FACULDADE DE ENGENHARIA DA UNIVERSIDADE DO PORTO

Roughness Characterisation of Laser-cut Metal Pieces using Vision-based Methods

Ricardo Manuel Cunha Meireles



Mestrado em Engenharia Mecânica

Supervisor: Prof. Armando Jorge Sousa

Second Supervisor: Prof. Jorge Wolfs Gil

September 18, 2026

© Ricardo Meireles, 2026

Roughness Characterisation of Laser-cut Metal Pieces using Vision-based Methods

Ricardo Manuel Cunha Meireles

Mestrado em Engenharia Mecânica

Approved in oral examination by the committee:

President: António Manuel Ferreira Mendes Lopes
Examiner: Daniel Queirós Silva
Member: Armando Jorge Miranda de Sousa

September 18, 2026

Resumo

O corte a laser é a tecnologia de referência para a transformação de chapa metálica na indústria moderna. Contudo, a qualidade do produto final ainda depende do julgamento do operador, de paragens na produção para perfilometria de contacto, ou da inspeção visual e tátil de um operador experiente. Esta dissertação, realizada em parceria com a ACN Cutting Systems no âmbito do projeto FiberNest XL, investiga se a rugosidade superficial de chapa metálica cortada a laser pode ser estimada a partir de uma única imagem RGB, com uma precisão útil para inspeção em ambiente industrial, abrangendo quatro materiais e várias espessuras de chapa metálica.

O trabalho combina duas modalidades complementares de caracterização da rugosidade superficial. A primeira consiste numa abordagem espetral que opera sobre imagens RGB da face de corte, com o objetivo de determinar o comprimento de onda entre os vales gerados pelo processo e analisar a sua variação ao longo da espessura da face. A segunda abordagem extrai perfis de rugosidade a partir de um *scanner* 3D, calibrado com um perfilómetro de contacto. Esta modalidade alcançou erros absolutos médios de aproximadamente 1.2 µm em R_a, 1.3 µm em R_q e 5.4 µm em R_z, validando o *scanner* como uma fonte prática de *labels* de rugosidade para modelos de inteligência artificial.

No âmbito deste trabalho, foi ainda projetado e montado um sistema automático de aquisição de imagens. As duas modalidades descritas foram combinadas num conjunto de dados multimodal, sobre o qual foram estudadas três arquiteturas de regressão dos parâmetros R_a, R_q e R_z. Adotou-se um protocolo de treino e validação robusto, livre de fugas de informação, que particiona os dados por face de corte. A partir destes dados, uma arquitetura *ResNet-18*, pré-treinada no *dataset ImageNet*, foi submetida a *fine-tuning*, atingindo um coeficiente de determinação, entre 0,73 e 0,76 e um erro absoluto médio próximo de 3.8 µm em R_a, uma precisão diretamente compatível com os requisitos industriais. Por fim, projetou-se uma arquitetura multimodal para consumir tanto o contexto tabular como as imagens do *dataset* criado, a qual se estabeleceu como o melhor preditor de R_z. No seu conjunto, os resultados validam a viabilidade da estimação da rugosidade em ambiente industrial a partir de uma única imagem da face de corte.

Palavras-chave: Corte laser • Rugosidade • Deep Learning • Luz estruturada • Análise espetral

Abstract

Laser cutting is the reference technology for sheet metal transformation in modern manufacturing. Nevertheless, the quality of the final product still depends on operator judgement, on production stops for contact profilometry, or on the visual and tactile inspection of an experienced operator. This dissertation, carried out in collaboration with ACN Cutting Systems within the FiberNest XL project, investigates whether the surface roughness of a laser-cut sheet metal edge can be estimated from a single RGB image with an accuracy that is useful for shop floor inspection, covering four materials and several plate thicknesses.

The work combines two complementary modalities of surface roughness characterisation. The first consists of a spectral approach that operates on the RGB images of the cut face, with the objective of determining the wavelength between the valleys generated by the process and of analysing its variation along the thickness of the face. The second modality extracts roughness profiles from a 3D scanner calibrated against a contact profilometer. This modality achieved mean absolute errors of approximately 1.2 µm on R_a, 1.3 µm on R_q and 5.4 µm on R_z, validating the scanner as a practical source of roughness labels for artificial intelligence models.

An automated image acquisition setup was also designed and assembled in the scope of this work. The two modalities described above were combined into a multimodal dataset over which three regression architectures for the parameters R_a, R_q and R_z were studied. A robust training and validation protocol, free of data leakage and partitioning the data by cut face, was adopted. From these data, a ResNet-18 architecture, pretrained on the ImageNet dataset, was fine-tuned, reaching a coefficient of determination between 0.73 and 0.76 and a mean absolute error close to 3.8 µm on R_a, an accuracy directly compatible with industrial requirements. Finally, a multimodal architecture was designed to consume both the tabular context and the images of the created dataset, and became the best predictor of R_z. Together, the results validate the feasibility of estimating roughness in an industrial environment from a single image of the cut face.

Keywords: Laser cutting • Surface roughness • Deep learning • Structured light scanning • Spectral analysis

The United Nations Sustainable Development Goals

The United Nations Sustainable Development Goals (SDG) provide a global framework for achieving a better and more sustainable future for all. This dissertation contributes to two of these goals by developing a vision-based, non-contact system for surface roughness estimation in laser-cut metal parts, supporting the technological upgrading of manufacturing processes and promoting more efficient and sustainable industrial production.

The specific Sustainable Development Goals addressed are the following:

SDG 8 Promote sustained, inclusive and sustainable economic growth, full and productive employment and decent work for all

SDG 9 Build resilient infrastructure, promote inclusive and sustainable industrialisation and foster innovation

Table 1: Dissertation contribution to the UN Sustainable Development Goals.

SDG	Target	Contribution	Performance Indicators and Metrics
8	8.2	The proposed inspection system enables real-time roughness assessment directly on the shop floor, eliminating production stoppages for laboratory measurements and reducing inspection time from minutes to seconds, thereby increasing manufacturing productivity.	• Reduction in inspection time per part • Increase in parts inspected per shift
	8.4	Early detection of out-of-specification cut surfaces reduces unnecessary secondary machining operations and material waste, contributing to more efficient resource use in production.	• Reduction in secondary machining operations • Decrease in non-conforming parts produced before detection
9	9.2	A low-cost RGB camera-based device replaces expensive metrological instruments, democratising access to quantitative quality control in laser cutting for small and medium-sized enterprises.	• Acquisition cost relative to conventional metrological instruments • Applicability across multiple materials and thicknesses without recalibration
	9.4	The adoption of a vision-based, non-contact inspection method upgrades the quality control process in laser cutting, replacing invasive procedures with a sustainable, low-energy alternative that supports cleaner industrial processes.	• Reduction in production line waiting time for inspection results • Number of contact-based inspection procedures replaced

Agradecimentos

Ao Professor Doutor Armando Jorge Sousa pela orientação do trabalho e disponibilidade, pela sua perspicácia e olhar crítico característicos.

Ao Professor Doutor Jorge Wolfs Gil, coorientador do trabalho, pela disponibilidade e por toda a ajuda prestada. A ambos pela confiança que depositaram em mim e por acreditarem no trabalho que desenvolvi.

Uma palavra especial para o Levy, o acompanhamento diário, a partilha de conhecimento técnico e a disponibilidade permanente para discutir qualquer problema foram decisivos. O meu muito obrigado.

Gostaria de agradecer ao projeto "FIBERNEST XL - Equipamento de corte a laser de fibra de grandes dimensões com auto-nesting", cujo número é 20436 e código de projeto COMPETE2030-FEDER-01919100, co-financiado pelo COMPETE 2030, Portugal 2030 e por fundos da União Europeia.

A todos os meus colegas e amigos que se cruzaram comigo durante todo o meu percurso académico, obrigado pelas horas de estudo partilhadas, pelas discussões dentro e fora das aulas e pelo companheirismo.

Por fim, aos meus pais e ao meu irmão, nada disto existiria sem o vosso apoio incondicional. Obrigado por acreditarem em mim, pelos sacrifícios que fizeram ao longo do caminho e por estarem sempre lá.

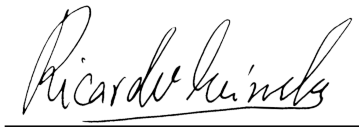
Ricardo Meireles

Declaração de Honra

Declaro que a presente dissertação é de minha autoria e não foi previamente apresentada e avaliada nesta ou em outra instituição. As referências a outros autores (afirmações, ideias, pensamentos), assim como a utilização de trabalhos publicados respeitam escrupulosamente as regras da atribuição, e encontram-se devidamente indicadas no texto e nas referências bibliográficas, de acordo com as normas de referenciação. Tenho consciência de que a prática de plágio ou de autoplágio constituem ilícitos académicos, conforme estabelecido no Código de Ética da U.Porto.

Mais declaro que incluí texto(s) publicado(s) em coautoria, tendo especificado, com transparência e rigor, a colaboração de cada coautor e informado sobre a eventual existência de alguma outra dissertação em que também tenha sido incluído. São também incluídos a autorização da revista em que se encontra publicado, assim como o acordo expresso assinado por todos os coautores para a inclusão do artigo.

Declaro, ainda, que utilizei ferramentas de Inteligência Artificial para a realização de parte(s) da presente dissertação, e essa utilização e os seus resultados estão devidamente assinalados e foram objeto de cuidada verificação para deteção de erros, imprecisões, atribuições infundadas ou outras formas de enviesamento de conteúdos e argumentos, bem como de determinação de autoria.



"Sempre que um homem sonha, o mundo pula e avança"

António Gedeão

Contents

1 Introduction **1**
- 1.1 Motivation and Context . . . 1
- 1.2 Proposed Solution and Objectives . . . 3
- 1.3 Dissertation Structure . . . 5

2 State of the Art and Fundamental Issues **6**
- 2.1 Laser cutting process fundamentals . . . 7
  - 2.1.1 Process principle and cutting regimes . . . 7
  - 2.1.2 Laser sources and beam characteristics . . . 8
  - 2.1.3 Process parameters and quality outputs . . . 10
  - 2.1.4 Formation of the cut surface and striations . . . 11
  - 2.1.5 Influence of material and thickness . . . 12
  - 2.1.6 Implications for image-based assessment . . . 13
- 2.2 Cut quality and surface roughness characterisation . . . 13
  - 2.2.1 Thermal cut quality - ISO 9013 . . . 13
  - 2.2.2 Profile roughness parameters . . . 16
  - 2.2.3 Profile filtering and evaluation length . . . 17
- 2.3 Assessment of laser-cut edges . . . 18
  - 2.3.1 AI-based quality assessment . . . 18
  - 2.3.2 Other assessment methods . . . 20
- 2.4 Image-based spectral analysis of surface texture . . . 21
- 2.5 Surface roughness measurement techniques . . . 23
- 2.6 AI models and evaluation . . . 25
  - 2.6.1 Convolutional neural networks . . . 25
  - 2.6.2 Transfer learning and fine-tuning . . . 26
  - 2.6.3 Multilayer perceptrons . . . 26
  - 2.6.4 Multimodal fusion . . . 27
  - 2.6.5 Training and validation . . . 27
- 2.7 Summary and research gap . . . 28

3 Materials and Methods **32**
- 3.1 Laser-Cut Specimens . . . 34
- 3.2 Sensors and Acquisition Systems . . . 34
  - 3.2.1 RGB Image Acquisition . . . 36
  - 3.2.2 Microscope Calibration . . . 38
  - 3.2.3 Surface Roughness Measurement . . . 41
  - 3.2.4 Structured Light Scanning . . . 42
- 3.3 Validation of ATOS Q Surface Roughness Measurements . . . 45

3.3.1 Profile Extraction from Point Cloud Data . . . 45
3.3.2 Spatial Registration and Roughness Calibration . . . 47
3.4 Striation Period Estimation: Initial Approach . . . 49
3.4.1 FFT-Based Striation Wavelength Estimation . . . 49
3.4.2 Validation via Laser Line Projection . . . 50
3.4.3 Harmonic Ambiguity and Motivation for a Robust Pipeline . . . 53
3.5 Spectral Feature Extraction from RGB Images . . . 53
3.5.1 Image Pre-processing . . . 54
3.5.2 One-Dimensional Profile and Spectral Analysis . . . 55
3.5.3 Harmonic Product Spectrum for Fundamental Identification . . . 56
3.5.4 Spectral Strength and Peak Confidence . . . 56
3.5.5 Top-Three Peak Extraction . . . 57
3.6 Dataset Construction . . . 58
3.6.1 Spatial Alignment, Strip Definition and Co-localisation . . . 58
3.6.2 Exclusion Criteria . . . 60
3.6.3 Roughness Label Assignment . . . 60
3.6.4 Data Pre-processing and Final Dataset . . . 61
3.7 Deep Learning Architectures . . . 62
3.7.1 Feature-driven MLP . . . 62
3.7.2 Image-driven ResNet-18 . . . 63
3.7.3 Multimodal ResNetFusion . . . 64
3.7.4 Training and Validation Protocol . . . 64
3.7.5 Feature Sets . . . 65

4 Results and Discussion **67**
4.1 Roughness Extraction and Calibration . . . 67
4.2 Spectral Analysis of Cut-Edge Images . . . 71
4.2.1 Variation of the Wavelength within a Face . . . 71
4.2.2 Effect of the Spectral Pipeline on the Extracted Wavelength . . . 73
4.2.3 Peak Strength and Peak Confidence . . . 75
4.2.4 Local versus Global Descriptors and Overall Structure . . . 76
4.3 Roughness Prediction with Deep Learning Models . . . 78
4.3.1 Training Behaviour . . . 78
4.3.2 General Models: Performance per Image . . . 79
4.3.3 Single-Material Models . . . 81
4.4 Discussion . . . 83
4.4.1 Roughness Labels from the Structured-Light Scanner . . . 83
4.4.2 Spectral Features from the RGB Microscope Images . . . 83
4.4.3 Deep Learning Models . . . 84

5 Conclusions **86**
5.1 Response to the Research Aim . . . 86
5.2 Contributions . . . 87
5.3 Limitations . . . 88
5.4 Future Work . . . 89

References **91**

List of Figures

2.1 Schematic of a fibre laser, adapted from [2]. . . . 9
2.2 Schematic of a CO_2 laser, adapted from [2]. . . . 9
2.3 Schematic of a Nd:YAG pulsed laser, adapted from [2]. . . . 10
2.4 Laser-cut edge showing the three characteristic zones of the cut face: Beam entry, shelf and washout zones. . . . 12
2.5 Cut geometry and reference dimensions, adapted from [24]. . . . 14
2.6 Angular characteristics of the cut face, adapted from [24]. . . . 15
2.7 Cross-sectional view of a laser-cut edge illustrating the perpendicularity or angularity tolerance u and the excluded zones Δa. Adapted from [24]. . . . 15
2.8 Surface roughness profile illustrating the calculation of the mean height of the profile R_{z5}. Adapted from [24]. . . . 16
2.9 Dross formation on the lower edge of a laser-cut face. Adapted from [24]. . . . 16
2.10 Operating principle of a contact stylus profilometer. . . . 24
2.11 Fringe pattern projected onto a metallic surface during structured-light scanning. 25

3.1 Overview of the methodological pipeline adopted in this dissertation. After multimodal acquisition, the workflow diverges into two parallel branches: validation of the ATOS Q system for surface roughness measurements, and spectral feature extraction from RGB data. Both branches converge into the dataset construction stage, which supports the subsequent AI-based regression. . . . 33
3.2 Laser-cut specimens considered in the experimental campaign. . . . 34
3.3 Overview of the two acquisition setups, sensor configurations, and respective data outputs. . . . 36
3.4 Calibration slide supplied by the manufacturer. . . . 37
3.5 RGB acquisition. . . . 37
3.6 Pinhole camera model used as the theoretical basis for calibration. . . . 38
3.7 Multi-view planar calibration principle adopted in Zhang's method. . . . 39
3.8 Feature-extraction pipeline used to detect the calibration grid intersections: grayscale input, binary mask, horizontal and vertical curves, merged lines, and final intersection points. . . . 40
3.9 Mitutoyo SJ-210 contact surface roughness tester and experimental setup. . . . 41
3.10 Zeiss ATOS Q during specimen scanning. . . . 44
3.11 Experimental setup for the laser line projection validation experiment. . . . 50
3.12 Validation pipeline for valley wavelength estimation using oblique laser line projection and HSV colour segmentation. . . . 51
3.13 Laser line projection result on a representative specimen face: original RGB image under oblique laser illumination (left), binary segmentation mask (centre), and connected-component overlay with annotated centroids (right). . . . 52

3.14 Illustration of the harmonic locking failure mode on the same specimen image. . 53
3.15 Pipeline for valley wavelength estimation from RGB microscope images via one-dimensional Fourier analysis. 54
3.16 Three-stage pipeline for training dataset construction. 58
3.17 Spatial alignment between RGB images and the ATOS Q Cloud Point. 59
3.18 Coordinate system adopted for the aligned ATOS Q mesh Cloud Point. 60
3.19 Architecture of the feature-driven MLP. 63
3.20 Architecture of the ResNet-18 baseline. 63
3.21 Architecture of the ResNetFusion multimodal model. 64

4.1 Cost landscape in the $(\mathrm{gap}_x, \mathrm{gap}_y)$ search space for the representative calibration acquisition of the reference specimen. The red star marks the optimal alignment adopted for the corresponding mesh. 68
4.2 Calibrated ATOS Q profiles versus SJ-210 reference, Strip 1. 69
4.3 Calibrated ATOS Q profiles versus SJ-210 reference, Strip 2. 69
4.4 Calibrated ATOS Q profiles versus SJ-210 reference, Strip 3. 69
4.5 Calibrated ATOS Q profiles versus SJ-210 reference, Strip 4. 70
4.6 Calibrated ATOS Q profiles versus SJ-210 reference, Strip 5. 70
4.7 Distribution of the typical spread of the extracted wavelength across photographs of the same face taken at the same normalised vertical position 74
4.8 Illustration of the peak strength S and the peak confidence C on two extreme photographs of the dataset. In each row, the left panel shows the original image and the right panel shows its FFT amplitude spectrum, with the dominant peak marked in red, the median of the search band shown as a dashed orange line, and the immediate spectral neighbourhood highlighted in red. 76
4.9 Correlation matrix computed at the level of the face over the final dataset. 78
4.10 Training and validation loss curves for a representative run of each architecture, all trained on the R_a target. The dashed grey vertical line marks the epoch of minimum validation loss, which is the epoch selected by early stopping. The validation loss reaches a plateau in all three cases. 79
4.11 Predicted versus measured roughness for the selected model of each parameter, over the pooled out-of-fold predictions. 81

List of Tables

1 Dissertation contribution to the UN Sustainable Development Goals. iii

2.1 Comparative overview of surface roughness assessment studies reviewed in this chapter. LC: laser cutting. SS: stainless steel. Al: aluminium. MS: mild steel. DED: directed energy deposition. DOE: design of experiments. RSM: response surface methodology. GA: genetic algorithm. ANN: artificial neural network. . . 31

3.1 Material properties and cutting parameters of the 17 scanned laser-cut specimens. 35
3.2 Technical specifications of the digital microscope. 38
3.3 Estimated camera parameters obtained during microscope calibration. 40
3.4 Technical specifications of the Mitutoyo SJ-210 surface roughness tester. 42
3.5 Technical specifications of the ATOS Q structured light scanner (MV 100 lens configuration). 43
3.6 Scan Point Computation parameters used in the *Zeiss Inspect 3D* software. 44
3.7 HSV threshold parameters for red laser stripe segmentation. 51
3.8 Composition of the final training dataset. 62

4.1 ATOS Q vs. Mitutoyo SJ-210 roughness parameters without positional alignment. 70
4.2 ATOS Q scan repeatability (standard deviation over four repeated scans), raw and calibrated roughness values, and SJ-210 reference ($k = 1.49$). 71
4.3 Roughness error metrics before and after applying the calibration factor $k = 1.49$ and positional alignment. 71
4.4 Mean wavelength variation over roughness 72
4.5 Distribution of the ratio between the naive and the pipeline wavelength across the dataset. 73
4.6 Per-image performance of each model and feature set, as the mean over the cross-validation folds and the three seeds. The MAE and RMSE are given in micrometres. The best value of each metric for each parameter is shown in bold. 79
4.7 Per-image performance of the specialised single-material models on aluminium, evaluated by leave-one-face-out. The image models are the mean over three seeds, and the MLP, reported for its Spectral and Process set, is a single reproducible run. The MAE and RMSE are in micrometres. 82

List of Acronyms

AI	Artificial Intelligence
AWJM	Abrasive Water Jet Machining
CCD	Charge-Coupled Device
CI	Change Index
CNN	Convolutional Neural Network
DTW	Dynamic Time Warping
FCNN	Fully Convolutional Neural Network
FDM	Fused Deposition Modelling
FFT	Fast Fourier Transform
FOV	Field of View
GLCM	Gray-Level Co-occurrence Matrix
GMM	Gaussian Mixture Model
HAZ	Heat-Affected Zone
HPS	Harmonic Product Spectrum
HSV	Hue, Saturation, Value
LBC	Laser Beam Cutting
LED	Light Emitting Diode
LOOCV	Leave-One-Out Cross-Validation
LSTM	Long Short-Term Memory
MAE	Mean Absolute Error
MAPE	Mean Absolute Percentage Error
MLP	Multilayer Perceptron
PBC	Plasma Beam Cutting
PCA	Principal Component Analysis
PSD	Power Spectral Density
RANSAC	Random Sample Consensus
ReLU	Rectified Linear Unit
ResNet	Residual Network
rFFT	real-valued Fast Fourier Transform
RGB	Red, Green, Blue

RMSE Root Mean Square Error
RSM Response Surface Methodology
SAM Surface Anisotropy Metric
SDG Sustainable Development Goals
SVD Singular Value Decomposition
VGG Visual Geometry Group

Chapter 1

Introduction

1.1 Motivation and Context

Laser cutting has become the reference technology for sheet metal transformation, combining high cutting speed, dimensional accuracy and the absence of mechanical tool wear. Fibre sources in particular have displaced CO_2 systems over most of the industrially relevant thickness range [2]. The economic value of the process, however, is not established by the speed at which a part is separated from the plate, but by the condition of the surface that is left behind. The cut edge is frequently a functional surface that is welded, assembled or exposed to the customer, so an edge with excessive roughness, with irregular striations or with adherent dross forces a secondary finishing operation, or leads to the rejection of the part altogether. In heavy metalworking, where the plates are thick and parts are large, the cost of detecting a defective cut only at the assembly stage is disproportionate to the cost of the cut itself.

Surface roughness is by a clear margin the most frequently studied of these indicators, since it condenses into a single quantity the thermal and gas dynamic conditions that prevailed at the cut front [2]. Roughness is also not a single value per part. It varies systematically along the thickness direction, being smaller near the beam entry and larger near the exit, where the melt film becomes unstable and the striations deepen [34, 53]. The reference measurement of these parameters is nonetheless still tactile. A stylus profilometer is dragged across the cut face along a prescribed evaluation length, which requires the part to be removed from the production flow, transported to a metrology bench, cleaned and carefully positioned, and which returns the roughness of a single line rather than of the surface. The procedure is accurate and traceable, but it is slow, it is operator-dependent, and it scales badly to large plates, so in practice it is applied to a small sample of the production or is replaced altogether by a visual and tactile judgement of the operator.

This dissertation was developed within an industrial collaboration with ACN Cutting Systems, a Portuguese seller of thermal cutting equipment, in the scope of the FiberNest XL project, carried out jointly with the Faculty of Engineering of the University of Porto. The project targets a large format fibre laser cutting machine for the heavy metalworking sector, and one of its innovation axes is precisely the introduction of an image-based system that monitors and certifies the quality

of the cut, replacing the manual and subjective inspection. The company therefore does not need a laboratory characterisation of the process. It needs a quality figure that can be obtained on the shop floor, on the part that has just been cut, quickly enough to be acted upon while the job is still running.

The application that motivates this work follows directly from that requirement. The intended outcome is a portable device that the operator places against the cut face and that returns an estimate of the surface roughness from a single Red, Green, Blue (RGB) image, without contact, without surface preparation and with little knowledge of how the part was produced. This deployment scenario is restrictive, and it shapes every methodological decision taken in the following chapters.

The literature has already established that this is plausible in principle. Convolutional networks have been shown to estimate R_z from images of laser-cut edges with useful accuracy [51, 53], and compact models operating on engineered intensity features have achieved comparable results with far fewer measurements [9, 33]. These works share, however, a set of restrictions that stand between them and the scenario described above. They are almost invariably confined to the mono-condition datasets that dominate the field, so the strong dependence of the striation signature on the material and on the plate thickness is never exercised. They rely on dedicated and carefully controlled illumination, whose geometry has been shown to be tightly coupled to the quality of the extracted features, so that a change in the acquisition setup degrades the accuracy unless a substantial amount of additional training data is collected [9]. The reference labels are obtained from line measurements, which the areal metrology literature has criticised as strongly position-dependent and insufficient to represent the topography of a laser-cut edge, proposing areal parameters such as S_a in their place [34]. The present work acknowledges this argument and is not blind to it. Because the reference topography is acquired here with a structured light scanner rather than with a stylus, the full surface is available, and areal parameters could in principle be computed from the same point cloud. The classical profile parameters R_a, R_q and R_z were nevertheless retained as the prediction targets, so that the results remain directly comparable with ISO 9013 and with the body of work reviewed in Chapter 2, and the areal characterisation is left as a natural extension of the same data.

Estimating roughness from an image of the cut edge is therefore neither an arbitrary nor an indirect choice. The striation pattern, the colouration and the burr morphology that are visible on a laser-cut surface are the visual manifestation of the same thermal and gas dynamic history that produced the roughness, which is what makes the mapping learnable at all [2]. What remains open is whether that mapping survives outside the controlled conditions in which it has been demonstrated, and whether the information that a network extracts from the raw image can be matched, or usefully complemented, by explicit and interpretable descriptors of the striation pattern. These two questions define the solution proposed in the following section.

1.2 Proposed Solution and Objectives

The solution pursued in this dissertation treats the estimation of surface roughness as a supervised regression problem, in which the input is an RGB image of the cut edge and the output is one of the profile roughness parameters R_a, R_q or R_z. The premise, established at the end of the previous section, is that the striation pattern left on the cut face is the visible trace of the same process history that determined the roughness, so a model trained on enough paired observations should be able to recover the second from the first. Turning that premise into a usable estimator, however, requires three things that the reviewed literature does not provide together, namely reference labels that describe the surface rather than a single line, a set of specimens broad enough for the model to be exercised across materials and thicknesses, and an evaluation protocol that reports what the model will actually do on a cutting condition it has never seen.

The experimental campaign was therefore built on seventeen laser-cut specimens supplied by the industrial partner, spanning four materials, thicknesses between 4 mm and 15 mm and seventeen distinct combinations of laser power, cutting speed, assist gas, gas pressure and focal position. Both faces of each specimen were analysed, and the resulting set covers a considerably wider portion of the process space than the single material and single thickness datasets that dominate the field. The counterpart of this breadth is that the specimens reflect production practice rather than a factorial design, which is discussed as a limitation in Chapter 3.

Three sensing modalities were combined over these specimens. RGB images were acquired with a low-cost digital microscope carried by the motion system of a Fused Deposition Modelling (FDM) printer, so that the sensor could be displaced over the cut face along a reproducible grid of positions and every image could be tagged with the coordinates at which it was taken. The reference topography was acquired with an ATOS Q structured light scanner, from which roughness profiles were extracted at controlled positions along the thickness direction. Because a structured light scanner is known to attenuate micro scale amplitudes relative to its lateral resolution, these estimates were not taken at face value, but were validated and calibrated against a Mitutoyo SJ-210 contact profilometer operated on the same bracket and over the same regions of the same specimen. The scanner therefore supplies the labels, and the stylus supplies the traceability that makes those labels defensible.

In parallel with the learned representation, an explicit and interpretable description of the striation pattern was extracted from the same images. A one-dimensional Fourier analysis of the intensity profile yields the wavelength of the dominant striations together with the spectral strength and a confidence measure of the corresponding peak, computed both over the whole image and over three horizontal strips, and made robust to harmonic locking through a Harmonic Product Spectrum (HPS). These descriptors were validated against an independent physical reference obtained by projecting a laser line onto the cut face and measuring the spacing of the illuminated grooves directly. Each image is finally associated with its normalised position within the plate thickness, which encodes whether it was taken near the smoother beam entry or the rougher exit.

Three architectures were then confronted on the resulting dataset. A Multilayer Perceptron (MLP) operating on the engineered features establishes the interpretable and data-efficient baseline, a Residual Network (ResNet)-18 pretrained on ImageNet regresses the roughness from the raw image alone, and a multimodal model fuses the image embedding with the tabular branch by late concatenation. Because images taken from the same face share the cutting condition and partially overlap, the models are validated by grouped cross-validation in which entire faces are held out, and the procedure is repeated over three random seeds, so that the reported performance corresponds to generalisation towards an entirely new cutting trial rather than towards a neighbouring photograph of a face already seen in training.

The general aim of this dissertation is thus to determine whether the surface roughness of a laser-cut edge can be estimated from a single RGB image with an accuracy that is useful for shop floor inspection, across several materials, thicknesses and cutting conditions compatible with the tolerances of ISO 9013. This aim is decomposed into the following specific objectives.

1. Establish a procedure for obtaining reference roughness parameters from structured light scanning of the cut face, and to validate and calibrate that procedure against contact profilometry, so that the labels used for training are quantitatively traceable.

2. Extract interpretable spectral descriptors of the striation pattern from RGB images of the cut edge, and to validate the estimated striation wavelength against an independent physical reference.

3. Construct a multimodal dataset that associates each image with its spectral descriptors, the cutting parameters of the corresponding specimen, its normalised position within the thickness and the roughness labels R_a, R_q and R_z.

4. Train and compare feature-driven, image-driven and multimodal regression models under an evaluation protocol grouped by cut face, quantifying how much of the roughness is recovered by the raw image and how much is added by the engineered descriptors and by the process context.

5. Assess whether specialising a model on a single material improves the estimate over a general model covering all materials, and to identify, on that basis, the configuration best suited to an inspection performed from a single acquisition.

1.3 Dissertation Structure

The remainder of this dissertation is organised into four chapters. Chapter 2 reviews the state of the art, starting from the fundamentals of the laser cutting process and the mechanisms through which the process parameters shape the condition of the cut edge, and proceeding to the methods that have been proposed for assessing that condition. Particular attention is given to the machine learning approaches applied to laser-cut edges, to the use of image-based spectral analysis for the characterisation of surface texture, and to non-contact three-dimensional measurement, since these three bodies of work correspond respectively to the models, to the engineered descriptors and to the reference labels adopted here. The chapter closes with the neural architectures and validation strategies that support the modelling stage.

Chapter 3 describes the materials and methods. It presents the laser-cut specimens and their cutting conditions, the three acquisition systems and the setups built to operate them, the validation and calibration of the structured light scanner against the contact profilometer, and the pipeline that extracts the striation wavelength and its associated spectral descriptors from the RGB images, together with its validation against the projected laser line. The chapter then details how the multimodal dataset is assembled and preprocessed, and specifies the three architectures, the feature sets and the training and validation protocol.

Chapter 4 reports the results. It first compares the roughness profiles obtained from the two instruments and establishes the calibration of the scanner, then confronts the spectral estimation of the striation period with the laser segmentation reference and examines the spatial stability of the wavelength and its relation to the roughness. The chapter concludes with the performance of the three models across the three roughness parameters and the four feature sets, the comparison between the general and the material-specific models, and the discussion of these results.

Chapter 5 summarises the main findings, revisits the objectives stated in Section 1.2, discusses the limitations of the work and outlines the developments required to advance from the present results towards the industrial device that motivated them.

Chapter 2

State of the Art and Fundamental Issues

Estimating the surface roughness of a laser-cut edge from an image draws on four bodies of knowledge that are usually kept apart. The first concerns the cutting process itself, since the roughness that is to be predicted is produced by the thermal and gas dynamic conditions at the cut front and inherits their dependence on the material, on the plate thickness and on the parameters chosen by the operator. The second concerns metrology, since roughness is not a single unambiguous quantity but a family of parameters defined over a filtered profile or over an area, each carrying its own conventions and its own standards. The third concerns the extraction of information from images of machined surfaces, whether through descriptors designed by hand or through representations learned from data. The fourth concerns the models themselves and the protocols used to establish that their reported performance is credible. This chapter reviews the four in turn and, in doing so, delimits what has already been demonstrated and what remains open.

The review begins in Section 2.1 with the fundamentals of the laser cutting process, covering the source technologies in industrial use, the parameters that govern the cut and the mechanisms through which those parameters shape the condition of the resulting edge. Particular attention is given to the role of the material and of the plate thickness, since both act as confounding variables for any method that attempts to infer quality from appearance alone. Section 2.2 then addresses how cut quality is formally defined and measured. It presents the geometrical characteristics and quality classes established by ISO 9013 for thermal cuts, the definition of the profile roughness parameters and the revision that the ISO 21920 series introduced when it superseded the standards on which that definition had rested, the filtering and scale selection rules that determine which wavelengths of the measured profile are retained as roughness, and the areal parameters that have been proposed as an alternative to describing a cut face by a single extracted line.

Section 2.3 surveys how the quality of laser-cut edges has been assessed in practice. The emphasis falls on approaches based on Artificial Intelligence (AI), which are examined both in the broader context of surface roughness prediction in machining and in the narrower set of studies that address laser-cut edges specifically, and the section closes with the alternative sensing strategies that have been proposed for the same purpose. Section 2.4 isolates one of these strategies, the description of surface texture through the spatial frequency content of its image, which is the

principle underlying the striation descriptors developed in Chapter 3. Section 2.5 then turns to the acquisition of reference topography without contact, comparing structured-light scanning with the tactile and optical instruments that constitute the traceable reference, and examining the resolution limits that determine whether such a system can be used as a source of roughness measurements at all.

