite finger, and the object exchanges actions with the finger only through the contact surface of the pad. The three actions are therefore applied directly on the pad through a connector defined on its surface, which reproduces the real load path and avoids introducing any fictitious moment. The Claw own mass is instead applied as a body load through the gravity and acceleration fields, which Inspire integrates over the geometry.

Claw: load cases. Three load cases are analysed, mirroring those defined for the Frame:

- **Gravity. The gripped mass m_p** = 170 [g] under g = 9,81 [m/s²] produces a weight

$$\mathbf{W = m\_p \cdot g = 0{,}170 \cdot 9{,}81 = 1{,}668\ [N]}$$

Clamped between the two opposing pads, the object is held by friction; by symmetry each finger reacts half of its weight as a force tangential to the pad:

$$\mathbf{F\_t = W / 2 = 1{,}668 / 2 = 0{,}834\ [N]}$$

The grip does not slip as long as this tangential force stays within the available friction. With the gripping force F = 7,755 [N] acting as the normal force, the condition is

$$\mathbf{F\_t \leq \mu \cdot F \ \rightarrow\ \mu \geq F\_t / F = 0{,}834 / 7{,}755 = 0{,}11}$$

A value of 0,11 is well within reach: an ASA on ASA contact already gives about 0,25 to 0,35 [18], and the rubber or elastomeric facing usually applied to gripper pads gives more. The gravity case is therefore not critical for retention. The corresponding load application on the Claw pad is shown in Figure 67.

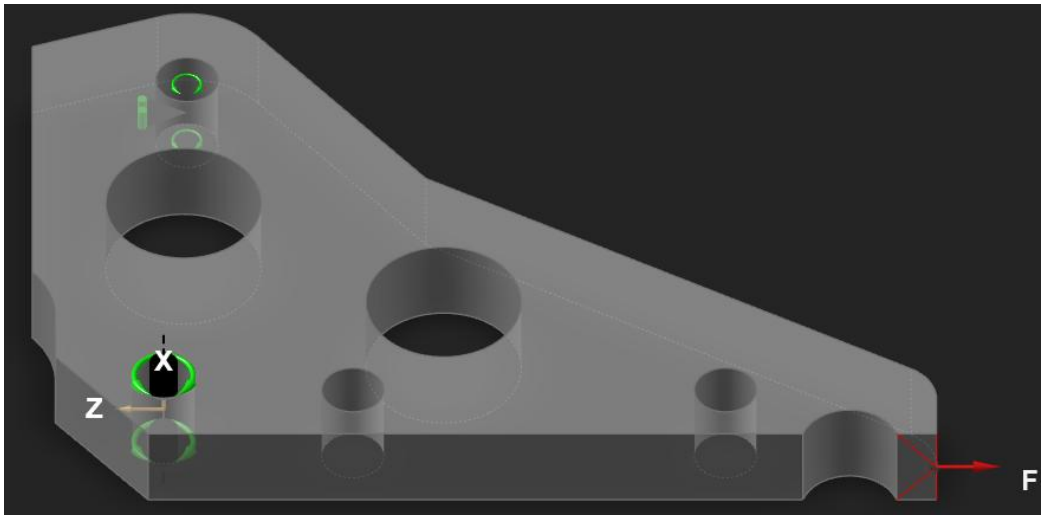

Figure 67: Application on the Claw pad of the gravitational force due to half the object weight.

- **Acceleration along X. The inertial action follows** the same logic. Under the peak acceleration a = 30 [m/s²] adopted for the Frame, the object inertia is

$$\mathbf{F\_in = m\_p \cdot a = 0{,}170 \cdot 30 = 5{,}10\ [N]}$$

shared between the two pads, so each finger reacts

$$\mathbf{F\_acc = F\_in / 2 = m\_p \cdot a / 2 = 2{,}55\ [N]}$$

along X, again tangential to the pad and therefore carried by friction. The no slip condition gives

$$\mathbf{F\_acc \leq \mu \cdot F \ \rightarrow\ \mu \geq F\_acc / F = 2{,}55 / 7{,}755 = 0{,}33}$$

The Claw is assumed to reach this friction coefficient without difficulty; in practice, at most a different polymeric facing would be applied to this surface alone to guarantee the required friction coefficient. The application of this load on the pad is shown in Figure 68.

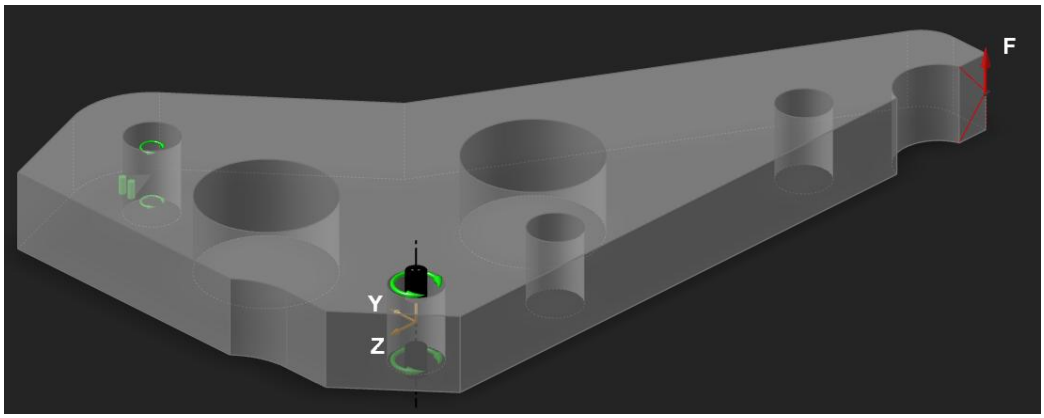

Figure 68: Application of the acceleration load along the X axis on the Claw pad.

- **Gripping (servo stall). The gripping force F = 7,755 [N], derived** in Appendix G, acts outward at the pad as the object resists the closing of the two fingers, while A and B react. It is the largest of the three pad actions and the load case that sizes the Claw, and therefore the one that drives its topology optimisation. Figure 69 shows this action applied to the Claw.

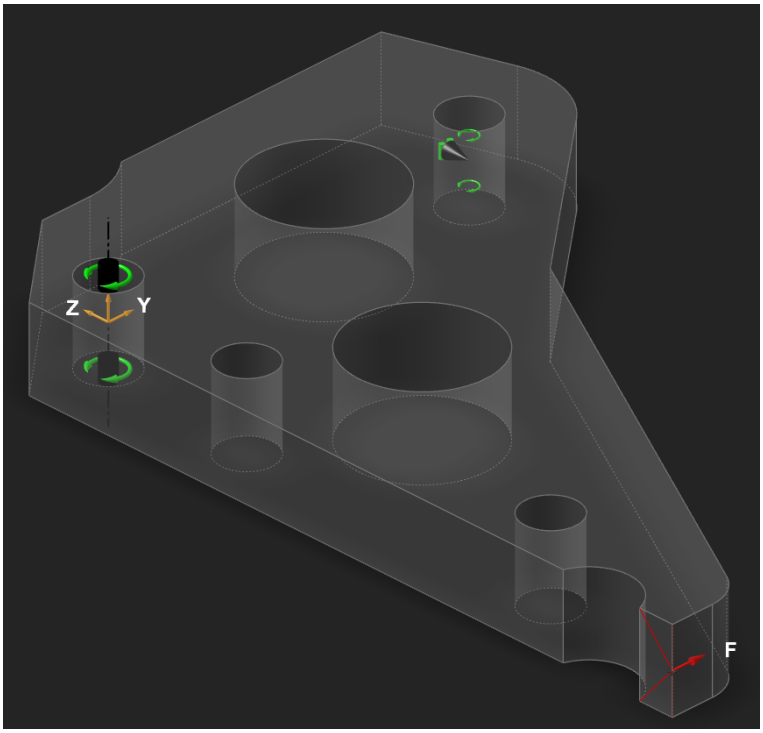

Figure 69: Application on the Claw of the object's reaction to the gripping (closing) force, computed in Appendix G.

Table 13 below summarises all the forces applied at the pad, referred to the Claw reference system.

Action	Symbol	Value / Unit	Direction
Gripping (servo stall)	F	7,755 [N]	Y, normal to the pad
Gravity (object share)	F_t	0,834 [N]	Z, tangential to the pad
Acceleration (object share)	F_acc	2,550 [N]	X, tangential to the pad

Table 13: Actions transmitted by the gripped object and applied at the Claw pad through the surface connector. The Claw own mass is applied separately as a gravity and an acceleration body load. The gripping force is normal to the pad; the gravity and acceleration shares are tangential and are carried by friction.

Reverse engineered models. The same procedure described above, identical in reference systems, supports, connectors, masses and load cases, was applied to the geometry of the Frame and of the Claw reconstructed by reverse engineering from the CT scans of Chapter 4. The nominal CAD and the reconstructed geometry are therefore analysed under the same conditions, so that the deviations reported in the next section reflect the geometry alone and not differences in the modelling. The reconstructed geometries analysed here are shown in Figure 70, with the Frame on the left and the Claw on the right.

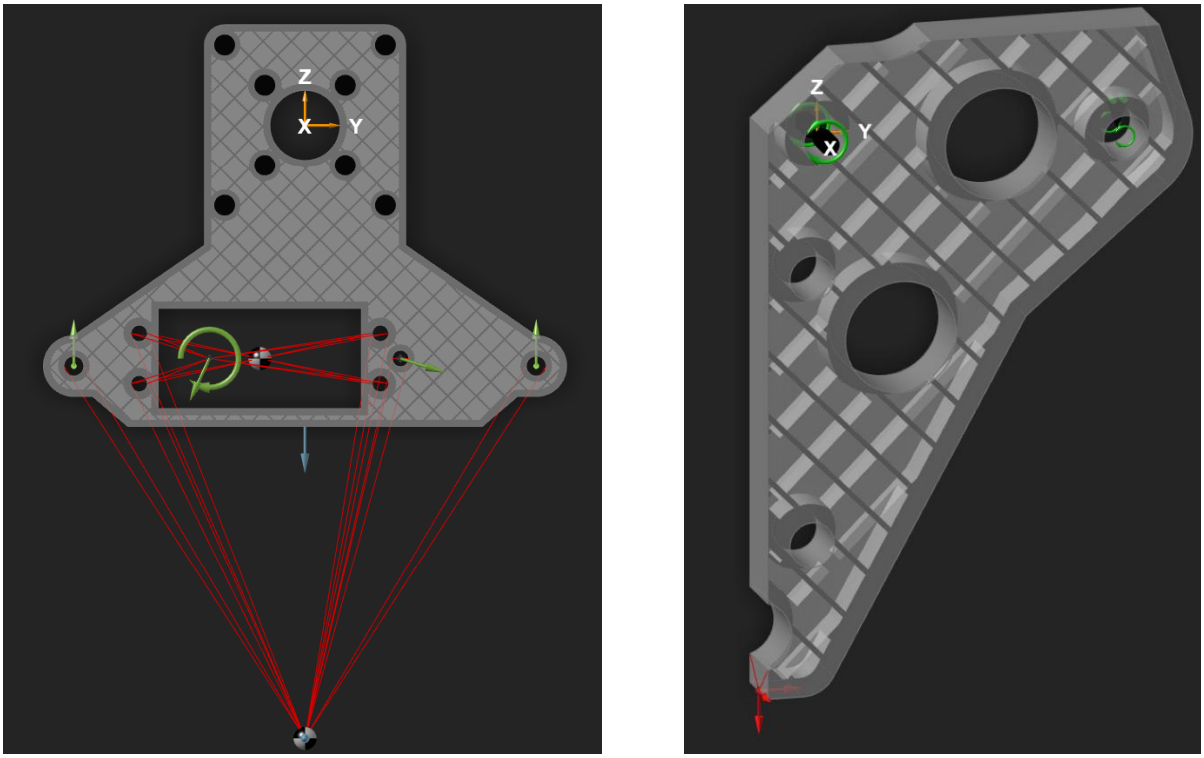

Figure 70: Reverse-engineered version of the Frame (left) and Claw (right), with the same loads and constraints applied as in the nominal models.

5.2 Results on the original CAD

The nominal CAD geometry of the Frame and of the Claw was analysed first, under the three load cases defined in Section 5.1, in order to establish a structural baseline before the geometric deviations of the reconstructed models are introduced in Section 5.3. The full displacement and SF fields are collected in Appendix H so as not to burden the discussion; here, for each component, only a summary table drawn from the appendix data is retained, together with an overall SF and displacement envelope map, in which every point takes the worst value among the three load cases.

Frame. The three load cases of Section 5.1 are applied first to the nominal CAD Frame. The overall SF and displacement envelopes are shown in Figure 71 and Figure 72.

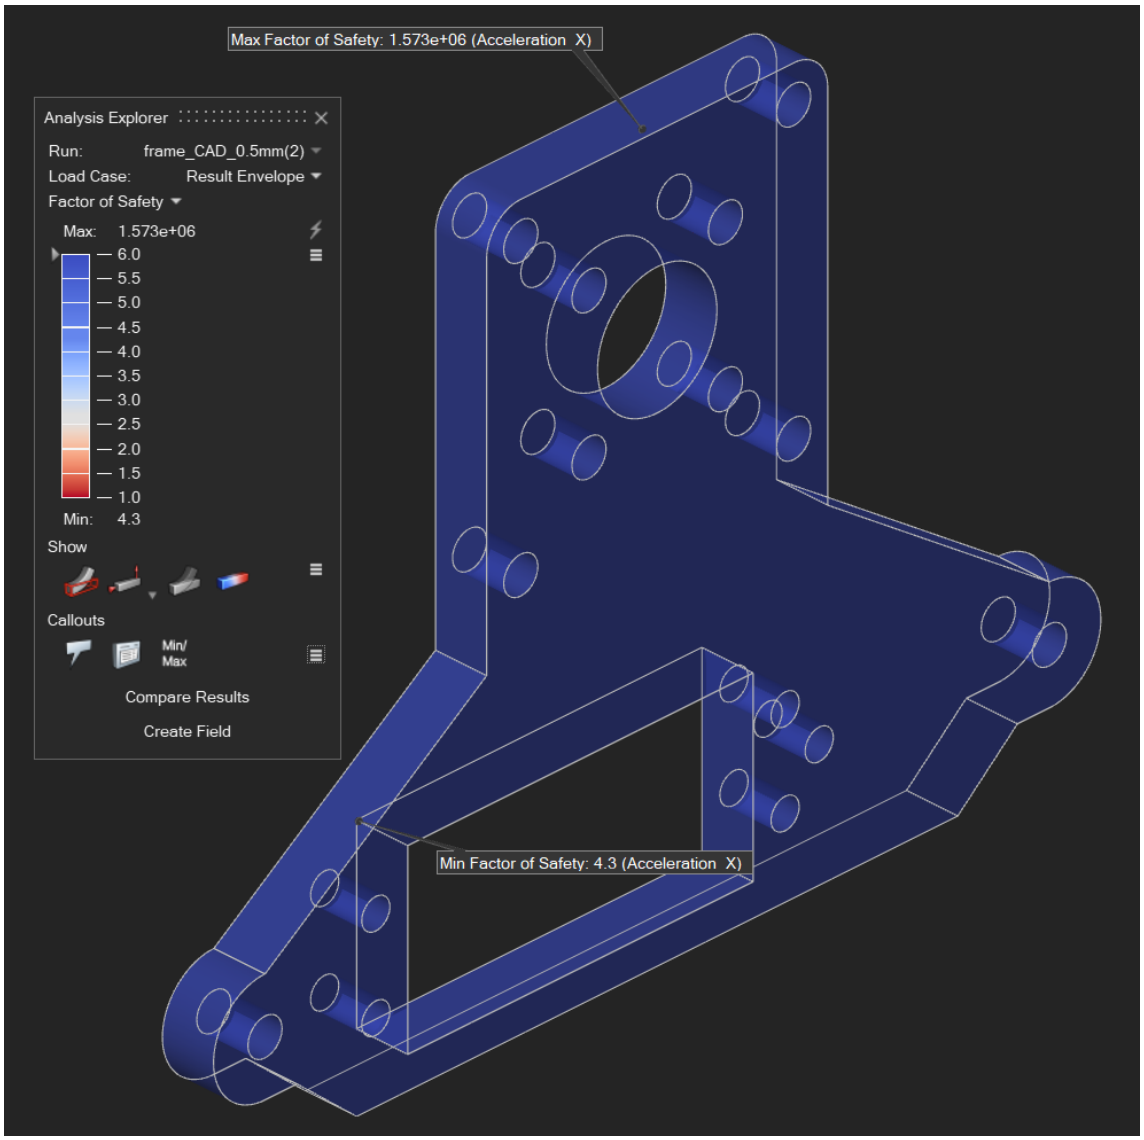

Figure 71: Overall envelope of the 3 cases on the solid Frame (CAD) in terms of SF. Min and Max SF are identified in the figure.

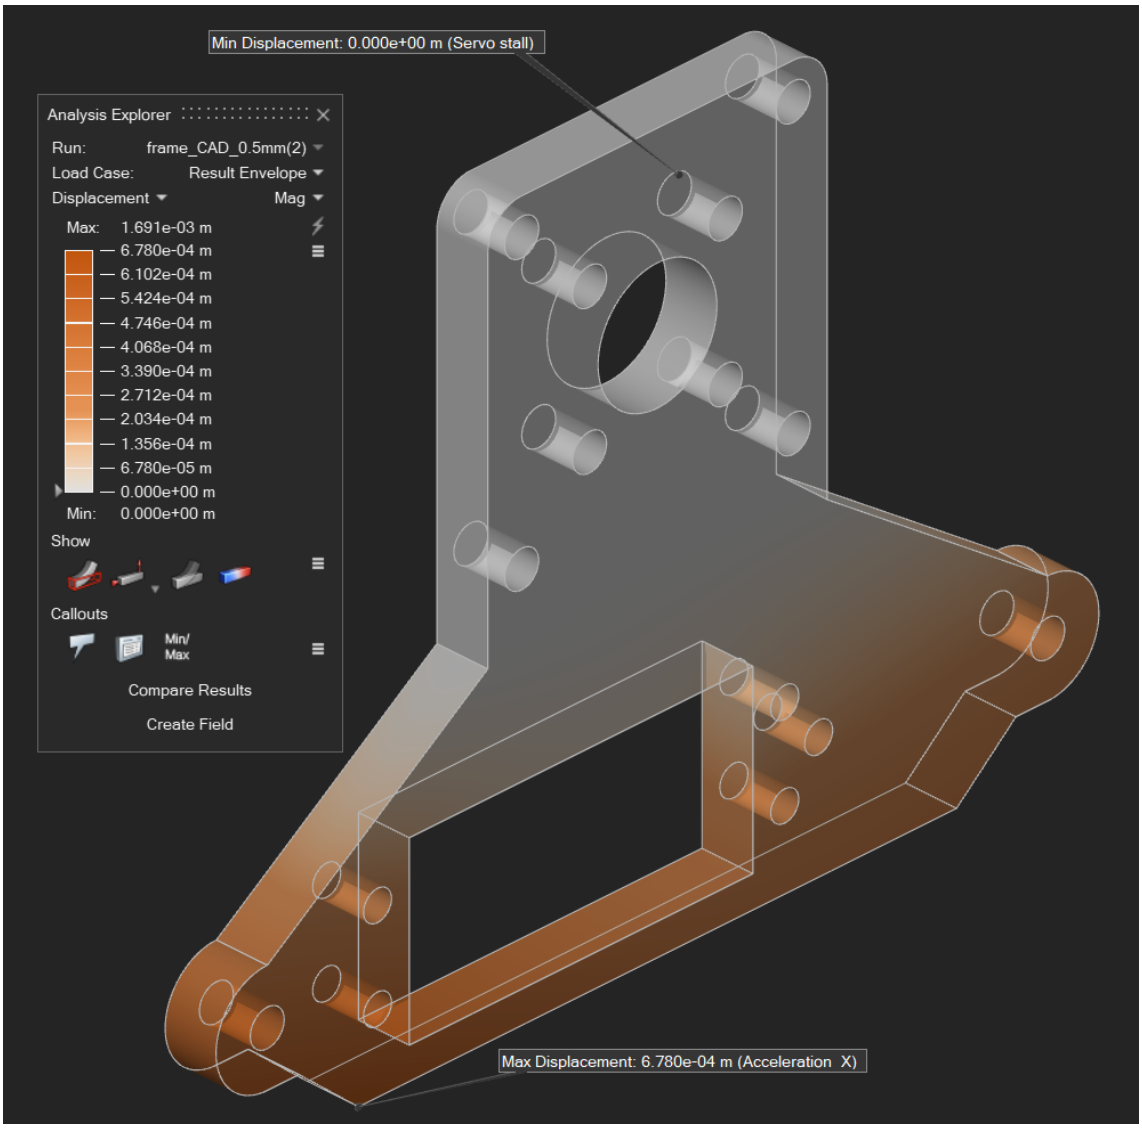

Figure 72: Overall envelope of the 3 cases on the solid Frame (CAD) in terms of displacement from the original position. Min and Max displacement are identified in the figure.

The three load cases for the Frame are consolidated in Table 14.

Load case	SF (min)	δmax [mm]
Gravity	181,4	$8{,}655 \times 10^{-3}$
Acceleration along X	4,3	0,678
Servo stall	11,6	0,213

Table 14: SF and peak displacement of the nominal CAD Frame under the three load cases. The minimum SF governs the verification, and the acceleration along X is the sizing case.

The sizing case for the Frame is the acceleration along X: its minimum SF of 4,3 is far lower than the gravity (181,4) and servo-stall (11,6) values, but still leaves a margin of more than four with respect to the yield strength of the material. The three load cases correspond to distinct operating states and are evaluated independently; in the envelope plots the minimum accordingly coincides with the acceleration case. The largest displacement of the whole set is 0,678 [mm], so the Frame is globally stiff.

Claw. The three load cases are now applied to the Claw; the overall envelopes are shown in Figure 73 and Figure 74.

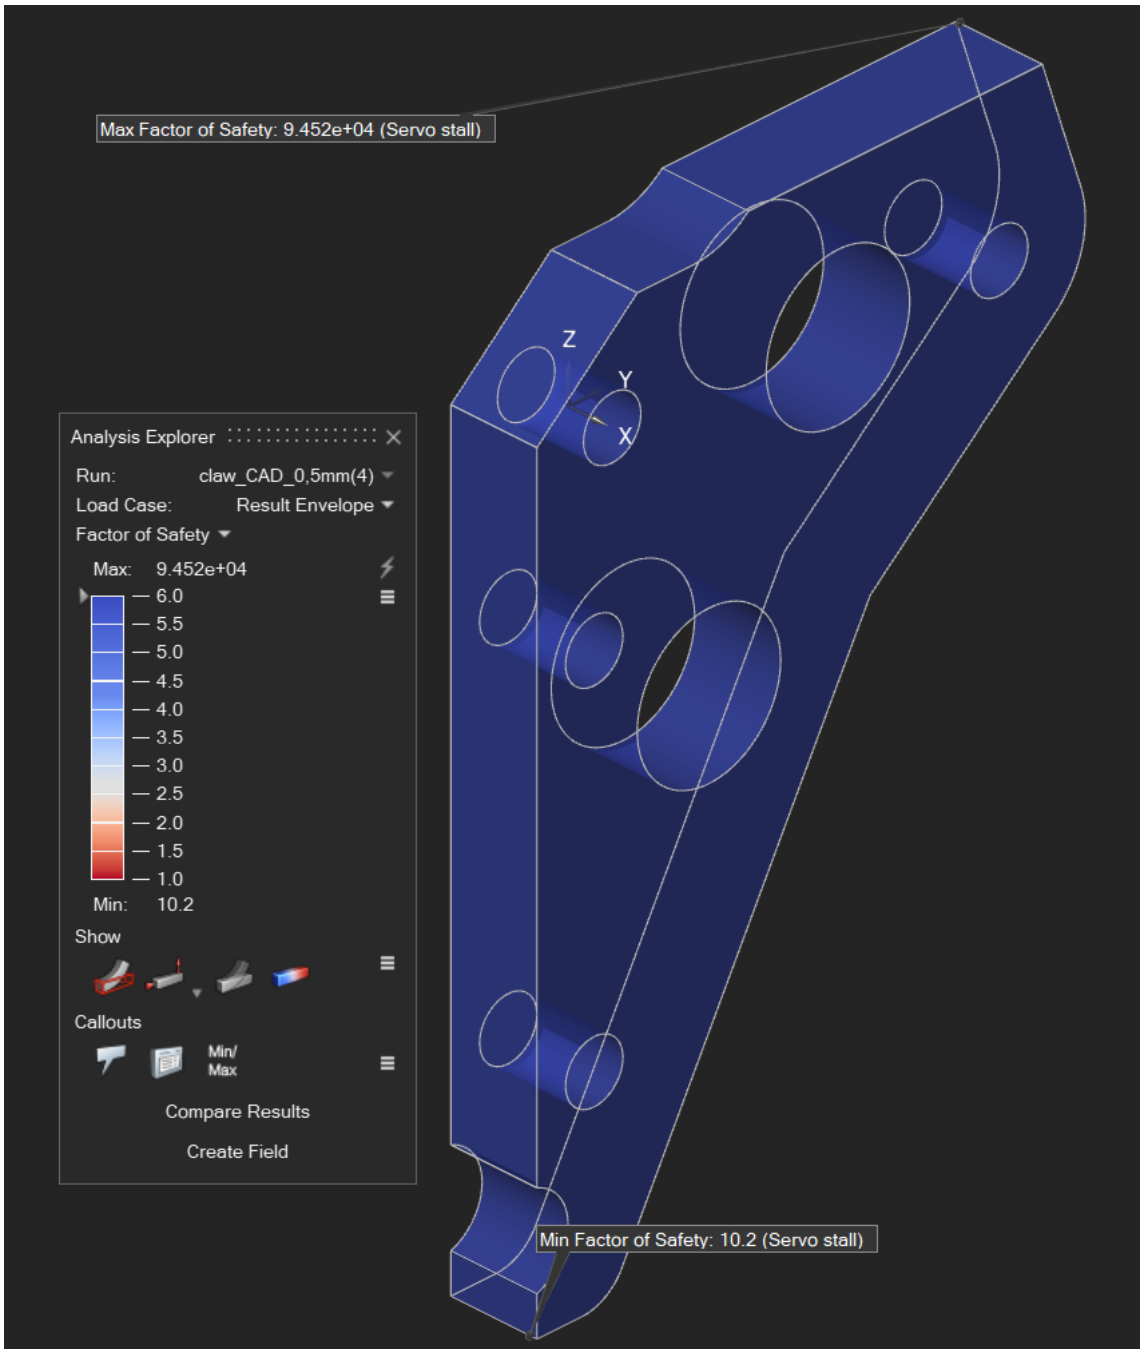

Figure 73: Overall envelope of the 3 cases on the solid Claw (CAD) in terms of SF. Min and Max SF are identified in the figure.

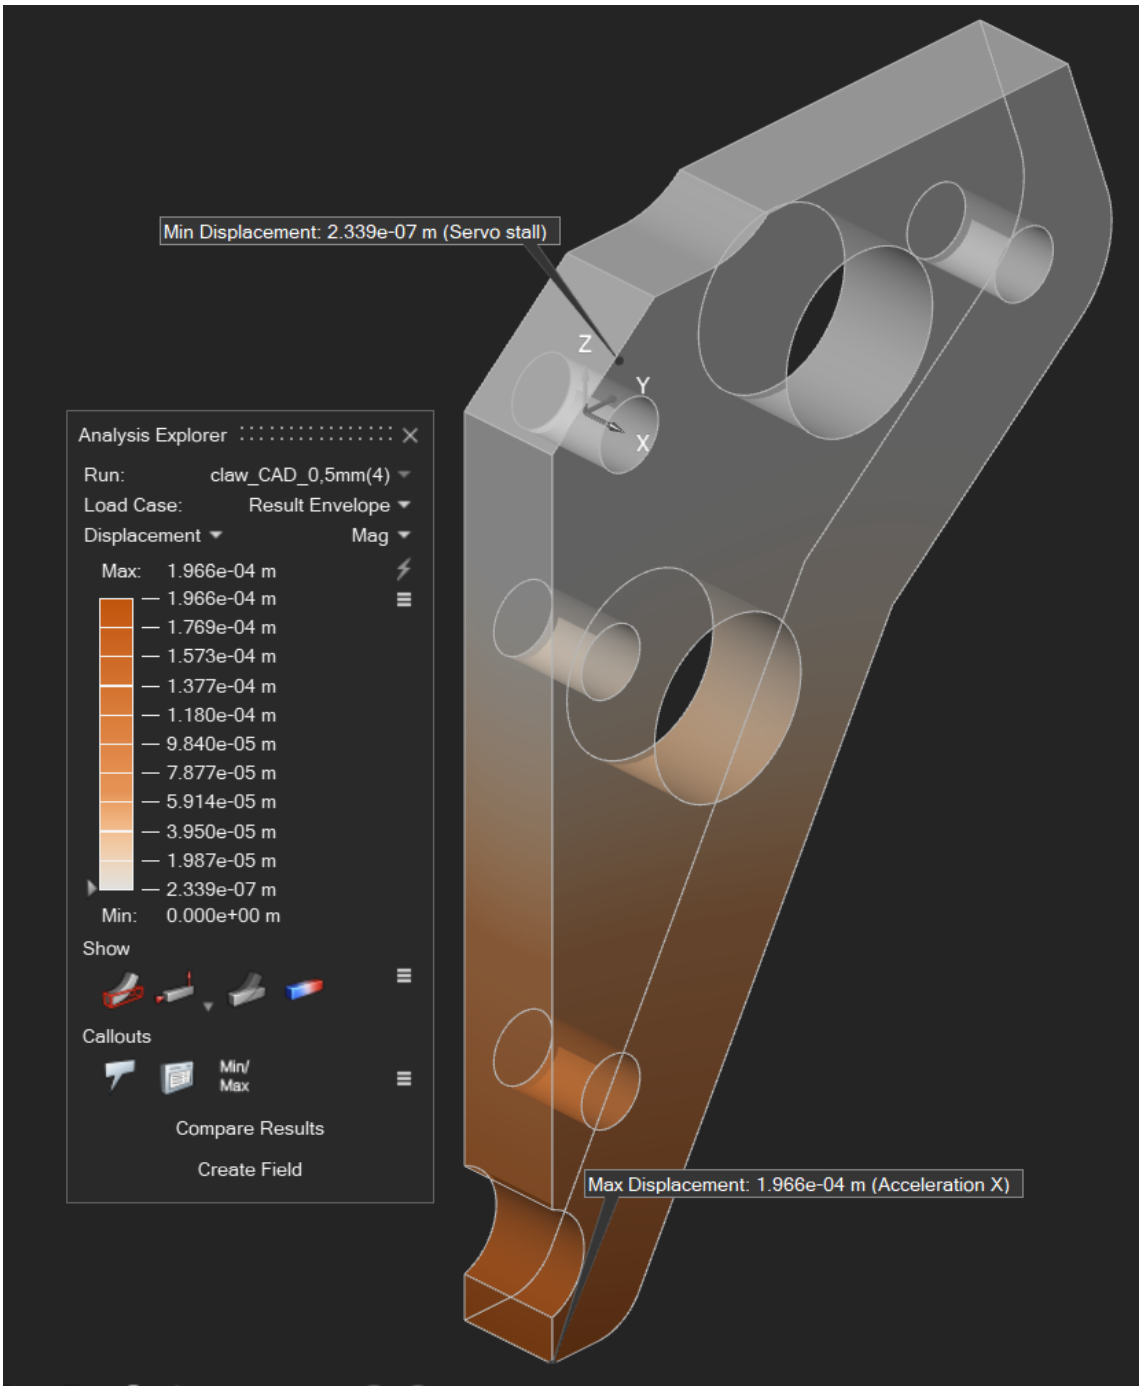

Figure 74: Overall envelope of the 3 cases on the solid Claw (CAD) in terms of displacement. Min and Max displacement are identified in the figure.

The three load cases for the Claw are consolidated in Table 15.

Load case	SF (min)	δmax [mm]
Gravity	271,1	$2{,}928 \times 10^{-3}$
Acceleration along X	26,6	0,197
Servo stall	10,2	0,057

Table 15: SF and peak displacement of the nominal CAD Claw under the three load cases. The servo-stall (gripping) case yields the lowest SF and drives the topology optimisation of Chapter 6; under acceleration the true limit is the pad–object friction rather than the strength of the finger.

Unlike the Frame, the lowest SF for the Claw occurs in the servo-stall (gripping) case, at 10,2, while the acceleration along X gives 26,6; both remain well above any risk of yielding or failure.

In the envelope the minimum coincides with the servo-stall value, and the largest displacement (0,197 [mm]) stays well below half a millimetre.

In every case it should be remembered that the true limit under acceleration is not the structural strength of the Claw but the claw–object friction force, which can be increased in three ways:

1. By operating in a less critical configuration, for example with the object resting against the flat (normal) faces of the gripper, so that the grip force is no longer the dominant action and the governing actions become instead the torsional moments acting on the Claw and the contact pressure over its surface.
2. By increasing the friction coefficient through an elastomeric facing on the pads. Even in the ideal limit, however, the available tangential capacity only tends to the full normal force of 7,755 [N]; this comfortably covers the 2,55 [N] demanded by the acceleration case, although the solution remains less robust than removing the dependence on friction altogether, as in the first option.
3. By replacing the servo with the original Proton unit, which would raise the available torque from 402 [N·mm] to 530 [N·mm] (+31,8 [%]). The gripping force would increase by the same proportion, lowering the SF of the servo-stall case, already the governing minimum, from 10,2 to approximately 7,8, a value that still leaves a comfortable margin.

In summary, for the solid parts the acceleration along the X axis sizes the Frame, whereas the minimum SF of the Claw occurs in the servo-stall case, with the acceleration case limited in practice by the claw–object friction. It is now of interest to examine how these same load cases affect a part that is not solid but built with an FDM infill, the subject of Section 5.3.

5.3 Results on RE CT and Discussion

The same load cases are now re-applied to the geometry of the Frame and of the Claw reconstructed from the CT data of Chapter 4, and the results are compared with the nominal CAD baseline of Section 5.2; the full maps are again collected in Appendix H. The comparison isolates the effect of the real, as-printed geometry, including the internal FDM infill captured by the scan, on both the peak displacement and the SF.

Frame. The displacement and SF results for the reconstructed Frame under the three load cases are summarised below; the overall SF and displacement envelopes are shown in Figure 75 and Figure 76.

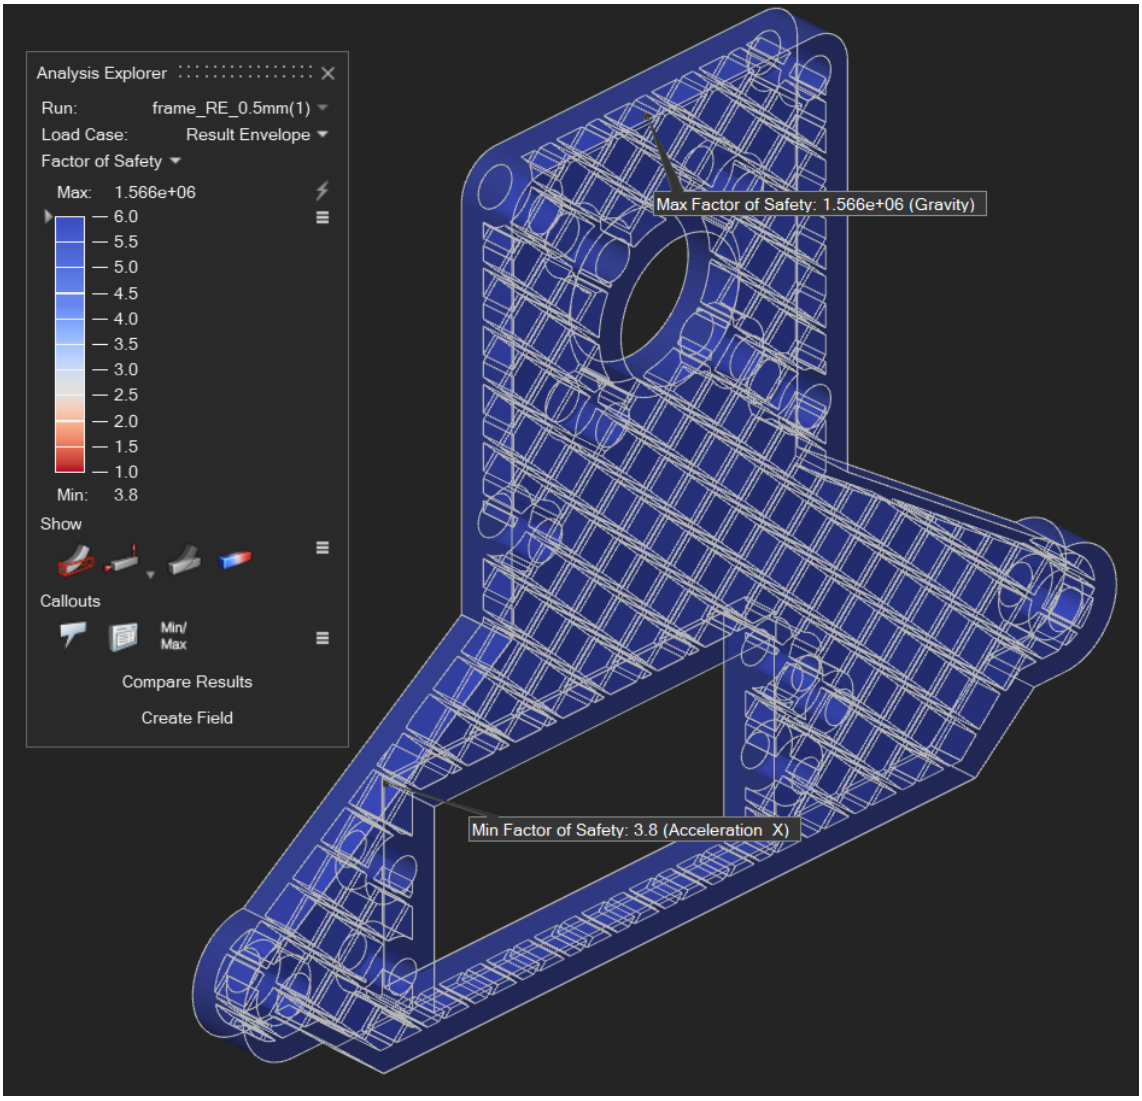

Figure 75: Overall envelope of the 3 cases on the reconstructed Frame (RE, with infill) in terms of SF. Min and Max SF are identified in the figure.

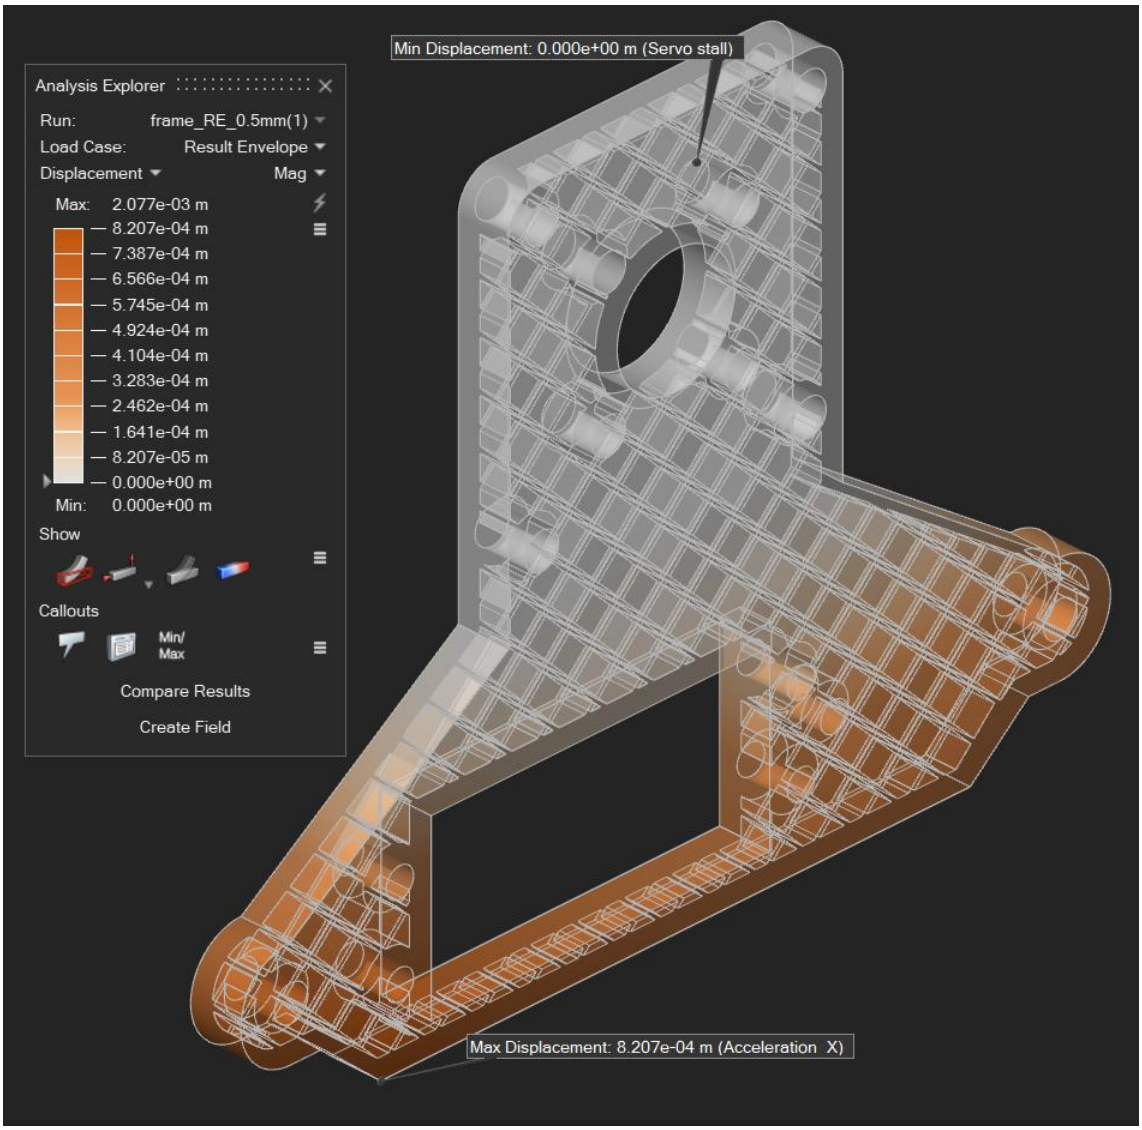

Figure 76: Overall envelope of the 3 cases on the reconstructed Frame (RE, with infill) in terms of displacement. Min and Max displacement are identified in the figure.

The CAD and reconstructed results for the Frame are compared in Table 16.

Load case	SF CAD	SF RE	ΔSF [%]	δmax CAD [mm]	δmax RE [mm]	Δδ [%]
Gravity	181,4	136,5	−24,8	$8{,}655 \times 10^{-3}$	$1{,}057 \times 10^{-2}$	+22,1
Acceleration along X	4,3	3,8	−11,6	0,678	0,821	+21,1
Servo stall	11,6	9,4	−19,0	0,213	0,256	+20,2

Table 16: CAD versus reconstructed (CT) comparison for the Frame, minimum SF and peak displacement under the three load cases, with the percentage deviation of the reconstructed model relative to the CAD baseline. Displacement deviations are approximate, owing to the coarse read-out of the contour plots.

As for the CAD model, the acceleration along X remains the sizing case for the reconstructed Frame, with a minimum SF of 3,8, close to the CAD value of 4,3 and still far from yielding. The overall envelope of the three cases again returns this value as its minimum, and the largest displacement of the set is 0,821 [mm], the largest of the whole study but still small in absolute terms.

Claw. The same comparison is repeated for the reconstructed Claw; the overall envelopes are shown in Figure 77 and Figure 78.

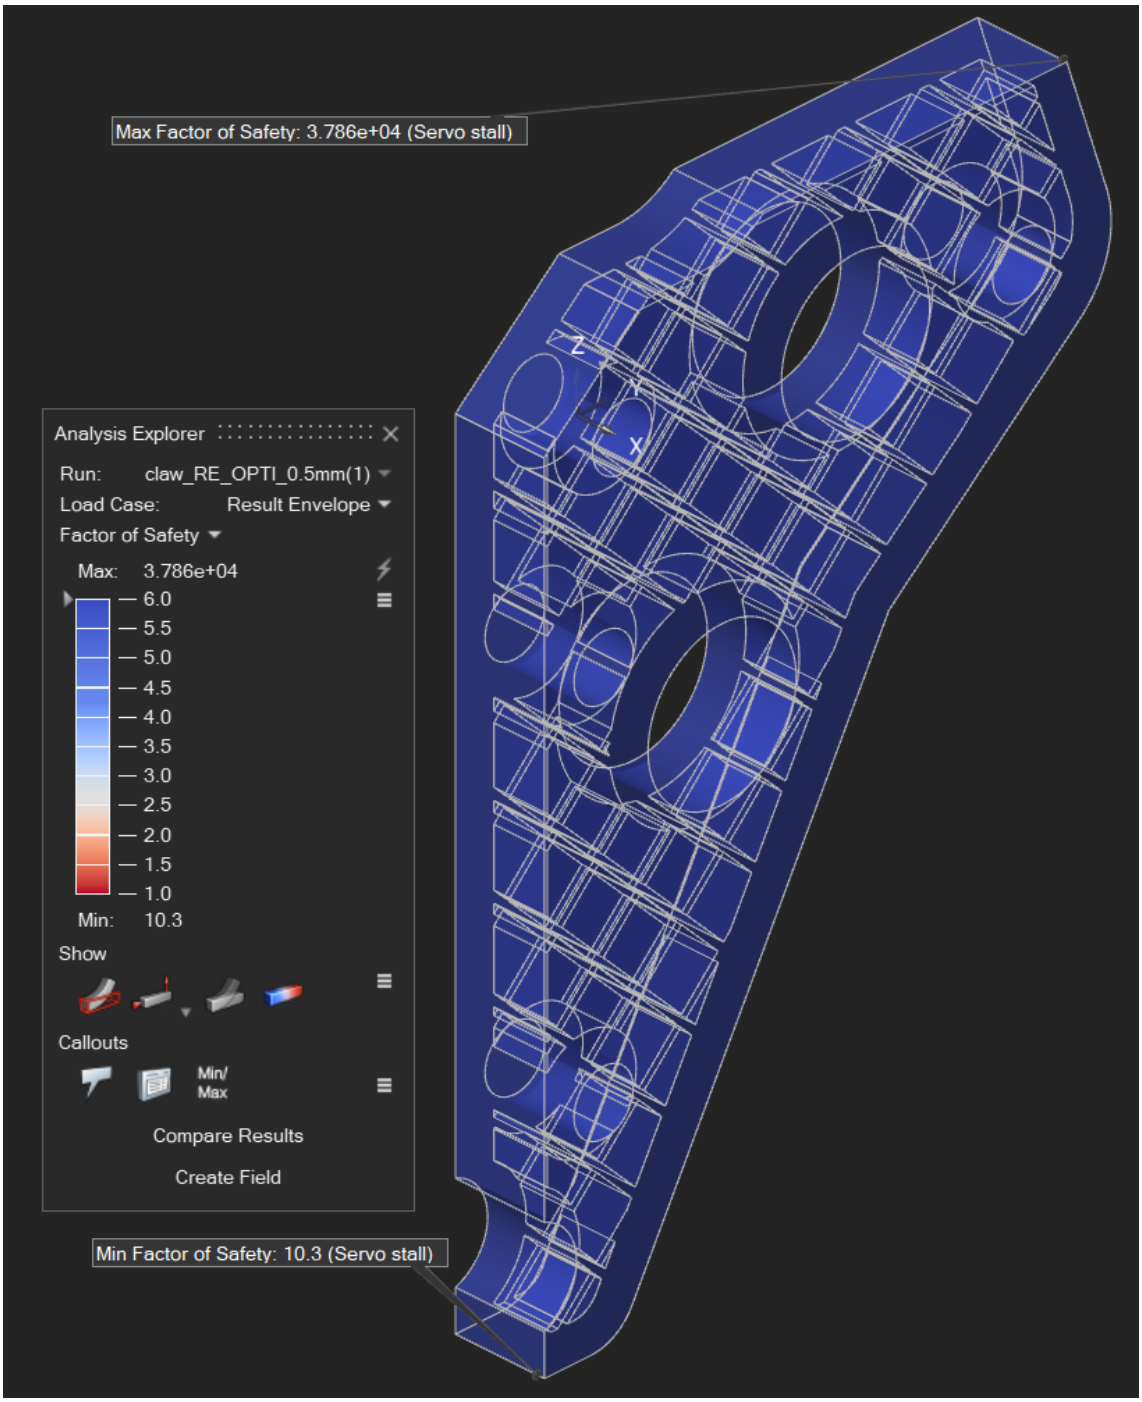

Figure 77: Overall envelope of the 3 cases on the reconstructed Claw (RE, with infill) in terms of SF. Min and Max SF are identified in the figure.

