

# **Micro-Manufacturing for Scientific Instrumentation**

by

Jiaming Liu

B.Eng. (Hons), Materials Science and Engineering  
Queen Mary University of London

Submitted to the Program in Media Arts and Sciences, School of Architecture and Planning, in partial fulfillment of the requirements for the degree of

Master of Science

at the

Massachusetts Institute of Technology

May 2026

© 2026 Jiaming Liu. All rights reserved.

The author hereby grants to MIT a nonexclusive, worldwide, irrevocable, royalty-free license to exercise any and all rights under copyright, including to reproduce, preserve, distribute and publicly display copies of the thesis, or release the thesis under an open-access license.

Authored by: Jiaming Liu  
Program in Media Arts and Sciences  
May 15th, 2026

Certified by: Neil Gershenfeld  
Director, Center for Bits and Atoms  
Professor, Massachusetts Institute of Technology  
Thesis Supervisor

Accepted by: Joseph Paradiso  
Academic Head, Program in Media Arts and Sciences

# Micro-Manufacturing for Scientific Instrumentation

by

Jiaming Liu

Submitted to the Program in Media Arts and Sciences, School of Architecture and Planning, on  
May 15th, 2026  
in partial fulfillment of the requirements for the degree of  
Master of Science

## ABSTRACT

Scientific instrumentation often advances at the pace allowed by its fabrication infrastructure. Laboratories can design complex experiments digitally, but micron-scale parts made from hard materials, high-aspect-ratio geometries, or tightly coupled opto-mechanical assemblies still require specialized machines, process knowledge, and metrology. This thesis develops a laboratory-scale framework for micron-scale manufacturing for scientific instrumentation: an integrated ecosystem of material selection, precision motion, laser micromachining, optical characterization, and automation.

The central case study is the development workflow for diamond magic-angle-spinning nuclear magnetic resonance rotors. Diamond rotors motivate the system because they combine demanding material choice, micron-scale geometry, internal bore quality, concentricity, low-damage machining, and functional spin performance. The thesis treats rotor fabrication as a closed loop: NMR requirements define the part, diamond stock and orientation constrain the process, laser machining modifies the geometry, and computed tomography, Raman spectroscopy, microscopy, dimensional metrology, and spin testing provide feedback.

Around this case study, the thesis documents supporting instrument-building capabilities: a low-cost piezo stage, a custom 532 nm confocal Raman platform, a low-cost spectrometer prototype, an in-house femtosecond laser micromachining platform, and an in-house waterjet-guided laser platform developed through optical, hydraulic, mechanical, and sensing design iterations.

The contribution is both technical and architectural: a set of fabrication, metrology, and process-development methods for difficult micron-scale scientific components, and a framework for designing, making, measuring, analyzing, and revising parts while preserving traceability between design intent, process conditions, material state, measured geometry, and functional performance.

Thesis supervisor: Neil Gershenfeld  
Title: Director, Center for Bits and Atoms  
Professor, Massachusetts Institute of Technology

# Micro-Manufacturing for Scientific Instrumentation

by

Jiaming Liu

This thesis has been reviewed and approved by the following committee members:

Thesis Supervisor: Professor Neil Gershenfeld  
Director, Center for Bits and Atoms  
Massachusetts Institute of Technology

Thesis Reader: Professor Robert Guy Griffin  
Professor of Chemistry; Director, Francis Bitter Magnet Laboratory  
Massachusetts Institute of Technology

Thesis Reader: Professor Debabrata Goswami  
HAG Professor of Chemistry  
Indian Institute of Technology Kanpur

# Acknowledgements

I would like to express my deepest gratitude to everyone who shared this spacetime with me during my two years at MIT. To the incredible community at the Center for Bits and Atoms (CBA), across the MIT campus, and to friends and colleagues around the world: without you, I could never have imagined such a splendid and transformative life experience. While there are too many of you to list individually, please know that the moments we spent together are deeply cherished. I hope that I was able to bring as much warmth and light into your days as you have brought into mine.

A special thank you goes to all my collaborators who contributed to the work underlying this thesis over the past two years. I invite you to find your names within these pages, and I hope I have done a decent job of representing our shared efforts. To the members of the staff at CBA: I am lucky to have you by my side, Dan, Kara, and Marissa; without you, CBA would not function as a lab.

To my advisor, Neil, words are simply insufficient to capture my gratitude. Thank you for making everything in my current life at MIT possible, and for believing in me during the tough times when I even doubted myself. Thank you for making the unprecedented hostility against scientific funding feel like smooth sailing, and for your unparalleled ability to create order out of chaos—and to occasionally bring the chaos back just to spark creativity.

To my thesis reader, Bob, thank you for your enduring encouragement and guidance. It has always been a profound comfort to walk into the Magnet Lab knowing you were there, ready to put down whatever was at hand to have a meaningful conversation with me.

To my thesis reader, Debu, thank you for your incredible guidance on the direction of these projects and for working alongside me. You graciously shared your valuable expertise and provided the critical, hands-on experience that was essential to bringing this work to life.

To my parents, your love and attention have been an ever-flowing spring. Your profound respect for knowledge and education laid the foundation for this journey and made me who I am today.

## **AI and Generative Tool Use**

There is no doubt that artificial intelligence has supercharged my progress throughout this work. Yet, paradoxically, moving toward an era of machines that know almost everything has only made me cherish my human interactions all the more.

AI and generative tools were used as drafting and editing assistants during the preparation of this thesis, including support for organizing source material, revising prose, formatting the manuscript, placing figures, generating draft charts, and checking document-build output. These tools were not used as scientific sources.

# Contents

<b>Abstract</b>	<b>2</b>
<b>Acknowledgements</b>	<b>4</b>
<b>1 Introduction: Open Micron-Scale Manufacturing as Scientific Infrastructure</b>	<b>16</b>
1.1 Motivation and Scope . . . . .	16
1.2 The Fabrication Bottleneck in Scientific Instrumentation . . . . .	16
1.3 Laboratory Access, Reproducibility, and Iteration . . . . .	17
1.4 Thesis Claim and Contributions . . . . .	17
1.5 Evidence Basis and Roadmap . . . . .	18
<b>2 Diamond MAS-NMR Rotors</b>	<b>19</b>
2.1 Chapter Overview . . . . .	19
2.2 Background and Motivation: The Instrumentation Bottleneck . . . . .	19
2.2.1 The Fabrication Bottleneck in Scientific Discovery . . . . .	19
2.2.2 Principles of MAS-NMR and the Resolution Challenge . . . . .	20
2.2.3 Homogeneous and Inhomogeneous Linewidths . . . . .	21
2.2.4 Sensitivity Requirement and Dynamic Nuclear Polarization . . . . .	22
2.2.5 Material Limits and the Transition to Diamond . . . . .	22
2.3 Diamond Selection for NMR Rotors . . . . .	23
2.3.1 Carbon Phase Stability and Synthesis Routes . . . . .	24
2.3.2 Defect Profiles and Dislocation Densities . . . . .	29
2.3.3 Crystallographic Orientation and Cleavage Mechanics . . . . .	29
2.3.4 Fracture Toughness and Brittle Failure . . . . .	30
2.3.5 Diamond Selection Criteria and Working Evidence . . . . .	30
2.3.6 State-of-the-Art Toughening Concepts . . . . .	32
2.4 Laser Machining for Diamond Rotor Drilling . . . . .	35
2.4.1 Optical Absorption Physics in Diamond . . . . .	35
2.4.2 Which Laser Is More Efficient for Diamond? . . . . .	37
2.4.3 Laser and Motion Architectures . . . . .	39
2.5 Lathe, Air-Bearing, and Clamping System Design . . . . .	40
2.5.1 First Generation: Rotary-Assisted Stage and Kinematic Flexure Mount . . . . .	40
2.5.2 Second Generation: Split OD and Centerless ID Stages . . . . .	41
2.5.3 Through-Bore Air-Bearing Flip-Spindle Micro-Lathe . . . . .	43
2.6 Characterization Methods for Single-Crystal Diamond . . . . .	43
2.6.1 Incoming Material Characterization . . . . .	43
2.6.2 Machining Feedback Loop . . . . .	46

2.7	Preliminary Results: NMR Spectra from Diamond Rotors . . . . .	48
2.8	Supporting Architecture . . . . .	51
2.8.1	Stator Material Selection and Manufacturing . . . . .	51
2.8.2	Aerostatic Bearing Evolution . . . . .	53
2.8.3	End Caps and Micro-Impellers . . . . .	54
2.8.4	Aerostatic-Bearing Modeling . . . . .	55
2.8.5	Thermal-Viscous Dynamics in Helium . . . . .	56
<b>3</b>	<b>Precision Motion and the Fabrication Ecosystem</b>	<b>59</b>
3.1	Why Accessible Sub- $\mu\text{m}$ Motion Matters . . . . .	59
3.2	Motion Architectures Used in the Fabrication Ecosystem . . . . .	59
3.3	Piezoelectric Stages for Accessible Fine Motion . . . . .	60
3.3.1	Operating Principle and Design Envelope . . . . .	60
3.3.2	Review of Piezo Stage Types . . . . .	61
3.3.3	Hysteresis, Creep, and Position Sensing . . . . .	62
3.4	In-House Low-Cost Piezo Stage . . . . .	62
3.4.1	Design Intent . . . . .	62
3.4.2	Mechanism, Flexure, and Material Considerations . . . . .	63
3.4.3	Bill of Materials . . . . .	65
3.4.4	Motion-Control Architecture and Measured Repeatability . . . . .	66
3.4.5	Key Commercial Components . . . . .	67
3.4.6	Characterization Data Products . . . . .	71
3.5	Long-Travel High-Speed Linear Motion for Laser Systems . . . . .	71
3.5.1	Why Long-Travel High-Speed Stages Are a Separate Class . . . . .	71
3.5.2	Ball-Screw Stages and the Oxford Laser System . . . . .	71
3.5.3	Direct-Drive Mechanical-Bearing Stages and Compact Auxiliary Axes . . . . .	72
3.5.4	Aerotech Stage Specifications as Design Targets . . . . .	73
3.5.5	Encoder, Bearing, and Servo Architecture Inside Commercial Stages . . . . .	74
3.5.6	Drive and Controller Architecture: Advantages and Drawbacks . . . . .	76
3.5.7	Assembly and Tuning Considerations . . . . .	78
3.5.8	Motion Error and Fabrication Outcomes . . . . .	78
3.5.9	Position-Synchronized Output . . . . .	79
3.6	Chapter Synthesis . . . . .	79
<b>4</b>	<b>Laser Micromachining Systems</b>	<b>80</b>
4.1	Femtosecond Laser Micromachining Systems . . . . .	80
4.1.1	Introduction to Femtosecond Laser . . . . .	80
4.1.2	State of the Art and Cold-Ablation Physics . . . . .	81
4.1.3	In-House Femtosecond Platform Integration . . . . .	82
4.1.4	Multi-Wavelength Machining Setup . . . . .	85
4.1.5	Beam Shaping and Non-Diffractive Bessel Methods . . . . .	86
4.1.6	Burst Mode and Wavelength Agility . . . . .	87
4.1.7	Femtosecond Source Characterization: Spectrum, Bandwidth, and Dispersion . . . . .	87
4.1.8	Autocorrelation for Daily Pulse-Duration Checks . . . . .	89
4.1.9	FROG/XFROG for Pulse-Field Retrieval . . . . .	90
4.1.10	d-scan for Dispersion-Aware Pulse Diagnostics . . . . .	90
4.1.11	Live Femtosecond Process Monitoring and Adaptive Control . . . . .	91
4.2	Waterjet-Guided Laser Micromachining Systems . . . . .	93