Section 2.6 presents the modelling background, covering the convolutional and fully connected architectures employed in this work, the strategies available for combining heterogeneous inputs into a single predictor, and the validation procedures that determine whether a reported accuracy reflects genuine generalisation or the reuse of information across the partition. Section 2.7 closes the chapter by consolidating the reviewed work into a comparative overview and by stating the gaps that the present dissertation addresses.

2.1 Laser cutting process fundamentals

Any method that infers the quality of a cut edge from its appearance depends on a relationship that the cutting process itself establishes, since the visual signature of the surface and the roughness that is to be estimated from it are two consequences of the same thermal and gas dynamic history. This section reviews that process to the extent required to make the relationship explicit. It begins with the principle of operation and with the regimes under which the material is separated, proceeds to the laser sources in industrial use and to the parameters that govern the cut, and then examines how the cut surface and its characteristic striation pattern are actually formed. The section closes by considering the influence of the material and of the plate thickness, which are the two variables that most strongly condition whether a relationship learned on one set of specimens transfers to another.

2.1.1 Process principle and cutting regimes

Laser cutting is a thermal, non-contact machining process in which a focused, coherent laser beam locally melts, vaporises or chemically alters the workpiece material, with the resulting molten or vaporised material subsequently expelled from the kerf by a coaxial high-pressure assist gas jet [2]. The process is widely adopted in sheet metal manufacturing due to its combination of high cutting speed, dimensional accuracy and the absence of mechanical tool wear, which distinguishes it from competing non conventional methods such as Abrasive Water Jet Machining (AWJM) and Plasma Beam Cutting (PBC). Compared to these alternatives, Laser Beam Cutting (LBC) achieves a kerf width below 1 mm and a surface roughness in the range of 1 µm to 10 µm, against 2 µm to 6.5 µm for AWJM at a substantially higher operating cost and above 6.5 µm for PBC [2].

Although the beam always acts by delivering energy to a localised region of the workpiece, the way in which the material is separated and removed differs according to the role played by the assist gas, giving rise to distinct cutting regimes [52]. In melt and blow cutting, an inert gas such as nitrogen or argon is delivered at high pressure and the beam alone supplies the energy needed to melt the material, with the gas serving the purely mechanical function of ejecting the

melt downwards through the kerf. Since the gas does not participate chemically, the cut surface remains comparatively free of oxide, and nitrogen is reported to yield better surface quality than oxygen, which is why it is the assist gas most frequently adopted in the reviewed literature and the standard choice for stainless steel and aluminium alloys [2]. In melt burn and blow cutting, oxygen is used instead and reacts exothermically with the heated material. The reaction adds heat to that supplied by the beam, which raises the material removal rate and allows greater thicknesses to be processed at a given power level, but it also thickens the oxide layer on the cut face and produces a generally rougher and more discoloured surface [2]. A third regime, vaporisation cutting, removes material directly without a liquid phase and is largely confined to thin sheets and non metallic materials, so it falls outside the scope of the present work [52].

The practical consequence of this distinction is most visible in the assist gas pressure. In melt and blow cutting, the gas alone must expel the entire melt volume, so pressures of the order of ten bar or more are required, whereas in melt burn and blow the oxygen acts primarily as a reactant and only needs to sustain the reaction and clear its products, so pressures below one bar are typical. This contrast is apparent across the studies compiled in the reviewed literature, where oxygen assisted cuts are consistently reported at pressures below 1.1 bar and nitrogen assisted cuts at pressures between 8 bar and 16 bar [2]. The two regimes therefore differ not by a single parameter but by the physical mechanism through which the cut face is produced, and consequently by the texture, the colouration and the oxide condition of the surface that results. This matters directly for any method that infers quality from the appearance of the edge, since two surfaces produced under different regimes may exhibit comparable roughness while presenting markedly different visual signatures.

2.1.2 Laser sources and beam characteristics

Three laser source technologies dominate industrial sheet metal cutting, each with a distinct optical architecture, as illustrated in Figures 2.1-2.2. The CO_2 laser, historically the most established source, generates a beam at a wavelength of approximately 10.6 µm within a gas filled optical resonator and delivers it to the cutting head through a system of mirrors. It offers high beam quality and output power suitable for a wide thickness range, and remains the reference technology for thick section cutting [2]. Pulsed Neodymium-doped Yttrium Aluminium Garnet (Nd:YAG) lasers employ a solid state crystal rod as the gain medium, pumped by flash lamps. Their shorter wavelength of approximately 1.06 µm improves absorption in metallic materials, and their ability to deliver high peak power in short pulse durations makes them the preferred choice for thin sheets, micro welding and fine cutting [2]. Fibre lasers use an ytterbium doped optical fibre as the gain medium, pumped by laser diodes, and deliver the beam to the cutting head through a flexible fibre optic cable. Operating at a wavelength of approximately 1.07 µm, the fibre source produces a beam of superior spatial quality that can be focused more tightly than either CO_2 or Nd:YAG beams, resulting in a narrower kerf and higher energy density at the workpiece. Combined with an electrical to optical conversion efficiency of up to 30 %, against approximately 10 % for CO_2

sources, this makes the fibre laser the preferred technology across much of the industrially relevant thickness range [2].

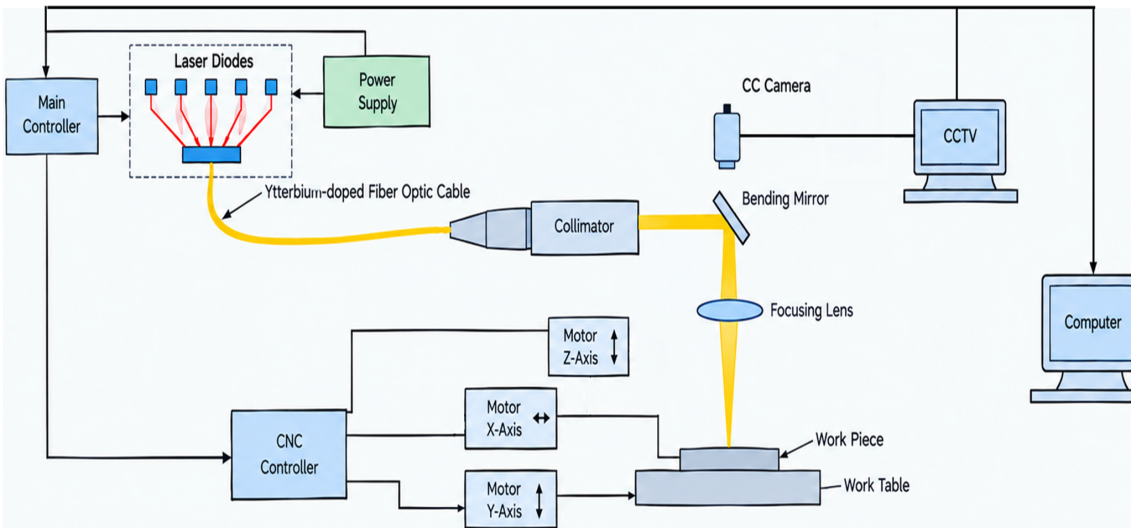


Figure 2.1: Schematic of a fibre laser, adapted from [2].

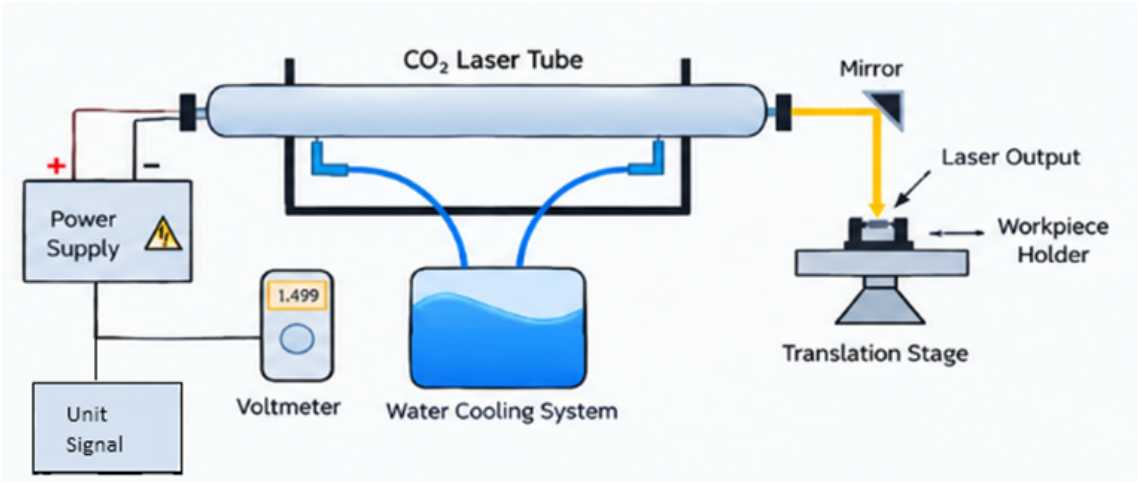


Figure 2.2: Schematic of a CO_2 laser, adapted from [2].

At lower power levels, the cut quality of fibre laser sources has been reported to degrade relative to CO_2 sources for thicknesses beyond approximately 4 mm, an effect attributed to reduced energy absorption at depth and to gas flow instabilities near the lower cut surface caused by the narrow kerf [2]. This limitation is, however, strongly dependent on available power, and high power fibre sources in the range of 8 kW to 12 kW have progressively extended the practical thickness range of fibre laser cutting well beyond this value in industrial practice [2].

The laser cutting parameter space is conventionally divided into beam parameters and process parameters [2]. Beam parameters describe the properties of the laser beam itself, including wavelength, power intensity, spot size, polarisation and, for pulsed sources, pulse duration and

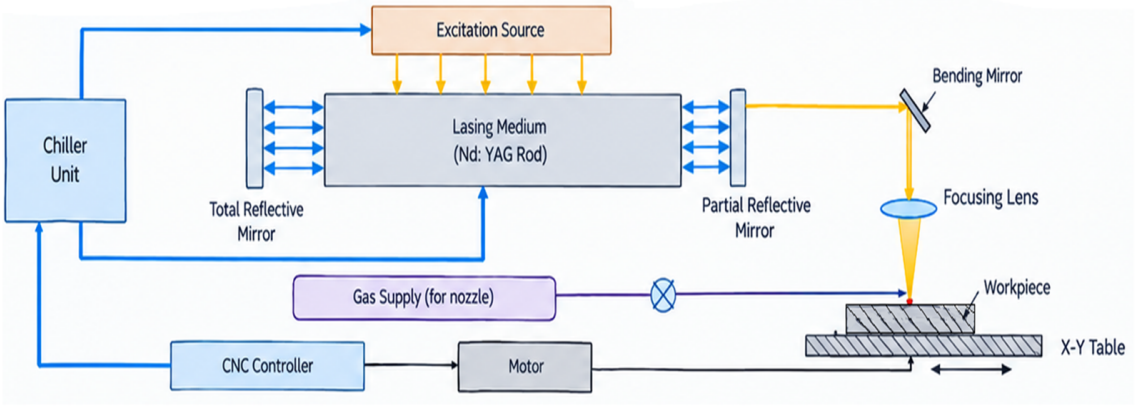


Figure 2.3: Schematic of a Nd:YAG pulsed laser, adapted from [2].

repetition frequency. Process parameters, which are more readily adjusted by the machine operator, comprise the focal position of the beam relative to the sheet surface, the assist gas type and pressure, the nozzle diameter, the stand-off distance between nozzle and workpiece, and the cutting speed. Among these, cutting speed, laser power and assist gas pressure are reported as the most frequently varied and controlled parameters across the literature, owing to their direct accessibility on industrial machines and their strong influence on the principal quality outputs [2].

2.1.3 Process parameters and quality outputs

The relationship between process parameters and cut quality outputs is broadly consistent across the reviewed materials, although its magnitude is material dependent. Surface roughness, kerf width, kerf taper, Heat-Affected Zone (HAZ) width and material removal rate are all proportional to laser power, pulse width and pulse frequency, and inversely proportional to cutting speed and standoff distance. Higher power combined with lower speed concentrates more thermal energy at the cut front, producing more melted material, a wider HAZ and rougher striations. Conversely, excessive cutting speed relative to power can also degrade roughness by distorting the striation pattern from a regular, vertical orientation to an inclined, irregular one. Assist gas pressure exerts a comparatively secondary but non-negligible influence, since it primarily governs the efficiency of molten material ejection from the kerf, with higher pressure generally improving surface finish by reducing the resolidified melt layer, though excessive pressure can introduce flow turbulence that degrades the cut. The choice of assist gas further modulates quality through its chemical interaction with the melt, as discussed in Section 2.1.1 [2].

This general behaviour is concretely illustrated by an experimental study on 10 mm thick S355J2 structural steel cut with a fibre laser using oxygen as the assist gas, in which cutting speed, beam power and gas pressure were varied according to a factorial design [31]. The results demonstrated that surface roughness R_{z5} is governed primarily by cutting speed, decreasing significantly as speed increases, whereas beam power and oxygen pressure did not reach statistical significance for this metric. Surface hardness and perpendicularity deviation, by contrast, proved

most sensitive to beam power, since higher power enlarges the HAZ, promoting martensitic or bainitic transformation that raises hardness while simultaneously degrading cut edge perpendicularity obtained for a sheet thickness above 5 mm, which is atypical in the laser cutting literature. These findings reinforce the conclusion that different process parameters dominate different quality metrics, underscoring the need for material and thickness specific characterisation.

2.1.4 Formation of the cut surface and striations

The characteristic appearance of a laser-cut edge, a set of approximately parallel grooves running from the beam entry surface towards the exit, is not an incidental blemish but the direct record of how the melt was removed from the kerf. Understanding this mechanism is central to the present work because it is precisely this pattern that an image-based quality assessment method observes, and because the mechanism explains why the roughness of a cut face is not uniform through the plate thickness.

During cutting, a thin film of molten material forms on the kerf front and is driven downwards by the shear exerted by the assist gas jet. As the cut rate increases, the beam rides ahead onto the unmelted material, raising the local power density on the unsloped surface and causing melting to proceed more rapidly. This material is swept down into the kerf as a step, and as the step is swept downward it leaves behind a mark on the cut edge, which is the striation [52]. The cause of striations is a subject of some dispute in the literature, and three mechanisms have been proposed to account for them. The step theory, just described, attributes the periodicity to the cyclic advance of the beam onto unmelted material. The critical droplet theory attributes it to the melt film pulsating in thickness before it can be blown free by the gas jet [52]. The hydrodynamic instability theory, proposed by Vicanek et al. [56] and subsequently supported by experimental observation of the melt flow dynamics at the kerf front [17], attributes the periodicity to an instability of the melt film under gas shear, which causes the lateral flow on the kerf sides to become discontinuous at intermediate cutting velocities while the central flow remains continuous. In reactive melt burn and blow cutting, a further mechanism is present. At speeds below the propagation speed of the burning reaction, the ignition temperature is reached and oxygen reacts outwards in all directions from the ignition point, introducing an additional and typically coarser striation pattern that is superimposed on the hydrodynamic one [52].

The cut face is not homogeneous through the plate thickness. Three distinct zones are generally observed, as illustrated in Figure 2.4. Near the beam entry, the direct penetration of the beam produces fine, nearly vertical striations and a comparatively smooth surface, although this region is also where the HAZ is widest, since the beam intensity and the thermal input to the surrounding material are highest at the top of the kerf [52]. Below this, in the shelf zone where most of the beam absorption occurs, the striations become more pronounced. In the washout region near the exit, the superheated melt flows out of the kerf, melting additional material on the way, and the combination of accumulated melt, reduced beam intensity and disturbed gas flow produces the deepest and most irregular grooves, together with the dross or burr that may adhere to the lower edge [52]. This depth dependence is empirically well established across materials, with surface

roughness increasing systematically from the beam entry towards the exit and with increasing plate thickness [2]. The depth position at which an image is taken is therefore not a nuisance variable but a determinant of the roughness value that will be observed there, which is why the normalised position within the plate thickness is included as an explicit input in the modelling approach described in Chapter 3.

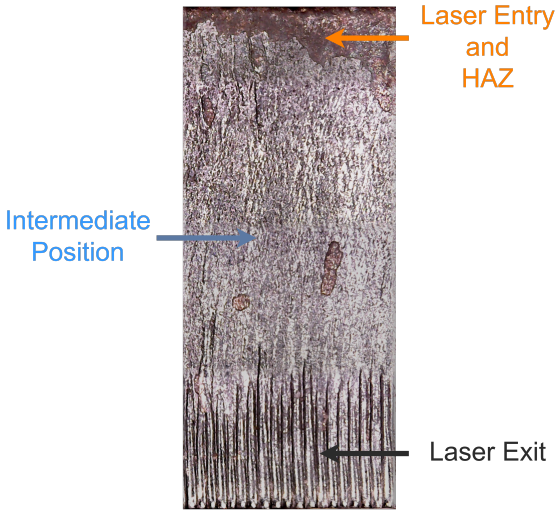


Figure 2.4: Laser-cut edge showing the three characteristic zones of the cut face: Beam entry, shelf and washout zones.

2.1.5 Influence of material and thickness

Material properties introduce a further layer of complexity that motivates material-specific process characterisation and, consequently, material-specific quality models. Titanium alloys combine high strength with low thermal conductivity, restricting heat dissipation and increasing the sensitivity of cut quality to small parameter deviations. Steel alloys exhibit high dynamic shear strength and a tendency towards accelerated thermal degradation adjacent to the cut, but otherwise respond predictably to the power and speed trade-off described in Section 2.1.3. Aluminium alloys present the most distinct behaviour. Their thermal conductivity is approximately four times that of steel, causing rapid heat dissipation that can prevent full depth melting, while their high reflectivity reduces beam energy absorption and requires proportionally higher laser power for equivalent penetration. The formation of a high melting point aluminium oxide layer further compounds these thermal isolation effects, distinguishing the relationship between parameters and quality for aluminium from that observed for ferrous alloys [2].

Increasing sheet thickness systematically elevates both surface roughness and kerf taper across all three material families, since achieving full penetration at greater depths requires higher power delivered at lower speed, concentrating additional thermal energy at the cut point. The review by Alsaadawy documents this trend across titanium, steel and aluminium alloys, confirming thickness as a primary confounding variable alongside material type [2]. The experimental study by Latka et al. on 10 mm S355J2 structural steel further illustrates the material and thickness dependence,

showing that the dominant parameter for surface roughness, cutting speed, differs from those governing hardness and perpendicularity, which are most sensitive to beam power [31]. These findings underscore that no single parameter setting optimises all quality metrics simultaneously, and that the relative importance of each process variable shifts with material type and plate thickness.

2.1.6 Implications for image-based assessment

The process fundamentals reviewed in the preceding subsections establish the physical basis for the relationship that image-based quality assessment methods seek to exploit. The striation pattern, the oxide colouration and the burr morphology visible on a laser-cut edge are direct visual manifestations of the thermal and gas dynamic conditions that prevailed at the cut front, themselves governed by laser power, cutting speed, assist gas pressure and focal position. This mechanistic link is what makes the inference of surface roughness from an image of the cut edge plausible in principle, since the same process history that determined the roughness also shaped the visual texture that the image records.

The review also sets the terms on which such a method must be judged. A model that observes only the appearance of the edge inherits the dependence of that appearance on the cutting regime, the material and the position within the cut. The difference in visual signature between a melt and blow surface, which is metallic and oxide-free, and a melt burn and blow surface, which carries a discoloured oxide layer, means that the same roughness value can correspond to markedly different image textures depending on the assist gas used. The systematic variation of roughness through the plate thickness, from smooth near the beam entry to rough near the exit, means that the depth position of an acquisition is as informative as the texture itself. And the strong dependence of the striation pattern on the material, documented across titanium, steel and aluminium alloys, means that a model trained on one material cannot be expected to transfer to another without explicit accommodation of the material as a variable.

2.2 Cut quality and surface roughness characterisation

Surface roughness is a consequence of the cutting process rather than an independent property of the material, and the standards that govern it reflect this by defining not only the measurement parameters but also the tolerances within which a cut is considered acceptable. This section introduces the normative framework for thermal cut quality, the profile roughness parameters most commonly adopted in the literature, and the filtering conventions on which their calculation depends.

2.2.1 Thermal cut quality - ISO 9013

The quality of thermally cut surfaces is governed by ISO 9013:2017 [24], which specifies geometrical product specifications and quality tolerances for thermal cuts in materials suitable for oxyfuel

flame cutting, plasma cutting and laser cutting [24]. For laser cutting, the standard applies to thicknesses between 0.5 mm and 32 mm. The quality of a cut surface is described by two mandatory characteristic values and two optional ones. The mandatory indicators are the perpendicularity or angularity tolerance, u, and the mean height of the profile, R_{z5}. The optional indicators are the drag, n, and the melting of the top edge, r.

Before defining the individual quality indicators, it is useful to introduce the geometric terminology adopted throughout the standard. Figure 2.5 and 2.6 illustrates the terms related to the cutting process and on the cut work piece.

In Figure 2.5, the numbered items identify the elements of the cut geometry: torch (1), nozzle (2), beam (3), kerf(4), start of cut (5) and the end of the cut (6). The lettered dimensions correspond to the workpiece thickness (a), the nozzle distance (b), the advance direction (c), the top kerf width (d), the cut thickness (e), the cut length (f), the bottom kerf width (g) and the cutting direction (h).

In Figure 2.6, the numbered items denote different elements, namely the upper edge of cut (1), the cut surface (2) and the lower edge of cut (3). The work piece thickness is denoted by (a), the cut thickness by (i), the depth of root face by (j), the length of cut (f) and (α) is the torch set angle and (β) is the cut angle.

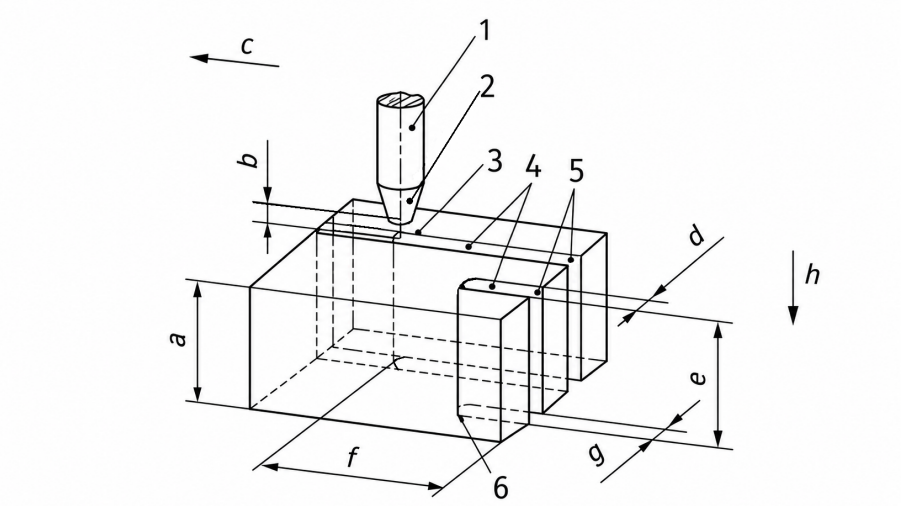


Figure 2.5: Cut geometry and reference dimensions, adapted from [24].

The perpendicularity or angularity tolerance, u, quantifies the deviation of the cut face from the ideal perpendicular or angled plane. As shown in Figure 2.7, it is measured within a restricted area of the cut surface, reduced from the upper and lower edges by a dimension Δa that accounts for the melting of the top edge, and it corresponds to the horizontal distance between the extreme positions of the cut face profile within that area. The standard defines five quality ranges for u, expressed as linear functions of the cut thickness a, ranging from a minimum tolerance applicable to precision cuts to a maximum corresponding to coarse industrial applications [24].

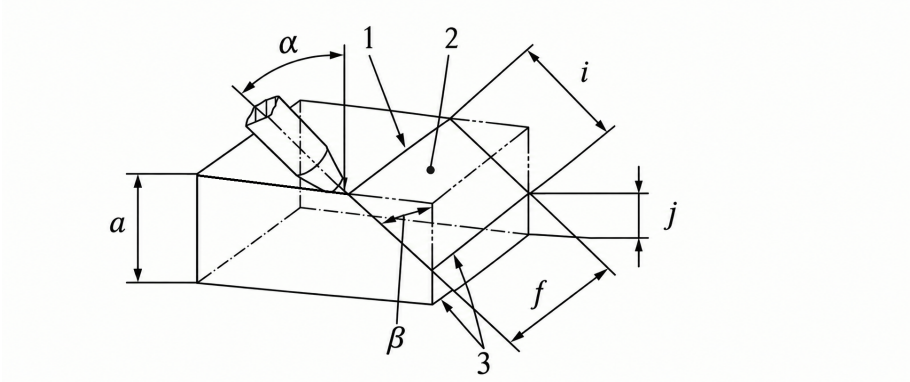


Figure 2.6: Angular characteristics of the cut face, adapted from [24].

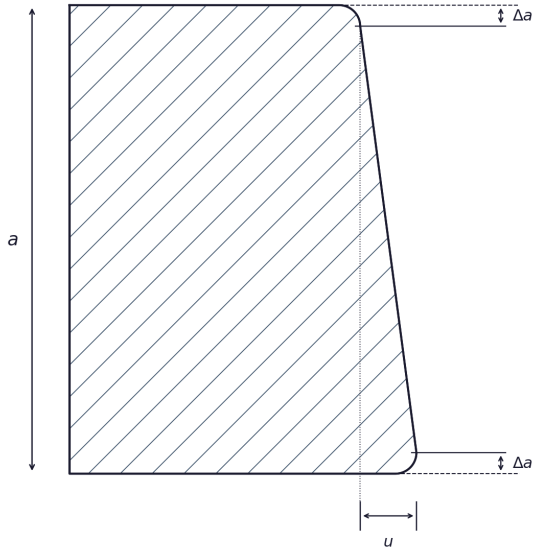


Figure 2.7: Cross-sectional view of a laser-cut edge illustrating the perpendicularity or angularity tolerance u and the excluded zones Δa. Adapted from [24].

The mean height of the profile, R_{z5}, describes the surface roughness of the cut face. As illustrated in Figure 2.8, the evaluation length l_n is divided into five consecutive sampling lengths l_r, and the maximum peak-to-valley height Z_{ti} is determined within each. R_{z5} is the arithmetic mean of these five individual heights [24],

$$R_{z5} = \frac{1}{5}\sum_{i=1}^{5} Z_{ti}, \tag{2.1}$$

where the subscript 5 was explicitly introduced by the standard to distinguish this arithmetic mean from the maximum of the five individual values [24]. The measurement is carried out at the point of maximum surface roughness of the cut face, at 15 mm from the start of the cut in the advance

direction, in accordance with ISO 4288 [23]. Four quality ranges are defined for R_{z5}, expressed as linear functions of the cut thickness a, covering values from the fine finishes achievable under optimised conditions to the coarser surfaces produced at the limits of the process [24].

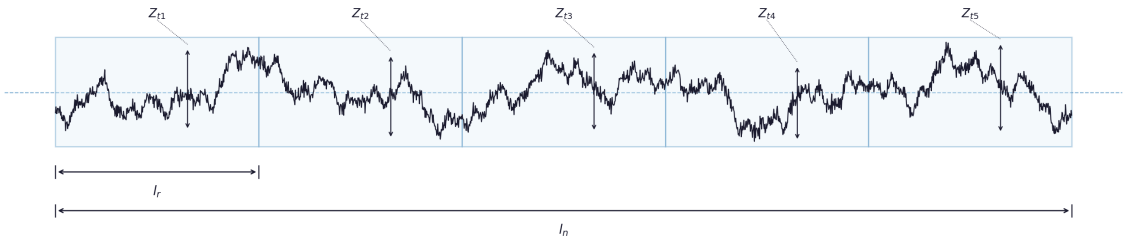


Figure 2.8: Surface roughness profile illustrating the calculation of the mean height of the profile R_{z5}. Adapted from [24].

The standard further requires that measurements be carried out on brushed, oxide-free cuts outside areas containing imperfections, and that the piercing area and the start and end sections of the cut be excluded from the evaluation. The optional indicator drag, n, quantifies the curvature of the striations in the cutting direction, which arises because the lower part of the kerf lags behind the upper part. The melting of the top edge, r, characterises the radius of rounding of the upper cut edge caused by thermal effects. Dross, defined as resolidified material adhering to the lower edge of the cut, may also be reported but is not included in the range classification, as illustrated in Figure 2.9.

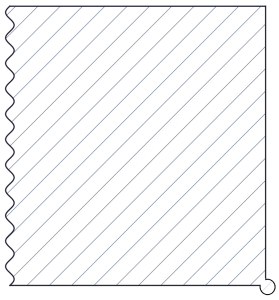


Figure 2.9: Dross formation on the lower edge of a laser-cut face. Adapted from [24].

The overall quality of a thermal cut is expressed as a pair of range values in the order u, R_{z5}, so that a cut classified as quality 2/2 satisfies both the perpendicularity and roughness tolerances of Range 2. Cuts that do not achieve the minimum specification of any range fall outside the scope of the standard [24].

2.2.2 Profile roughness parameters

The profile roughness parameters most widely adopted in the surface quality literature, R_a, R_q and R_z, are currently defined by ISO 21920-2:2021, which superseded ISO 4287:1997 [20]. The

transition between the two standards introduced a change in the calculation of R_a and R_q that is relevant in practice. Under ISO 4287, both parameters were computed within each sampling length individually and the reported value was the mean over the five sampling lengths. Under ISO 21920-2, they are computed directly over the entire evaluation length without subdivision [20].

The arithmetic mean deviation of the assessed profile, R_a, is defined as

$$R_a = \frac{1}{l_r} \int_0^{l_r} |Z(x)| \, \mathrm{d}x, \tag{2.2}$$

where $Z(x)$ is the height of the filtered profile at position x and l_r is the evaluation length. The root mean square deviation, R_q, is defined as

$$R_q = \sqrt{\frac{1}{l_r} \int_0^{l_r} Z^2(x) \, \mathrm{d}x}, \tag{2.3}$$

and weights larger deviations more heavily than R_a, making it more sensitive to occasional deep valleys or high peaks [20]. The maximum height of the profile, R_z, retains the structure of the older standard, being defined as the mean of the individual maximum peak-to-valley heights over the sampling lengths, in direct analogy with the R_{z5} of ISO 9013.

The three parameters capture complementary aspects of the surface profile. R_a is the most widely adopted indicator in the roughness prediction literature and provides an amplitude average that is robust to isolated outliers [11]. R_q is related to the root mean square of the height distribution and is relevant in functional contexts where energy dissipation or contact mechanics are of concern. R_z is sensitive to the extreme events in the profile, capturing the peak-to-valley range that governs whether a surface will fit against a mating part or whether a fatigue crack will initiate at a valley. Together they describe the average, the energy and the extreme excursion of the filtered profile.

2.2.3 Profile filtering and evaluation length

A measured surface profile is the superposition of roughness, waviness and form. Roughness parameters are computed not from the raw profile but from a filtered component obtained by removing the longer wavelength contributions. The Gaussian profile filter, specified in ISO 16610-21:2011, is the standard filter for this separation, attenuating contributions above a specified cut-off wavelength λ_c while preserving those below it with minimal phase distortion [19].

The cut-off wavelength and the evaluation length are selected as a function of the expected roughness amplitude according to the rules specified in ISO 4288. For R_z between 10 µm and 50 µm, a cut-off of 2.5 mm with an evaluation length of 12.5 mm is prescribed. For R_z above 50 µm, the recommended cut-off increases to 8 mm with a corresponding evaluation length of 40 mm. ISO 9013 explicitly prescribes that the measurement of R_{z5} be carried out in accordance with ISO 4288, linking the thermal cut quality standard directly to the profile filtering rules [24].

2.3 Assessment of laser-cut edges

The assessment of cut edge quality has been approached from two broad directions. The first, reviewed in Section 2.3.1, applies machine learning and computer vision to infer quality indicators directly from images of the edge or from signals recorded during the cut. The second, reviewed in Section 2.3.2, employs statistically grounded experimental methods, process coupled sensing and areal metrology to characterise or predict quality without relying on learned representations. Both directions are relevant to the present work, since the choice of roughness as the primary quality indicator, and the decision to pursue image based estimation, are motivated in part by what each approach has demonstrated and where it falls short.

2.3.1 AI-based quality assessment

The application of AI methods to surface roughness prediction in machining has evolved substantially over the past decade, progressing from traditional machine learning methods towards increasingly dominant deep learning approaches. Yang et al. provide a comprehensive taxonomy of this evolution, identifying three main methodological paradigms: machine learning combined with process parameters, machine learning combined with image processing, and end-to-end deep learning [57]. In parameter based methods, algorithms such as support vector machines and artificial neural networks establish mappings between cutting conditions and roughness outputs, offering stable predictions within specific process environments but exhibiting limited generalisation when facing variable surface textures or unseen machining conditions. Image based machine learning methods extract texture features from surface photographs, typically through operators such as the Gray-Level Co-occurrence Matrix (GLCM) or wavelet transforms, and correlate them with roughness values, offering non-contact measurement but remaining sensitive to illumination conditions and image quality. Deep learning methods, particularly Convolutional Neural Networks (CNNs), bypass manual feature engineering by learning high-level representations directly from raw image data, achieving average prediction accuracies above 94 % and relative errors below 0.1 µm in the broader machining literature, albeit at the cost of larger annotated datasets and greater computational demands. Prior to 2020, parameter based and image processing methods dominated, whereas from 2022 onwards deep learning has become the preferred approach, driven by its superior generalisation capabilities and the growing availability of image data in industrial settings [57].