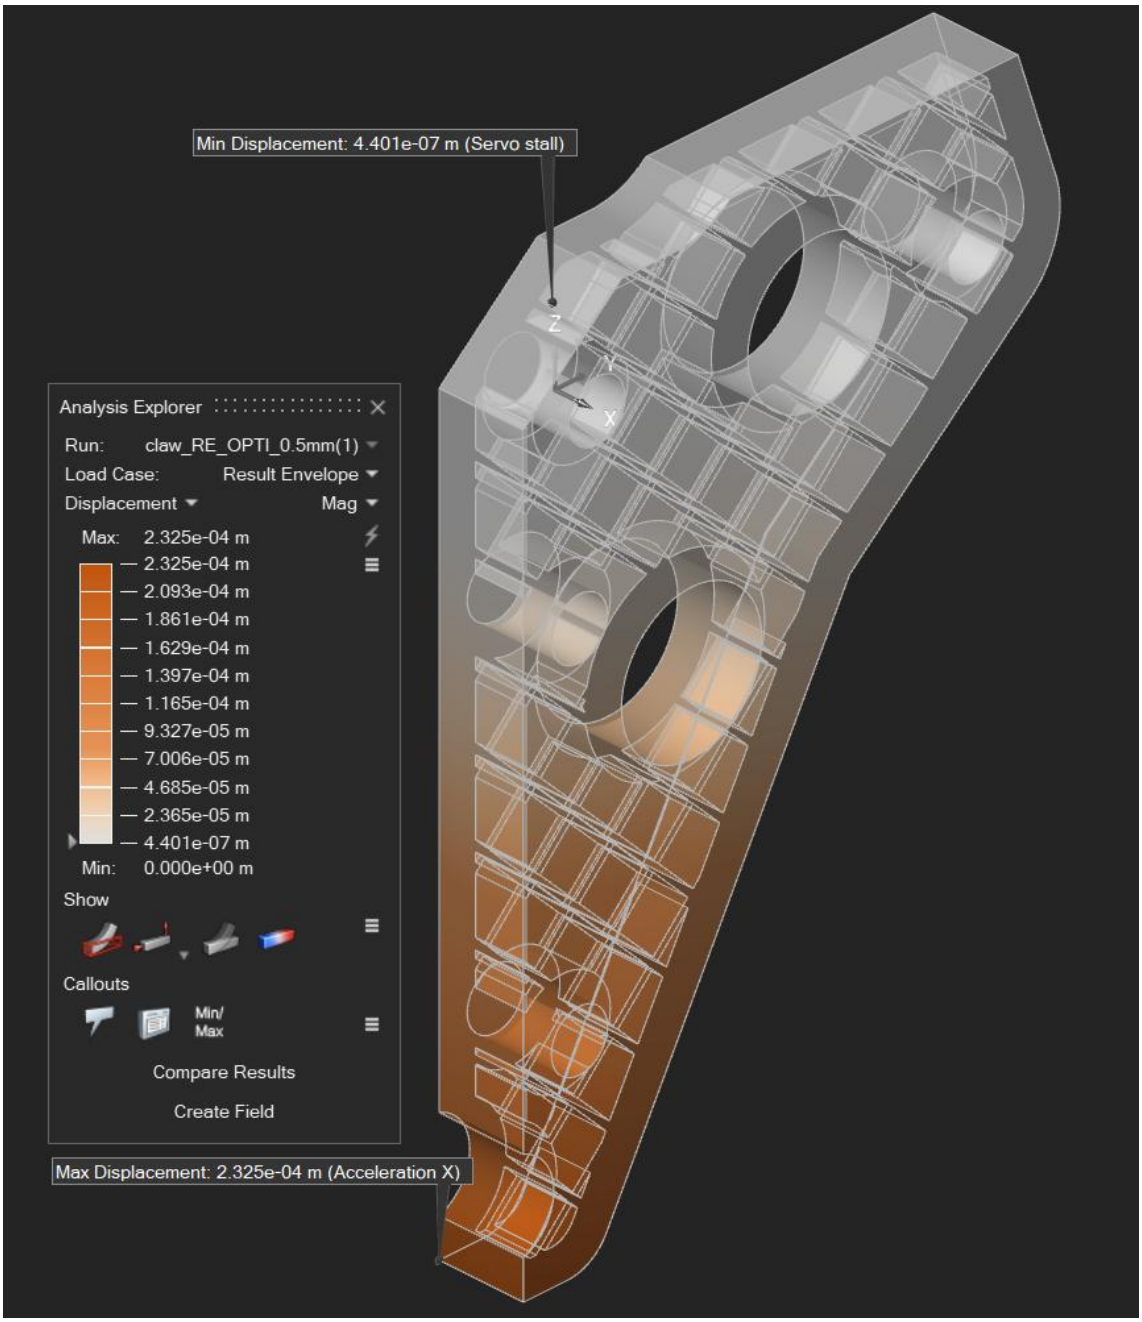

Figure 78: Overall envelope of the 3 cases on the reconstructed Claw (RE, with infill) in terms of displacement. Min and Max displacement are identified in the figure.

The CAD and reconstructed results for the Claw are compared in Table 17.

Load case	SF CAD	SF RE	ΔSF [%]	δmax CAD [mm]	δmax RE [mm]	Δδ [%]
Gravity	271,1	10,3	−96,2	2,928 × 10^{-3}	3,766 × 10^{-3}	+28,6
Acceleration along X	26,6	25,0	−6,0	0,197	0,233	+18,3
Servo stall	10,2	10,3	+1,0	0,057	0,077	+34,8

Table 17: CAD versus reconstructed (CT) comparison for the Claw, minimum SF and peak displacement under the three load cases, with the percentage deviation relative to the CAD baseline. The Claw is essentially insensitive to the infill, with deviations within a few per cent except under gravity.

Comparing the two sets of results in Table 16 and Table 17 explains the observed trends:

- Under gravity the reduction in SF was expected: the load acts in tension along the weakest direction of the part, effectively stretching all the internal square infill cells, which are inevitably less resistant than solid material.
- Under acceleration the practically identical result was also expected. From the neutral-axis principle for a part in bending, the fibres in tension and compression are those on the two faces normal to the X axis; the closer a fibre lies to the neutral axis, the less it is stressed. This is why the critical points lie on the outer faces and not in the infill, and why the infill does not change the result.
- Under servo stall a torsional moment is essentially applied to the fibres, which redistribute the load elastically through the interior. Paradoxically, the fibres most "at risk" remain in every case the longest ones rather than the supports, which tend to be solid surfaces; here too the infill is of little relevance.
- Finally, the finite-element discretisation itself introduces an approximation. As noted earlier, these values should be read with some judgement: a residual numerical error exists which is difficult to quantify but remains within a relatively small neighbourhood of the reported results.

It can therefore be concluded that the infill is not the principal weakness of the 3D-printed parts of this end-effector.

What the analysis cannot capture, however, are the sources of error and defect that affect the real parts:

- First, the fusion of the deposited material in FDM is only partial and generally non-homogeneous, so the resulting bulk material is imperfectly formed.
- Voids and printing imperfections cannot easily be modelled, yet they introduce uncertainty and produce questionable local load distributions that may fail.
- Because of this partial fusion and the layer-by-layer nature of FDM, the part tends to be weaker in the direction normal to the build plate, in this case the X direction.

These defects act on the material, not on the geometry: the reconstruction captures the printed shape, including the infill, but the material card remains the isotropic bulk one. Of them, the inter-layer weakness is the one that can be bounded quantitatively: tensile tests on standardised printed specimens (coupons) reported in the literature place the inter-layer strength some 20 [%] to 40 [%] below the bulk value, a knock-down that is not represented and matters chiefly in the acceleration case, where the load crosses the layers. Even reading the results as upper bounds in that direction, a 30 [%] reduction would still leave the governing safety factors at about 3,0 for the CAD Frame and 2,7 for the reconstructed one, with every other case above 6,5. The analyses are therefore considered sufficient to proceed.

6. Topology Optimisation

6.1 Fundamentals of Topology Optimisation

Structural efficiency is a recurring requirement in mechanical design, driven by the need to reduce cost, resource consumption and environmental impact. As shown by the structural analysis carried out in the previous chapter with Altair OptiStruct, the initial geometry of the component leaves ample room for improvement. The SF results (4,3 for the solid CAD model, 3,8 for the characterised-infill configuration) confirm that the part is heavily oversized with respect to its actual operating requirements. Physically, not all the mass of a body contributes equally to carrying the applied loads: there are entire regions in which the material is only lightly stressed and represents nothing more than “dead weight”. The objective therefore becomes to remove this superfluous mass, redistributing it only where strictly necessary to provide the required stiffness and strength, bringing the SF closer to more efficient, targeted values.

To meet this need for weight reduction, topology optimisation is adopted: a computational design methodology that completely redefines the connectivity of the space within an initial volume, known as the design domain. Rather than merely adjusting the external dimensions or surface details of an existing design, this approach gives the algorithm broad geometric freedom (subject only to the constraints explicitly imposed), allowing it to determine where to place voids and solid material according to the actual internal force flows. The result is the generation of an entirely new, optimised geometry, shaped directly by the load paths of the component.

Being a computational, mathematical process, topology optimisation does not provide an immediate result but proceeds iteratively. The solver runs a series of successive simulations to evaluate how the stresses propagate through the material at each geometric variation. This approach typically requires many successive trials: at each step the algorithm evaluates the structural response, progressively removes or reduces the density of the least efficient material, and recomputes the new stress distribution. This iterative cycle converges to an optimised final configuration, often with organic shapes that are hard to obtain by conventional manufacturing but well suited to additive manufacturing.

6.2 Design-Space Setup

Once the need to lighten the component has been established, the first operational step within the optimisation environment is the strict separation of the geometric volumes. To ensure that the algorithm can operate freely without compromising the functionality of the part, the geometric domain is divided into two fundamental macro-categories: the **Design Space** (the design/optimisation domain) and the **Non-Design Space** (the constrained domain).

The Design Space is the volumetric region within which the Altair OptiStruct algorithm is allowed to remove or redistribute material. This volume was identified starting from the maximum envelope of the original component. At this stage, for the Claw in particular, any holes or geometric cavities deemed superfluous for the global structural behaviour were closed beforehand (Figure 79); sealing these minor features gives the optimiser a continuous, larger volume, maximising the ideal paths for the force flows.

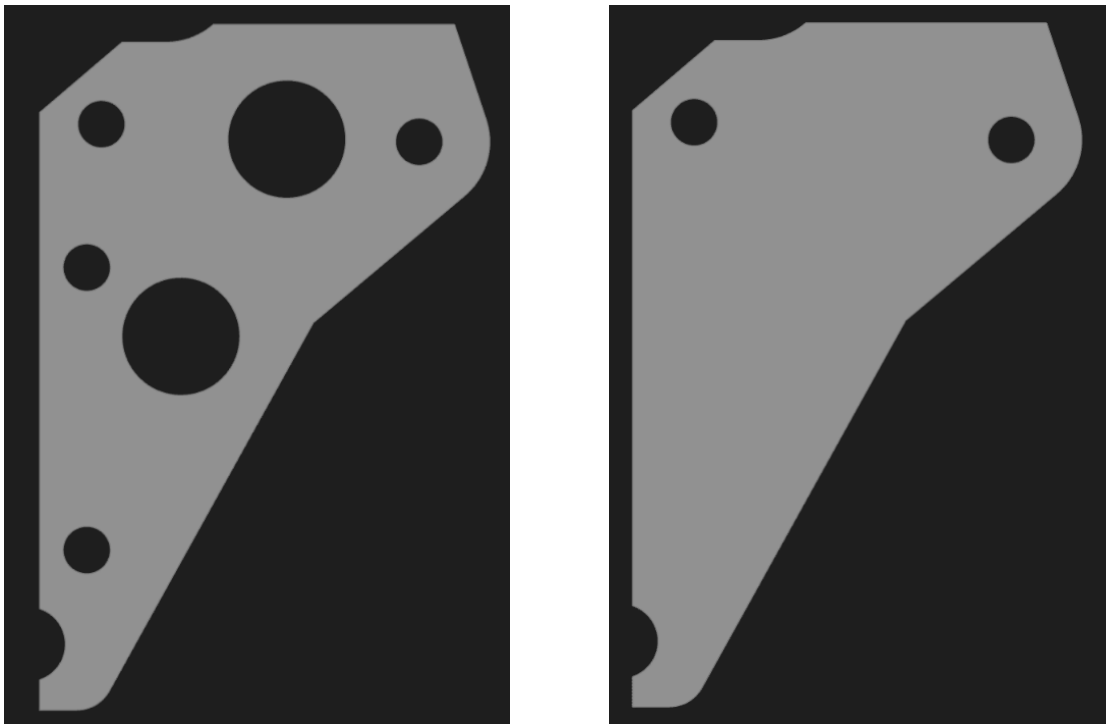

Figure 79: (left) original Claw part; (right) the same part with all non-functional holes filled, where the algorithm might remove material to topologically optimise the part.

Conversely, the Non-Design Space encompasses all the regions of the component that cannot be modified for functional, kinematic or assembly reasons. In this specific case, the fixing holes, the structural coupling systems and the geometric housing (slot) for the servomotor were defined as Non-Design Space (Figure 80). Keeping these regions unaltered is essential to preserve the mating tolerances and to ensure that the final part remains perfectly interchangeable with the original version.

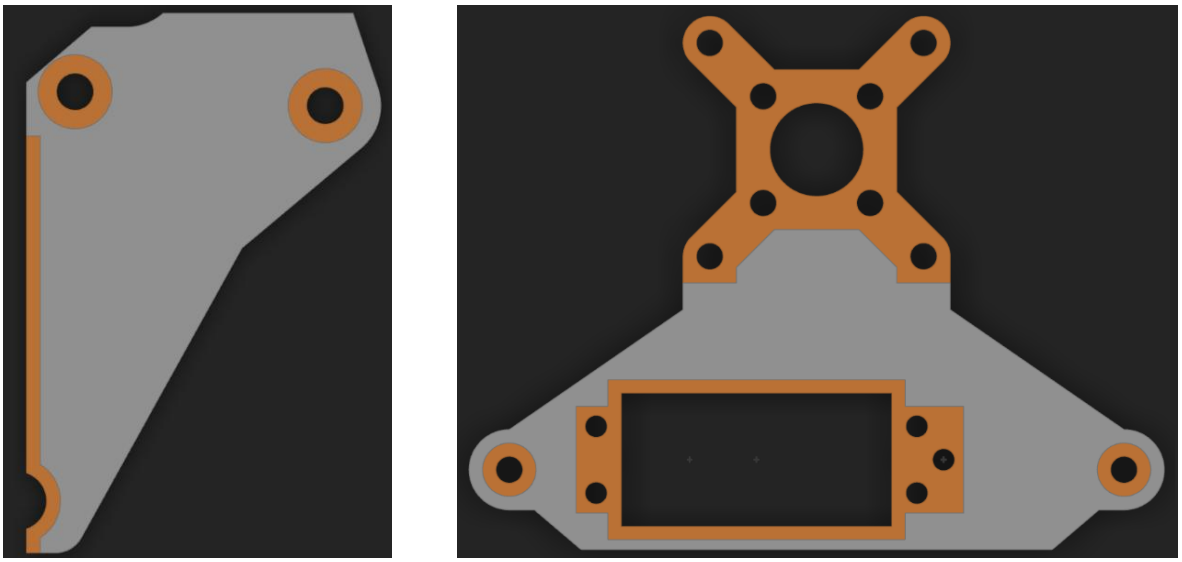

Figure 80: Identification of the Non-Design Space regions (orange) to be excluded from the optimisation, at the functional attachment holes and the gripping surface, whose original shapes are preserved.

To complete the preparation of the mathematical model, the contacts between the Non-Design Space and Design Space regions had to be defined. In Inspire, four different types of connection between bodies can be assigned:

- Bonded (blue): a direct contact that joins two bodies as if they were welded, allowing no relative movement; loads are transmitted fully across the interface.
- Sliding (yellow): allows the two bodies to slide relative to each other along the contact plane, while preventing both interpenetration and separation normal to the plane. This type of contact would have been useful in a multibody system with parts in relative motion, such as the entire gripper system.
- Separating (green): the two bodies are in contact but can separate and slide relative to each other; the interface transmits only compressive loads, and only in the regions actually in contact. This is a non-linear contact, useful for representing bodies that, as they deform, move into or out of mutual contact.
- No contact: all interfaces that do not fall into the three previous cases, where no contact is present between the parts.

In the present case, all the interfaces between Non-Design Space and Design Space were defined as bonded contacts (blue), so as to rigidly connect the functional regions to the optimisable volume. This ensures that the loads applied to the functional zones are correctly transmitted to the optimisable volume, avoiding numerical singularities during the calculation iterations. To keep the visualisation and management of the model clear and standardised, a specific colour convention was adopted: the constrained regions (Non-Design Space) are shown in grey, while the algorithm's action volume (Design Space) is marked in burgundy (Figure 81); the corresponding bonded contacts on the interfaces are highlighted in blue (Figure 82).

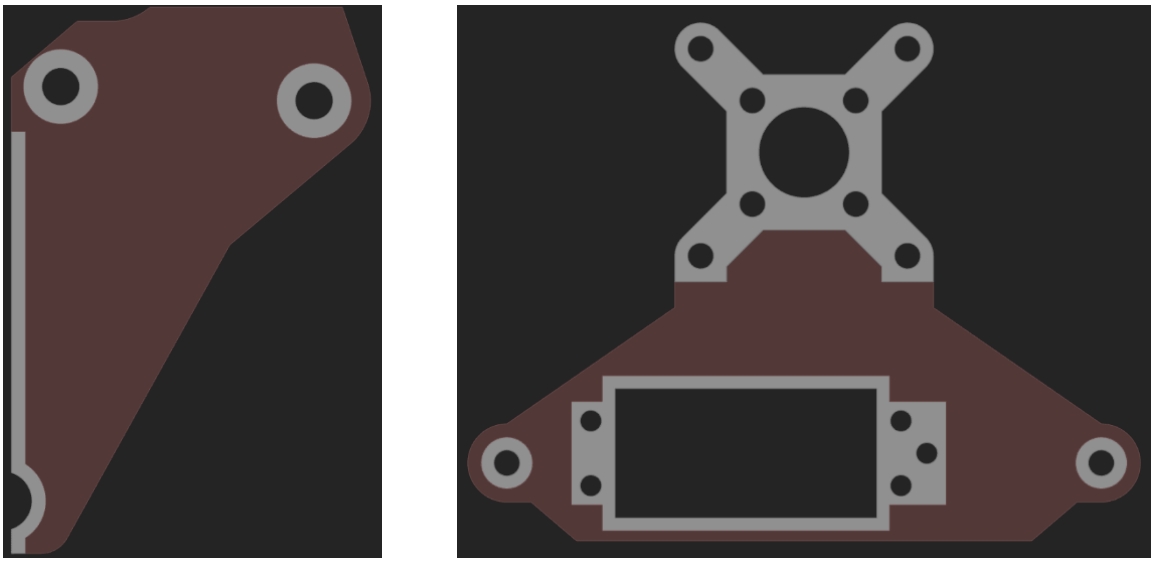

Figure 81: (left) Design Space (burgundy) and Non-Design Space (grey) of the Claw; (right) Design Space (burgundy) and Non-Design Space (grey) of the Frame.

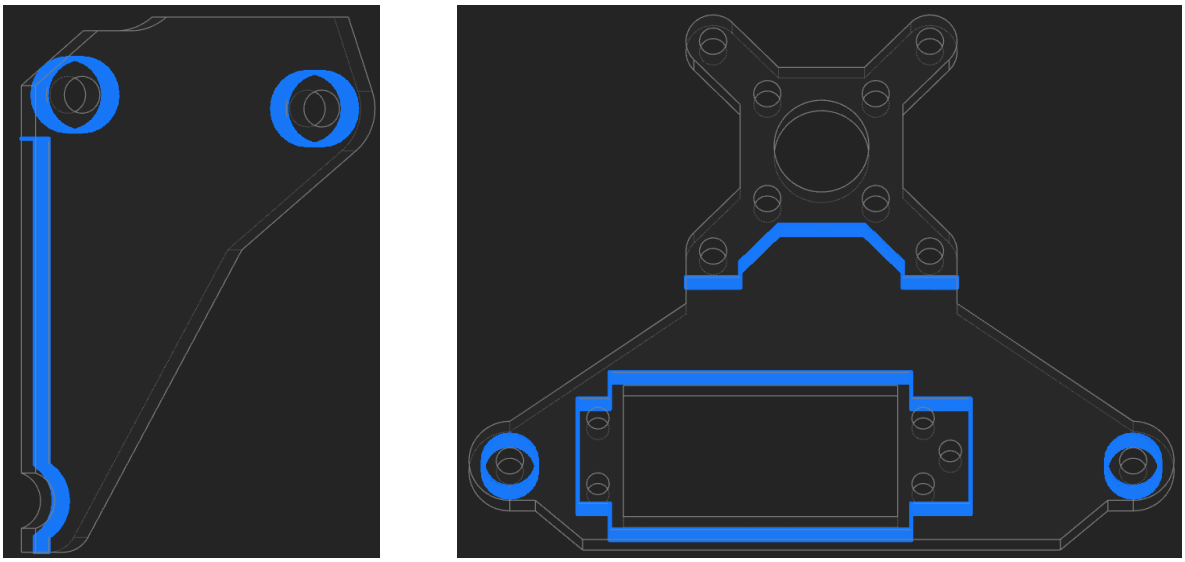

Figure 82: Bonded contacts (blue) defined on the interfaces between Design Space and Non-Design Space; (left) Claw, (right) Frame, represented by hundreds of blue dots.

6.3 Raw Optimisation Results

The topology optimisation was set up for both components with the objective of maximising their structural stiffness (the maximise stiffness option in Inspire), setting as a target a reduction of the design-space volume to about 20 [%] of the original. To make the parts geometrically suitable for 3D printing and to avoid complex undercuts that would have required excessive support material, an extrusion constraint was applied. This forced the algorithm to develop the geometries mainly along the vertical build direction of the part on the plate. The calculation was then launched on the basis of the load cases previously defined for the Claw and the Frame.

Once the calculation was complete, Altair Inspire returned the raw optimisation results, i.e. the theoretical material-density distribution in the form of a coarse mesh. Using the software's density slider, the neighbourhood of the ideal solution was explored to clean the model of isolated elements and to define the main force flows. Once the optimal configuration was found, the actual geometric reconstruction was carried out through a PolyNURBS fit. This operation made it possible to blend and smooth the shapes realistically, transforming the raw software output into a smooth geometric model that serves as the starting point for the redesign of Chapter 7.

The visual steps of this transition, from the initial volume to the final optimised model, are reported below: first the raw optimisation results for both components (Figure 83, Claw; Figure 84, Frame); then, for each component, the PolyNURBS fit and its comparison with the original design space (Figure 85 and Figure 86 for the Claw, Figure 87 and Figure 88 for the Frame).

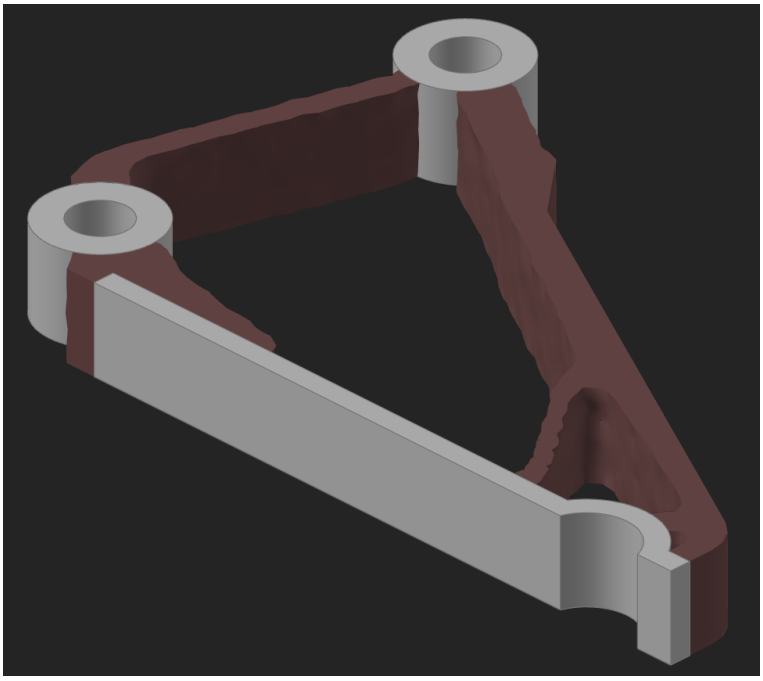

Figure 83: Raw topology-optimisation result of the Claw with the extrusion constraint.

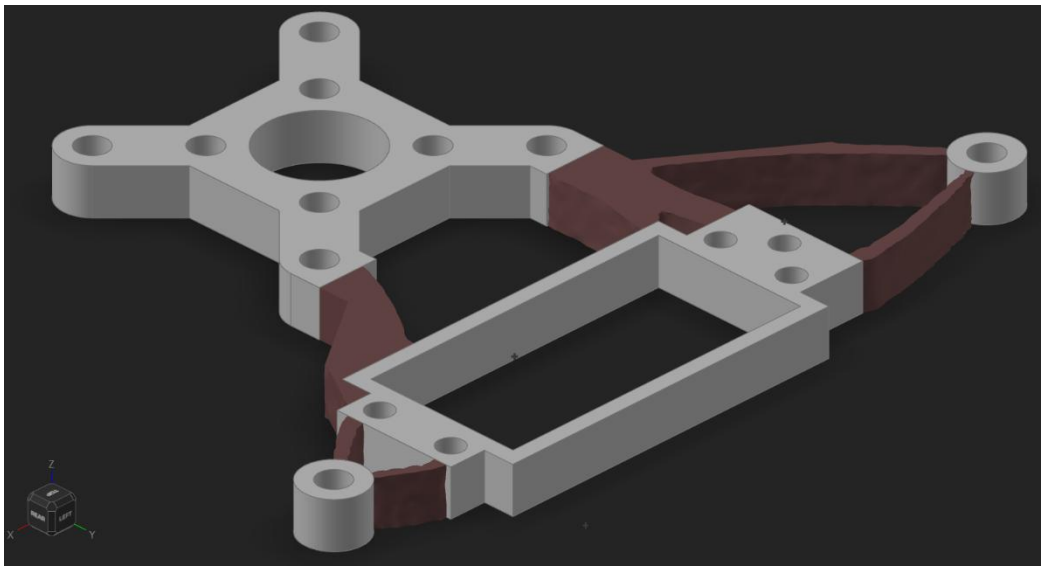

Figure 84: Raw topology-optimisation result of the Frame with the extrusion constraint.

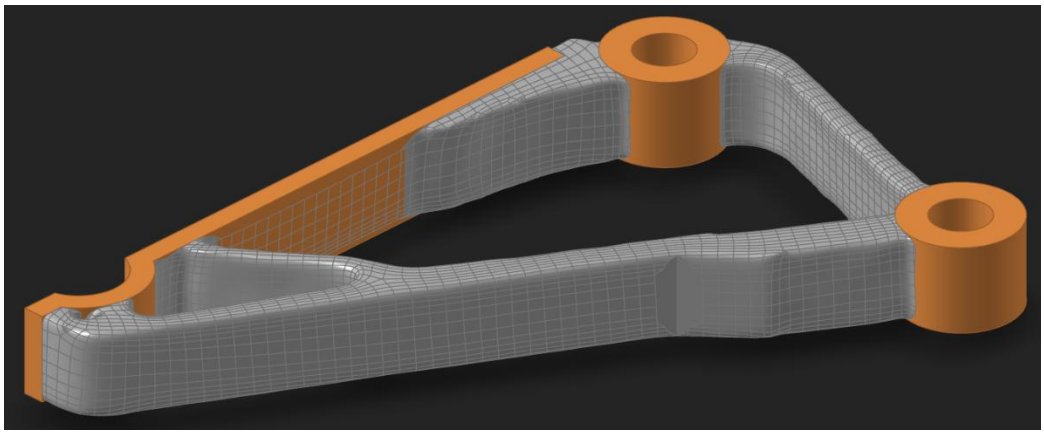

Figure 85: PolyNURBS fit on the optimised surface (grey) with the connected Non-Design Space parts (orange), Claw.



Figure 86: The PolyNURBS part compared with the previously defined design space for the Claw.

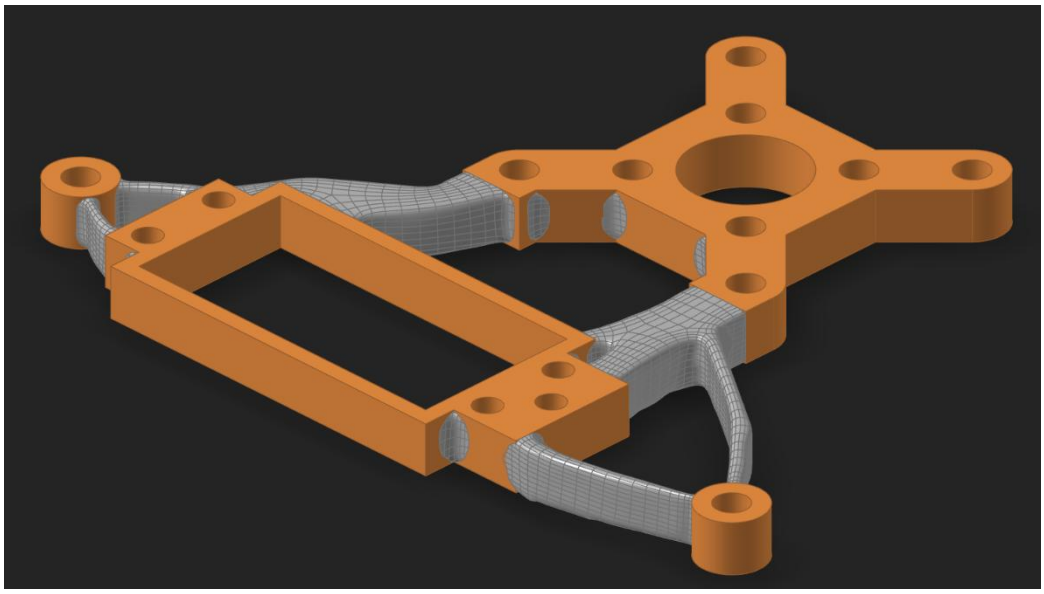

Figure 87: PolyNURBS fit on the optimised surface (grey) with the connected Non-Design Space parts (orange), Frame.

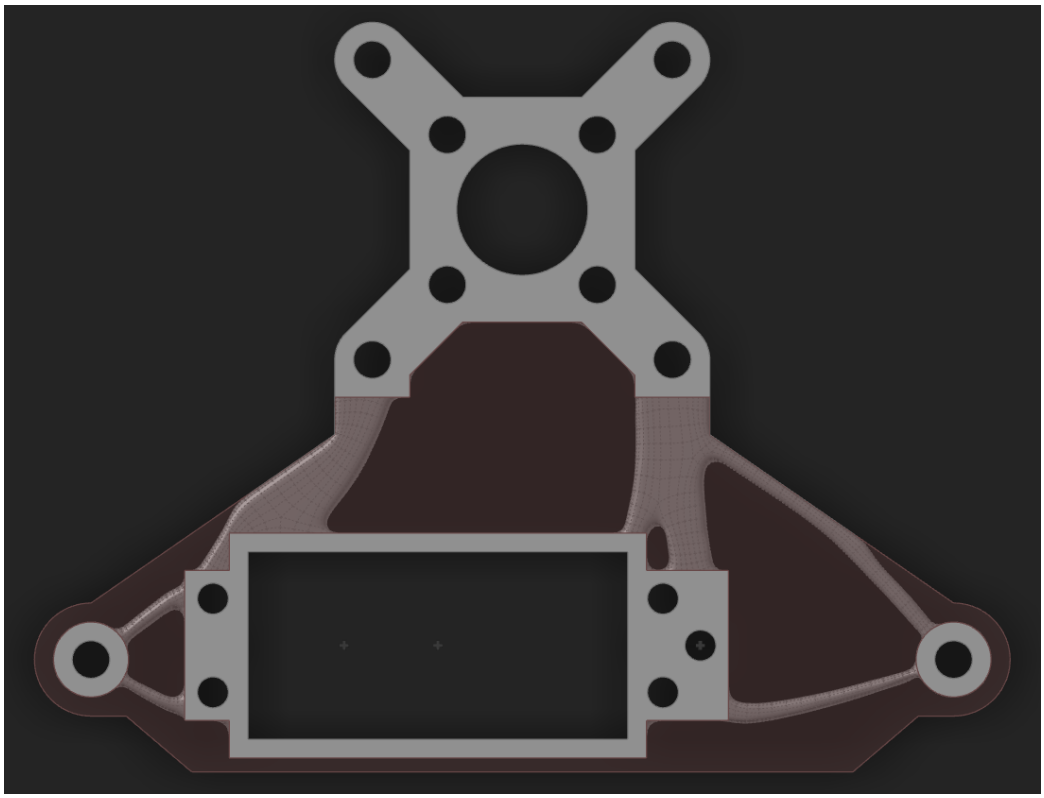

Figure 88: The PolyNURBS part compared with the previously defined design space for the Frame.

7. Redesign and Manufacturing

7.1 Redesign Based on Topology Optimisation

In principle it would be perfectly possible to launch a simulation and verify that the optimised parts meet the requirements; however, this would ignore several factors that the raw optimisation does not account for:

1. The optimisation model joins the Design and Non-Design volumes through bonded contacts, which transmit load as a tie between distinct bodies rather than through continuous material: the stresses read at those interfaces are polluted by this artificial constraint, whereas the real part is a single continuous piece.
2. The automatic PolyNURBS fit tends to create negative (re-entrant) fillets near the contact surfaces with the Non-Design Space regions. If left unaddressed, these notches would act as stress raisers exactly where the load transfers from the Non-Design regions into the part; correcting them by hand is straightforward, and it is one of the goals of the manual remodelling.
3. The model was studied under three significant and critical load cases; real service, however, also includes conditions that were not modelled (assembly preloads, accidental contacts, small impacts), so it is advisable to add further connections and slightly greater thicknesses where additional loads might be expected to act.

Using the PolyNURBS result as a guideline, the Frame and the Claw were recreated as shapes compatible with these considerations, modelled as single monolithic parts, without re-entrant fillets at the contacts and with slightly increased thicknesses and, where useful (for example on the Claw), additional branching. The final geometries are shown in Figure 89 and Figure 90 for the Claw and the Frame respectively.

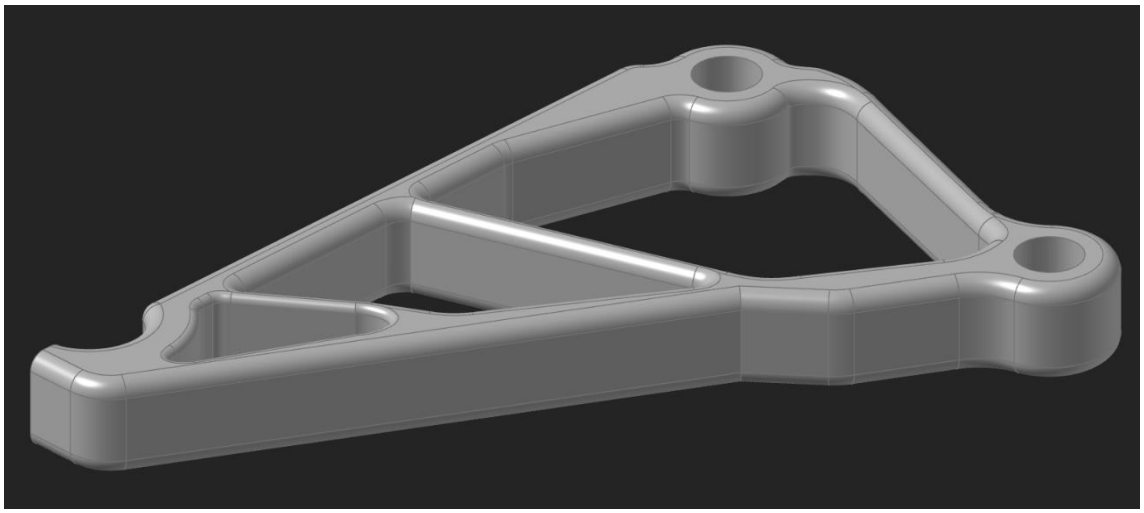

Figure 89: Redesign of the Claw following the PolyNURBS prototype, ensuring full contact between the optimised volume and the remaining components.

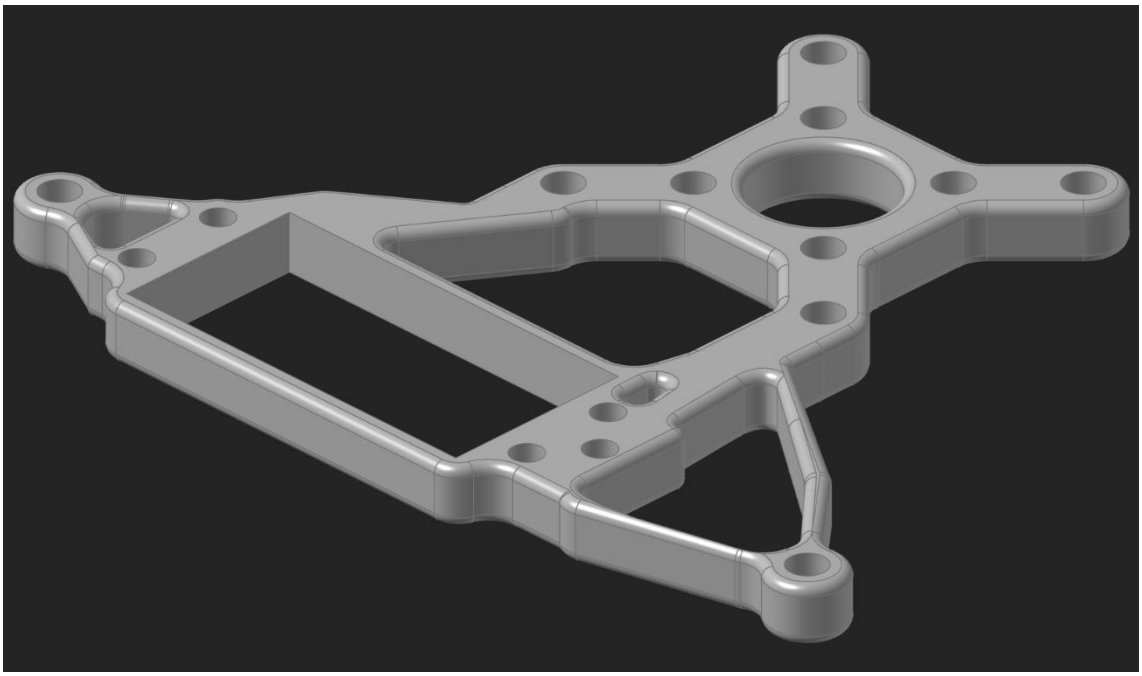

Figure 90: Redesign of the Frame following the PolyNURBS prototype, ensuring full contact between the optimised volume and the remaining components.

7.2 Structural Verification of Redesigned Parts

To verify that the redesigned geometries are structurally adequate, the three load cases of Section 5.1 were re-applied, unchanged, to the final models of Section 7.1. Since these are now single monolithic bodies, the results are free of the bonded-contact artefact discussed above and can be read with the same confidence as those of Chapter 5. The SF envelopes are shown in Figure 91 for the Claw and Figure 92 for the Frame: the governing case is the acceleration along X, with minima of 6,0 for the Claw and 3,4 for the Frame, the latter localised at a sharp, unfilleted corner of the servo slot, which acts as a stress raiser. The displacement envelopes follow in Figure 93 and Figure 94: under the same acceleration case the Claw reaches 1,457 [mm], measured relative to its anchoring holes, and the Frame 1,112 [mm].

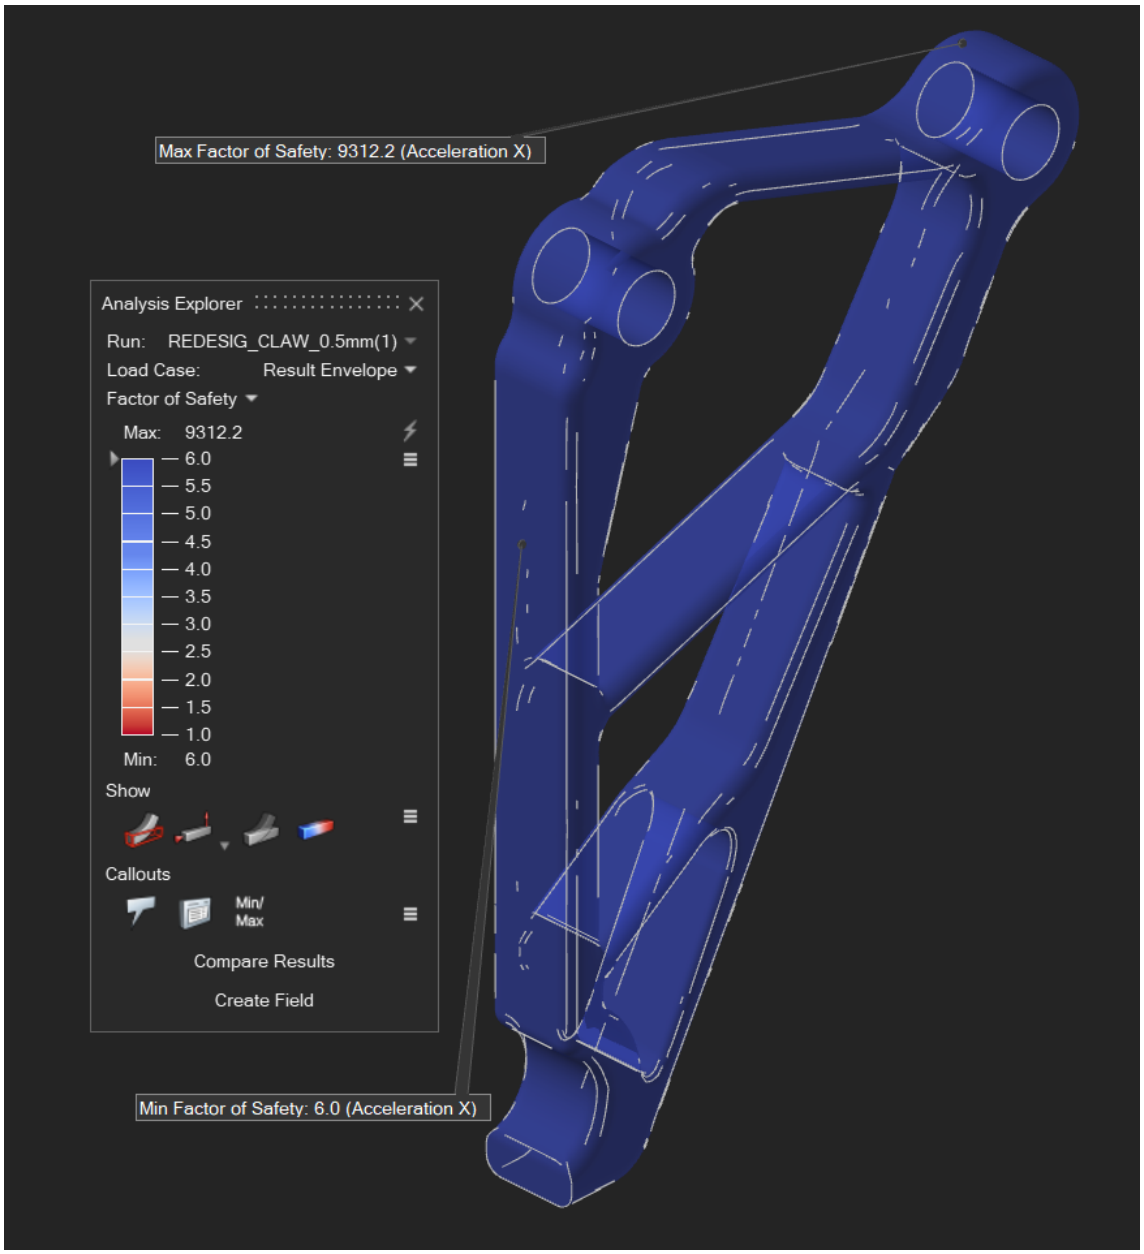

Figure 91: Result envelope of the gravity, X-axis acceleration and servo-stall cases on the redesigned Claw. The minimum SF, 6,0, is reached under linear acceleration along the X axis.

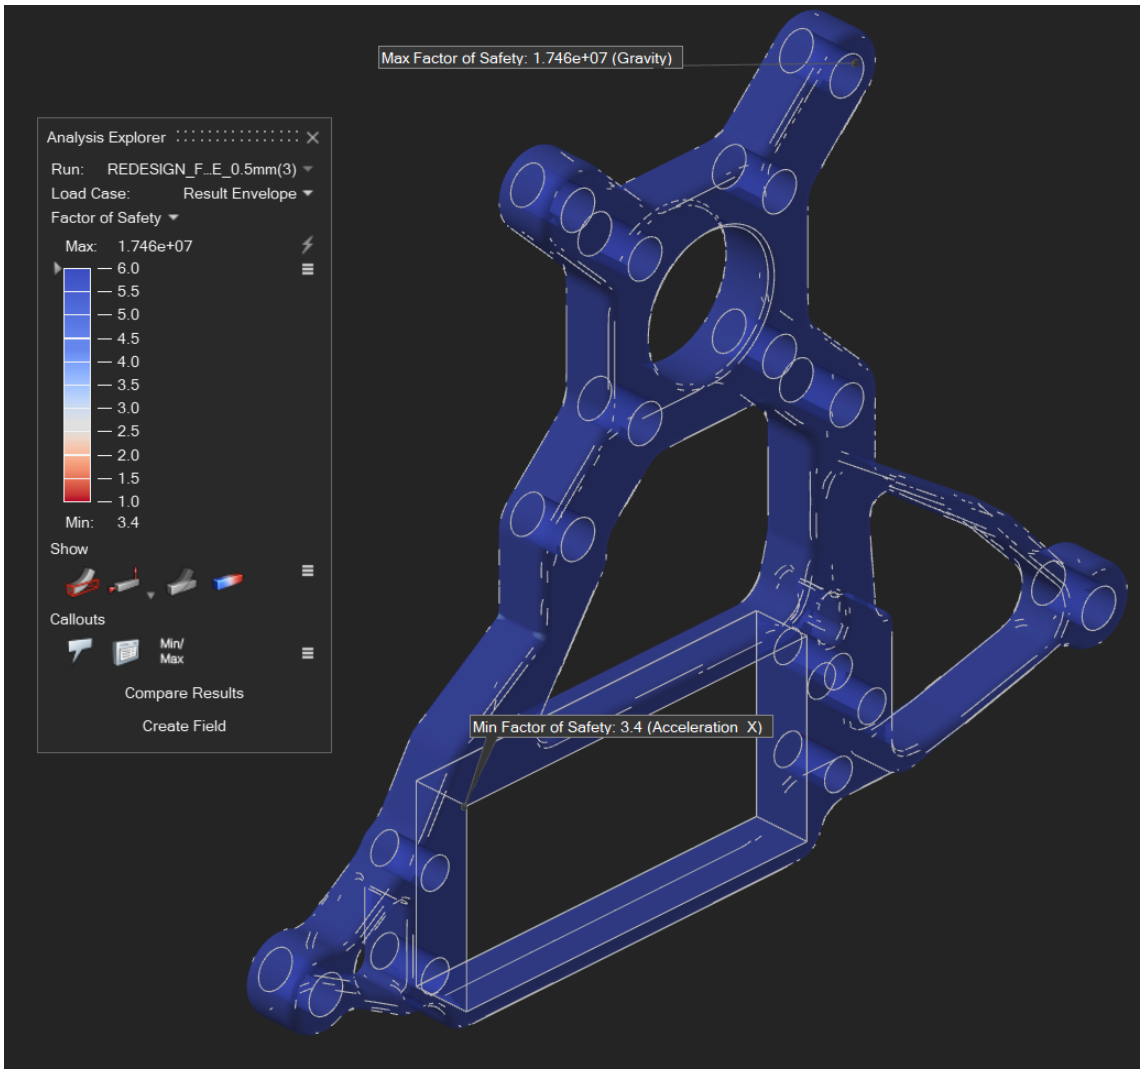

Figure 92: Result envelope of the gravity, X-axis acceleration and servo-stall cases on the redesigned Frame. The minimum SF, 3,4, is reached under linear acceleration along the X axis, at a sharp unfilleted corner of the servo slot.

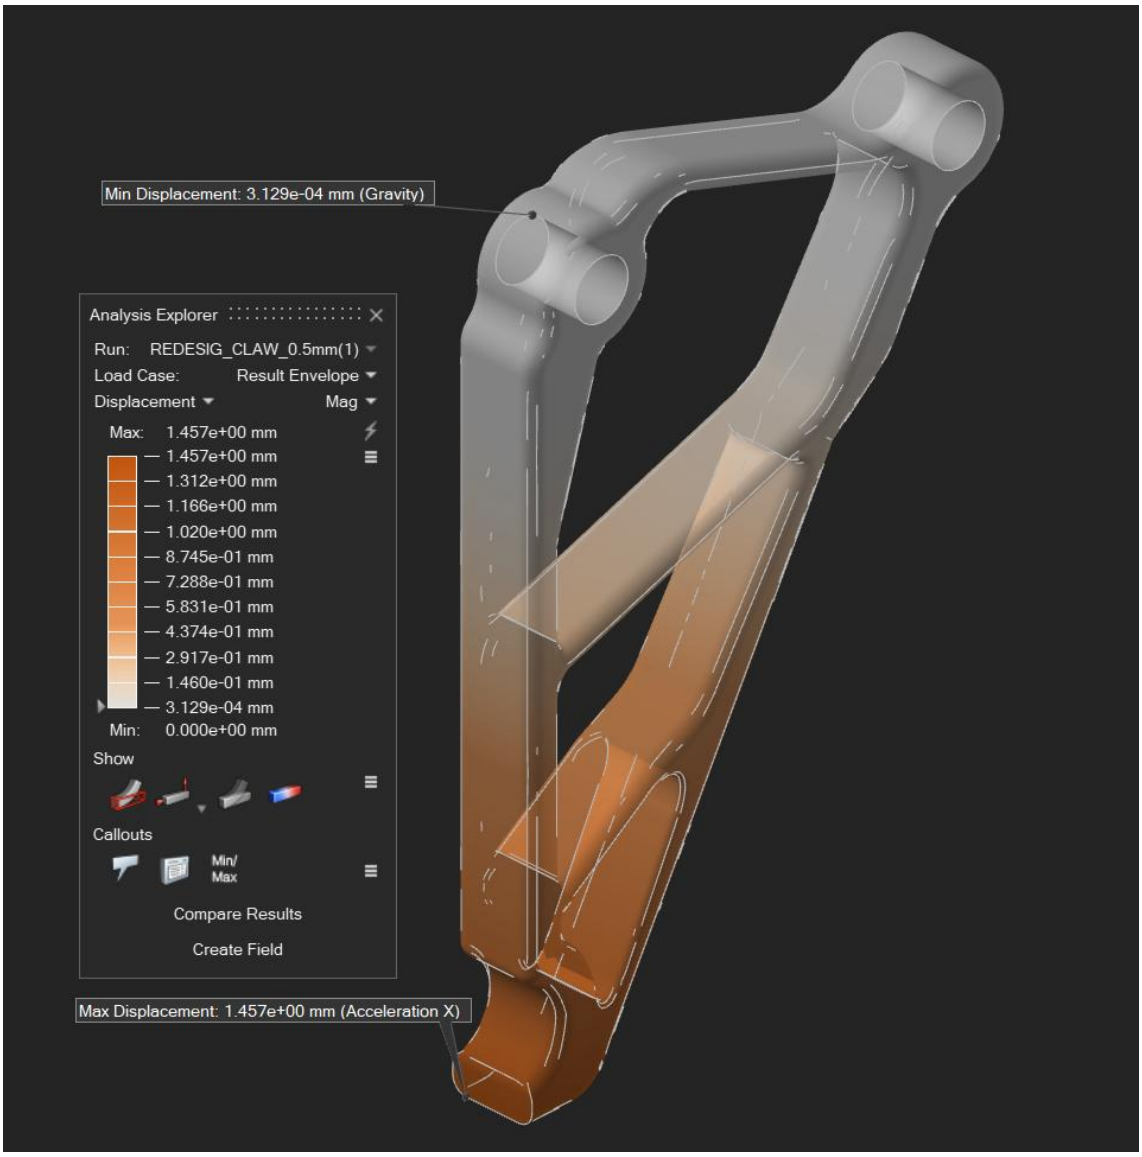

Figure 93: Result envelope of the gravity, X-axis acceleration and servo-stall cases on the redesigned Claw. The maximum displacement, 1,457 [mm], is reached under linear acceleration along the X axis. This displacement is measured relative to the anchoring holes with the link and the gear, not with respect to the global reference system placed at the Frame–robot attachment.