4.2.1	Brief History and Current Direction . . . . .	93
4.2.2	Working Principle . . . . .	94
4.2.3	System Architecture . . . . .	94
4.2.4	External Industry Benchmark . . . . .	95
4.2.5	In-House WJGL Platform Integration . . . . .	95
4.2.6	First-Generation In-House WJGL Design . . . . .	97
4.2.7	Second-Generation In-House WJGL Design . . . . .	98
4.2.8	Third-Generation In-House WJGL Design . . . . .	99
4.2.9	Dual-Beam Coupling as a Future WJGL Upgrade . . . . .	100
4.2.10	Dual-Beam Coupling Performance Metrics . . . . .	100
4.2.11	Optical-Field Regulation . . . . .	101
4.2.12	Gas Assistance for WJGL Head Development . . . . .	101
4.2.13	Diamond-Oriented WJGL Direction . . . . .	103
4.2.14	WJGL Validation Workflow . . . . .	105
<b>5</b>	<b>Characterization and Spectroscopy Tools</b>	<b>106</b>
5.1	Chapter Overview . . . . .	106
5.2	Motivation for Open Spectroscopy . . . . .	106
5.3	Single-Laser 532 nm Confocal Raman Platform . . . . .	107
5.3.1	Design Target . . . . .	107
5.3.2	Laser Wavelength Selection . . . . .	107
5.3.3	Optical Architecture . . . . .	108
5.3.4	Component Selection Criteria . . . . .	111
5.3.5	Spectrometer and Detector Selection . . . . .	112
5.3.6	Raman Prototype Bill of Materials . . . . .	112
5.3.7	Reference Off-the-Shelf Confocal Raman BOM . . . . .	113
5.3.8	Build Workflow . . . . .	114
5.3.9	Scan Control and Data Products . . . . .	115
5.3.10	Spectrum Processing, Noise, and SNR . . . . .	115
5.3.11	Calibration and Validation . . . . .	116
5.4	Published Custom and Open Raman Systems . . . . .	116
5.5	Commercial Raman Benchmarks and Robotic Integration . . . . .	117
5.6	Low-Cost Personal Spectrometer . . . . .	118
5.6.1	Collaborative Baseline . . . . .	118
5.6.2	Design Choices . . . . .	121
5.6.3	Calibration and Validation . . . . .	121
5.7	Commercial Spectrometer Benchmarks . . . . .	122
5.8	Open and Self-Build Spectrometer Landscape . . . . .	123
5.9	Instrument-Building Continuum . . . . .	123
5.10	Detailed Build and Validation Protocols . . . . .	124
5.10.1	Raman Build Record . . . . .	124
5.10.2	Raman Alignment Checklist . . . . .	124
5.10.3	Raman Acquisition Template . . . . .	124
5.10.4	Personal Spectrometer Build and Validation Sequence . . . . .	125
5.10.5	Chapter Summary . . . . .	125

<b>6 Synthesis and Future Work</b>	<b>127</b>
6.1 Synthesis . . . . .	127
6.2 Contributions . . . . .	127
6.3 Limitations . . . . .	128
6.4 Implications . . . . .	128
6.5 Roadmap . . . . .	128
<b>Appendix A: Lessons Learned</b>	<b>130</b>
<b>References</b>	<b>132</b>

# List of Figures

2.1	Bruker BioSpin AVANCE NEO console and Bruker Ascend 800 MHz NMR magnet at Francis Bitter Magnet Lab, used as the high-field experimental context for the diamond MAS-NMR rotor work. . . . .	20
2.2	Diamond type chart from GIA. Top to bottom: type Ia, type Ib, type IIa, type IIb [15]. . . . .	24
2.3	GIA schematic showing how nitrogen and boron atoms can replace carbon atoms in the diamond lattice. The diagram links type Ia and Ib diamonds to nitrogen aggregation state, type IIa to absence of measurable nitrogen or boron, and type IIb to boron substitution [15]. . . . .	24
2.4	Typical rough single-crystal diamond morphologies from HPHT and CVD growth. HPHT rough is commonly faceted; CVD rough is commonly square or plate-like because it grows on a planar seed. . . . .	25
2.5	Carbon pressure-temperature diagram showing graphite, diamond, liquid carbon, and representative artificial synthesis windows. Adapted from Bundy [16]. . . . .	25
2.6	Cubic HPHT diamond press. The left panel shows the exterior of a six-anvil press, and the right panel shows the opposed-anvil interior. Adapted from Deljanin [29] and Diamond Land [30]. . . . .	26
2.7	Internal HPHT growth assembly. The right panel, from Fang et al. [24], labels the growth-cell stack: 1, steel cap; 2, Cu + Mo sheet; 3, NaCl + ZrO <sub>2</sub> sleeve; 4, graphite heater; 5, carbon source; 6, metal catalyst; 7, seed crystal; 8, ZrO <sub>2</sub> + MgO pillar; 9, pyrophyllite. . . . .	27
2.8	Microwave-plasma CVD diamond reactor architecture, showing gas/vacuum handling, microwave coupling, plasma generation, substrate support, and multi-reactor production hardware. Adapted from Diamond Land [34]. . . . .	28
2.9	CVD diamond plates during plasma exposure and after removal from the reactor. The square plate geometry follows from homoepitaxial growth on planar seeds. . . . .	28
2.10	Diamond machining and cleavage-risk evidence. Panels show SEM-visible cleavage, micro-CT evidence of internal chipping in CVD stock, and the 54.76° relation between [100] machining and [111]-related cleavage in the diamond lattice. Adapted from Schaffer et al. [2]. . . . .	30
2.11	Representative diamond MAS rotor geometry and feature scale. The figure shows the 0.7 mm-class rotor envelope, physical rotor assembly, and tomographic views used to verify the internal bore and wall geometry. This scale illustrates why the material must combine high strength with laser machinability, low taper, and low ID/OD eccentricity. Adapted from Golota et al. [1]. . . . .	31
2.12	Summary of diamond rotor fabrication outcomes from the internal parsed rotor log. The chart summarizes 108 rotor-like records and groups material classes from the vendor/method and source-group fields. Rows previously labeled as unspecified are counted as HPHT. The classification is used as working material-selection evidence rather than as a controlled supplier-comparison study. . . . .	32
2.13	Nanotwinned diamond deformation model from Wen et al., showing twin-boundary-related deformation fields and dislocation interaction with the nanotwin architecture [41]. . . . .	34

2.14	Diamond wavelength absorption and laser-coupling evidence map. . . . .	36
2.15	First-generation rotary-assisted stage used for early diamond rotor laser machining. The figure documents the initial rotary-stage concept that preceded the OD-first and centerless-ID machining workflow. Adapted from patent US20210146475A1 [55]. . . . .	40
2.16	Second-generation outer-diameter air-bearing stage. The workpiece-defined rotational axis reduces transferred spindle runout during OD shaping. Custom-built high-precision air bearing lathe used for machining the OD of the diamond rotors. (a) A wire frame diagram showing the overall assembly, highlighting the DC drive motor, drive belt, precision ground spindle, and the custom 4-jaw chuck. (b) A sectional 3D CAD view reveals internal components, including the graphite air bearing pads that provide low-friction, high-precision rotation, the magnetic axial preload mechanism, and the positioning of the diamond stock within the chuck. (c) A photograph shows the physical assembled lathe. Adapted from Schaffer et al. [2]. . . . .	41
2.17	Second-generation centerless inner-diameter stage. (a) Wire-frame and top-view layout showing the drive wheel, follower wheels, and diamond stock. (b) Cross-sectional CAD view showing the drive wheel on a linear rail and support bearings that guide the rotor. (c) Assembled fixture on its positioning stage. The already-machined OD acts as the datum for bore machining, improving ID/OD concentricity. Adapted from Schaffer et al. [2]. . . . .	42
2.18	Through-bore air-bearing flip-spindle micro-lathe. . . . .	43
2.19	Bruker Raman comparison of GW and E6 diamond samples. The spectra were smoothed, baseline corrected, and normalized to the first-order diamond band. Relative to E6, the GW sample shows lower intensity in the second-order diamond band near $2600\text{ cm}^{-1}$ and an additional weak feature near $614\text{ cm}^{-1}$ , consistent with a different impurity or defect population in the two diamond feedstocks. . . . .	45
2.20	THz-TDS absorption-coefficient comparison for sapphire, single-crystal diamond, and zirconia. Diamond exhibits low absorption comparable to sapphire and substantially lower absorption than zirconia across most of the measured THz range, supporting its use as a DNP-compatible MAS rotor material. Measurements were performed with the Nelson laboratory in Chemistry. . . . .	45
2.21	Fracture-mechanics-derived relative angular frequencies for candidate MAS rotor materials, normalized to YSZ [1]. . . . .	46
2.22	Micro-CT reconstruction of a laser-machined diamond rotor section. The three-dimensional reconstruction makes the bore geometry and internal surface condition visible without destructive sectioning, allowing chips, taper, discontinuities, and residual debris to be identified before additional processing or spin testing. . . . .	47
2.23	SEM end-view of a laser-machined diamond rotor. The measurement overlay shows an outer diameter of approximately $699\text{ }\mu\text{m}$ and an inner diameter of approximately $471\text{ }\mu\text{m}$ , while the end-face morphology provides a direct check for bore shape, edge chipping, and surface damage after machining. . . . .	47
2.24	$^1\text{H}$ -detected spectra of $\text{U-}^{15}\text{N},^{13}\text{C-N-f-MLF-OH}$ recorded using a diamond rotor. (A) (H)CH correlation spectrum. (B) (H)NH correlation spectrum [1]. . . . .	49
2.25	Three-dimensional $^1\text{H}$ -detected RFDR (hCCH) spectrum of $\text{U-}^{13}\text{C},^{15}\text{N-A}\beta_{1-40}$ recorded using an improved diamond rotor spinning at $\omega_r/2\pi = 100\text{ kHz}$ at $277\text{ K}$ in an $800\text{ MHz}$ spectrometer [2]. . . . .	50
2.26	Formlabs Ceramic Resin stator prototype for MAS hardware development. The print captures the overall stator envelope and gas-feature layout, while visible support-contact artifacts and surface roughness indicate that post-print machining, finishing, and dimensional verification are required before the geometry can be used as a precision bearing component. . . . .	53

2.27	BMF-printed air-bearing and stator prototype components used to evaluate sub-millimeter bearing geometry and feed-channel resolution. . . . .	54
2.28	Micro-CT cross-section of a Bruker ceramic air-bearing reference geometry, showing the central ID and radial feed passages used to support a sub-millimeter MAS rotor. . . . .	54
2.29	Rotor-closure and gas-drive prototypes. (A) Two-part diamond drive-tip prototype fabricated by Zach Fredin, showing turbine-fin geometry, central post, and laser-machined surface texture; the machining marks and assembled construction motivate improved finishing and single-piece diamond closure fabrication. (B) Micro-printed drive-tip and impeller-style prototypes developed with Nick Wiesner and Samuel Strymis for sub-millimeter rotor closure development; additive microfabrication can generate turbine-like features for rapid iteration, but final use still depends on balance, cap fit, NMR background, thermal expansion, and spin survivability. . . . .	55
2.30	Helium-bearing flow simulations from preliminary modeling by Nick Wiesner at the MIT Francis Bitter Magnet Laboratory. (A) CFD velocity field for gas injection into the sub-millimeter rotor-bearing clearance, showing gas acceleration through the restrictor and turning into the rotor-bearing film. (B) Recess/source pressure ratio as a function of gauge-pin size for nitrogen and helium bearing-flow cases; gauge-pin size is used as a rotor-diameter proxy in a fixed bearing geometry, so increasing gauge-pin size reduces clearance and changes pressure recovery. Together, the panels show that rotor-bearing gap, restrictor geometry, and gas choice must be optimized as a coupled aerodynamic system. . . . .	57
3.1	Multilayer piezo stack concept from the laboratory slide deck. Alternating ceramic and electrode layers convert applied electric field into small axial strain; many active layers increase usable displacement at a practical drive voltage. . . . .	61
3.2	Front CAD view of the in-house low-cost piezo micro-motion stage. The stage was designed by Bowen Zeng, with work from Max Aalto and Jake Read, and was inspired by the low-cost open-source XYZ nanopositioner architecture reported by Liao et al. [83]. The assembly combines piezo stack actuation, miniature ball-slide guidance, magnetic encoder boards, magnetic scale tape, and a compact fixture stack. . . . .	63
3.3	Flexure-stage fabrication evidence from the laboratory slide deck. The waterjet trial was unexpectedly clean for the geometry: the flexure used approximately 1 mm internal radii and 0.5 mm wall sections through 1/4 in (6.35 mm) stock, dimensions that would normally push abrasive-waterjet cutting toward a coupon-validation regime. . . . .	64
3.4	PiezoDrive PDU150 driver used to command the prototype piezo stacks. The module turns low-voltage command signals into high-voltage piezo drive outputs while providing the boost conversion and offset-voltage structure needed for practical bench operation [106]. . . . .	68
3.5	Thorlabs PC4WM piezo stack actuator selected for the low-cost stage. The actuator gives compact fine motion, but the usable stage performance depends on preload, amplifier current, guide friction, flexure stiffness, and closed-loop sensing [107]. . . . .	69
3.6	AS5311 magnetic encoder and multipole-scale concept used for compact feedback in the in-house piezo stage. The scale converts linear travel into a repeating magnetic field pattern, while the encoder reads the field without adding sliding contact to the motion axis [108]. . . . .	70
3.7	THK LS852 miniature linear ball slide used as the prototype guideway. This component makes the stage mechanically accessible and easy to assemble, but the final motion behavior must include rolling-element friction, straightness error, and cage behavior [109]. . . . .	70