A complementary perspective on this evolution is provided by Ghosh et al., who conducted a systematic review of 52 publications on AI-aided surface roughness evaluation in precision manufacturing from 2014 to 2024 [11]. The review confirms that CNNs and MLPs have emerged as the preferred modelling architectures since 2020, accounting for 61 % of recent publications, while R_a remains the most widely adopted roughness target in more than 85 % of works. Critically, the review identifies laser cutting as under-represented among the manufacturing processes studied, and highlights multimodal sensor fusion, combining surface images with process parameters and additional derived features, as an insufficiently explored research direction [11].

Within the narrower literature on laser-cut edges specifically, Tatzel and León demonstrated that the surface roughness R_z of stainless steel edges could be estimated directly from RGB images using a CNN, without recourse to contact or optical profilometry [53]. Their dataset comprised 3336 edges cut from 3 mm stainless steel on a fibre laser, generated by fully factorial variation of four dominant process parameters. A lightweight CNN inspired by the Visual Geometry Group (VGG)-16 architecture but trained from scratch was applied to 224×224 pixel patches obtained by dividing each edge image into 15 non-overlapping regions. Nine regression targets were defined, corresponding to R_z measurements at different positions across the sheet thickness. The model achieved Mean Absolute Errors (MAEs) ranging from 3.3 µm at the topmost line, where reflectivity artefacts affected the ground truth, to 4.8 µm at the lowest position, highlighting both the potential of the image-based approach and the practical difficulty of obtaining reliable reference measurements on highly specular metallic surfaces [53].

Building directly on this foundation, Tatzel et al. extended the scope of CNN-based inference to the inverse problem of recovering cutting parameters from a single edge image [54]. Using the same dataset and acquisition setup, a custom CNN simultaneously regressed five outputs: feed rate, gas pressure, two nozzle distances and focus position. RGB images outperformed grayscale and depth-augmented inputs, with the network achieving an MAE of 1.1 $\mathrm{m\,min^{-1}}$ for feed rate and 0.2 mm for focus position. A comparative expert survey showed that experienced human operators could not solve the same regression task reliably, achieving accuracy near the random baseline of 33 %, which underscores the complexity of the process appearance relationship and the unique capacity of deep learning to exploit it. Layer-wise relevance propagation further revealed that striation morphology is particularly informative for focus position, while coloured regions near the burr are more relevant for gas pressure [54].

Stahl et al. subsequently presented a more comprehensive quality assessment [51], their methodology combined a relatively simple RGB image-acquisition system with incident and transmitted illumination to estimate four quality parameters: average roughness depth $Rz5$, edge-slope tolerance, groove tracking and burr height. The dataset comprised 3,140 cut edges obtained from 785 stainless-steel samples produced using systematically varied cutting parameters, while reference measurements were acquired with a 3D optical profilometer according to DIN EN ISO 9013. The VGG16 and Xception networks were adapted for multi-output regression and trained using transfer learning, hyperparameter optimisation, and cross-validation. The best-performing configuration achieved an average coefficient of determination of $R^2 = 0.87$, with individual values of 0.86 for roughness, 0.79% for edge-slope tolerance, 87.7% for groove tracking and 0.96 for burr height. These results demonstrate the potential of convolutional neural networks to provide rapid, non-contact and cost-effective quality assessment of laser-cut edges, although the authors highlighted limitations related to imbalanced data distributions, uncertainty in the reference roughness measurements and the tendency of the models to underestimate extreme values.

García et al. took a complementary direction by investigating whether hand-crafted intensity descriptors of the striation pattern could predict roughness with fewer labelled examples than a CNN requires [9]. Working on stainless steel sheets, the authors extracted features from the in-

tensity distribution of edge images and trained compact regression models, achieving competitive accuracy. A key finding was that the illumination geometry is tightly coupled to the quality of the extracted features, so that a change in the acquisition setup substantially degrades predictive accuracy unless additional calibration data are collected.

Levichev et al. extended this feature-based approach to a broader experimental validation, demonstrating that machine learning models trained on image-derived descriptors can quantify roughness in laser cutting across multiple process conditions [33]. The study confirmed that compact models operating on handcrafted features remain competitive with deeper architectures when the training set is small, reinforcing the case for approaches that combine interpretable descriptors with learned representations.

Kristijan et al. addressed the related problem of process parameter optimisation for stainless steel laser cutting, using a neural network trained on a Latin hypercube design of experiments to predict surface roughness and kerf geometry as a function of cutting speed, laser power, gas pressure and focal position [48]. While not primarily an image-based method, this work illustrates the broader interest in data-driven approaches for quality control and parameter selection in laser cutting, and documents the roughness ranges achievable on stainless steel under varied process conditions.

2.3.2 Other assessment methods

Beyond AI-based approaches, the literature documents a range of complementary methods for predicting and characterising cut quality, spanning statistical modelling, process-coupled sensing and areal metrology.

Response Surface Methodology (RSM) has been widely adopted as a statistically grounded framework for relating process parameters to cut quality outputs. Shamlooei et al. applied RSM together with a genetic algorithm to optimise R_z in fibre laser cutting of 12 mm thick mild steel, considering six input parameters and finding that nozzle diameter, standoff distance, gas pressure and focus position were the most influential factors, with a prediction error of 4.79 % against experimental validation [47]. Rodić et al. extended this paradigm by integrating RSM with a regression tree in a hybrid modelling strategy for R_a prediction in laser cutting of EN 10130 steel [43]. The RSM component achieved an R^2 of 0.82, while the hybrid model improved this to 0.89 with a 20.85 % reduction in Root Mean Square Error (RMSE), demonstrating that combining interpretable statistical models with machine learning corrections can meaningfully enhance predictive accuracy even with limited datasets.

A distinct approach to in-process quality monitoring was proposed by Santolini et al., who addressed the problem of detecting cut interruption in industrial laser tube cutting machines in near real time [46]. Rather than characterising the post-process surface, the method fused multimodal process parameters, including laser power, feed rate, gas pressure and the light intensity collected through the optical fibre, and classified each cut as good, bad or missed using Gaussian Mixture Model (GMM), a Long Short-Term Memory (LSTM) network and a Fully Convolutional Neural Network (FCNN). On the publicly released ALCIDE dataset of 20 tube pieces, the FCNN

achieved an average classification accuracy of 76.65 %, outperforming the GMM and the LSTM. Although the task is binary classification of cut continuity rather than quantitative roughness estimation, the work establishes the feasibility of near online quality monitoring from machine embedded sensors, and the multimodal fusion strategy it proposes is directly related to the combined image and process parameter approach adopted in this dissertation.

A further distinct approach exploits workpiece vibration signals as quality indicators. Venkata Rao et al. instrumented stainless steel sheets with an accelerometer during CO_2 laser beam machining and developed mathematical models to predict surface roughness, kerf width and material removal rate from the amplitude and frequency of sheet vibration [55]. The vibration based models achieved an RMSE of 0.615 µm for R_a, demonstrating that process coupled vibration signals carry sufficient information to serve as indirect quality indicators without requiring post-process surface measurement.

A complementary perspective on quality characterisation is offered by Librera et al., who argued that the two-dimensional roughness parameters R_a and R_{z5} measured along a single line as prescribed by ISO 9013 are insufficient to capture the full surface topography of laser-cut edges [34]. Using focus variation microscopy to acquire three-dimensional surface data from CO_2 and fibre laser cuts of AISI 304 stainless steel at thicknesses of 6 mm and 10 mm, the authors showed that two-dimensional measurements are strongly position dependent along the thickness direction and fail to reproduce the visual differences between the two laser technologies. Areal roughness parameters S_a and S_{10z}, computed with appropriately selected cut-off lengths in both the cutting and thickness directions, were proposed as more representative descriptors, a suggestion that this dissertation acknowledges while retaining profile parameters as prediction targets for the reasons discussed in Chapter 3.

Taken together, these studies illustrate that the prediction and characterisation of laser-cut quality remains a multi-faceted problem. RSM based and hybrid statistical models provide interpretable frameworks for parameter optimisation, near online sensor fusion enables cut interruption detection without surface imaging, vibration signals offer a non-contact route to in-process quality inference, and areal metrology highlights the limitations of the linear metrics that most modelling studies use as their target outputs.

2.4 Image-based spectral analysis of surface texture

The surface texture of a machined part is the spatial superposition of the kinematic tool marks imposed by the cutting geometry, the dynamic disturbances introduced by vibrations and tool wear, and the waviness inherited from the machine tool structure. All of these components leave imprints at characteristic spatial frequencies, which means that the frequency content of a surface image can, in principle, reveal information about process stability and surface quality without requiring contact or interferometric measurement. This principle motivates the spectral approach to surface characterisation reviewed here.

A representative image-based frequency-analysis workflow was presented by Söderberg et al. for end-milled spheroidal graphite iron surfaces imaged with a Charge-Coupled Device (CCD) digital microscope [49]. Their procedure comprised three main steps. First, a one-dimensional intensity profile was extracted from each image by row-averaging across a central band of interest, which reduces random intensity fluctuations while preserving the periodic structure of the machining marks. Second, the detrended profile was transformed to the frequency domain using the Fast Fourier Transform (FFT), and the Power Spectral Density (PSD) was estimated using Welch's method with a Hann window and 50% overlap to reduce spectral leakage. Third, descriptors were derived from the resulting spectrum and from the intensity profile itself. These included the dominant spatial wavelength, obtained from the frequency of maximum amplitude in the magnitude spectrum, together with statistical moments of the intensity profile, namely the root mean square amplitude, the skewness and the kurtosis. In addition, two composite spectral indices were introduced: the Change Index (CI), defined as the ratio between the integrated PSD above $1.5 f_{\mathrm{TPF}}$ and the total spectral power, where f_{TPF} denotes the tooth-passing frequency; and the Surface Anisotropy Metric (SAM), which quantifies texture directionality from the histogram of image-gradient orientations computed using Sobel operators [49].

Söderberg et al. compared the image-derived metrics with calibrated White Light Interferometry measurements performed according to ISO 25178. They reported agreement in the direction of change of several image-based descriptors and the corresponding areal surface parameters across optimised and aggressive machining conditions, supporting their use as complementary indicators of relative surface-quality variation for production monitoring, rather than as direct substitutes for calibrated roughness parameters [49]. The CI increased substantially under more aggressive cutting conditions and showed a strong correlation with the machining parameters, indicating sensitivity to high-frequency spectral content associated with process disturbances. The SAM, in contrast, was largely independent of the cutting parameters and was interpreted as reflecting machine and setup geometry rather than process conditions [49].

Despite these promising results, Söderberg et al. also identified a key limitation of the dominant-wavelength estimator based on the arg max of the FFT spectrum. In their experimental configuration, the detected dominant wavelength was approximately 5.7–6.4 mm, whereas the theoretical feed-per-tooth spacing was only 0.04–0.065 mm. The spectral peak was therefore dominated by larger-scale texture components, which the authors associated primarily with form deviation or waviness rather than with the individual kinematic tool marks [49]. This observation shows that a single-peak search can be strongly influenced by low-frequency image content when the spatial resolution, field of view and contrast do not adequately resolve the fundamental striation period.

This limitation is directly relevant to the present work. In laser-cut edges, the striation pattern is similarly quasi-periodic and may be accompanied by substantial low-frequency intensity variations arising from reflectivity changes, edge geometry and illumination non-uniformity. Relying exclusively on the largest spectral peak can therefore lead to dominant-wavelength estimates that are inconsistent with the physical striation spacing. To address this issue, Chapter 3 introduces the Harmonic Product Spectrum (HPS) as a complementary estimator. Rather than selecting a

single peak, the HPS combines information from harmonic-related frequency components and enhances frequencies whose harmonics are consistently present, thereby improving robustness to the dominance of low-frequency components and to variability in the image formation process.

The broader family of transformed-domain texture descriptors includes wavelet transforms, which provide multi-resolution decomposition at different scales, and filter-based methods such as Gabor filters, which are selective in both spatial frequency and orientation. As surveyed by Yang et al., these descriptors remain sensitive to illumination conditions and image quality, with changes in the acquisition geometry capable of substantially degrading the extracted features [57]. This sensitivity, documented specifically for laser-cut edge imagery by García et al., who showed that the grazing angle and the relative orientation between the illumination direction and the striations strongly affect image contrast and roughness-prediction performance, motivates the flat-field illumination correction applied in the present work before any spectral extraction is performed [9].

2.5 Surface roughness measurement techniques

Surface roughness can be measured by a wide range of instruments that the ISO 25178 series classifies into three broad categories: line profiling methods, which measure a height profile along a single trace; areal topography methods, which capture the surface over a two-dimensional field; and area-integrating methods, which produce scalar quantities that depend on integrated surface properties [32]. The two families most relevant to surface roughness characterisation are line profiling and areal topography, and within each family both contact and non-contact implementations exist.

The contact stylus profilometer, the instrument assumed by ISO 21920 and its predecessor ISO 4287, drags a diamond-tipped stylus of defined radius across the surface while a transducer records its vertical displacement, producing a calibrated one-dimensional height profile from which any profile roughness parameter can be computed [32]. The method is accurate, traceable and independent of surface optical properties, but it is inherently limited to the single line swept by the stylus tip and requires the part to be removed from the production flow, positioned and carefully cleaned before measurement. Among non-contact line profiling methods, optical profilometers using laser triangulation or phase-shifting interferometry achieve comparable or superior vertical resolution without surface contact, but share the single-trace limitation and are sensitive to surface reflectivity and environmental disturbances [32]. Both contact and non-contact line profiling instruments are the standard tools assumed by surface roughness specification standards and are the instruments used as the traceable reference in the present work. Figure 2.10 illustrates the stylus principle.



Figure 2.10: Operating principle of a contact stylus profilometer.

Areal topography methods extend this capability to two dimensions. Structured-light projection, also known as fringe projection or phase-shift profilometry, is one such method and is explicitly classified within the areal topography category in ISO 25178 [32]. It projects a sinusoidal fringe pattern onto the surface and recovers the three-dimensional geometry from the phase deformation of that pattern as observed by one or more cameras, as illustrated in Figure 2.11. Dense point clouds covering areas of hundreds of square millimetres are captured in a single acquisition, at throughput rates far exceeding those of line scanning instruments, which makes the technique attractive for acquiring reference topography at a large number of positions across an extended surface without repositioning.

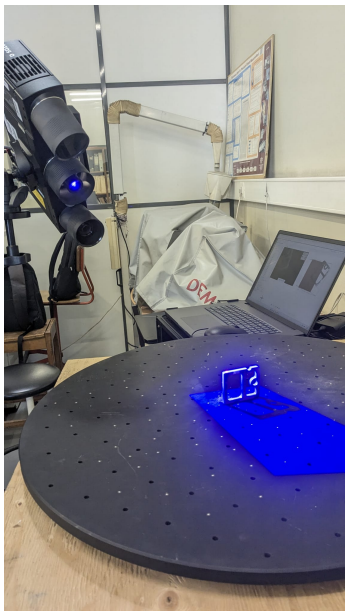


Figure 2.11: Fringe pattern projected onto a metallic surface during structured-light scanning.

Despite this throughput advantage, structured-light scanning is primarily a dimensional metrology tool and its use for roughness assessment is subject to a fundamental resolution constraint. Muratović et al. applied a GOM Scan 1 scanner with fringe projection and blue illumination to evaluate the surface topography of polymer gears, reporting a nominal point distance of 0.06 mm over measuring areas of 200 mm × 125 mm with repeatability of 0.02 mm to 0.03 mm [38]. The study identified that the lateral resolution and optical spot size of the scanner average the features of small surfaces and smooth out peak heights and valley depths that fall below the instrument's resolution, leading to a systematic attenuation of micro-scale roughness parameters [38]. Validation against stylus profilometry as a traceable contact reference was recommended.

2.6 AI models and evaluation

This section introduces the relevant architectural principles, the transfer learning strategy that underpins CNN-based approaches in small-dataset settings, the multimodal fusion paradigm, and the evaluation protocol that determines whether reported performance reflects genuine generalisation.

2.6.1 Convolutional neural networks

CNNs learn hierarchical feature representations directly from raw image data through alternating convolutional and pooling operations, extracting local spatial patterns at progressively larger receptive fields without requiring manual feature engineering [11, 29]. Their principal advantage is the ability to learn texture, edge orientation and spatial structure directly from data, making them particularly suited to image-based inspection of machined surfaces [11, 44]. CNNs have been applied to surface roughness regression, quality classification and defect detection in manufacturing,

achieving average prediction accuracies consistently above those of hand-crafted feature methods in settings where sufficient training data are available [11, 29].

The architecture employed in this dissertation is ResNet-18, introduced by He et al. [16]. The key innovation is the shortcut connection, which adds the block input directly to its output according to

$$\mathbf{y} = \mathscr{F}(\mathbf{x}, \{W_i\}) + \mathbf{x}, \tag{2.4}$$

where $\mathbf{x}$ is the block input, $\mathbf{y}$ is the block output and $\mathscr{F}(\mathbf{x}, \{W_i\})$ is the residual mapping learned by the stacked convolutional layers. This reformulation addresses the degradation problem observed in very deep plain networks, where adding more layers can increase rather than reduce training error, by making it easier for the solver to learn near-identity mappings when those are optimal [16]. The shortcut connections add neither extra parameters nor computational complexity. ResNet-18 consists of one initial 7×7 convolutional layer followed by four residual stages, each comprising two 3×3 convolutional building blocks, global average pooling and a fully connected output layer, for a total of 18 trainable layers. Batch normalisation [25] is applied after each convolutional operation, stabilising the distribution of activations and accelerating convergence by reducing internal covariate shift.

2.6.2 Transfer learning and fine-tuning

Training a deep CNN from random initialisation requires large annotated datasets that are difficult to obtain in industrial inspection scenarios. Transfer learning addresses this constraint by initialising the network with weights pre-trained on a large general-purpose dataset and then adapting them to the target task [58].

Yosinski et al. quantified the layer-by-layer generality of CNN features trained on ImageNet [6], showing that the first layers learn general low-level descriptors such as Gabor-like filters and colour blobs that transfer well across tasks, while the last layers learn task-specific representations [58]. A key finding is that initialising a network with transferred features and then fine-tuning them on the target dataset boosts generalisation performance even after substantial retraining, because the network has been preconditioned to represent natural image structure [58]. The choice of how many layers to fine-tune depends on the size of the target dataset and the similarity between the source and target domains. When the target dataset is small and the domain is dissimilar, freezing the lower layers and fine-tuning only the upper layers reduces the risk of overfitting, while still benefiting from the general low-level representations encoded in the earlier layers.

2.6.3 Multilayer perceptrons

A MLP is a feedforward neural network that maps an input vector to an output through one or more hidden layers, each applying a linear transformation followed by a non-linear activation function [15]. For an input $\mathbf{x}$, the output of the l-th hidden layer is

$$\mathbf{h}^{(l)} = \sigma\left(W^{(l)}\mathbf{h}^{(l-1)} + \mathbf{b}^{(l)}\right), \tag{2.5}$$

where $W^{(l)}$ and $\mathbf{b}^{(l)}$ are the weight matrix and bias vector of layer l, σ is a non-linear activation function such as the Rectified Linear Unit (ReLU), and $\mathbf{h}^{(0)} = \mathbf{x}$. MLPs are trained by back-propagating the gradient of a loss function with respect to the network weights [15]. The Adam optimiser [28] is widely adopted for this purpose, maintaining per-parameter adaptive learning rates based on first and second moment estimates of the gradients, providing faster and more stable convergence than vanilla stochastic gradient descent on problems with heterogeneous input scales.

MLPs are a natural choice when inputs are already expressed as compact numerical descriptors, since they can model nonlinear dependencies without the computational overhead of convolutional processing [11, 29]. Their limitation is that they do not exploit spatial structure and therefore depend entirely on the quality of the input representation, which makes them less effective than CNNs when rich spatial information remains embedded in raw images [11, 14].

2.6.4 Multimodal fusion

Multimodal architectures combine inputs from different sensing modalities into a single prediction, exploiting complementary information that no single modality captures alone [40]. In manufacturing applications, visual data and process metadata are naturally complementary: images capture the local surface texture, while numerical variables encode the operating conditions and spatial context that shape it [3, 29]. More recently, Ferreira et al. [7] proposed a generative joint embedding architecture for monitoring directed energy deposition, combining on-axis and off-axis imaging with process parameters through a contrastive alignment loss, which illustrates the growing interest in aligning multiple sensing modalities with process metadata for manufacturing monitoring. The fusion strategy most commonly adopted is late fusion by feature concatenation, in which each branch produces an embedding vector independently and the embeddings are concatenated before the final prediction head [40]. This design allows each branch to be pre-trained or fine-tuned independently and avoids the instabilities that can arise when gradients from very different input spaces are combined too early in the network.

2.6.5 Training and validation

The performance of a supervised model should be estimated using data that were not used during model fitting. This requirement is particularly important for small datasets and for datasets containing correlated observations, such as multiple images or measurements obtained from the same specimen or experimental condition. In these cases, a naive random split may place highly similar or near-duplicate observations in both the training and test sets, resulting in an overly optimistic estimate of generalisation performance. This situation is commonly discussed as data leakage, namely the unintended use of information during training that would not be legitimately available when predicting new observations [27].

Cross-validation is a commonly used approach for estimating generalisation performance when the available dataset is limited. It allows the data to be reused for both training and validation while maintaining a separation between the observations used to fit the model and those used to evaluate each fold. The choice of validation strategy involves a compromise between bias, variance and computational cost [1, 30]. Leave-One-Out Cross-Validation (LOOCV) uses nearly all observations for training in each iteration and can provide estimates with relatively low bias. However, it may also exhibit high variance and is not necessarily the most reliable strategy for model selection [30]. In comparison, k-fold cross-validation generally provides a more practical compromise. Kohavi et al. found that, for the datasets and classifiers investigated, stratified ten-fold cross-validation provided a favourable balance for model selection, whereas small values of k could increase bias and larger values, including leave-one-out, could produce less stable estimates [30]. More recent analyses also show that the choice of k affects the bias–variance behaviour of the estimate and should be considered in relation to the dataset size and model complexity [1].

Repeating k-fold cross-validation with different partitions can reduce the dependence of the reported result on a single random split and allows the variability between repetitions to be quantified. However, repetition alone does not prevent leakage if correlated observations are distributed across the training and validation folds. When several observations originate from the same specimen, cutting condition, manufacturing batch or spatial region, the partition should therefore be performed at the level of the corresponding experimental unit rather than at the individual image or measurement level. In this way, all observations belonging to a common source remain in the same fold, providing a more realistic estimate of the model's ability to generalise to previously unseen experimental units [27].

2.7 Summary and research gap

This chapter has reviewed the four bodies of knowledge on which the present dissertation draws. The laser cutting process, examined in Section 2.1, establishes that the striation pattern, the oxide colouration and the burr morphology visible on a cut edge are direct visual manifestations of the thermal and gas dynamic conditions at the cut front, which explains why image-based roughness estimation is physically plausible. The dependence of these features on the cutting regime, the material and the position within the plate thickness sets the terms on which any such method must be evaluated. Section 2.2 introduced the normative framework that defines what roughness means in this context. ISO 9013 specifies the quality of thermal cuts through two mandatory indicators, u and R_{z5}, with the latter requiring an extended traverse of the cut face that is incompatible with a single image acquisition. The profile roughness parameters R_a, R_q and R_z defined by ISO 21920-2 provide computable alternatives that together describe the average, the energy and the extreme excursion of the filtered profile.

Section 2.3 surveyed the methods that have been applied to assess the quality of laser-cut edges, with particular attention to the image-based and machine learning approaches that are most

directly related to the present work. Table 2.1 consolidates the principal studies in this area, together with selected works from the broader machining roughness prediction literature that provide methodological context. The patterns in the table make the research gap visible without further argument. Among the six studies directly addressing laser-cut edges, four are confined to a single material and a single thickness, invariably 3 mm to 12 mm stainless steel. Only Levichev et al. exercise more than two materials, and only Kristijan et al. varies thickness within a single material. None combines more than one material across multiple thicknesses. The broader machining literature shares the single-material restriction but has begun to explore multimodal fusion and generalisation across process conditions, directions that remain largely unexplored in the laser cutting domain.

Section 2.4 reviewed image-based spectral analysis as an alternative to end-to-end learned representations. The Söderberg framework showed that one-dimensional Fourier analysis of intensity profiles can recover process-relevant spectral signatures from low-cost microscope images and that these signatures correlate with independently measured areal roughness, validating the approach against a traceable reference [49]. The critical limitation identified was that the dominant wavelength estimated by the arg-max of the power spectrum locks onto low-frequency form and waviness components rather than the kinematic striation period, motivating the use of a harmonic product spectrum for robust fundamental frequency estimation.

Section 2.5 addressed the acquisition of reference topography. Structured-light scanning was identified as the only practically available instrument that can produce dense areal coverage of a laser-cut face at the throughput required to label a training dataset, but its lateral resolution introduces a systematic attenuation of micro-scale roughness amplitudes that must be corrected against a traceable contact reference before the data can be used as training labels.

Section 2.6 introduced the architectural building blocks of the modelling approach. The residual learning framework of ResNet enables stable training at depths that would cause plain networks to degrade, transfer learning from ImageNet reduces the data requirement for the image branch, and grouped cross-validation with face-level partitioning is required to prevent the leakage that would arise from placing overlapping images of the same face on both sides of the partition.

Taken together, the reviewed literature leaves four gaps that the present dissertation addresses. The first is the absence of datasets covering more than two materials and more than one thickness at a time, which limits the ability of existing models to generalise across the process space. The second is the lack of reference labels acquired from areal measurements rather than single-line traverses, which leaves open the question of how representative the training targets are of the actual surface. The third is the absence of a validated and interpretable spectral descriptor of the striation pattern that is robust to the dominant low-frequency content of the raw Fourier spectrum. The fourth is the absence of a leakage-safe evaluation that holds out entire cutting conditions from training, which means that the reported performance figures in the existing literature may overstate generalisation to truly unseen cuts.

Table 2.1: Comparative overview of surface roughness assessment studies reviewed in this chapter. LC: laser cutting. SS: stainless steel. Al: aluminium. MS: mild steel. DED: directed energy deposition. DOE: design of experiments. RSM: response surface methodology. GA: genetic algorithm. ANN: artificial neural network.

Reference	Method	Material	Thickness (mm)	Target	Best result
Laser cutting — AI-based					
Tatzel and León [53]	CNN (VGG-inspired, from scratch)	SS 1.4301	3	R_z	MAE 3.3 µm to 4.8 µm across nine depth positions
Tatzel et al. [54]	CNN (RGB input)	SS 1.4301	3	Process parameters	MAE 1.1 m min^{-1} (feed rate), 0.2 mm (focus)
García et al. [9]	MLP (image intensity features)	SS 304L	12	R_{z5}	$R^2 = 0.964$, relative error 12.2 %
Levichev et al. [33]	MLP (image intensity features)	SS, Al, MS	6	R_{z5}	Relative error 11 % to 19 % (Al and SS)
Stahl et al. [51]	CNN (VGG16, Xception)	SS	3	R_{z5}, u, n, b	$R^2 = 0.861$ for R_{z5}
Kristijan et al. [48]	ANN (process parameters)	SS	1–6	R_a, kerf width	$R^2 = 0.973$ for R_a
Laser cutting — other methods					
Latka et al. [31]	Factorial DOE	S355J2 steel	10	R_{z5}, u, hardness	Cutting speed identified as dominant parameter for R_{z5}
Shamlooei et al. [47]	RSM + GA	Mild steel	12	R_z	Prediction error 4.79 %
Rodić et al. [43]	RSM + regression tree	EN 10130 steel	—	R_a	$R^2 = 0.89$ (hybrid model)
Other machining processes — AI-based					
Ross et al. [44]	CNN, LSTM, MLP	Ti alloy	various	R_a	CNN best performer across architectures
Guo et al. [14]	CNN + data augmentation	Steel	various	R_a	Augmentation improves generalisation
Asadi et al. [3]	CNN + tabular (multimodal)	Various	—	Layer height, S_a	Multimodal outperforms single branch
Reviews					
Ghosh et al. [11]	Systematic review (52 papers)	Various	various	R_a (majority)	CNN and MLP dominant since 2020, LC under-represented
Ko and Yin [29]	Literature review	Various	various	R_a	Deep learning outperforms traditional Machine Learning consistently

Chapter 3

Materials and Methods

This chapter describes the methodological pipeline adopted in this dissertation for the characterisation and prediction of surface roughness on laser-cut edges. The proposed approach integrates multiple sensing modalities, namely RGB imaging, surface roughness measurement and structured light scanning. Classical spectral analysis is applied to the RGB images to extract interpretable descriptors of the striation pattern, which serve both as standalone predictors and as inputs to the AI-based regression models developed in the final stage of the pipeline.

The overall workflow is structured as a sequence of interdependent stages, as illustrated in Figure 3.1. The pipeline begins with the definition of the specimen set and the collection of multimodal data through three independent acquisition systems. A cross-validation stage follows, in which the ATOS Q structured-light scanner serves as geometric ground truth for the roughness measurements obtained from the Mitutoyo SJ-210. Based on this validated dataset, the pipeline diverges into two main analytical branches. The first applies classical computer vision techniques directly to the raw RGB data, extracting surface-related features without any learning-based component. The second constructs a labelled dataset and employs AI-based classification models.

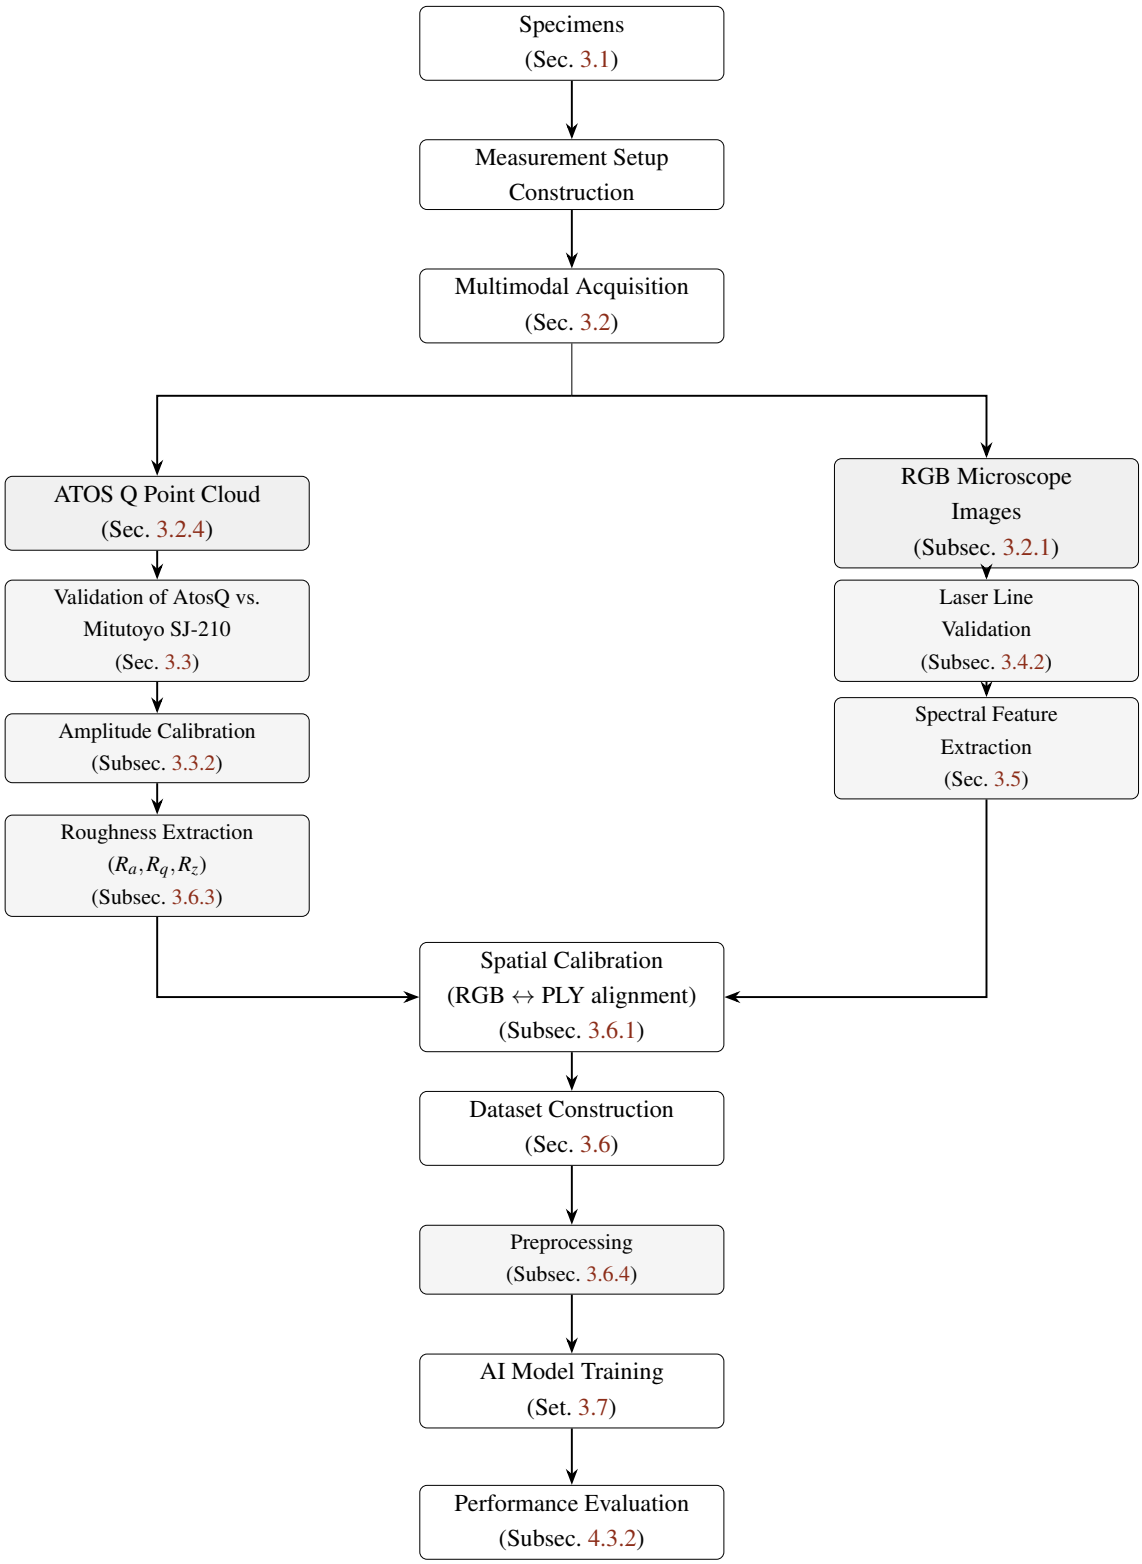
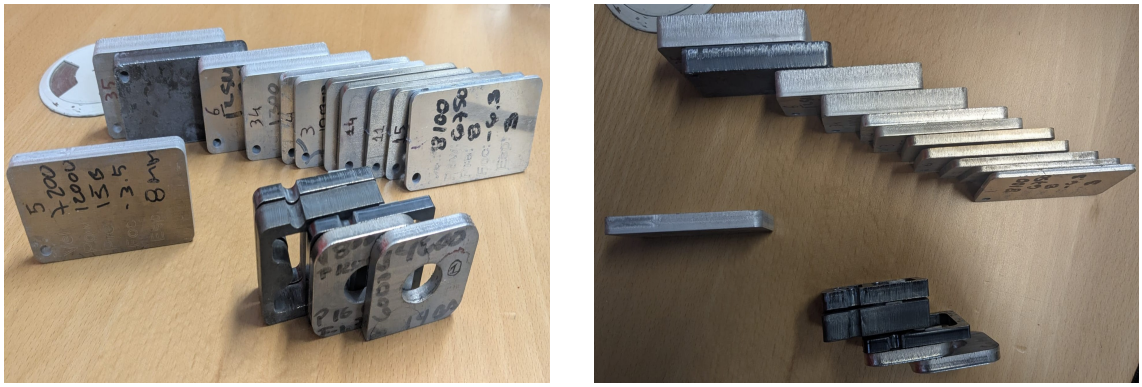


Figure 3.1: Overview of the methodological pipeline adopted in this dissertation. After multimodal acquisition, the workflow diverges into two parallel branches: validation of the ATOS Q system for surface roughness measurements, and spectral feature extraction from RGB data. Both branches converge into the dataset construction stage, which supports the subsequent AI-based regression.