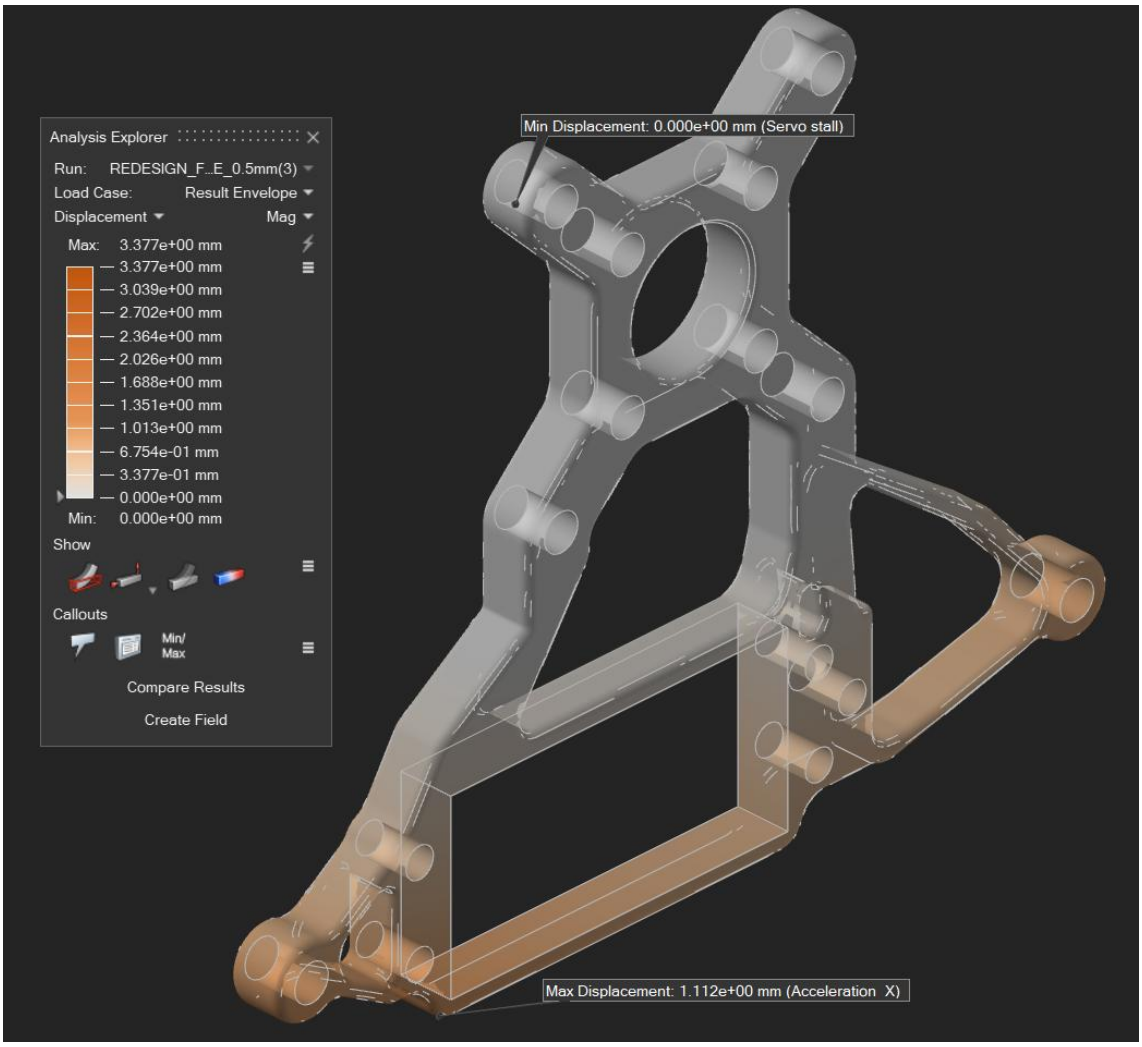

Figure 94: Result envelope of the gravity, X-axis acceleration and servo-stall cases on the redesigned Frame. The maximum displacement, 1,112 [mm], is reached under linear acceleration along the X axis.

Based on the available data, the topologically optimised parts, even once 3D-printed and therefore not fully solid, are expected to remain above the minimum SF target of 2,0 set in Section 2.5. Passing from the solid CAD to the characterised-infill model cost the Frame about 12 [%] of its minimum SF (from 4,3 to 3,8) and left the Claw practically unchanged (10,2 to 10,3); applying a comparable penalty to the redesigned values suggests as-printed minima of roughly 3,0 for the Frame and 6,0 for the Claw, both comfortably above the target. Table 18 consolidates the three configurations.

Component	Metric	CAD	RE (infill)	TO redesign
Frame	Min SF	4,3	3,8	3,4
Frame	Max δ [mm]	0,678	0,821	1,112
Claw	Min SF	10,2	10,3	6,0
Claw	Max δ [mm]	0,197	0,233	1,457

Table 18: Governing minimum SF and maximum displacement of the Frame and Claw across the CAD, reverse-engineered (RE) and topologically optimised (TO) models.

7.3 Mass Reduction Assessment

The redesigned components were then 3D-printed using the same parameters, printer and material as the original specimens. Figure 95 and Figure 96 compare the original and redesigned Frame and Claw as printed in ASA.

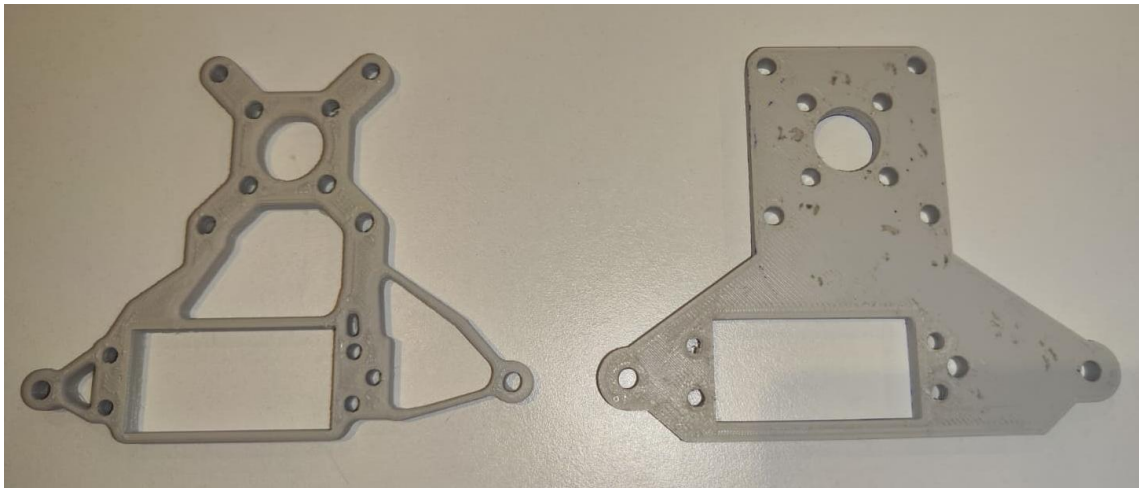

Figure 95: Comparison between the original and redesigned Frame, both printed in ASA.

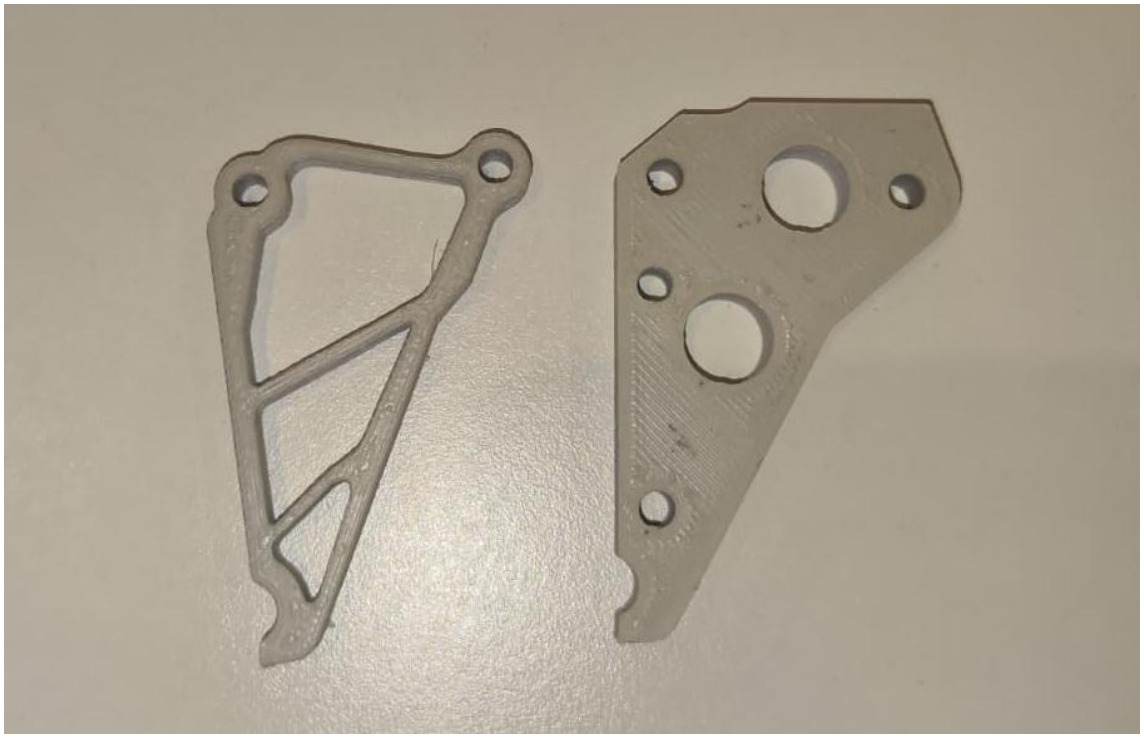

Figure 96: Comparison between the original and redesigned Claw, both printed in ASA.

It remains to quantify the actual mass reduction achieved. Table 19 reports the mass of the Frame and the Claw for the printed original and the optimised printed part, together with the resulting percentage reduction with respect to the printed original. The solid CAD and the reverse-engineered model are omitted from the comparison: the former is not directly comparable with a specimen printed at 20 [%] infill, while the latter reproduces the printed geometry, and therefore its mass, by construction.

Component	Original printed [g]	TO printed [g]	Reduction [%]
Frame	15,6	9,3	≈ 40,4
Claw	5,2	2,4	≈ 53,8

Table 19: Mass of the Frame and Claw for the printed original and the TO printed part, with the percentage mass reduction with respect to the printed original.

In this way, parts were obtained that weigh significantly less than their original counterparts while still comfortably meeting the structural requirements established in Section 2.5.

Extending the comparison to the complete end-effector, Table 20 summarises the mass at each stage of the workflow. Starting from the 133 [g] catalogue mass of the original commercial gripper, in-house FDM fabrication had already reduced the assembled specimen to 101,4 [g]; replacing the original printed Frame and the two Claws with their topology-optimised versions removes a further 11,9 [g], bringing the fully assembled optimised end-effector to approximately 89,5 [g], an overall reduction of about 32,7 [%] with respect to the original commercial unit.

Stage	Mass [g]	Reduction vs original [%]
Original commercial gripper	133	–
FDM as-built specimen	101,4	23,8
Optimised specimen (post-TO)	89,5	32,7

Table 20: End-effector mass at each stage of the workflow, from the original commercial unit to the topology-optimised specimen.

8. Cost and Time Analysis

8.1 Scope and Methodology of Cost Estimation

In line with the standard practice at the UPV, the economic valuation of the project accounts for the engineering hours invested across all of its phases (literature review, calculations, design, simulation, testing and report writing), as well as for the cost of the equipment, software and materials employed. For a 30-ECTS Master's thesis project the nominal student workload is of the order of 750 hours (25 hours per credit); the valuation presented here prices the complete comparative campaign as it was actually carried out, amounting to roughly 700 hours of engineering work, consistent with that envelope. A substantial share of those hours served the comparative aims of the study, rather than the delivery of the optimised parts, and is discussed in the following sections. All figures should be read as indicative estimates.

The cost is built up under two complementary scenarios. Scenario A prices the project as if it had been carried out by an engineering firm: the labour is obtained by multiplying the hours dedicated to each phase (Table 21) by a realistic net hourly rate for a junior engineer, and to this the cost of the resources actually used is added at market conditions: 3D printing (machine time and material), the three scanning technologies (X-ray computed tomography and structured-light scanning with the EinScan and the GOM ATOS) and the allocated share of the commercial software licences. The sum of labour and resources gives the direct cost of the project.

On this direct cost a percentage for contingencies is applied, to cover unforeseen events, followed by a percentage of industrial profit. Finally, as required in Spain, Value Added Tax (VAT/IVA) at 21 [%] is added to obtain the final price of the project. Scenario B, by contrast, reports what the same campaign actually cost within the academic environment: the student's labour is part of the degree, the software runs under educational licences and the scanning equipment is accessed at internal university rates. Being a cost account rather than a commercial quotation, contingencies, profit and VAT do not apply. Both scenarios are summarised side by side in Table 22.

8.2 Time Breakdown by Phase

Table 21 breaks down the engineering time by phase. The values include the printing and assembly of the original specimen and the effort of each subsequent stage of the workflow, up to the writing of the present report. The table also makes the comparative overhead of the study visible: scanning the same parts with three different technologies, inspecting the results twice, reconstructing an infill that a manufacturer would already own as native CAD, and repeating the FEM verification on two models are all activities that a streamlined production workflow would perform once, or not at all. Roughly half of the total effort belongs to this comparative layer rather than to the direct path from the original part to its optimised reprint.

Phase	Time [hours]
Literature review and state of the art	45
Component selection and characterisation	15
Printing and assembly of the original part	15
Structured-light scanning (EinScan + ATOS)	25
X-ray computed tomography	20
Mesh comparison and inspection (GOM Inspect)	100
CAD reconstruction in SolidWorks	80
FEM structural analysis	80
Topology optimisation	50
Redesign and structural verification	70
Reprinting of the optimised parts	15
Report writing	160
Meetings and supervision	25
Total	**700**

Table 21: Engineering time by phase.

8.3 Cost Summary Table

Table 22 gathers the cost items of both scenarios. In Scenario A the engineering labour is obtained from the 700 h of Table 21 at an assumed net rate of 15 €/h; the CT campaign is priced at typical service rates for a job of this size (400–700 €), and the software cost is the allocated share of the commercial subscriptions of the CAD and optimisation packages over the duration of the project. Contingencies of 10 [%] and an industrial profit of 6 [%] are then applied, and finally VAT at 21 [%] is added to reach the total price. In Scenario B only the internal university rate of the CT system (15 €/h) generates actual expenditure (about 200 € for the machine time of the campaign), and even that charge is borne by the hosting department rather than by the student, while the printing filament was supplied by the laboratory.

Item	Scenario A – Company [€]	Scenario B – Academic [€]
Engineering labour (700 h × 15 €/h)	10 500	0
3D printing (machine time + ASA material)	30	0
X-ray computed tomography	600	200
Structured-light scanning (EinScan + ATOS)	200	0
Software licences (allocated share)	2 500	0
Direct cost	**13 830**	**200**
Contingencies (10 [%])	1 383	–
Industrial profit (6 [%])	913	–
Subtotal before tax	**16 126**	**–**
VAT / IVA (21 [%])	3 386	–
Total project cost	**19 512**	**200**

Table 22: Cost summary of the project under the company (A) and academic (B) scenarios.

Read as a price, the Scenario A total of Table 22 describes a one-off engineering study, not a recurring production cost. Most of it is labour, and most of that labour served comparison and documentation. In a production setting the workflow would be standardised: one scanning technology, chosen once for the class of parts at hand, an established optimisation-and-redesign pipeline, and no duplicate inspection. Its engineering cost would then be amortised over the series, while the benefits repeat with every unit produced, namely less material per part for the manufacturer and lighter end-effector components, which allow the robot to reach higher accelerations or the same cycle times with smaller actuators. Under those conditions the up-front cost of introducing the process is recovered over time, and the figure above should be read as the exploratory step that makes that standardisation possible. The gap between the two scenarios is itself informative: the same campaign that a firm would quote at roughly 19 500 € was executed academically at a residual internal cost of about 200 €, borne by the hosting department, the difference being the unpaid student labour and the educational access to equipment, materials and software.

9. Conclusions

9.1 Encountered Issues and Limitations

The comparison of the three acquisition techniques highlighted the intrinsic limits of each. Structured-light scanning with the GOM ATOS and X-ray computed tomography (CT) proved essentially equivalent on the external surfaces, their global deviations differing by no more than the resolution of the instruments themselves, with the handheld EinScan the noisiest of the three; only the CT, however, could resolve the internal cavities of the gripper, and it therefore had to be used as the master for the reverse engineering.

The CT campaign, however, proved to be the most demanding stage of the whole workflow. From the outset, component selection was constrained by the physical volume of the CT chamber and by the X-ray penetration achievable through the material, which limited the parts that could realistically be digitised. The scanner initially available at the Universitat Politècnica de València could not be used for the reasons explained earlier, so the scans were carried out on the Politecnico di Torino system instead; arranging this between the two institutions took time and set the schedule back. A further setback came at the start of July, when the workstation developed a RAM fault that corrupted part of the data already written to disk and forced a full reset of the machine, which slowed the work down again.

Even with a dedicated campaign, the CT reconstruction did not go below an accuracy of about ±0,3 [mm]. Part of this deviation does not stem from the technique itself but from the parameters chosen when exporting the STL meshes (triangle size, chord error and total number of triangles), and from the deviation of the 3D-printed specimen with respect to the original CAD.

A second family of difficulties concerned data interoperability and mesh quality. Moving the geometry across the different proprietary environments repeatedly caused format conflicts, and the raw STL files carried extensive topological errors (open loops, non-manifold edges and disconnected fragments) that required intensive manual healing before the meshes could be used. The subsequent generation of the STEP solids introduced errors of its own, most notably inverted face normals that corrupted the solid topology and had to be corrected by hand. The custom Python routines written to automate parts of this processing were themselves a source of debugging effort, with algorithmic exceptions that had to be traced and resolved case by case.

A third group of limitations concerns the mechanical modelling. The FEM analyses treated the material as an ideal, isotropic solid, whereas the FDM parts are intrinsically anisotropic and contain internal defects (small voids and non-uniform infill) that are not captured by the idealised model and that make the printed part weaker, particularly in the direction normal to the build plate. Because the safety factors obtained remained comfortably above the design target, these effects do not overturn the conclusions, but they must be kept in mind when interpreting the absolute values.

Finally, the optimisation environment introduced both practical and methodological limitations. Altair Inspire suffered from chronic instability, with sudden crashes during the denser FEM runs and the topology-optimisation cycles, and from GPU rendering artefacts that made the on-screen interpretation of stress concentrations unreliable; a specific solver bug also prevented validation analyses from being run directly on the automatically generated PolyNURBS geometries. On the methodological side, the bonded contacts between design and non-design space do not perfectly reproduce a monolithic body, and the automatic PolyNURBS reconstruction tended to introduce re-entrant fillets. For these reasons, and compounded by the validation constraints just mentioned, the optimised geometry was not adopted as generated but

redesigned by hand, with slightly increased thicknesses and additional connections that trade a small part of the theoretical mass saving for a more robust and genuinely manufacturable component.

9.2 Summary of Results

This thesis has demonstrated a complete physical-to-digital-to-physical workflow for the lightweight redesign of a robotic end-effector, showing that the combination of 3D scanning, reverse engineering, finite-element analysis, topology optimisation and additive manufacturing makes it possible to reduce the mass of the moving components appreciably while preserving their structural integrity. The objective set out in Chapter 1, a mass reduction of the order of 10–30 [%] on the structural parts, was met on the redesigned Frame and Claw, whose figures are collected in Table 19. At the level of the complete end-effector, the combination of in-house fabrication and topology optimisation lowered the assembled mass from the 133 [g] catalogue value of the original commercial unit to approximately 89,5 [g], an overall reduction of about 32,7 [%] (Table 20).

The metrological comparison is a result in its own right. The same physical parts were digitised with three different technologies, giving a like-for-like measure of their accuracy and of their suitability for external and internal geometry: the global standard deviation against the nominal model was 0,147 [mm] for the handheld EinScan, 0,117 [mm] for the GOM ATOS and 0,096 [mm] for the CT. Only the CT, however, could see inside the part, and it was this capability, rather than external-surface accuracy, on which it matched the ATOS, that made it the master dataset for the reconstruction.

Several intermediate results also stand on their own. The CT data proved detailed enough to rebuild the components in a parametric CAD environment, printed infill included, and to quantify how that infill changes the mechanical response with respect to the solid model: the minimum safety factor of the Frame moved from 4,3 on the nominal CAD to 3,8 on the infill-based reconstruction. The topology-optimisation study was not left as a raw density map but carried through to a manufacturable, structurally verified redesign, with minimum safety factors of 3,4 for the Frame and 6,0 for the Claw, in both cases above the target of 2,0 set in Section 2.5.

The project was also framed in economic terms. The effort was estimated at 700 hours, consistent with the 30 ECTS credits of the Master's thesis, and costed under two scenarios: a full commercial estimate (including contingency, profit and Spanish VAT) of about 19 512 €, and an academic scenario in which university licences and departmental equipment reduce the out-of-pocket cost to roughly 200 €. The gap between the two figures is itself informative, since it isolates the value of the institutional resources that made the work possible.

Not every step proved strictly necessary to reach the final result. The three-way scanning comparison, in particular, was partly a methodological exercise: a single well-chosen technique would have sufficed to drive the redesign, and the handheld and fixed-head structured-light acquisitions, though instructive, did not feed the final CAD. Recognising which activities were essential and which were exploratory is itself part of the engineering judgement the work was meant to develop.

More broadly, the project confirmed the cross-disciplinary nature of mechatronic engineering, combining metrology, structural mechanics, computational optimisation and additive manufacturing in a single pipeline; algorithmic tools developed for the outer-shell extraction and the colour-based infill filtering helped to reduce otherwise intractable problems to a form

that standard engineering software could handle. The resulting mass reduction is directly relevant to industrial practice, where a lighter end-effector means lower energy consumption or higher attainable working speeds. Natural next steps are the experimental validation of the printed optimised parts and the refinement of the FEM with an anisotropic material model that reflects the FDM process. The fully reassembled manipulator, fitted with the optimised parts, is shown in Figure 97.

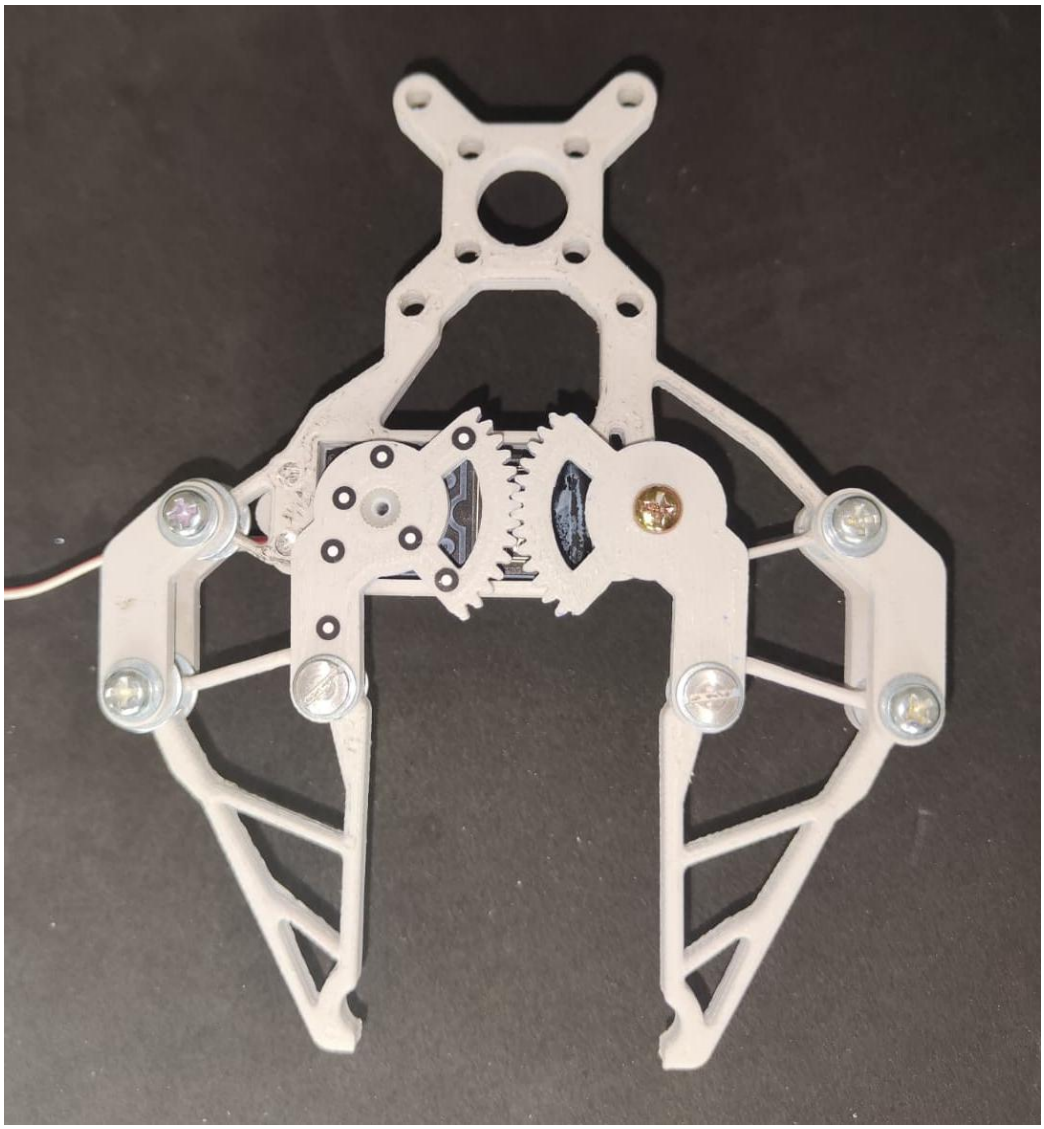

Figure 97: The manipulator fully assembled with the redesigned and optimised parts.

Bibliography

[1] ServoCity. (n.d.). *Assembly instructions for Servo-Driven Parallel Gripper Kit (Servo Included) (SKU: 3219-0001-0002).* ServoCity.com. https://cdn.robotshop.com/media/S/Sct/RB-Sct-2039/pdf/servocity_servo_driven_parallel_gripper_kit_servo_assembly_instructions.pdf

[2] Sha, L., Lin, A., Zhao, X., & Kuang, S. (2020). A topology optimization method of robot lightweight design based on the finite element model of assembly and its applications. *Advances in Mechanical Engineering*, *12*(7). https://doi.org/10.1177/0036850420936482

[3] Jia, J., & Sun, X. (2023). Structural optimization design of a six-degrees-of-freedom serial robot with integrated topology and dimensional parameters. *Sensors,* 23(16), 7183. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC10459507/

[4] Yahya, S., Moghavvemi, M., & Almurib, H. A. F. (2012). Joint torque reduction of a three dimensional redundant planar manipulator. *Sensors*, *12*(6), 6869–6892. https://doi.org/10.3390/s120606869

[5] Carabin, G., Wehrle, E., & Vidoni, R. (2017). A review on energy-saving optimization methods for robotic and automatic systems. *Robotics*, *6*(4), 39. https://doi.org/10.3390/robotics6040039

[6] Clavel, R. (1988). DELTA, a fast robot with parallel geometry. *Proceedings of the 18th International Symposium on Industrial Robots*, 91–100. Springer-Verlag.

[7] Clavel, R. (1990). Device *for the movement and positioning of an element in space* (US Patent No. 4,976,582). United States Patent and Trademark Office. https://patents.google.com/patent/US4976582A/en

[8] Nabat, V., Rodriguez, M. de la O., Company, O., Krut, S., & Pierrot, F. (2005). Par4: Very high speed parallel robot for pick-and-place. *Proceedings of the 2005 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 553–558. https://doi.org/10.1109/IROS.2005.1545143

[9] ETT Co., Ltd. (n.d.). ET-S3003 *standard servo motor datasheet*. ES Electronics. https://www.es.co.th/schemetic/pdf/et-servo-s3003.pdf

[10] 3DJAKE. (2024). 3DJAKE ASA: Technical Data Sheet (Version 2.0). Niceshops GmbH. https://3d.nicecdn.com/upload/file/Technical_Data_Sheet_ASA_V2.pdf

[11] 3DJAKE. (2024). 3DJAKE ABS: Technical Data Sheet (Version 1.0). Niceshops GmbH. https://3d.nice-cdn.com/upload/file/3DJake_ABS.pdf

[12] SHINING 3D. (2020). EinScan Pro 2X 2020: Multi-functional handheld 3D scanner – Technical data sheet (Version 0.7). https://www.einscan.com

[13] GOM mbH. (2013). ATOS Core – Optical 3D measuring system: Product brochure (Rev. B). https://www.gom.com

[14] GE Sensing & Inspection Technologies GmbH. (2009). v|tome|x s – Versatile and flexible CT system: Product brochure. https://www.phoenix-xray.com

[15] Brooks, R. A., & Di Chiro, G. (1976). Beam hardening in X-ray reconstructive tomography. Physics in Medicine & Biology, 21(3), 390–398. https://doi.org/10.1088/0031-9155/21/3/004

[16] Katz, S., Tal, A., & Basri, R. (2007). Direct visibility of point sets. ACM Transactions on Graphics, 26(3), Article 24. https://doi.org/10.1145/1276377.1276407

[17] Fernández-Vicente, M., Calle, W., Ferrándiz, S., & Conejero, A. (2016). Effect of infill parameters on tensile mechanical behavior in desktop 3D printing. *3D Printing and Additive Manufacturing,* 3(3), 183–192. https://doi.org/10.1089/3dp.2015.0036

[18] Acrylonitrile styrene acrylate (ASA). (n.d.). MakeItFrom.com. https://www.makeitfrom.com/material-properties/Acrylonitrile-Styrene-Acrylate-ASA

[19] Altair Engineering Inc. (n.d.). Altair Inspire: Run analysis (element size and solver options) [Software documentation]. Altair Help. https://help.altair.com/inspire/en_us/topics/inspire/structure/analysis_run_options_r.htm

Appendices

The appendices collect the material that supports the main text without interrupting it. Appendix A gathers the technical drawings of the components, included for repeatability purposes. Appendices B, C and D report the complete inspection results of the three scanning routes compared in Chapter 3: handheld EinScan, fixed-head ATOS and X-ray CT respectively. Appendices E and F list the Python scripts developed for the outer-shell extraction and for the HSV-based infill segmentation. Appendix G derives the servo stall loads applied in the FEM analyses of Chapter 5, Appendix H collects the corresponding safety-factor and displacement maps, and Appendix I documents the hardware and software configuration with which the whole campaign was carried out.

A - Technical drawings of the components

This appendix collects the dimensioned technical drawings of the components, provided so that the geometry described in the main text can be reproduced independently of the CAD files. Each sheet reports the principal dimensions of the corresponding part together with its title block.

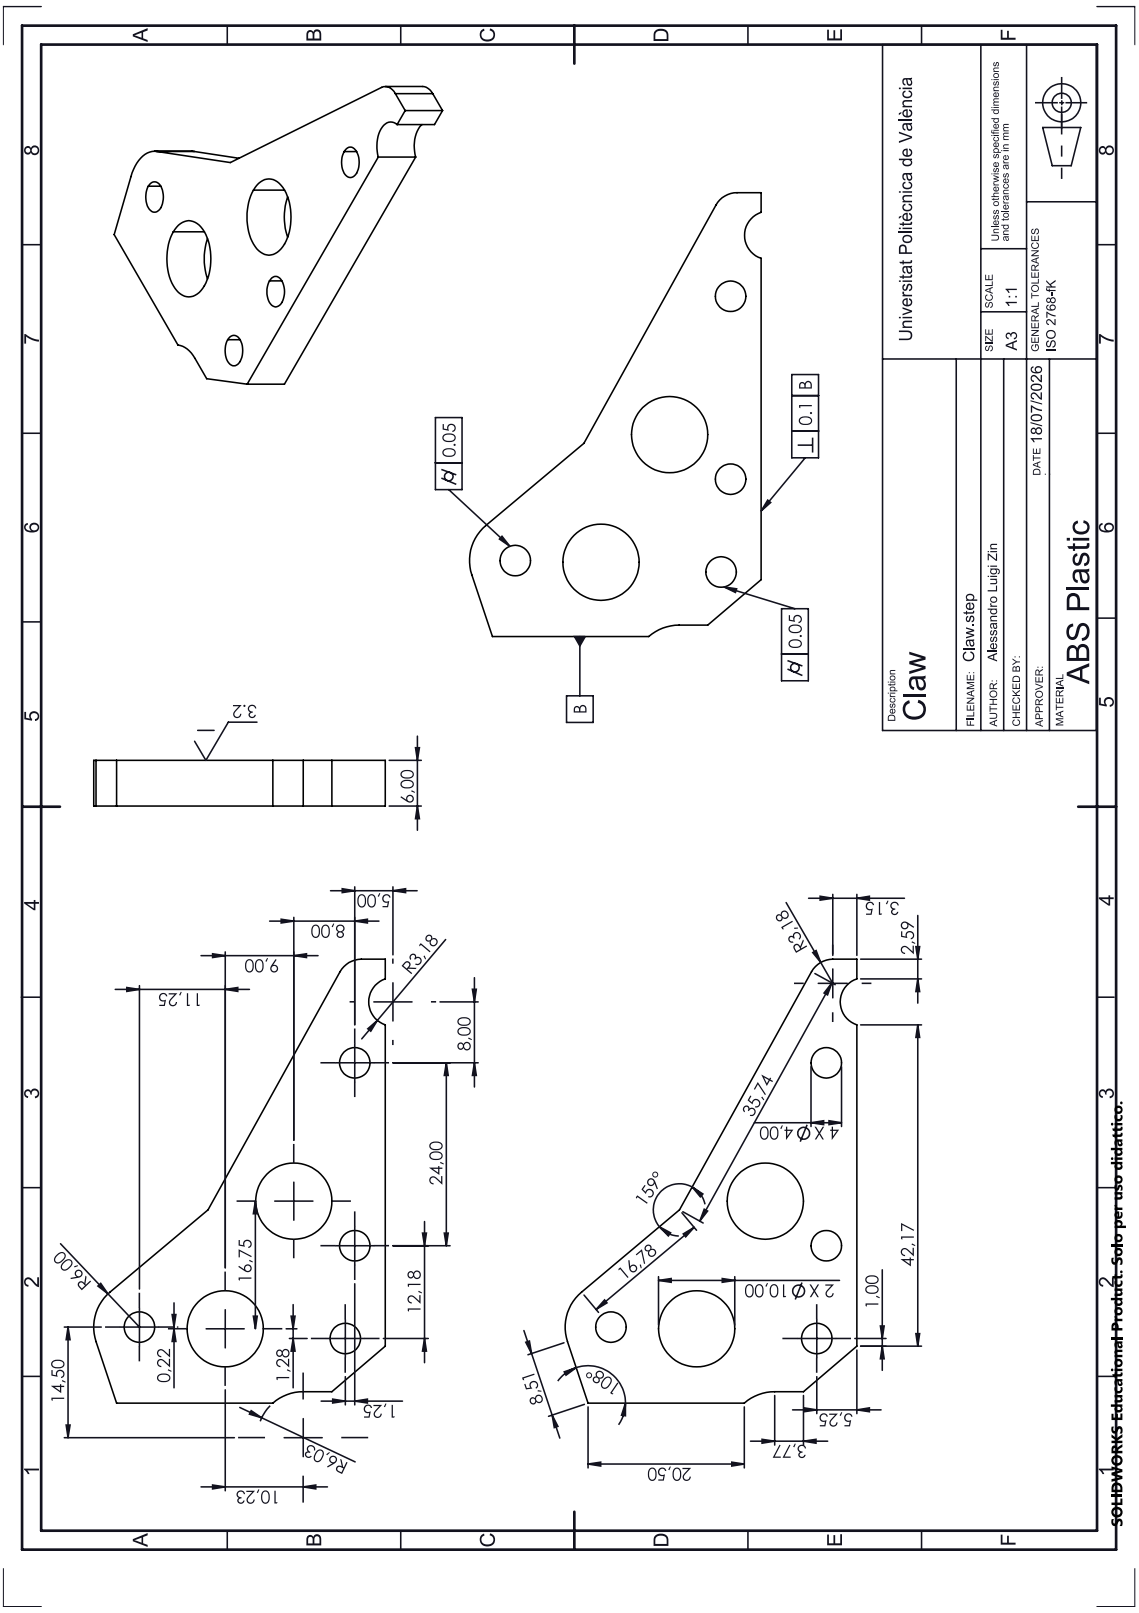

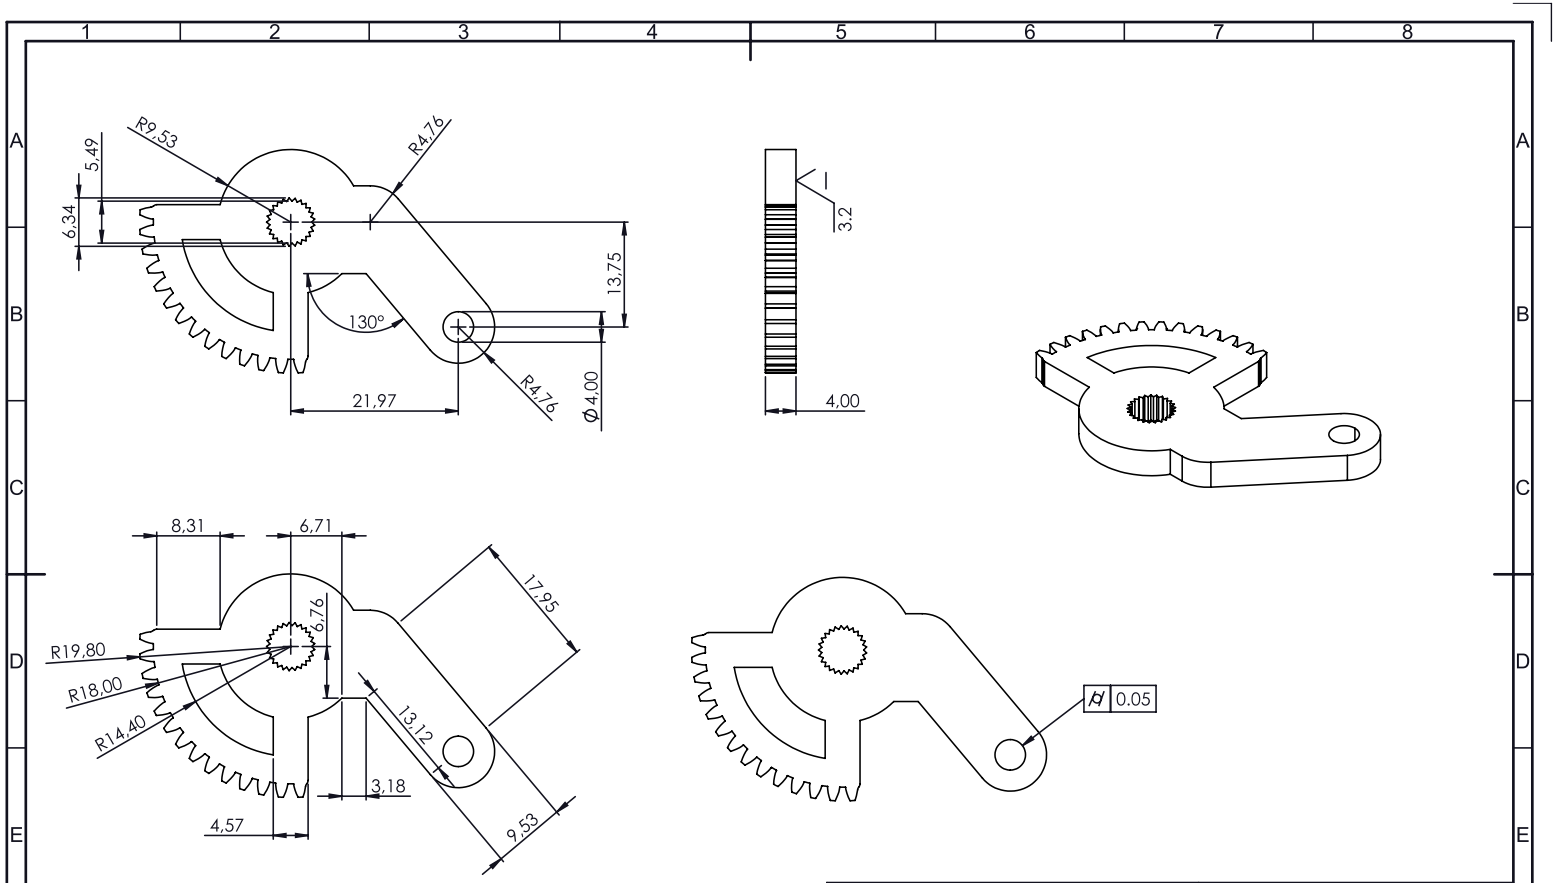

GENERAL NOTES:
EXTERNAL SECTOR CORRESPONDS TO A 54-TOOTH FULL SPUR GEAR WITH A 20° PRESSURE ANGLE AND 19.00 mm

PITCH RADIUS. INTERNAL HOLE FEATURES A STANDARD 25-TOOTH (25T) STRAIGHT SPLINE PROFILE FOR SERVO MOTOR COUPLING.

Description: Drive Gear			Universitat Politècnica de València
FILENAME: Drive_Gear.step			
AUTHOR: Alessandro Luigi Zin		SIZE: A3	SCALE: 1:1
CHECKED BY:			Unless otherwise specified dimensions and tolerances are in mm
APPROVER:	DATE: 18/07/2026	GENERAL TOLERANCES: ISO 2768-fK	
MATERIAL: Acetal Plastic			

SOLIDWORKS Educational Product. Solo per uso didattico.

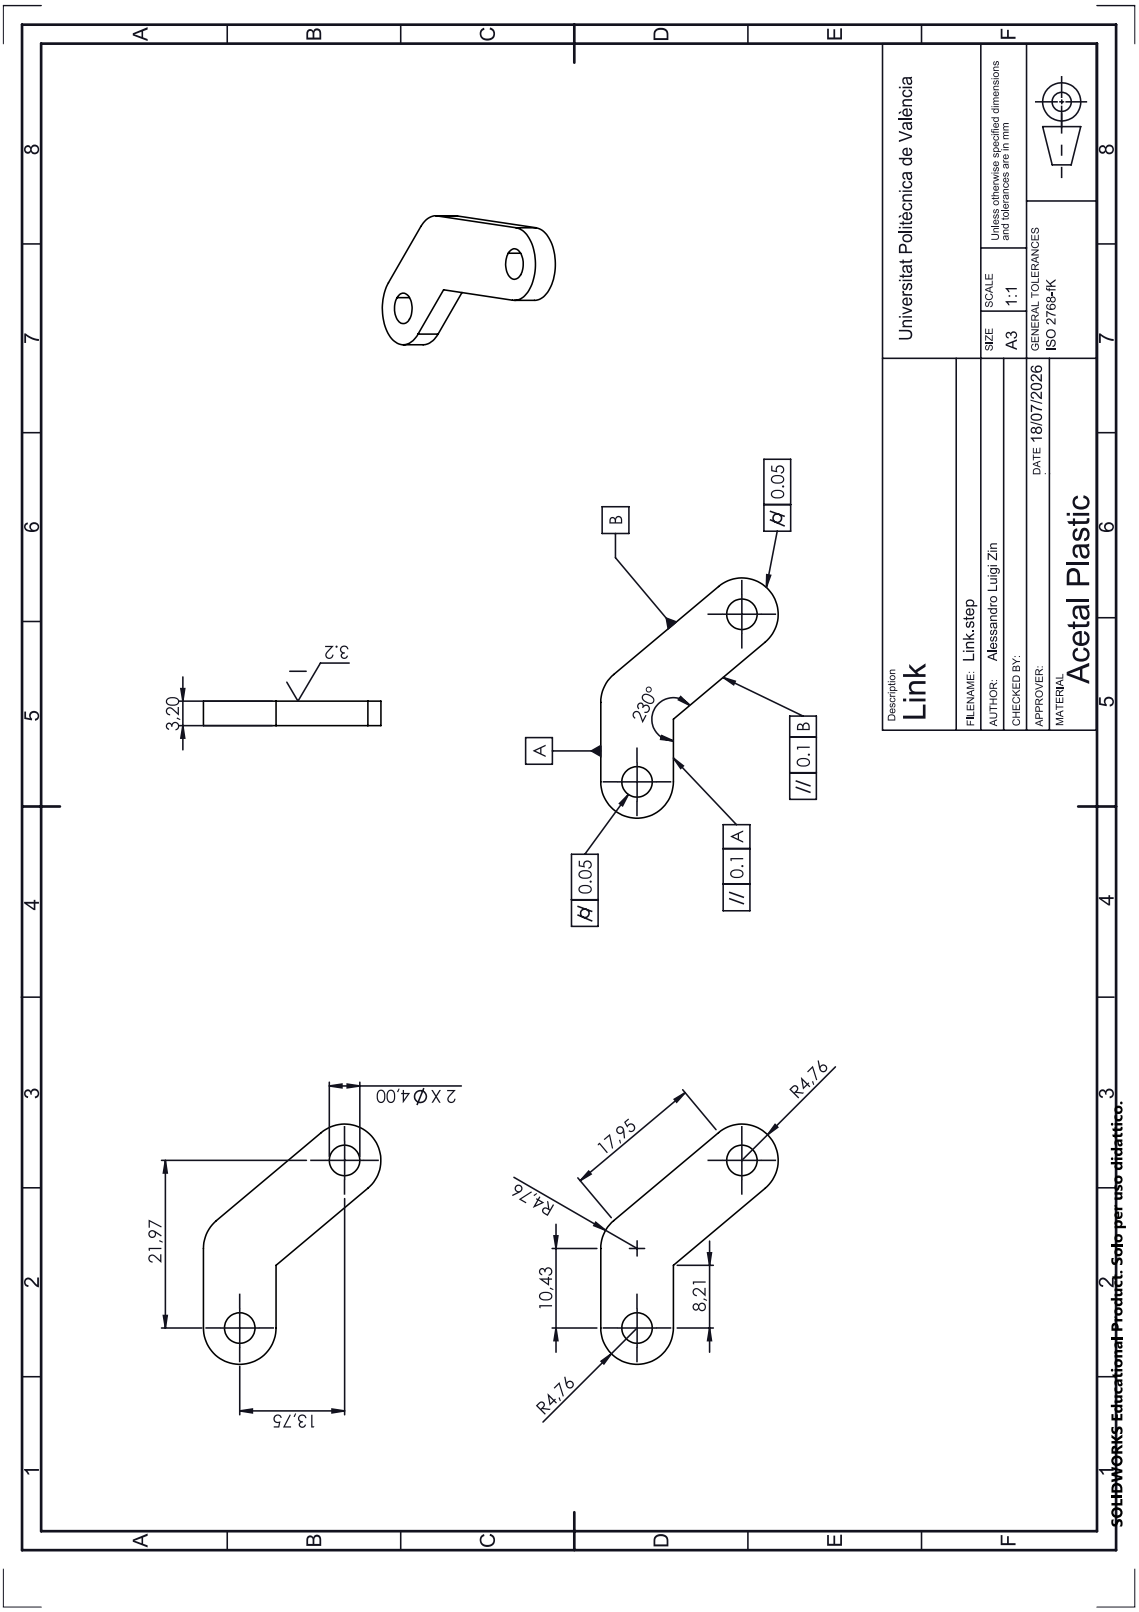
Universitat Politècnica de València
Description
Link
FILENAME: Link.step
AUTHOR: Alessandro Luigi Zin
CHECKED BY:
APPROVER:
MATERIAL
Acetal Plastic
DATE 18/07/2026
SIZE
A3
SCALE
1:1
Unless otherwise specified dimensions and tolerances are in mm
GENERAL TOLERANCES
ISO 2768-fK
3,20
3,2
230°
0.05
0.1
A
B
2 X Ø4,00
21,97
13,75
17,95
R4,76
10,43
8,21
SOLIDWORKS Educational Product. Solo per uso didattico.