3.8	Commercial Aerotech reference stages used to anchor the long-travel stage comparison. The images show the physical distinction between a compact direct-drive crossed-roller stage, a protected ball-screw mechanical-bearing stage, and a compact nanopositioning direct-drive stage [87, 111, 113]. . . . .	74
4.1	Cold-ablation comparison from Lin and Hong, illustrating long-pulse thermal damage and femtosecond-pulse localized removal [136]. . . . .	81
4.2	In-house femtosecond micromachining platform, showing the ultrafast source, optical table, beam-delivery optics, vertical optical breadboard, objective/imaging path, motion stage, chiller, and supporting electronics. . . . .	82
4.3	Simplified femtosecond laser micromachining setup from Owusu-Ansah and Dalton for normal-incidence machining of a 0.5 mm tungsten sheet, showing the source, shutter, attenuator, safety interlock, steering optics, focusing lens, and CNC stage [140]. . . . .	83
4.4	iPhone 15 OLED display pixels imaged through the in-house femtosecond laser micro-vision path. The specified 2556 by 1179 display at 460 ppi gives an approximately 55 $\mu\text{m}$ display-pixel pitch; the separated RGB subpixel structure demonstrates $\mu\text{m}$ -scale inspection capability rather than machining calibration [147]. . . . .	84
4.5	Three-output wavelength-routing region of the in-house femtosecond platform, showing the infrared, green, and ultraviolet output geometry and the fixed/kinematic mirror mounts used to route one selected wavelength at a time onto the common downstream machining axis. . . . .	85
4.6	Non-diffractive beam comparison from Nguyen et al.: (a) Bessel-Gauss beam with 3° conical half-angle, 11 $\mu\text{m}$ core diameter, and 43 mm Bessel length; (b) Gaussian beam with similar core diameter and 0.24 mm confocal range; (c) longitudinal and transverse intensity sections [155]. . . . .	87
4.7	Experimental set-up of single-shot autocorrelator. RM: retro mirror, FM: folding mirror, BS: beam splitter, CL: cylindrical lens, F: filters, A: aperture. From Ramaiah et al. [161]. . . . .	89
4.8	FROG/XFROG diagnostic setup from Nakano et al., showing reference- and test-pulse routing, half-wave plate, dichroic mirror, beam splitters, concave mirrors, BBO crystal, removable mirror, delay line, fiber coupling, and multichannel spectrometer for cross-correlation, SHG-FROG, and self-diffraction FROG measurements [165]. . . . .	91
4.9	Self-calibrating d-scan setup from Alonso et al., using prism-compressor insertion as the scan variable, an off-axis parabola, thin BBO crystal, blue filtering of the residual fundamental, and spectrometer acquisition of the SHG trace over 51 prism positions [167]. . . . .	92
4.10	Measured and retrieved d-scan results from Alonso et al., including d-scan traces, pulse spectrum, retrieved spectral phase, temporal intensity, temporal phase, measurement uncertainty, and retrieved compressor GDD/TOD per unit insertion [167]. . . . .	92
4.11	Basic working principle of waterjet-guided laser machining: a laser is focused into a pressurized water jet, total internal reflection guides the beam along the jet, and water cooling/debris removal occur at the workpiece. . . . .	94
4.12	Waterjet-guided laser system architecture from He et al., identifying the laser/optical unit, coupling unit, water-supply unit, air-supply unit, work platform, CCD observation path, beam expansion/collimation optics, dichroic element, focusing lens, optical window, nozzle, gas jet, water layer, and total-reflection guided path [173]. . . . .	95
4.13	Avonisys finished-material demonstration showing 0.25 mm-pitch features in 18 mm aluminum; the black material on the left is silicon. . . . .	96

4.14	First-generation in-house WJGL design and hardware: 1064 nm MOPA laser, tip-tilt mount, XY stage, CMOS camera path, camera focusing lens, dichroic element, aspheric focusing optic, and initial laser-water coupling unit with protective lens, water chamber, and nozzle. Designed and built by Max Aalto, Vineet Padia, and John Zhang. . . . .	98
4.15	Second-generation in-house WJGL design: 1064 nm MOPA source, XY stage, high-magnification imaging path, ring light, prism mirror, and revised head/lens section with axicon lens, focusing ring, aspheric lens, O-ring seal, and 50 $\mu\text{m}$ nozzle. Designed and developed by Max Aalto. . . . .	99
4.16	Third-generation in-house WJGL design: compact custom head CAD, assembled vertical head in the machine frame, and granite bridge/base machine concept with workpiece stage and catcher region. Developed with Max Aalto, Leo Cheng, and Abigail Suk. . . . .	100
4.17	Dual-beam WJGL coupling setup from Wang et al., combining two 532 nm solid-state laser channels with polarization optics and time-domain pulse control before focusing into the micro water jet [177]. . . . .	101
4.18	Dual-beam WJGL coupling-performance measurements from Wang et al., showing micro-jet transmission behavior, coupling efficiency versus jet diameter and processing distance, coupled power versus coupled jet length, and pulse energy versus coupled jet length [177]. . . . .	102
4.19	Optical-field regulation method for WJGL from Zhang et al., comparing centered and eccentric beam conditions, cross-section energy-density distribution, and groove-depth response for taper reduction in deep processing [176]. . . . .	102
4.20	Gas-assisted WJGL reference from Cao et al., showing the processing head, simulation geometry, focusing lens, pulsed laser, water inlet, helium inlet, outer coaxial gas inlet, water jet, gas nozzle, workpiece, mixture outlet, and deep-hole flow model [178]. . . . .	103
4.21	WJGL diamond machining system, metrology, and optimization workflow from Wang et al., linking the processing system, water-light coupling, coupling-power measurement, diamond morphology measurement, and NSGA-III process optimization [179]. . . . .	104
4.22	Energy-dependent WJGL groove response in single-crystal diamond from Chen et al., showing underwater power attenuation for initial powers of 8 W to 20 W, effective jet length at 10 W and 20 W, and high-power water-jet disturbance [180]. . . . .	105
5.1	Laser and power-supply selection criteria for the local Raman platform. Wavelength sets the Raman-efficiency/background tradeoff, linewidth limits usable spectral resolution, beam quality controls confocal focus size, and power stability determines whether repeated spectra can be compared quantitatively. . . . .	108
5.2	Component-level layout used to explain the local 532 nm confocal Raman path. The design separates the laser delivery path, microscope objective, dichroic or beam-splitting region, imaging branch, confocal aperture, Rayleigh-rejection filter, fiber coupling, and Raman spectrometer so each subsystem can be aligned and validated independently. . . . .	109
5.3	CAD rendering of the local confocal Raman architecture, showing the excitation source, microscope optics, collection path, and compact spectrometer connection. . . . .	109
5.4	Bench implementation of the local Raman optical path during development. The system uses cage-mounted optics so the excitation, imaging, filtering, and spectrometer-coupling paths can be modified and realigned during iteration. . . . .	110
5.5	Spectrometer selection criteria for the Raman platform. Grating groove density and blaze wavelength set the resolution/coverage/efficiency tradeoff; slit width and focal length set line sharpness and throughput; detector quantum efficiency, pixel sampling, cooling, dark current, and read noise determine how weak Raman bands can be measured. . . . .	112

5.6	OpenRAMAN project spectra used as processing benchmarks during system development. The key quality metric is the SNR of assigned Raman bands after dark subtraction, spike removal, and baseline correction, with smoothing limited to display so that narrow vibrational features are not artificially broadened or shifted [190]. . . . .	116
5.7	Robotic Raman integration concept for remote or hazardous-site sensing, reproduced from an MIT News report on a 2026 Raman spectroscopy workshop [206]. In this chapter it serves as a future-facing comparison to the local benchtop confocal Raman platform rather than as a completed local capability. . . . .	118
5.8	Published MPMS images showing the baseline open spectrometer concept that we replicated and are improving in the lab [207]. . . . .	119
5.9	MPMS monochromator design, including the ray-traced geometry and the corresponding 3D-printed optical body. These figures support the collaborative low-cost spectrometer section: we replicated this concept in the lab and are improving the mechanical stiffness, mount repeatability, and alignment interfaces [207]. . . . .	119
5.10	Cyclohexane Raman comparison from the MPMS project, showing how the modular spectrometer was evaluated against a commercial measurement while remaining an educational/access instrument [207]. . . . .	122

# List of Tables

2.1	Representative systems motivating ultrafast MAS-NMR, based on applications reviewed in modern solid-state NMR and ultrafast MAS literature [3, 4]. . . . .	21
2.2	Surface-speed and Mach-number calculation for a 0.7 mm rotor at room temperature. The calculation uses $U = \pi D\nu_r$ , $\alpha_{\text{air}} \approx 343$ m/s, and $\alpha_{\text{He}} \approx 1007$ m/s. . . . .	23
2.3	Comparison of candidate MAS rotor materials relevant to high-frequency and DNP-compatible rotor design. Diamond values are material- and defect-state-dependent, so the table is a design comparison rather than a universal supplier specification [1, 2, 14]. . . . .	23
2.4	Ranking of state-of-the-art diamond toughening methods for possible MAS-NMR rotor use.	35
2.5	Estimated dimensional change for a 0.700 mm support or closure feature. The $\alpha$ column reports the linear thermal expansion coefficient in $\mu\text{m m}^{-1} \text{K}^{-1}$ ; the next columns translate that normalized material value into actual micrometer-scale diameter change [62–73]. . . . .	52
3.1	Bill of materials for the in-house XY piezo stage. . . . .	65
5.1	Design criteria for the major Raman microscope components. . . . .	111
5.2	Costed Raman prototype BOM from the local original system. Prices are reference values from the 2025/2026 build period. . . . .	113
5.3	Peer-reviewed off-the-shelf CRM component reference from Díaz Barrero et al. . . . .	114
5.4	MPMS reference BOM and module set. . . . .	120

# Chapter 1

## Introduction: Open Micron-Scale Manufacturing as Scientific Infrastructure

### 1.1 Motivation and Scope

Experimental science depends on instruments that can make, move, measure, and control matter at the scale of the phenomena under investigation. Progress in structural biology, materials science, optics, microfluidics, and precision metrology is frequently limited not only by theoretical understanding, but by the availability of physical apparatus with the required geometry, material properties, tolerances, and measurement feedback. Instrumentation is therefore not an auxiliary activity; it is a primary mechanism by which new experimental regimes become accessible.

At the macro scale, accessible digital fabrication tools have substantially changed instrument development. Desktop 3D printers, computer-aided design tools, low-cost motion controllers, laser cutters, and hobby-grade CNC machines allow many laboratories to prototype fixtures, enclosures, mechanical interfaces, and experimental accessories without relying entirely on centralized shops. The corresponding capability is much less mature at the micron scale. Scientific parts such as high-aspect-ratio holes, optically aligned assemblies, precision stages, miniature gas bearings, micro-rotors, and embedded metrology structures are often too specialized for commodity manufacturing, too precise for ordinary desktop fabrication, and too iterative for slow outsourcing.

This thesis argues that the missing capability is best understood as a manufacturing ecosystem rather than as a single machine. Micron-scale fabrication requires coordinated control of motion, energy delivery, material selection, fixturing, inspection, and process knowledge. If these elements are isolated, a laboratory may possess capable tools while still lacking a reliable route from design to functioning part. If they are integrated into a documented workflow, the same tools become a reusable experimental capability.

### 1.2 The Fabrication Bottleneck in Scientific Instrumentation

The frontiers of scientific inquiry are often defined by the physical and engineering capabilities of instruments. This limitation is especially visible when a scientific question depends on a component that is small, fragile, material-specific, and not commercially standardized. A laboratory may understand the physics of an experiment and may possess the digital design tools needed to describe the required part, but still be unable to fabricate, inspect, and revise the part on a useful time scale.