3.1 Laser-Cut Specimens

The experimental dataset considered in this work comprised 17 laser-cut specimens. Since both sides of each specimen were analysed, the dataset included a total of 34 individual faces. The selected specimens covered different materials, thicknesses, and laser cutting conditions, which were considered relevant due to their potential influence on the resulting cut-edge quality and surface characteristics. The specimens analysed were not selected by the authors but were provided by the company as part of an industrial collaboration, meaning that the variation in material, thickness, geometry, and cutting conditions reflects practical production conditions rather than a fully controlled experimental design. This lack of control over specimen selection constitutes a limitation of the present study.

A representative view of the analysed specimens is presented in Figure 3.2. The main material properties and cutting parameters associated with the selected parts are summarised in Table 3.1. These parameters include the material type, thickness, laser power, cutting speed, assist gas, gas pressure, and focal position.



(a) Set of laser-cut specimens used in this work.

(b) Top view of specimens showing the analysed faces.

Figure 3.2: Laser-cut specimens considered in the experimental campaign.

3.2 Sensors and Acquisition Systems

Three sensors were used across two independent acquisition setups. The first setup combines an automated positioning platform with two interchangeable sensor configurations: *Configuration 1* accommodates a digital microscope for RGB image acquisition of the faces of laser-cut specimens, while *Configuration 2* mounts a surface roughness tester drive unit, in the latter, the microscope remains installed and serves as a visual aid for aligning the stylus on the specimen surface.

To allow these configurations of the first setup, a bracket was designed and manufactured specifically for this work using FDM.

Both configurations share a Creality Ender 3 desktop 3D printer [5], repurposed as an automated positioning platform, ensuring repeatable spatial referencing across measurement sessions.

Table 3.1: Material properties and cutting parameters of the 17 scanned laser-cut specimens.

ID	Material	Thickness [mm]	Power [W]	Speed [mm/min]	Gas	Pressure [bar]	Focal position [mm]
S01	Steel	6	3800	2500	O2	0.60	6.0
S02	Steel	8	3800	1250	O2	0.59	6.5
S03	Steel	10	3800	1100	O2	0.57	6.5
S04	Steel	12	3800	1000	O2	0.67	5.0
S05	Stainless steel	6	6000	4600	N2	14.0	3.3
S06	Stainless steel	6	12000	8000	N2	16.0	1.7
S07	Aluminium	5	8000	10800	N2	14.0	-0.7
S08	Aluminium	6	8000	7500	N2	14.0	-1.0
S09	Aluminium	8	12000	7200	N2	15.0	-3.5
S10	Aluminium	12	12000	1250	N2	15.0	-4.2
S11	Aluminium	3	6750	8100	N2	8.0	-6.3
S12	Galvanised carbon steel	4	7900	8800	N2	12.0	-0.8
S13	Stainless steel	5	8000	7000	N2	14.0	-0.5
S14	Galvanised carbon steel	3	7900	10000	N2	12.0	-1.0
S15	Steel	12	8000	1900	O2	0.7	9.0
S16	Aluminium	10	12000	1300	N2	10.0	-5.0
S17	Aluminium	15	12000	900	N2	14.8	-7.5

The second setup consists of a stand-alone structured-light 3D scanner operated independently, dedicated to full-surface three-dimensional geometry acquisition. The relationship between the two setups and the overall acquisition workflow is illustrated in Figure 3.3.

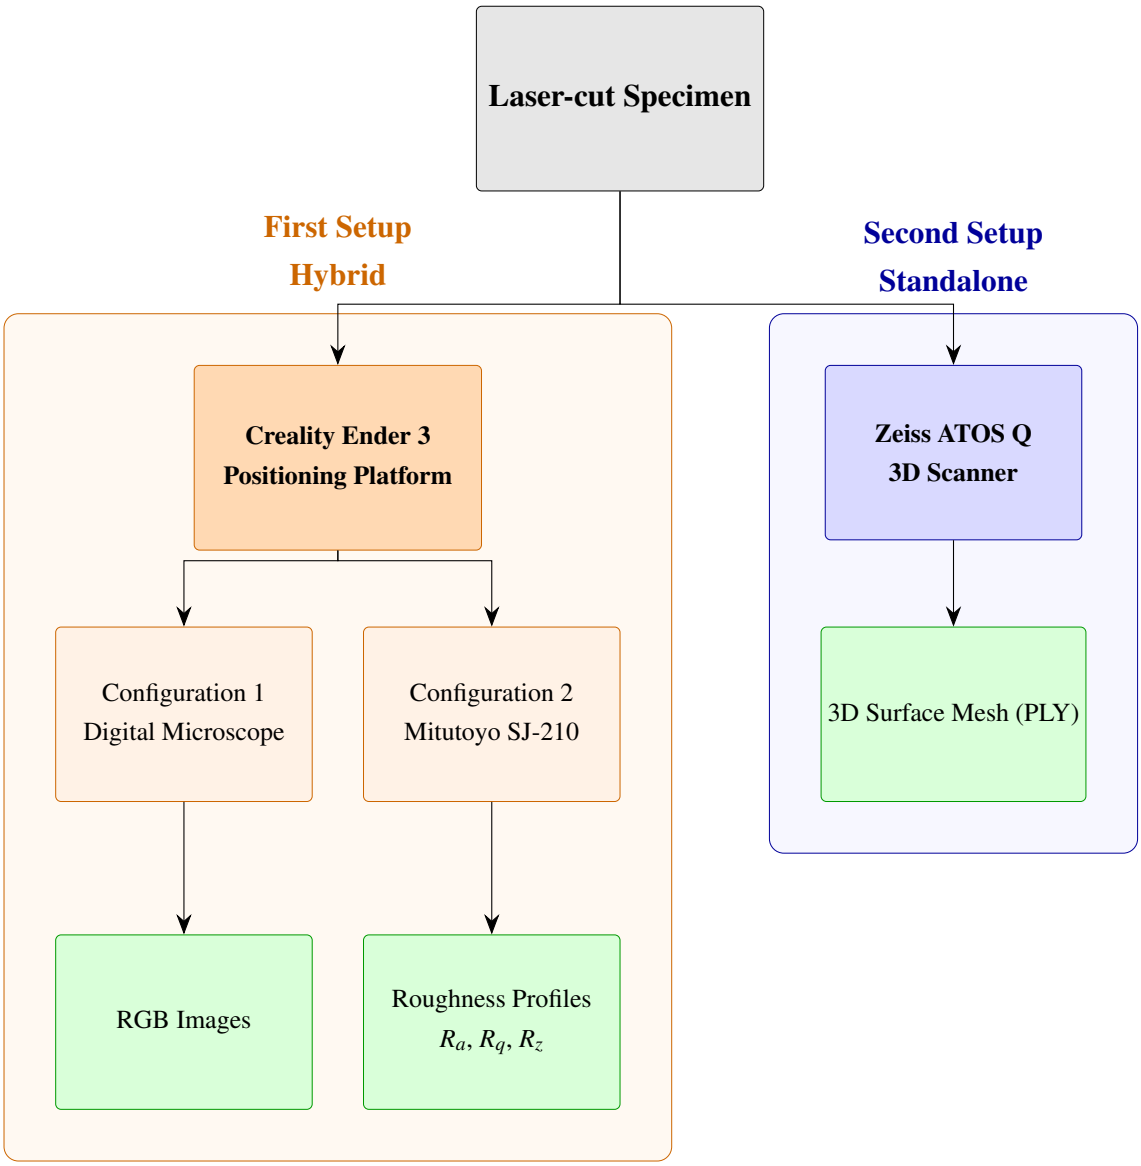
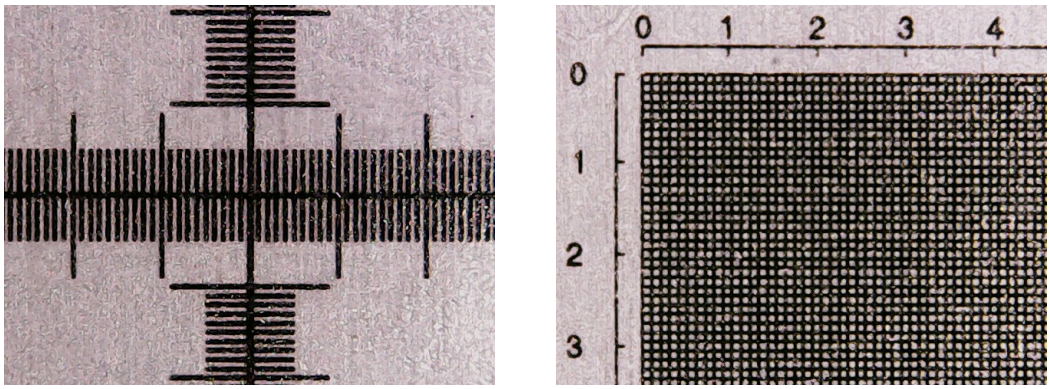


Figure 3.3: Overview of the two acquisition setups, sensor configurations, and respective data outputs.

3.2.1 RGB Image Acquisition

In the first setup, RGB images were acquired using a digital microscope mounted vertically on a custom bracket attached to the carriage of the Creality Ender 3 3D printer. This configuration, referred to as Configuration 1, exploits the printer's motion system to perform controlled planar displacement of the sensor over the specimen surface.

Before the acquisition, the microscope focus was adjusted manually for each specimen, and the Field of View (FOV) was verified using a calibration slide supplied by the sensor manufacturer that can be seen in Figure 3.4a, allowing the pixel-to-millimetre scale factor to be established accurately.

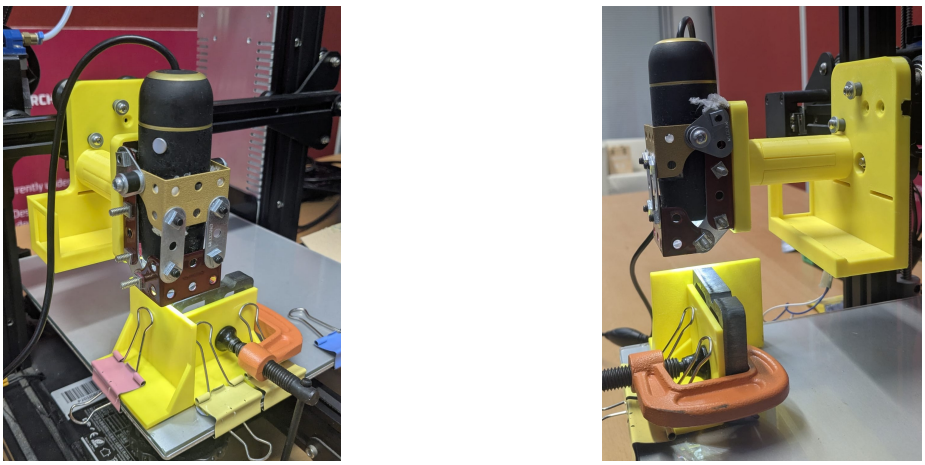


(a) Ruler used for FOV verification. (b) Calibration grid.

Figure 3.4: Calibration slide supplied by the manufacturer.

Image acquisition was performed in an automated manner through a custom Python script that communicates with the 3D printer via G-code commands over a serial interface, using the Marlin firmware protocol [36]. The script coordinates the sequential displacement of the carriage across predefined XY positions over the specimen surface, triggering image capture at each location, and recording the corresponding printer axis coordinates as metadata for each acquired image. This approach ensures reproducible positioning and eliminates operator-induced variability during the acquisition process.

Figure 3.5 shows the digital microscope mounted on the FDM bracket, Table 3.2 lists the main technical specifications of the digital microscope used in this work .



(a) Digital microscope acquiring RGB images over a specimen face. (b) Modular FDM bracket in Configuration 1.

Figure 3.5: RGB acquisition.

Table 3.2: Technical specifications of the digital microscope.

Parameter	**Value**
Sensor type	CMOS
Native resolution	640 × 480 px (0.3 MPx)
Interpolated resolution	(2.0 MPx)
Optical magnification range	5×
Digital magnification range	Up to 100×
Interface / Communication	USB 2.0
Output format	JPG

3.2.2 Microscope Calibration

Camera calibration establishes the mathematical relationship between image coordinates and real-world coordinates, enabling accurate pixel-to-millimetre conversion for quantitative surface measurements. The calibration problem is formulated through the pinhole camera model, in which the transformation from 3D world points to 2D image points is governed by the camera intrinsic parameters (focal length and principal point), lens distortion coefficients, and extrinsic parameters describing the camera pose.

Zhang's calibration method [60] was adopted for its flexibility in planar calibration scenarios. This approach estimates the homography for each view by observing a planar target from multiple orientations, eliminating the need for a dedicated 3D calibration rig. The theoretical basis of the method and the pinhole projection model are illustrated in Figures 3.6 and 3.7.

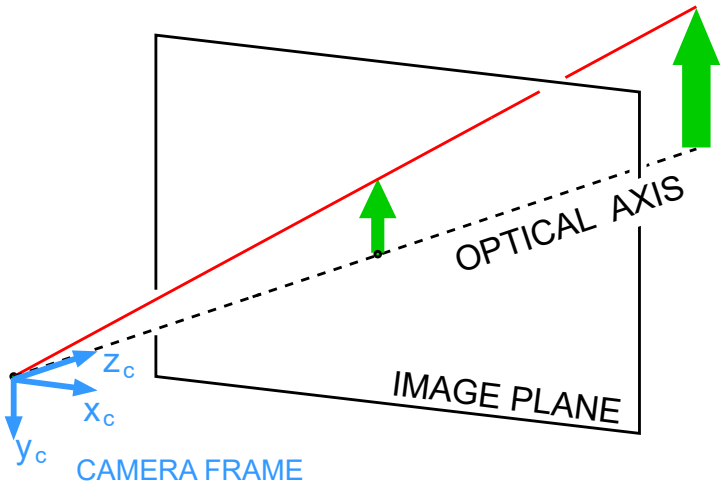


Figure 3.6: Pinhole camera model used as the theoretical basis for calibration.

Calibration target and procedure - Given the micro-scale nature of the imaging setup, a standard checkerboard target was impractical. Instead, a transparent microscope calibration slide with

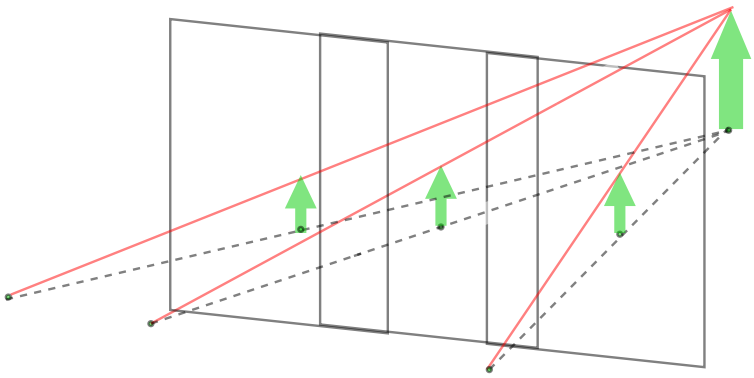


Figure 3.7: Multi-view planar calibration principle adopted in Zhang's method.

a precise grid matrix was used to provide the required geometric point correspondences, ensuring compatibility with the small FOV and close working distance of the digital microscope.

The calibration procedure comprised two stages: (i) extraction of grid intersections from calibration images, and (ii) estimation of the intrinsic camera matrix, lens distortion coefficients, and reprojection error using OpenCV [42]. The use of a planar grid makes the method particularly suitable for repeated laboratory calibration, as it is simple to reproduce and does not require highly specialised hardware.

Figure 3.4a shows the calibration slide used in the experimental setup, while Figure 3.4b presents the grid target used to establish the point correspondences required by the calibration algorithm.

Feature extraction pipeline - A multi-step image-processing pipeline was developed to detect grid intersections from the calibration slide. The raw calibration image was first converted to grayscale and binarized to isolate the grid pattern from the background. Horizontal and vertical structural curves were then extracted independently from the binary mask. Since the observed grid lines could be fragmented or partially overlapping, a grouping and merging stage was applied to consolidate them into continuous line segments.

After line refinement, the intersections between the horizontal and vertical grid structures were computed. These intersection coordinates served as the final 2D image points used for calibration. This pipeline is illustrated in Figure 3.8, which summarizes the intermediate outputs from grayscale conversion to the final intersection map.

Calibration results - The extracted 2D image points were matched with their known 3D coordinates on the calibration slide and passed to the OpenCV `calibrateCamera` function. The algorithm estimated the intrinsic parameters, radial and tangential distortion coefficients, and overall reprojection error.

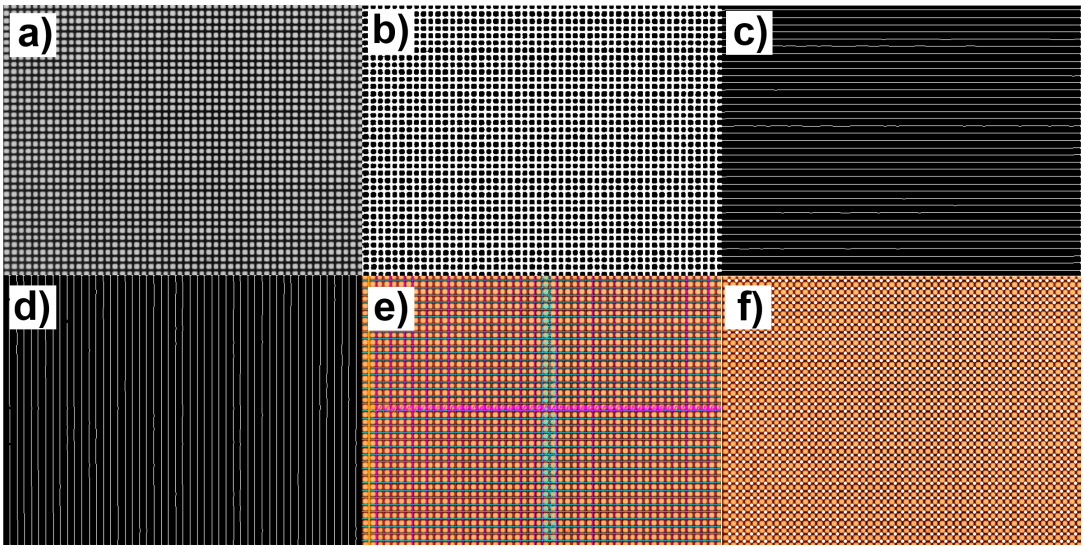


Figure 3.8: Feature-extraction pipeline used to detect the calibration grid intersections: grayscale input, binary mask, horizontal and vertical curves, merged lines, and final intersection points.

The total reprojection error was 0.6748 pixel, which is below the one-pixel threshold and indicates very good agreement between the mathematical model and the physical imaging system. The estimated camera parameters are summarised in Table 3.3. The focal lengths were 19 799.48 pixel in the x direction and 19 861.84 pixel in the y direction, with a principal point located at $(616.76, 278.79)$ pixels. The radial distortion coefficients were estimated as $k_1 = 3.3138$, $k_2 = -2117.75$, and $k_3 = -2.9675$, while the tangential distortion coefficients were $p_1 = -0.0040$ and $p_2 = 0.0220$.

Table 3.3: Estimated camera parameters obtained during microscope calibration.

Parameter	**Value**
Principal focal length f_x (pixels)	19799.48
Principal focal length f_y (pixels)	19861.84
Principal point c_x (pixels)	616.76
Principal point c_y (pixels)	278.79
Radial distortion (k_1, k_2, k_3)	$(3.3138, -2117.75, -2.9675)$
Tangential distortion (p_1, p_2)	$(-0.0040, 0.0220)$
Total reprojection error (pixels)	0.6748

Given the sub-pixel reprojection error, optical distortion across the FOV can be considered negligible for the intended metrological application. Consequently, direct pixel-to-length conversion can be used reliably without a separate undistortion stage.

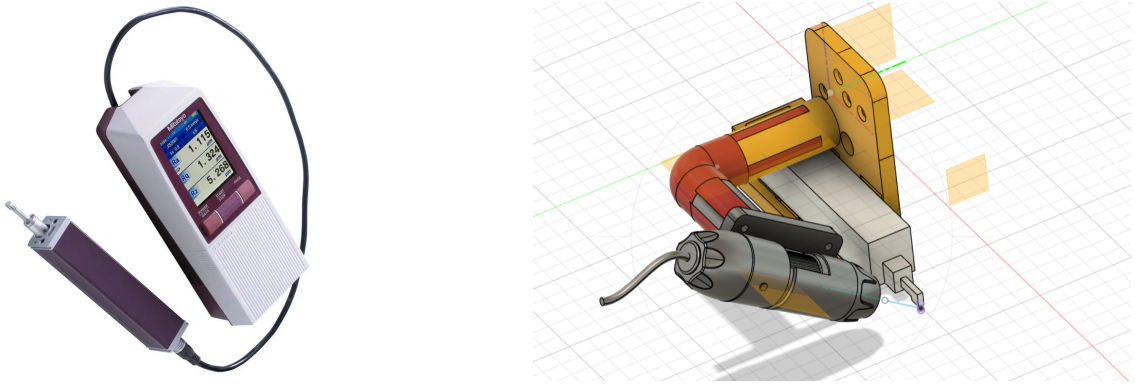
3.2.3 Surface Roughness Measurement

Surface roughness measurements were performed using a Mitutoyo SJ-210 surface roughness tester [37], mounted on the same custom FDM bracket described in Section 3.2. This configuration, referred to as Configuration 2, accommodates the surface roughness tester drive unit and enables a more precise and accurate control of the position of the stylus over the specimen surface through the printer's motion system.

Measurements were conducted on a single specimen, at five positions distributed along the specimen thickness direction. At each position, three independent measurements were performed, yielding a total of fifteen roughness profiles. The stylus traversal direction was set perpendicular to the laser cutting direction.

Roughness profiles were extracted and exported in Microsoft Excel format using the SJ-Communication-Tool software.

It should be noted that this configuration was employed as a validation step for surface measurements obtained from the structured-light scanner, which will be discussed in Section 3.3. Figure 3.9 shows the Mitutoyo SJ-210 surface roughness tester and its installation in Configuration 2, Table 3.4 lists the main technical specifications of the Mitutoyo SJ-210 surface roughness tester.



(a) Mitutoyo SJ-210 surface roughness tester. (b) Modular FDM bracket in Configuration 2.

Figure 3.9: Mitutoyo SJ-210 contact surface roughness tester and experimental setup.

Table 3.4: Technical specifications of the Mitutoyo SJ-210 surface roughness tester.

Parameter	**Value**
Model	Mitutoyo Surftest SJ-210
Measuring range (X axis)	16.0 mm
Measuring range (Z axis)	360 µm (−200 µm to 160 µm)
Resolution (Z axis)	0.02 µm
Measuring force	0.75 mN
Stylus tip radius	2 µm
Stylus tip angle	60°
Measuring speed	0.25, 0.50, 0.75 mm s^{-1}
Filter standard	ISO Gaussian
Cut-off (λ_c)	2.5 mm
Short-wavelength (λ_s)	8 µm
Evaluation length (l_t)	$1\times$, $2\times$, $3\times$, $5\times$, $10\times\lambda_c$
Applicable standards	JIS, ISO, ANSI, VDA

3.2.4 Structured Light Scanning

Three-dimensional surface geometry was acquired using the ATOS Q structured light scanner [59], equipped with a 100 mm measuring volume lens set (MV 100). The scanner was mounted on the dedicated Zeiss floor stand supplied with the system and operated via a laptop running the Zeiss Inspect 3D proprietary software [4].

The ATOS Q is an active structured light scanner [10] based on the Triple Scan principle. The sensor head comprises three optical units: a central blue Light Emitting Diode (LED) projector and two laterally mounted stereo cameras. During acquisition, the projector casts a sequence of precise fringe patterns onto the specimen surface. The two cameras capture the deformed fringe patterns simultaneously, and 3D coordinates are reconstructed via triangulation from three independent ray intersections: left camera/right camera, left camera/projector, and right camera/projector. This triple redundancy improves data reliability on challenging surfaces such as reflective areas, deep cavities, and sharp edges. The blue narrow-band LED light source further mitigates the influence of ambient illumination by enabling optical filtering during image acquisition.

Table 3.5 lists the main technical specifications of the ATOS Q scanner used in this work.

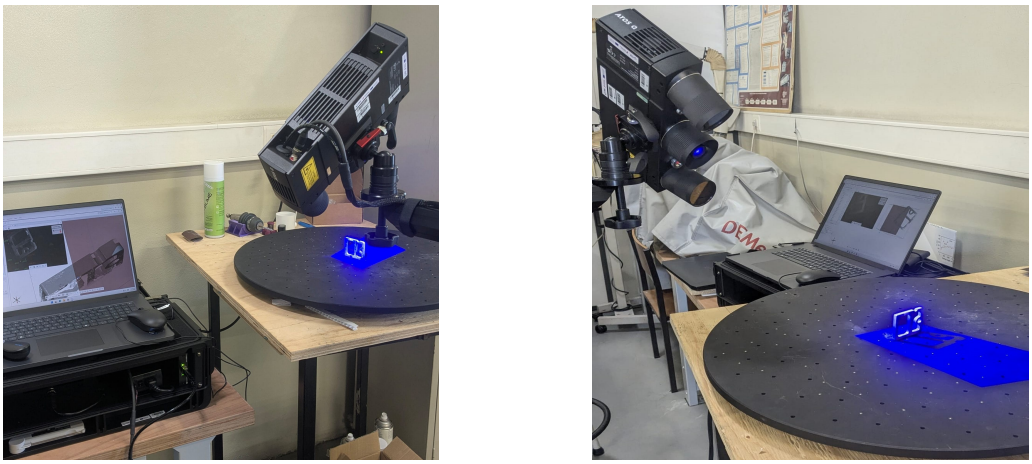
Table 3.5: Technical specifications of the ATOS Q structured light scanner (MV 100 lens configuration).

Parameter	**Value**
Model	ATOS Q (8 MP)
Light source	Blue LED (narrow-band)
Measuring volume (MV 100)	$100 \times 70 \times 60\,\text{mm}$
Measuring area	$100 \times 70\,\text{mm}^2$
Point distance (MV 100)	0.035 mm
Working distance	490 mm
Dimensions	$340 \times 240 \times 83\,\text{mm}$
Weight	approx. 4 kg
Software	*Zeiss Inspect 3D*
Output format	PLY, STL

Since the specimens consist of metallic sheets, their surfaces exhibit high reflectivity, which can introduce artefacts and data loss during structured light acquisition. To mitigate this effect, a thin uniform coat of ATTBLIME® ABP anti-reflective scanning spray [13] (Graichen Produktions- und Vertriebs GmbH, Germany) was applied to each specimen surface before scanning. The spray was applied from a distance of 15 cm to 25 cm, producing a uniform matte grey film with a nominal coating thickness of 7 µm, sufficient to diffuse the projected fringe pattern without masking fine surface details.

Specimens were placed on a Zeiss rotary table, whose surface was equipped with pre-attached adhesive reference markers. These markers allow the software to register and align successive scans into a common coordinate frame. Multiple scans were acquired at different rotational positions until maximum mesh coverage was achieved, minimising surface holes in the final model.

Scan point computation was performed using *Zeiss Inspect 3D*, with the parameters listed in Table 3.6. A conservative configuration was adopted to maximise the point cloud reliability: points associated with strong brightness differences, shiny surfaces, scan borders, shadow areas, and groove edges were set to *Avoid*, ensuring that only high-confidence points were retained. The minimum fringe contrast threshold was set to 15, the maximum residual to 0.05 pixel, and the maximum viewing angle between sensor and surface to 60°. The final surface mesh was exported in PLY format for subsequent processing. Figure 3.10 shows the ATOS Q sensor and a specimen positioned on the rotary table during acquisition.



(a) Overall setup.

(b) Structured light fringe projection.

Figure 3.10: Zeiss ATOS Q during specimen scanning.

Table 3.6: Scan Point Computation parameters used in the *Zeiss Inspect 3D* software.

Parameter	**Value**
Min. fringe contrast	15
Points at strong brightness differences	Avoid
Points on shiny surfaces	Avoid
Points on borders in scan area	Avoid
Points in shadow areas	Avoid
Points on groove edges	Avoid
Triple Scan points	Avoid
Triple Scan points at strong brightness diff.	Avoid
Max. residual	0.05 pixel
Max. viewing angle sensor/surface	60°
Depth limitation	Automatic
Scan surface far away from sensor	−200 mm
Scan surface close to sensor	150 mm
Reference points far away from sensor	−200 mm
Reference points close to sensor	150 mm
Threshold for movement check	0.05 pixel
Threshold for ambient light check	0.50 gray values
Quality check Triple Scan points	Intersection deviation
Threshold for intersection deviation	0.30 pixel

3.3 Validation of ATOS Q Surface Roughness Measurements

The ATOS Q structured-light scanner was not designed as a roughness metrology instrument. Its spatial resolution, while sufficient to capture the dominant surface texture features of a laser-cut edge, is known to attenuate micro-scale peak and valley amplitudes relative to a contact stylus, as discussed in Section 2.5. Before the scanner data could be used as a source of training labels, it was therefore necessary to establish that the extracted roughness profiles are a faithful representation of the surface, and to correct any systematic bias relative to a traceable contact reference. This section describes the validation procedure against the Mitutoyo SJ-210 contact profilometer and the calibration that follows from it. The validation was conducted on a single representative specimen, chosen for its broad range of roughness across the five measurement positions, which span from the laser entry side to the opposite face and cover a roughness range representative of the full dataset.

3.3.1 Profile Extraction from Point Cloud Data

Extracting a roughness profile from a structured-light point cloud requires transforming an unordered set of three-dimensional coordinates into a one-dimensional height signal comparable to that produced by a contact stylus. The procedure follows five stages: face identification, coordinate alignment, height projection, resampling, and Gaussian filtering.

The point cloud exported by the ATOS Q contains contributions from all visible faces of the specimen. The face of interest was isolated by iterative plane segmentation, in which a Random Sample Consensus (RANSAC) algorithm [8], detects one plane at a time, removes its inlier points, and proceeds to the next. For each candidate plane, a final equation was refined by least-squares fitting via Singular Value Decomposition (SVD) [12], over all inlier points. This two-step approach, RANSAC for robustness against outliers followed by SVD for precision, was applied iteratively until ten planes had been identified or too few points remained.