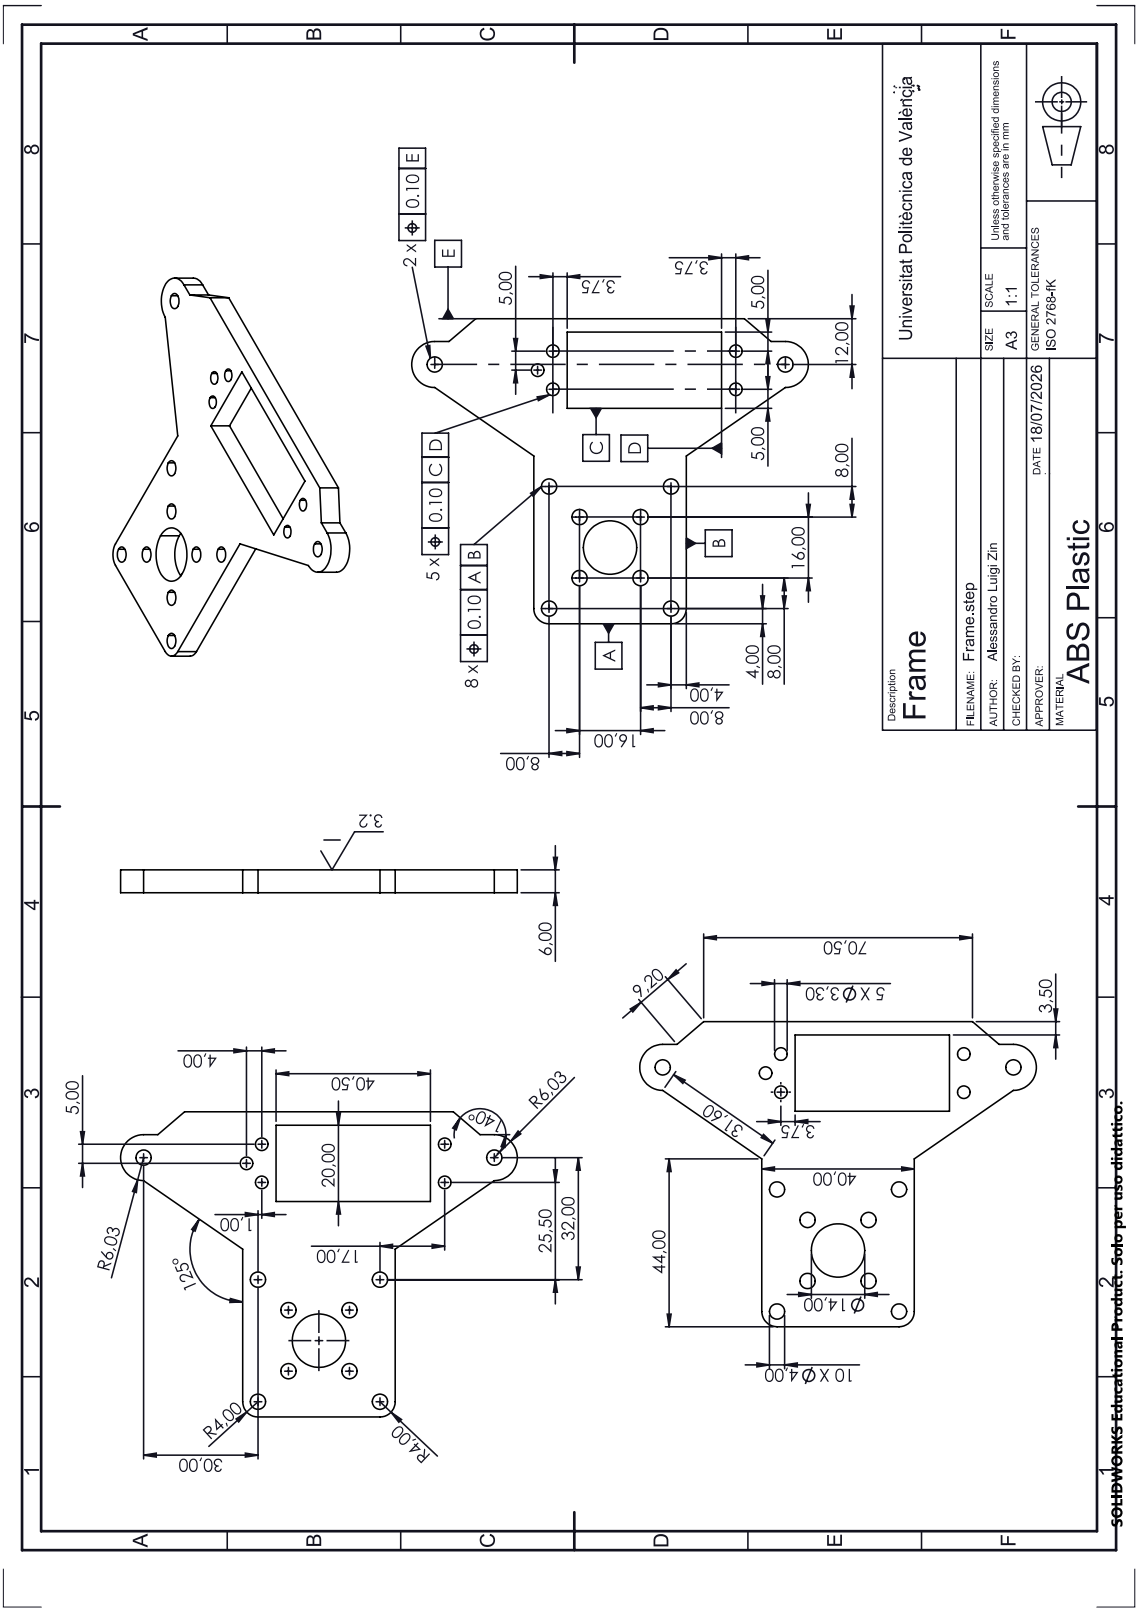
Universitat Politècnica de València
Description
Frame
FILENAME: Frame.step
AUTHOR: Alessandro Luigi Zin
CHECKED BY:
APPROVER:
MATERIAL
ABS Plastic
DATE 18/07/2026
SIZE A3
SCALE 1:1
Unless otherwise specified dimensions and tolerances are in mm
GENERAL TOLERANCES ISO 2768-fK
SOLIDWORKS Educational Product. Solo per uso didattico.

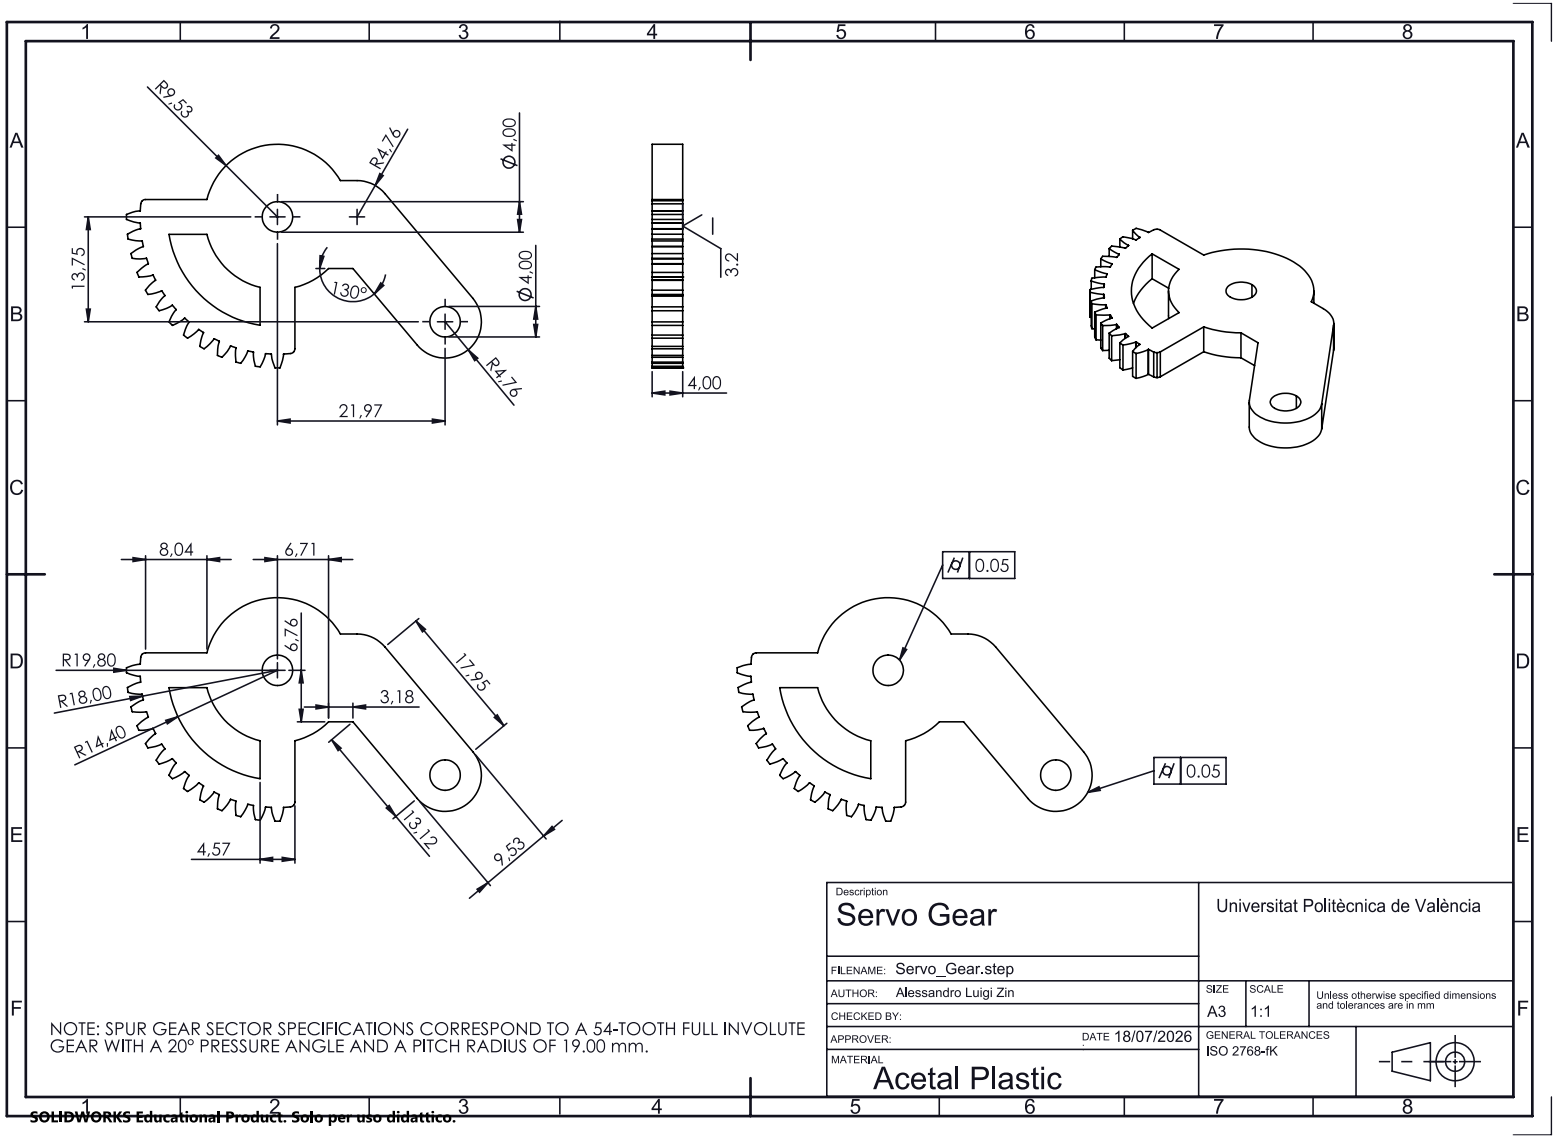

SOLIDWORKS Educational Product. Solo per uso didattico.

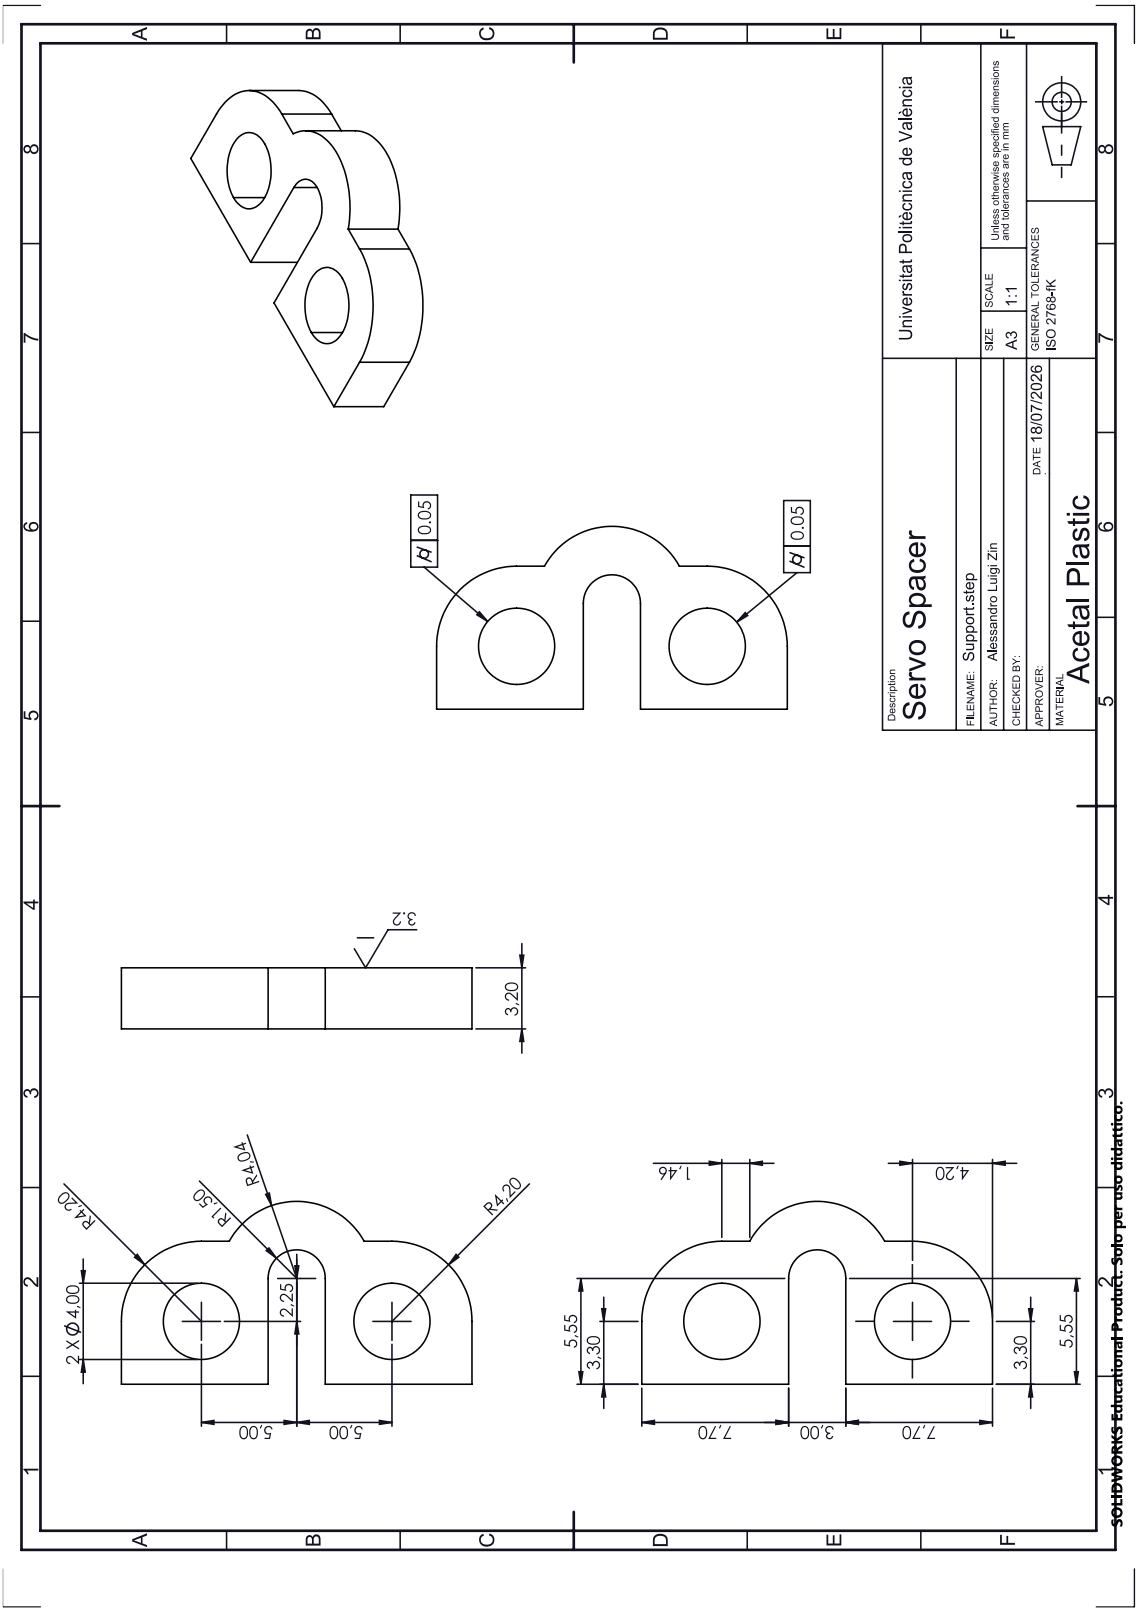
Description
Servo Spacer
Universitat Politècnica de València
FILENAME: Support.step
AUTHOR: Alessandro Luigi Zin
CHECKED BY:
APPROVER:
DATE 18/07/2026
SIZE A3
SCALE 1:1
Unless otherwise specified dimensions and tolerances are in mm
GENERAL TOLERANCES ISO 2768-fK
MATERIAL Acetal Plastic
0.05
0.05
3.2
3.20
R4.20
R4.04
R1.50
R4.20
2 X Ø4.00
2.25
5,00
5,00
1,46
4,20
5,55
3,30
3,30
5,55
7,70
3,00
7,70
SOLIDWORKS Educational Product. Solo per uso didattico.

B - EinScan Pro 2X Inspected Results

This appendix presents the results obtained in GOM Inspect after isolating each relevant component (Frame, Gear Arm, Link and Claw) individually and refining the corresponding mesh in Meshmixer. The values for the Frame are not repeated here, as they have already been discussed in Section 3.2. In the following images the colour scale represents the signed dimensional deviation in millimetres, projected onto the mesh surface.

The surface comparison for every component was performed in GOM Inspect with a maximum search distance of 1,5 [mm], applied consistently across all three scanning technologies considered in this thesis (EinScan, ATOS, CT) to ensure methodological homogeneity of the comparison. The choice is constrained by the geometry of the parts: the thinnest component in the comparison has a minimum wall thickness of 4 [mm], so the search distance was set to less than half this value to prevent the nearest-point algorithm from associating points on opposite surfaces of the same part, which would inflate the apparent deviations on thin features. The chosen value is at the same time large enough to encompass the actual surface deviations observed, which fall approximately within ±0,5 [mm] across all measured datasets, so that no real measurement is rejected as an outlier.

It should also be noted that, only to ease the partial reconstruction of the teeth, the spare Gear Arm was scanned instead of the one mounted on the servomotor. This was done to obtain the best possible result while keeping the data reasonably comparable with the other techniques.

Gear Arm

Figure 98 maps the signed deviation of the Gear Arm mesh against the nominal CAD, and Table 23 summarises its statistics. The mean deviation of +0,04 [mm] indicates no systematic offset; the starred maximum (+1,50) corresponds to points that saturate the 1,5 [mm] search distance in the partially reconstructed tooth region, and should therefore be read as unmatched areas rather than as true surface error.

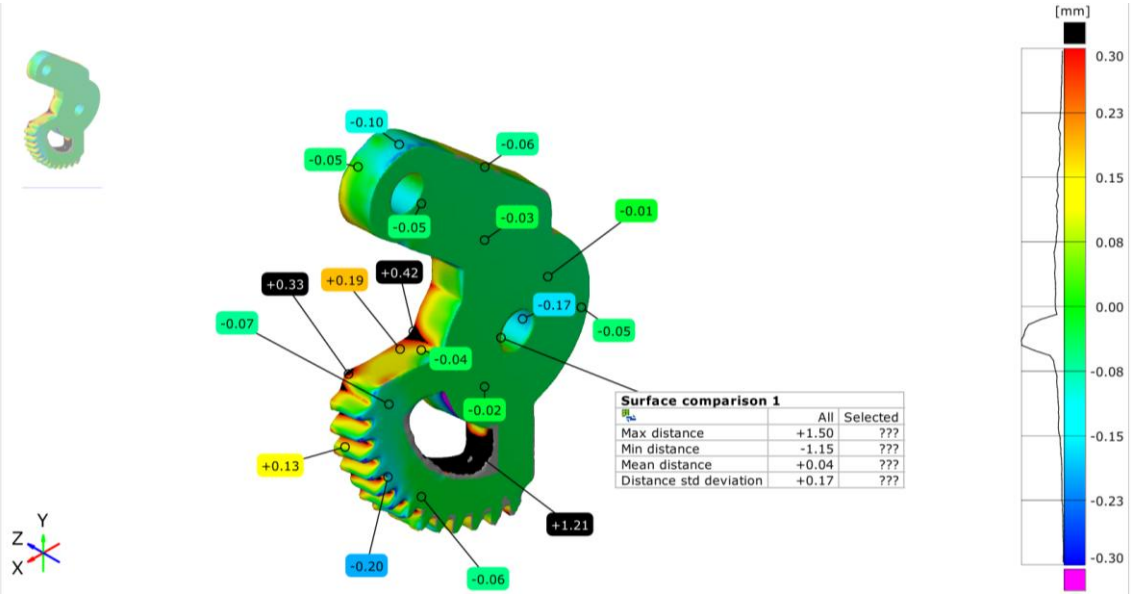

Figure 98: Surface deviation map of the EinScan-derived Gear Arm mesh with respect to the nominal CAD model, obtained after pre-alignment and global best-fit registration in GOM Inspect.

Gear Arm Surface Info	Value [mm]
Max distance	+1,50
Min distance	−1,15
Mean distance	+0,04
Distance std deviation	0,17

Table 23: Maximum, minimum, mean and standard deviation of the surface-distance values shown in Figure 98, for the Gear Arm mesh.

Link

Figure 99 and Table 24 report the same analysis for the Link, the best-behaved component of the handheld campaign: deviations remain within +0,21/−0,40 [mm], with a mean of −0,03 [mm] and a standard deviation of only 0,06 [mm], in line with the fixed-head results of Appendix C.

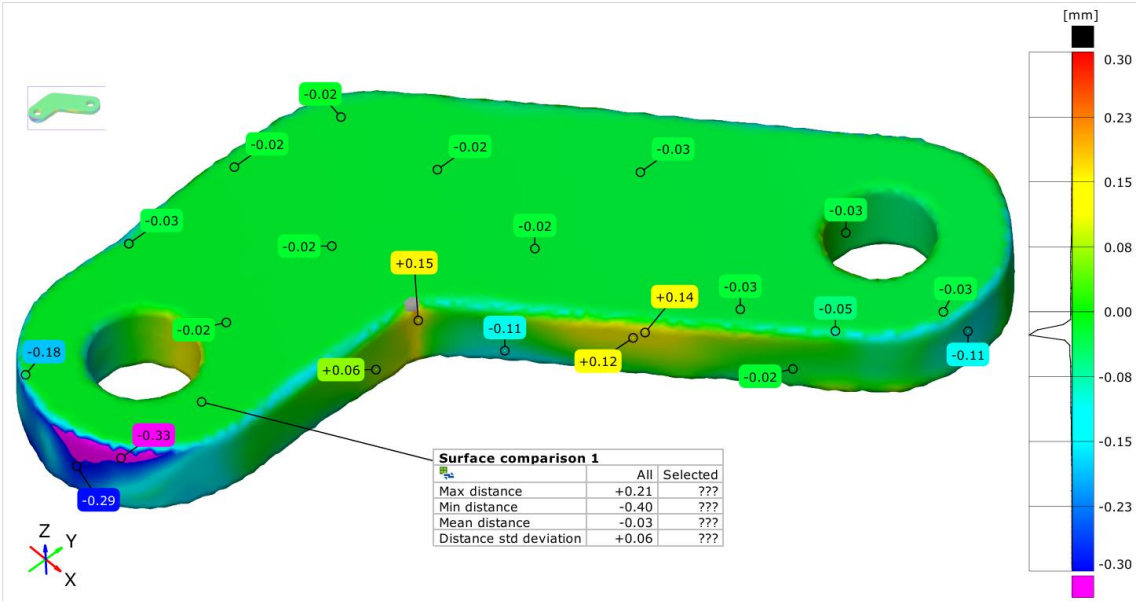

Figure 99: Surface deviation map of the EinScan-derived Link mesh with respect to the nominal CAD model, obtained after pre-alignment and global best-fit registration in GOM Inspect.

Link Surface Info	Value [mm]
Max distance	+0,21
Min distance	−0,40
Mean distance	−0,03
Distance std deviation	0,06

Table 24: Maximum, minimum, mean and standard deviation of the surface-distance values shown in Figure 99 for the Link mesh.

Claw

Figure 100 and Table 25 refer to the Claw. The deviation band is wider here (+0,96/−0,73 [mm]) and the mean of −0,15 [mm] reveals a slight systematic undersize, consistent with the difficulty of capturing its small, partially occluded features with a handheld scanner.

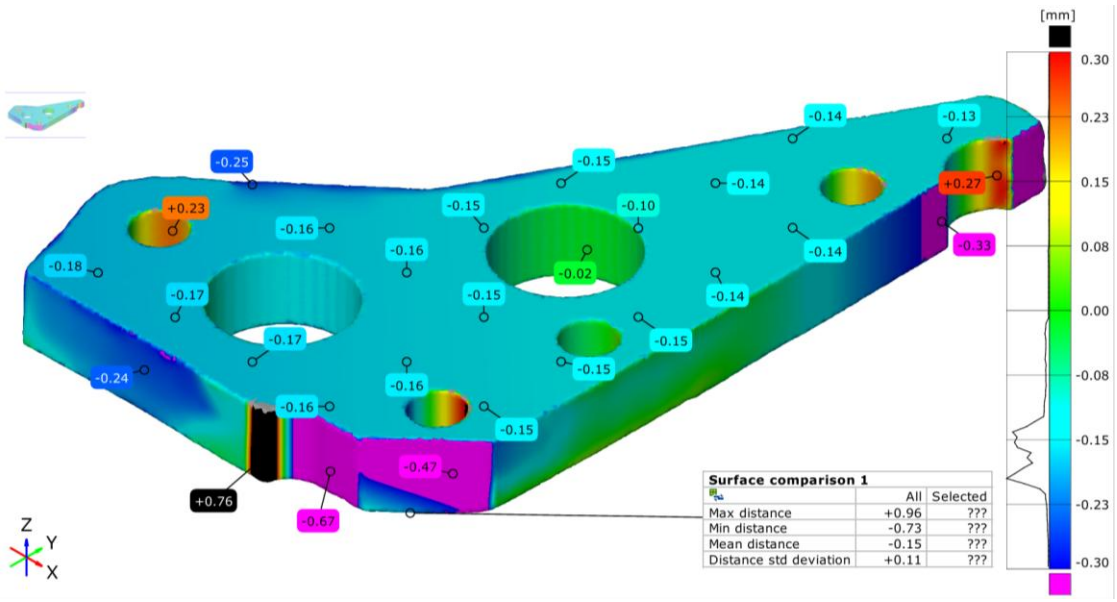

Figure 100: Surface deviation map of the EinScan-derived Claw mesh with respect to the nominal CAD model, obtained after pre-alignment and global best-fit registration in GOM Inspect.

Claw Surface Info	Value [mm]
Max distance	+0,96
Min distance	−0,73
Mean distance	−0,15
Distance std deviation	0,11

Table 25: Maximum, minimum, mean and standard deviation of the surface-distance values shown in Figure 100 for the Claw mesh.

To obtain a single global indicator that is approximately comparable with the other scanning techniques, the four components are aggregated as follows. The global mean is the unweighted arithmetic mean of the per-component mean deviations μ_i. The global standard deviation is computed as a pooled standard deviation:

$$\sigma_{pooled} = \sqrt{(\text{mean}(\sigma_i^2) + \text{var}(\mu_i))}$$

This is the standard formula for the dispersion of a mixture of distributions when each component is given equal weight (i.e. assuming an approximately equal number of evaluation points per mesh). A simple arithmetic mean of the σ_i values would systematically underestimate the global dispersion, because standard deviations are not additive, but variances are. The values reported in Table 26 are computed accordingly.

EinScan Global Info	Value [mm]
Mean distance	−0,04
Distance std deviation	0,147

Table 26: Aggregated EinScan handheld-SLS error indicators (mean and std.) across the four components. See the introductory paragraphs of this appendix for the underlying assumptions.

One final caveat: beyond being only indicative, these values are also affected by several sources of error: most notably the fact that the end-effector was acquired as a single assembly, unlike with the other techniques, and then segmented through cuts and Boolean operations; and that the Gear Arm measured here is the spare unit, not the one mounted on the servomotor. Despite these caveats, the data obtained are reasonably good considering how they were produced and are used in Section 3.5 as the basis for a qualitative, rather than strictly numerical, comparison with the other scanning methods.

C - GOM ATOS Core 300 Inspected Results

This appendix presents the results obtained in GOM Inspect of each relevant component (Frame, Gear Arm, Link and Claw) individually. The values for the Link are not repeated here, as they have already been discussed in Section 3.3. In the following images the colour scale represents the signed dimensional deviation in millimetres, projected onto the mesh surface.

The surface comparison for every component was performed in GOM Inspect with a maximum search distance of 1,5 [mm], applied consistently across all three scanning technologies considered in this thesis (EinScan, ATOS, CT) to ensure methodological homogeneity of the comparison. The choice is constrained by the geometry of the parts: the thinnest component in the comparison has a minimum wall thickness of 4 [mm], so the search distance was set to less than half this value to prevent the nearest-point algorithm from associating points on opposite surfaces of the same part, which would inflate the apparent deviations on thin features. The chosen value is at the same time large enough to encompass the actual surface deviations observed, which fall approximately within ±0,5 [mm] across all measured datasets, so that no real measurement is rejected as an outlier.

Gear Arm

Figure 101 maps the signed deviation of the ATOS-derived Gear Arm mesh against the nominal CAD, and Table 27 summarises it. The distribution is well centred (mean −0,01 [mm], standard deviation 0,16 [mm]), with the extreme values (+0,86 / −0,87 [mm]) confined to the tooth flanks and to the thin edges of the arm.

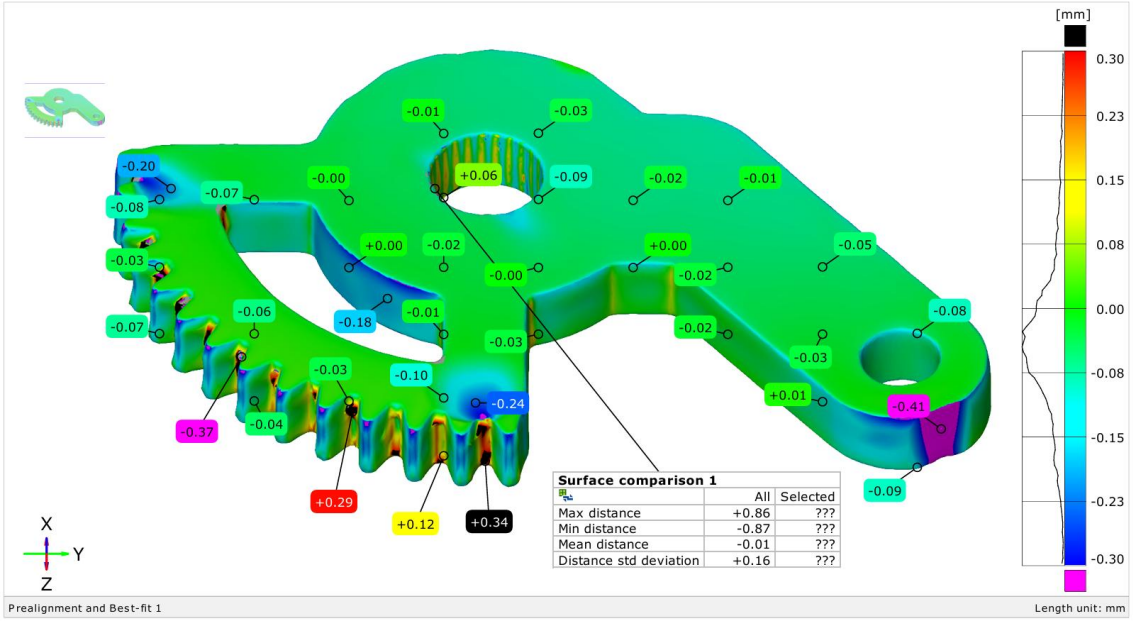

Figure 101: ATOS deviation map of the Gear Arm against the nominal CAD model. The map visualises the local distance between the acquired mesh and the reference geometry, expressed in millimetres.

Gear Arm Surface Info	Value [mm]
Max distance	+0,86
Min distance	−0,87
Mean distance	−0,01
Distance std deviation	0,16

Table 27: Maximum, minimum, mean and standard deviation of the surface-distance values shown in Figure 101, for the Gear Arm mesh.

Frame

Figure 102 and Table 28 report the same comparison for the Frame, which shows the widest band of the three parts (maximum +1,41 [mm]) but only a small systematic offset (mean −0,07 [mm], standard deviation 0,13 [mm]); the largest local deviations fall on the tall, thin walls that are hardest for the fixed-head sensor to reach.

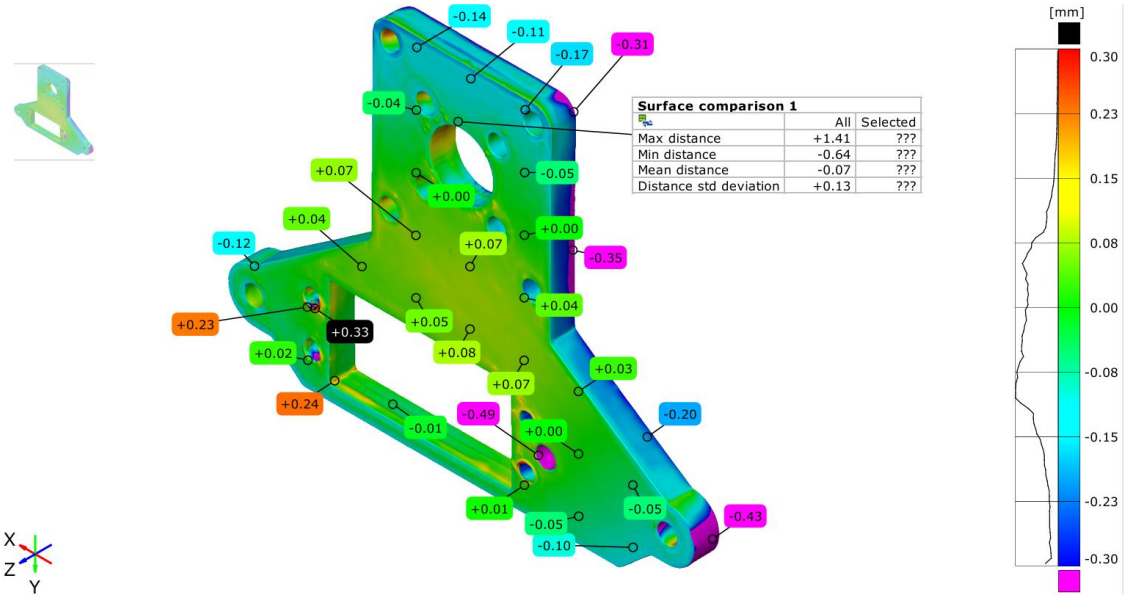

Figure 102: ATOS deviation map of the Frame against the nominal CAD model. The map visualises the local distance between the acquired mesh and the reference geometry, expressed in millimetres.

Frame Surface Info	Value [mm]
Max distance	+1,41
Min distance	−0,64
Mean distance	−0,07
Distance std deviation	0,13

Table 28: Maximum, minimum, mean and standard deviation of the surface-distance values shown in Figure 102 for the Frame mesh.

Claw

Figure 103 and Table 29 show the Claw, the most accurate of the three components measured with the ATOS: the deviation stays within a narrow band (+0,27 / −0,44 [mm]), with a mean of −0,03 [mm] and a standard deviation of 0,08 [mm].

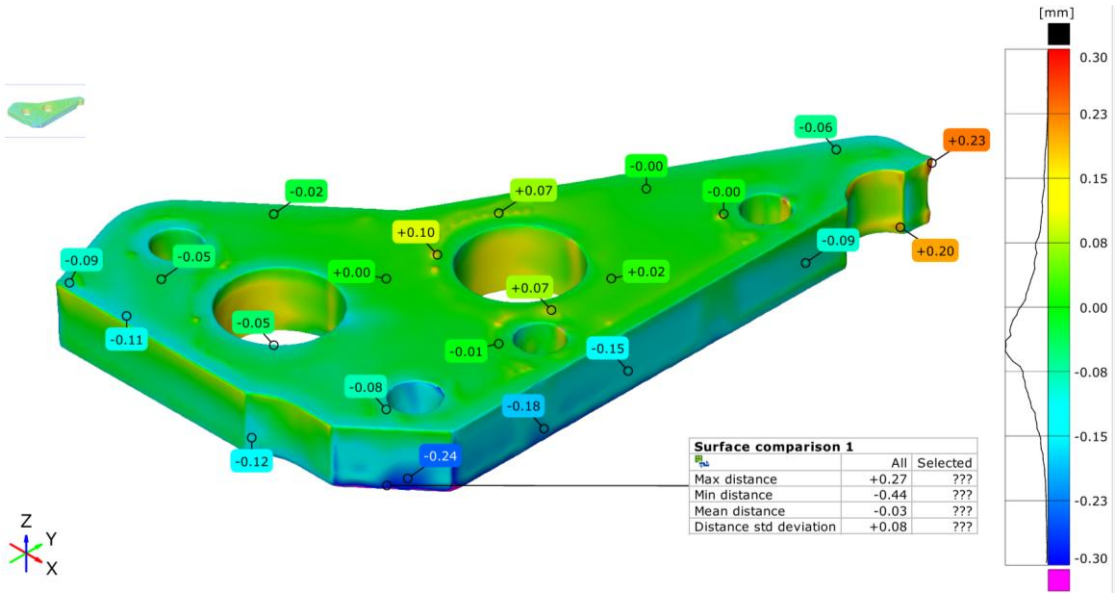

Figure 103: ATOS deviation map of the Claw against the nominal CAD model. The map visualises the local distance between the acquired mesh and the reference geometry, expressed in millimetres.

Claw Surface Info	Value [mm]
Max distance	+0,27
Min distance	−0,44
Mean distance	−0,03
Distance std deviation	0,08

Table 29: Maximum, minimum, mean and standard deviation of the surface-distance values shown in Figure 103 for the Claw mesh.

The four components are aggregated into a single global indicator exactly as described for the EinScan case in Appendix B: the global mean is the unweighted mean of the per-component means, and the global standard deviation is obtained as the pooled standard deviation of the individual distributions. The resulting values are reported in Table 30.

ATOS Global Info	Value [mm]
Mean distance	−0,045
Distance std deviation	0,117

Table 30: Aggregated GOM ATOS Core 300 fixed-head SLS error indicators (mean and std.) across the four components. See the introductory paragraphs of this appendix for the underlying assumptions.

Again, these values are only indicative and are affected by additional sources of error specific to this technique. In particular, each component was acquired separately, after disassembling the end-effector, so the individual scans do not share a common reference frame and small part-to-part alignment differences add to the measured deviations; moreover, the Gear Arm measured here is the spare unit, not the one mounted on the servomotor. Despite these caveats, the data are reasonably good for the intended purpose and are used in Section 3.5 as the basis for a qualitative, rather than strictly numerical, comparison with the other scanning methods.

D - Phoenix v|tome|x s Results

Given the interest that motivated the use of this technique, this appendix collects, for each component:

- All the datos|x images (radiographs), their grey-value histograms, and the values set on (or returned by) the machine, to guarantee the reproducibility of the experiment;
- all the views of the meshes once cleaned in VG Studio software, so that the reconstructed geometry can be inspected;
- all the GOM Inspect comparisons, presented in two states: first the surface comparison and the section of the part without removing the internal material and without closing the holes, and then the surface comparison and the section of the part after removing the internal material with the HPR algorithm described in Appendix E and closing as many holes as possible;
- a discussion, as already carried out in the main text, of the overall mean and standard deviation across the components.

The servo support is included here as well. It is reported for completeness in the images, but it is NOT taken into account in the standard-deviation and mean calculations, to keep the same number of samples for all three processes in the cross-scanner comparison of Section 3.5.

A general remark on image quality: all the following images are clearer than those reported for the Gear Arm. This is because the Gear Arm was scanned first, with a tube current of 120 [µA], whereas for all the subsequent parts (starting with the Frame) this value was raised to 160 [µA] to image the regions with abundant material more sharply. This setting was then kept for the Claw, the servo support, and the Link.

A note on the operating principle, useful before reading the per-part results: as already established in Section 3.4, a CT scan reconstructs not only the outer surface of the part but also its internal geometry (inner shell, infill, and the cavities left by the layered FDM deposition). When the raw mesh is compared against the nominal CAD with a finite search distance, these internal features are wrongly associated by the nearest-point algorithm and bias both the mean and the standard deviation. For this reason, every component is shown in two states, before and after the HPR treatment, so that the improvement can be appreciated directly.

Link

Figure 104 to Figure 109 and Table 32 document the Link and the servo support.

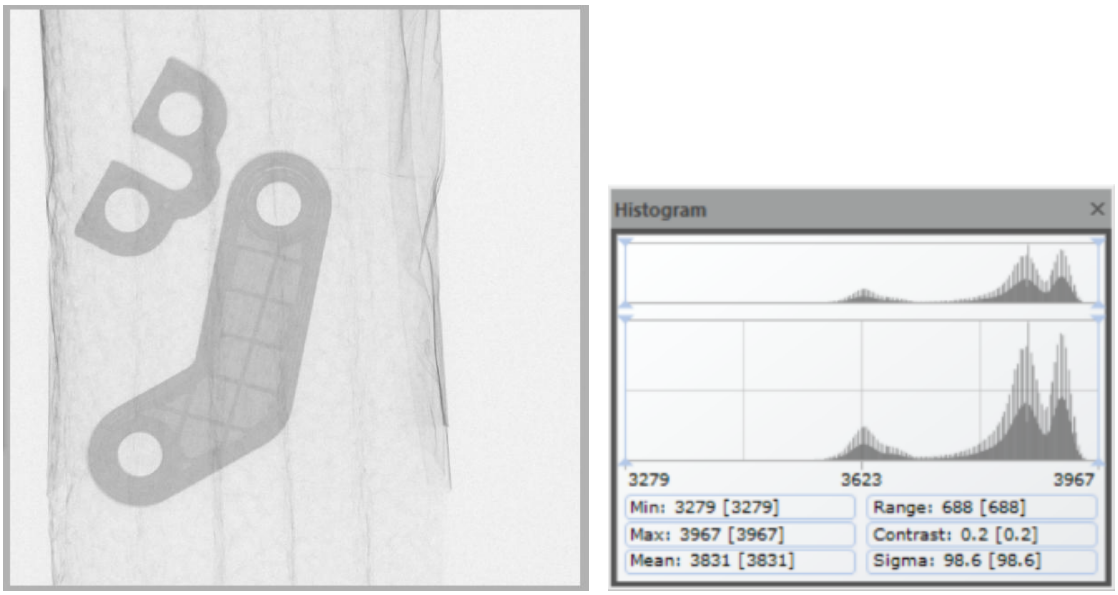

Figure 104: (left) one of the 1500 radiographs acquired for the Link and servo support; (right) the corresponding three-peak grey-value histogram. The three peaks correspond to the background (air), the low-density infill, and the solid material.

Scan information	Value / Unit
Geometry	
Magnification	3,874563
Voxel size	51,618744 [µm]
FOD	210 [mm]
Acquisition	
Number of images	1500
Resolution Images	990 × 1000 pixels
Detector	
Type	DXR-250 RT
Timing	200 [ms]
Skip frames	4
X-Ray	
Voltage	160 [kV]
Current	160 [µA]
Filter	0,1 Cu [mm]
Focal spot size	25,6 [µm]

Table 31: Values set on the Phoenix v|tome|x s machine for the acquisition of the Link and servo support.

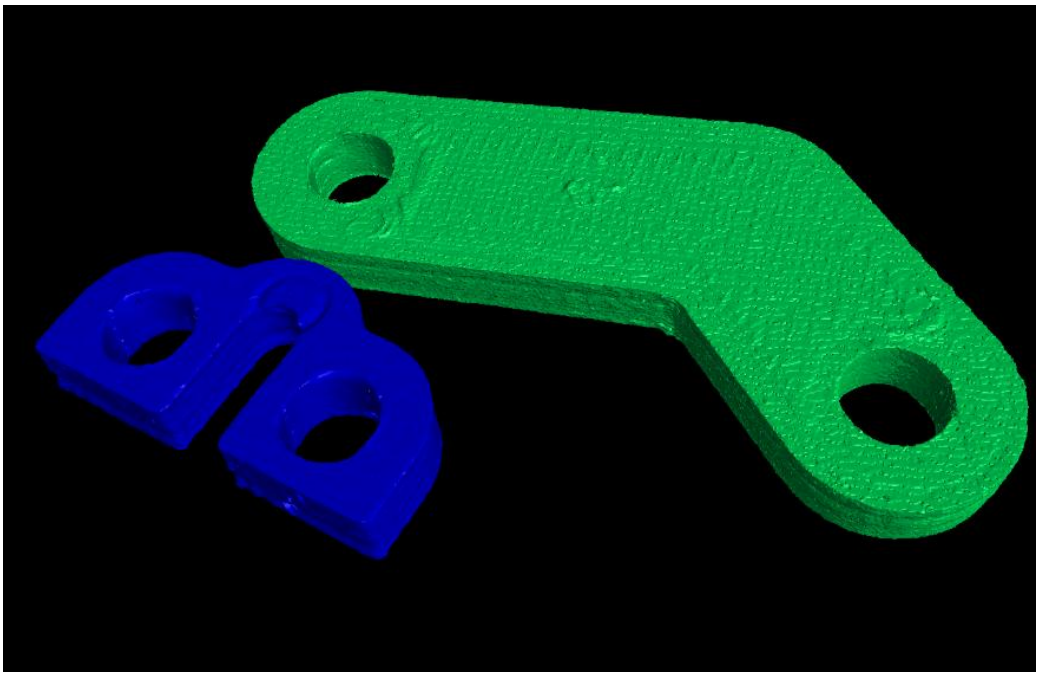

(a)

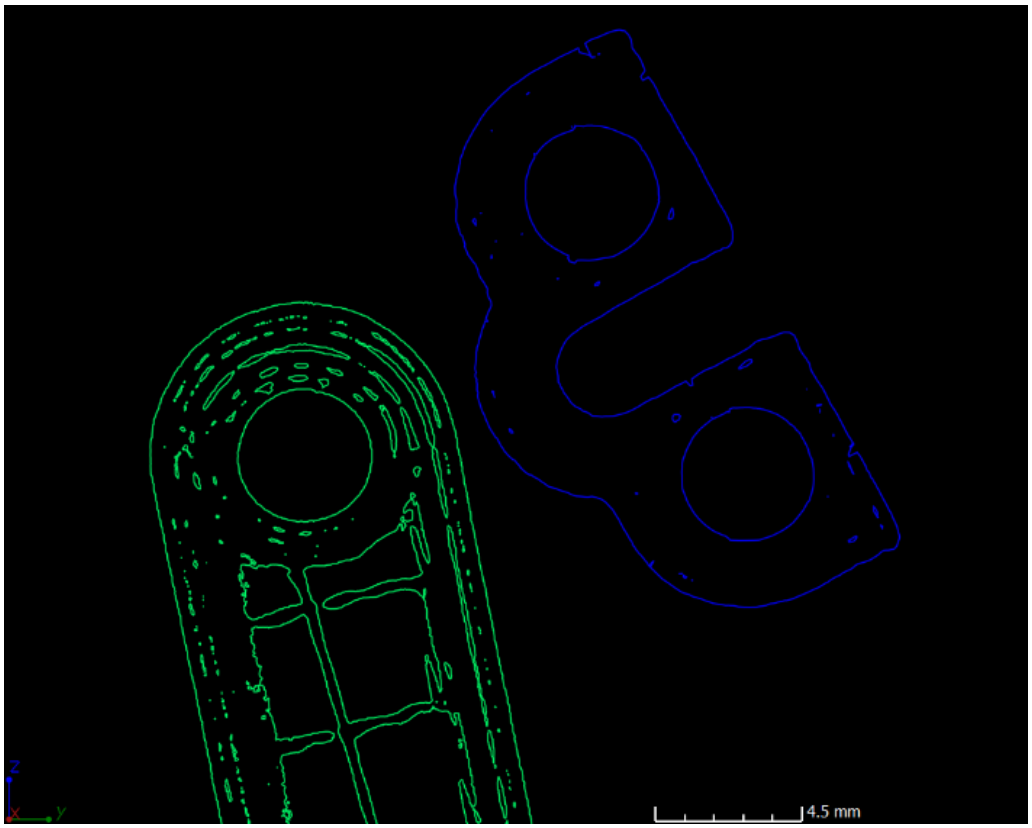

(b)

Figure 105: (a) full view of the Link (green) and servo support (blue) reconstructed after cleaning the grey-value histogram of the background and the low-density material; (b) view of the Link (green) and servo support (blue) mesh cleaned in VG Studio and sectioned to reveal the infill structure.

In the figure above, the servo support, being 3D-printed as well, presents clearly visible cavities even though it should in principle be 100 [%] solid (fully dense). These cavities introduce errors once the mesh is imported into GOM Inspect.

As obtained from VG Studio, the mesh imported into GOM Inspect contained 1526 holes for the Link and 217 for the support. The comparison is therefore carried out between this raw, perforated mesh and the version corrected through the HPR algorithm and the GOM Inspect hole-closing routine.