The problem is structural. Outsourcing can produce individual components, but it rarely gives the laboratory direct control over the process variables that determine success. A one-off outsourced part may answer the question of whether a design can be made, but it does not necessarily reveal why one material failed, why

one fixture generated runout, why one laser setting produced damage, or why one inspection method missed the defect. For research, that missing causal process knowledge is often as important as the part itself.

Open micron-scale manufacturing is therefore not simply about lowering cost. It is about changing where iteration occurs. If the design, fabrication, measurement, and revision loop can occur inside the laboratory, then the laboratory can accumulate process knowledge as part of the experiment. The thesis uses this idea to organize a set of technical systems: precision stages, laser micromachining platforms, characterization tools, data infrastructure, and automation.

### **1.3 Laboratory Access, Reproducibility, and Iteration**

The broader motivation is access to experimental capability. At larger length scales, local digital fabrication has changed the practical structure of early-stage instrument development. A research group can often design, fabricate, test, and revise fixtures or mechanical interfaces without transferring each iteration to a centralized facility. The scientific value of that local loop is not that it matches industrial production in throughput or ultimate precision. Its value is that design intent, process error, measurement, and revision remain close enough to support rapid causal learning.

Open micron-scale manufacturing pursues the same experimental structure at smaller length scales. The objective is not to reproduce the throughput or precision envelope of mature industrial production lines in every laboratory. Rather, the objective is to make micron-scale instrument development more local, inspectable, and iterative. A laboratory should be able to specify a functional tolerance, select or build an appropriate fabrication process, measure the resulting part, identify the dominant failure mode, and modify the process on the basis of evidence.

Micron-scale parts are common in modern scientific instruments, but access to their fabrication remains uneven. A rotor for solid-state nuclear magnetic resonance, a microfluidic chip, a confocal scan mechanism, a precision optical alignment stage, and a laser-machined diamond component may appear to be unrelated artifacts. In practice, they share a set of manufacturing constraints: small features, difficult materials, high sensitivity to alignment, and the need for feedback between fabrication and characterization.

The access gap becomes especially clear when the required part must be iterated. A local fabrication ecosystem can record which material failed, which fixture produced runout, which laser setting created damage, which inspection method detected the defect, and which process change improved yield. The value of the ecosystem is therefore not only part production, but the accumulation of traceable process knowledge. In this sense, micron-scale fabrication is also a reproducibility problem: the relevant output is not only a finished part, but a documented sequence of design decisions, process parameters, inspection results, and failure modes that can be repeated, challenged, and improved.

### **1.4 Thesis Claim and Contributions**

The claim of this thesis is that open micron-scale manufacturing can be built as an integrated laboratory stack. The stack contains four coupled layers:

1. Precision motion and fixturing that make micron-scale positioning and repeatable workholding possible.
2. Laser micromachining systems that deliver controlled energy into difficult materials.
3. Characterization tools that close the loop between designed geometry, manufactured geometry, material damage, and functional performance.

4. Automation and software infrastructure that make these workflows repeatable, inspectable, and eventually self-improving.

The thesis contributes a synthesis of this stack through a concrete case study: diamond magic-angle-spinning nuclear magnetic resonance rotors. Diamond rotors require material selection, high-aspect-ratio micromachining, sub- $\mu\text{m}$ -scale concentricity, careful reversal and clamping, gas-bearing support, and multiple forms of characterization before spin testing. Prior work demonstrated the feasibility of diamond MAS rotors [1], and more recent work advanced the fabrication and spin performance of 0.7 mm diamond rotors [2]. This thesis uses that application to define the machinery, metrology, data structure, and workflow required for a laboratory to develop micron-scale scientific parts as part of its own research process.

## 1.5 Evidence Basis and Roadmap

This thesis combines completed experimental work, ongoing machine development, literature review, and proposed engineering workflows. The evidentiary status of each component is therefore stated explicitly so that completed results are not conflated with design intent, external benchmarks, or future work.

The diamond MAS-NMR rotor project is the primary workflow and case study. It defines the central process chain of the thesis: NMR motivation, diamond material selection, laser and mechanical machining, metrology feedback, and supporting hardware. The 532 nm confocal Raman platform was built in this work as a local characterization instrument, although its quantitative performance characterization remains incomplete; it is therefore treated as a demonstrated instrument-building effort and as a framework for the calibrated Raman measurements still required. The low-cost piezo stage was built or prototyped as an access-oriented precision-motion platform, and it motivates the calibration, repeatability, and cost documentation needed for reproducible micron-scale positioning.

Other components occupy different evidentiary categories. The MPMS-based personal spectrometer was replicated and improved as an educational and access-oriented spectroscopy effort, complementing the Raman system by exposing spectrometer architecture, calibration practice, and validation limits. The waterjet-guided laser system and the femtosecond machining system are in development using existing laser sources with custom optics, fluidics, and mechanical design; they are presented as in-house micromachining platforms and process-development frameworks rather than as completed production tools. Commercial stages, commercial lasers, CT, SEM, THz tools, and robotic-alignment literature provide comparison standards and design constraints unless local measurements are explicitly reported.

The remainder of the thesis follows this evidence structure. Chapter 2 develops the diamond rotor case study. Chapter 3 describes the precision motion layer. Chapter 4 develops the laser-machining layer, including femtosecond and waterjet-guided systems. Chapter 5 describes inspection, Raman/confocal systems, and spectrometer tools. Chapter 6 summarizes the thesis as an open micron-scale manufacturing ecosystem.

## Chapter 2

# Diamond MAS-NMR Rotors

### 2.1 Chapter Overview

This chapter treats the diamond magic-angle-spinning nuclear magnetic resonance (MAS-NMR) rotor as the central manufacturing case study of the thesis. The object is small, but it combines the major constraints that make scientific-instrument fabrication difficult: a brittle anisotropic material, sub-millimeter geometry, a high-aspect-ratio internal bore, low runout, low ID/OD eccentricity, gas-bearing compatibility, heat management, DNP compatibility, and repeated nondestructive inspection. A rotor is therefore not only a sample holder. It is the part where material science, laser machining, precision motion, microscopy, and NMR performance meet.

The chapter is organized around seven technical questions: why MAS-NMR requires new rotor materials, what published NMR spectra demonstrate about diamond-rotor usability, how HPHT and CVD diamond differ as rotor substrates, which laser processes can machine high-aspect-ratio diamond bores, how fixture generations improved OD/ID registration, how diamond stock and machined rotors should be characterized, and how stators, gas bearings, caps, and helium operation constrain the rotor as a system.

### 2.2 Background and Motivation: The Instrumentation Bottleneck

#### 2.2.1 The Fabrication Bottleneck in Scientific Discovery

The frontiers of scientific inquiry are often defined not only by imagination, theory, or computation, but by whether the required instrument can be fabricated with the necessary precision and material properties. Structural biology, materials science, catalysis, and energy-storage research all depend on tools that can probe matter at molecular or atomic length scales. When the available hardware reaches its physical limits, discovery slows because the next scientific question requires a part that cannot yet be made repeatably.

This bottleneck is acute in solid-state NMR. MAS-NMR can determine structure and dynamics in systems that are non-crystalline, insoluble, heterogeneous, or embedded in native-like environments: amyloid fibrils, membrane proteins in lipid bilayers, viral assemblies, catalysts, battery materials, and other materials that are difficult for solution NMR, X-ray crystallography, or cryo-EM alone [3, 4]. The analytical power of MAS-NMR is therefore tied directly to the mechanical performance of the small rotor that contains the sample. If the rotor cannot spin fast enough, if it heats too much, if it fails mechanically, or if it reduces microwave delivery during DNP, then the experiment is capped by the sample hardware rather than by the spectrometer.



Figure 2.1: Bruker BioSpin AVANCE NEO console and Bruker Ascend 800 MHz NMR magnet at Francis Bitter Magnet Lab, used as the high-field experimental context for the diamond MAS-NMR rotor work.

## 2.2.2 Principles of MAS-NMR and the Resolution Challenge

Nuclear magnetic resonance exploits the spin angular momentum of nuclei such as  $^1\text{H}$ ,  $^{13}\text{C}$ , and  $^{15}\text{N}$ . When these nuclei are placed in a strong static magnetic field  $B_0$ , their spin-state degeneracy is broken by the nuclear Zeeman interaction, and each nucleus resonates at a Larmor frequency determined by its gyromagnetic ratio and local chemical environment. In solution NMR, rapid isotropic tumbling averages orientation-dependent interactions and produces sharp spectral lines.

In solids, proteins in fibrils, membrane assemblies, and rigid materials, that rapid tumbling is absent or incomplete. The spin Hamiltonian is then broadened by orientation-dependent terms, including chemical-shift anisotropy and heteronuclear or homonuclear dipolar couplings. In densely protonated biological samples, the  $^1\text{H}$ - $^1\text{H}$  dipolar network can produce broad, unresolved spectra unless it is averaged mechanically.

Magic-angle spinning restores resolution by rotating the sample about an axis inclined at the magic angle relative to  $B_0$  [5, 6]. The reason this particular angle matters is that the dominant chemical-shift anisotropy and dipole-dipole terms are second-rank tensor interactions. Their orientation dependence contains the factor  $3 \cos^2 \theta - 1$ , conventionally written as the second Legendre polynomial:

$$P_2(\cos \theta_m) = \frac{1}{2}(3 \cos^2 \theta_m - 1). \quad (2.1)$$

When this angular factor is zero, the first-order contribution from these second-rank anisotropic interactions averages away during rapid rotation. Setting  $P_2(\cos \theta_m) = 0$  gives:

$$\theta_m = \cos^{-1} \left( \frac{1}{\sqrt{3}} \right) \approx 54.74^\circ. \quad (2.2)$$

At this angle, first-order second-rank anisotropic interactions are averaged in the rotating frame. The averaging is only effective when the spinning frequency is high enough relative to the interaction strength. Dense proton networks in folded proteins can have homonuclear couplings on the order of tens of kilohertz, and simulations and experiments on selectively methyl-protonated proteins show that MAS frequencies above 300 kHz may be required to recover more than 80% of the maximum proton-detected sensitivity in some spin systems [7]. This does not mean every solid-state NMR experiment requires 300 kHz. It means rotor speed is a quantitative experimental variable, and future high-resolution proton-detected methods are limited by the materials and manufacturing of the rotor.

### 2.2.3 Homogeneous and Inhomogeneous Linewidths

The push toward ultra-fast MAS is strongest in biological targets with severe spectral crowding. The observed linewidth contains inhomogeneous and homogeneous contributions. Inhomogeneous broadening comes from distributions of local structure, conformation, magnetic susceptibility, or chemical environment. Faster spinning cannot make a heterogeneous sample homogeneous. Homogeneous broadening, by contrast, is associated with active spin-spin interactions and transverse relaxation. Increasing MAS frequency can reduce this mechanically averageable component and reveal spectral information otherwise hidden under dipolar broadening.

Table 2.1: Representative systems motivating ultrafast MAS-NMR, based on applications reviewed in modern solid-state NMR and ultrafast MAS literature [3, 4].

Biological or material target	Significance and structural challenge	NMR resolution constraint
Alzheimer’s fibrils	Pathogenic amyloid aggregates that are insoluble and polymorphic	Severe spectral overlap; high resolution is needed to distinguish polymorphic structural variations
Influenza M2 channels	Viral tetrameric proton channels in lipid bilayers; drug target	Spectra depend on lipid environment and are broadened by dynamics and proton networks
Alpha-synuclein	Fibrillar protein aggregates associated with Parkinson’s disease	Rigid fibril cores and disordered regions complicate assignment without fast MAS
Pf1 bacteriophage	Filamentous viral assembly used as a model macromolecular system	Cryogenic experiments show strong relaxation and linewidth constraints
Lithium-ion battery materials	Paramagnetic and heterogeneous microstructures in energy storage	Paramagnetic and electron-nuclear interactions broaden spectra and motivate fast spinning

Cryogenic operation introduces an important linewidth dilemma. At low temperature, molecular motion freezes, conformational distributions become more static, and apparent linewidths may broaden even when refocused homogeneous decay times improve. A Pf1 bacteriophage comparison reviewed by Sergeyev and co-workers reports that, at 11–12 kHz MAS, cooling from 273 K to 100 K shortened apparent linewidth-limited  $T_2^*$  values from approximately 4.5 to 1.7 ms for  $^{13}\text{C}$  and from approximately 2.6 to 1.0 ms for  $^{15}\text{N}$ . At the same time, refocused homogeneous  $T_2'$  values increased from approximately 5.1 to 10.5 ms for  $^{13}\text{C}$  and from approximately 5.5 to 41.7 ms for  $^{15}\text{N}$  [8]. The point is that cooling can improve homogeneous relaxation

while worsening the apparent spectrum through frozen disorder, so cryogenic DNP does not remove the need for faster and more stable MAS hardware.