The resulting planes were grouped by dominant axis orientation into three reference surfaces: the cut face, the lateral face, and the top face. Their unit normals $\hat{\mathbf{x}}$, $\hat{\mathbf{y}}$, and $\hat{\mathbf{z}}$ define a local coordinate frame aligned with the thickness direction, the cutting direction, and the surface normal, respectively. The point cloud was transformed into this frame by

$$\mathbf{P}' = (\mathbf{P} - \mathbf{o})\,\mathbf{R}^{\top}, \tag{3.1}$$

where $\mathbf{R} = \begin{bmatrix} \hat{\mathbf{x}} & \hat{\mathbf{y}} & \hat{\mathbf{z}} \end{bmatrix}$ and $\mathbf{o}$ is the intersection of the three reference planes, computed by solving the corresponding linear system.

For each individual strip, the local surface normal was subsequently refined by SVD on the strip points alone. The signed distance of each point to the resulting local plane then gives the surface height deviation,

$$\delta_i = (\mathbf{p}_i - \bar{\mathbf{p}}) \cdot \hat{\mathbf{n}}, \tag{3.2}$$

where $\hat{\mathbf{n}}$ is the unit strip normal and $\bar{\mathbf{p}}$ is the strip centroid. The residuals δ_i, expressed in micrometres, constitute the raw surface deviation profile from which all subsequent processing is performed.

The profile was resampled in two stages. In the first stage, the point cloud was projected onto a coarse uniform grid whose spacing is governed by the ATOS Q point density, approximately 10 µm. Each grid cell was assigned the height of its nearest neighbouring point, without averaging, so that peak amplitudes are preserved, a property that is particularly important for the accurate estimation of R_z. Cells for which no point lay within five times the mean inter-point spacing were filled by linear interpolation between the two nearest valid cells,

$$\delta_k = \delta_{k_1} + \frac{k - k_1}{k_2 - k_1}\left(\delta_{k_2} - \delta_{k_1}\right), \quad k_1 < k < k_2, \tag{3.3}$$

where k_1 and k_2 are the indices of the nearest valid cells on either side of the gap at position k. In the second stage, the coarse profile was resampled onto a fixed grid with spacing $\Delta y = 1.55\,\mu\text{m}$ by linear interpolation, matching the nominal sampling interval of the SJ-210 and yielding approximately 8000 samples over the 12.5 mm evaluation length. Short-wavelength content introduced by this interpolation step was subsequently removed by the λ_s short-cutoff filter of the ISO 11562:1997 [21].

The resampled profile was filtered in accordance with ISO 11562 using a Gaussian filter with cut-off wavelength $\lambda_c = 2.5\,\text{mm}$ and profile factor $\alpha = 0.4697$, which ensures 50 % transmission at λ_c. The roughness profile was then obtained as the difference between the primary profile and the waviness returned by the filter,

$$z_r(y) = z(y) - z_w(y), \tag{3.4}$$

where $z(y)$ is the primary profile and $z_w(y)$ is the waviness component. A cut-off of 2.5 mm was applied uniformly to all profiles and to both instruments, following the practice established in the laser-cut edge literature [33, 53]. Although ISO 4288:1996 [23] recommends $\lambda_c = 8\,\text{mm}$ for surfaces with $R_z > 50\,\mu\text{m}$, the shorter cut-off was retained to maintain comparability between the two instruments and to remain consistent with the evaluation length available from the microscope FOV, as discussed in Section 2.2.3.

The roughness parameters were computed from the filtered profile in the manner prescribed by ISO 4287:1997 [22], which the Mitutoyo SJ-210 implements internally. Each parameter is the arithmetic mean of its values over $N_c = 5$ consecutive sampling lengths,

$$R_a = \frac{1}{N_c}\sum_{i=1}^{N_c} \left|z_r(y_i)\right|, \tag{3.5}$$

$$R_q = \sqrt{\frac{1}{N_c}\sum_{i=1}^{N_c} z_r^2(y_i)}, \tag{3.6}$$

$$R_z = \frac{1}{N_c}\sum_{k=1}^{N_c}\left(\max_{i\in I_k} z_r(y_i) - \min_{i\in I_k} z_r(y_i)\right), \tag{3.7}$$

where I_k is the set of sample indices within the k-th sampling length. This formulation was adopted for both instruments to ensure that any difference in the computed parameters reflects only the difference in measurement principle and not in the integration scheme.

The agreement between the two instruments was quantified by three scalar metrics computed on the roughness parameters across the five strips. The MAE and the RMSE measure the average and the deviation-weighted discrepancy between the ATOS Q and the SJ-210,

$$\mathrm{MAE} = \frac{1}{N_s}\sum_{j=1}^{N_s}\left|R_j^{\mathrm{ATOS}} - R_j^{\mathrm{SJ210}}\right|, \tag{3.8}$$

$$\mathrm{RMSE} = \sqrt{\frac{1}{N_s}\sum_{j=1}^{N_s}\left(R_j^{\mathrm{ATOS}} - R_j^{\mathrm{SJ210}}\right)^2}, \tag{3.9}$$

and the Mean Absolute Percentage Error (MAPE) expresses the same deviation in relative terms,

$$\mathrm{MAPE} = \frac{1}{N_s}\sum_{j=1}^{N_s}\left|\frac{R_j^{\mathrm{ATOS}} - R_j^{\mathrm{SJ210}}}{R_j^{\mathrm{SJ210}}}\right| \times 100\,\%, \tag{3.10}$$

where $N_s = 5$ is the number of strips. All three metrics operate on the same aggregated scalar roughness parameters. The RMSE penalises large deviations more heavily than the MAE, making it sensitive to outlier strips, while the MAPE facilitates comparison across parameters with different magnitudes.

3.3.2 Spatial Registration and Roughness Calibration

Before the calibration factor can be estimated, the five-strip measurement array extracted from each ATOS Q scan must be placed in correspondence with the reference positions measured by the SJ-210 stylus. This registration problem has two independent components: alignment along the cutting direction and positioning along the specimen thickness. Each component is solved by a different method, chosen to ensure that the amplitude of the roughness signal plays no role in determining the measurement position, which would otherwise introduce circularity into the subsequent calibration.

Longitudinal Alignment by Dynamic Time Warping - Along the cutting direction, the evaluation window was aligned to the stylus trace by Dynamic Time Warping (DTW), [45]. DTW finds the optimal non-linear correspondence between two sequences by minimising an accumulated cost, defined by the recurrence

$$D(i,j) = d(i,j) + \min\{D(i-1,j),\, D(i,j-1),\, D(i-1,j-1)\}, \tag{3.11}$$

where $d(i,j) = |a_i - b_j|$ is the pointwise difference between samples a_i and b_j. Warping was constrained by a Sakoe-Chiba band of width $w = 0.08\,L$, where L is the profile length, to prevent physically unreasonable alignments.

Before computing the DTW cost, both profiles were standardised by z-score,

$$\tilde{z}(y) = \frac{z(y) - \mu_z}{\sigma_z}, \tag{3.12}$$

so that the registration is driven entirely by the shape of the roughness profile rather than by its amplitude. This design choice is deliberate: by decoupling shape registration from amplitude estimation, the position of the measurement window is not influenced by the calibration error, which preserves the independence between registration and calibration.

In practice, the DTW cost was minimised over a two-dimensional grid spanning both the longitudinal offset Δy and the transverse offset Δx, with the total score being the sum of DTW distances over the five strips,

$$(\Delta x^\star, \Delta y^\star) = \arg\min_{\Delta x, \Delta y} \sum_{s=1}^{5} \mathrm{DTW}\big(\tilde{z}_s^{\mathrm{ATOS}}(\Delta x, \Delta y),\, \tilde{z}_s^{\mathrm{SJ210}}\big)\,. \tag{3.13}$$

The search was conducted in two stages. A coarse pass covered offsets of ± 1 mm in Δx and ± 2 mm in Δy in steps of 0.20 mm and 0.30 mm, respectively. A fine pass then refined the solution around the coarse optimum in steps of 0.05 mm. The resulting longitudinal window and the mean surface height were stored for each acquisition and passed to the calibration stage.

Transverse Positioning by Roughness Pattern Matching - The along-strip profile carries limited information about the position in the thickness direction, because all strips share a similar oscillatory roughness pattern that differs only in amplitude. The transverse position was therefore determined in the calibration step by an independent sweep over the specimen thickness.

The five-strip comb was translated across the full thickness range in steps of 0.05 mm and the transverse offset was selected as the one for which the normalised R_a profile of the ATOS Q across the strip centres best matched the normalised across-width pattern of the SJ-210. Formally, the selected offset minimises the shape mismatch

$$\Delta x^\star = \arg\min_{\Delta x} \sum_{s=1}^{5} \left(\frac{R_a^{\mathrm{ATOS}}(x_s)}{\overline{R_a^{\mathrm{ATOS}}}} - \frac{R_a^{\mathrm{SJ210}}(s)}{\overline{R_a^{\mathrm{SJ210}}}} \right)^2, \quad x_s = \Delta x + (s-1)\,\Delta x_{\mathrm{gap}}, \tag{3.14}$$

where $\Delta x_{\mathrm{gap}} = 2\,\mathrm{mm}$ is the fixed inter-strip spacing. Because the criterion depends only on the shape of the transverse roughness trend and not on its amplitude, the transverse position remains independent of the calibration error. The longitudinal window obtained from DTW was kept fixed throughout this step.

Amplitude Calibration Factor - Structured-light scanners are known to underestimate roughness amplitudes on laser-cut surfaces, because their lateral resolution is insufficient to resolve the

steep walls of the cut grooves, leading to a systematic loss of peak-to-valley height. To correct for this effect, a linear amplitude gain factor k was estimated from the registered strip pairs. The calibrated roughness parameters are

$$R_\alpha^{\text{cal}} = k \cdot R_\alpha^{\text{ATOS}}, \qquad \alpha \in \{a, q, z\}, \tag{3.15}$$

where k is the mean ratio of the SJ-210 measurements to the ATOS Q measurements across all registered strip pairs,

$$k_\alpha = \frac{1}{N_s N_m} \sum_{m=1}^{N_m} \sum_{s=1}^{N_s} \frac{R_{\alpha,s,m}^{\text{SJ210}}}{R_{\alpha,s,m}^{\text{ATOS}}}, \tag{3.16}$$

with $N_s = 5$ strips and $N_m = 4$ acquisitions, yielding $N_s \times N_m = 20$ registered pairs in total. The estimated factors were $k_a = 1.49$, $k_q = 1.50$, and $k_z = 1.51$. Given the small variation across parameters, a single representative value $k = 1.49$ was adopted uniformly for all three, as the differences fall within the inter-acquisition variability observed across the four calibration scans.

3.4 Striation Period Estimation: Initial Approach

The thermal interaction between the laser beam and the workpiece produces a quasi-periodic dragline pattern on the cut-edge surface. The striations manifest as alternating dark valleys and bright crests oriented perpendicular to the cutting direction, and are clearly discernible in standard RGB microscope images. The first step of the present work was to verify whether this pattern could be reliably characterised by its spatial period from image data alone.

An initial spectral analysis pipeline was developed to estimate the dominant striation wavelength λ_v directly from the grayscale intensity profile of each RGB image, using a one-dimensional real-valued Fast Fourier Transform (rFFT). To verify that the extracted λ_v corresponded to a genuine physical periodicity of the cut-edge surface rather than to an image artefact or a residual illumination gradient, an independent geometric measurement was performed using oblique laser line projection.

3.4.1 FFT-Based Striation Wavelength Estimation

Each RGB image was converted to an 8-bit grayscale representation $I_g(x,y)$ using the standard ITU-R BT.601 luminance weighting:

$$I_g = 0.299\,R + 0.587\,G + 0.114\,B \tag{3.17}$$

where R, G, and B denote the red, green, and blue channel intensities in the range $[0, 255]$. A one-dimensional intensity profile was constructed by averaging image rows over a central band spanning the vertical extent $[0.45\,H,\ 0.95\,H]$, where H denotes the image height in pixels. This band was chosen to cover the region of highest striation contrast, which, as visible in Figure 3.13,

corresponds to the lower portion of the image. The resulting signal was detrended by subtracting a first-order polynomial least-squares fit and multiplied by a Hann window to suppress spectral leakage. The windowed signal was then transformed via the rFFT and the dominant frequency bin k^* was identified as the bin of maximum magnitude within a search band restricted to physically plausible striation frequencies, excluding the DC component and high-frequency noise. The corresponding striation wavelength was obtained as $\lambda_v = 1/f(k^*)$, where $f(k^*)$ denotes the dominant bin frequency in cycles per pixel. No harmonic disambiguation or illumination correction was applied at this stage.

3.4.2 Validation via Laser Line Projection

To verify that the extracted λ_v corresponded to a genuine physical periodicity of the cut-edge surface rather than to an image artefact or a residual illumination gradient, an independent geometric measurement was performed using oblique laser line projection. A construction-grade cross-line laser level was directed onto the cut-edge surface at a shallow oblique angle of incidence. This illumination geometry causes the visible red light to preferentially fill the recessed striation valleys while leaving the raised crests largely unilluminated. The resulting chromatic contrast enables direct segmentation of the striation pattern by colour, yielding a valley count and an inter-valley spacing that is entirely independent of the FFT analysis and anchored in the physical three-dimensional geometry of the surface. The experimental setup is shown in Figure 3.11.



Figure 3.11: Experimental setup for the laser line projection validation experiment.

The complete validation pipeline is illustrated in Figure 3.12 and comprises four sequential stages: segmentation of the laser-illuminated image in the Hue, Saturation, Value (HSV) colour space, morphological post-processing to consolidate valley detections, connected-component analysis to extract individual valley segments, and computation of the mean inter-valley spacing.

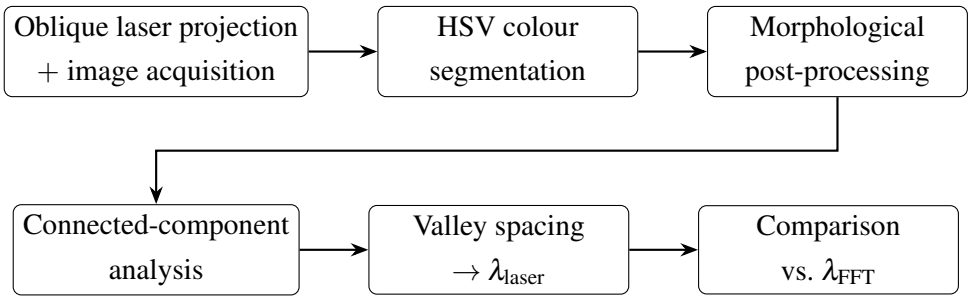


Figure 3.12: Validation pipeline for valley wavelength estimation using oblique laser line projection and HSV colour segmentation.

HSV colour segmentation - Each laser-illuminated image was converted to the HSV colour space, which decouples chromatic information from luminance and makes the red-stripe detection insensitive to illumination gradients. In the OpenCV [42], HSV convention ($H \in [0, 180]$, $S, V \in [0, 255]$), the red spectral region maps to hue values near both 0 and 180 due to the circular wrap-around of the hue axis. A binary mask was therefore formed as the union of two complementary threshold ranges:

$$M_{\text{raw}}(x,y) = \mathbf{1}[H \in [H_{\min,1}, H_{\max,1}]] \vee \mathbf{1}[H \in [H_{\min,2}, H_{\max,2}]] \tag{3.18}$$

subject to $S(x,y) \geq S_{\min}$ and $V(x,y) \geq V_{\min}$, where the saturation lower bound suppresses achromatic background pixels and the value lower bound excludes dark shadow regions. Table 3.7 lists the threshold values adopted.

Table 3.7: HSV threshold parameters for red laser stripe segmentation.

Parameter	**Symbol**	**Value**
Hue lower bound — primary range	$H_{\min,1}$	0
Hue upper bound — primary range	$H_{\max,1}$	10
Hue lower bound — wrap-around	$H_{\min,2}$	165
Hue upper bound — wrap-around	$H_{\max,2}$	180
Saturation lower bound	$S_{\min}$	120
Value lower bound	$V_{\min}$	50

Morphological post-processing - A sequential morphological pipeline is applied to consolidate fragmented detections into coherent, vertically elongated valley segments: a morphological opening (2×2 px, 3 iterations) removes small isolated noise clusters, a vertical dilation (1×15 px, 9 iterations) bridges intra-valley discontinuities caused by local variations in surface reflectivity, a Union-Find algorithm merges components belonging to the same physical striation that remain separated by horizontal gaps subject to a maximum centroid separation of 40 px and a maximum vertical bounding-box gap of 150 px, and an additional splitting step separates components

wider than 40 px that may represent two adjacent valleys merged by over-dilation, using a column-density peak-finding procedure.

Connected-component analysis and wavelength extraction - Valley segments were extracted from the post-processed mask using 8-connectivity labelling. Components were retained only if their bounding-box area satisfied $A_i \geq 1700\,\mathrm{px}^2$ and their aspect ratio satisfied $h_i/w_i \geq 2.1$, ensuring that only vertically elongated structures consistent with laser-cut striations were counted. The retained components were sorted by horizontal centroid abscissa c_x, and the mean inter-valley spacing was computed as:

$$\bar{\Delta}_X = \frac{c_{x,N} - c_{x,1}}{N-1} \quad [\mathrm{px}] \tag{3.19}$$

The laser-based valley wavelength was then obtained as:

$$\lambda_{\mathrm{laser}} = \bar{\Delta}_X \cdot s_{\mathrm{px}} \quad [\mathrm{mm}] \tag{3.20}$$

where $s_{\mathrm{px}} = \mathrm{FOV}_X/W$ is the physical size of one pixel in $\mathrm{mm\,px}^{-1}$. A representative result is shown in Figure 3.13, where 18 valley segments are detected with a mean inter-centroid spacing of 36.1 px.

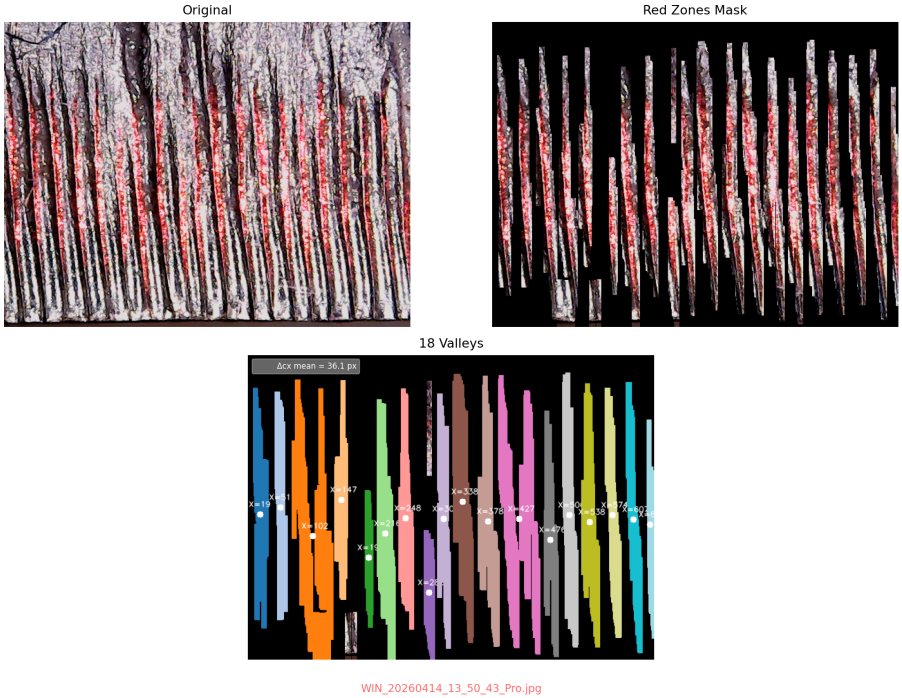


Figure 3.13: Laser line projection result on a representative specimen face: original RGB image under oblique laser illumination (left), binary segmentation mask (centre), and connected-component overlay with annotated centroids (right).

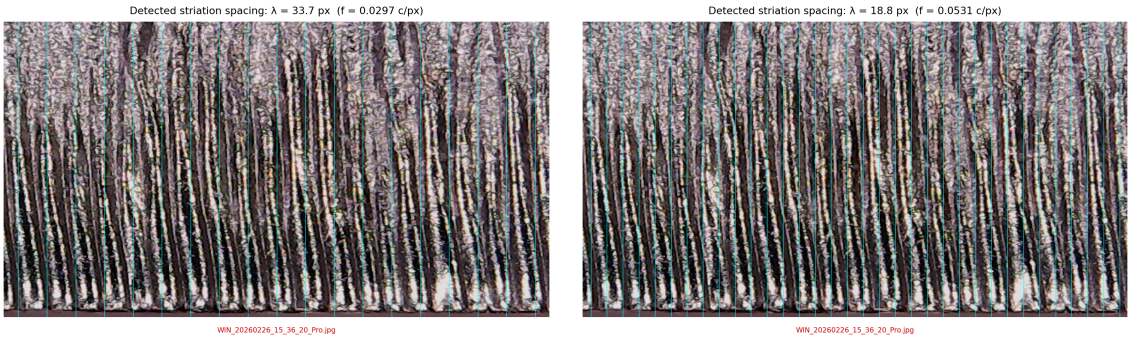
3.4.3 Harmonic Ambiguity and Motivation for a Robust Pipeline

A direct comparison between the FFT-derived wavelengths λ_{FFT} and the laser-based references λ_{laser} on the same set of images confirmed that the striation periodicity is captured by the grayscale intensity spectrum for the majority of images. However, a subset of images yielded λ_{FFT} values equal to approximately half the laser reference, indicating that the dominant FFT peak had locked onto the second harmonic of the true striation frequency rather than the fundamental. Figure 3.14 illustrates this failure mode on the same specimen image: in one run the pipeline correctly identifies $\lambda = 33.7\,\mathrm{px}$, whilst in another it locks onto $\lambda = 18.8\,\mathrm{px}$, approximately half the true period.

This harmonic locking arises because surface striations are not purely sinusoidal: the asymmetric intensity profile of valleys and crests generates spectral energy at integer multiples of the fundamental frequency, and in images with sharp or high-contrast striations the second harmonic can exceed the fundamental in magnitude.

A further limitation of the initial approach was its sensitivity to low-frequency illumination gradients. Without flat-field correction, smooth intensity drifts across the FOV introduced spurious low-frequency content into the spectrum, affecting the reliability of strip-level estimates in particular.

These two limitations, harmonic locking and illumination sensitivity, motivated the development of the robust spectral pipeline presented in Section 3.5, which incorporates anisotropic flat-field correction and the Harmonic Product Spectrum to resolve both issues systematically.



(a) FFT captures the fundamental striation period.

(b) FFT locks onto the second harmonic, approximately half the true period.

Figure 3.14: Illustration of the harmonic locking failure mode on the same specimen image.

3.5 Spectral Feature Extraction from RGB Images

The preliminary study described in Section 3.4 established that the striation period is accessible from RGB images via spectral analysis, but identified two systematic failure modes in the plain rFFT approach: harmonic locking and sensitivity to illumination gradients. The pipeline described in this section resolves both issues through anisotropic flat-field correction, which suppresses illumination artefacts prior to profile construction, and the HPS [41], which ensures that the identified

dominant frequency corresponds to the physical striation fundamental rather than one of its harmonics.

The estimation pipeline, illustrated in Figure 3.15, comprises five sequential stages: grayscale conversion, anisotropic flat-field correction, construction of a one-dimensional intensity profile, spectral analysis via the rFFT combined with the HPS to identify the fundamental striation frequency, and extraction of the three dominant wavelengths together with their spectral strengths.

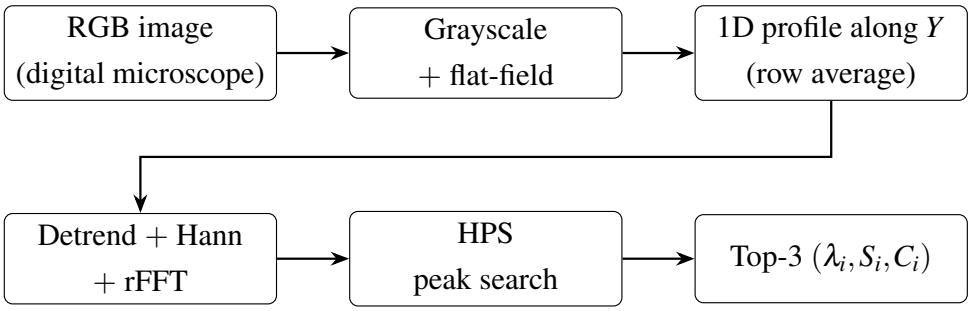


Figure 3.15: Pipeline for valley wavelength estimation from RGB microscope images via one-dimensional Fourier analysis.

3.5.1 Image Pre-processing

Each RGB image was converted to an 8-bit grayscale representation $I_g(x,y)$ following Equation 3.17.

The physical scale is set by the pixel sampling interval $s_{\mathrm{px}} = \mathrm{FOV}_H/W$, where $\mathrm{FOV}_H = 5.081\,\mathrm{mm}$ is the calibrated horizontal FOV the microscope and $W = 640\,\mathrm{px}$ is the image width, yielding $s_{\mathrm{px}} \approx 0.007\,94\,\mathrm{mm\,px^{-1}}$. The corresponding vertical FOV is $\mathrm{FOV}_V = H \cdot s_{\mathrm{px}} \approx 3.81\,\mathrm{mm}$, where $H = 480\,\mathrm{px}$.

Anisotropic flat-field correction - Even under controlled illumination, microscope images exhibit low-frequency intensity variations caused by light source non-uniformity, specular highlights on raised crests, and gradual shadowing across the FOV. These artefacts contaminate the Fourier spectrum with spurious low-frequency content and introduce bias when sub-regions of the image are analysed independently. To suppress them, an anisotropic flat-field correction is applied prior to constructing the 1D profile.

The local background $B_g(x,y)$ is estimated by Gaussian blurring with markedly different standard deviations along the two axes:

$$B_g(x,y) = G_{\sigma_X,\sigma_Y}[I_g(x,y)], \qquad \sigma_X = 25\,\mathrm{px}, \quad \sigma_Y = 3\,\mathrm{px} \tag{3.21}$$

The large σ_X strongly smooths the image along the striation direction, averaging out the periodic intensity modulation, whilst the small σ_Y preserves spatial resolution along the cutting direction. The corrected image is then obtained by normalisation:

$$I'_g(x,y) = \frac{I_g(x,y)}{\max(B_g(x,y),\,1)} \cdot 128 \tag{3.22}$$

The denominator is clipped from below at unity to prevent division by zero in dark regions, and the factor of 128 re-centres the output around a mid-grey level for numerical stability of subsequent operations.

3.5.2 One-Dimensional Profile and Spectral Analysis

The flat-fielded image I'_g is reduced to a one-dimensional signal by averaging across image rows. For an analysis region indexed by the row set $\mathscr{R}$, the profile is:

$$p(y) = \frac{1}{|\mathscr{R}|} \sum_{x \in \mathscr{R}} I'_g(x,y), \qquad y = 0,\,1,\,\ldots,\,W-1 \tag{3.23}$$

This produces a 1D signal of length W that captures the periodic alternation between dark valleys and bright crests along Y.

Detrending and windowing - A linear trend is removed from $p(y)$ by subtracting the first-order polynomial least-squares fit, yielding the detrended signal $\tilde{p}(y)$. This eliminates any residual low-frequency component not captured by the flat-field correction and prevents the DC and near-DC bins from dominating the spectrum. A Hann window is then applied to suppress spectral leakage arising from the finite signal length:

$$q(y) = \tilde{p}(y) \cdot w(y), \qquad w(y) = \frac{1}{2}\left[1 - \cos\left(\frac{2\pi y}{W-1}\right)\right] \tag{3.24}$$

Real-valued FFT - The windowed signal $q(y)$ is transformed via the rFFT, which exploits the conjugate symmetry of a real input to compute only the $\lfloor W/2 \rfloor + 1$ unique non-negative frequency bins:

$$Q(k) = \sum_{y=0}^{W-1} q(y)\, \mathrm{e}^{-\mathrm{j}\,2\pi k y / W}, \qquad k = 0,\,1,\,\ldots,\,\lfloor W/2 \rfloor \tag{3.25}$$

The physical frequency axis is obtained through the sampling interval s_{px}:

$$f(k) = \frac{k}{W \cdot s_{\mathrm{px}}} \quad \left[\mathrm{mm}^{-1}\right] \tag{3.26}$$

Search band restriction - The peak search is restricted to a physically meaningful frequency band, excluding the DC component and high-frequency noise:

$$k_{\min} = \max(1,\ \lfloor 0.010\,W \rfloor), \qquad k_{\max} = \min(\lfloor 0.15\,W \rfloor,\ \lfloor W/2 \rfloor) \tag{3.27}$$

For $W = 640\,\text{px}$ and $\text{FOV}_H = 5.081\,\text{mm}$, this corresponds to a wavelength range of approximately $\lambda_v \in [0.053\,\text{mm}, 0.847\,\text{mm}]$, which encompasses the expected striation period across all cutting conditions investigated.

3.5.3 Harmonic Product Spectrum for Fundamental Identification

Surface striations are not purely sinusoidal: their asymmetric intensity profile generates spectral energy not only at the true striation frequency but also at integer multiples of it. In images where the striations are sharp or high-contrast, a direct search over the magnitude spectrum can lock onto a harmonic rather than the fundamental, yielding a wavelength that is an integer fraction of the true value. To suppress this failure mode, the HPS is employed, defined as the pointwise product of the magnitude spectrum with its downsampled copies:

$$\text{HPS}(k) = \prod_{h=1}^{n_h} |Q(h \cdot k)|\,, \qquad k = 0, 1, \ldots, \left\lfloor \frac{W}{2 n_h} \right\rfloor \tag{3.28}$$

with $n_h = 3$ in this work. The downsampling by integer factor h collapses the h-th harmonic of any fundamental onto the fundamental bin itself, so bins that correspond to a true fundamental with strong harmonics are reinforced multiplicatively, whilst isolated harmonic bins are not. The dominant bin is then identified as:

$$k^* = \arg \max_{k \in [k_{\min}, k_{\max}]} [\text{HPS}(k)] \tag{3.29}$$

and the corresponding striation wavelength is obtained in physical units as:

$$\lambda_v = \frac{1}{f(k^*)} = \frac{W \cdot s_{\text{px}}}{k^*} \quad [\text{mm}] \tag{3.30}$$

3.5.4 Spectral Strength and Peak Confidence

Two complementary scalar descriptors are computed for each identified peak to quantify the reliability of the corresponding wavelength estimate.

The spectral strength S measures the saliency of the dominant peak relative to the background spectral noise level, defined as the ratio of the peak magnitude to the median of the rFFT magnitude spectrum over the search band:

$$S = \frac{|Q(k^*)|}{\text{median}_{k \in [k_{\min}, k_{\max}]} |Q(k)|} \tag{3.31}$$

A high value of S indicates that the dominant frequency stands well above the spectral floor, suggesting a regular and well-defined striation pattern. Note that S is computed on the original rFFT magnitude spectrum rather than on the HPS, making it more physically interpretable as a measure of absolute peak prominence.

The peak confidence C measures the sharpness of the peak relative to its immediate spectral neighbourhood, defined as the ratio of the peak magnitude to the mean of the surrounding bins within a window of $\pm\delta_{\mathcal{N}}$ bins, excluding the $\pm\delta_{\text{ex}}$ bins closest to the peak:

$$C = \frac{|Q(k^*)|}{\text{mean}_{k\in\mathcal{N}(k^*)}\,|Q(k)|} \tag{3.32}$$

where the neighbourhood $\mathcal{N}(k^*)$ is defined as:

$$\mathcal{N}(k^*) = [k^* - \delta_{\mathcal{N}},\, k^* + \delta_{\mathcal{N}}] \setminus [k^* - \delta_{\text{ex}},\, k^* + \delta_{\text{ex}}] \tag{3.33}$$

with $\delta_{\mathcal{N}} = 8$ bins and $\delta_{\text{ex}} = 2$ bins. Whilst S captures global spectral saliency, C captures local peak sharpness: a broad, flat spectral lobe yields a low C even if S is high, making C a more reliable indicator of whether the identified wavelength corresponds to a genuine periodic structure rather than a diffuse frequency band.

3.5.5 Top-Three Peak Extraction

Rather than retaining only the dominant striation wavelength, the pipeline extracts the three most prominent independent spectral peaks, providing a richer feature set for the subsequent regression models and mitigating the risk of feature collapse when the spectrum exhibits two peaks of comparable amplitude.

The extraction is performed by a greedy peak-picking procedure on the HPS values within the search band. After each peak is identified, two complementary suppression steps are applied before searching for the next one: a local neighbourhood of $\pm\delta_s = 5$ bins around the found peak is set to $-\infty$, preventing adjacent bins of the same spectral lobe from being selected as independent peaks, and the neighbourhoods ($\pm\,\delta_s$ bins) around the harmonic and sub-harmonic frequencies $h \cdot f^*$ for $h \in \{2, 3, 4, \frac{1}{2}, \frac{1}{3}\}$ are likewise suppressed. This harmonic-aware strategy ensures that λ_2 and λ_3 represent genuinely independent spectral components of the surface texture rather than integer multiples or fractions of the dominant striation frequency λ_1.

For each of the three peaks, the wavelength λ_i ($i = 1,2,3$, ordered by HPS amplitude) is computed via Equation 3.30, and the corresponding spectral strength S_i and peak confidence C_i are evaluated via Equations 3.31 and 3.32 against the same search-band statistics. The triplet $\{(\lambda_i, S_i, C_i)\}_{i=1,2,3}$ constitutes the Fourier-based feature set extracted per analysis region. This extraction is performed at two spatial scales: once over the full image, yielding a global feature set, and once per sub-region of the image aligned with the specimen thickness direction, yielding local strip-level features. The definition of these sub-regions and their alignment with the structured-light scan coordinate system are described in Section 3.6.