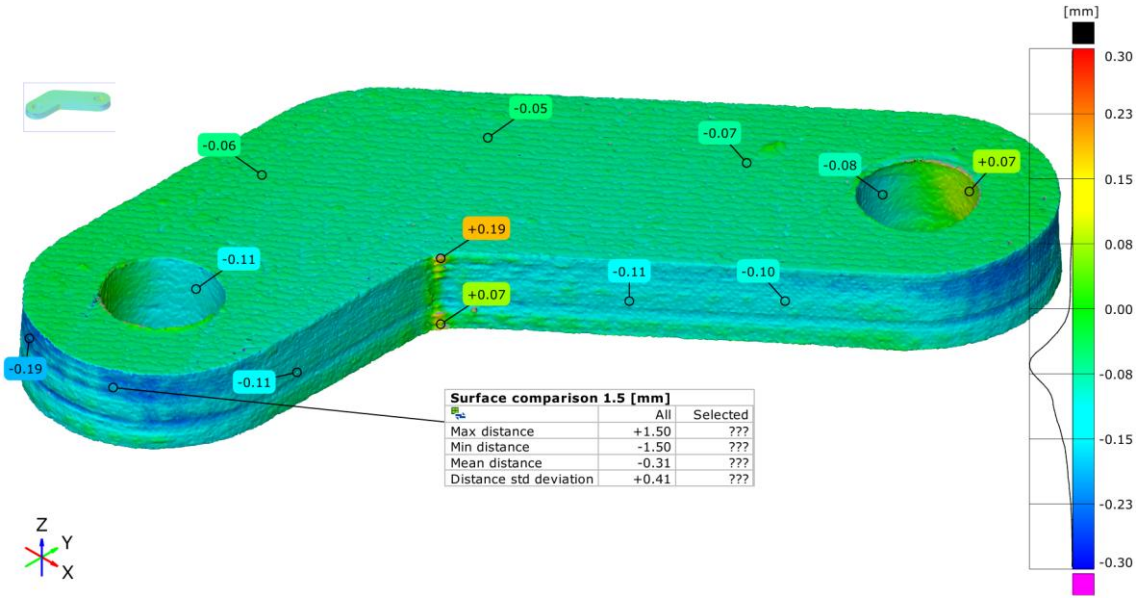

(a)

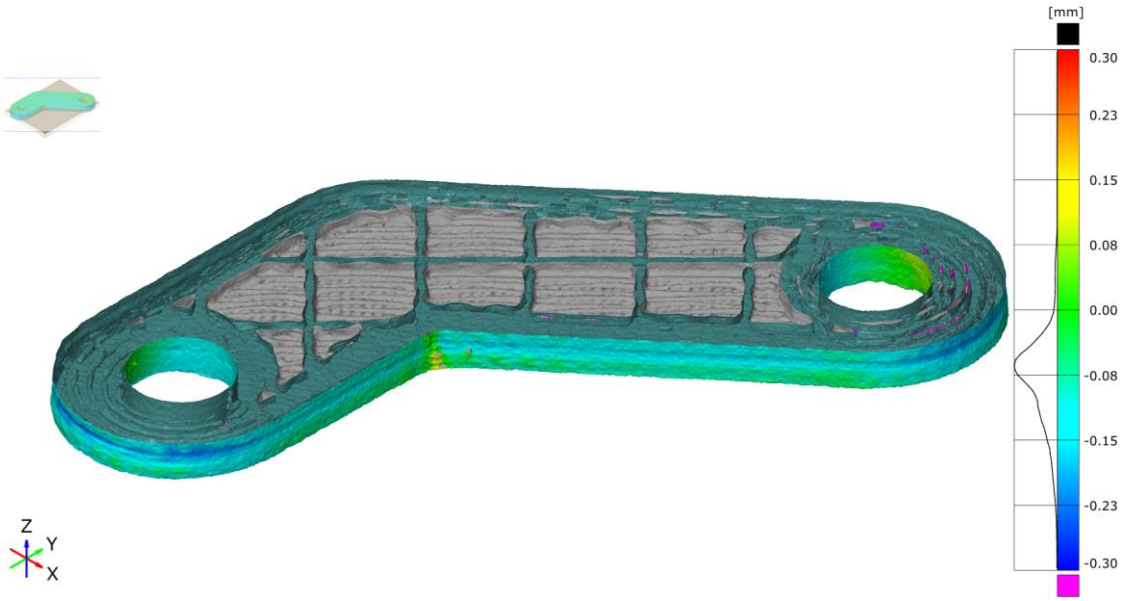

(b)

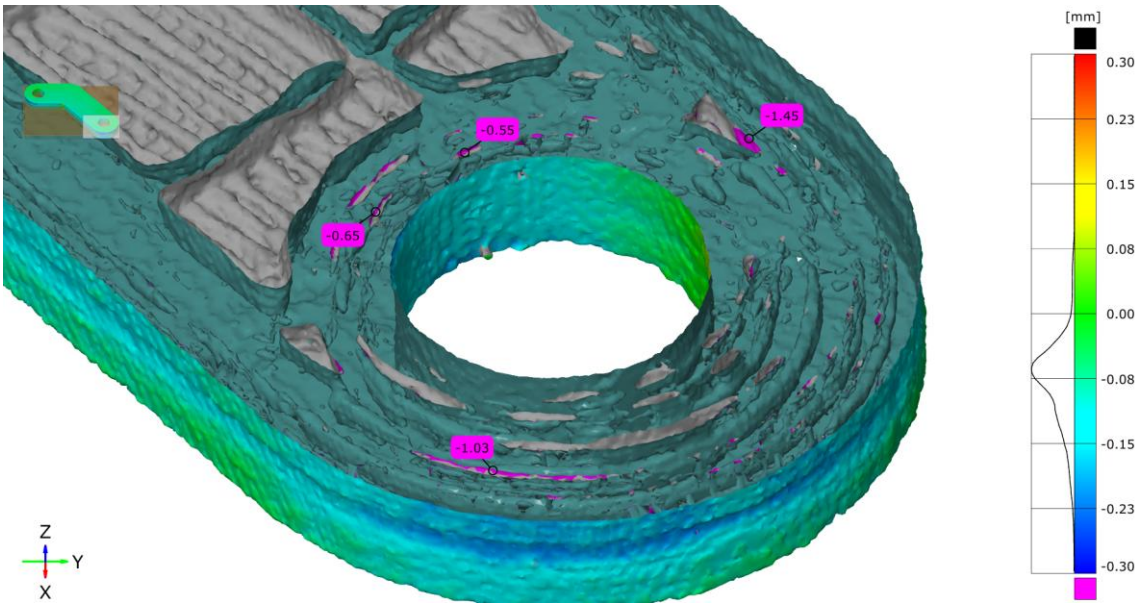

(c)

Figure 106: (a) CT deviation map of the original full Link against the nominal CAD model; (b) section view of the original full Link; (c) internal surfaces wrongly recognised and considered by the comparison.

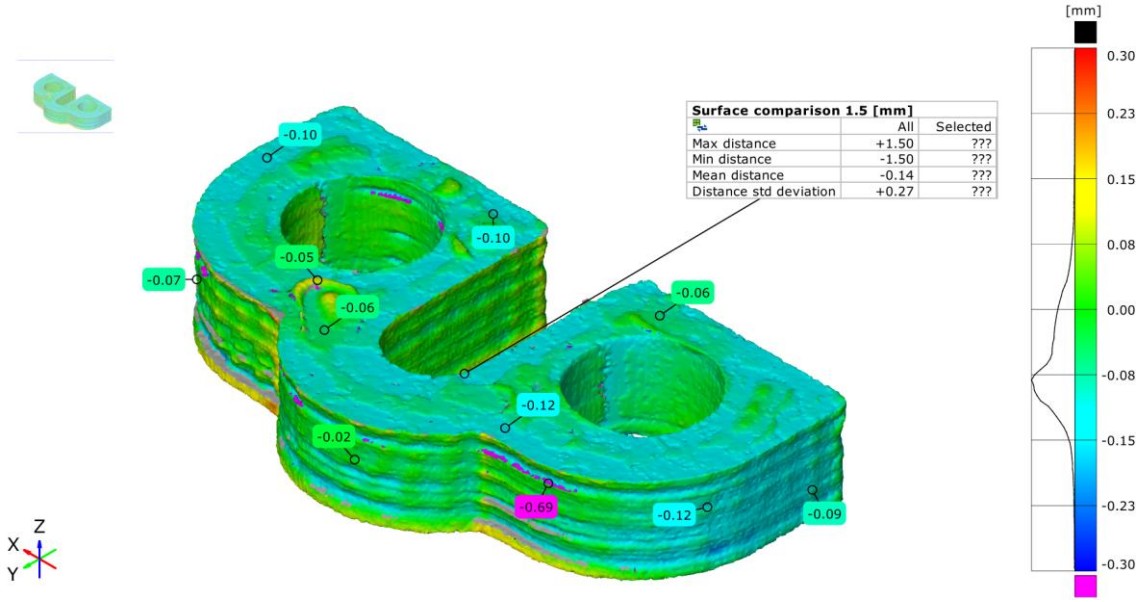

(a)

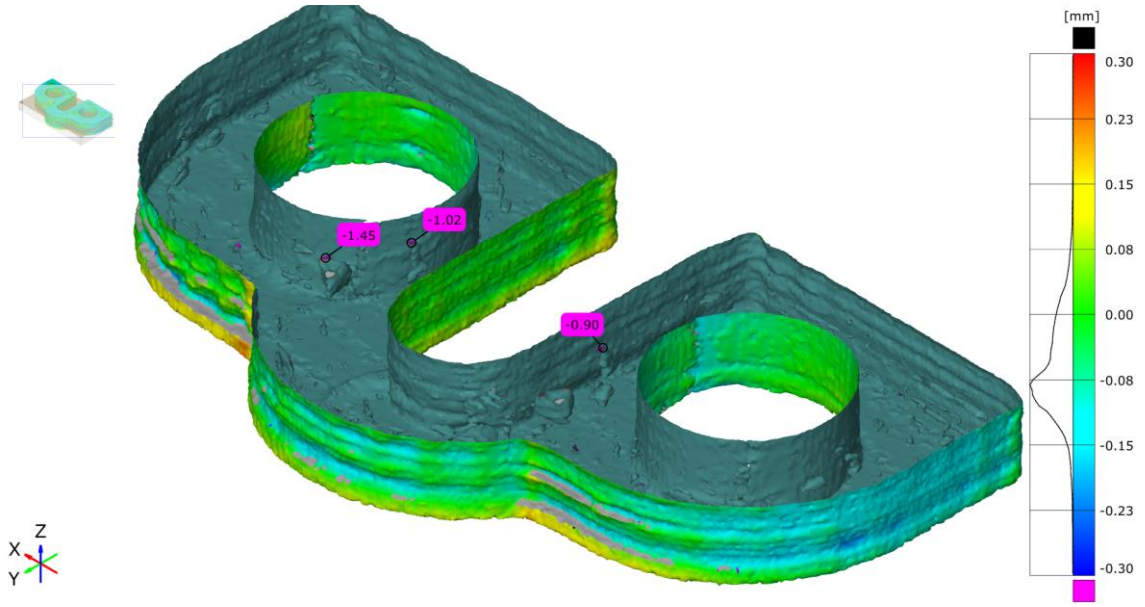

(b)

Figure 107: (a) CT deviation map of the original full servo support against the nominal CAD model; (b) support section with some cavities marked, wrongly considered by the comparison.

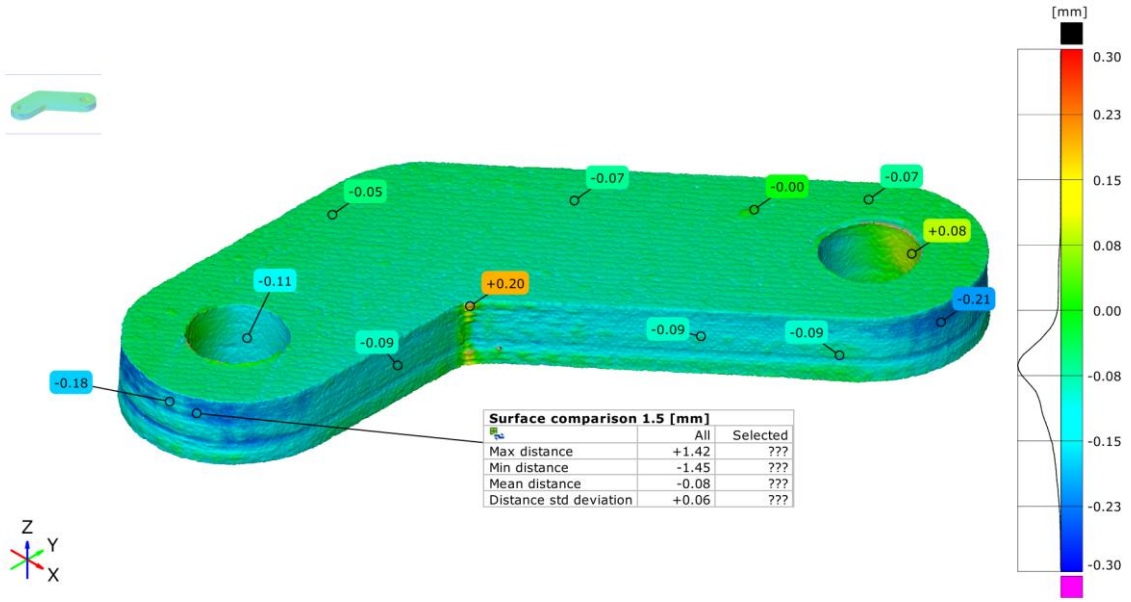

(a)

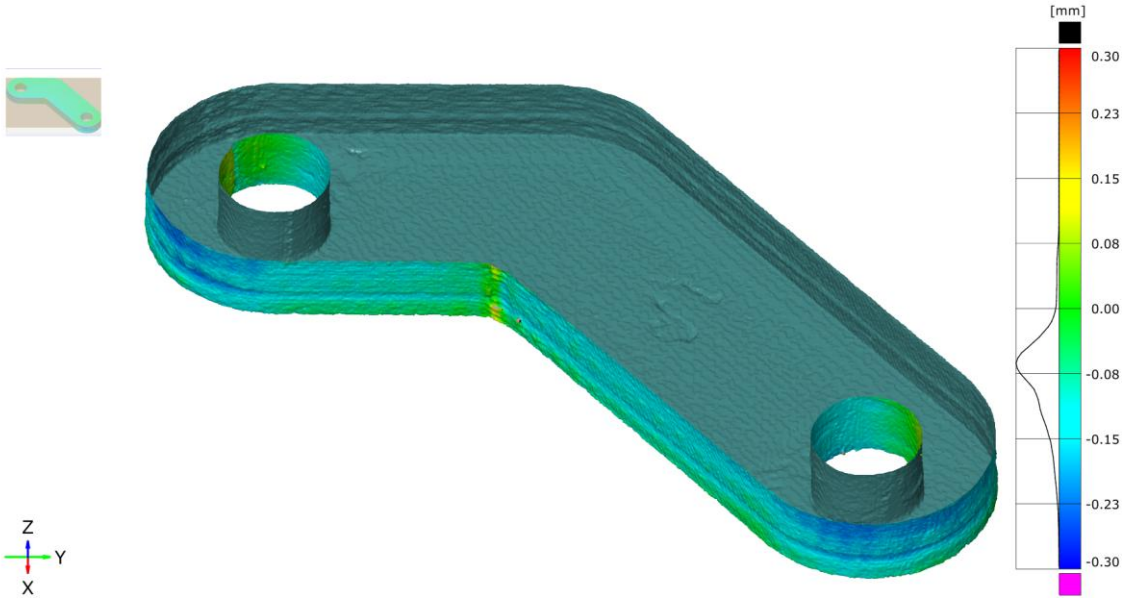

(b)

Figure 108: (a) CT deviation map of the emptied and repaired Link against the nominal CAD model; (b) section of the emptied and repaired Link (note the difference with Figure 106-b).

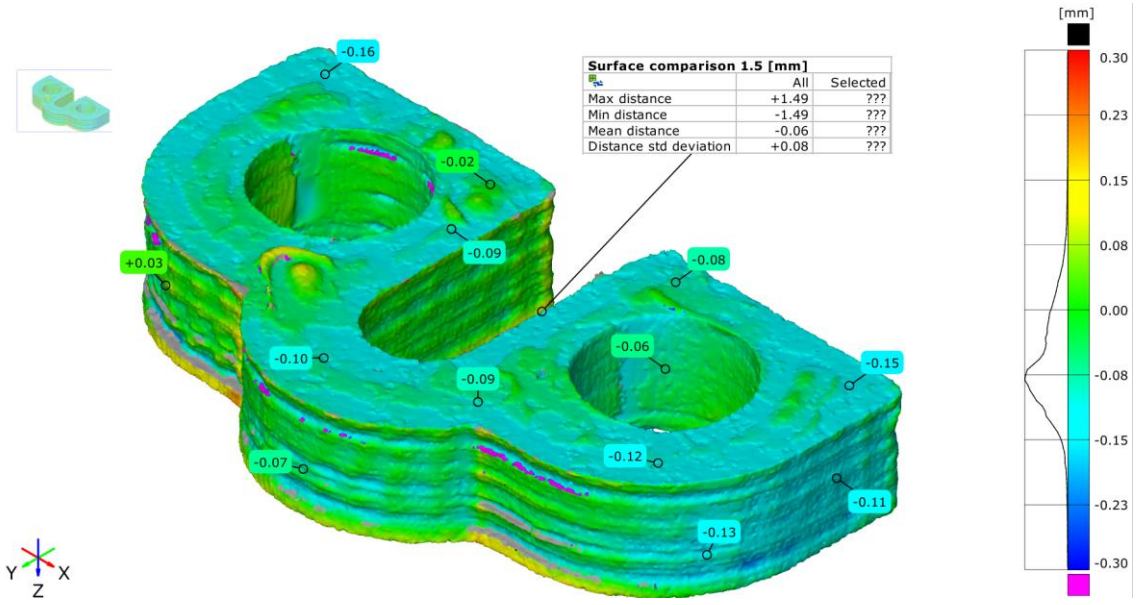

(a)

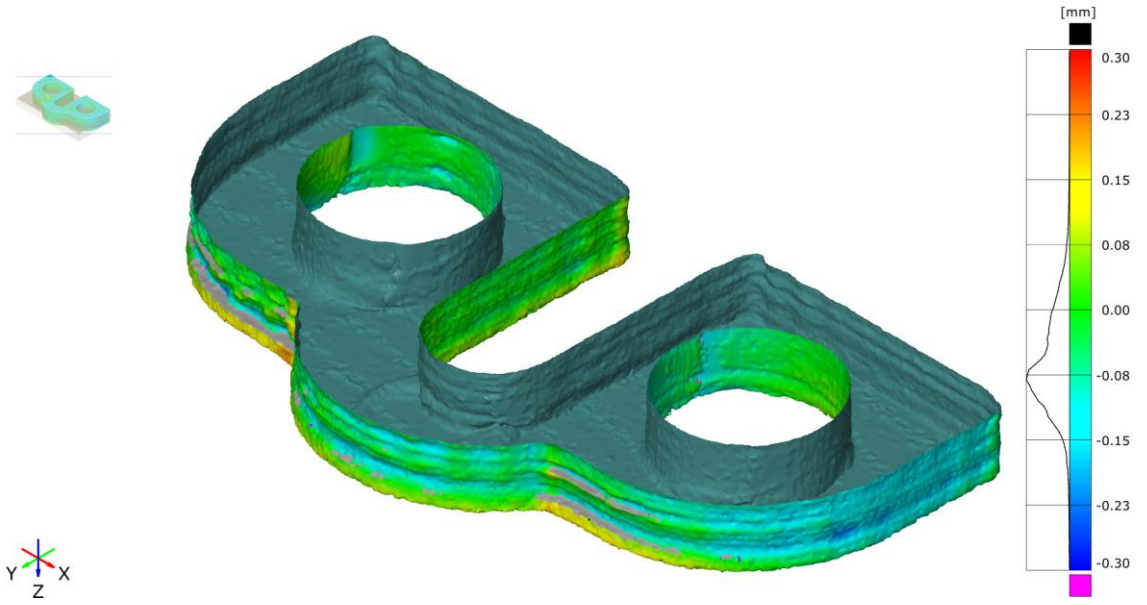

(b)

Figure 109: (a) CT deviation map of the emptied and repaired servo support against the nominal CAD model; (b) section of the emptied and repaired servo support (note the difference with Figure 107-b).

We recall that the support data serve only to show that cavities and printing defects are present even in nominally solid parts; these mean and standard-deviation values are not included in the global calculation.

Component	Max [mm]	Min [mm]	Mean [mm]	St.D. [mm]	# of Holes
Link					
Original	+1,50	−1,50	−0,31	0,41	1526
After HPR and Holes treatments	+1,42	−1,45	−0,08	0,06	0
Support					
Original	+1,5	−1,5	−0,14	0,27	217
After HPR and Holes treatments	+1,49	−1,49	−0,06	0,08	0

Table 32: Comparison between the original Link and servo support and their state after application of the HPR algorithm and hole closing in GOM Inspect.

The effect of the treatment is striking for the Link: the mean deviation moves from −0,31 to −0,08 [mm] and the standard deviation collapses from 0,41 to 0,06 [mm], a sevenfold reduction. This confirms that the bulk of the original scatter was due to the internal geometry rather than to the outer surface, which is the only surface of metrological interest for the comparison with the other methods. The same trend, though milder, is visible on the support.

Claw

Figure 110 to Figure 113 and Table 34 document the Claw.

A short note: in this figure and in the one for the Frame, the radiographs shown are not complete. This is because datos|x, although it holds the scan data at different Z levels, displays only one selected slice during this phase. This phase serves mainly to check, in the radiographs, whether the part has moved during acquisition.

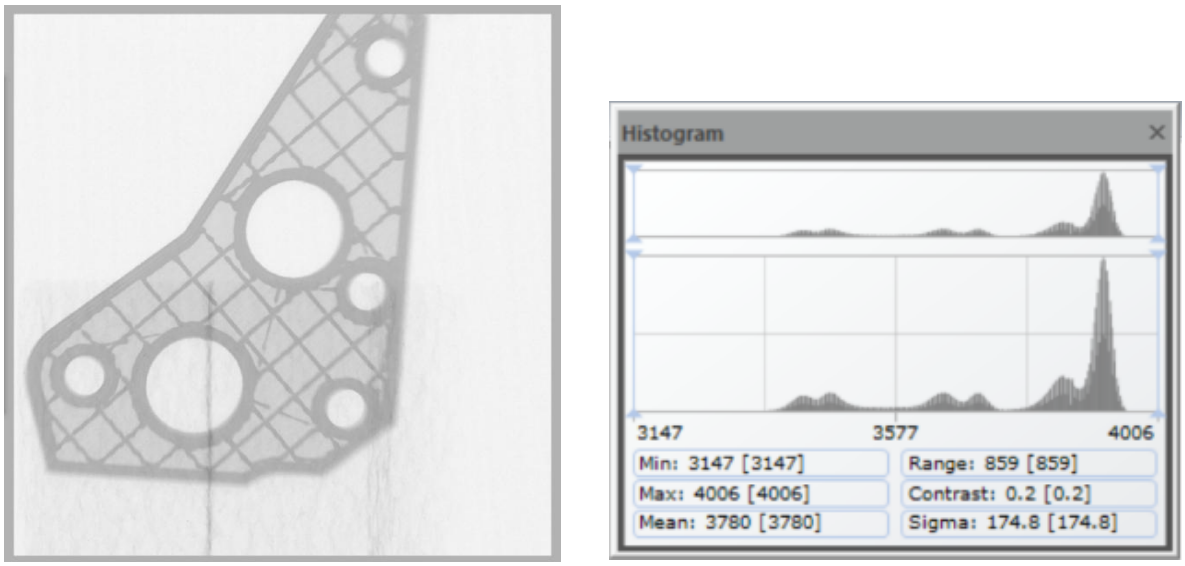

Figure 110: (left) one of the 1500 radiographs acquired for the Claw; (right) the corresponding three-peak grey-value histogram.

Scan information	Value / Unit
Geometry	
Magnification	3,616259
Voxel size	55,305779 [µm]
FOD	225 [mm]
Acquisition	
Number of images	1500
Resolution Images	990 × 1000 pixels
Detector	
Type	DXR-250 RT
Timing	200 [ms]
Skip frames	4
X-Ray	
Voltage	160 [kV]
Current	160 [µA]
Filter	0,1 Cu [mm]
Focal spot size	25,6 [µm]

Table 33: Values set on the Phoenix v|tome|x s machine for the acquisition of the Claw.

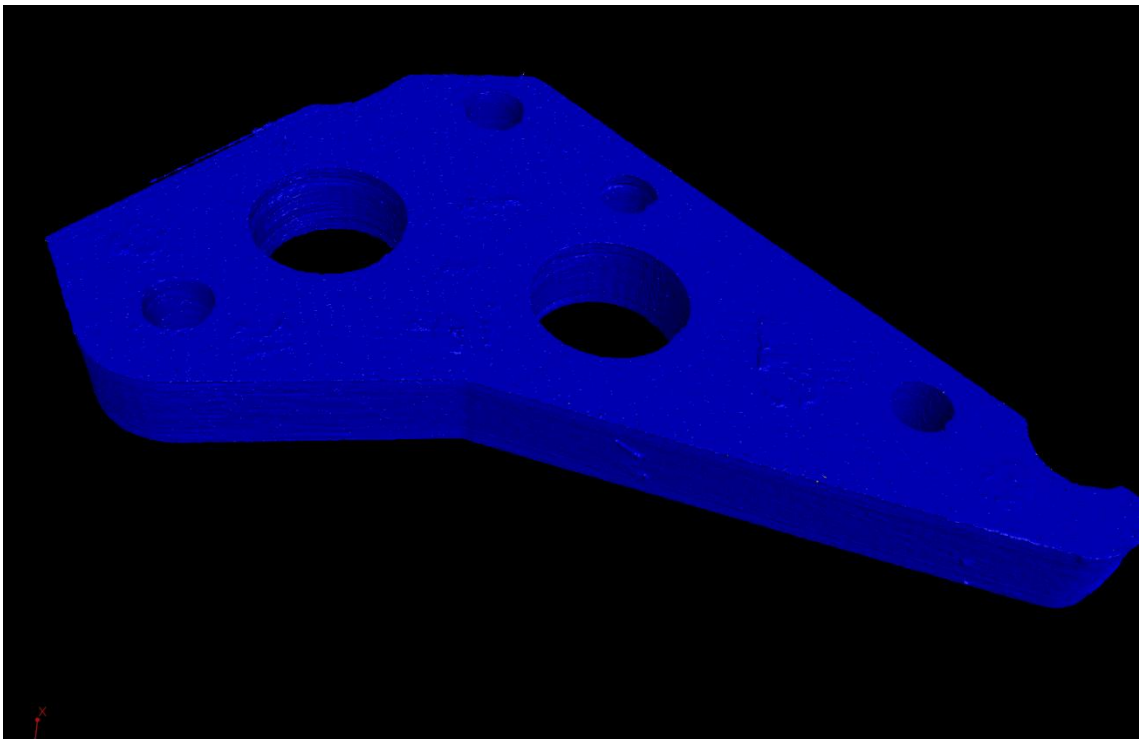

(a)

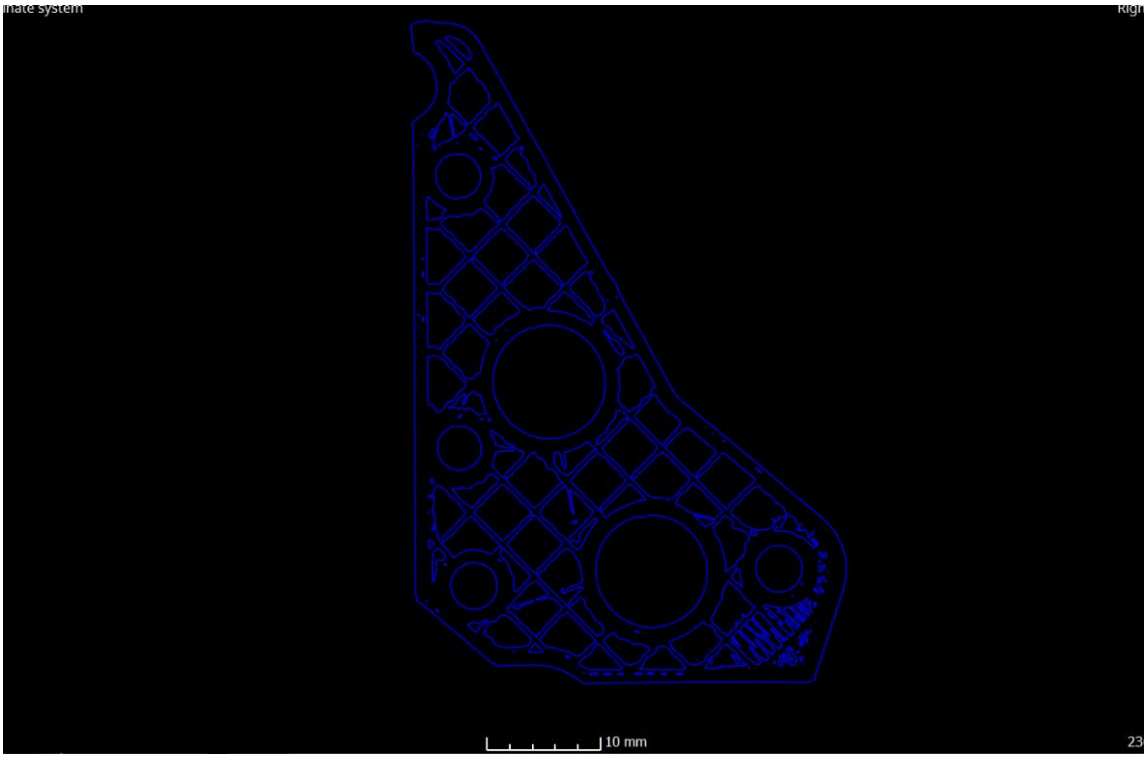

(b)

Figure 111: (a) full view of the Claw reconstructed after cleaning the grey-value histogram of the background and the low-density material; (b) view of the Claw cleaned in VG Studio and sectioned to reveal the infill structure.

The raw mesh presented 1650 holes; the reworked mesh, after HPR and hole closing, presented none.

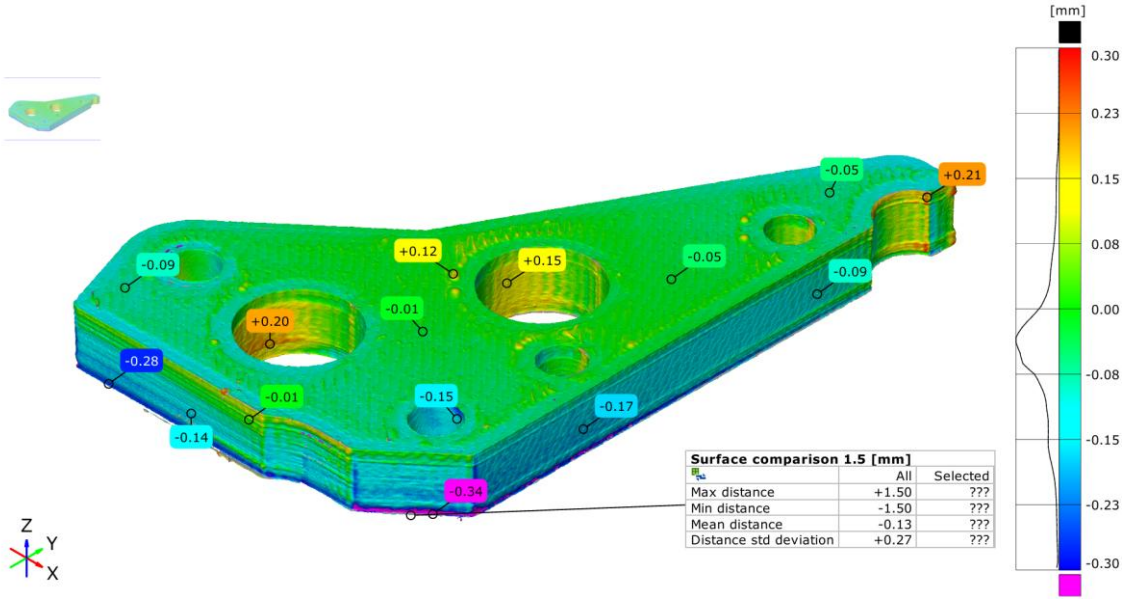

(a)

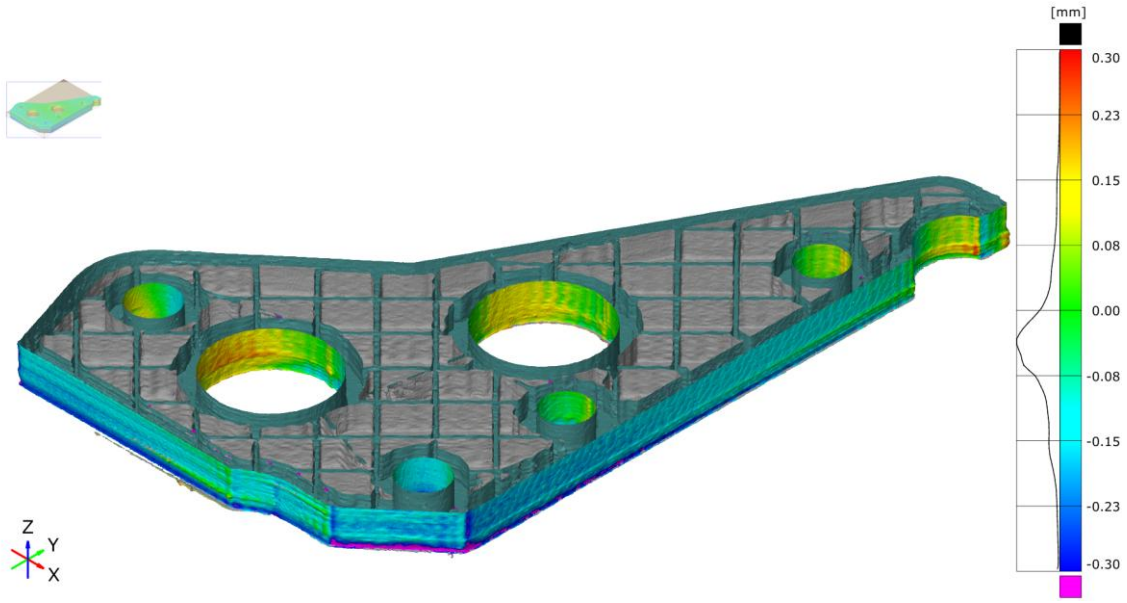

(b)

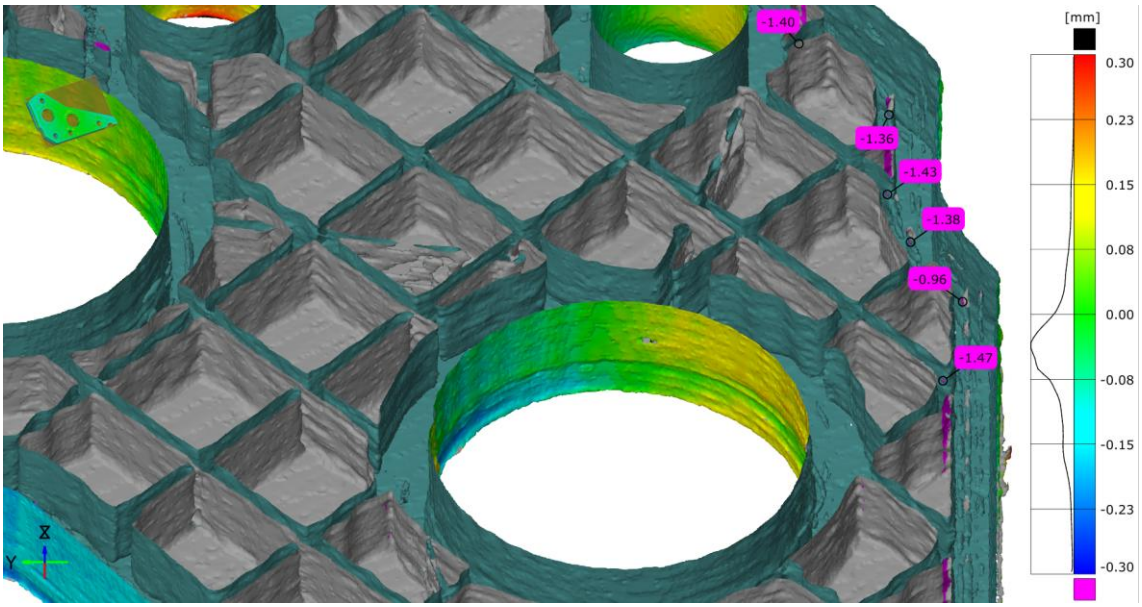

(c)

Figure 112: (a) CT deviation map of the original full Claw against the nominal CAD model; (b) section view of the original full Claw; (c) internal surfaces wrongly recognised and considered by the comparison.

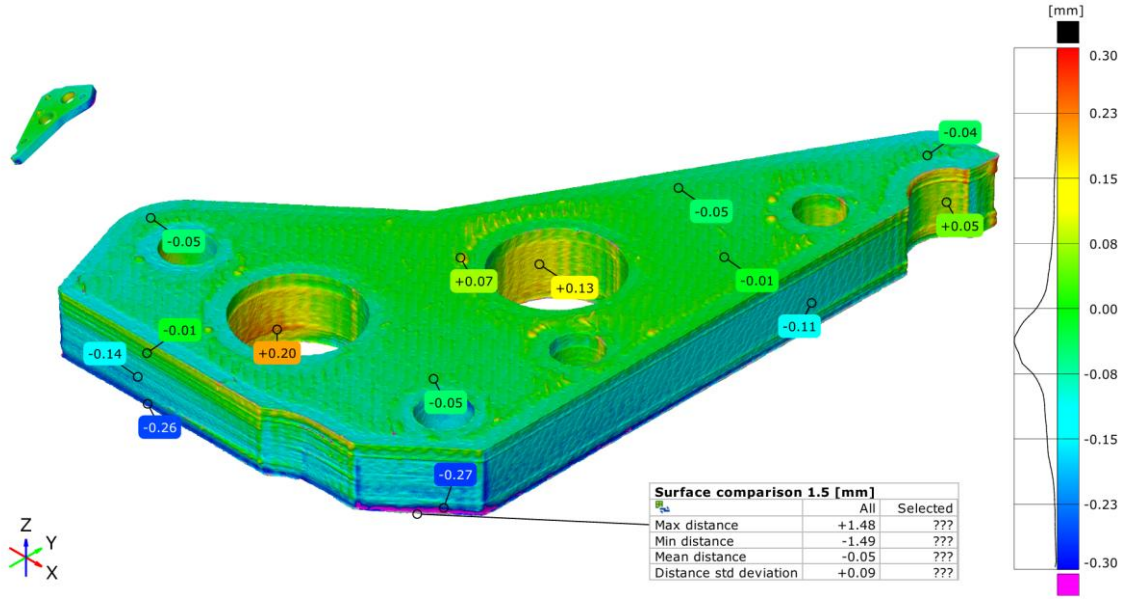

(a)

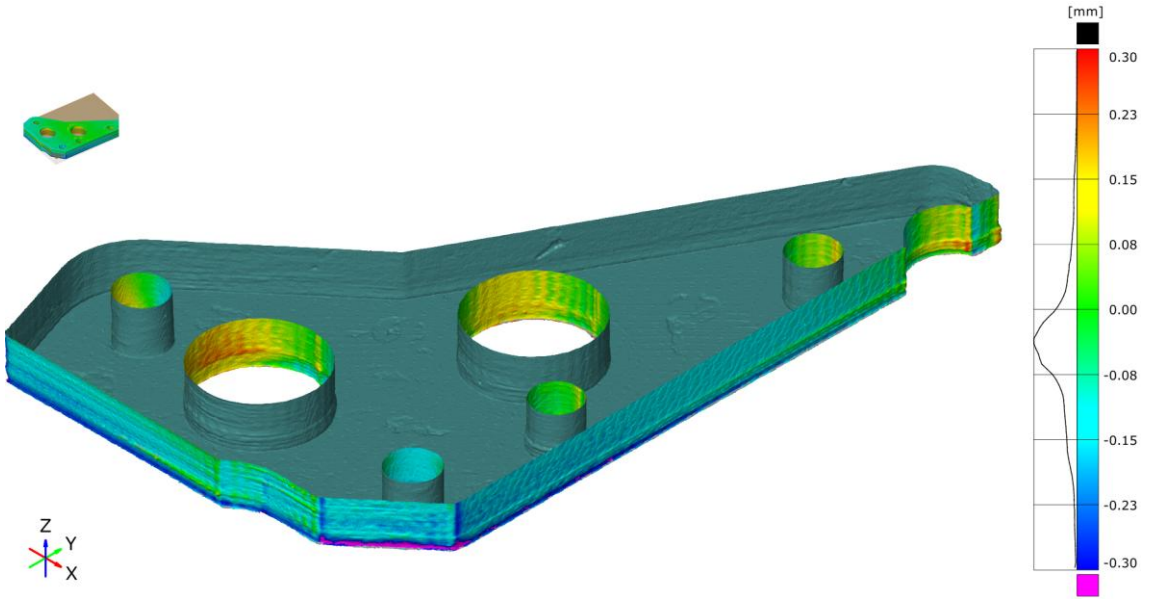

(b)

Figure 113: (a) CT deviation map of the emptied and repaired Claw against the nominal CAD model; (b) section of the emptied and repaired Claw (note the difference with Figure 112-b).

Component	Max [mm]	Min [mm]	Mean [mm]	St.D. [mm]	# of Holes
Claw					
Original	+1,50	−1,50	−0,13	0,27	1650
After HPR and Holes treatments	+1,48	−1,49	−0,05	0,09	0

Table 34: Comparison between the original Claw and its state after application of the HPR algorithm and hole closing in GOM Inspect.

Again, the standard deviation drops by a factor of three (0,27 to 0,09 [mm]) and the mean goes from −0,13 to −0,05, confirming the same mechanism observed on the Link, the Support, the Gear Arm.

Frame

Figure 114 to Figure 117 and Table 36 document the Frame.

A short note: as for the Claw, the radiographs shown here are not complete, since datos|x displays only one selected Z slice during this phase. This phase serves mainly to check whether the part moved during acquisition.

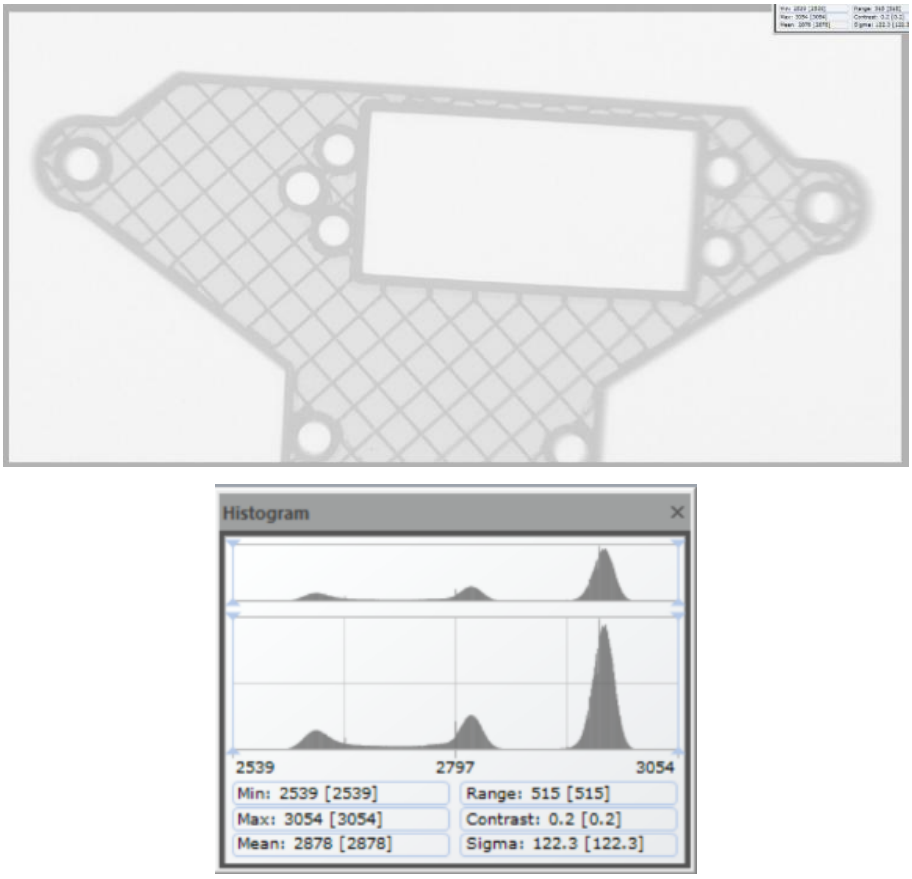

Figure 114: (top) one of the 1500 radiographs acquired for the Frame; (bottom) the corresponding three-peak grey-value histogram.

Scan information	Value / Unit
Geometry	
Magnification	3,616259
Voxel size	55,305779 [µm]
FOD	225 [mm]
Acquisition	
Number of images	1500
Resolution Images	1981 × 1000 pixels
Detector	
Type	DXR-250 RT
Timing	200 [ms]
Skip frames	4
X-Ray	
Voltage	160 [kV]
Current	160 [µA]
Filter	0,1 Cu [mm]
Focal spot size	19,2 [µm]

Table 35: Values set on the Phoenix v|tome|x s machine for the acquisition of the Frame.

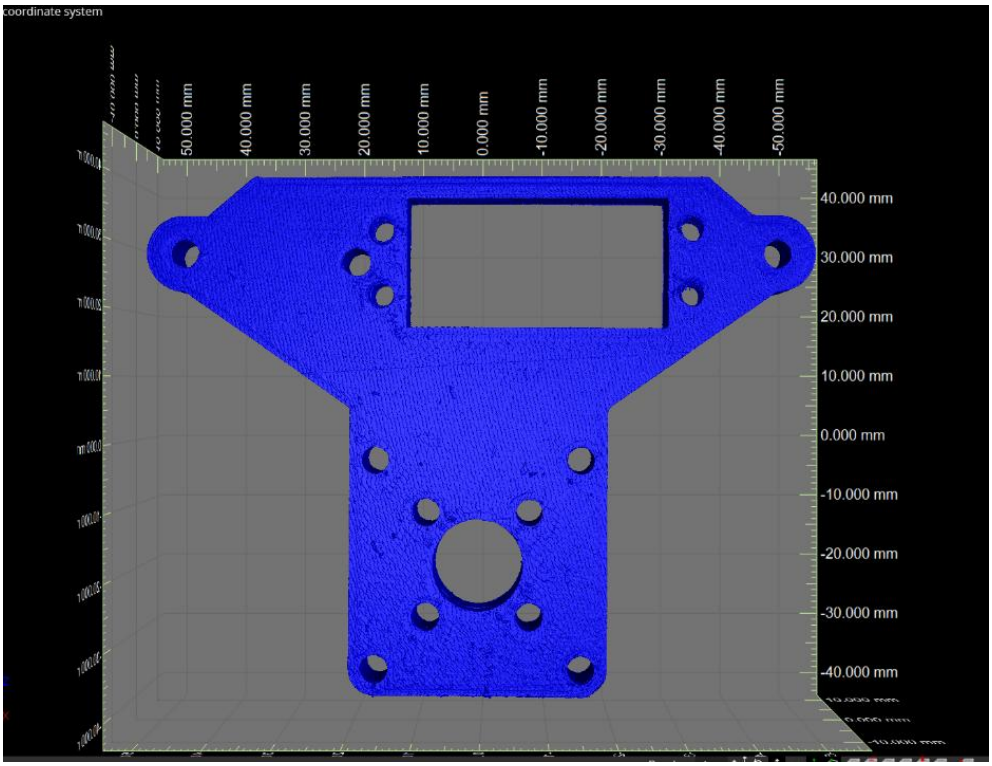

(a)

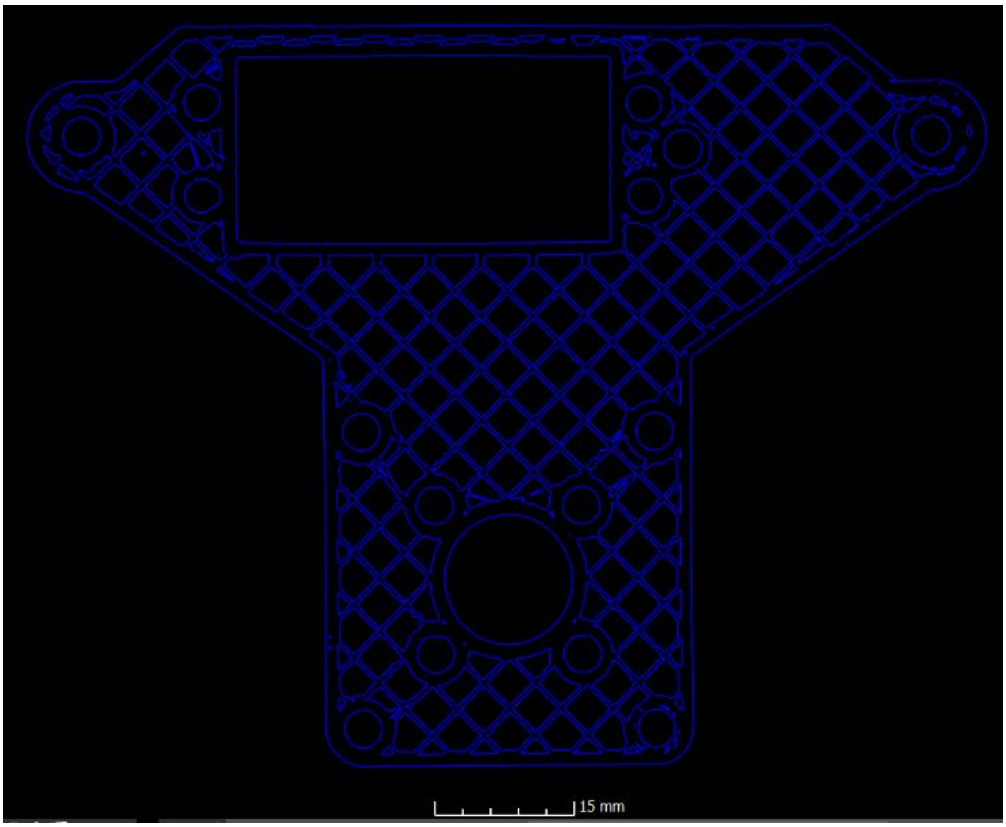

(b)

Figure 115: (a) full view of the Frame reconstructed after cleaning the grey-value histogram of the background and the low-density material; (b) view of the Frame cleaned in VG Studio and sectioned to reveal the infill structure.