## 2.2.4 Sensitivity Requirement and Dynamic Nuclear Polarization

The same rotor miniaturization that enables higher MAS frequencies also reduces sample volume. NMR is intrinsically insensitive because nuclear Zeeman energy splittings are small relative to thermal energy. Only a tiny population difference contributes to observable magnetization. When the rotor volume is reduced to a fraction of a microliter, the number of spins drops sharply, so signal averaging can become impractically long.

Dynamic Nuclear Polarization (DNP) addresses this sensitivity crisis by transferring polarization from electron spins to nuclear spins. The electron gyromagnetic ratio is approximately 660 times that of the proton, so electron spins possess much larger thermal polarization under the same field and temperature conditions [9]. In high-field DNP, samples are doped with paramagnetic polarizing agents and irradiated with high-frequency microwaves, often under cryogenic MAS conditions. This couples the sensitivity problem back to rotor material selection.

DNP-relevant rotors must satisfy three coupled material requirements:

**Terahertz (THz) and Microwave Transparency.** The rotor material must transmit the microwave or THz field efficiently enough that the sample receives the polarization-transfer drive field. Strong absorption or reflection attenuates the field and can produce local heating.

**Thermal Conductivity.** Gas-bearing friction and aerodynamic drag generate heat at high MAS frequencies. High thermal conductivity helps dissipate heat and preserve cryogenic sample conditions.

**Low-Temperature Mechanical Stability.** The rotor must withstand centrifugal stress at cryogenic temperature, where many materials become less tolerant of flaw growth or brittle fracture.

Diamond is attractive in all three categories. It combines high thermal conductivity, low dielectric loss in relevant ranges, and high strength. HPHT type Ib diamond adds another possible advantage: substitutional nitrogen defects, or P1 centers, can act as electron-spin reservoirs. Recent work on P1-rich diamond particles demonstrated microwave-driven DNP of bulk  $^{13}\text{C}$  spins with large room-temperature enhancements, establishing P1 diamond as a scientifically credible intrinsic-polarization platform [10]. For rotor design, that result motivates P1-rich HPHT diamond while leaving the analyte-polarization problem as an open engineering question.

## 2.2.5 Material Limits and the Transition to Diamond

For decades, yttria-stabilized zirconia (YSZ) has been a standard MAS rotor material. Commercial 0.7 mm YSZ rotors are highly successful, but rotor literature emphasizes that pushing far beyond the present ultrafast regime introduces both gas-dynamic and material limits [1, 2]. The tangential velocity of a rotor is:

$$U = \pi D \nu_r, \quad (2.3)$$

so a 0.7 mm rotor reaches approximately 244 m/s at 111 kHz, 330 m/s at 150 kHz, 440 m/s at 200 kHz, and 660 m/s at 300 kHz. At 293 K, the ideal-gas sound speed is approximately 343 m/s in air, 349 m/s in nitrogen, and 1007 m/s in helium, using  $a = \sqrt{\gamma R_s T}$  [11, 12]. The same rotor therefore approaches transonic surface speeds in air or nitrogen by roughly 150 kHz, while remaining at much lower Mach number in helium. Helium does not make the rotor stable because it reaches the speed of sound; rather, its higher sound speed

keeps the bearing and drive flow farther from sonic choking, shock-like losses, and severe compressibility penalties at the same rotor surface velocity [13]. The mechanical stress also grows rapidly with angular velocity, and microscopic flaws can become burst-initiation sites.

Table 2.2: Surface-speed and Mach-number calculation for a 0.7 mm rotor at room temperature. The calculation uses  $U = \pi D\nu_r$ ,  $a_{\text{air}} \approx 343$  m/s, and  $a_{\text{He}} \approx 1007$  m/s.

MAS frequency	Surface speed	Mach number in air	Mach number in helium
111 kHz	244 m/s	0.71	0.24
150 kHz	330 m/s	0.96	0.33
200 kHz	440 m/s	1.28	0.44
300 kHz	660 m/s	1.92	0.66

Single-crystal HPHT diamond is therefore introduced as a functional material choice for the combined speed, heat, microwave, and fracture limits of conventional rotor materials. The material choice must still be cautious, because diamond is extremely hard but also brittle and anisotropic.

Table 2.3: Comparison of candidate MAS rotor materials relevant to high-frequency and DNP-compatible rotor design. Diamond values are material- and defect-state-dependent, so the table is a design comparison rather than a universal supplier specification [1, 2, 14].

Material	Density	Fracture-toughness range	THz transparency	Feasibility for > 150 kHz MAS
YSZ	$\sim 6.0$ g/cm <sup>3</sup>	3–8 MPa m <sup>1/2</sup>	Poorer than diamond	Baseline, but burst-limited at extreme speed
Sapphire	3.98 g/cm <sup>3</sup>	< 3 MPa m <sup>1/2</sup>	High	Too fragile for this rotor application
CVD diamond	3.52 g/cm <sup>3</sup>	ordinary colorless single-crystal CVD: $\sim 4$ –12 MPa m <sup>1/2</sup> [14]	Excellent	Chemically attractive but defect- and cleavage-limited in rotor trials
HPHT diamond	3.52 g/cm <sup>3</sup>	type IIa: 4.2–5.6; type Ib yellow: $\sim 8$ –12 MPa m <sup>1/2</sup> [14]	Excellent	Preferred material candidate for functional diamond rotors

The fracture-toughness entries in Table 2.3 are deliberately conservative. Liang et al. report ordinary colorless single-crystal CVD diamond near  $8 \pm 4$  MPa m<sup>1/2</sup>, synthetic type Ib yellow diamond near  $10 \pm 2$  MPa m<sup>1/2</sup>, and literature type IIa values of 4.2–5.6 MPa m<sup>1/2</sup> [14]. The much higher boron/nitrogen co-doped CVD values and annealed-CVD values are discussed later as toughening mechanisms, not assumed as normal purchasable rotor stock.

## 2.3 Diamond Selection for NMR Rotors

To use diamond as a rotor material, the chapter must separate point-defect chemistry from extended-defect structure. Chemical purity alone is not enough. The rotor must survive laser machining, OD/ID registration, cap insertion, packing, and spin testing. The selection problem therefore depends on synthesis method, dislocation density, crystallographic orientation, fracture mechanics, and the intended NMR or DNP function of the rotor.

Diamond type is a useful shorthand for point-defect chemistry, but it is not a complete rotor-selection label. Type Ia diamonds contain nitrogen in aggregated clusters and account for approximately 95% of natural diamonds; they commonly range from near-colorless to light yellow and are historically associated with “cape” diamonds. Type Ib diamonds contain isolated substitutional nitrogen, are often bright yellow or “canary” in color, and are rare in nature but important in synthetic HPHT material. Type IIa diamonds have no measurable nitrogen or boron impurities and are usually colorless, although gray, brown, yellow, or pink

variants can occur; this is the chemically purest diamond type. Type IIb diamonds contain boron, conduct electricity, and are commonly blue or gray-blue because boron-related absorption modifies the visible spectrum [15]. For MAS rotors, the practical interpretation is that color and type help flag impurity state, but rotor performance still depends on extended defects, orientation, residual stress, and machining response.

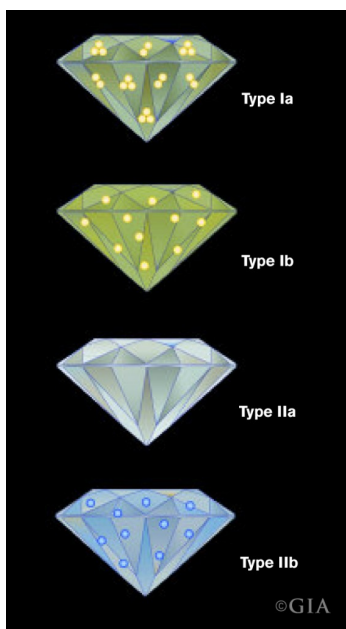


Figure 2.2: Diamond type chart from GIA. Top to bottom: type Ia, type Ib, type IIa, type IIb [15].

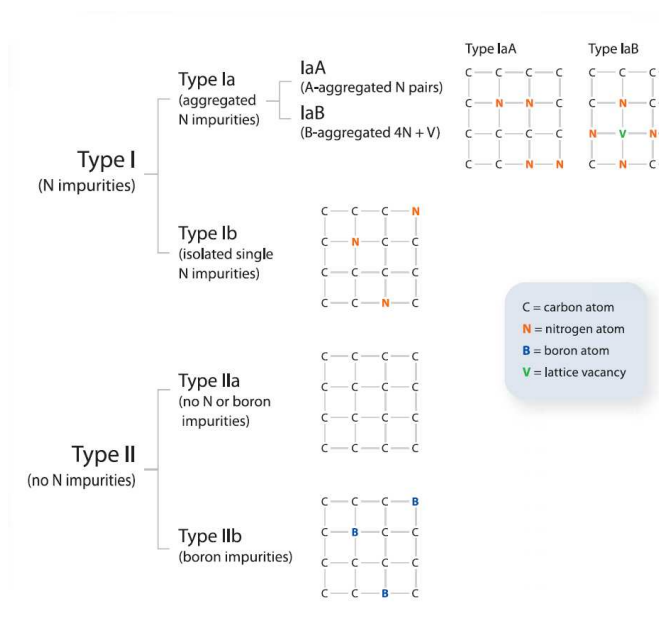


Figure 2.3: GIA schematic showing how nitrogen and boron atoms can replace carbon atoms in the diamond lattice. The diagram links type Ia and Ib diamonds to nitrogen aggregation state, type IIa to absence of measurable nitrogen or boron, and type IIb to boron substitution [15].

### 2.3.1 Carbon Phase Stability and Synthesis Routes

Synthetic single-crystal diamond is typically produced by High-Pressure High-Temperature (HPHT) growth or by Chemical Vapor Deposition (CVD). The two methods make visually and structurally different starting materials, as shown in Fig. 2.4, and those differences matter because the rotor blank must later be oriented, held, laser drilled, inspected, and spun at high frequency. For a rotor, the synthesis method is not background material; it determines the balance among purity, dislocation density, growth-sector structure, inclusion risk, size, price, and machining survivability.

**Carbon phase stability.** The diamond material problem begins with the carbon phase diagram. At ambient pressure, graphite is the thermodynamically stable solid form of carbon, while diamond is metastable. Diamond becomes the stable solid phase only at elevated pressure, so HPHT synthesis uses multi-anvil equipment to enter or approach the diamond-stability field. The phase diagram is not a manufacturing recipe by itself, because graphite-to-diamond conversion is controlled by both equilibrium and kinetics. Bundy’s updated carbon phase-diagram review emphasizes that large activation energies and path dependence allow metastable and hybrid carbon forms to persist after quenching [16, 17]. For rotor feedstock, this means the same nominal “diamond” phase can arrive as faceted HPHT single crystal, polycrystalline direct-conversion diamond, nanodiamond powder, or CVD plate, each with different defect structure and different machining risk.



Figure 2.4: Typical rough single-crystal diamond morphologies from HPHT and CVD growth. HPHT rough is commonly faceted; CVD rough is commonly square or plate-like because it grows on a planar seed.




Figure 2.5: Carbon pressure-temperature diagram showing graphite, diamond, liquid carbon, and representative artificial synthesis windows. Adapted from Bundy [16].