3.6 Dataset Construction

The training dataset pairs RGB microscope images and their associated spectral features with calibrated roughness labels derived from the ATOS Q structured-light scanner. Its construction involves three interdependent stages: mesh pre-processing and face alignment, image acquisition with spatial metadata recording, and roughness label extraction. The complete pipeline is illustrated in Figure 3.16.

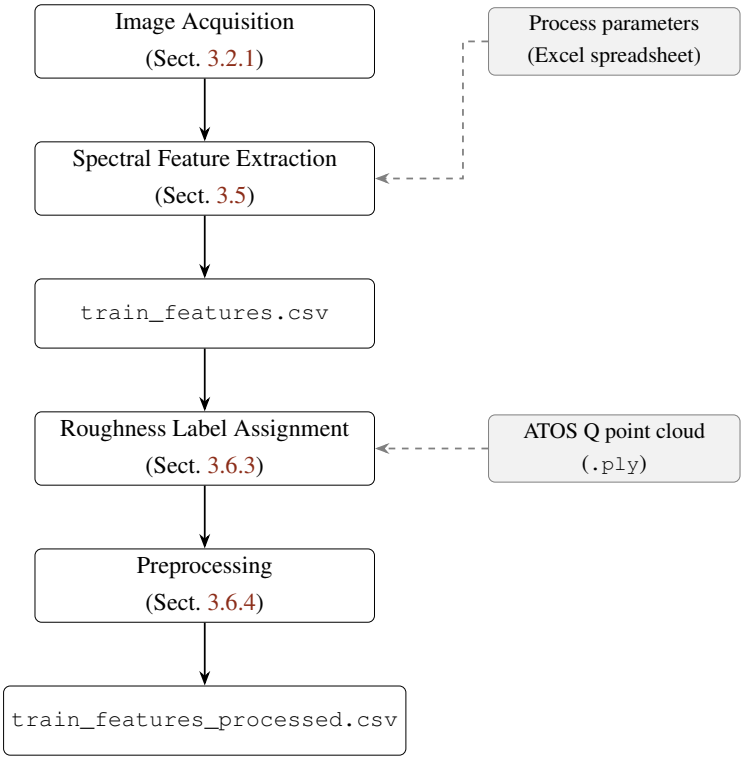


Figure 3.16: Three-stage pipeline for training dataset construction.

From the initial set of 34 specimen faces, four faces with thickness below 4 mm were excluded prior to dataset construction. The remaining 30 faces span four materials and 15 distinct cutting conditions, yielding a total of 1041 observations after roughness label assignment and quality filtering.

3.6.1 Spatial Alignment, Strip Definition and Co-localisation

A fundamental requirement of the dataset is that the spectral features extracted from each RGB image and the roughness labels assigned to it refer to the same physical region of the specimen surface. This co-localisation is achieved through the combination of a structured-light mesh pre-processing pipeline and a spatial metadata system embedded in the image acquisition process.

Each ATOS Q point cloud was pre-processed to isolate the cut-edge face of interest and place it in a consistent coordinate system. The face was first cropped from the full scan using a manually defined bounding box, adjusted interactively in cases where localised voids were present. A

least-squares plane was then fitted to the cropped point cloud using SVD to remove residual tilt introduced during specimen fixturing, and Principal Component Analysis (PCA) [26], was applied to align the face with the target coordinate system: X along the specimen thickness direction, Y along the cutting direction, and Z along the surface normal, the axis along which roughness depth is measured. The resulting point cloud, expressed in the aligned coordinate system, was exported as a `.ply` file for downstream processing.

RGB images were acquired using a motorised stage that traverses the specimen face in a regular grid pattern. For each acquired image, the controller automatically records a metadata entry in a per-face CSV file containing the image filename, its grid position indices, and the physical coordinates of its footprint in the specimen face coordinate system (`x_start`, `x_end`, `y_start`, `y_end`), expressed in millimetres. Since both the acquisition stage and the structured-light scanner share a common physical coordinate system anchored to the specimen, these coordinates are directly compatible with the aligned mesh, without requiring any image-to-mesh registration step. The physical footprint of each image is used to retrieve the corresponding mesh points by a simple spatial query, from which the roughness profile is extracted.

Within each image footprint, three horizontal sub-regions are defined along the X axis, referred to as strips. Each strip has a half-width of 0.15 mm and is centred at positions X_1, X_2, and X_3, placed at an inset of 0.7 mm from the face edges and at the face centre respectively. Strip 1 is positioned closest to the laser exit side of the cut, where striations are most well-defined, and strip 3 closest to the laser entry side. Spectral features and roughness labels are extracted independently for each strip, capturing the spatial variation of surface texture along the thickness direction, in addition to the global features extracted from the full image. The spatial alignment between the image grid, the strip positions, and the underlying point cloud is illustrated in Figure 3.17 and the coordinate system adopted for the aligned mesh is illustrated in Figure 3.18.

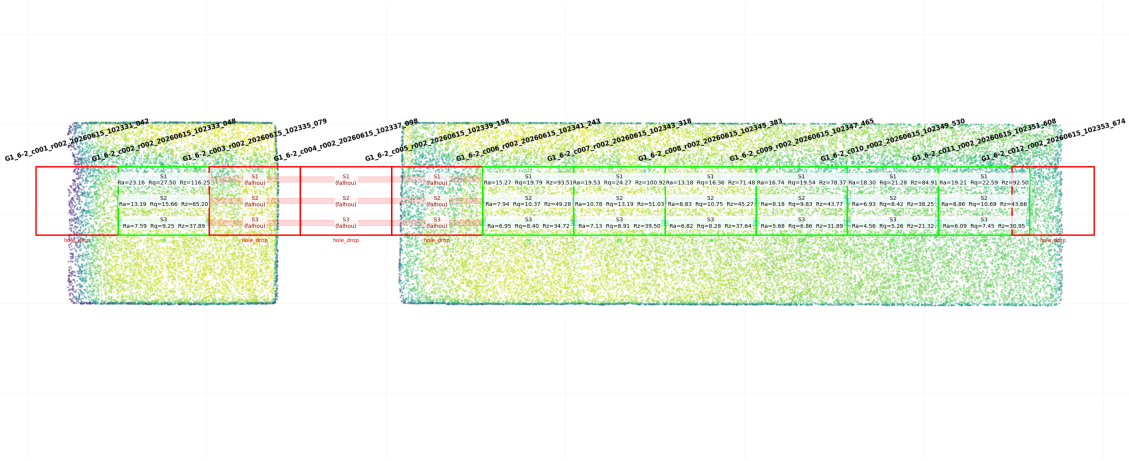


Figure 3.17: Spatial alignment between RGB images and the ATOS Q Cloud Point.

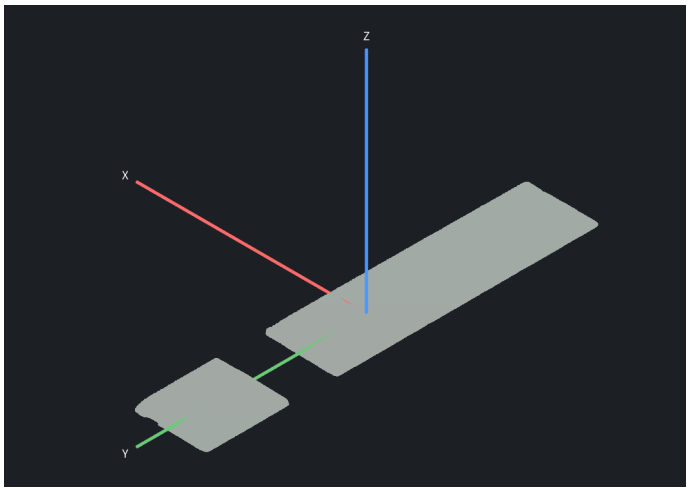


Figure 3.18: Coordinate system adopted for the aligned ATOS Q mesh Cloud Point.

3.6.2 Exclusion Criteria

Two categories of observations were excluded from the dataset prior to model training.

Thin specimen faces - Four faces with thickness below 4 mm were excluded because the vertical FOV of the microscope ($\mathrm{FOV}_Y = 3.81\,\mathrm{mm}$) exceeds the face height, meaning that a single image may capture regions beyond the physical cut-edge boundary. Since no automatic cropping to the face extent is applied during acquisition, these faces were removed entirely to avoid introducing boundary artefacts into the spectral features.

Point-cloud voids - Images whose physical footprint overlaps a localised void in the aligned mesh, arising from specular reflections, shadowing, or surface contamination during scanning, are automatically detected and excluded. The void detection is performed by projecting the mesh onto a uniform planar grid and identifying cells with insufficient point density. Any image for which the overlapping void exceeds a minimum area threshold is removed from the dataset, as the corresponding roughness profile would be incomplete and unreliable.

3.6.3 Roughness Label Assignment

The roughness parameters R_a, R_q, and R_z are extracted from each strip following the same filtering and evaluation procedure described in Section 3.3, with the exception that only two evaluation lengths are used instead of five. This reduction is imposed by the physical length of the microscope FOV along Y, which limits the available profile length. As a consequence, the ISO 9013 parameter R_{z5}, defined as the arithmetic mean of five independent sampling lengths, cannot be computed from a single image acquisition and is therefore not included as a regression target.

The raw ATOS Q roughness values are multiplied by the calibration factor $k = 1.49$ derived in Section 3.3.2, which corrects for the systematic amplitude underestimation of the structured-light scanner relative to the contact profilometer. The calibrated values are stored as the regression

targets for each of the three strip positions. For images where a strip overlaps a point-cloud void, the corresponding label is stored as a missing value and excluded from training.

3.6.4 Data Pre-processing and Final Dataset

The raw feature CSV produced by the spectral extraction and roughness label assignment pipelines is processed by a dedicated pre-processing script prior to model training. Four sequential operations are applied.

Removal of unlabelled observations - Rows for which no valid roughness value is available at any of the three strip positions are removed. This affects images whose associated region of the point cloud was entirely voided or for which roughness profile extraction failed at all three positions.

Face-level mean roughness - Three summary columns $\overline{R_a}$, $\overline{R_q}$, and $\overline{R_z}$ are added, computed as the arithmetic mean of the available per-strip values:

$$\overline{R_a} = \frac{1}{|\mathscr{P}|} \sum_{j \in \mathscr{P}} R_{a,j} \tag{3.34}$$

where $\mathscr{P} \subseteq \{1,2,3\}$ is the set of strip positions for which a valid measurement exists. This ensures that all retained observations carry a valid face-level roughness estimate even when only one strip position is available.

Assist gas encoding - The categorical assist-gas column is binarised: $gas\_encoded = 0$ for oxygen and $gas\_encoded = 1$ for nitrogen.

Material one-hot encoding - The four material categories are expanded into binary indicator columns: aluminium, steel, stainless steel, and galvanised steel. One-hot encoding is adopted because the material categories carry no intrinsic ordinal relationship.

Normalised image position. A scalar feature p_{norm} encodes the physical position of each image along the specimen thickness direction. It is defined as the ratio of the image centre coordinate x_c to the face half-thickness $t/2$:

$$p_{\text{norm}} = \frac{x_c}{t/2} \tag{3.35}$$

where x_c is derived from the acquisition metadata. This feature provides explicit positional context along the thickness, which correlates with the variation in striation regularity observed across the cut depth.

The composition of the final processed dataset is summarised in Table 3.8.

Table 3.8: Composition of the final training dataset.

Property	**Value**
Overall	
Total observations	1041
Specimen faces	30
Distinct cutting conditions	15
Observations by material	
Aluminium	539
Steel	326
Stainless steel	116
Galvanised steel	60
Observations by assist gas	
N_2	715
O_2	326
Feature columns	
Process parameters	8
Material one-hot encoding	4
Spectral descriptors (global + p/ strip)	18
Roughness labels (R_a, R_q, R_z)	12

The processed dataset is stored as a semicolon-delimited CSV file and used directly as input to the model training pipelines described in Section 3.7.

3.7 Deep Learning Architectures

Three architectures were implemented to contrast explicit feature engineering with end-to-end image learning. All of them were trained for single-output regression of R_a, R_q or R_z. The design follows the target application, an on-field and real-time cut-quality inspection in which the operator acquires an RGB image of the cut edge. For this reason, two of the three models consume that image directly, while the tabular input, when present, has a dimension N that depends on the feature set under evaluation.

3.7.1 Feature-driven MLP

The MLP is the interpretable and feature-based baseline. It receives the engineered feature vector and processes it through two fully connected layers of 64 and 32 neurons. Each hidden layer is

followed by Batch Normalization [25], a ReLU activation [39], and dropout [50] ($p = 0.3$), the network ends in a single linear output. The resulting architecture is shown in Figure 3.19.

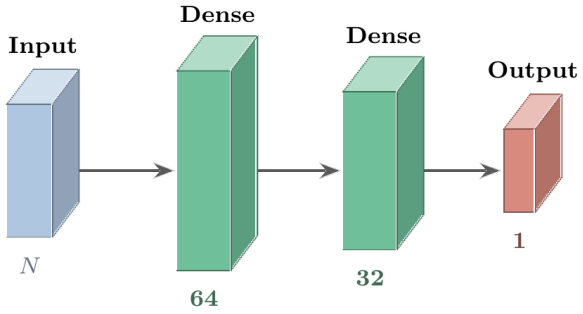


Figure 3.19: Architecture of the feature-driven MLP.

3.7.2 Image-driven ResNet-18

This model regresses roughness directly from the RGB edge image using a ResNet-18 backbone [16], pretrained on ImageNet [6]. ResNet-18 comprises an initial convolutional stem followed by four stages of residual blocks. The residual, or skip, connections let the gradients flow through the deep network and allow the model to learn hierarchical representations of the edge morphology, from the local striation texture in the early stages to the global patterns in the deeper ones. To adapt the ImageNet features to the cut-edge domain without overfitting the small dataset, the backbone is kept frozen except for its last residual stage, which is fine-tuned. A Global Average Pooling layer then collapses the final feature maps into a 512-dimensional embedding. This embedding is fed to a regression head of 64 and 32 neurons, with Batch Normalization, ReLU and dropout after each hidden layer, ending in a single linear output. The complete architecture is shown in Figure 3.20.

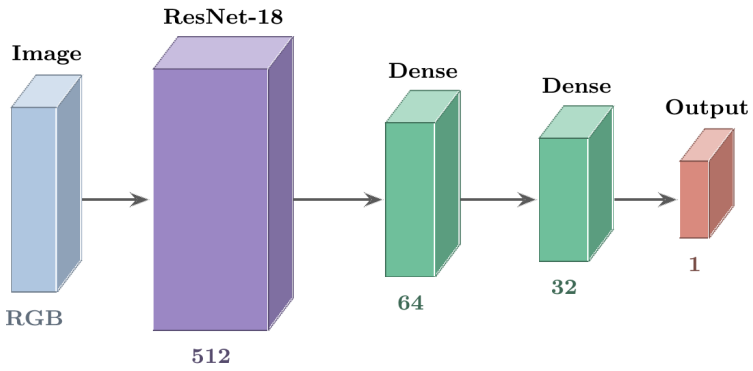


Figure 3.20: Architecture of the ResNet-18 baseline.

3.7.3 Multimodal ResNetFusion

The fusion model combines both sources of information by late fusion. The image-driven branch is the same ResNet-18 backbone, producing a 512-dimensional embedding. In parallel, a feature-driven branch maps the engineered features through fully connected layers of 128 and 64 neurons to a 64-dimensional vector. The two embeddings are then concatenated, giving 576 dimensions, and passed to a regression head of 64 and 32 neurons, with Batch Normalization, ReLU and dropout throughout the dense blocks. This lets the model complement the learned visual texture with process context that a normalized image cannot convey, most notably the plate thickness and the cutting speed, which set the absolute scale of the striations. The complete architecture is shown in Figure 3.21.

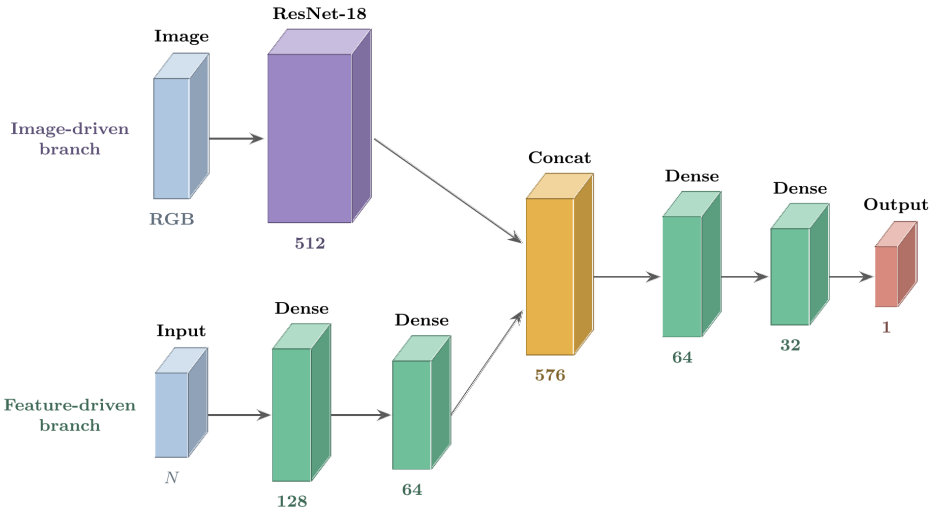


Figure 3.21: Architecture of the ResNetFusion multimodal model.

3.7.4 Training and Validation Protocol

Consistent with the deployment scenario, in which the device produces a single reading from a single acquisition, the unit of training and evaluation is the individual image. The models were trained and scored per image, so that the training distribution matches what the device observes on-field.

- **Optimization -** All models were trained with the Huber loss [18], which is robust to the occasional outliers in the roughness measurements. Optimization used the AdamW optimizer [35], with a decoupled weight decay of 10^{-4} and an initial learning rate of 10^{-3}, together with a ReduceLROnPlateau scheduler (factor 0.5, patience 8). Training ran for up to 150 epochs with early stopping (patience 25), gradient clipping at 1.0 and a batch size of 32. A dropout of 0.3 was applied as an important regularizer, given the limited dataset size.
- **Target standardization -** The three roughness parameters differ by roughly an order of magnitude, with R_z being far larger than R_a. When trained on their raw scale, a single Huber

transition point and a single learning rate cannot suit all three at once, and the larger-scale target dominates the gradient and converges poorly. The regression target was therefore standardized with a z-score, using the training statistics only. All reported errors are then inverted back to micrometres, so that the three parameters are learned under a comparable loss scale while remaining physically interpretable.

- **Face-balanced weighting -** Each cut face is described by an unequal and non-constant number of images, following its acquisition grid of columns and rows, and ranging from a handful to several dozen. If left unaddressed, the faces photographed more often would dominate the loss and bias the model towards their cutting conditions. Each image was therefore weighted by $1/n_{\text{face}}$, so that every face, and hence every cutting condition, contributes equally regardless of how many times it was imaged.

- **Leakage-safe, stratified cross-validation -** Images from the same face are strongly correlated, since they share the same cutting parameters and some images have overlap. A conventional per-sample split would place near-duplicate images in both the training and the test sets, yielding over-optimistic and unrealistic scores. This overlap, however, is not discarded but exploited. The partially overlapping images of a face act as a natural augmentation of the framing and placement variability that the handheld device will encounter in production, so all images are retained. This, in turn, makes grouping by face mandatory in order to avoid leakage. Accordingly, the data were partitioned by face using Stratified Group K-Fold with five folds. The groups were defined by the face identifier, and the folds were stratified by R_a quantile bins so that each fold spans the full roughness range. A fraction of the training faces was held out for validation and early stopping. Because the number of independent faces is small, a single split is unstable, so the cross-validation was repeated over three random seeds. The results are reported as the mean $\pm$ one standard deviation across folds and seeds, together with the pooled out-of-fold predictions. This protocol emulates the realistic requirement that the model must generalize to an entirely new cutting trial, never seen during training.

3.7.5 Feature Sets

The tabular input feeds the MLP and the feature-driven branch of the fusion model, whereas the image-only ResNet-18 does not use it. To assess the contribution of each source of information, the tabular features were organized into four nested sets, each evaluated independently across the three roughness parameters.

Process - The cutting parameters together with the material and the normalized position of the image within the thickness, namely the thickness, the laser power, the cutting speed, the gas pressure, the focal position, the assist gas, the material identity and the normalized vertical position. These are the quantities known at the moment of the acquisition, without looking

at the surface content, and they form the baseline that a process-monitoring approach could reach.

Spectral - The global spectral descriptors extracted from the striation pattern, that is the wavelength, the strength and the confidence of the three dominant peaks of the whole-image spectrum. This set contains only information derived from the image, and it tests whether the striation signature on its own is sufficient to estimate the roughness.

Spectral and Process - The union of the two previous sets. Since it differs from the Process set only by the addition of the spectral descriptors, the comparison between the two isolates the added value of the spectral information over the process baseline.

All - The complete set, which extends Spectral and Process with the strip-level descriptors, that is the wavelength, strength and confidence measured separately in the three horizontal strips. It therefore tests whether the finer spatial detail of the local descriptors adds anything beyond the global description and the context.

The sets are nested, so that Process and Spectral are both contained in Spectral and Process, which is in turn contained in All. This structure makes the ablation directly interpretable. The step from Process to Spectral and Process measures the contribution of the spectral features, while the step from Spectral and Process to All measures the contribution of the local, strip-level detail. The comparison between the four sets, reported in Section 4.3, indicates which representation is most useful for each roughness parameter.

Chapter 4

Results and Discussion

This chapter presents and discusses the results obtained at each stage of the methodology described in Chapter 3. The chapter is organised into three main parts. The first addresses the roughness characterisation pipeline, covering the initial validation of the ATOS Q structured-light scanner against the Mitutoyo SJ-210 contact profilometer and the derivation of the amplitude calibration factor. The second examines the spectral features extracted from the RGB microscope images, analysing their spatial stability across the specimen set and their relationship with the roughness parameters. The third evaluates the three deep learning architectures trained to predict surface roughness directly from image data, comparing their performance across feature sets and materials. A consolidated discussion is presented in Section 4.4, which synthesizes the main findings and their implications for the target application.

4.1 Roughness Extraction and Calibration

This section quantifies the systematic amplitude bias between the ATOS Q structured-light scanner and the Mitutoyo SJ-210 contact profilometer, and derives the calibration factor that renders the two instruments consistent.

As a first assessment, the roughness parameters extracted from the ATOS Q point cloud were compared directly with the SJ-210 reference at five strip positions across the specimen thickness, without any positional alignment. The strip positions were placed at fixed, equally spaced locations anchored to a reference offset determined by a prior search that minimised the overall metric error, as described in Section 3.3.1. The scanner reproduces the general shape of the contact profilometer profiles over the full evaluation length, confirming that the point cloud density is sufficient to capture the overall surface form. However, the ATOS Q systematically underestimates peak amplitudes across all strip positions, a direct consequence of the finite lateral point spacing of the structured-light sensor: even after resampling to a finer grid by interpolation, the sharp peaks and deep valleys of the laser-cut striations are not captured at the original acquisition resolution and cannot be recovered by interpolation alone. The quantitative comparison is summarised in Table 4.1.

The same specimen face was then scanned four times under identical conditions to assess the repeatability of the ATOS Q measurements. The standard deviations across the four repeated scans, reported in Table 4.2, remain below 0.25 µm for R_a and R_q and below 1.2 µm for R_z across all positions. This high scan-to-scan consistency confirms that the bias is stable and systematic, which is a necessary condition for it to be correctable by a single multiplicative factor.

Before the calibration is applied, the strip positions of the ATOS Q comb must be registered to the corresponding SJ-210 measurements. The two-stage alignment procedure described in Section 3.3.2 was executed independently for each of the four repeated acquisitions, and Figure 4.1 shows the resulting cost landscape in the $(\text{gap}_x, \text{gap}_y)$ search space for a representative acquisition. The minimum, marked by the red star, falls in a narrow, well-defined region of the search space, which supports the reliability of the shape-based registration before the amplitude correction is applied.

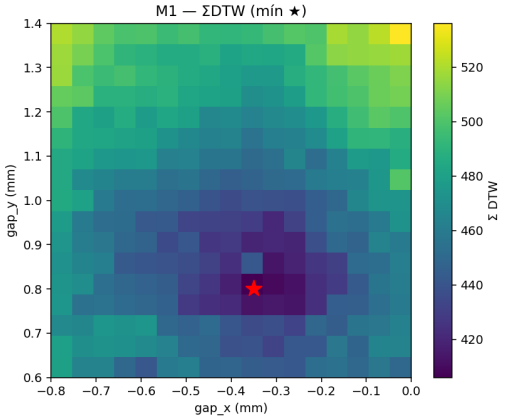


Figure 4.1: Cost landscape in the $(\text{gap}_x, \text{gap}_y)$ search space for the representative calibration acquisition of the reference specimen. The red star marks the optimal alignment adopted for the corresponding mesh.

With the registration in place, the calibration factor $k = 1.49$ was applied uniformly to all raw ATOS Q roughness values. Table 4.2 reports the raw and calibrated values alongside the SJ-210 reference, and Table 4.3 summarises the error metrics before and after the correction. The MAPE drops from approximately 32 % to 9 % across all three parameters, demonstrating that the calibrated ATOS Q labels are sufficiently accurate to serve as regression targets for the deep learning models. Figures 4.2 to 4.6 illustrate the effect of the calibration strip by strip, with the SJ-210 reference plotted in solid black and the four repeated ATOS Q acquisitions shown together, one as a solid coloured line representing the acquisition used for alignment and the remaining three as dashed lines to convey the scan-to-scan consistency.

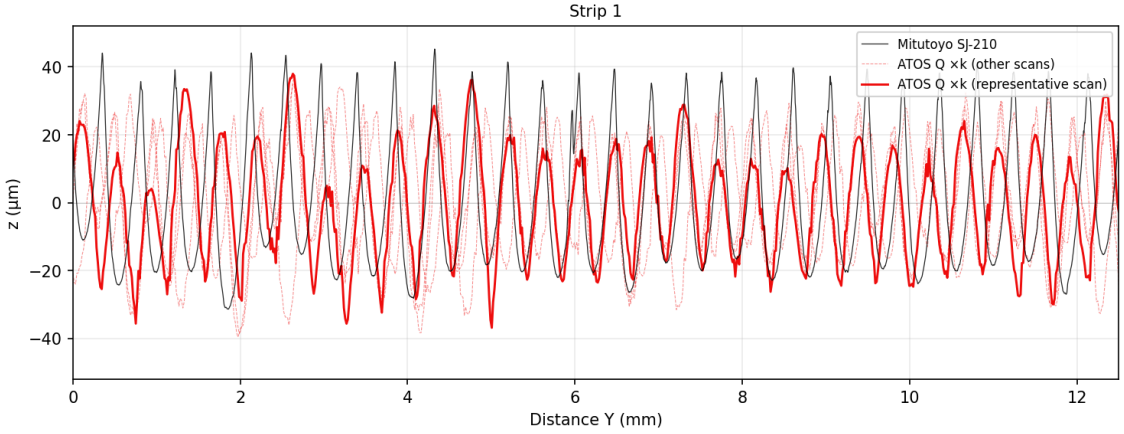


Figure 4.2: Calibrated ATOS Q profiles versus SJ-210 reference, Strip 1.

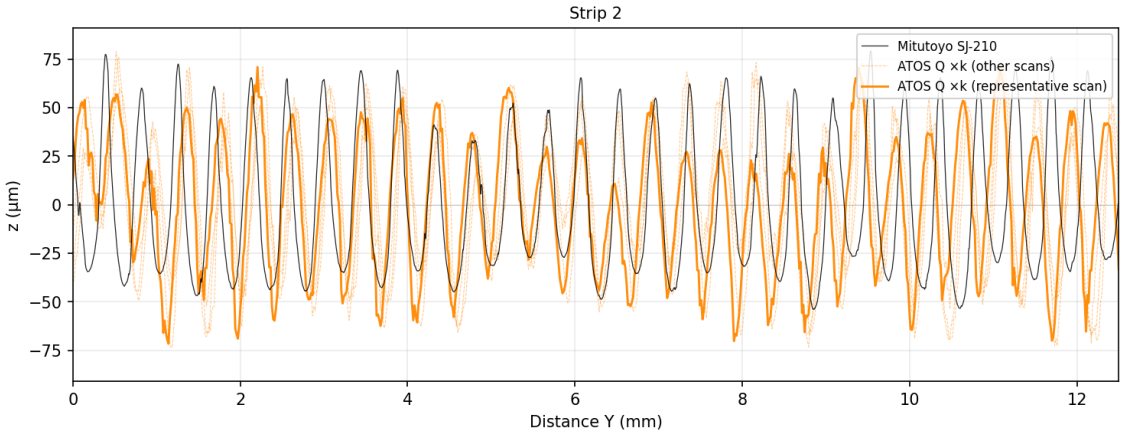


Figure 4.3: Calibrated ATOS Q profiles versus SJ-210 reference, Strip 2.

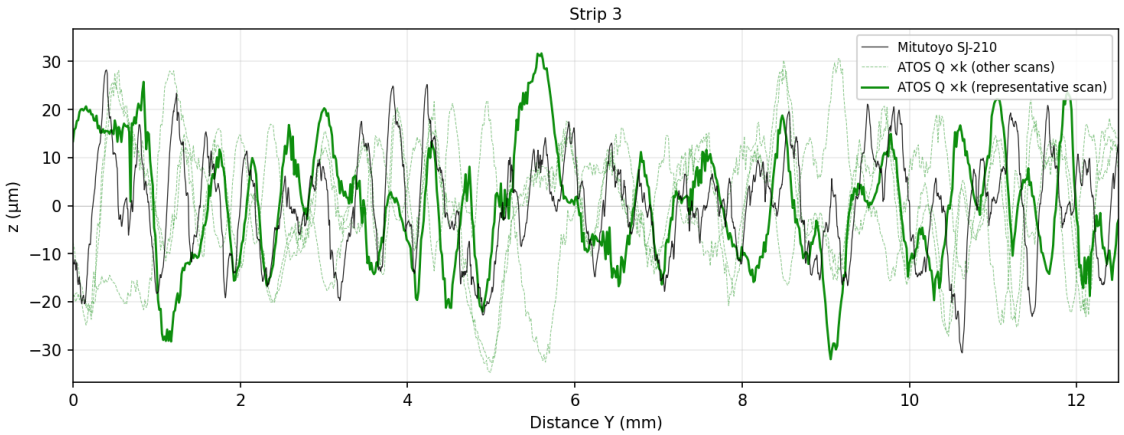


Figure 4.4: Calibrated ATOS Q profiles versus SJ-210 reference, Strip 3.

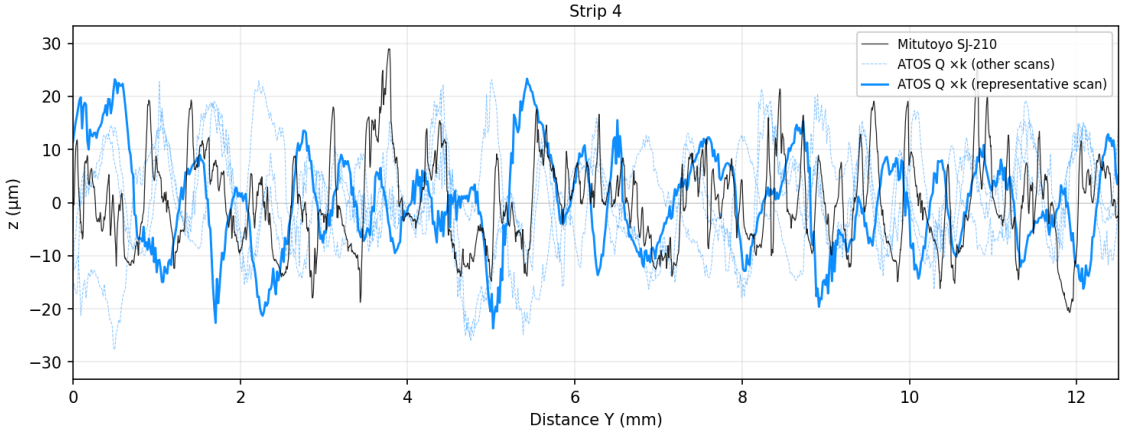


Figure 4.5: Calibrated ATOS Q profiles versus SJ-210 reference, Strip 4.

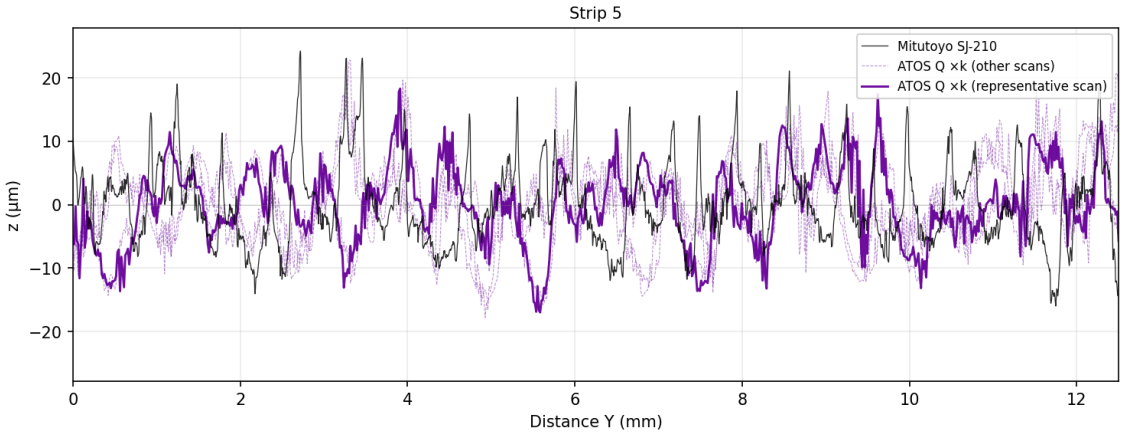


Figure 4.6: Calibrated ATOS Q profiles versus SJ-210 reference, Strip 5.