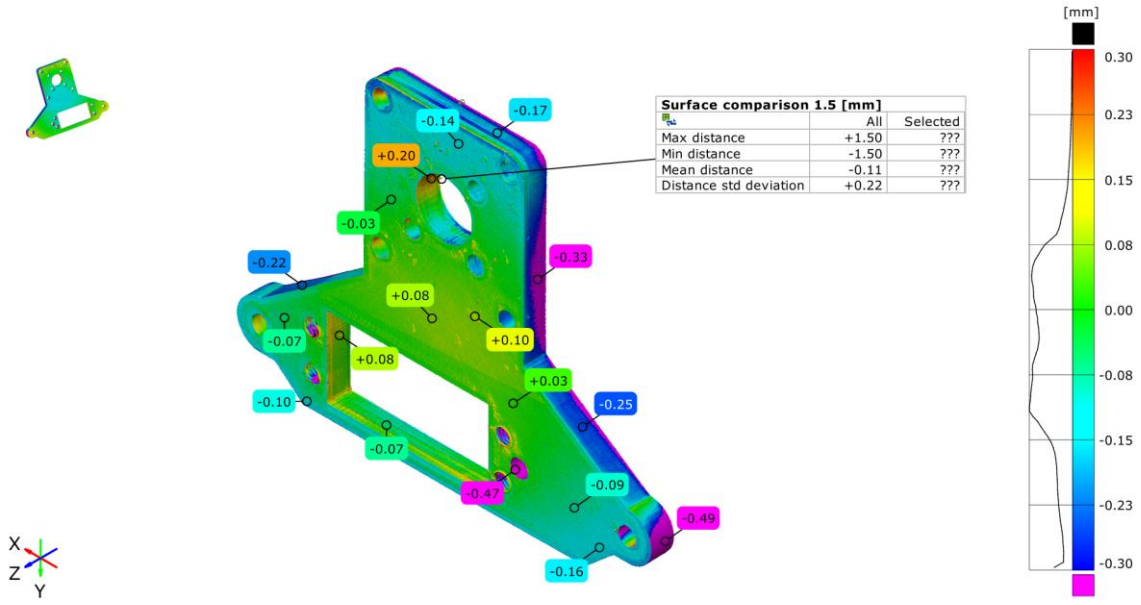
Surface comparison 1.5 [mm]
All Selected
Max distance +1.50 ???
Min distance -1.50 ???
Mean distance -0.11 ???
Distance std deviation +0.22 ???
-0.14
-0.17
+0.20
-0.03
-0.22
-0.33
+0.08
-0.07
+0.08
+0.10
+0.03
-0.10
-0.25
-0.07
-0.47
-0.09
-0.49
-0.16
X
Z
Y
[mm]
0.30
0.23
0.15
0.08
0.00
-0.08
-0.15
-0.23
-0.30

(a)

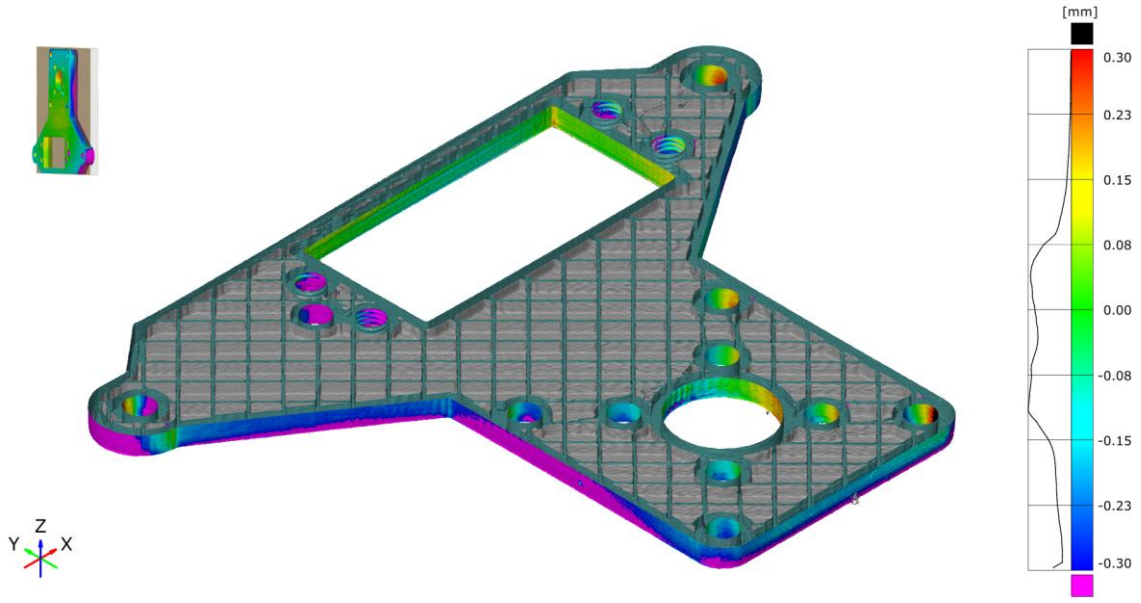
Z
Y
X
[mm]
0.30
0.23
0.15
0.08
0.00
-0.08
-0.15
-0.23
-0.30

(b)

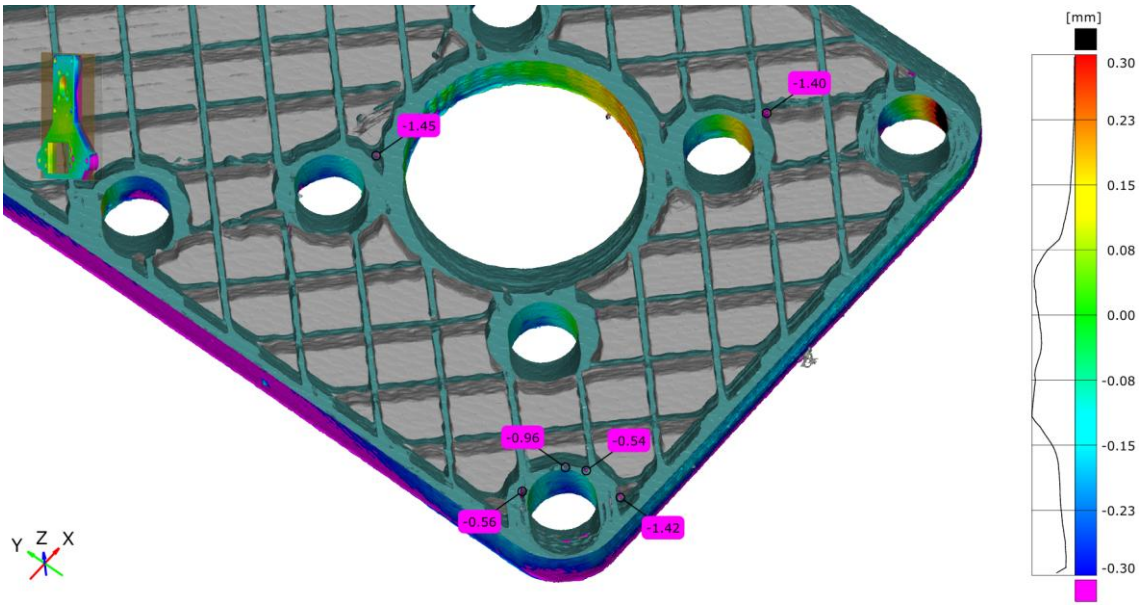

(c)

Figure 116: (a) CT deviation map of the original full Frame against the nominal CAD model; (b) section view of the original full Frame; (c) internal surfaces wrongly recognised and considered by the comparison.

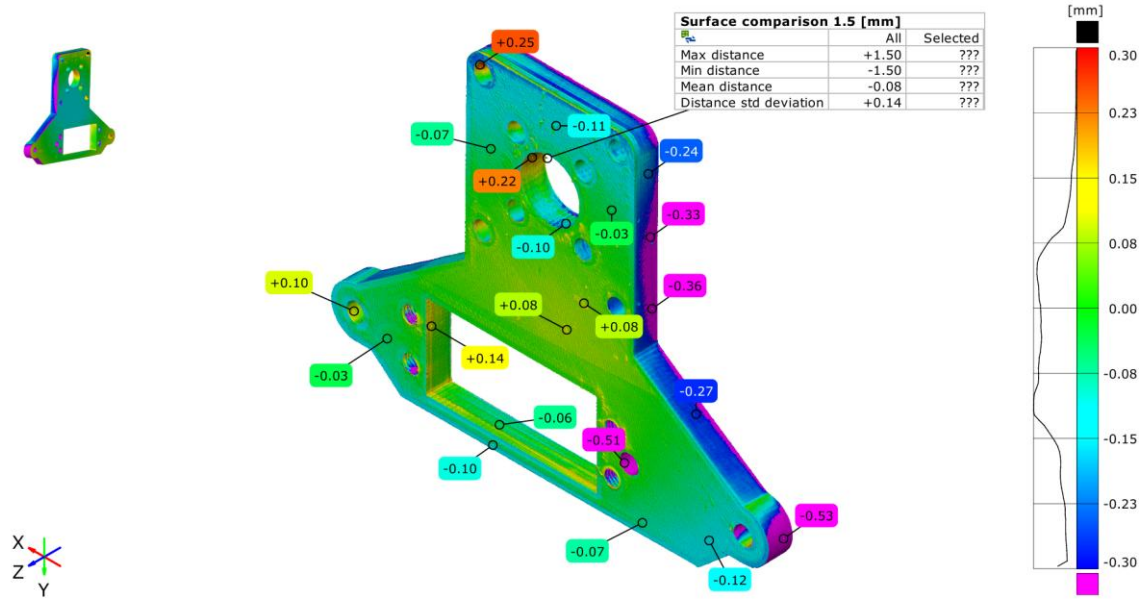

(a)

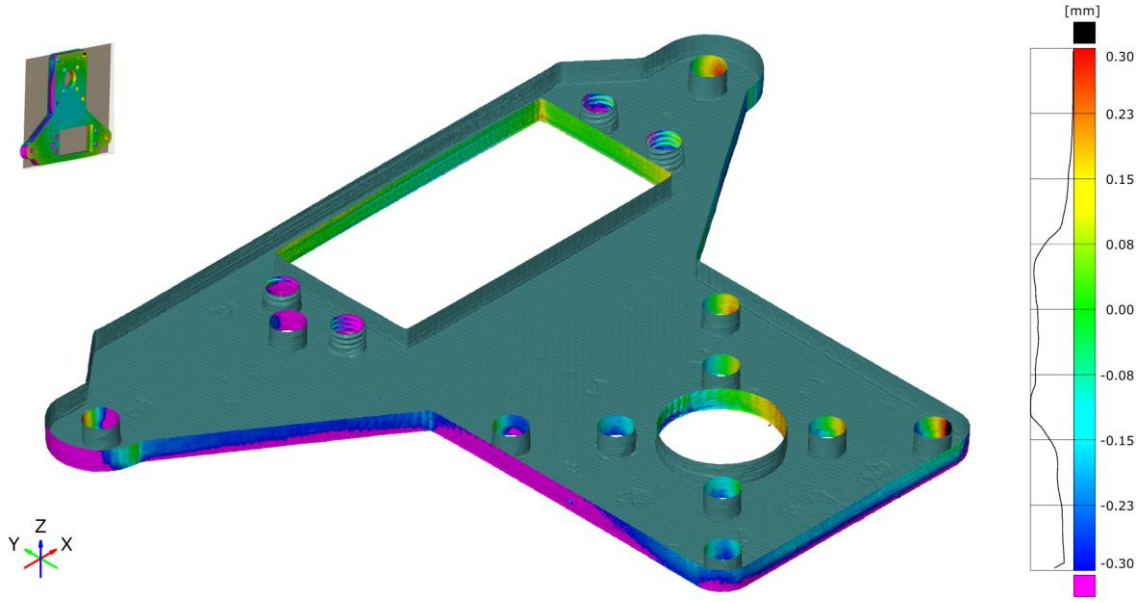

(b)

Figure 117: (a) CT deviation map of the emptied and repaired Frame against the nominal CAD model; (b) section of the emptied and repaired Frame (note the difference with Figure 116-b).

Component	Max [mm]	Min [mm]	Mean [mm]	St.D. [mm]	# of Holes
Frame					
Original	+1,5	−1,5	−0,11	0,22	2102
After HPR and Holes treatments	+1,5	−1,5	−0,08	0,14	0

Table 36: Comparison between the original Frame and its state after application of the HPR algorithm and hole closing in GOM Inspect.

The Frame confirms the trend seen on the other components, although the improvement is more contained: the mean deviation moves from −0,11 to −0,08 [mm] and the standard deviation from 0,22 to 0,14 [mm]. The residual standard deviation, higher than for the Link or the Claw, is consistent with the Frame being the largest and most geometrically complex part, with thin walls and through-windows that are intrinsically harder to reconstruct and to empty completely.

Global CT results

This final section aggregates the per-component results of the CT route, following the same equal-weighting approach adopted for the EinScan and ATOS acquisitions: each component is treated as an independent test case and contributes equally, regardless of its triangle count. The servo support is excluded from this aggregation, as stated above, to preserve the same set of four components (Gear Arm, Link, Claw, Frame) used for the other two technologies in Section 3.5. The global mean and pooled standard deviation are computed with the same formula reported in Appendix B.

The aggregation is reported in two states. The first uses the original, full meshes (with infill and internal geometry), and is included to quantify, at the global level, the bias introduced by the internal features. The second uses the post-HPR, hole-treated meshes, which represent the outer surface alone and are therefore the only ones suitable for the cross-scanner comparison.

CT Global Info	Value [mm]
Original	
Mean distance	−0,178
Distance std deviation	0,316
HPR and hole treatment	
Mean distance	−0,070
Distance std deviation	0,096

Table 37: Global CT indicators computed on the original full meshes and on the post-HPR, hole-treated meshes in equal weighting.

The comparison between the two states is the clearest possible justification of the whole HPR procedure. On the original meshes the global standard deviation is 0,316 [mm], dominated by the internal geometry and by the Link in particular (σ = 0,41 [mm]); after emptying, it falls to 0,096 [mm], more than a threefold reduction, while the global mean moves from −0,178 to −0,070 [mm]. In other words, roughly two thirds of the apparent global dispersion of the raw CT data was an artefact of the infill rather than a property of the outer surface.

E - Outer-Shell Extraction Algorithm (HPR)

The following Python script removes the internal geometry (inner shell, infill, sealed cavities) from a CT mesh, retaining only the outer-shell triangles in their original measured positions, by means of the Hidden Point Removal operator. It is reproduced here in full, with syntax highlighting for readability.

```
#!/usr/bin/env python3
# ============================================================================
#  outer_shell.py
#  ----------------------------------------------------------------------
#  Outer-shell extraction from a CT mesh of an FDM-printed part, based on
#  Hidden Point Removal (HPR). Developed for the reverse-engineering chapter
#  of this thesis to make the CT acquisitions metrologically comparable with
#  the structured-light acquisitions (EinScan, ATOS).
#
#  ----------------------------------------------------------------------
#  1. PROBLEM
#  A CT scan of an FDM component does not reconstruct only the outer surface
#  of the part: it also captures the internal geometry, namely the inner
#  shell (the perimeter wall separating the infill cavities from the outer
#  skin) and the infill structure itself, together with the small cavities
#  caused by the non-ideally homogeneous layered deposition. When such a mesh
#  is compared against the nominal CAD in GOM Inspect with a finite search
#  distance, the nearest-point algorithm also associates these interior
#  features, biasing both the mean deviation and, more severely, the standard
#  deviation. The resulting statistics no longer describe the outer surface
#  of the part and are therefore not comparable with those of the optical
#  techniques, which natively reconstruct the outer surface only.
#
#  ----------------------------------------------------------------------
#  2. WHY NOT A DISTANCE-TO-CAD CRITERION
#  A natural first attempt is to discard triangles whose distance to the CAD
#  exceeds a threshold. This conflates two physically different situations:
#  a triangle may be far from the CAD because it is interior geometry (which
#  must be discarded) or because it is genuine outer surface that deviates
#  strongly from nominal, e.g. through local shrinkage or poor deposition
#  (which must be kept). Reducing the threshold to remove the interior
#  therefore also removes high-deviation outer wall, truncating exactly the
#  tail of the distribution that the measurement is meant to capture. The
#  correct discriminating question is not "how far is this triangle from the
#  nominal geometry" but "is this triangle exposed to the outside" - an
#  intrinsic geometric property of the scanned part that requires no CAD
#  reference and does not bias the deviation statistics.
```

```
#
#  ------------------------------------------------------------------------
#  3. METHOD: HIDDEN POINT REMOVAL (HPR)
#  The Hidden Point Removal operator [1] determines, for a given viewpoint,
#  which points of a point set are directly visible (not occluded by other
#  points), using a spherical-flip transformation followed by a convex-hull
#  construction, without requiring surface normals or mesh connectivity.
#  In this script the operator is applied from a large number of viewpoints
#  uniformly distributed over a sphere enclosing the part. A vertex is
#  classified as belonging to the outer shell if it is visible from at least
#  one external viewpoint (the union of the per-viewpoint visible sets);
#  interior vertices (inner shell, infill, sealed cavities) are never visible
#  from outside and are discarded.
#
#  Because HPR reasons on point POSITIONS rather than on rays traced through
#  the mesh, it does not "leak" through the small holes of a fragmented CT
#  mesh. This is the key advantage over the two simpler strategies that were
#  tested and discarded during development: (i) direct visibility ray-casting,
#  where rays entered through the holes of the fragmented mesh and struck the
#  interior, and (ii) voxel flood-fill from the exterior, where no voxel size
#  simultaneously sealed the holes and preserved the thin outer wall.
#
#  The output mesh contains the ORIGINAL triangles in their ORIGINAL
#  positions: the procedure performs only a binary keep/discard selection.
#  No vertex is moved, smoothed or reconstructed, so the subsequent surface
#  comparison in GOM Inspect operates on genuinely measured data.
#
#  ------------------------------------------------------------------------
#  4. VALIDATION
#  On a synthetic fragmented double-shell test part containing infill and a
#  through-hole, with per-face ground-truth labels, the procedure achieved:
#     full resolution         -> 100.0 % outer-shell recall, 0.0 % interior
#                                contamination;
#     downsampled + KNN prop.  ->  99.5 % outer-shell recall, 0.1 % interior
#                                contamination.
#
#  ------------------------------------------------------------------------
#  5. PARAMETERS
#  Because the parts differ in size and morphology, the two main parameters
#  (number of viewpoints and HPR radius) were adapted per component and are
#  reported in the corresponding result section. Thin, plate-like parts in
#  particular require a larger radius, since their broad flat faces are
```

```
#  otherwise classified as occluded.
#
#     --views N            number of viewpoints distributed over the
#                          enclosing sphere (default 150). More viewpoints
#                          give more complete coverage and fewer residual
#                          holes, at a higher computational cost.
#     --vote {any,majority,all}
#                          a triangle is retained when this many of its three
#                          vertices are visible (default 'majority' = 2/3).
#                          'all' (3/3) is stricter and reduces interior
#                          leakage at the cost of more boundary holes; 'any'
#                          (1/3) is the most permissive and is preferable on
#                          heavily fragmented meshes, where outer and interior
#                          triangles rarely share a vertex.
#     --radius-factor R    the HPR radius is set to (bounding-box diagonal)*R
#                          (default 100). A larger radius makes more points
#                          visible and bridges larger surface holes; it is
#                          required for thin plate-like components.
#     --cam-factor C       viewpoint distance from the part centre, expressed
#                          as (bounding-box diagonal)*C (default 3).
#     --max-points N       maximum number of points fed to HPR (default
#                          1,200,000). Above this the unique-vertex cloud is
#                          randomly downsampled, the visibility test is run on
#                          the subset, and the labels are propagated to all
#                          vertices by nearest neighbour (validated above).
#                          Set above the number of unique vertices to disable
#                          downsampling entirely.
#     --min-cluster N      after selection, connected fragments smaller than N
#                          triangles are discarded (default 0 = keep all). For
#                          a metrological comparison 0 is recommended, so that
#                          no genuine surface region is removed.
#     -o OUT               output path (default <input>_outer.stl).
#
#  ------------------------------------------------------------------------
#  6. DEPENDENCIES / EXECUTION
#  Requires Python 3.11 with numpy, scipy and open3d. Using the uv runner,
#  no manual installation is needed:
#
#     uv run --python 3.11 --with open3d --with numpy --with scipy \
#         outer_shell.py CT_part.stl --views 300 --vote any --radius-factor 3000
#
#  Output: <input>_outer.stl, to be imported into GOM Inspect, best-fit
```

```
#  aligned onto the nominal CAD, and surface-compared with the same maximum
#  search distance used for the other scanning technologies.
#
#  ------------------------------------------------------------------------
#  7. REFERENCE
#  [1] S. Katz, A. Tal and R. Basri, "Direct Visibility of Point Sets",
#      ACM Transactions on Graphics (Proc. SIGGRAPH), vol. 26, no. 3, 2007.
# ============================================================================

import argparse, sys
import numpy as np
import open3d as o3d
from scipy.spatial import cKDTree


def log(m):
    """Print a status message with an identifying prefix, so that each run
    can be followed step by step and reported reproducibly."""
    print(f"[outer] {m}", flush=True)


def fib_sphere(n):
    """
    Return n approximately uniform directions on the unit sphere, generated
    with the Fibonacci (golden-spiral) distribution.

    A naive latitude/longitude grid would cluster points at the poles; the
    Fibonacci construction instead spaces the points uniformly in surface
    area. The polar angle is distributed uniformly in cos(phi) (equal-area
    bands, by Archimedes' theorem), while the azimuth advances by the golden
    angle so that successive points fall into the largest remaining gaps.
    This guarantees gap-free, repeatable coverage for a given n.
    """
    i = np.arange(0, n, dtype=float) + 0.5
    phi = np.arccos(1 - 2 * i / n)                # polar angle (equal area)
    th = np.pi * (1 + 5 ** 0.5) * i               # azimuth (golden angle)
    return np.column_stack((np.cos(th) * np.sin(phi),
                            np.sin(th) * np.sin(phi),
                            np.cos(phi)))


def visible_vertices_hpr(Vu, views, radius_factor, cam_factor, max_points):
```

```
    """
    Determine which of the unique vertices Vu belong to the outer shell.

    Returns a boolean mask over Vu, True where a vertex is visible from at
    least one external viewpoint (union of the per-viewpoint HPR results).
    For very large clouds the visibility test is performed on a random subset
    and the resulting labels are propagated to the full set by nearest
    neighbour, which was verified to give essentially identical results.
    """
    # Geometry of the viewpoint sphere, derived from the part bounding box.
    center = (Vu.max(0) + Vu.min(0)) / 2.0
    diag = float(np.linalg.norm(Vu.max(0) - Vu.min(0)))
    radius = diag * radius_factor                  # HPR spherical-flip radius
    cam_dist = diag * cam_factor                   # viewpoint distance
    dirs = fib_sphere(views)

    # Optional downsampling for tractability on very large clouds.
    n = len(Vu)
    if n > max_points:
        sel = np.random.choice(n, max_points, replace=False)
        Ps = Vu[sel]
        log(f"  HPR on {max_points:,} sampled points (of {n:,}); "
            f"visibility will be propagated by nearest neighbour")
    else:
        Ps = Vu
        log(f"  HPR on all {n:,} unique vertices")

    pcd = o3d.geometry.PointCloud()
    pcd.points = o3d.utility.Vector3dVector(Ps)

    # Accumulate visibility over all viewpoints (logical OR).
    vis_s = np.zeros(len(Ps), dtype=bool)
    for k, d in enumerate(dirs):
        cam = center + d * cam_dist                # viewpoint position
        try:
            # hidden_point_removal returns the indices of the visible points
            # for this viewpoint, given the spherical-flip radius.
            _, idx = pcd.hidden_point_removal(cam.tolist(), radius)
            if len(idx):
                vis_s[np.asarray(idx, dtype=np.int64)] = True
        except Exception:
            # For a few degenerate viewpoints the convex-hull step may fail;
```

```
            # such viewpoints are simply skipped (the union over the many
            # remaining viewpoints is unaffected in practice).
            pass
        if (k + 1) % 50 == 0:
            log(f"    {k+1}/{views} viewpoints, visible so far "
                f"{vis_s.sum():,}/{len(Ps):,}")

    if len(Ps) == n:
        return vis_s

    # Propagate the sampled labels to every unique vertex via a single
    # vectorised nearest-neighbour query (fast even for millions of points).
    log("  propagating visibility to the full vertex set (batch KNN) ...")
    tree = cKDTree(Ps)
    _, nn = tree.query(Vu, k=1)
    return vis_s[nn]


def main():
    # ------------------------- command-line interface ---------------------
    ap = argparse.ArgumentParser()
    ap.add_argument("stl", help="CT mesh of the printed part (STL)")
    ap.add_argument("--views", type=int, default=150)
    ap.add_argument("--vote", choices=["any", "majority", "all"],
                    default="majority")
    ap.add_argument("--radius-factor", type=float, default=100.0)
    ap.add_argument("--cam-factor", type=float, default=3.0)
    ap.add_argument("--max-points", type=int, default=1_200_000)
    ap.add_argument("--min-cluster", type=int, default=0)
    ap.add_argument("-o", "--out", default=None)
    args = ap.parse_args()

    # ------------------------------ loading --------------------------------
    log(f"loading {args.stl} ...")
    m = o3d.io.read_triangle_mesh(args.stl)
    if len(m.vertices) == 0:
        sys.exit("ERROR: mesh empty")
    V = np.asarray(m.vertices)
    F = np.asarray(m.triangles)
    log(f"  {len(V):,} verts, {len(F):,} tris")
    log(f"  bbox (mm): {np.round(V.max(0) - V.min(0), 2)}")
```

```
# The STL format stores each vertex once per incident face, so the same
# physical point appears repeated. The visibility test is run on the set
# of UNIQUE positions; 'inverse' maps every original vertex back to its
# unique position, so the per-vertex labels can be expanded afterwards.
key = np.round(V, 4)
Vu, inverse = np.unique(key, axis=0, return_inverse=True)
log(f"  {len(Vu):,} unique vertex positions")

# ------------------------ visibility (HPR) -----------------------------
log(f"running HIDDEN POINT REMOVAL ({args.views} viewpoints, "
    f"vote='{args.vote}') ...")
vis_u = visible_vertices_hpr(Vu, args.views, args.radius_factor,
                             args.cam_factor, args.max_points)
vis_v = vis_u[inverse]                          # label per original vertex
log(f"  visible vertices: {vis_v.sum():,}/{len(V):,}")

# -------------------------- triangle vote ------------------------------
# Count how many of each triangle's three vertices are visible (0..3),
# then keep the triangle according to the chosen voting rule.
vcount = vis_v[F].sum(axis=1)
if args.vote == "any":
    keep = vcount >= 1
elif args.vote == "all":
    keep = vcount == 3
else:                                           # majority
    keep = vcount >= 2
log(f"  kept {keep.sum():,}/{len(F):,} "
    f"({100*keep.sum()/len(F):.1f}%); dropped {(~keep).sum():,} interior")

# Rebuild the output mesh from the ORIGINAL vertex positions and the
# selected subset of triangles (no geometry is modified).
out_mesh = o3d.geometry.TriangleMesh(
    o3d.utility.Vector3dVector(V),
    o3d.utility.Vector3iVector(F[keep]))
out_mesh.remove_unreferenced_vertices()

# --------------------- optional fragment removal ------------------------
# Remove small isolated connected components if requested. Kept disabled
# by default for metrology, since discarding fragments could selectively
# remove the more fragmented surface regions and bias the comparison.
if args.min_cluster > 0:
    lbl, counts, _ = out_mesh.cluster_connected_triangles()
```

```
        lbl = np.asarray(lbl); counts = np.asarray(counts)
        small = counts[lbl] < args.min_cluster
        out_mesh.remove_triangles_by_mask(small)
        out_mesh.remove_unreferenced_vertices()
        log(f"  removed {int(small.sum()):,} tris in fragments "
            f"< {args.min_cluster}")

    # ------------------------------ export ---------------------------------
    out_mesh.compute_vertex_normals(); out_mesh.compute_triangle_normals()
    out = args.out or args.stl.rsplit('.', 1)[0] + "_outer.stl"
    o3d.io.write_triangle_mesh(out, out_mesh)
    log(f"DONE -> {out}")
    log(f"  {len(out_mesh.triangles):,} outer-shell triangles "
        f"(positions UNCHANGED). No CAD used.")
    log("Import into GOM, best-fit onto CAD, surface-compare with the same")
    log("max search distance used for the other scanning technologies.")


if __name__ == "__main__":
    main()
```

F - HSV and Filters Algorithm

The script below implements the HSV and morphology segmentation introduced in Section 4.2. The pipeline stops at the morphological mask, which is the image carried into the SolidWorks reconstruction of the infill. The threshold bounds and the morphology settings are exposed through trackbars and tuned interactively on the cross-section.

```
"""
Interactive HSV and morphology segmentation.

This tool isolates the infill walls in a CT cross-section. Through trackbars,
the operator tunes in real time:

  1. an HSV colour threshold, which produces a binary mask of the candidate
     wall pixels;
  2. a single morphological stage (erosion and/or dilation). The order can be
     switched between an opening (erosion then dilation) and a closing
     (dilation then erosion).

The pipeline stops at the morphological mask, which is the image used for the
infill reconstruction described in Chapter 4. Press "s" to save the masks and
the current parameters, and "Esc" to quit.
"""

import cv2
import numpy as np
import os
from datetime import datetime

# ------------------------------------------------------------ configuration
IMAGE_PATH = "infill_section.png"   # CT cross-section to be segmented
SAVE_DIR = "save"                   # output folder for masks and parameters
KERNEL = np.ones((3, 3), np.uint8)  # 3 × 3 structuring element

os.makedirs(SAVE_DIR, exist_ok=True)


def nothing(x):
    """Empty callback required by the OpenCV trackbar API."""
    pass


# ------------------------------------------------------------------ windows
```

```
cv2.namedWindow("HSV", cv2.WINDOW_NORMAL)          # HSV threshold trackbars
cv2.namedWindow("MORPH", cv2.WINDOW_NORMAL)        # morphology trackbars
cv2.namedWindow("MASK_HSV", cv2.WINDOW_NORMAL)     # mask after HSV threshold
cv2.namedWindow("MASK_MORPH", cv2.WINDOW_NORMAL)   # mask after morphology
cv2.namedWindow("OUTPUT", cv2.WINDOW_NORMAL)       # segmented colour image

# ------------------------------------------------------------- trackbars
# HSV bounds. Hue on the 0 to 179 OpenCV scale; S and V on 0 to 255.
cv2.createTrackbar("H_min", "HSV", 0, 179, nothing)
cv2.createTrackbar("H_max", "HSV", 179, 179, nothing)
cv2.createTrackbar("S_min", "HSV", 0, 255, nothing)
cv2.createTrackbar("S_max", "HSV", 255, 255, nothing)
cv2.createTrackbar("V_min", "HSV", 0, 255, nothing)
cv2.createTrackbar("V_max", "HSV", 255, 255, nothing)

# Morphology controls. Each operation can be enabled and given an iteration
# count; ORDER selects opening (0) or closing (1).
cv2.createTrackbar("ERODE_ON", "MORPH", 0, 1, nothing)
cv2.createTrackbar("ERODE_IT", "MORPH", 1, 50, nothing)
cv2.createTrackbar("DILATE_ON", "MORPH", 0, 1, nothing)
cv2.createTrackbar("DILATE_IT", "MORPH", 1, 50, nothing)
cv2.createTrackbar("ORDER", "MORPH", 0, 1, nothing)


def apply_morph(mask):
    """Apply erosion and/or dilation to a binary mask.

    ORDER = 0 gives an opening (erosion then dilation).
    ORDER = 1 gives a closing (dilation then erosion).
    """
    mask = mask.copy()
    order = cv2.getTrackbarPos("ORDER", "MORPH")
    erode_it = cv2.getTrackbarPos("ERODE_IT", "MORPH")
    dilate_it = cv2.getTrackbarPos("DILATE_IT", "MORPH")

    if order == 0:
        if cv2.getTrackbarPos("ERODE_ON", "MORPH"):
            mask = cv2.erode(mask, KERNEL, iterations=erode_it)
        if cv2.getTrackbarPos("DILATE_ON", "MORPH"):
            mask = cv2.dilate(mask, KERNEL, iterations=dilate_it)
    else:
```

```
        if cv2.getTrackbarPos("DILATE_ON", "MORPH"):
            mask = cv2.dilate(mask, KERNEL, iterations=dilate_it)
        if cv2.getTrackbarPos("ERODE_ON", "MORPH"):
            mask = cv2.erode(mask, KERNEL, iterations=erode_it)
    return mask


# ------------------------------------------------------------ load the image
img_original = cv2.imread(IMAGE_PATH)
if img_original is None:
    raise FileNotFoundError(f"Could not read '{IMAGE_PATH}'.")

# ------------------------------------------------------------- main loop
while True:
    img = img_original.copy()

    # Read the HSV bounds from the trackbars.
    h_min = cv2.getTrackbarPos("H_min", "HSV")
    h_max = cv2.getTrackbarPos("H_max", "HSV")
    s_min = cv2.getTrackbarPos("S_min", "HSV")
    s_max = cv2.getTrackbarPos("S_max", "HSV")
    v_min = cv2.getTrackbarPos("V_min", "HSV")
    v_max = cv2.getTrackbarPos("V_max", "HSV")

    # Stage 1: HSV threshold, giving the binary mask of candidate wall pixels.
    hsv = cv2.cvtColor(img, cv2.COLOR_BGR2HSV)
    lower = np.array([h_min, s_min, v_min])
    upper = np.array([h_max, s_max, v_max])
    mask_hsv = cv2.inRange(hsv, lower, upper)

    # Stage 2: single morphological stage (opening or closing).
    mask_morph = apply_morph(mask_hsv)

    # Segmented colour image, kept only as visual feedback.
    output = cv2.bitwise_and(img, img, mask=mask_morph)

    # Display the intermediate and final results.
    cv2.imshow("MASK_HSV", mask_hsv)
    cv2.imshow("MASK_MORPH", mask_morph)
    cv2.imshow("OUTPUT", output)
```

```
    key = cv2.waitKey(30) & 0xFF

    # Press "s" to save the masks and the current parameters.
    if key == ord("s"):
        ts = datetime.now().strftime("%Y%m%d_%H%M%S")
        cv2.imwrite(f"{SAVE_DIR}/mask_hsv_{ts}.png", mask_hsv)
        cv2.imwrite(f"{SAVE_DIR}/mask_morph_{ts}.png", mask_morph)
        cv2.imwrite(f"{SAVE_DIR}/output_{ts}.png", output)

        with open(f"{SAVE_DIR}/params_{ts}.txt", "w") as f:
            f.write(f"HSV: {h_min},{h_max},{s_min},{s_max},{v_min},{v_max}\n")
            for p in ["ERODE_ON", "ERODE_IT", "DILATE_ON", "DILATE_IT", "ORDER"]:
                f.write(f"{p}: {cv2.getTrackbarPos(p, 'MORPH')}\n")
        print(f"Data saved in the '{SAVE_DIR}' folder")

    # Press "Esc" to quit.
    if key == 27:
        break

cv2.destroyAllWindows()
```

G - FEM Servo Stall forces

This appendix derives the complete set of forces and moments generated by the servo stall condition on the two structural components of the end-effector, the Frame and the Claw. These actions constitute the servo stall load case applied to the FEM models of Chapter 5 and summarised there in Table 12. As anticipated in Section 5.1, the most convenient strategy is to traverse the kinematic chain body by body: the static equilibrium of each component is written in turn, and the joint reactions are propagated backwards from the gripping interface to the Frame.

Assumptions and modelling principles. The derivation rests on the following assumptions:

- **Gripped object.** The object held between the two Claws is treated as rigid and incompressible. Being clamped from both sides, it is in equilibrium and, by Newton's third law, transmits two equal and opposite forces F to the right and left Claw.
- **Gripping configuration.** The object width of 27,5 [mm] sets the clear separation between the two Claw pads and therefore defines the configuration of maximum extension of the mechanism along the global Z axis; this usable width is smaller than the roughly 38 [mm] measured between the two extremes of the gear arm, because the Claw bodies and the offset between the pivot A and the contact pad take up part of the opening. All geometric quantities used below (l, m, n, z, g_r, g_l) are measured from the CAD model in this configuration.
- **Joint convention.** Every revolute pair connecting two bodies is represented by two coincident points: a point on the first body (e.g. B on the Claw) and its counterpart on the second (B′ on the Link). Newton's third law imposes that the force vectors at the two points balance, $\mathbf{F_B + F_{B'} = 0}$: every force acting at B is accompanied by an equal and opposite force at B′.
- **Planar reduction.** Since the mechanism is symmetric with respect to its mid plane and all revolute axes are parallel to the global X axis, the assembly is collapsed onto a single plane and analysed as a 2D structure. The planar axes are denoted x and y: the planar x axis corresponds to the global Y axis and the planar y axis to the global Z axis of the FEM reference system of Section 5.1, while moments about the revolute axes act about the global X axis. The planar reduction serves only for the derivation of the loads; the FEM itself is performed on the full 3D geometry.
- **Static condition.** At stall the mechanism is motionless, so each body satisfies pure static equilibrium. All joints are modelled as frictionless revolute pairs, each with a single degree of freedom (DOF).
- **Gear mesh.** The tooth contact at the pitch point P is described by a tangential force F_t and a radial force F_r related through the nominal pressure angle of the gear pair introduced in Chapter 2, $\mathbf{F_r = F_t \cdot \tan \alpha}$ with $\boldsymbol{\alpha = 20°}$.
- **Symmetry.** Owing to the mirror symmetry of the mechanism, the right and left Claw and Link pairs are statically identical up to a sign inversion of the x components. The Claw and Link equilibria are therefore derived once, and the mirrored forces are introduced directly in the equilibrium of the driven gear arm.

The analysis starts from the tip of the kinematic chain and proceeds backwards: Claw, Link, slave gear arm and finally the servo driven gear arm, where the servo torque closes the problem.

Claw. The free-body diagram is shown in Figure 118; the same situation is mirrored on the opposite Claw. The Claw is connected to the gear arm at the revolute pair A and to the Link at

the revolute pair B, and receives the gripping force F from the object. The pin reactions F_{xA}, F_{yA}, F_{xB} and F_{yB} are drawn, purely as a convention, in the positive direction of the planar axes. With the Claw immobile at stall, equilibrium reads:

$$\Sigma F_x: \quad F + F_{xA} + F_{xB} = 0$$
$$\Sigma F_y: \quad F_{yA} + F_{yB} = 0 \quad \rightarrow \quad F_{yA} = -F_{yB}$$
$$\Sigma M(A): \quad F \cdot l + F_{yB} \cdot m + F_{xB} \cdot n = 0$$

These are three equations in four unknowns; the system is closed by the equilibrium of the Link, which fixes F_{xB}. The sign of the n term follows the offset between A and B; as shown below, the term vanishes in any case. By the joint convention, the counterpoints A′ (on the gear arm) and B′ (on the Link) satisfy $\mathbf{F}_A + \mathbf{F}_{A'} = \mathbf{0}$ and $\mathbf{F}_B + \mathbf{F}_{B'} = \mathbf{0}$.

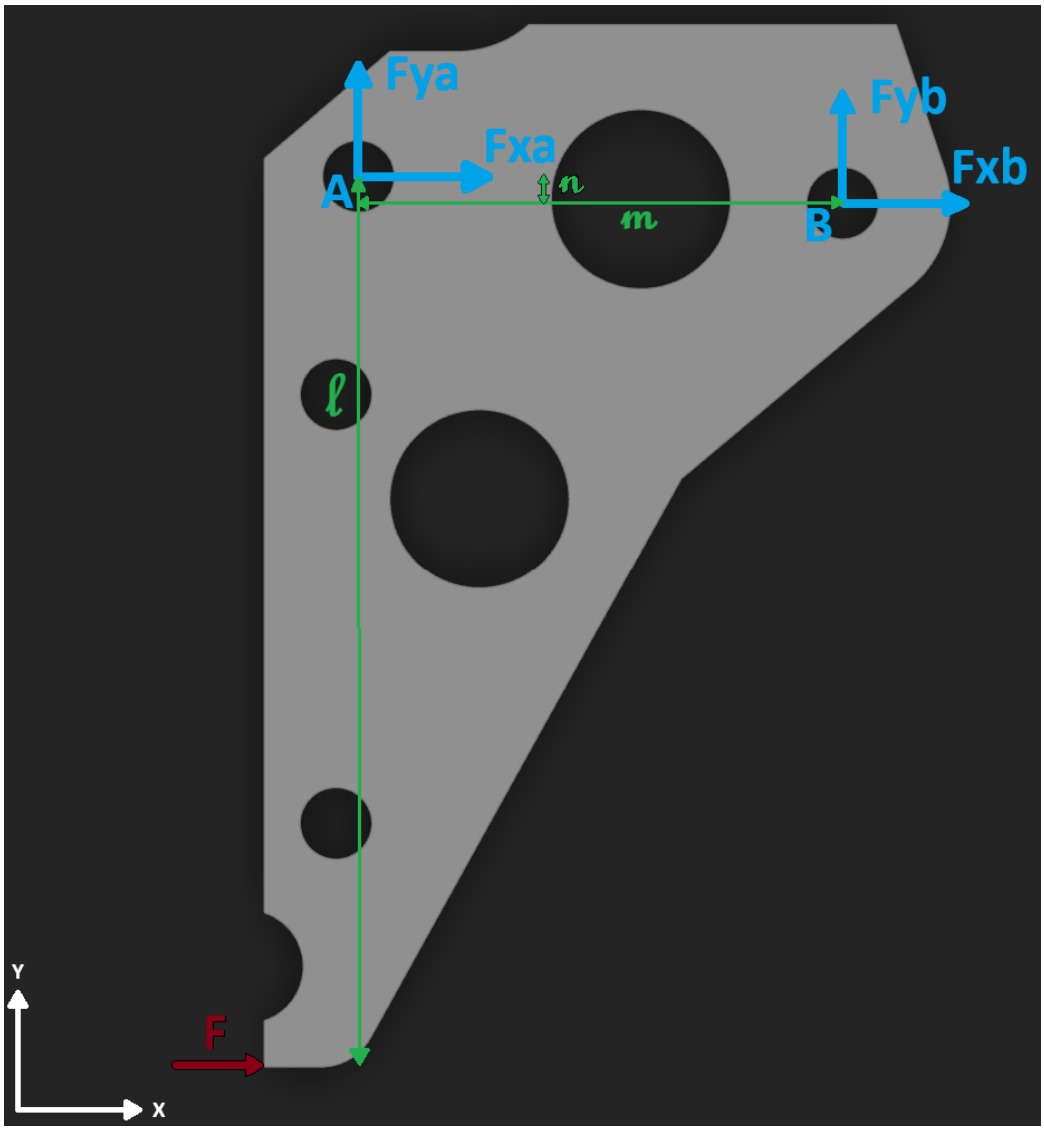

Figure 118: Free-body diagram of the Claw in the planar reference system. F is the gripping force transmitted by the object; A and B are the revolute pairs connecting the Claw to the gear arm and to the Link; l, m and n are the lever arms used in the moment equilibrium about A.

Link. The free-body diagram is shown in Figure 119; the same situation is mirrored on the opposite Link. The Link connects the Claw (B′) to the Frame (C) through two revolute pairs and carries no other load: it is therefore a two-force member. In the configuration under analysis the holes B′ and C are aligned along the planar y axis at a distance z, so the moment equilibrium about C directly eliminates the x component:

$$\Sigma F_x: \quad F_{xB'} + F_{xC} = 0 \quad \rightarrow \quad F_{xB'} = -F_{xC}$$
$$\Sigma F_y: \quad F_{yB'} + F_{yC} = 0 \quad \rightarrow \quad F_{yB'} = -F_{yC}$$
$$\Sigma M(C): \quad F_{xB'} \cdot z = 0 \quad \rightarrow \quad F_{xB'} = F_{xC} = 0$$

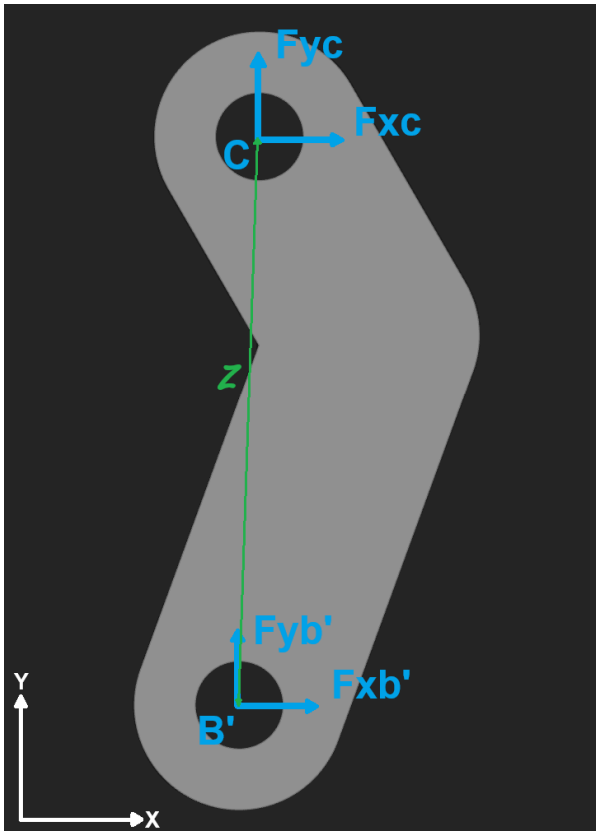

Figure 119: Free-body diagram of the Link. The two revolute pairs B' (to the Claw) and C (to the Frame) are aligned along the planar y axis at a distance z; with no other applied load the Link behaves as a two-force member.

Back substitution to the Claw. Propagating $\mathbf{F}_{xB} = -\mathbf{F}_{xB'} = \mathbf{0}$ into the Claw equilibrium yields the closed-form solution of the Claw and Link subchain as a function of the still unknown gripping force F:

$$F \cdot l + F_{yB} \cdot m + 0 \cdot n = 0 \quad \rightarrow \quad F_{yB} = -F_{yA} = -F \cdot l/m$$
$$F + F_{xA} + 0 = 0 \quad \rightarrow \quad F_{xA} = -F$$
$$F_{yC} = -F_{yB'} = F_{yB} = -F \cdot l/m$$

Through the joint convention, the forces acting on the gear arm at A′ and on the Frame at C′ follow immediately:

$$F_{xA'} = F \qquad F_{yA'} = -F \cdot l/m$$
$$F_C + F_{C'} = 0 \quad \rightarrow \quad F_{xC'} = 0, \quad F_{yC'} = +F \cdot l/m$$

Slave gear arm. The free-body diagram is shown in Figure 120. The slave gear arm rotates about the revolute pair D (the slave gear axis on the Frame) and is loaded by the mesh forces F_t and F_r at the pitch contact point P, at a distance g_r from D, and by the Claw reaction at A′, at a distance g_l from D:

$$\Sigma F_x: \quad F_{xD} + F_r + F_{xA'} = 0 \quad \rightarrow \quad F_{xD} = -F_r - F$$
$$\Sigma F_y: \quad F_{yD} + F_t + F_{yA'} = 0 \quad \rightarrow \quad F_{yD} = -F_t + F \cdot l/m$$
$$\Sigma M(D): \quad F_{xA'} \cdot g_l - F_t \cdot g_r = 0 \quad \rightarrow \quad F_t = F \cdot g_l / g_r$$

The pin reaction is transferred to the Frame at the slave gear hole through the joint convention, $\mathbf{F}_D + \mathbf{F}_{D'} = \mathbf{0}$.

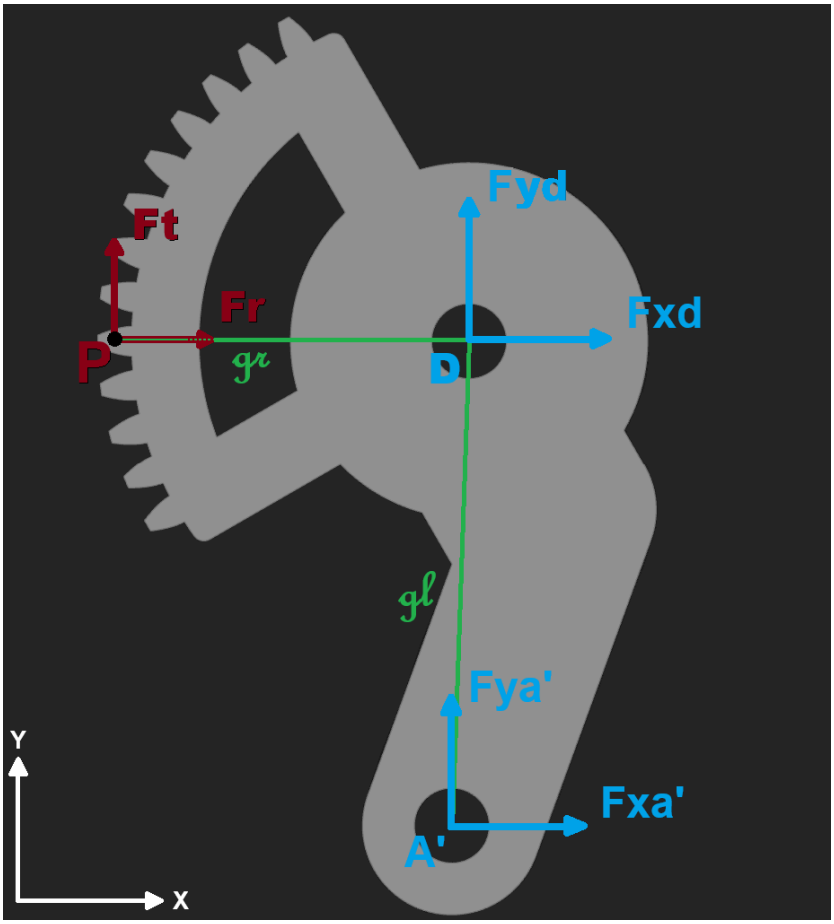

Figure 120: Free-body diagram of the slave gear arm. The mesh forces Ft and Fr act at the pitch contact point P at a distance gr from the axis D; the Claw reaction acts at A' at a distance gl from D.

Driven gear arm. The free-body diagram is shown in Figure 121. The servo driven gear arm rotates about the revolute pair E (the servo output axis) and is loaded by the servo stall torque M_s, by the mesh reaction at the pitch point (equal and opposite to the forces applied to the slave gear, by Newton's third law) and, exploiting the symmetry of the mechanism, by the mirrored Claw reaction at its own arm hole. The pin reaction is transferred to the Frame through $\mathbf{F}_E + \mathbf{F}_{E'} = \mathbf{0}$. Equilibrium reads:

$$\Sigma F_x: \quad F_{xE} - F_r - F = 0 \quad \rightarrow \quad F_{xE} = F_r + F$$

$$\Sigma F_y: \quad F_{yE} - F_t - F \cdot l/m = 0 \quad \rightarrow \quad F_{yE} = F_t + F \cdot l/m$$

$$\Sigma M(E): \quad M_s - F \cdot g_l - F_t \cdot g_r = 0$$

Substituting the mesh relation $\mathbf{F}_t = \mathbf{F} \cdot \mathbf{g}_l / \mathbf{g}_r$ found above, the servo torque closes the problem:

$$M_s = F \cdot g_l + F_t \cdot g_r = 2 \cdot F \cdot g_l \quad \rightarrow \quad F = M_s/(2 \cdot g_l)$$

The gripping force applied to the object therefore depends exclusively on the servo torque and on the length g_l of the link connecting the gear arm to the Claw. The relation also implies $F_t = M_s/(2 \cdot g_r)$: only half of the servo torque crosses the gear mesh, the other half being routed directly to the Claw on the driven side through its own gear arm. A single mesh estimate $F_t = M_s/g_r$ would therefore overestimate the tooth load by a factor of two.