The routes in Fig. 2.5 should therefore be read as process families. Catalytic HPHT growth uses pressure and temperature to make diamond thermodynamically favorable, but it also uses a molten metal solvent-catalyst to reduce the kinetic barrier for breaking  $sp^2$  graphite bonds and precipitating  $sp^3$  diamond. Direct catalyst-free graphite-to-diamond conversion removes the metal solvent but requires much more severe conditions; Bundy reported that spontaneous graphite-to-diamond transformation occurs at pressures and temperatures roughly twice those required with conventional catalyst-solvents [18]. Recent direct-conversion work also shows that graphite-to-diamond transformation is governed by nanoscale graphite-diamond interfaces and can produce complex mixed or defected structures during partial transformation [19]. These routes

can produce scientifically important ultrahard or nanostructured materials, such as nano-polycrystalline or nanotwinned diamond, but they are not automatically suitable for long, low-defect, sub-millimeter single-crystal rotor blanks [20, 21].

Shock-wave or explosive synthesis occupies a still different region of the map. The pressure and temperature excursion is extreme but brief, and the rapid quench typically produces nanodiamond powder rather than large single-crystal plates or rods; detonation nanodiamonds are commonly nanometer-scale in representative reports [22, 23]. These powders are important for nanomaterials, polishing, and surface chemistry, but they are not appropriate feedstock for 0.7 mm-class single-crystal MAS rotors. HPHT and CVD therefore remain the relevant bulk single-crystal routes for the present rotor workflow.

**HPHT synthesis methodology: catalytic growth and press design.** HPHT synthesis reproduces the thermodynamic diamond-stability regime using cubic, belt, or related high-pressure presses. Carbon, often graphite, is exposed to approximately 5.0–6.5 GPa and temperatures near 1300–1600 °C in representative literature [24–27]. Metal solvent catalysts such as Fe, Ni, Co, or alloys dissolve carbon and lower the activation barrier for crystallization. Wang et al. show that, for an iron-based catalyst, the eutectic melt between carbon and catalyst is central to the graphite-to-diamond transformation; the catalyst lowers the environmental conditions needed to destroy graphite C–C bonds and form  $sp^3$  carbon in the diamond-stability field [28]. A thermal gradient then drives carbon precipitation onto an oriented diamond seed, and growth can proceed slowly over tens or hundreds of hours.

The press architecture in Fig. 2.6 is important because pressure is not only a thermodynamic variable; it is a machine-design and cost variable. For a fixed sample-cell area, the required load scales with pressure. Raising pressure also increases demands on anvils, gaskets, pressure media, heater materials, thermal insulation, and cell reproducibility. Higher temperature adds chemical compatibility and heater-life constraints. Catalytic HPHT growth is therefore attractive industrially because it moves diamond growth into a lower-pressure/lower-temperature window than direct conversion, allowing larger sample volumes and more repeatable equipment. Large-volume cubic presses exploit six opposed anvils to compress a roughly cubic cell from all orthogonal directions, which is why this geometry has become important for lower-cost industrial synthetic-diamond production.




Figure 2.6: Cubic HPHT diamond press. The left panel shows the exterior of a six-anvil press, and the right panel shows the opposed-anvil interior. Adapted from Deljanin [29] and Diamond Land [30].

The same process that makes HPHT scalable also introduces material tradeoffs. Catalyst-assisted growth can introduce metallic inclusions or growth-sector heterogeneity, and ordinary HPHT diamond often incorporates substitutional nitrogen to form type Ib material with P1 centers. Nitrogen getters such as Ti, Zr, or Al can bind nitrogen as stable nitrides, allowing lower-nitrogen type IIa material. Type IIa HPHT growth can produce very high crystalline perfection, but the process is more specialized and less broadly available than ordinary type Ib HPHT production [27, 31]. For the present thesis, type Ib HPHT is attractive when P1 centers are useful for DNP-oriented concepts, while type IIa HPHT is attractive when the rotor should behave as a more passive container.




Figure 2.7: Internal HPHT growth assembly. The right panel, from Fang et al. [24], labels the growth-cell stack: 1, steel cap; 2, Cu + Mo sheet; 3, NaCl + ZrO<sub>2</sub> sleeve; 4, graphite heater; 5, carbon source; 6, metal catalyst; 7, seed crystal; 8, ZrO<sub>2</sub> + MgO pillar; 9, pyrophyllite.

The labeled stack in Fig. 2.7 shows how HPHT growth is controlled locally inside the press. The steel cap transfers load and closes the reaction cell. The Cu + Mo sheet helps distribute heat and current while separating reactive components. The NaCl + ZrO<sub>2</sub> sleeve acts as a pressure medium and thermal/electrical insulator around the hot zone. The graphite heater supplies resistive heat. The carbon source provides feedstock, and the metal catalyst dissolves carbon and enables diamond precipitation at lower pressure and temperature than direct graphite conversion. The seed crystal defines the growth orientation, while the ZrO<sub>2</sub> + MgO pillar supports the seed and helps maintain the thermal gradient. The surrounding pyrophyllite behaves as a gasket and pressure/thermal insulation material. These parts are included because diamond quality is determined not only by nominal pressure and temperature, but by gradients, containment, chemical compatibility, and seed positioning inside the cell.

**CVD synthesis methodology: growth and defect inheritance.** CVD diamond is grown from a carbon-containing gas, commonly methane diluted in hydrogen, activated by plasma in a low-pressure reactor. Literature process windows vary by reactor and growth target, but substrate temperatures of roughly 800–1100 °C are common for microwave-plasma CVD diamond growth [32, 33]. CVD does not follow the bulk equilibrium route in Fig. 2.5. Instead, diamond is stabilized metastably at the growth surface by activated carbon species and hydrogen-rich chemistry that suppresses graphitic sp<sup>2</sup> carbon. This is the source of CVD’s value: it avoids metal-catalyst contamination, allows excellent chemical purity and isotopic or dopant control, and can produce plate geometries that are difficult by HPHT.

The same low-pressure, surface-kinetic growth mode also creates a different defect problem. CVD does not use pressure to equilibrate the whole crystal volume; it builds the crystal layer by layer on a seed. Dislocations, stacking faults, polishing damage, and strain fields in the seed can propagate upward into the CVD layer, and new dislocations can multiply during growth when surface morphology, impurity incorporation,

or thermal stress is not controlled. Misra's review of X-ray topography on single-crystal diamond reports that type IIa HPHT material can be nearly free of dislocations in optimized cases, whereas CVD single crystals often contain aggregate dislocations and dislocation bundles that originate from the substrate and propagate through angular cones during growth [31]. For MAS rotors, this means CVD's chemical purity is not enough. A CVD blank with dislocation bundles or residual stress can crack during high-aspect-ratio laser drilling even if point-defect impurity levels are low.





Figure 2.8: Microwave-plasma CVD diamond reactor architecture, showing gas/vacuum handling, microwave coupling, plasma generation, substrate support, and multi-reactor production hardware. Adapted from Diamond Land [34].




Figure 2.9: CVD diamond plates during plasma exposure and after removal from the reactor. The square plate geometry follows from homoepitaxial growth on planar seeds.

In Fig. 2.9, the dark edge material around the CVD plates should not be treated as clean single-crystal diamond without verification. Edge overgrowth, local non-diamond carbon, or morphology changes can occur where plasma density, substrate temperature, gas transport, and carbon supersaturation differ from the central growth region. Raman mapping, birefringence inspection, X-ray topography, SEM, and pilot machining are therefore more important for CVD than for simply verifying chemical purity.

**Post-growth annealing and HPHT treatment.** Post-growth treatment can improve diamond, but it should be framed as defect engineering rather than as a universal repair step. Low-pressure/high-temperature

(LPHT) annealing of high-growth-rate single-crystal CVD diamond has been shown to reduce optical absorption and change nitrogen-vacancy-hydrogen related defect structure without graphitization under hydrogen plasma conditions up to 2200 °C [35]. HPHT annealing can also modify CVD material inside the diamond-stability field. Charles et al. reported that a dark, highly absorbing nitrogen-doped CVD layer became clear after a 1 h anneal at 7 GPa and 2200 °C, with substantial changes in infrared, EPR, and photoluminescence defect signatures [36]. Ueda et al. found that HPHT annealing of CVD homoepitaxial diamond increased free-exciton emission, reduced defect-related cathodoluminescence bands, and improved room-temperature hole mobility from 826 to 1030 cm<sup>2</sup> V<sup>-1</sup> s<sup>-1</sup> [37].

These results matter for rotor selection because they show that CVD diamond is not a single fixed material state. Annealing can reduce or transform optically active point defects, relax some growth-related damage, and improve optical or electronic metrics. It cannot be assumed to remove macroscopic cracks, inclusions, poor seed inheritance, dislocation bundles, or laser-induced damage. The practical workflow should therefore treat annealed CVD as a candidate material class: useful enough to test, but still requiring Raman, birefringence or X-ray topography, micro-CT, SEM, and machining trials before it is considered equivalent to optimized HPHT stock for a functional MAS rotor.

### 2.3.2 Defect Profiles and Dislocation Densities

Diamond selection requires a distinction between chemical and structural defects. HPHT growth occurs through a liquid-solid process in a metal solvent-catalyst environment. It can produce metallic inclusions and nitrogen incorporation, but the thermodynamic growth method can also produce high structural continuity. CVD diamond can achieve exceptional chemical purity, but its dislocation density and residual stress depend strongly on seed quality, homoepitaxial growth conditions, and post-growth treatment [26, 32, 33]. For rotor fabrication, crack-initiating extended defects can be more dangerous than point-defect impurity concentration alone.

CVD diamond can be chemically purer because it avoids metal flux inclusions and allows very low nitrogen incorporation. That chemical purity is useful for optics, electronics, sensors, and some NMR contexts. However, commercial CVD stock can contain dense dislocation bundles and residual stress fields caused by homoepitaxial growth imperfections or thermal mismatch. During nanosecond laser micromachining, those defects can become stress concentrators that initiate cleavage or chipping.

HPHT type Ib diamond contains P1 centers, often at tens to hundreds of ppm in ordinary material. For passive rotors, those defects may be a tradeoff. For DNP-oriented diamond rotors, they become a possible asset because the electron spins in the diamond lattice may provide polarization. The design rule is therefore conditional: type Ib HPHT is favored when P1 centers are part of the experiment; type IIa HPHT is favored when the rotor should be as spectroscopically passive and structurally clean as possible.

### 2.3.3 Crystallographic Orientation and Cleavage Mechanics

Diamond is mechanically anisotropic. Theoretical work predicts ideal tensile strength near 225 GPa along  $\langle 100 \rangle$ , lower strength along  $\langle 110 \rangle$ , and a pronounced minimum along  $\langle 111 \rangle$  [38]. This directionality explains why the  $\{111\}$  family is the dominant cleavage family in diamond. The rotor literature therefore machines along the  $[100]$  direction, so the bore and long axis avoid direct coupling to the weakest cleavage geometry [1, 2].




Figure 2.10: Diamond machining and cleavage-risk evidence. Panels show SEM-visible cleavage, micro-CT evidence of internal chipping in CVD stock, and the  $54.76^\circ$  relation between  $[100]$  machining and  $[111]$ -related cleavage in the diamond lattice. Adapted from Schaffer et al. [2].

For rotor blanks,  $[100]$  alignment reduces one major failure mode but does not remove cleavage risk. The  $\{111\}$  planes remain present at oblique angles to the bore axis, so residual stress, laser-induced graphitization, chips, or dislocation bundles can still trigger fracture along crystallographic planes. This is why orientation control must be paired with defect screening and post-machining micro-CT or SEM feedback rather than treated as a complete solution by itself.

### 2.3.4 Fracture Toughness and Brittle Failure

Single-crystal diamond has exceptional hardness, but it is still a brittle covalent solid. Published indentation-based measurements show that fracture toughness varies strongly with diamond type, orientation, doping, and residual stress; Liang et al. report typical type IIa values around  $4.2\text{--}5.6 \text{ MPa m}^{1/2}$ , type Ib synthetic yellow diamond around  $10(2) \text{ MPa m}^{1/2}$ , and boron/nitrogen co-doped CVD diamond between 22 and 34  $\text{MPa m}^{1/2}$  under their measurement conditions [14]. Because diamond cannot rely on room-temperature plastic deformation to blunt a crack tip, failure is governed by stress concentration and Griffith-type brittle fracture. A microscopic flaw, laser-induced surface roughness, an inclusion, or a dislocation cluster can release strain energy by forming new cleavage surfaces, especially along  $\{111\}$ .