Table 4.1: ATOS Q vs. Mitutoyo SJ-210 roughness parameters without positional alignment.

Strip	**ATOS Q** (µm)			**SJ-210** (µm)			Δ (µm)		
	R_a	R_q	R_z	R_a	R_q	R_z	ΔR_a	ΔR_q	ΔR_z
Strip 1	12.51	14.48	57.06	16.22	18.96	68.95	−3.72	−4.48	−11.89
Strip 2	25.42	28.82	102.68	30.94	35.41	122.71	−5.53	−6.58	−20.02
Strip 3	8.19	9.81	40.13	8.54	10.45	44.41	−0.34	−0.64	−4.28
Strip 4	4.72	5.68	25.13	6.80	8.33	39.14	−2.07	−2.65	−14.02
Strip 5	4.01	5.09	24.00	4.94	6.28	32.96	−0.93	−1.19	−8.96
MAE (µm)	2.52	3.11	11.83						
RMSE (µm)	3.15	3.80	12.94						
MAPE (%)	18.8	19.8	21.2						

Table 4.2: ATOS Q scan repeatability (standard deviation over four repeated scans), raw and calibrated roughness values, and SJ-210 reference ($k = 1.49$).

Strip	Std dev (µm)			ATOS Q raw (µm)			ATOS Q $\times k$ (µm)			SJ-210 (µm)		
	σ_{R_a}	σ_{R_q}	σ_{R_z}	R_a	R_q	R_z	R_a	R_q	R_z	R_a	R_q	R_z
Strip 1	0.135	0.175	1.189	9.253	10.770	42.133	13.788	16.047	62.779	16.224	18.959	68.950
Strip 2	0.083	0.127	1.092	19.828	23.093	90.235	29.544	34.409	134.451	30.942	35.407	122.707
Strip 3	0.133	0.084	1.152	6.707	8.099	32.919	9.994	12.068	49.049	8.539	10.453	44.405
Strip 4	0.230	0.225	1.023	4.950	5.945	25.708	7.375	8.858	38.306	6.796	8.331	39.144
Strip 5	0.029	0.027	0.577	3.303	4.084	19.576	4.922	6.086	29.169	4.942	6.283	32.958

Table 4.3: Roughness error metrics before and after applying the calibration factor $k = 1.49$ and positional alignment.

Parameter	MAE (µm)		RMSE (µm)		MAPE (%)	
	Raw	Cal.	Raw	Cal.	Raw	Cal.
R_a	4.680	1.178	6.026	1.438	32.1	9.1
R_q	5.488	1.250	6.852	1.575	32.8	8.6
R_z	19.518	5.437	21.284	6.521	33.2	8.5

4.2 Spectral Analysis of Cut-Edge Images

This section evaluates the spectral features extracted from the RGB microscope images of the cut edge. The analysis is organised in four parts. The physical behaviour of the striation wavelength across the cutting parameters and along the thickness of the cut is first characterised, since it defines the expectation against which the extractor should be assessed. The effect of the pipeline on the recovered wavelength is then measured by comparison with a naive FFT baseline. The two supporting spectral descriptors, peak strength and peak confidence, are examined next. Finally, the relationship between the global descriptors and their local counterparts is analysed, together with the overall correlation structure that motivates the feature sets used in Section 4.3.

4.2.1 Variation of the Wavelength within a Face

Within a single face, the calibrated roughness varies substantially along the thickness of the cut, since the rough side of the specimen is on one edge of the face and the smoother side on the opposite edge. The question addressed in this subsection is whether the dominant striation wavelength follows that same variation, tracking the roughness from the rough side to the smoother one, or whether it remains essentially the same along the thickness of the face regardless of how rough that region is.

To compare positions across faces of different thicknesses on a common scale, each photograph is described by its normalized vertical position on the face, p_{norm}, defined in Section 3.6.4.

This descriptor is comparable between a thin and a thick plate, and separates the rough side of the cut from the smoother one regardless of the total number of images acquired for the face. The images of the dataset were sorted by p_{norm} and grouped into four quartiles of equal population, so that each quartile contains approximately one quarter of the photographs and covers a specific range of vertical positions on the face. Table 4.4 reports the intervals of p_{norm} defining each quartile, together with the mean wavelength and the mean calibrated roughness inside them.

Table 4.4: Mean wavelength variation over roughness

Quartile of p_{norm}	**Q1**	**Q2**	**Q3**	**Q4**
Interval of p_{norm}	$+0.52,\ +0.75$	$+0.09,\ +0.52$	$-0.37,\ +0.09$	$-0.75,\ -0.37$
λ (mm)	0.616	0.660	0.626	0.640
R_a (µm)	27.1	12.0	10.7	7.5

The contrast between the two rows is immediate. The calibrated roughness varies by a factor of 3.6 across the thickness of the face, from 27.1 µm on the rough side to 7.5 µm on the smoother one, whereas the wavelength stays within a 4% margin of its mean value across all four quartiles. The correlation between the wavelength and the normalized position, computed directly at the level of the individual photograph, is $r = 0.001$, essentially zero. The transition from the rougher to the smoother region of the cut therefore appears as a reduction in the depth of the striations rather than in their spacing. The grooves keep the same pitch and simply become shallower along the thickness of the face. Roughness and the striation wavelength encode different aspects of the cut, the roughness reflecting the amplitude of the surface texture and the wavelength reflecting its periodicity, and they do not track each other within the same face.

The analysis reported above describes how the wavelength varies along the thickness of a face, from its rough side to its smoother one. The wavelength also shows some variation between faces of different materials and cutting conditions, but modest and partially confounded, since the material, the assist gas and the cutting regime are aligned in the dataset. At the level of face, the wavelength is longest for aluminium, at 0.70 mm, which is cut with nitrogen at high speed, and shortest for carbon steel, at 0.57 mm, which is cut with oxygen at low speed, with the other two materials falling in between. Among the individual process parameters, only the cutting speed shows a moderate positive association with the wavelength at level face, ($r = 0.32$), consistent with the physical expectation that faster cuts allow less time for fine striations to form. The associations with the plate thickness and with the calibrated roughness are weak at face level and are reported quantitatively in the correlation matrix of Section 4.2.4. These observations are compatible with the earlier conclusion that the wavelength is essentially a signature of the cutting regime rather than a predictor of the roughness magnitude.

Two consequences follow from these observations. First, the wavelength cannot on its own be a strong predictor of roughness, since the same value is observed in regions whose roughness differs by a factor of more than three, which anticipates the modest performance of purely spectral

feature sets in Section 4.3. Second, all images of a face taken at the same normalized vertical position should return essentially the same wavelength, since they observe the same region of the cut and differ only in their location along the length of the specimen. This second consequence provides the physical baseline against which the reproducibility of the extractor is evaluated in the next subsection.

4.2.2 Effect of the Spectral Pipeline on the Extracted Wavelength

The previous subsection established that photographs of a given face taken at the same normalized vertical position observe the same region of the cut and should therefore return the same wavelength. This provides a direct way to measure the reliability of the extractor. Photographs are grouped by face and by their position on the face, and any dispersion of λ inside such a group is attributed either to a genuine instability of the process along the length of the cut, which is expected to be small, or to noise introduced by the extractor itself.

To measure how much of that dispersion comes from the extractor, each photograph was processed by two different extractors and the results compared. The first is a naive rFFT applied to the row-averaged intensity profile, without flat-field correction, without the Harmonic Product Spectrum, and without the harmonic-suppression stage. The second is the pipeline described in Section 3.5.

Two aspects were assessed. The first is the frequency with which the naive extractor captures a harmonic of the true wavelength rather than the fundamental. This was quantified by the ratio between the naive and the pipeline wavelength for each photograph, using a tolerance of 10 % around each of the multiples $\frac{1}{3}$, $\frac{1}{2}$, 2 and 3. Table 4.5 summarises the resulting distribution over the 1041 photographs of the dataset. The naive extractor coincides with the pipeline in 36 % of the images and falls on one of the four harmonic multiples in 23 % of them, with the classical sub-harmonic $\lambda/2$ accounting for the largest share at 18 %. The remaining 40 % correspond to disagreements that do not fall close to an integer multiple and are dominated by spurious peaks introduced by illumination gradients and reflections. The pipeline, by construction, avoids all of these failure modes.

Table 4.5: Distribution of the ratio between the naive and the pipeline wavelength across the dataset.

Ratio (naive / pipeline)	**n**	**%**
Match ($\sim 1\times$)	377	36.2
Half ($\sim 0.5\times$)	184	17.7
One third ($\sim \frac{1}{3}\times$)	41	3.9
Double ($\sim 2\times$)	15	1.4
Triple ($\sim 3\times$)	4	0.4
Other	420	40.3

The second aspect is the typical spread of the extracted wavelength across photographs of the same face taken at the same normalized vertical position, quantified as the standard deviation of λ inside each group divided by its mean, and expressed as a percentage of the average value. Since the physical expectation is a constant wavelength inside such a group, this spread is a direct measure of the reliability of the extractor. Its median value over the 84 groups with at least three photographs is 31 % for the naive extractor and 22 % for the pipeline, with the pipeline improving on the naive baseline in 67 of those 84 groups. Figure 4.7 overlays the two distributions and makes the shift visible.

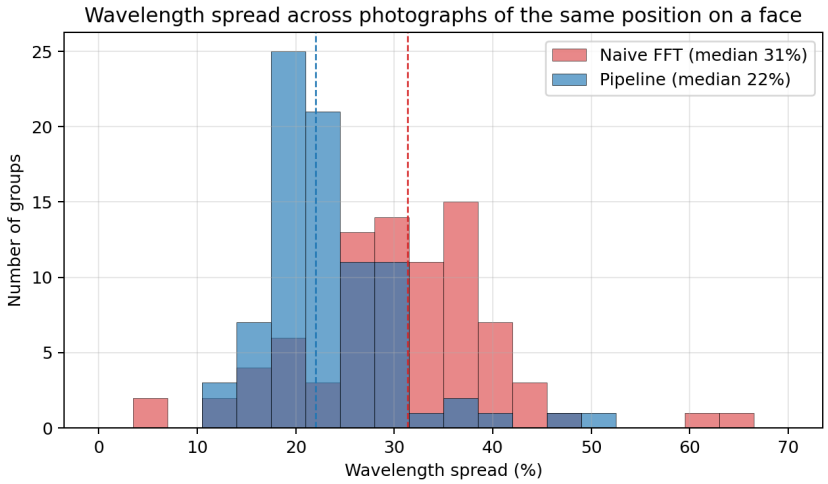


Figure 4.7: Distribution of the typical spread of the extracted wavelength across photographs of the same face taken at the same normalised vertical position

.

The reduction in spread, from 31 % to 22 %, is moderate. Its meaning is, however, qualitatively distinct from a simple noise reduction, because the pipeline also removes the harmonic captures that would otherwise enter the dataset with values equal to half or twice the true wavelength and would act as heavy outliers during the training of the downstream models. The pipeline therefore contributes on two fronts, a moderate reduction of the dispersion and a systematic elimination of the harmonic failure modes.

Some dispersion nevertheless remains, with the pipeline still delivering a typical spread close to 22 %. This residual variability is attributed to three factors of the acquisition setup, namely the uncontrolled ambient illumination, the specular character of the cut metallic surfaces, and the narrow FOV of the digital microscope, which frames only a small region of the face and captures the striation pattern with a different phase and different fractions of complete grooves in each shot. Under controlled illumination and a sensor of higher spatial resolution, the same pipeline would be expected to deliver a substantially lower residual dispersion. The reported values reflect the compromise imposed by the industrial conditions of the acquisition.

4.2.3 Peak Strength and Peak Confidence

The peak strength S and the peak confidence C complement the wavelength as descriptors of the striation pattern. The strength quantifies how much the dominant spectral peak stands out above the background of the spectrum, and the confidence quantifies how sharply that peak is defined against its immediate spectral neighbourhood. The two descriptors capture, respectively, the prominence and the definition of the periodic texture.

The peak strength shows a moderate positive association with the calibrated roughness at face level, with a Pearson correlation of $r = 0.30$ against R_a. This association is however largely indirect, since the pronounced striations that give rise to a high S are also produced by the cutting of thick plates at low cutting speeds, a regime that yields high roughness values on its own. When the plate thickness and the cutting speed are held constant through a partial correlation, the association drops to $r_{\text{partial}} = 0.21$. The peak strength is therefore an image-derived proxy for the cutting regime rather than a direct predictor of amplitude, but it carries a component that is not fully explained by the process parameters.

The peak confidence plays a different role. Along the thickness of the face, its mean stays essentially unchanged from the rough side to the smoother one, so it does not resolve the systematic gradient in the way that was initially expected. At the level of the individual photograph, however, C varies over more than one order of magnitude, from 0.28 to 10.72, separating images where the striation pattern is well defined from images where the pattern is faint or interrupted. Its value is therefore that of a per-image quality indicator, telling the model how much confidence to place in the spectral estimate extracted from a given photograph, rather than a systematic descriptor of the cut geometry. The two descriptors are also moderately related to each other at photo level, with $r = 0.62$ between S and C, since a pronounced peak tends to also be sharp, which is a natural consequence of their definitions.

Figure 4.8 illustrates the two descriptors on the two extreme cases of the dataset, one photograph with a prominent and sharply defined striation pattern and one photograph without a periodic pattern. In the first case the dominant peak stands well above the median of the spectrum (high S) and is clearly isolated from its neighbours (high C). In the second case the peak is barely distinguishable from the background of the spectrum, giving low values for both descriptors. The two features therefore quantify, in a directly observable way on the spectrum, how much and how cleanly the periodic signature of the striation dominates the amplitude spectrum of the image.

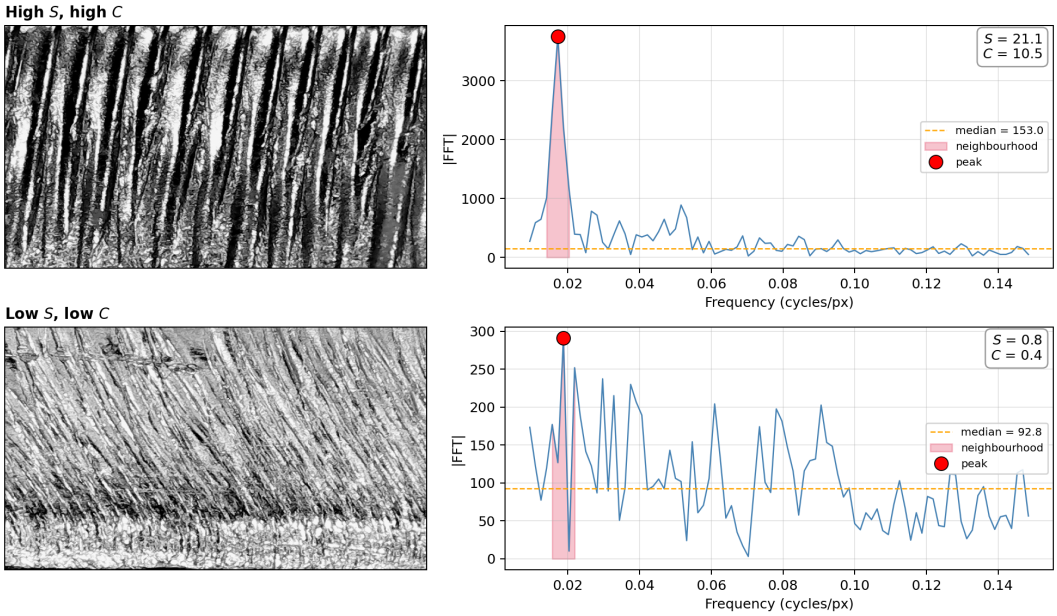


Figure 4.8: Illustration of the peak strength S and the peak confidence C on two extreme photographs of the dataset. In each row, the left panel shows the original image and the right panel shows its FFT amplitude spectrum, with the dominant peak marked in red, the median of the search band shown as a dashed orange line, and the immediate spectral neighbourhood highlighted in red.

4.2.4 Local versus Global Descriptors and Overall Structure

Before examining the overall correlation structure, a distinction between two analysis levels is necessary. Some of the relationships reported in this chapter are measured at the level of the individual photograph, over the 1041 images of the filtered dataset, and describe how the descriptors of a single photograph relate to each other. Others are measured at the level of the face, that is over the 30 specimens faces of the dataset, after aggregating the photographs of each face by the median of the spectral descriptors and by the mean of the calibrated roughness. Face level correlations describe how a whole cutting condition relates to the process parameters and to the roughness, which is the scale that matters for a model that has to generalise across new cuts. The heatmap discussed at the end of this subsection is computed at face level for this reason.

The extraction pipeline computes the three descriptors both on the whole image and separately on three horizontal sub-regions of the frame, one capturing the upper part of the FOV, one the middle and one the lower part. The intended purpose of these local descriptors is to expose the vertical evolution of the striation pattern within each single photograph, since the rough and the smoother sides of the cut can appear in different regions of the frame. Whether these local descriptors carry information beyond the global ones is however an empirical question, and is examined here.

The correlation between the local and the global descriptors of the same photograph is only moderate. At photo level, across the 1041 images of the dataset, the upper sub-region and the global values correlate at $r = 0.24$ for the wavelength, $r = 0.28$ for the peak strength, and $r = 0.49$ for the peak confidence. The three wavelengths reported per photograph, that is λ_1, λ_2 and λ_3, are

also weakly correlated with one another, with $r = 0.19$ between λ_1 and λ_2 and $r = 0.03$ between λ_2 and λ_3, which confirms that the harmonic-suppression stage of the pipeline produces genuinely independent spectral components rather than harmonic repetitions of a single dominant peak. The local descriptors therefore capture something that the global ones do not, but they are also noisier, being computed on smaller image areas that contain fewer periods of the striation pattern.

Figure 4.9 summarises the overall correlation structure of the face-level dataset and provides the natural bridge to the regression models of Section 4.3. Three blocks are visible: the spectral descriptors, the process parameters, and the calibrated roughness. At face level, the roughness is strongly determined by the process, with $r = 0.86$ against the plate thickness and $r = -0.70$ against the cutting speed, and the three roughness parameters carry almost identical information, with correlations very close to unity between them. The block linking the spectral descriptors to the roughness is consistently weak. The wavelength is essentially uncorrelated with the roughness at face level, as had already been observed in Section 4.2.1. The peak strength shows the highest association among the three descriptors at face level, but this association is moderate at $r = 0.30$ and is further reduced when the process parameters are controlled for. The peak confidence at face level is weakly correlated with all the roughness parameters, consistent with its interpretation as a per-image quality indicator rather than as a systematic descriptor of the cut.

The frequency-domain descriptors therefore characterise the striation pattern well but have limited standalone power to predict the roughness magnitude, which is governed by the process parameters. This does not make them redundant. As descriptors of the surface texture, they provide the learning model with complementary information about the cut pattern, capturing regularities that the scalar process parameters do not express, and they are expected to help the network recognise cut-quality patterns even when, on their own, they are not strong roughness predictors. The usefulness of each combination of descriptors is quantified directly in the next section, through the performance of the regression models trained on the four different feature sets.

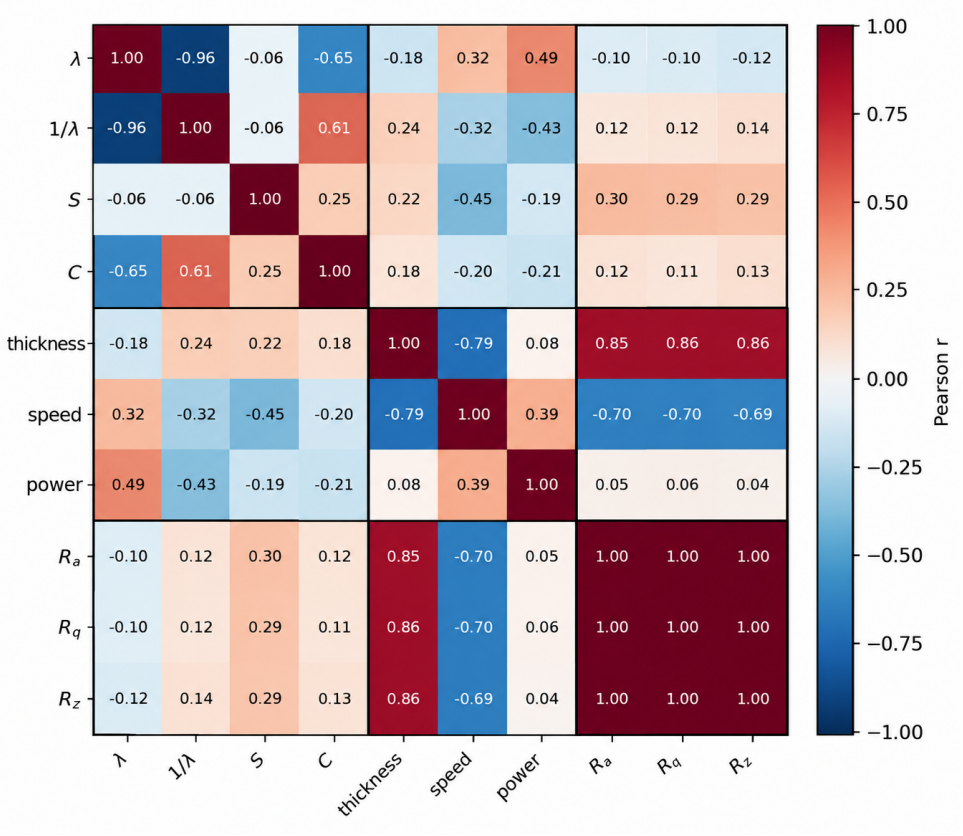


Figure 4.9: Correlation matrix computed at the level of the face over the final dataset.

4.3 Roughness Prediction with Deep Learning Models

This section reports the performance of the three architectures introduced in Section 3.7 for the prediction of the three roughness parameters. The evaluation is organised in three parts. The training behaviour is first inspected to confirm that the reported figures reflect genuine generalisation rather than optimisation artefacts. The general models are then compared under the four feature sets defined in Section 3.7.5, and their performance is discussed by architecture. Finally, the effect of specialising a model on a single material is examined.

4.3.1 Training Behaviour

All models were trained under the leakage-safe cross-validation protocol described in Section 3.7, with five folds by face repeated over three random seeds, and with early stopping on the validation loss. Figure 4.10 shows the training and validation loss curves for a representative run of each architecture, all trained on the R_a target. In the three cases the validation loss reaches a plateau rather than rising, and the epoch of minimum validation loss selected by early stopping is marked with a dashed grey line. The residual gap between the training and the validation curves reflects the expected difficulty of the task on a small dataset and does not translate into optimistic scores,

since all reported metrics are computed on faces held out of the training folds. The same qualitative behaviour is observed across the other targets and feature sets.

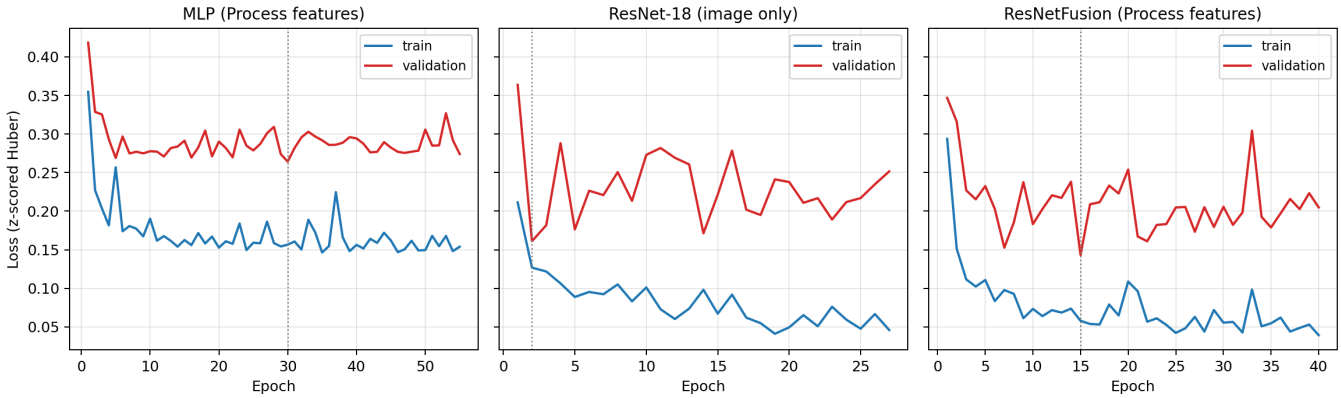


Figure 4.10: Training and validation loss curves for a representative run of each architecture, all trained on the R_a target. The dashed grey vertical line marks the epoch of minimum validation loss, which is the epoch selected by early stopping. The validation loss reaches a plateau in all three cases.

4.3.2 General Models: Performance per Image

Every model is trained and scored per image, in line with the deployment scenario in which a decision is made from a single acquisition. Performance is measured by three complementary metrics, the coefficient of determination R^2, the MAE, and the RMSE, the latter two given in micrometres. All values are reported as the mean over the five cross-validation folds and the three seeds. Table 4.6 gathers the results, which are discussed by architecture in the following.

Table 4.6: Per-image performance of each model and feature set, as the mean over the cross-validation folds and the three seeds. The MAE and RMSE are given in micrometres. The best value of each metric for each parameter is shown in bold.

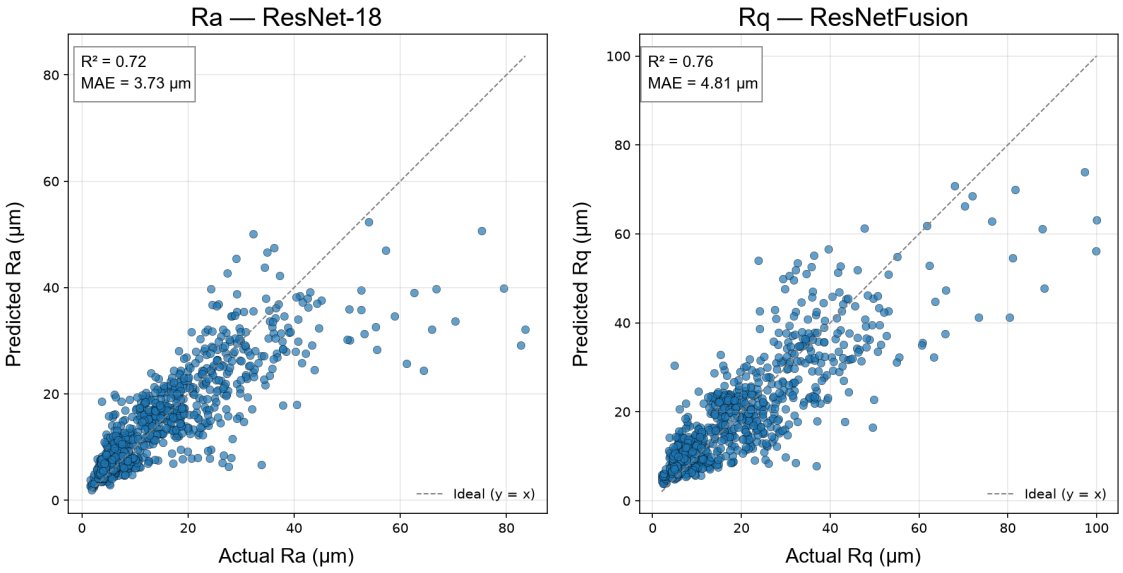
		R_a			R_q			R_z		
Model	Features	R^2	MAE	RMSE	R^2	MAE	RMSE	R^2	MAE	RMSE
MLP	Process	0.49	5.03	8.49	0.47	6.29	10.58	0.45	24.59	42.34
	Spectral	0.13	7.69	11.20	0.12	9.25	13.63	0.12	36.87	53.80
	Spectral+Process	0.50	5.17	8.42	0.48	6.12	10.39	0.48	23.75	41.31
	All	0.48	5.31	8.61	0.46	6.43	10.63	0.45	25.77	42.51
ResNet-18	Image only	0.73	3.76	6.19	**0.76**	4.63	**7.02**	0.74	19.26	29.12
Fusion	Process	0.74	3.77	**6.10**	0.75	**4.38**	7.15	**0.77**	**17.54**	**27.74**
	Spectral	0.67	4.03	6.57	0.74	4.63	7.31	0.75	19.39	28.99
	Spectral+Process	**0.74**	**3.70**	6.14	0.71	4.74	7.58	0.76	18.47	28.36
	All	0.69	3.95	6.49	0.74	4.68	7.34	0.74	19.49	29.44

Feature-driven MLP - The feature-driven MLP establishes how far the roughness can be predicted from the engineered descriptors alone. Using only the spectral features it barely exceeds a constant predictor, with a coefficient of determination around 0.12 for the three parameters, in direct agreement with the weak spectral-to-roughness correlations reported in Section 4.2.1. The wavelength does not vary with the local roughness along the thickness of the face, and the peak strength carries only an indirect signal that is mostly explained by the process parameters. When the process parameters and the normalized vertical position of the image are added to the input, the coefficient of determination rises to about 0.50 for the three parameters. The position feature is decisive on its own, because it indicates whether the image was taken near the rough side or near the smoother side of the cut, which is the dominant source of variation across images of the same face. Even in its best configuration, however, the MLP stays well below the image-based models, with a MAE close to 5 µm for R_a against about 3.8 µm for the models that receive the image.

Image-driven ResNet-18 - Feeding the raw image to a pretrained ResNet-18, with only its last residual stage fine-tuned, produces a large improvement over any tabular model. The image-only network reaches a coefficient of determination of 0.73, 0.76 and 0.74 for R_a, R_q and R_z, with a mean absolute error of about 3.8, 4.6 and 19 µm. It extracts substantially more information from the edge texture than the hand-crafted descriptors are able to encode, and it does so from the image alone, without requiring any tabular input at inference time. This makes it the most convenient model for the target application, in which the operator only needs to point the device at the cut edge.

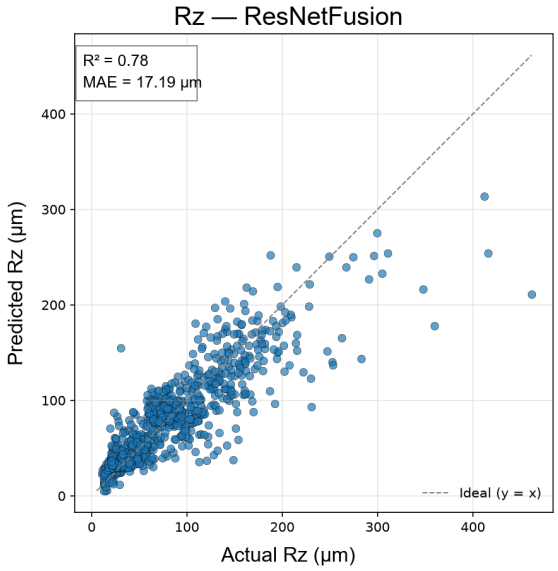
Multimodal ResNetFusion - The multimodal ResNetFusion adds the tabular context to the image branch. With the process parameters and the position feature in the tabular branch, the fusion matches the image-only model on R_a and R_q and becomes the best of all the models on R_z, reaching a coefficient of determination of 0.77 and a MAE of 17.5 µm, against 0.74 and 19.3 µm for the image-only ResNet. This gain is consistent with the nature of R_z, the peak-to-valley height, which is governed by the absolute scale of the deepest striations. That scale is set by the plate thickness and by the position along the depth of the cut, both carried by the tabular branch but not directly available to a normalized image. On R_a and R_q, where the image already captures most of the signal, the tabular context adds little and the two image-based models are equivalent.

Taken together, the image-only ResNet-18 is the reference model for R_a and R_q, as it matches the fusion while requiring only the image, whereas the fusion with the process and position context is preferred for R_z, where the scale information yields a small but consistent improvement. Figure 4.11 shows the predicted against the measured values for the selected model of each parameter, where the points lie close to the identity line across the full roughness range.



(a) R_a, ResNet-18.

(b) R_q, ResNet-18.



(c) R_z, ResNetFusion.

Figure 4.11: Predicted versus measured roughness for the selected model of each parameter, over the pooled out-of-fold predictions.

4.3.3 Single-Material Models

To assess whether a model specialised in a single material outperforms the general one, the three architectures were also trained on a single material. The analysis focuses on aluminium, which is the only material with enough data for a reliable evaluation. Steel, with only ten faces, produced estimates that varied strongly between runs and is therefore not reported. Given the small number of faces, the models were evaluated by leave-one-face-out, in which each face is held out in turn and predicted by a model trained on the remaining faces, and the image models are averaged

over three seeds. Each specialised model is compared with the general model restricted to the aluminium faces. Table 4.7 reports the per-image results, discussed by architecture.

Table 4.7: Per-image performance of the specialised single-material models on aluminium, evaluated by leave-one-face-out. The image models are the mean over three seeds, and the MLP, reported for its Spectral and Process set, is a single reproducible run. The MAE and RMSE are in micrometres.

	R_a			R_q			R_z		
Model	R^2	MAE	RMSE	R^2	MAE	RMSE	R^2	MAE	RMSE
MLP	0.65	4.98	7.52	0.64	6.15	9.38	0.65	23.66	35.49
ResNet-18	0.77	3.94	6.09	0.77	4.83	7.40	0.79	18.53	27.44
Fusion	0.78	3.84	5.99	0.79	4.70	7.14	0.79	18.45	27.38

Feature-driven MLP - On aluminium the feature-driven MLP reaches a coefficient of determination of about 0.65 with the Spectral and Process features, close to the general model evaluated on aluminium. A notable effect is that specialising on aluminium turns the purely spectral features from almost useless into informative, raising their coefficient of determination from about 0.13 in the general setting to about 0.35. This confirms once more that the spectral signature is conditional on the material. Even so, the tabular MLP stays clearly below the models that receive the image.