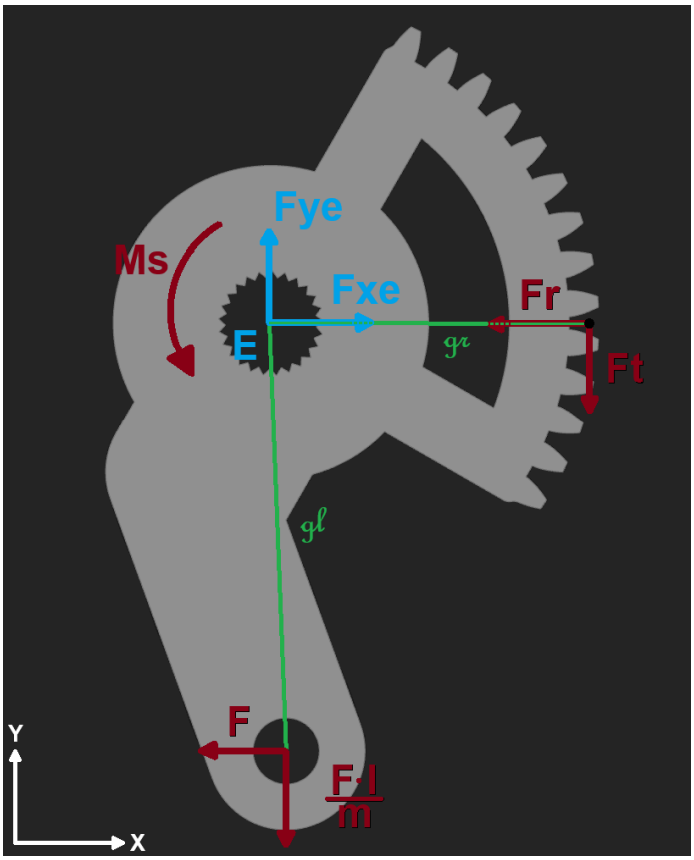

Figure 121: Free-body diagram of the driven gear arm. The servo torque Ms acts about the axis E together with the mesh reaction at the pitch point and the mirrored Claw reaction at the arm hole.

Numerical evaluation. The numerical values are obtained by introducing the maximum nominal servo torque and the geometric quantities measured from the CAD model in the gripping configuration, collected in Table 38.

Variable	Value / Unit
Ms (max nominal servo torque at 6 [V], Chapter 2)	402,0 [N·mm]
gr (gear pitch radius to the contact point P)	18,94 [mm]
gl (gear link, distance between gear axis and Claw hole)	25,92 [mm]
l (distance along y between hole A and the line of action of F)	49,78 [mm]
m (distance along x between holes A and B of the Claw)	27,00 [mm]
α (nominal pressure angle of the gear pair)	20 [°]
Object width (Claw separation in the analysed configuration)	27,5 [mm]

Table 38: Input data for the numerical evaluation of the servo stall load set. Lengths measured from the CAD model in the gripping configuration.

Substituting these values into the closed-form relations derived above yields, with results rounded to three decimal figures:

$$F = M_s/(2 \cdot g_l) = 402{,}0/(2 \times 25{,}92) = 7{,}755 \text{ [N]}$$

$$F \cdot l/m = 7{,}755 \times 49{,}78/27{,}00 = 14{,}298 \text{ [N]}$$

$$F_t = M_s/(2 \cdot g_r) = 402{,}0/(2 \times 18{,}94) = 10{,}612 \text{ [N]}$$

$$F_r = F_t \cdot \tan \alpha = 10{,}612 \times \tan 20° = 3{,}863 \text{ [N]}$$

from which the pin reactions on the two gear arms follow:

$$F_{xD} = -3{,}863 - 7{,}755 = -11{,}618 \text{ [N]} \qquad F_{yD} = -10{,}612 + 14{,}298 = +3{,}686 \text{ [N]}$$

$$F_{xE} = 3{,}863 + 7{,}755 = +11{,}618 \text{ [N]} \qquad F_{yE} = 10{,}612 + 14{,}298 = +24{,}910 \text{ [N]}$$

Resulting load set on the Frame. Negating the pin reactions through the joint convention and mapping the planar axes back to the global reference system of Section 5.1 (planar x → global Y, planar y → global Z, moments about the revolute axes → global X), the complete load set acting on the Frame is obtained, as collected in Table 39. The servo reaction torque $-M_s$ acts on the Frame through the four servo mounting holes.

Point on the Frame	Component	Value / Unit
C′ = C″ (Link holes, one per side)	FZ	+14,298 [N]
D′ (slave gear axis)	FY	+11,618 [N]
	FZ	−3,686 [N]
E′ (servo driven gear axis)	FY	−11,618 [N]
	FZ	−24,910 [N]
	TX	−402,0 [N·mm]

Table 39: Servo stall load set transmitted to the Frame, expressed in the global reference system of Section 5.1. The set coincides with Table 12 of Chapter 5.

As a consistency check, the set is self-equilibrated: $\boldsymbol{\Sigma F_Z = 2 \times 14{,}298 - 3{,}686 - 24{,}910 = 0\ [N]}$ and $\boldsymbol{\Sigma F_Y = 11{,}618 - 11{,}618 = 0\ [N]}$, confirming that the applied loads represent an internal equilibrium state of the mechanism.

For the Claw analysis of Section 5.1 the relevant load set is the gripping force F = 7,755 [N] at the contact pad together with the pin reactions at A ($F_x = -7{,}755$ [N], $F_y = +14{,}298$ [N]) and at B ($F_y = -14{,}298$ [N]), likewise self-equilibrated in forces and in moments about A.

Limitations. The derivation assumes rigid bodies, frictionless revolute pairs and a single mesh contact at the pitch point; the lever arms are measured in one specific configuration (maximum extension at gripping) and the planar reduction neglects any out-of-plane action. These simplifications are consistent with, and on the conservative side of, the rigid connector modelling philosophy adopted for the FEM in Section 5.1.

H - FEM SF & Displacement Results

To keep the main discussion compact, the full safety-factor and displacement fields of the nominal CAD and reconstructed (CT) parts are collected here, for each component and for each of the three load cases. Of all the information these maps contain, two indicators matter for the present verification: the minimum factor of safety, which measures how close the part comes to yielding, and the maximum displacement, which measures its global stiffness. Both are reported in the text accompanying each pair of maps and are consolidated, together with the CAD-versus-CT comparison, in the summary tables of Sections 5.2 and 5.3. The safety-factor maps locate the most stressed region of the part, while the displacement contours give the magnitude of the total displacement in [mm], with the peak value called out on the part.

Frame - nominal CAD

Gravity

Under self-weight the Frame is very lightly loaded: the safety-factor distribution stays everywhere far above the yield limit, with a minimum of 181,4 confined to a small region around the mounting holes, while the displacement is negligible, peaking at $8{,}655 \times 10^{-3}$ [mm]. Gravity alone therefore has no structural relevance for this component.

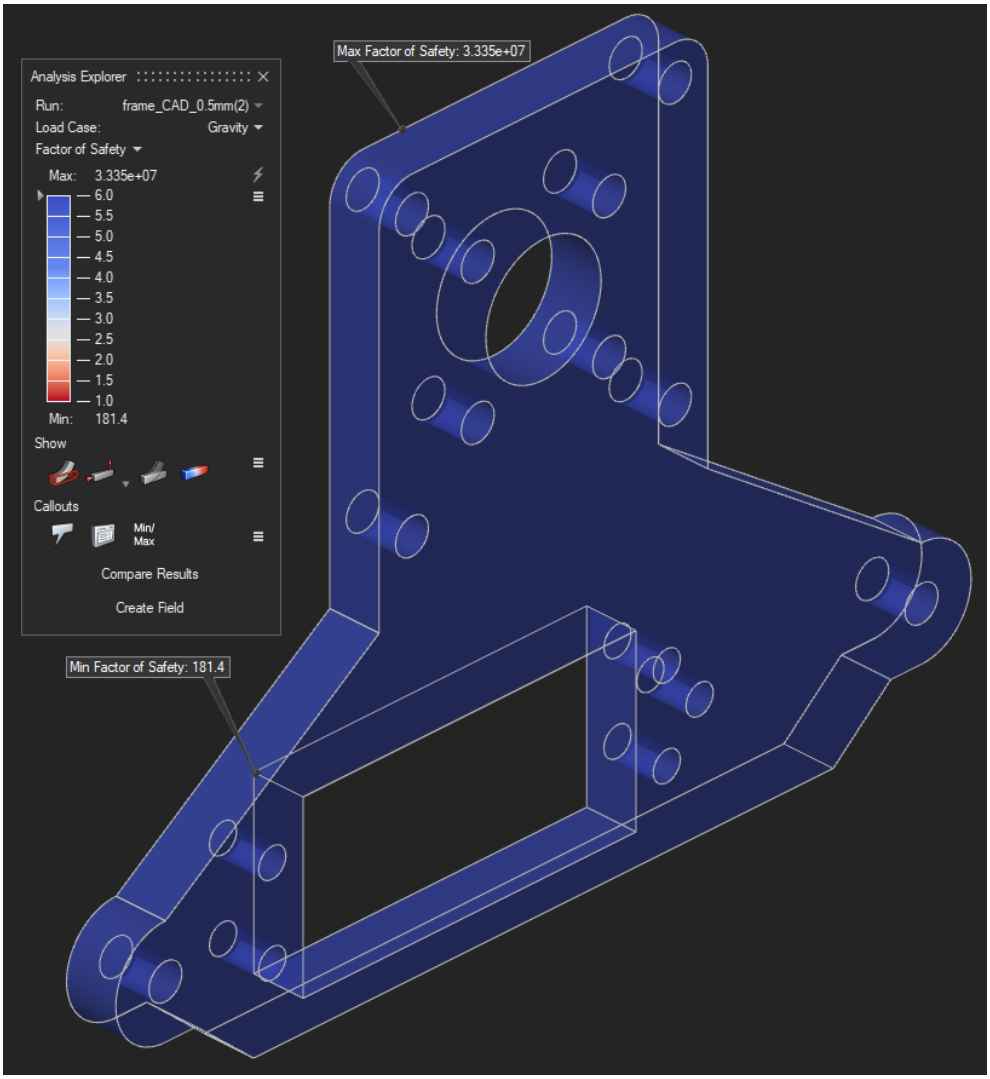

Figure 122: Factor-of-safety distribution on the nominal CAD Frame under the gravity load case.

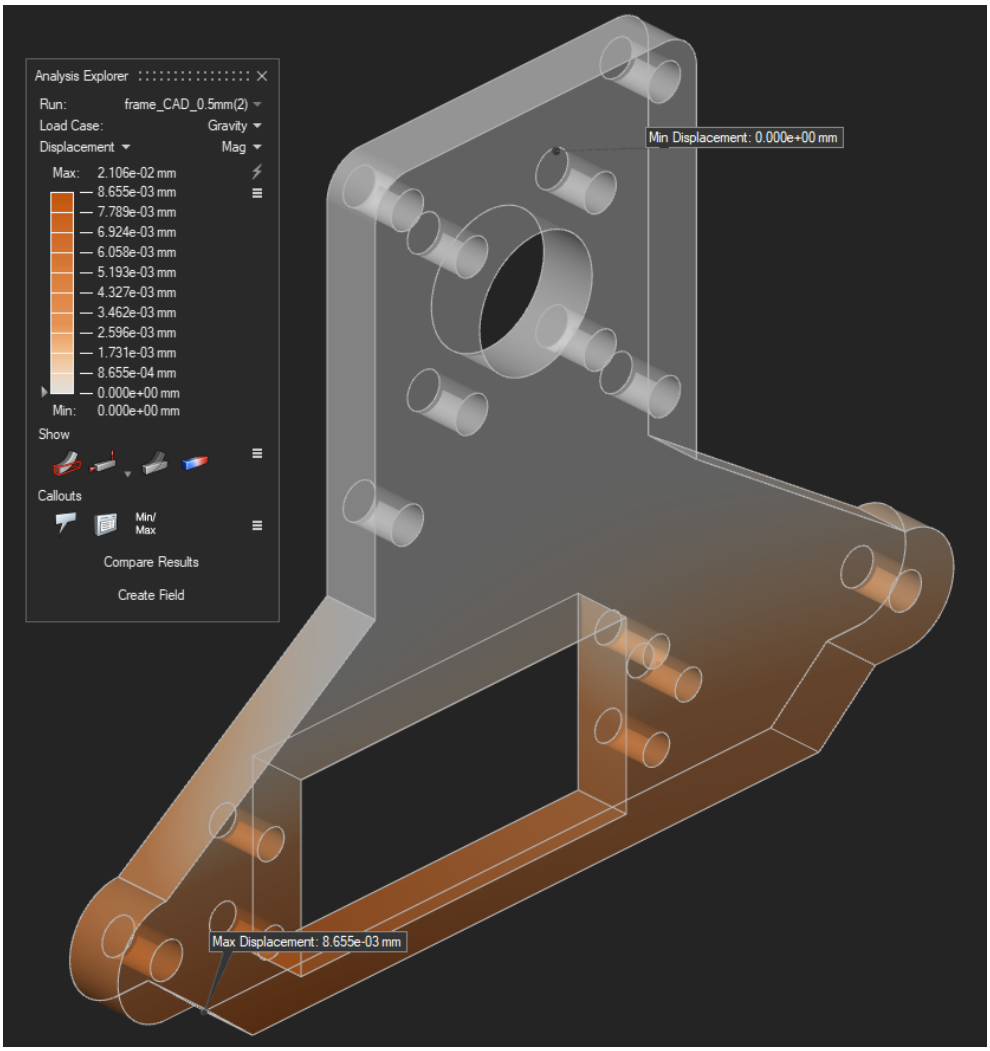

Figure 123: Displacement-magnitude contour of the nominal CAD Frame under the gravity load case; peak displacement 8,655 × 10⁻³ [mm].

Acceleration along X

The acceleration along X is by far the most demanding case for the Frame. The safety factor drops to a minimum of 4,3, localised where the bracket meets the supports and the bending moment from the offset masses is greatest, so this is the value that sizes the component. The displacement field also peaks here, at 6,780 × 10^{-1} [mm] on the free end – an order of magnitude above the other two cases, yet still small in absolute terms.

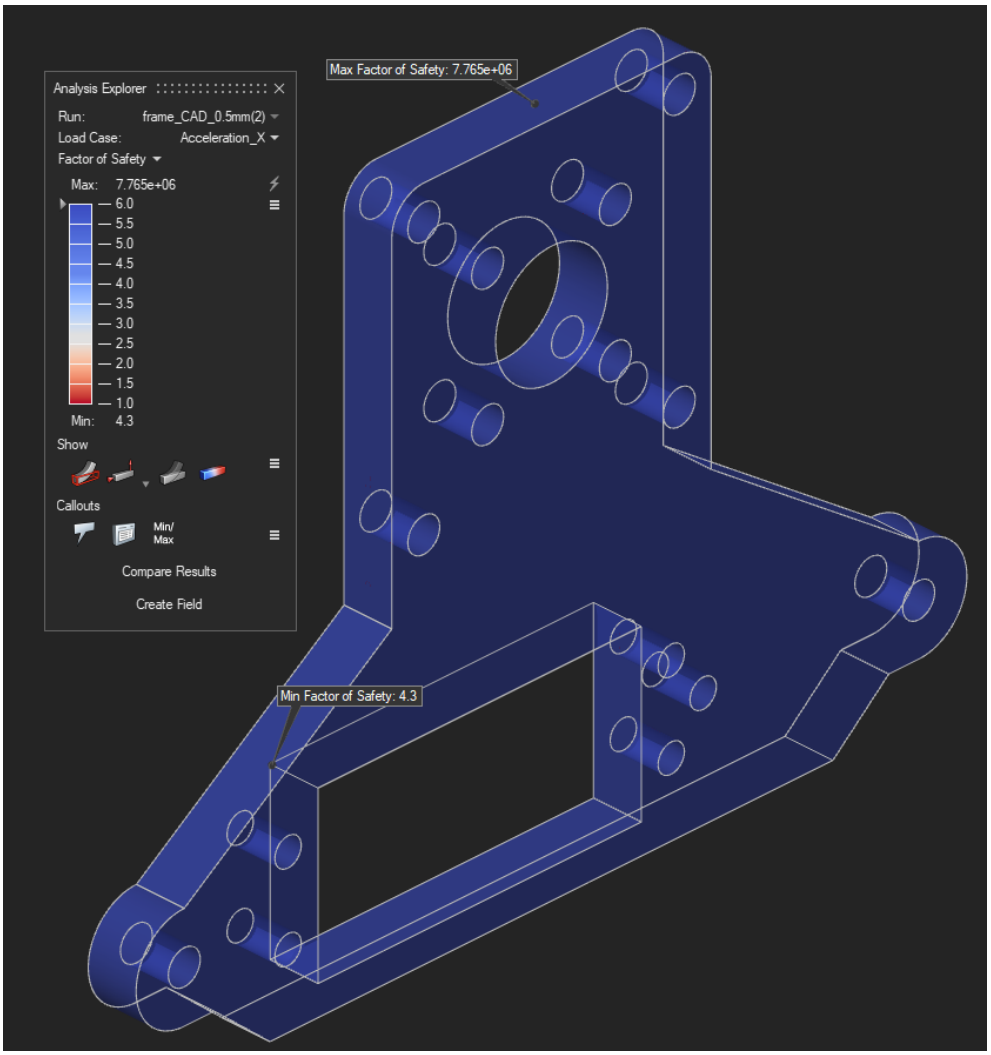

Figure 124: Factor-of-safety distribution on the nominal CAD Frame under the acceleration-along-X load case.

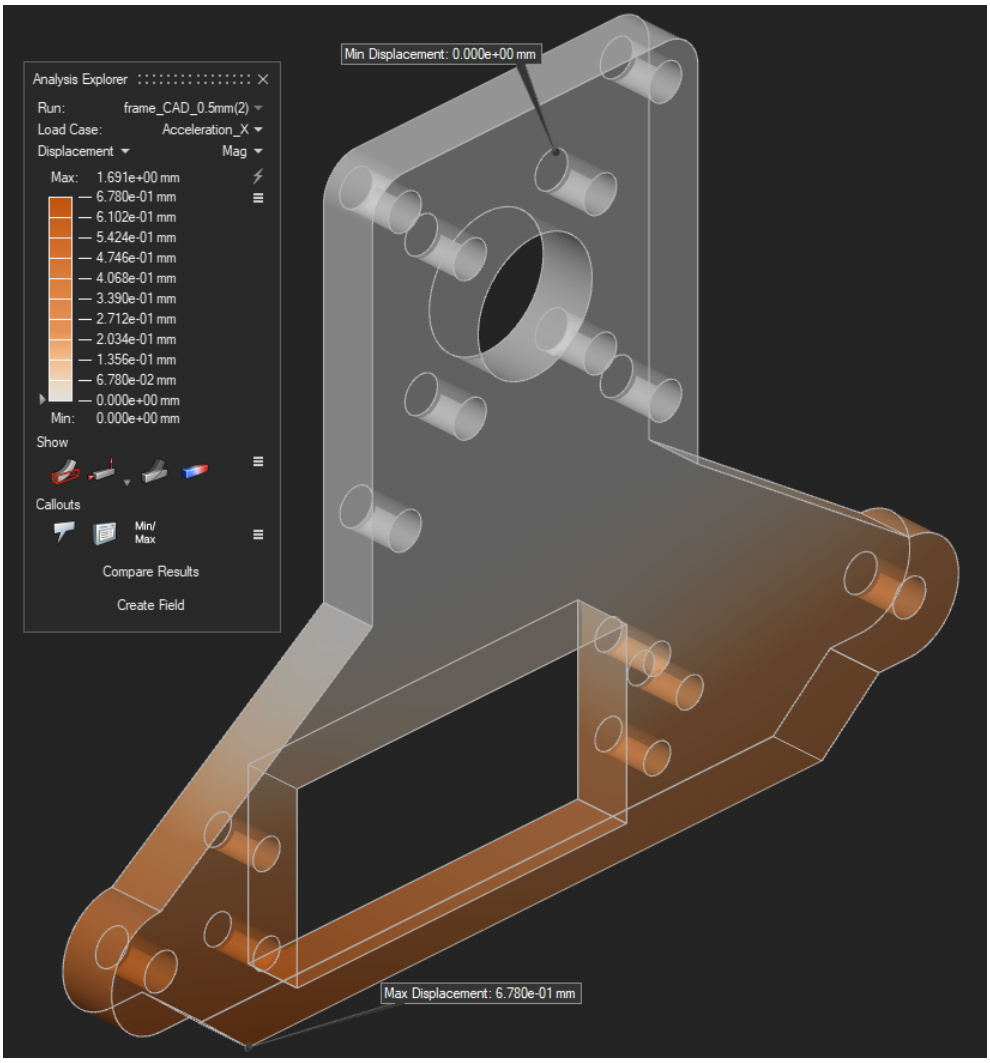

Figure 125: Displacement-magnitude contour of the nominal CAD Frame under the acceleration-along-X load case; peak displacement 0,678 [mm], at the free end of the bracket.

Servo stall (gripping)

In the servo-stall case the self-equilibrated gear and link forces load the Frame internally. The minimum safety factor of 11,6 appears near the gear-axis seats, where the reaction torque concentrates, and the peak displacement remains modest at 2,133 × 10^{-1} [mm]. The part is safe under the stall loads.

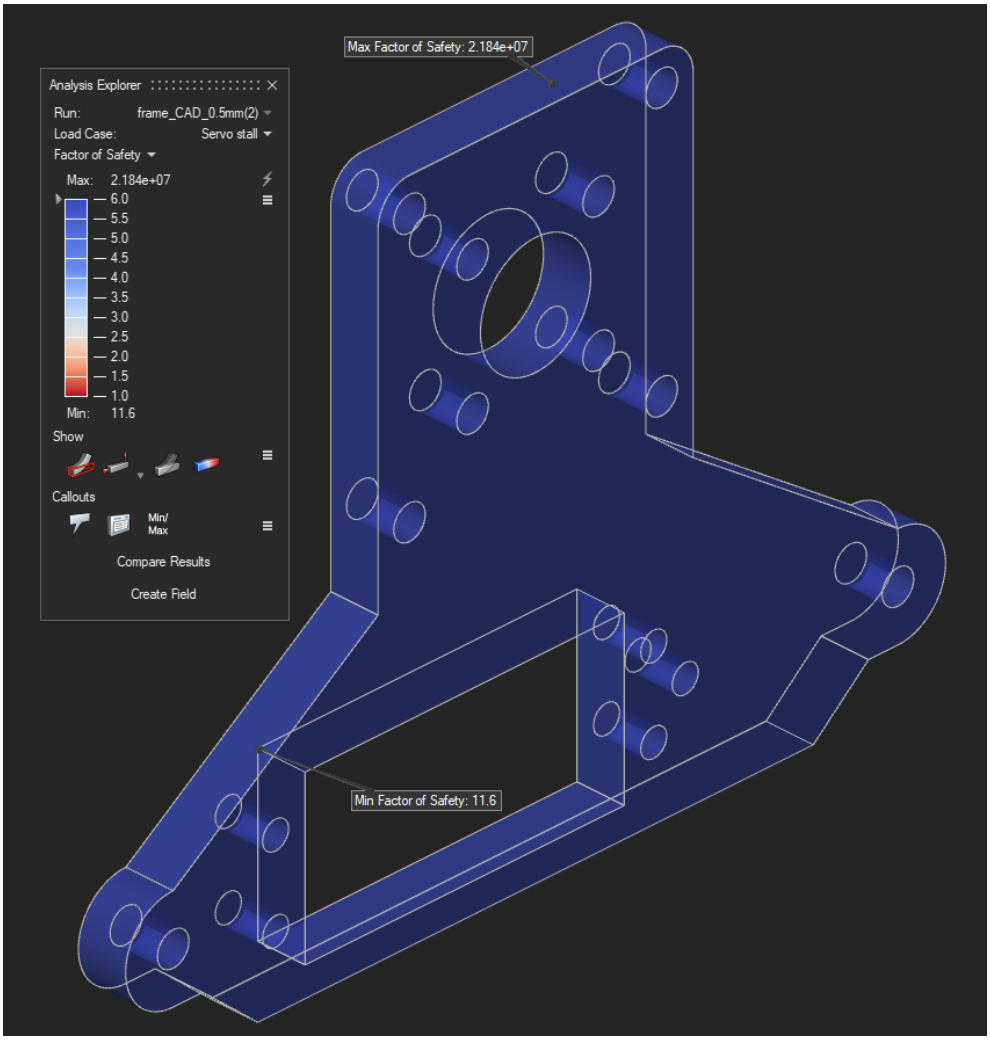

Figure 126: Factor-of-safety distribution on the nominal CAD Frame under the servo-stall load case.

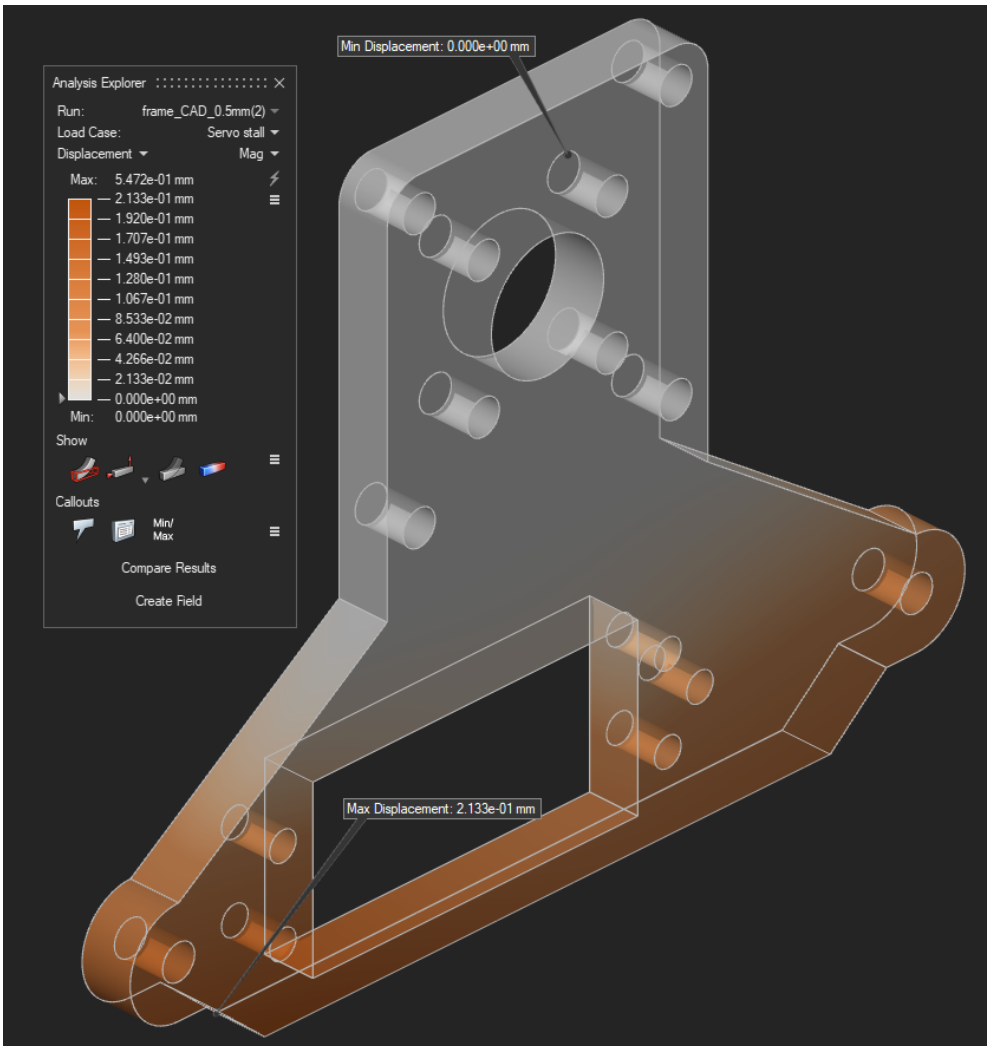

Figure 127: Displacement-magnitude contour of the nominal CAD Frame under the servo-stall load case; peak displacement 0,213 [mm].

Claw - nominal CAD

Gravity

Gravity is entirely non-critical for the solid Claw: the safety factor never falls below 271,1 and the displacement is negligible, with a peak of 2,928 × 10^{-3} [mm]. The finger behaves as a rigid body under its own weight and the small share of the object it carries.

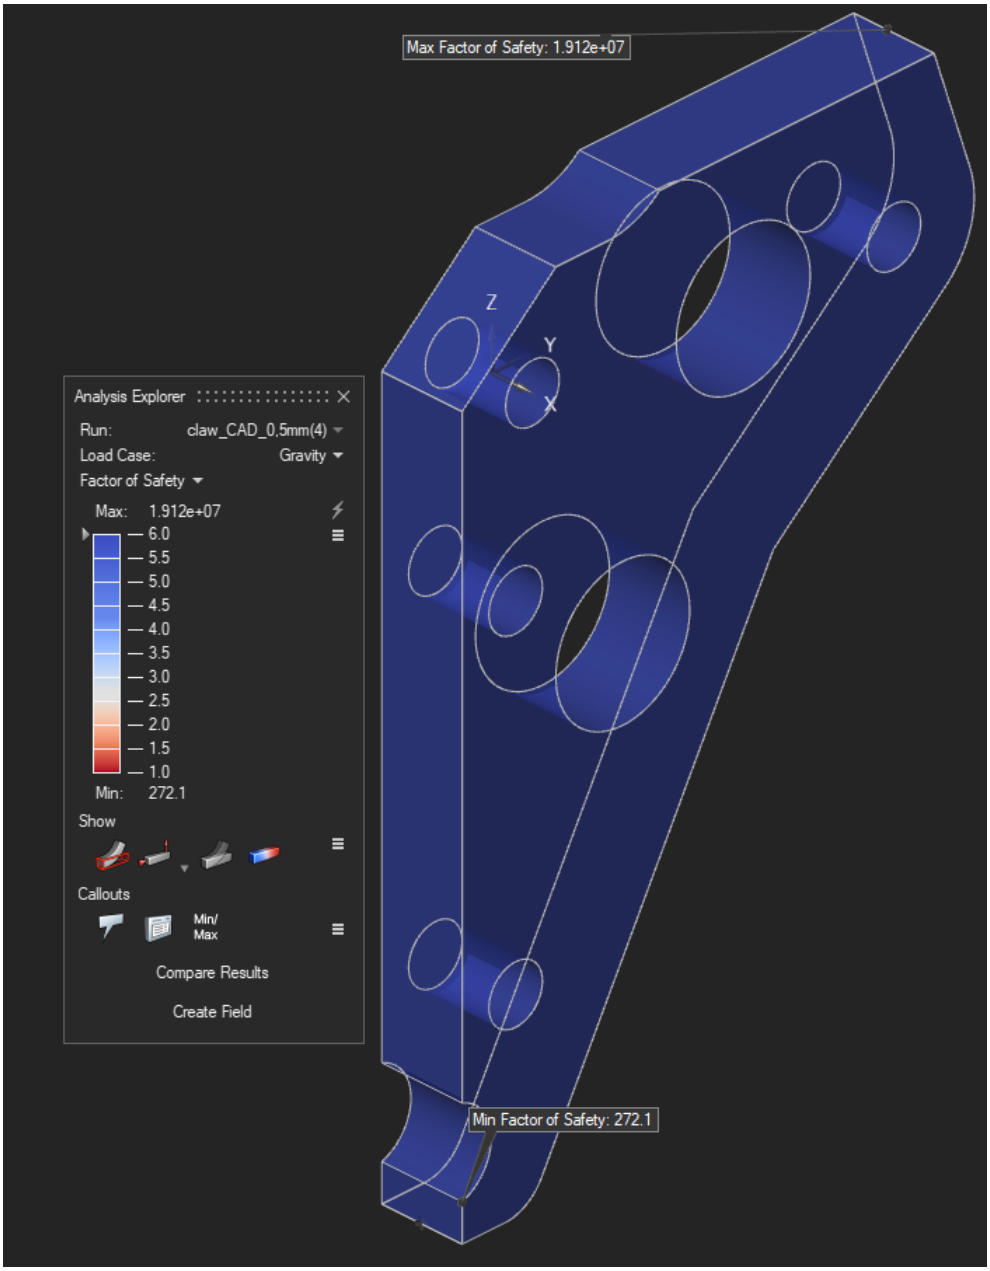

Figure 128: Factor-of-safety distribution on the nominal CAD Claw under the gravity load case.

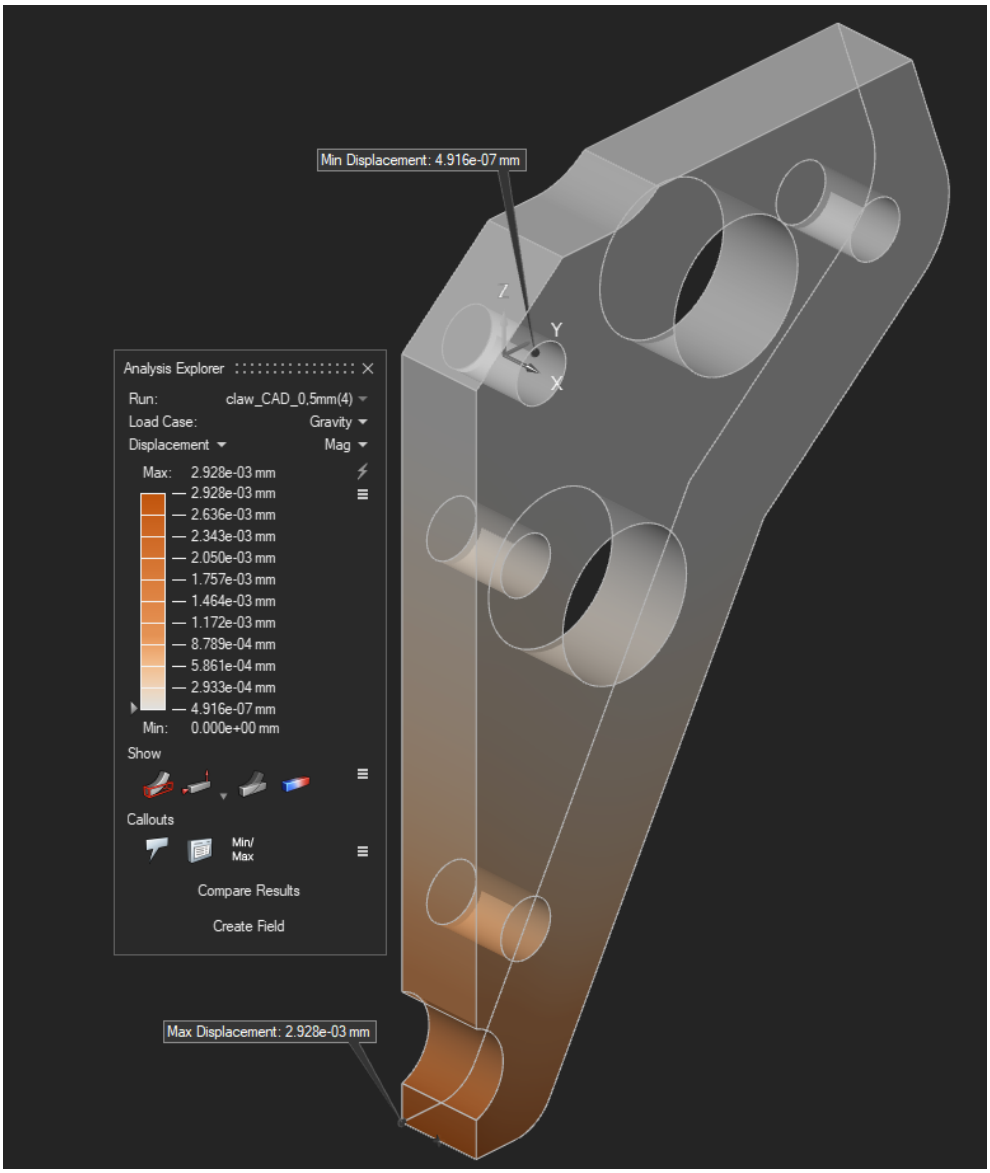

Figure 129: Displacement-magnitude contour of the nominal CAD Claw under the gravity load case; peak displacement 2,928 × 10⁻³ [mm].

Acceleration along X

Under the X acceleration the solid Claw reaches its lowest safety factor, 26,6, at the root of the finger near pin A, where the inertial load produces the largest bending; the displacement peaks at 1,966 × 10^{-1} [mm] at the pad end. The margin is still large, confirming that the real limit here is the pad–object friction discussed in Section 5.1, not the strength of the finger.

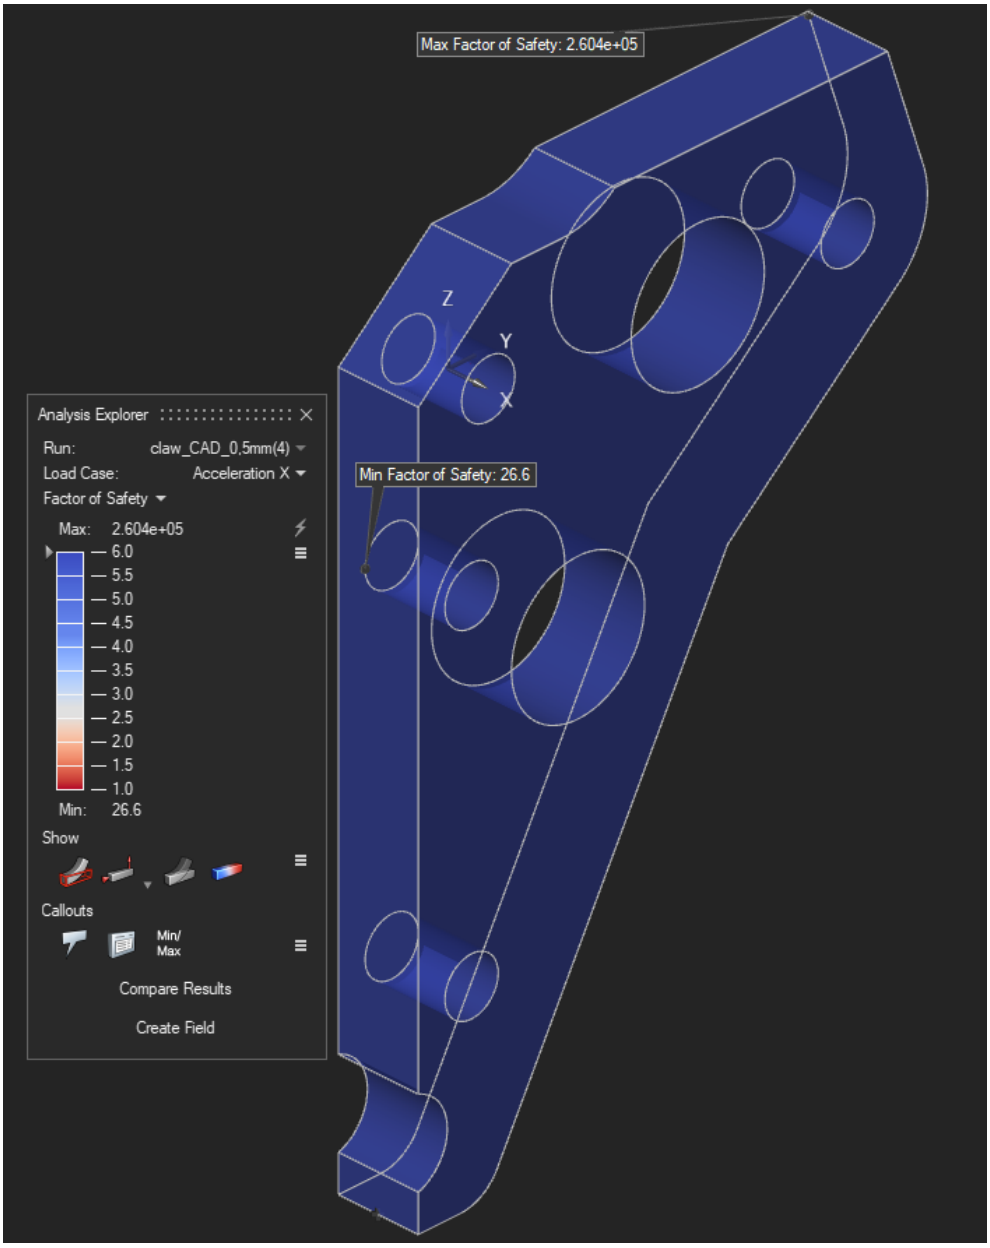

Figure 130: Factor-of-safety distribution on the nominal CAD Claw under the acceleration-along-X load case.

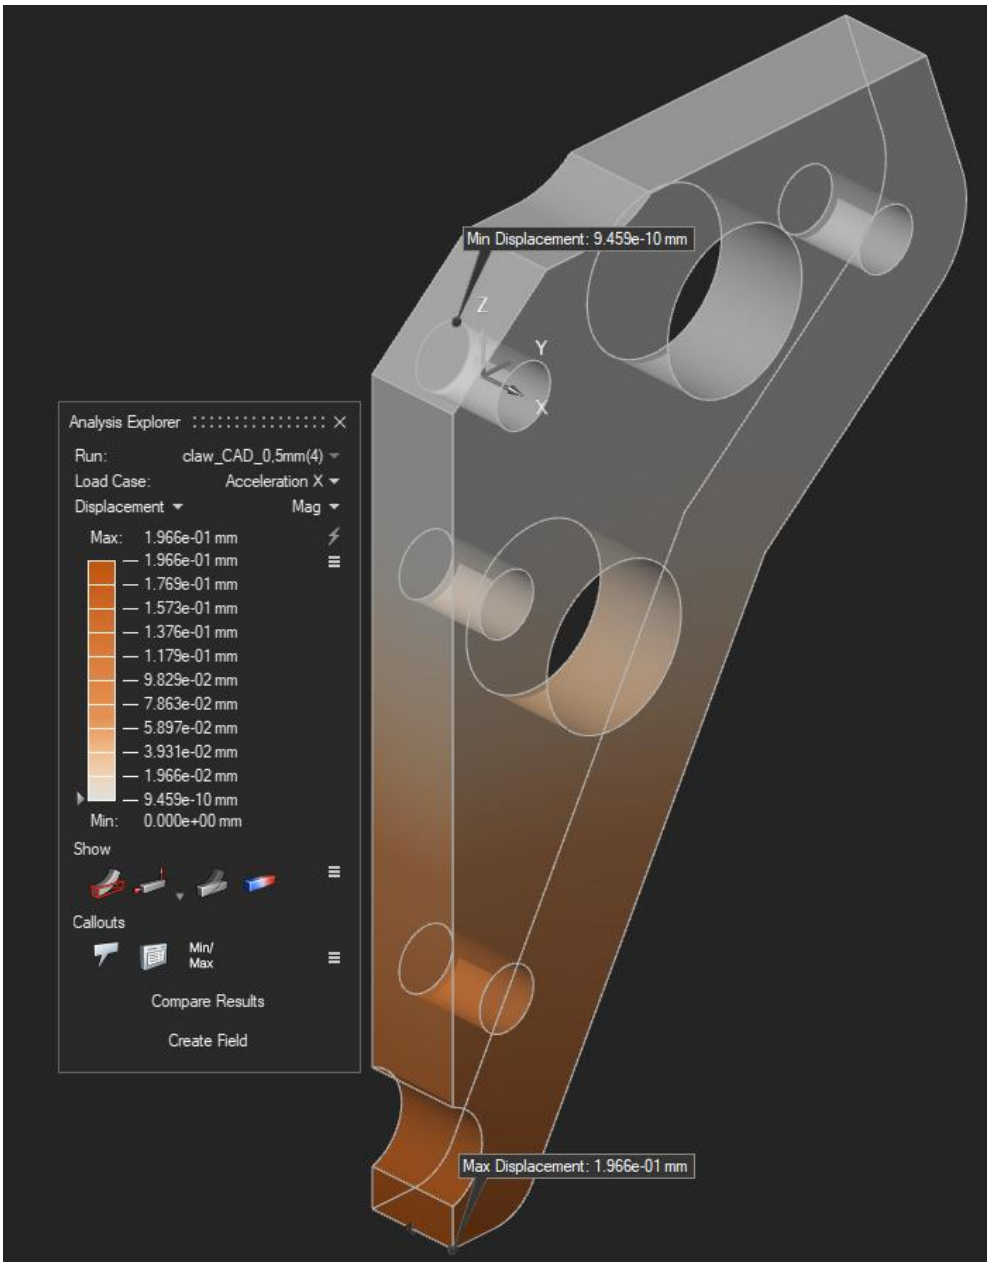

Figure 131: Displacement-magnitude contour of the nominal CAD Claw under the acceleration-along-X load case; peak displacement 0,197 [mm].

Servo stall (gripping)

The gripping (servo-stall) action is the design-driving load case for the Claw, being the largest force applied at the pad. The safety factor falls to 10,2 around the pin seats that react the closing force, while the displacement stays very small at 5,706 × 10^{-2} [mm]. This is the case that governs the topology optimisation of Chapter 6.

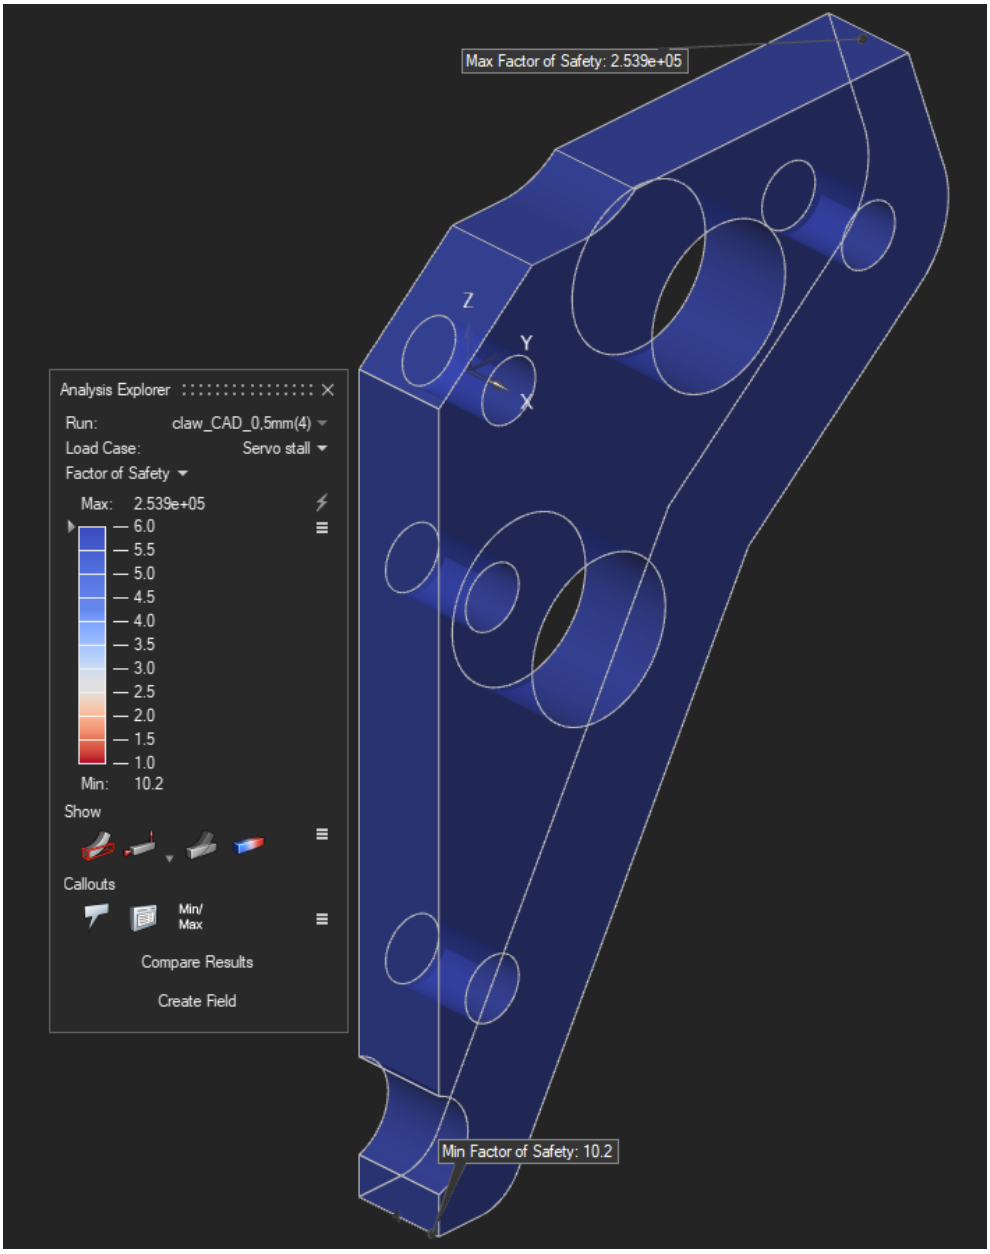

Figure 132: Factor-of-safety distribution on the nominal CAD Claw under the servo-stall (gripping) load case.

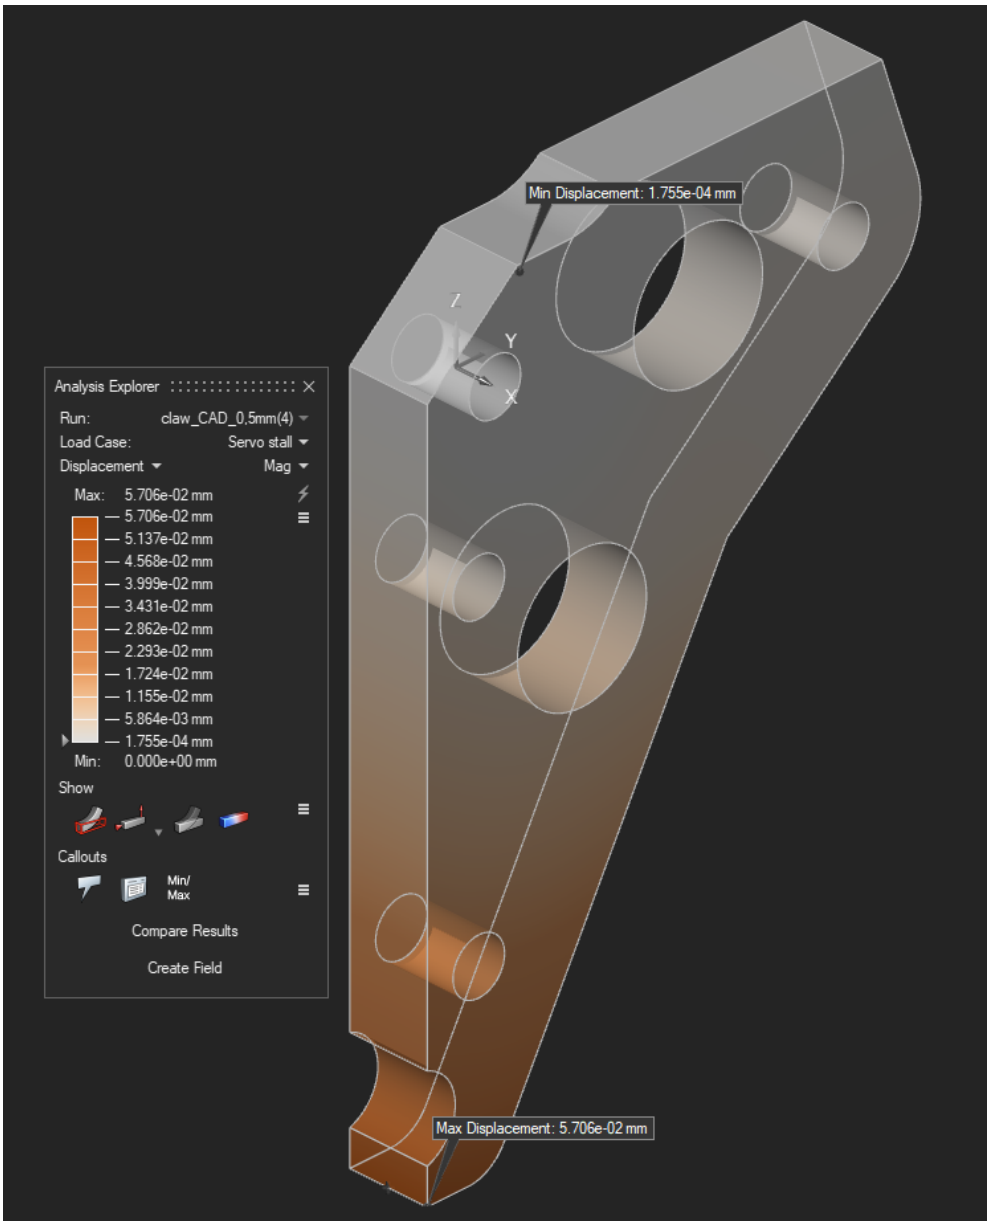

Figure 133: Displacement-magnitude contour of the nominal CAD Claw under the servo-stall (gripping) load case; peak displacement 0,057 [mm].