This is the reason material selection and machining cannot be separated. A diamond with better point-defect purity but more extended dislocation bundles may fail sooner than a less chemically pure HPHT type Ib crystal with a more continuous lattice. It also explains why CT, Raman, SEM, and spin-test feedback must be part of the fabrication process rather than post hoc documentation.

### 2.3.5 Diamond Selection Criteria and Working Evidence

The diamond-selection criteria are grounded in the three Center for Bits and Atoms / Francis Bitter Magnet Laboratory diamond-rotor papers: the first diamond rotor demonstration, the high-aspect-ratio nanosecond laser-machining study, and the later high-frequency rotor fabrication paper [1, 2, 39]. Together, these papers make the selection logic more specific than "diamond is strong." The material has to be THz-compatible for DNP, thermally conductive enough for high-speed operation, mechanically strong enough for centrifugal

loading, available as [100]-oriented single-crystal stock, and machinable into a 0.7 mm-class rotor with a long internal bore.

The rotor-geometry evidence below shows why material selection cannot be separated from manufacturability. The demonstrated diamond rotor is not a simple tube; it must combine a 0.7 mm outer diameter, a sub-millimeter internal bore, millimeter-scale length, cap interfaces, low taper, and sufficient concentricity for stable MAS. The high-aspect-ratio machining paper establishes that type Ib HPHT diamond with four-point [100] orientation can be machined using 532 nm nanosecond laser processing, with rotary-assisted drilling reaching average aspect ratios near 40:1 and maximum reported aspect ratios up to 66:1 [39]. The later high-frequency rotor paper adds the process evidence that OD-first machining, micro-CT inspection, and centerless ID machining improve ID/OD concentricity and produce rotors compatible with commercial 0.7 mm MAS hardware [2].




Figure 2.11: Representative diamond MAS rotor geometry and feature scale. The figure shows the 0.7 mm-class rotor envelope, physical rotor assembly, and tomographic views used to verify the internal bore and wall geometry. This scale illustrates why the material must combine high strength with laser machinability, low taper, and low ID/OD eccentricity. Adapted from Golota et al. [1].

The internal rotor records provide working evidence for applying these criteria to project material choices. They are process-development evidence rather than a controlled supplier study. They show how the technical problem evolved from feasibility machining to material screening, OD/ID registration, centerless fixturing, cap fit, and spin testing. The parsed log contains 108 rotor-like records: 44 survived or were usable, 8 were marginal or not cleanly usable, 55 were broken or had major defects, and 1 remained unknown. Grouped by material class, HPHT accounts for 88 records with 43 usable outcomes, CVD accounts for 18 records with only 1 usable outcome, and polycrystalline stock accounts for 2 records with no usable outcome. In these records, HPHT single-crystal diamond is therefore the only material class that is repeatedly viable, while CVD and polycrystalline stock show more frequent chipping, cracking, source defects, and brittle failure.

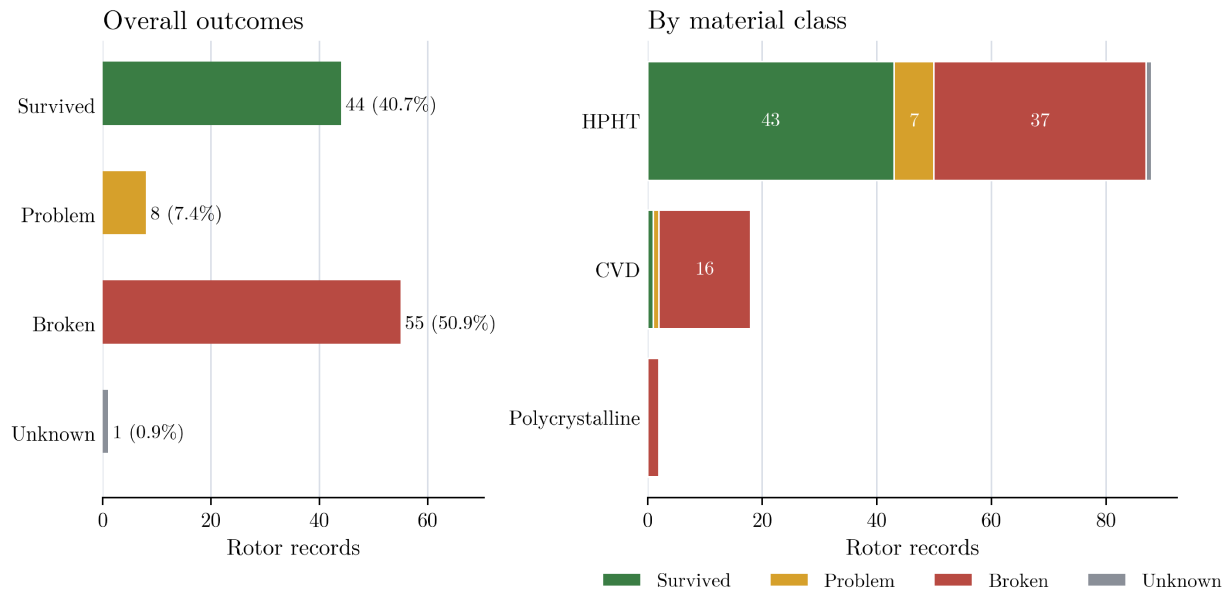


Figure 2.12: Summary of diamond rotor fabrication outcomes from the internal parsed rotor log. The chart summarizes 108 rotor-like records and groups material classes from the vendor/method and source-group fields. Rows previously labeled as unspecified are counted as HPHT. The classification is used as working material-selection evidence rather than as a controlled supplier-comparison study.

**Material-selection summary.** The working material decision is therefore conditional rather than a simple HPHT/CVD ranking. YSZ remains the commercial baseline, but its density, burst-speed limits, and poorer DNP thermal/microwave margin make it a comparison material rather than the main direction for this thesis case study. Use [100]-oriented single-crystal diamond unless a specific experiment justifies another orientation, and screen extended defects and residual stress before assuming chemical purity implies mechanical reliability. HPHT type Ib is preferred when P1 centers are useful and the stock passes orientation and defect screening; HPHT type IIa is the cleaner passive-container direction when intrinsic electron-spin density is not desired. CVD single crystal remains promising because of its chemical purity, but it should be treated as stock-dependent and only advanced after defect mapping, pilot machining, and CT/SEM confirmation. Nanotwinned or nano-polycrystalline diamond are material-science development targets rather than established functional rotor candidates, and secondary-phase PCD is excluded from the functional rotor path because binder and secondary phases introduce NMR, RF/microwave, and laser-machining risks.

### 2.3.6 State-of-the-Art Toughening Concepts

Diamond toughening methods can be grouped by the mechanism that resists crack initiation or changes the crack path: impurity engineering, residual-stress management, coherent interface toughening, nanograin crack deflection, twin-boundary strengthening, and secondary-phase toughening. For MAS-NMR rotors, these mechanisms must be evaluated against a stricter standard than hardness or indentation toughness alone. The material must also be available as a dense sub-millimeter blank, must tolerate high-aspect-ratio laser machining, must avoid magnetic or conductive secondary phases, and must remain compatible with RF, microwave, and DNP operation.

**CVD method 1: boron/nitrogen co-doped single-crystal CVD diamond.** The highest reported toughness numbers for single-crystal diamond are associated with impurity-engineered CVD growth. Liang et al.

modified high-density plasma CVD growth to incorporate boron and nitrogen and reported boron/nitrogen co-doped single-crystal CVD diamond with fracture toughness between 22 and 34 MPa m<sup>1/2</sup> while retaining hardness near 78 GPa [14]. In the same study, low-pressure/high-temperature post-growth annealing produced hardness of at least roughly 125 GPa with fracture toughness of 12–16 MPa m<sup>1/2</sup>. Mechanistically, the improvement is attributed to impurity-related changes in residual stress, local bonding, and crack propagation, rather than to a binder phase. For MAS rotors, this method has the best performance potential among single-crystal CVD options because it preserves a continuous diamond lattice and avoids the NMR and microwave penalties of metallic binders. The main feasibility risk is that the reported material is not automatically equivalent to commercially available, long, low-defect, [100]-oriented 0.7 mm rotor stock. Co-doped CVD would still require Raman, CT, EBSD, and spin-test validation because CVD growth can also generate dislocation bundles, residual-stress gradients, or growth-sector boundaries that cause laser-machining failure.

**CVD method 2: periodic nitrogen-doped nano-multilayers.** A second CVD method is to grow periodic nitrogen-doped homoepitaxial multilayers. Zhao et al. report microwave-plasma CVD multilayers in which alternating tensile and compressive stresses increase fracture toughness; the highest reported value is 18.2 MPa m<sup>1/2</sup> when the single-layer thickness is about 96 nm, approximately 2.1 times the HPHT seed substrate used in that study [40]. The concept is attractive because it deliberately creates coherent interfaces that consume energy during crack initiation and propagation. It is also more structurally controllable than random defect incorporation. Its rotor feasibility is lower than co-doped bulk CVD, however, because a multilayer is inherently directional. A MAS rotor experiences hoop stress, machining stress, bore-wall stress, and cap-fit stress in multiple directions. If the multilayer architecture is not aligned to the dominant crack path, it may toughen one failure mode while leaving another unchanged. A thick, uniform, defect-free multilayer stack also has to be grown at rotor-blank scale and then survive high-aspect-ratio drilling.

**CVD method 3: post-growth annealing and stress management.** CVD diamond often contains residual stresses from thermal mismatch, high growth rate, and defect incorporation. Annealing can relax some of these stresses and improve hardness or toughness, as shown in the low-pressure/high-temperature treatment reported by Liang et al. [14]. For rotor fabrication, annealing is a feasible auxiliary method because it can be applied after material procurement and before or after machining trials. Its performance ceiling is lower than engineered co-doping or multilayers because it does not remove all extended defects or eliminate {111} cleavage. It is best understood as a stress-management method: it may reduce residual stress enough to improve machining yield, but it is unlikely to transform ordinary commercial CVD stock into a reliable rotor material by itself.

**NPD.** Nano-polycrystalline diamond (NPD) is a binderless nanograined diamond produced from graphite or other carbon precursors under static HPHT conditions. Irifune et al. reported ultrahard sintered polycrystalline diamond from graphite, and later work describes NPD as a material with high hardness and no single macroscopic {111} cleavage plane [20]. The toughening mechanism is nanograin crack deflection: a crack cannot propagate through one continuous single-crystal cleavage plane and must instead pass through many grain boundaries. For MAS NMR rotors, that fracture isotropy is attractive, but feasibility is uncertain. The material must be available as long, dense, pore-free sub-millimeter blanks; it must tolerate laser machining without grain-boundary chipping; and its grain-boundary network must not introduce NMR background, microwave loss, or debris-related failure. NPD therefore ranks high in conceptual fracture performance but low to moderate in rotor feasibility.




Figure 2.13: Nanotwinned diamond deformation model from Wen et al., showing twin-boundary-related deformation fields and dislocation interaction with the nanotwin architecture [41].

**Nanotwinned diamond.** Nanotwinned diamond represents one of the highest-performance diamond materials reported. Huang et al. synthesized bulk nanotwinned diamond from onion-carbon nanoparticles under HPHT conditions and reported an average twin thickness near 5 nm, Vickers hardness up to approximately 200 GPa, and substantially improved thermal stability [21]. The toughening and strengthening mechanism is twin-boundary strengthening: nanoscale twin boundaries obstruct crack growth and plastic flow. Figure 2.13 is useful because it shows the mechanism rather than only reporting a hardness value. In the simulated nanotwinned diamond microstructure, the twin boundaries act as repeated internal barriers. Under compression, shuffle-set dislocations and local displacement fields are interrupted by the twin architecture instead of passing through a large uninterrupted crystal volume. For a MAS rotor, this would be powerful if it could be produced as long, homogeneous, [100]-referenced stock and machined without introducing surface damage. Its feasibility remains low because the synthesis conditions are extreme, the available stock geometry is not naturally matched to rotor manufacturing, and no high-aspect-ratio MAS rotor process has been demonstrated on nanotwinned diamond [41].