Image-driven ResNet-18 - The image-only ResNet-18 reaches a coefficient of determination of about 0.77 to 0.79 for the three parameters, stable across the three seeds, and does not overfit despite the limited number of images. This performance is essentially equal to that of the general model evaluated on the aluminium faces, which confirms two things. First, aluminium has enough images to support a material-specific model. Second, the general model already exploits the aluminium data as well as a dedicated model would.

Multimodal ResNetFusion - The multimodal ResNetFusion behaves like the ResNet, reaching about 0.78 to 0.79 for the three parameters, and matches the general model restricted to aluminium. Adding the tabular context to a model that is already trained on a single material brings no further gain, because fixing the material removes most of the ambiguity that the process and position features resolve in the pooled setting.

In summary, specialising on aluminium is viable and does not overfit, but it does not outperform the general model, which already predicts aluminium as well. The image-based models remain clearly ahead of the tabular MLP, and the small improvement that the fusion brings in the pooled study disappears once the material is fixed.

4.4 Discussion

The three preceding sections addressed the three technical building blocks of the target application, the derivation of reliable roughness labels from the structured-light scanner, the extraction of spectral descriptors from the RGB microscope images, and the deep learning models that predict roughness from those inputs. This section synthesises the main findings of each block, comments on their limitations, and draws the practical implications for the intended field-inspection use case.

4.4.1 Roughness Labels from the Structured-Light Scanner

The comparison between the ATOS Q and the Mitutoyo SJ-210 shows that the structured-light scanner reproduces the general shape of the reference roughness profile but systematically underestimates its amplitude across all measurement positions. This bias is repeatable across independent acquisitions of the same specimen, with standard deviations well below the systematic error, which is the necessary condition for a single multiplicative correction to be defensible. The adopted factor $k = 1.49$ reduces the MAPE from roughly one third to a value close to nine percent for the three roughness parameters, bringing the ATOS Q into practical agreement with the contact reference.

The residual disagreement after calibration is texture-dependent, being larger on the sharpest striations of the laser-entry region than on the smoother positions. This reflects the fundamental limit of a scalar correction against a bias whose magnitude depends on the local surface sharpness. A complementary observation confirms this point, since the same specimen scanned at a higher point density requires a smaller correction than the standard acquisition, showing that the factor is specific to the resolution of the acquisition and would be reduced by a denser scan.

The practical outcome of this block is that the ATOS Q, once calibrated, provides labels of sufficient accuracy to serve as regression targets for the downstream models. The calibration is nonetheless tied to the acquisition protocol adopted here, and a scanner of higher spatial resolution would relax the need for the correction while narrowing the residual disagreement observed on the sharpest textures.

4.4.2 Spectral Features from the RGB Microscope Images

The spectral analysis of the cut-edge images revealed two complementary findings. On the physical side, the dominant striation wavelength behaves as a constant of the cutting process. It does not vary along the thickness of the face, even when the local roughness changes by a factor of more than three between the rough side of the cut and the smoother one, and it varies only modestly between materials and cutting regimes, with a moderate positive association with the cutting speed as the only appreciable link to the process. The transition from the rough to the smoother region therefore manifests as a reduction in the depth of the striations rather than in their spacing, and the two descriptors of the periodic texture, its frequency and its amplitude, encode different aspects of the cut.

On the extraction side, the adopted pipeline improves on a naive Fourier transform in two qualitatively distinct ways. The reduction in the typical spread of the recovered wavelength across photographs of the same face is moderate, from about one third to about one fifth of the mean value, but the elimination of the harmonic captures that the naive extractor produced in roughly one quarter of the photographs is categorical. This second contribution is arguably the more important of the two, since a harmonic capture translates into a value equal to half or twice the true wavelength and would enter the downstream models as a heavy outlier. A residual dispersion nevertheless remains, attributable to the specular character of the cut metallic surfaces, to the uncontrolled ambient illumination, and to the narrow FOV of the digital microscope, which frames only a small region of the face with a variable phase and a variable number of grooves. Under a controlled illumination setup and a sensor of higher spatial resolution the pipeline would be expected to deliver a substantially lower residual dispersion, and the improvement over the naive baseline would concentrate almost entirely on the elimination of the harmonic outliers.

The correlations between the spectral descriptors and the roughness are consistently weak. The wavelength is essentially independent of the roughness magnitude, and the peak strength, which is the only descriptor with an appreciable association, is largely a proxy for the cutting regime rather than a direct predictor of amplitude. The peak confidence behaves as a per-image quality indicator, telling the model how much confidence to place in each individual spectral estimate. Taken together, the spectral descriptors characterise the striation pattern well but have limited standalone predictive power over the roughness magnitude, which is a direct consequence of the physical decoupling between periodicity and amplitude observed on the cut.

4.4.3 Deep Learning Models

The performance of the three architectures follows the picture that the spectral analysis anticipated. The feature-driven MLP trained on the purely spectral set barely exceeds a constant predictor, in direct agreement with the weak spectral-to-roughness correlations reported in the previous section. When the process parameters and the normalized vertical position of the image are added, the same architecture reaches about half of the total variance, driven above all by the position feature, which carries the strongest single signal for the local roughness inside a face. The tabular MLP nevertheless remains clearly below the image-based models, indicating that the engineered descriptors do not capture all the information that the raw edge texture contains. This effect is worth highlighting on its own, since it identifies the normalized vertical position of the image, a simple geometric descriptor computed directly from the acquisition, as the single most informative tabular input for the roughness prediction task.

The image-only ResNet-18 recovers a substantial part of that missing information. It reaches a coefficient of determination between 0.73 and 0.76 for the three parameters and a MAE close to 3.8 µm on R_a, which is a useful accuracy for the target application. The multimodal ResNetFusion matches the ResNet on R_a and on R_q and becomes the best of all the models on R_z, where the added tabular context of thickness and position resolves the absolute scale of the deepest striations, which is not directly available to a normalized image. The specialisation of a model on a single material

was shown to be viable and free of overfitting, but it does not outperform the general model, which already exploits the aluminium data as well as a dedicated one and covers the other materials at the same time.

For the intended field-inspection use case, in which the operator points a handheld device at the cut edge and expects a single reading, these results support two practical recommendations. The image-only ResNet-18 is the reference model for R_a and R_q, since it requires only the image at inference time and matches the fusion on those two parameters. The ResNetFusion with process context is preferred when R_z is the quantity of interest, because it exploits the plate thickness that a normalized image cannot express. In both cases, acquiring a small number of independent images of the same face and averaging their predictions would further reduce the impact of the residual per-image dispersion identified in the spectral analysis, exploiting the same aggregation effect that transforms noisy per-image spectral features into stable face-level descriptors.

Two limitations should be kept in mind when reading these results. The first is the size of the dataset, with only thirty independent cut faces, which limits the confidence of the estimates and is reflected in the standard deviations reported across folds and seeds. The second is a mild tendency of the models to under-predict the roughest surfaces, which are the most under-represented in the dataset and where the label uncertainty from the structured-light scanner is also the largest. Both limitations point to the collection of additional data, in particular of rougher specimens and of a larger number of carbon steel faces, and to the use of a controlled acquisition setup with better illumination and a sensor of higher spatial resolution, as the natural next steps for the continuation of this line of work.

Chapter 5

Conclusions

This chapter closes the dissertation by returning to the question that motivated the work, by revisiting the objectives declared in Section 1.2, and by identifying the limitations and the natural continuations of the study.

5.1 Response to the Research Aim

The general aim of this dissertation was to determine whether the surface roughness of a laser-cut edge can be estimated from a single RGB image with an accuracy that is useful for shop floor inspection, across several materials and thicknesses. The results reported in Chapter 4 answer that question in the affirmative. Trained on 1041 images of 30 independent cut faces spanning four materials and 15 distinct cutting conditions, and evaluated by cross-validation grouped by cut face, an ImageNet pretrained ResNet-18 recovers between 0.73 and 0.76 of the variance of the three roughness parameters and reaches a MAE of about 3.8 µm on R_a. A multimodal extension of the same network, which adds the plate thickness, the cutting parameters and the normalised position of the image on the face to the tabular branch, matches the image-only model on R_a and on R_q and improves the estimate of R_z to a coefficient of determination of 0.77 and a MAE of 17.5 µm. These figures were obtained under an evaluation protocol that never places two images of the same face on different sides of the partition, and therefore correspond to generalisation towards a cutting trial that the model has not seen in training.

The setup that produced the results in the literature [9, 51, 53] is typically restricted to a single material and a single thickness, uses controlled illumination and a wider FOV, and predicts several roughness values along the thickness of the face simultaneously, so that the network receives the vertical position as part of the target rather than having to infer it from the image alone. The present work operates in the opposite regime, with a broader material and thickness range, an uncontrolled acquisition, a narrow field of view, and a single output per image, which makes the mapping harder to learn and moves the practical value of the model in a different direction. Rather than a metrologically precise reading of the roughness, the model reliably distinguishes rough

regions of a cut from smooth ones across the whole thickness of the face, which is precisely what the intended industrial use case requires.

The five specific objectives declared in Section 1.2 were addressed as follows. The procedure for extracting roughness parameters from the ATOS Q structured light scanner was implemented and validated against the Mitutoyo SJ-210 contact profilometer, and a single amplitude factor $k = 1.49$ was estimated over four repeated acquisitions of a representative specimen. Its application reduces the MAPE against the reference from about a third to about nine percent for the three parameters. The interpretable spectral descriptors of the striation pattern were extracted through a pipeline that combines a flat field illumination correction, a HPS and a harmonic suppression stage, and were validated against the spacing of grooves illuminated by a projected laser line. The multimodal dataset associates every image with those descriptors, with the process parameters of the corresponding specimen, with the normalised position of the image on the face, and with the calibrated roughness targets, and its scope in materials and thicknesses is broader than that of the datasets that dominate the field. The three architectures were trained on that dataset and compared under the same protocol, which quantified precisely how much of the roughness is recovered by the raw image and how much is added by the engineered descriptors and by the process context. Specialising a model on a single material was finally evaluated for aluminium, and was shown to be viable and free of overfitting, but not to outperform the general model restricted to that material.

5.2 Contributions

The specific contributions of the work are the following.

1. A validated and calibrated procedure that turns a structured light scanner not intended for roughness metrology into a practical source of areal referenced labels, with a residual error against a contact reference of about 1.2 µm on R_a, 1.3 µm on R_q and 5.4 µm on R_z. Because the roughness pipeline is implemented directly on the full three-dimensional point cloud, it can be applied to arbitrary regions of the cut face rather than to fixed strips alone, and can be readily extended from the profile parameters R_a, R_q and R_z adopted here to areal parameters such as S_a, without changes to the underlying methodology.

2. A spectral characterisation of the striation pattern that is robust to the harmonic locking observed in a naive Fourier extraction, which reduces the typical dispersion of the recovered wavelength inside the same region of a face from about a third to about a fifth of its mean value, and that eliminates in construction the harmonic outliers that the naive extractor produces in roughly one quarter of the images of the dataset.

3. A multimodal characterisation of the cut edge that combines, by design, two complementary natures of information over the same acquisitions. The image-based spectral pipeline describes the periodic signature of the striation pattern in an interpretable way, and the structured light pipeline provides the traceable roughness labels against which the models

are trained and evaluated. The two modalities operate independently but are aligned image by image, which allows them to be studied both in isolation and as inputs to a common multimodal model.

4. A multimodal dataset of laser-cut edges spanning four materials, seven plate thicknesses between 4 mm and 15 mm, and fifteen distinct combinations of laser power, cutting speed, assist gas, gas pressure and focal position, associating every image with its spectral descriptors, its process parameters, its normalised position on the face and its calibrated roughness. The dataset covers a considerably wider portion of the process space than the single material and single thickness datasets that dominate the field.

5. An automated acquisition setup, designed and built for the purpose of this work, in which a low-cost digital microscope and a Mitutoyo SJ-210 stylus profilometer are alternately mounted on a common bracket carried by the motion system of a fused deposition modelling printer, coordinated through a custom G-code interface that displaces the sensor along a reproducible grid and tags every image with its position.

6. A leakage-safe evaluation protocol for image-based roughness prediction that groups the observations by cut face through stratified group cross-validation repeated over three random seeds. The partition ensures that every reported figure measures generalisation towards a cutting trial never seen in training, in contrast with the conventional random split, which places images with the same information on both sides of the partition and produces optimistic scores that do not reflect field behaviour.

7. Empirical evidence that the raw image outperforms the engineered descriptors on this task and that those descriptors are conditional on the material, together with a characterisation of when multimodal fusion and material specialisation are worth their cost. The image-based models improve over the best tabular configuration by nearly 25 % in coefficient of determination and by more than one micrometre in MAE on R_a. The fusion of the image with the tabular context leaves R_a and R_q essentially unchanged but improves R_z, the parameter that carries the absolute scale of the deepest striations. Specialising a model on a single material was shown to be viable and free of overfitting, but not to outperform the general model, which supports the deployment of a single model that encodes the material as an input and covers all of them at once.

5.3 Limitations

Before listing the limitations that condition the generality of the results, three constraints of the experimental setting deserve explicit acknowledgement, since they were imposed by the industrial collaboration rather than chosen. The specimen set was supplied by the industrial partner and reflects production practice rather than a factorial design, so material, assist gas and cutting regime are partially aligned in the sample. The laser profilometer that would naturally have served as the

roughness reference became unavailable during the campaign, and the ATOS Q structured light scanner was adopted as its replacement, at the cost of a more elaborate extraction pipeline and of the residual bias reported in Section 3.3. The RGB acquisition was performed with a low cost digital microscope whose FOV frames only a small region of the face, which is more restrictive than the wide-angle acquisitions found in the literature [9, 53] and contributes to the residual per-image dispersion identified in the spectral analysis. These constraints do not invalidate the results but should inform the reading of the comparison with related work.

On top of these experimental constraints, four further limitations condition the generality of the results.

1. The dataset size. With only thirty independent cut faces after the field of view filter, the estimates carry a non-negligible uncertainty, which the standard deviations across folds and seeds report explicitly. The tail of the target distribution is also represented by few observations, which shows a mild tendency of the models to underpredict the roughest surfaces.

2. The partial confounding of process variables. The material, the assist gas and the cutting regime are partially aligned in the specimen set, which reflects industrial practice rather than a factorial design and prevents the effects of the three factors from being disentangled at face level.

3. The material dependence of the visual signature. The illumination provided by the digital microscope is constant across acquisitions, but the four materials of the dataset respond differently to it, with the aluminium and the stainless steel surfaces producing stronger specular returns than the carbon and galvanised steels. Combined with the narrow field of view, which frames only a small region of the face with a variable phase, this material-dependent response contributes to the residual per-image dispersion that survives the spectral pipeline.

4. The material coverage of the calibration factor. The calibration factor $k = 1.49$ was derived from four repeated acquisitions of a single characterised specimen, and its applicability across the four materials of the dataset was not independently verified. Different materials interact differently with the structured light of the ATOS Q scanner, in particular through their reflectivity and their response to the anti-reflective coating, so a per-material calibration could reveal a residual bias that the adopted single factor does not correct.

5.4 Future Work

The path from the present results towards the industrial device that motivates the work suggests four directions of continuation.

The first is the enlargement and rebalancing of the dataset. Additional specimens and additional cutting conditions on the rougher end of the process would strengthen the estimates in the region where the models currently underpredict, and would allow the material and regime dependence of the visual signature to be studied without the confounding observed here.

The second is the acquisition head. A dedicated optical head with controlled diffuse illumination, an anti-reflective treatment adapted to specular metallic surfaces, and a sensor of higher spatial resolution would remove most of the sources of residual dispersion identified in the spectral analysis. Under such an acquisition the spectral descriptors would recover a substantially larger portion of their theoretical value.

The third is the deployment of the estimator on the industrial device itself. The image-only ResNet-18 is the natural candidate, since it requires only the RGB image at inference time and matches the fusion on R_a and R_q. Embedding the model on the target platform, integrating it with the operator interface, and validating the whole assembly on the shop floor are the engineering steps that separate the present feasibility result from a deployable instrument. A practical strategy for the field use is to acquire a small number of independent images of the same face and to average their predictions, which exploits the aggregation effect that transforms noisy per-image estimates into stable face-level ones.

The fourth is the extension of the reference targets to areal roughness parameters, in particular the areal arithmetic mean height S_a, which the areal metrology literature [32, 34] has argued to be more representative of the topography of a laser-cut edge than the profile parameters retained here. Because the structured light scanner already delivers the full surface, the point cloud used in this work supports the direct computation of areal parameters, and revisiting the same regression protocol with S_a as target would test whether the image based models retain their accuracy when the label itself becomes a better descriptor of the surface.

The evidence gathered in this dissertation supports the pursuit of these four directions. The feasibility of estimating surface roughness from a single RGB image of a laser-cut edge across several materials and thicknesses is established. The remaining path is one of engineering and of consolidation rather than one of demonstration.

References

[1] Tahsinul Abedin, Haoming Xu, and Shahadat Uddin. “The impact of K selection in K-fold cross-validation on bias and variance in supervised learning models”. en. In: *Scientific Reports* 16.1 (Jan. 2026), p. 6084. ISSN: 2045-2322. DOI: 10.1038/s41598-026-37247-x. URL: https://www.nature.com/articles/s41598-026-37247-x.

[2] Muhammad Alsaadawy et al. “A comprehensive review of studying the influence of laser cutting parameters on surface and kerf quality of metals”. In: *The International Journal of Advanced Manufacturing Technology* (2023). DOI: 10.1007/S00170-023-12768-1.

[3] Reza Asadi et al. “Multimodal deep learning-based on and off-axis melt pool monitoring for layer height and surface metrology predictions in directed energy deposition”. en. In: *Journal of Intelligent Manufacturing* (Oct. 2025). ISSN: 0956-5515, 1572-8145. DOI: 10.1007/s10845-025-02714-1. URL: https://link.springer.com/10.1007/s10845-025-02714-1.

[4] Carl Zeiss AG. *ZEISS INSPECT 3D: Metrology Software for 3D Scanning and Analysis*. Software application. Version used for structured-light scan processing and roughness extraction. 2024.

[5] Creality 3D Technology Co., Ltd. *Ender-3 3D Printer*. https://www.creality.com/products/ender-3-3d-printer. Accessed: 2026-09-10. 2024.

[6] Jia Deng et al. “ImageNet: A large-scale hierarchical image database”. In: *2009 IEEE Conference on Computer Vision and Pattern Recognition*. ISSN: 1063-6919. June 2009, pp. 248–255. DOI: 10.1109/CVPR.2009.5206848. URL: https://ieeexplore.ieee.org/document/5206848/.

[7] José Ferreira et al. “Gen-JEMA: enhanced explainability using generative joint embedding multimodal alignment for monitoring directed energy deposition”. In: *Journal of Intelligent Manufacturing* (2025). DOI: 10.1007/S10845-025-02614-4.

[8] Martin A. Fischler and Robert C. Bolles. “Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography”. In: *Commun. ACM* 24.6 (June 1981), pp. 381–395. ISSN: 0001-0782. DOI: 10.1145/358669.358692. URL: https://dl.acm.org/doi/10.1145/358669.358692.

[9] Alberto Tomás García et al. “Roughness prediction of laser cut edges by image processing and artificial neural networks”. en. In: *Procedia Manufacturing* 54 (2021), pp. 257–262. ISSN: 23519789. DOI: 10.1016/j.promfg.2021.07.040. URL: https://linkinghub.elsevier.com/retrieve/pii/S2351978921001761.

[10] Jason Geng. “Structured-light 3D surface imaging: a tutorial”. en. In: *Advances in Optics and Photonics* 3.2 (June 2011), p. 128. ISSN: 1943-8206. DOI: 10.1364/AOP.3.000128. URL: https://opg.optica.org/abstract.cfm?URI=aop-3-2-128.

[11] Sourish Ghosh et al. “Towards AI driven surface roughness evaluation in manufacturing: a prospective study”. en. In: *Journal of Intelligent Manufacturing* 36.7 (Oct. 2025), pp. 4519–4548. ISSN: 0956-5515, 1572-8145. DOI: 10.1007/s10845-024-02493-1. URL: https://link.springer.com/10.1007/s10845-024-02493-1.

[12] Gene H. Golub and Charles F. Van Loan. *Matrix Computations*. 4th. Baltimore: Johns Hopkins University Press, 2013.

[13] Graichen Produktions- und Vertriebs GmbH. *ATTBLIME ABP: Anti-Reflective 3D Scanning Spray – Technical Datasheet*. Manufacturer datasheet. Available on request from the manufacturer. 2024.

[14] Miaoxian Guo et al. “Prediction of surface roughness using deep learning and data augmentation”. en. In: *Journal of Intelligent Manufacturing and Special Equipment* 5.1 (Aug. 2024), pp. 221–241. ISSN: 2633-6596, 2633-660X. DOI: 10.1108/JIMSE-10-2023-0010. URL: http://www.emerald.com/jimse/article/5/1/221-241/1237608.

[15] Simon S. Haykin and Simon S. Haykin. *Neural networks and learning machines*. en. 3rd ed. New York: Prentice Hall, 2009. ISBN: 978-0-13-147139-9.

[16] Kaiming He et al. “Deep Residual Learning for Image Recognition”. en. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. Las Vegas, NV, USA: IEEE, June 2016, pp. 770–778. ISBN: 978-1-4673-8851-1. DOI: 10.1109/CVPR.2016.90. URL: http://ieeexplore.ieee.org/document/7780459/.

[17] Koji Hirano and Remy Fabbro. “Experimental investigation of hydrodynamics of melt layer during laser cutting of steel”. en. In: *Journal of Physics D: Applied Physics* 44.10 (Mar. 2011), p. 105502. ISSN: 0022-3727, 1361-6463. DOI: 10.1088/0022-3727/44/10/105502. URL: https://iopscience.iop.org/article/10.1088/0022-3727/44/10/105502.

[18] Peter J. Huber. “Robust Estimation of a Location Parameter”. In: *Breakthroughs in Statistics: Methodology and Distribution*. Ed. by Samuel Kotz and Norman L. Johnson. Springer Series in Statistics. Originally published in Annals of Mathematical Statistics, 35(1):73–101, 1964. New York, NY: Springer, 1992, pp. 492–518.

[19] International Organization for Standardization. *Geometrical product specifications (GPS) — Filtration — Part 21: Linear profile filters: Gaussian filters*. Tech. rep. ISO 16610-21:2011. Geneva, Switzerland: ISO, 2011.

[20] International Organization for Standardization. *Geometrical product specifications (GPS) — Surface texture: Profile — Part 2: Terms, definitions and surface texture parameters*. Tech. rep. ISO 21920-2:2021. Geneva, Switzerland: ISO, 2021.

[21] International Organization for Standardization. *ISO 11562:1997 – Geometrical Product Specifications (GPS) – Surface texture: Profile method – Metrological characteristics of phase correct filters*. Tech. rep. Geneva, Switzerland: ISO, 1997.

[22] International Organization for Standardization. *ISO 4287:1997 – Geometrical Product Specifications (GPS) – Surface texture: Profile method – Terms, definitions and surface texture parameters*. Tech. rep. Geneva, Switzerland: ISO, 1997.

[23] International Organization for Standardization. *ISO 4288:1996 – Geometrical Product Specifications (GPS) – Surface texture: Profile method – Rules and procedures for the assessment of surface texture*. Tech. rep. Geneva, Switzerland: ISO, 1996.

[24] International Organization for Standardization. *Thermal cutting — Classification of thermal cuts — Geometrical product specification and quality tolerances*. Tech. rep. ISO 9013:2017. Geneva, Switzerland: International Organization for Standardization, 2017.

[25] Sergey Ioffe and Christian Szegedy. "Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift". In: *Proceedings of the 32nd International Conference on Machine Learning (ICML)*. PMLR, 2015, pp. 448–456.

[26] Ian T. Jolliffe. *Principal Component Analysis*. 2nd. New York: Springer, 2016. DOI: 10.1007/b98835.

[27] Shachar Kaufman et al. "Leakage in data mining: Formulation, detection, and avoidance". en. In: *ACM Transactions on Knowledge Discovery from Data* 6.4 (Dec. 2012), pp. 1–21. ISSN: 1556-4681, 1556-472X. DOI: 10.1145/2382577.2382579. URL: https://dl.acm.org/doi/10.1145/2382577.2382579.

[28] Diederik P. Kingma and Jimmy Ba. "Adam: A Method for Stochastic Optimization". In: *3rd International Conference on Learning Representations (ICLR)*. 2015.

[29] Jeong Hoon Ko and Chen Yin. "A review of artificial intelligence application for machining surface quality prediction: from key factors to model development". en. In: *Journal of Intelligent Manufacturing* 37.2 (Feb. 2026), pp. 775–798. ISSN: 0956-5515, 1572-8145. DOI: 10.1007/s10845-025-02571-y. URL: https://link.springer.com/10.1007/s10845-025-02571-y.

[30] Ron Kohavi. "A Study of Cross-Validation and Bootstrap for Accuracy Estimation and Model Selection". In: *Proceedings of the 14th International Joint Conference on Artificial Intelligence (IJCAI)*. Vol. 2. Morgan Kaufmann, 1995, pp. 1137–1143.

[31] L Łatka, E Jonda, and T Kiełczawa. "The influence of the laser cutting process parameters on the quality of the cut edge". en. In: *Journal of Physics: Conference Series* 2676.1 (Dec. 2023), p. 012001. ISSN: 1742-6588, 1742-6596. DOI: 10.1088/1742-6596/2676/1/012001. URL: https://iopscience.iop.org/article/10.1088/1742-6596/2676/1/012001.

[32] Richard Leach, ed. *Optical Measurement of Surface Topography*. en. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011. ISBN: 978-3-642-12011-4. DOI: 10.1007/978-3-642-12012-1. URL: http://link.springer.com/10.1007/978-3-642-12012-1.

[33] Nikita Levichev et al. "Experimental validation of a machine learning algorithm for roughness quantification in laser cutting". en. In: *Procedia CIRP* 113 (2022), pp. 564–569. ISSN: 22128271. DOI: 10.1016/j.procir.2022.09.175. URL: https://linkinghub.elsevier.com/retrieve/pii/S2212827122013762.

[34] Erica Librera et al. "On the use of Areal Roughness Parameters to Assess Surface Quality in Laser Cutting of Stainless Steel with CO2 and Fiber Sources". In: *Procedia CIRP* (2015). DOI: 10.1016/J.PROCIR.2015.06.069.

[35] Ilya Loshchilov and Frank Hutter. "Decoupled Weight Decay Regularization". en. In: *7th International Conference on Learning Representations (ICLR)*. arXiv:1711.05101 [cs.LG]. Jan. 2019. DOI: 10.48550/arXiv.1711.05101. URL: http://arxiv.org/abs/1711.05101.

[36] Marlin Firmware. *Marlin Firmware*. URL: https://marlinfw.org/.

[37] Mitutoyo Corporation. *Surftest SJ-210 — Portable Surface Roughness Tester*. https://www.mitutoyo.com/products/small-tool-instruments-and-data-management/surface-roughness/surftest-sj-210/. Accessed: 2026-09-10. 2024.

[38] Enis Muratović et al. "Surface Topography and Tolerance Quality Evaluation of Polymer Gears Using Non-Contact 3D Scanning Method". en. In: *Materials* 19.7 (Mar. 2026), p. 1324. ISSN: 1996-1944. DOI: 10.3390/ma19071324. URL: https://www.mdpi.com/1996-1944/19/7/1324.

[39] Vinod Nair and Geoffrey E. Hinton. "Rectified linear units improve restricted Boltzmann machines". In: *Proceedings of the 27th International Conference on Machine Learning (ICML)*. 2010, pp. 807–814.

[40] Jiquan Ngiam et al. "Multimodal Deep Learning". en. In: *Proceedings of the 28th International Conference on Machine Learning (ICML)*. 2011.

[41] A. Michael Noll. "Pitch Determination of Human Speech by the Harmonic Product Spectrum, the Harmonic Sum Spectrum and a Maximum Likelihood Estimate". In: *Proceedings of the Symposium on Computer Processing in Communications*. Vol. XIX. Symposium held April 8–10, 1969. Brooklyn, NY: Polytechnic Press, 1970, pp. 779–797.

[42] OpenCV Team. *OpenCV: Open Source Computer Vision Library*. https://github.com/opencv/opencv. Version 4.8 used for image processing and camera calibration. 2024.

[43] Dragan Rodić et al. "Integration of RSM and Machine Learning for Accurate Prediction of Surface Roughness in Laser Processing". In: *Applied Sciences* (2025). DOI: 10.3390/APP15137064.

[44] Nimel Sworna Ross et al. "Measuring Surface Characteristics in Sustainable Machining of Titanium Alloys Using Deep Learning-Based Image Processing". In: *IEEE Sensors Journal* 23.12 (June 2023), pp. 13629–13639. ISSN: 1558-1748. DOI: 10.1109/JSEN.2023.3269529. URL: https://ieeexplore.ieee.org/document/10110326/.

[45] H. Sakoe and S. Chiba. "Dynamic programming algorithm optimization for spoken word recognition". In: *IEEE Transactions on Acoustics, Speech, and Signal Processing* 26.1 (1978), pp. 43–49. DOI: 10.1109/TASSP.1978.1163055.

[46] Giorgio Santolini et al. "Cut Quality Estimation in Industrial Laser Cutting Machines: A Machine Learning Approach". In: *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. 2019. DOI: 10.1109/CVPRW.2019.00052.

[47] M. Shamlooei et al. "Optimization of Surface Roughness in Laser Cutting Process of Mild Steel Using RSM and GA Method". In: *Lecture notes in mechanical engineering* (2023). DOI: 10.1007/978-3-031-38241-3_91.

[48] Kristijan Sket et al. "Optimizing laser cutting of stainless steel using latin hypercube sampling and neural networks". In: *Optics & Laser Technology* (2024). DOI: 10.1016/J.OPTLASTEC.2024.112220.

[49] Vilhelm Söderberg et al. "Towards On-Machine Surface Metrology Using Image-Based Frequency Analysis for Surface Variation Analysis". en. In: *Journal of Manufacturing and Materials Processing* 10.2 (Feb. 2026), p. 69. ISSN: 2504-4494. DOI: 10.3390/jmmp10020069. URL: https://www.mdpi.com/2504-4494/10/2/69.

[50] Nitish Srivastava et al. "Dropout: A Simple Way to Prevent Neural Networks from Overfitting". en. In: *Journal of Machine Learning Research* 15.1 (2014), pp. 1929–1958.

[51] Janek Stahl, A. Frommknecht, and Marco F. Huber. “Comprehensive Quantitative Quality Assessment of Thermal Cut Sheet Edges using Convolutional Neural Networks”. In: *Proceedings of the British Machine Vision Conference (BMVC)*. 2023.

[52] William M. Steen and Jyotirmoy Mazumder. *Laser Material Processing*. en. London: Springer London, 2010. ISBN: 978-1-84996-061-8. DOI: 10.1007/978-1-84996-062-5. URL: http://link.springer.com/10.1007/978-1-84996-062-5.

[53] Leonie Tatzel and Fernando Puente León. “Image-based roughness estimation of laser cut edges with a convolutional neural network”. In: *Procedia CIRP* (2020). DOI: 10.1016/J.PROCIR.2020.09.166.

[54] Leonie Tatzel et al. “Image-based modelling and visualisation of the relationship between laser-cut edge and process parameters”. en. In: *Optics & Laser Technology* 141 (Sept. 2021), p. 107028. ISSN: 00303992. DOI: 10.1016/j.optlastec.2021.107028. URL: https://linkinghub.elsevier.com/retrieve/pii/S003039922100116X.

[55] K. Venkata Rao et al. “Modelling of kerf width and surface roughness using vibration signals in laser beam machining of stainless steel using design of experiments”. en. In: *Optics & Laser Technology* 169 (Feb. 2024), p. 110146. ISSN: 00303992. DOI: 10.1016/j.optlastec.2023.110146. URL: https://linkinghub.elsevier.com/retrieve/pii/S0030399223010393.

[56] M Vicanek et al. “Hydrodynamical instability of melt flow in laser cutting”. en. In: *Journal of Physics D: Applied Physics* 20.1 (Jan. 1987), pp. 140–145. ISSN: 0022-3727, 1361-6463. DOI: 10.1088/0022-3727/20/1/021. URL: https://iopscience.iop.org/article/10.1088/0022-3727/20/1/021.

[57] Huguang Yang, Han Zheng, and Taohong Zhang. “A review of artificial intelligent methods for machined surface roughness prediction”. en. In: *Tribology International* 199 (Nov. 2024), p. 109935. ISSN: 0301679X. DOI: 10.1016/j.triboint.2024.109935. URL: https://linkinghub.elsevier.com/retrieve/pii/S0301679X2400687X.

[58] Jason Yosinski et al. “How transferable are features in deep neural networks?” en. In: *Advances in Neural Information Processing Systems (NeurIPS)*. Vol. 27. Nov. 2014. DOI: 10.48550/arXiv.1411.1792. URL: http://arxiv.org/abs/1411.1792.

[59] ZEISS / GOM GmbH. *ATOS Q — Compact 3D Scanner*. https://www.zeiss.com/metrology/products/systems/optical-systems/3d-scanning/atos-q.html. Accessed: 2026-09-10. 2024.

[60] Z. Zhang. “A flexible new technique for camera calibration”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 22.11 (Nov. 2000), pp. 1330–1334. ISSN: 1939-3539. DOI: 10.1109/34.888718. URL: https://ieeexplore.ieee.org/document/888718/.