Frame - reconstructed (CT)

Gravity

On the reconstructed geometry the gravity case shows the first visible effect of the printed infill: the minimum safety factor falls from 181,4 to 136,5, as the self-weight stretches the internal rhomboidal cells, which are weaker than solid material. The peak displacement rises only slightly, to 1,057 × 10^{-2} [mm], and both values remain far from any limit.

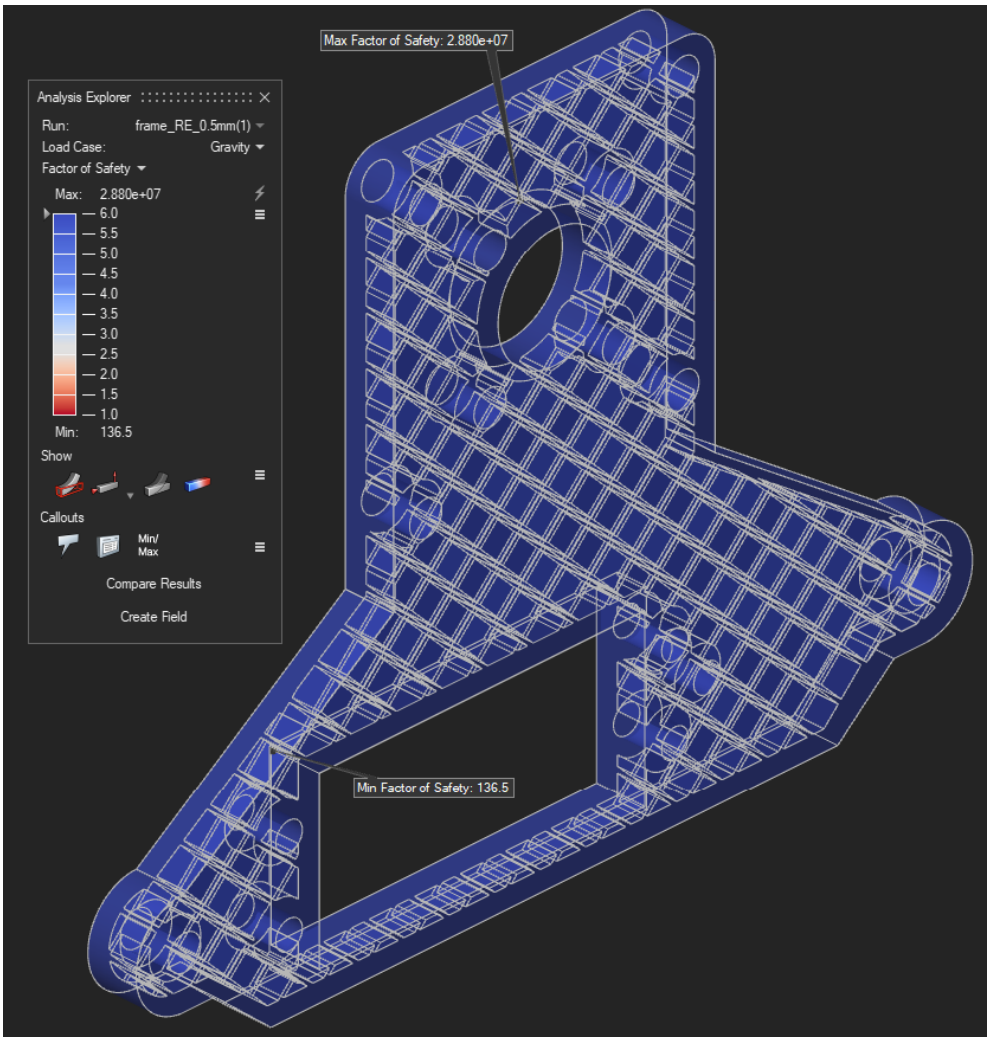

Figure 134: Factor-of-safety distribution on the reconstructed (CT) Frame under the gravity load case.

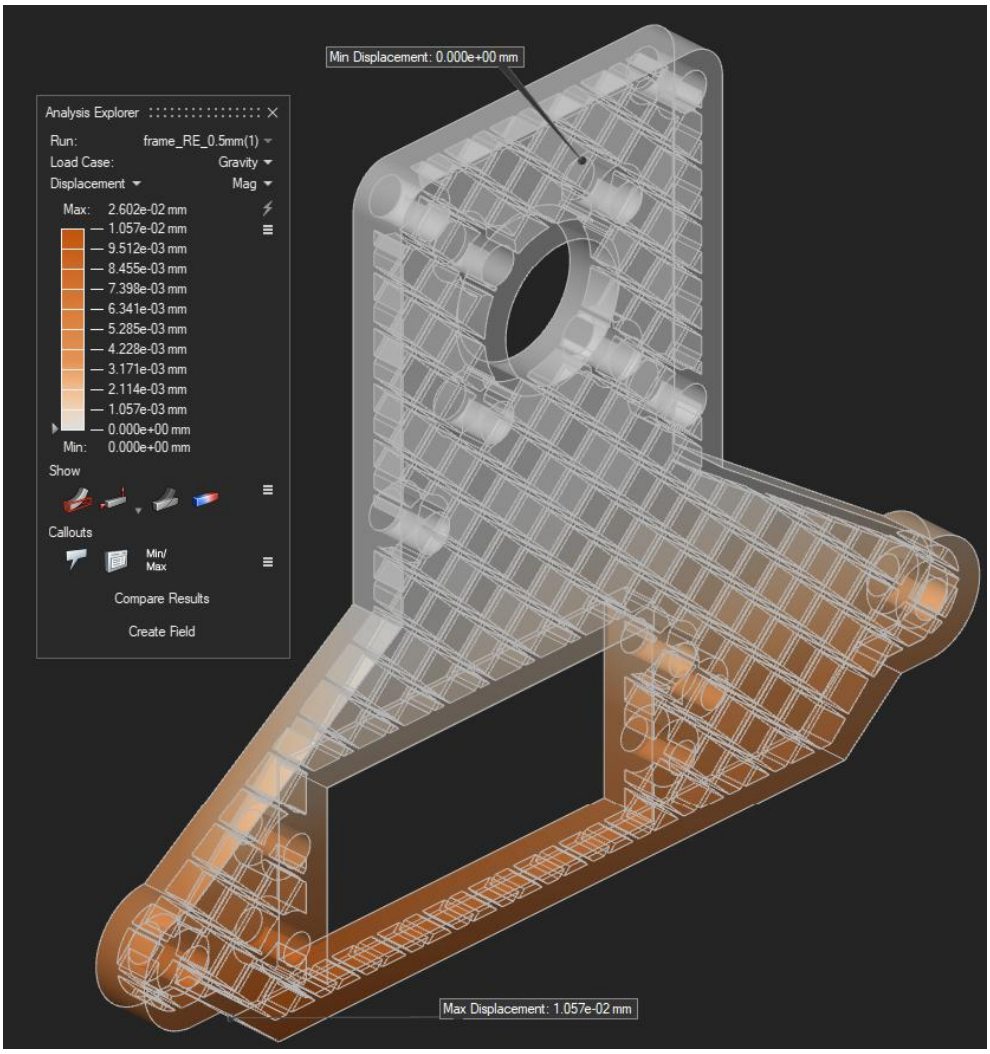

Figure 135: Displacement-magnitude contour of the reconstructed (CT) Frame under the gravity load case; peak displacement 1,057 × 10⁻² [mm].

Acceleration along X

The acceleration case stays critical for the reconstructed Frame, with the safety factor essentially unchanged at 3,8 against the 4,3 of the solid model. This confirms that the governing stress lies on the outer faces normal to X, away from the infill that sits near the neutral axis in bending. The displacement grows to 8,207 × 10^{-1} [mm], the largest of the whole study, consistent with the slightly more compliant real geometry.

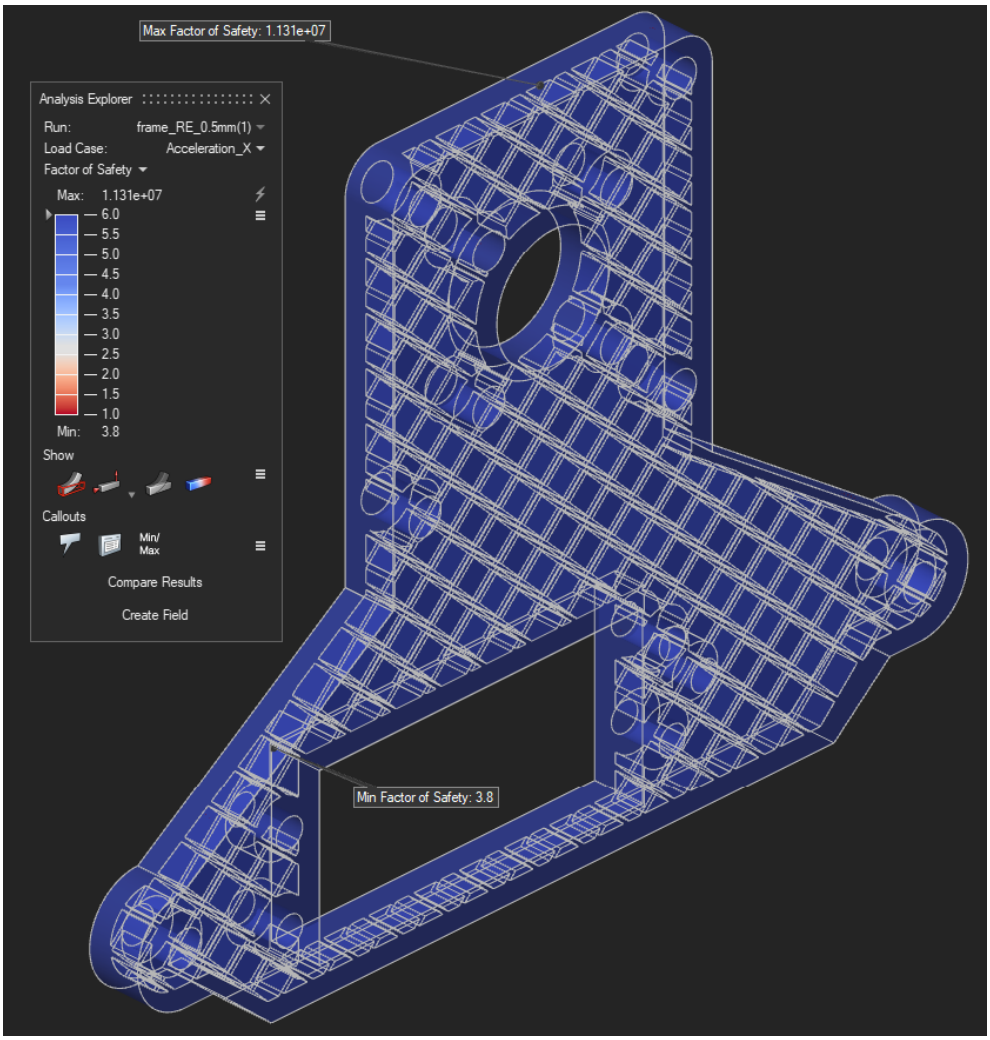

Figure 136: Factor-of-safety distribution on the reconstructed Frame under the acceleration-along-X load case.

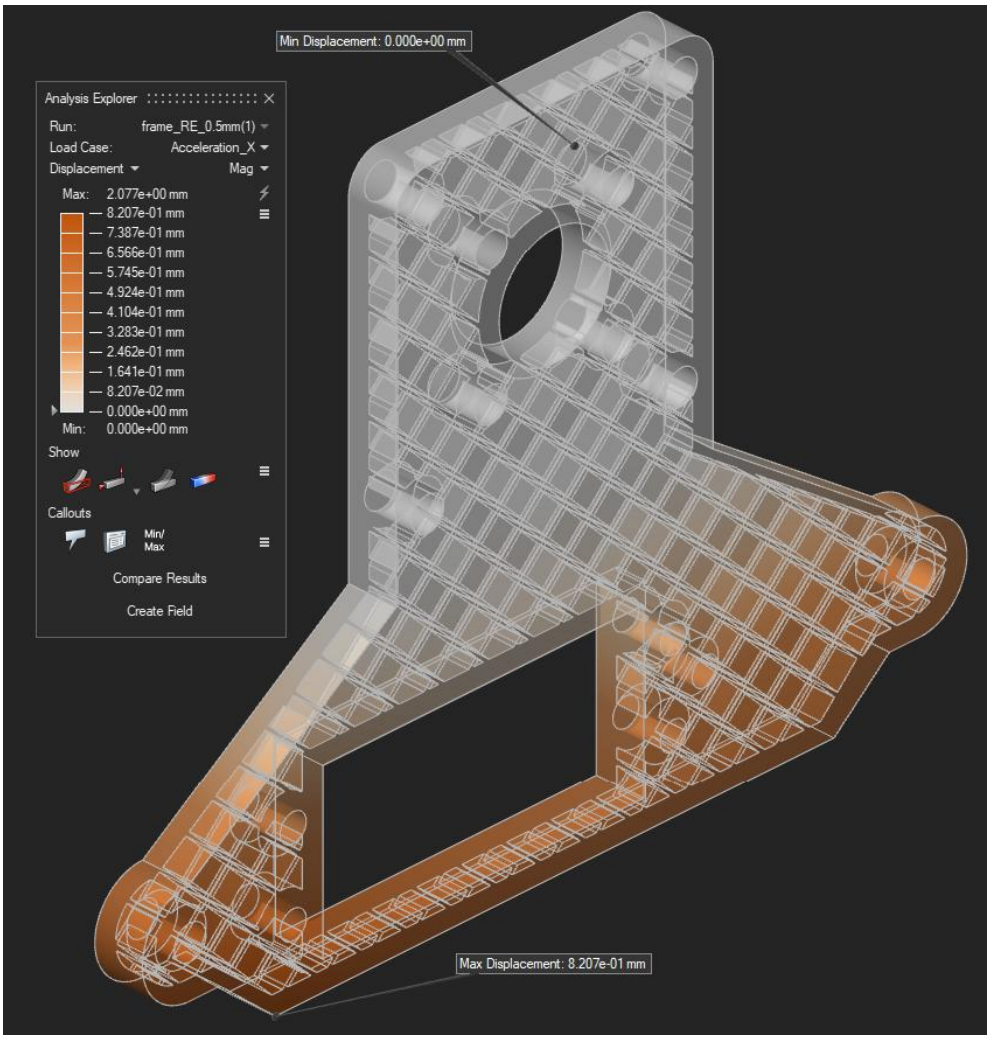

Figure 137: Displacement-magnitude contour of the reconstructed Frame under the acceleration-along-X load case; peak displacement 0,821 [mm].

Servo stall (gripping)

Under servo stall the reconstructed Frame keeps a minimum safety factor of 9,4, close to the solid value of 11,6; the small reduction reflects the torsional load redistributing through the infilled interior. The peak displacement is 2,564 × 10^{-1} [mm], and the component remains well within its safe range.

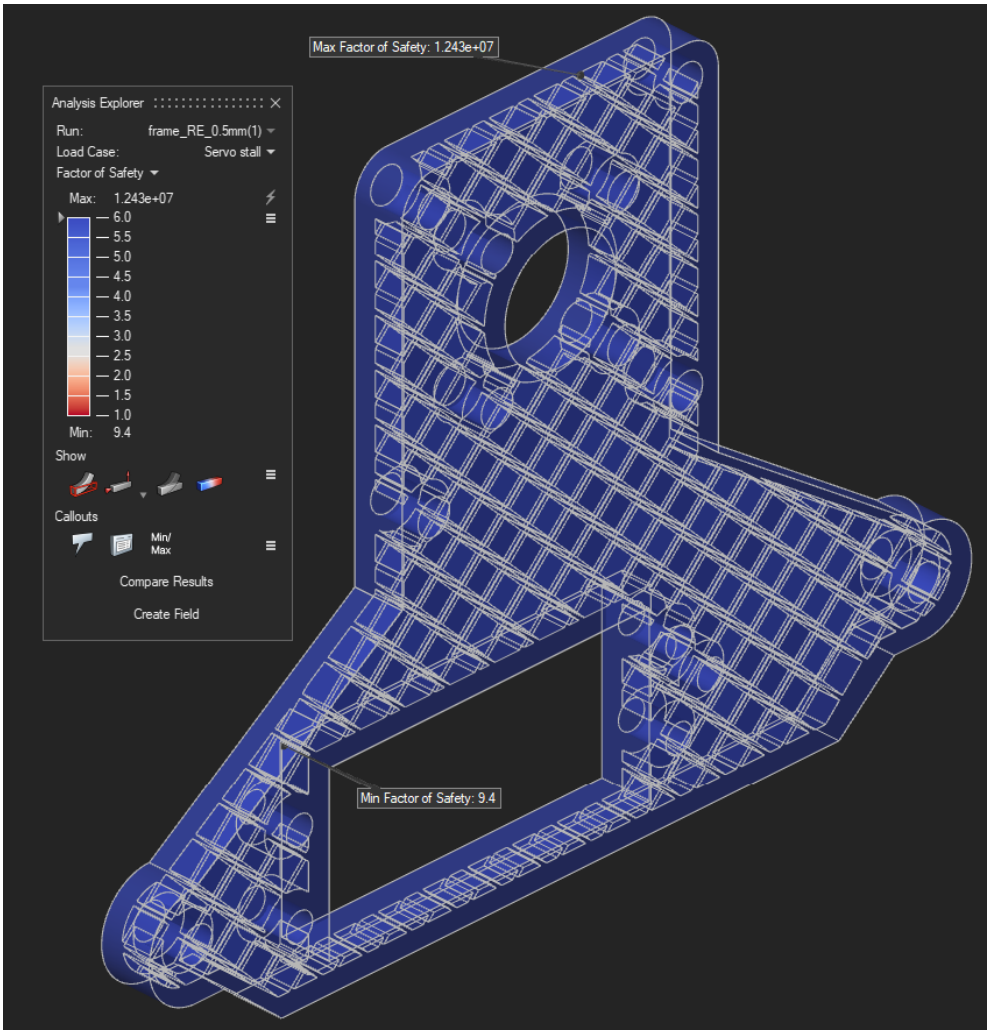

Figure 138: Factor-of-safety distribution on the reconstructed Frame under the servo-stall load case.

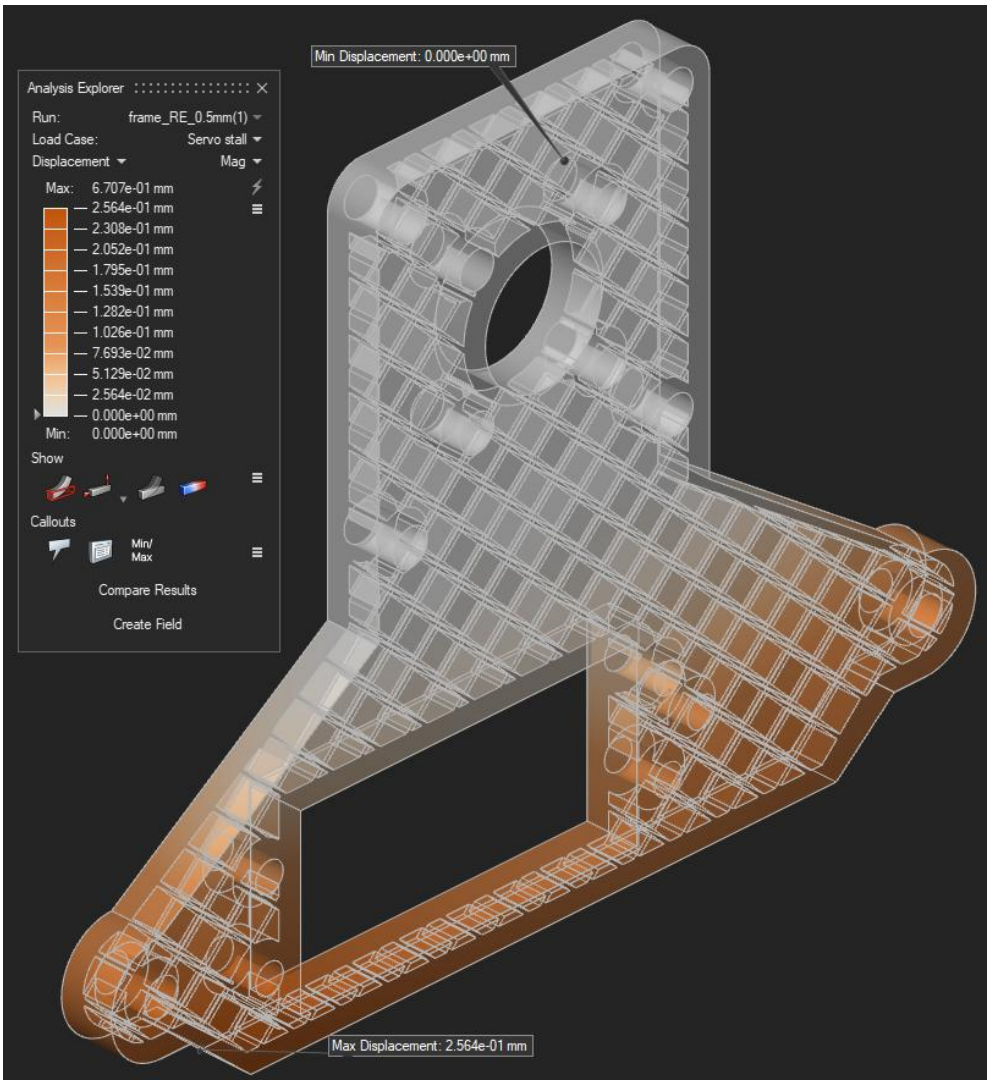

Figure 139: Displacement-magnitude contour of the reconstructed Frame under the servo-stall load case; peak displacement 0,256 [mm].

Claw - reconstructed (CT)

Gravity

This is the most pronounced infill effect of the whole study: on the reconstructed Claw the gravity safety factor collapses from 271,1 to 10,3, because the load acts in tension across the weakest direction of the printed cells, making the self-weight case the one most sensitive to the real geometry. The displacement, however, stays negligible at $3{,}766 \times 10^{-3}$ [mm], showing that the drop is a local strength effect and not a loss of global stiffness.

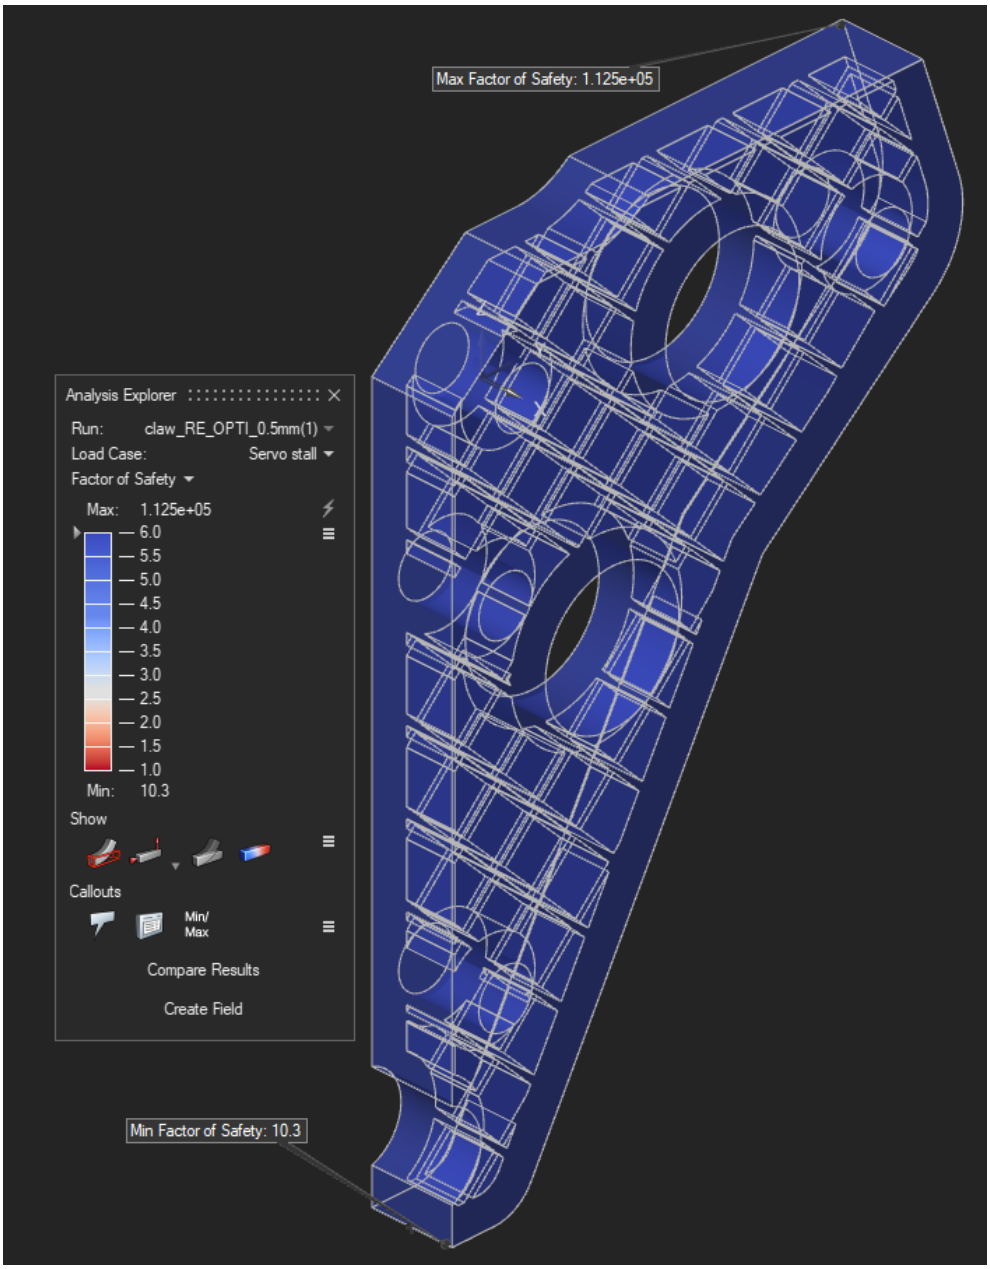

Figure 140: Factor-of-safety distribution on the reconstructed Claw under the gravity load case.

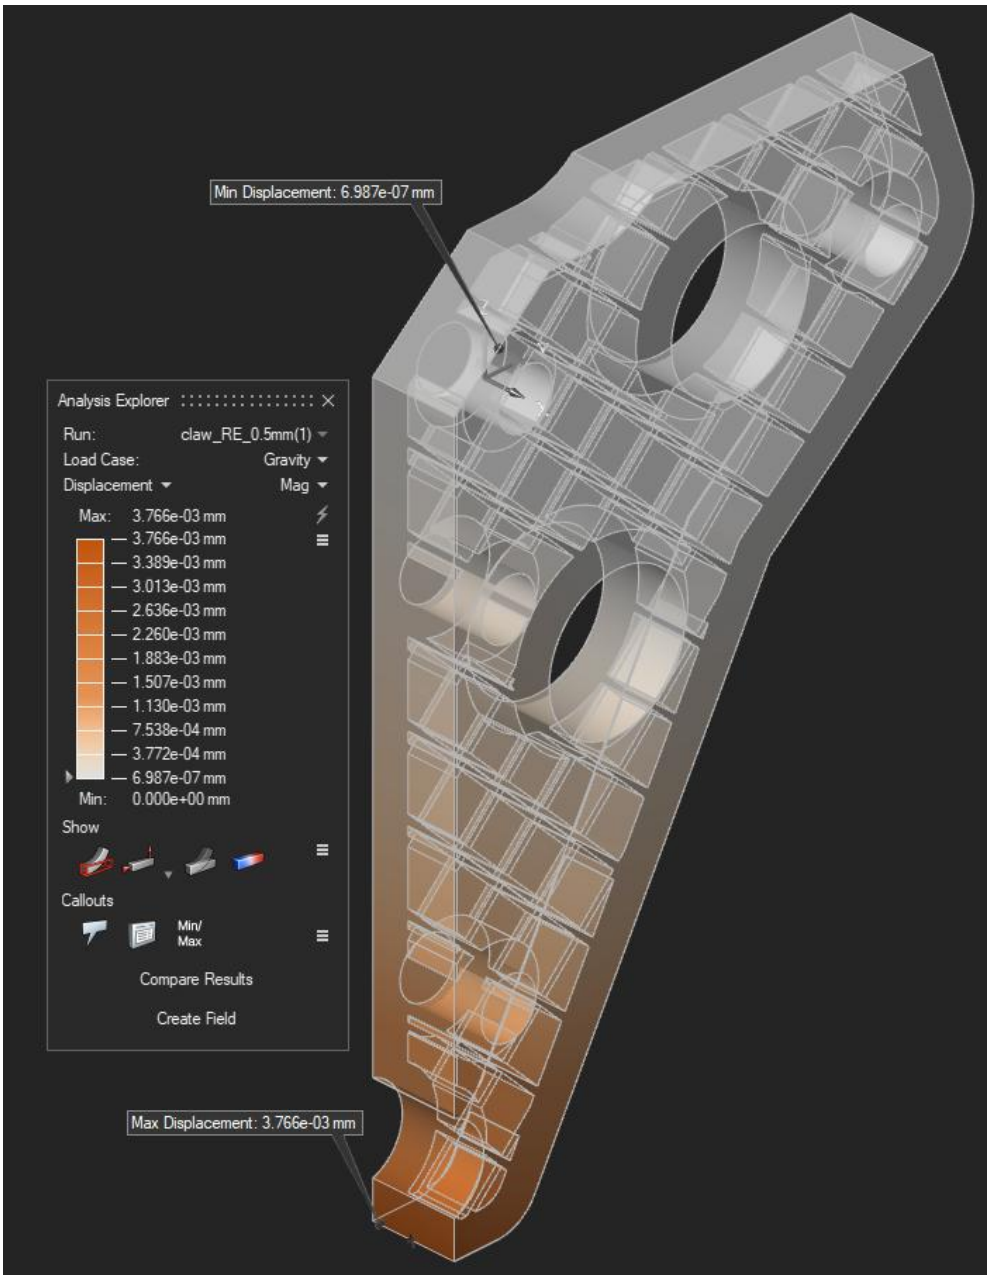

Figure 141: Displacement-magnitude contour of the reconstructed (CT) Claw under the gravity load case; peak displacement 3,766 × 10^{-3} [mm].

Acceleration along X

The acceleration case is almost insensitive to the infill, with the safety factor moving only from 26,6 to 25,0, again because the critical fibres lie on the outer faces rather than in the cellular interior. The peak displacement is 2,325 × 10^{-1} [mm], marginally above the solid model.

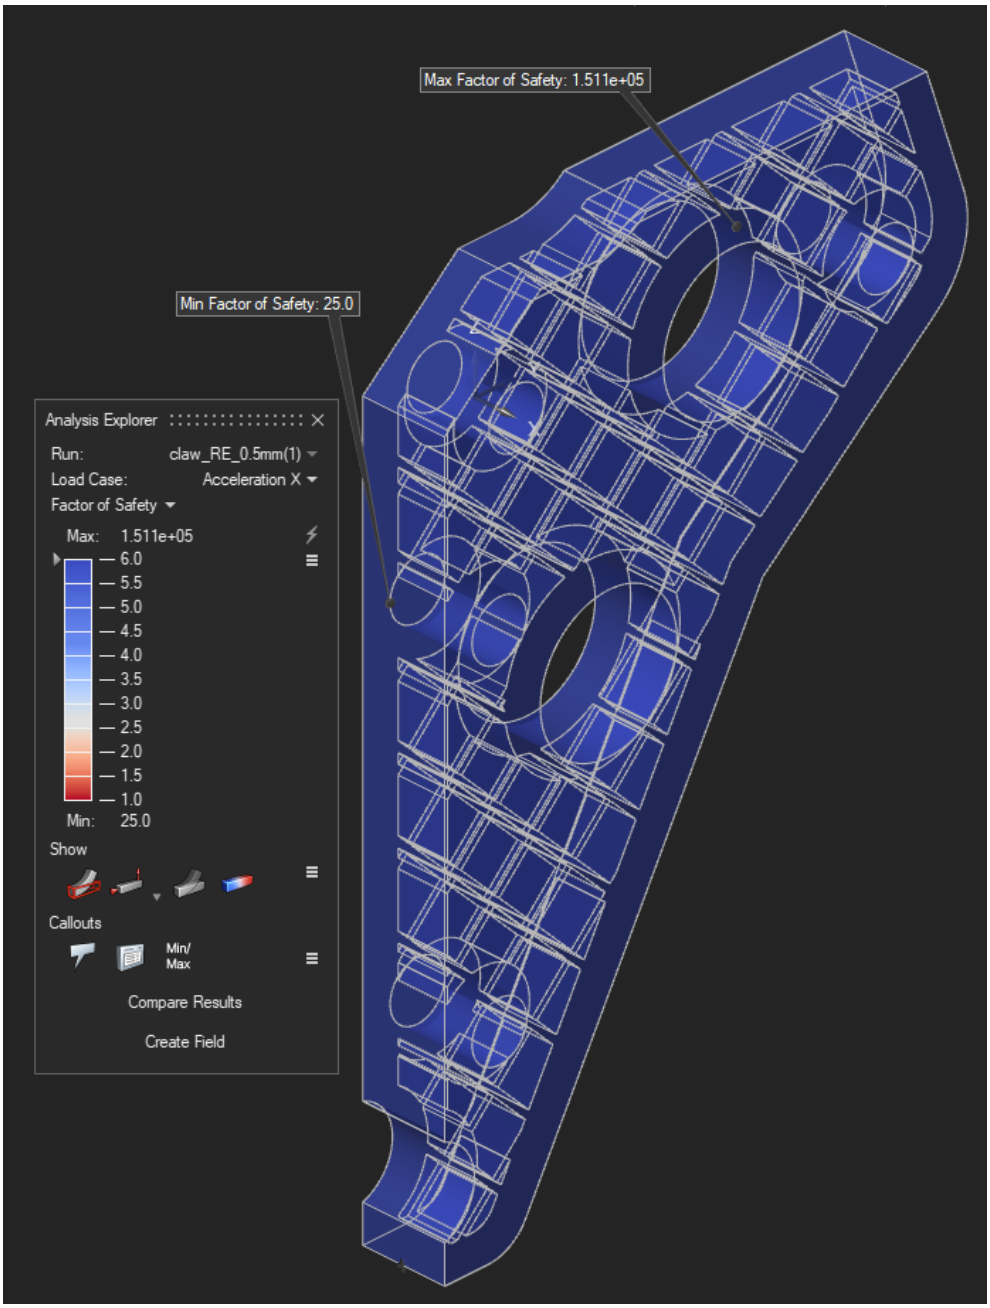

Figure 142: Factor-of-safety distribution on the reconstructed Claw under the acceleration-along-X load case.

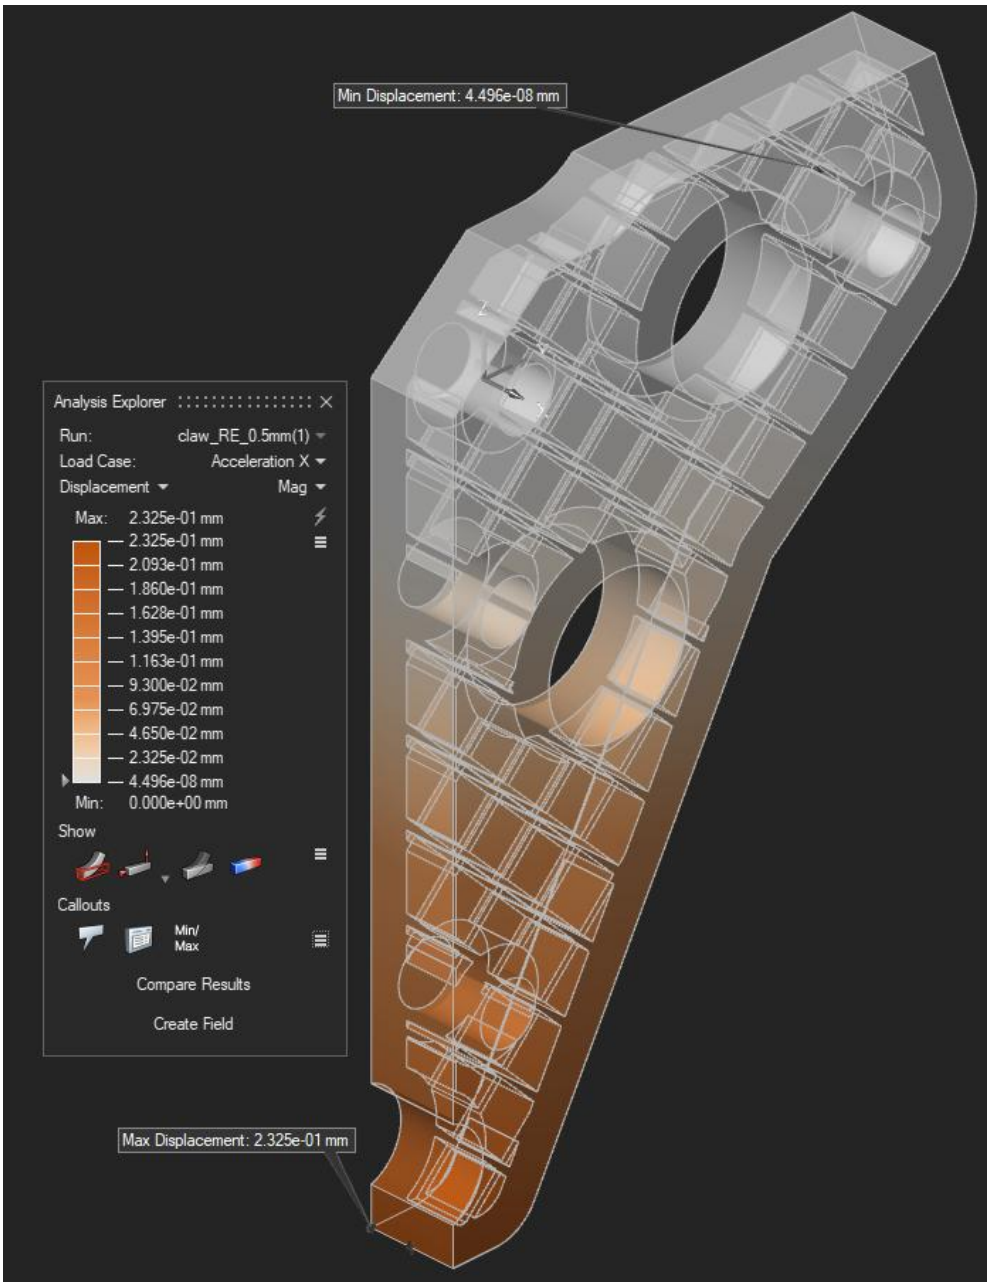

Figure 143: Displacement-magnitude contour of the reconstructed Claw under the acceleration-along-X load case; peak displacement 0,233 [mm].

Servo stall (gripping)

For the gripping case the reconstructed Claw retains a minimum safety factor of 10,3, practically identical to the solid value of 10,2, since the closing force is reacted mainly at the solid pin seats. The displacement peak of 7,693 × 10^{-2} [mm] is likewise close to the nominal result, so the infill is of little relevance to this design-driving case.

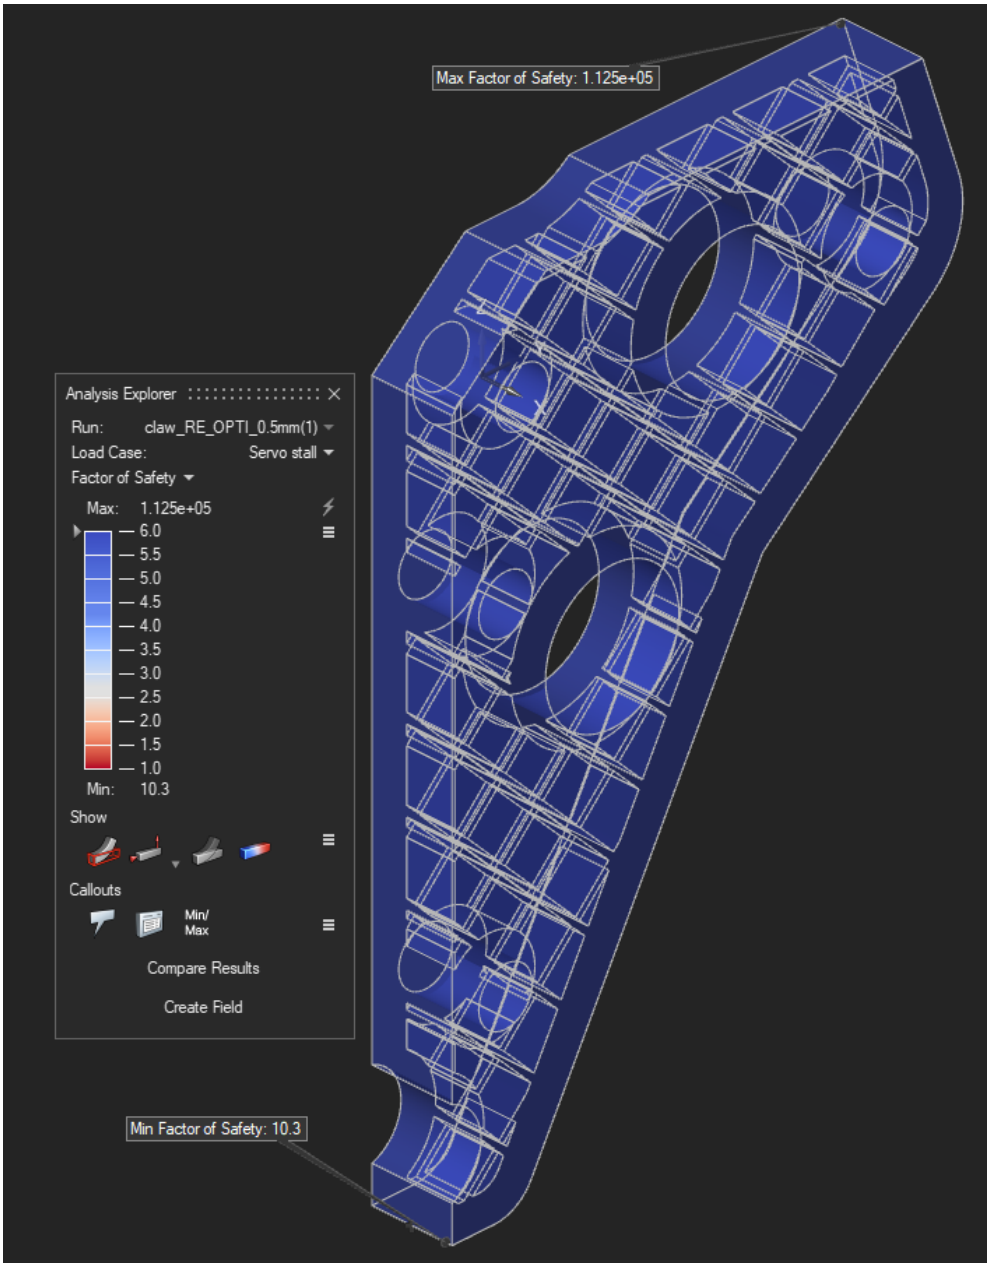

Figure 144: Factor-of-safety distribution on the reconstructed Claw under the servo-stall (gripping) load case.

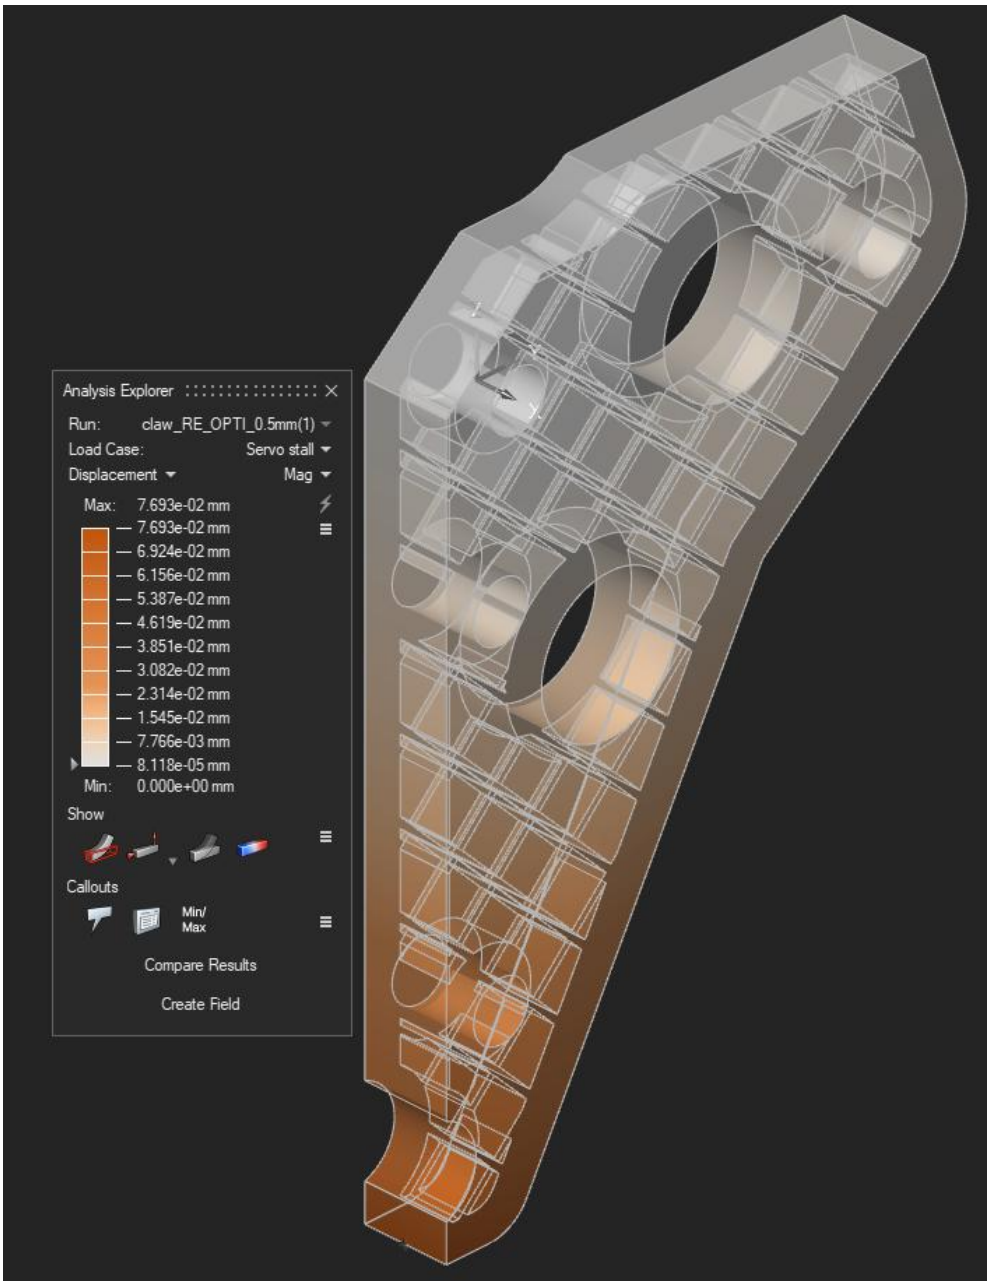

Figure 145: Displacement-magnitude contour of the reconstructed Claw under the servo-stall (gripping) load case; peak displacement 0,077 [mm].

I - Technical Specifications Document

This appendix summarises the hardware and the software configuration used throughout the project, so that the acquisition, reconstruction and optimisation campaign can be reproduced under equivalent conditions. Version fields refer to the releases available in the university laboratories during the academic year 2025-2026.

Hardware

Item	Model / key specification	Role in the workflow
FDM 3D printer	Ultimaker S3, 0,4 [mm] nozzle	Manufacturing of the replica and of the optimised parts (Sections 2.4, 7.3)
Handheld structured-light scanner	Shining 3D EinScan Pro 2X	External-surface acquisition (Section 3.2)
Fixed structured-light scanner	GOM ATOS Core 300	External-surface acquisition (Section 3.3)
X-ray CT system	GE Phoenix v|tome|x s, microfocus tube	Internal and external acquisition (Section 3.4)
Servo actuator	Robbe/Futaba S3003, 402 N·mm stall torque	Gripper actuation (Section 2.2)
End-effector kit	ServoCity Servo-Driven Parallel Gripper Kit (SKU 3219-0001-0002)	Reference assembly (Section 2.1)
Workstation	Acer Nitro V16, AMD Ryzen 7 8845HS, 16 GB DDR5-5600, NVIDIA GeForce RTX 4060 Laptop	CAD, FEM and topology-optimisation runs

Software

Software	Version	Licence	Use
SolidWorks	2025	Educational	Parametric CAD reconstruction and drawings (Section 4.2)
Altair Inspire	2026	Educational	FEM analyses and topology optimisation (Chapters 5-6)
GOM Inspect	2019	Free	Mesh inspection and deviation maps (Chapter 3)
Ultimaker Cura	5.10.2	Free	Slicing of original and redesigned parts (Section 2.4)
Autodesk Meshmixer	3.5.0	Free	Mesh healing of handheld scans (Section 3.2)
Python	3.11	Open source	Custom processing scripts (Appendices E-F)
Open3D / NumPy	1.24.3/0.17.0	Open source	HPR outer-shell extraction (Appendix E)
OpenCV	4.8.0	Open source	HSV infill segmentation (Appendix F)
phoenix datos|x	3.2	Departmental	CT volume reconstruction (Section 3.4)