**Secondary-phase method: toughened polycrystalline diamond composites.** Industrial polycrystalline diamond can be toughened by adding binder or secondary phases that deflect cracks and transfer load. In a high-thermostability Si-B polycrystalline diamond system, carbide-particle toughening has been used to improve fracture behavior during HPHT sintering [42]. This method is effective for cutting-tool materials, but it is not a feasible method for functional MAS-NMR rotors. The binder or secondary phases can introduce magnetic susceptibility, conductivity, dielectric loss, and unwanted NMR background. In addition, grain boundaries and binder-rich regions would be vulnerable during high-aspect-ratio laser drilling, where graphitization, debris retention, and local thermal mismatch already limit diamond survival. Secondary-phase PCD is therefore useful as a comparison case for fracture mechanics and tool-material design, but it should be excluded from the functional rotor material set.

This ranking is limited to material-science toughening mechanisms. Co-doped CVD, stress-managed CVD, and CVD nano-multilayers are the most relevant single-crystal directions because they can in principle preserve a continuous diamond body while improving resistance to crack initiation or propagation. Nano-polycrystalline and nanotwinned diamond offer high conceptual fracture performance, but their stock geometry, orientation control, and high-aspect-ratio machining behavior are not yet matched to MAS rotor fabrication. Secondary-phase PCD is excluded from the functional rotor path because binder and secondary-phase compatibility problems conflict with the MAS-NMR system requirements.

Table 2.4: Ranking of state-of-the-art diamond toughening methods for possible MAS-NMR rotor use.

Rank	Method	Main toughening mechanism	NMR compatibility
1	Boron/nitrogen co-doped single-crystal CVD	Bulk impurity engineering and stress modification; reported 22–34 MPa m <sup>1/2</sup> toughness	Potentially compatible if low-loss, low-dislocation stock is validated
2	CVD post-growth annealing / stress management	Relaxes residual stress and can raise hardness while retaining moderate toughness	Compatible as a treatment, but stock still requires screening
3	Periodic nitrogen-doped CVD nano-multilayers	Coherent interfaces and alternating stress fields deflect cracks; reported 18.2 MPa m <sup>1/2</sup> toughness	Potentially compatible, but layer direction and scale-up must be proven
4	Nano-polycrystalline diamond	Nanograin crack deflection and removal of continuous single-crystal cleavage planes	Unproven for NMR; grain-boundary background and machining must be tested
5	Nanotwinned diamond	Nanoscale twin-boundary strengthening; reported hardness up to ~200 GPa	Unproven for NMR; stock geometry is the main barrier
6	Secondary-phase toughened PCD	Carbide/binder load transfer and crack deflection	Not compatible because secondary phases can affect NMR, RF, and microwave behavior

## 2.4 Laser Machining for Diamond Rotor Drilling

Following material selection, the next question is how to machine a high-aspect-ratio internal bore without destroying the diamond. Wavelength, pulse duration, pulse energy, focus geometry, and motion architecture determine sidewall taper, heat-affected zone (HAZ), graphitization, debris evacuation, and cleavage risk.

### 2.4.1 Optical Absorption Physics in Diamond

Ideal diamond has a wide indirect bandgap of approximately 5.5 eV, so it is transparent through much of the visible and near-infrared spectrum. A simple one-photon bandgap argument would say that 1064 nm, 532 nm, and 355 nm light should not couple strongly to pristine diamond because their photon energies are about 1.17, 2.33, and 3.49 eV. That statement is only the starting point. Real laser machining couples through defects, surface roughness, contamination, multiphoton ionization, plasma formation, and the graphitic carbon produced during earlier pulses [39, 43, 44].

The wavelength map in Fig. 2.14 is not a measured absorption spectrum for a specific diamond supplier. It is an evidence map that separates intrinsic diamond absorption, impurity-assisted absorption, and process-induced graphitized carbon. This distinction matters because a plot of raw intrinsic diamond optical constants is visually misleading in the visible range: the Phillip and Taft extinction-coefficient record, as distributed by refractiveindex.info, has nonzero  $k$  only near the deep-UV band edge in the 200–1200 nm window and reports  $k = 0$  through most of the visible and near-IR [45, 46]. That supports the physical statement that high-purity diamond is weakly absorbing at 355, 532, and 1064 nm by simple one-photon linear absorption, but it does not by itself explain practical machining.

The solid colored curves therefore show literature-supported coupling regimes rather than quantitative absorption coefficients. Type IIa/high-purity diamond is treated as a low-linear-absorption window unless defects, surfaces, or ultrafast nonlinear intensity provide coupling. Type Ia diamond is included as the aggregated-nitrogen case, with visible-center behavior drawn from the optical-absorption review literature and GIA diamond-type descriptions [15, 47]. Type Ib HPHT diamond contains isolated nitrogen and is commonly yellow; nitrogen-related photoconductivity and absorption thresholds reported near 4.2, 3.0, 2.3, and 1.7 eV correspond to approximately 295, 413, 539, and 729 nm, so nitrogen and other defects can help initiate absorption at UV/visible machining wavelengths [15, 49]. Type IIb diamond contains boron, and boron-related absorption extends from the infrared toward the visible, often producing blue or gray-

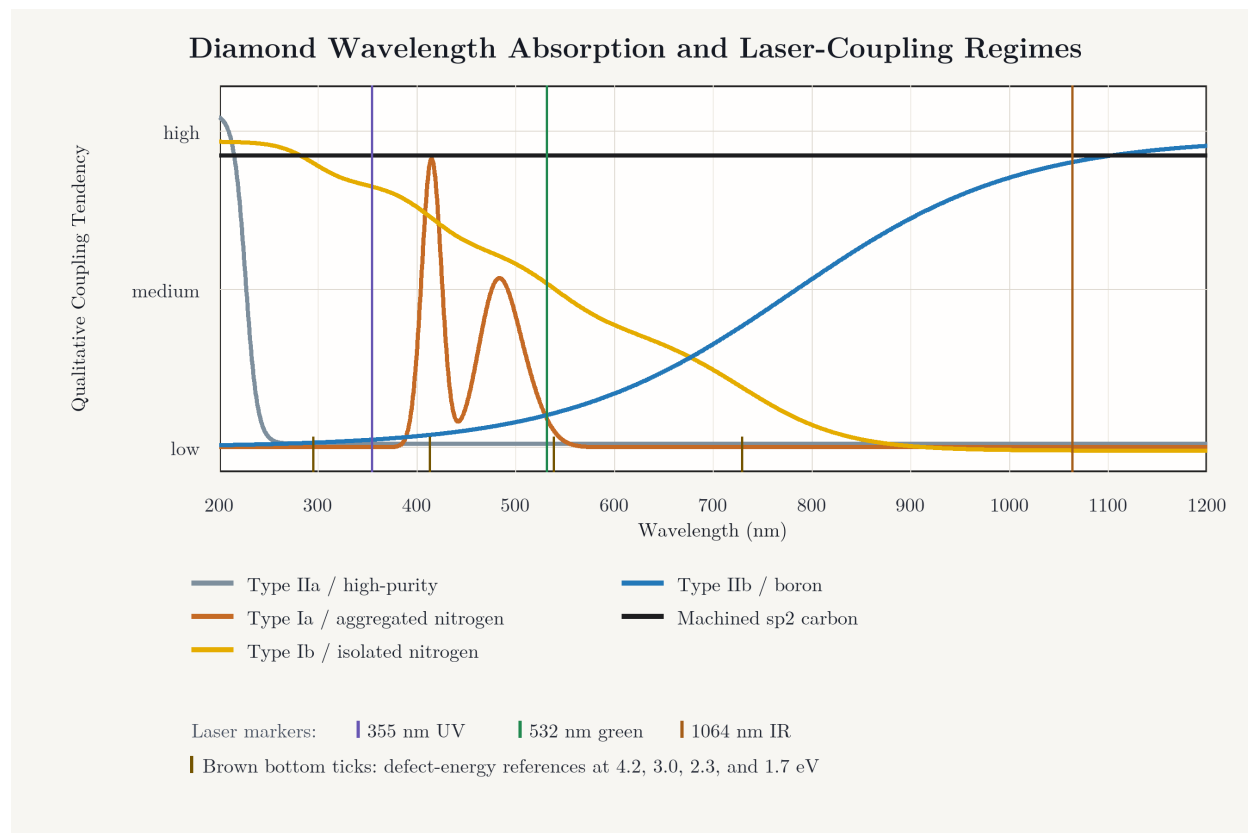


Figure 2.14: Qualitative wavelength-coupling map for diamond machining. Intrinsic or type IIa diamond is weakly absorbing through most visible/NIR wavelengths; nitrogen, boron, and process-induced graphitized carbon add coupling pathways. Vertical lines mark 355, 532, and 1064 nm [15, 39, 45–49].

blue color by preferentially removing red/NIR light [47, 48]. During nanosecond machining, the decisive practical transition is the formation of graphitized  $sp^2$  carbon, which absorbs broadly and allows subsequent pulses to couple much more efficiently than pristine diamond [39]. The reliability notes embedded below the plot separate high-confidence source-backed statements, such as intrinsic band-edge transparency and process-induced graphitization, from sample-dependent impurity strengths.

**Ultraviolet (UV - 355 nm / 3.49 eV).** UV photons do not bridge the full diamond bandgap in one photon, but they couple more readily than green or infrared light through defect states, surface states, and nitrogen-related absorption. In HPHT type Ib diamond, P1 centers and other defects can assist initial absorption. UV nanosecond tools can therefore initiate machining more easily than 1064 nm tools, but they still produce a thermal ablation process with HAZ and cracking risks.

**Green (532 nm / 2.33 eV).** Pristine diamond is largely transparent at 532 nm. In practice, nanosecond green machining works because absorption starts at defects or surface imperfections and rapidly produces a graphitized  $sp^2$  layer. Subsequent pulses couple efficiently to that dark, absorbing layer. The published high-aspect-ratio rotor machining work uses 532 nm nanosecond rotary-assisted drilling, making it the best-supported process for deep bore formation in HPHT diamond [2, 39].

**Infrared (1064 nm / 1.17 eV).** Near-infrared nanosecond machining couples weakly to polished diamond without an absorptive coating or pre-damaged surface. Once graphitization begins, removal can become aggressive, but thermal gradients can be severe. For rotor bores, 1064 nm is therefore better framed as a roughing or coating-assisted candidate than as the cleanest direct-machining method.

**Femtosecond Multiphoton Absorption (343 nm, 515 nm, 1035 nm).** Ultrashort pulses can drive nonlinear absorption because the peak intensity is high enough for multiphoton ionization and dense electron-hole plasma formation. This can reduce thermal diffusion and graphitic damage, making femtosecond machining attractive for finishing and low-damage structuring. It is not automatically the fastest route for millimeter-deep, sub-millimeter bores because deep holes introduce focus, debris, plasma-shielding, and beam-propagation problems.

## 2.4.2 Which Laser Is More Efficient for Diamond?

There is no single most efficient laser for diamond because "efficiency" can mean different things: low ablation threshold, high volumetric removal rate, low heat-affected zone, low graphitic residue, low taper in a deep bore, or minimum post-processing. Wavelength matters, but it is not the only control variable. Pulse duration, fluence relative to threshold, pulse energy, repetition rate, focal spot size, beam quality, polarization, numerical aperture, debris evacuation, and focus control can change the result as much as the nominal wavelength.

**Pulse Duration and Thermal Diffusion.** Longer pulses tend to remove material through thermal graphitization, melting/sublimation, plasma formation, and redeposition. Kononenko et al. showed that nanosecond UV-IR pulses can ablate CVD diamond, but the mechanism remains strongly coupled to graphitization and thermal modification [43]. Chichkov et al. compared femtosecond, picosecond, and nanosecond ablation of solids and showed why shorter pulses are favored for precision: energy is deposited before substantial heat diffusion into the surrounding material [50]. For diamond specifically, femtosecond processing can drive nonlinear absorption and dense electron-hole plasma formation with much less overheating around the focal zone, which makes it attractive for low-damage finishing and microstructuring [44, 51]. The tradeoff is

throughput: femtosecond machining may give the cleanest local damage profile, but it does not automatically produce the fastest deep-hole process.

**Fluence, Pulse Energy, and Repetition Rate.** The useful control variable is not only average power. A pulse must exceed the local modification or ablation threshold at the focus, while avoiding fluence so high that plasma shielding, cracking, excessive graphitization, or debris redeposition dominate. Pulse energy and spot size determine fluence; pulse duration determines peak intensity; repetition rate determines both throughput and heat accumulation. In transparent or weakly absorbing materials, femtosecond pulses can couple through multiphoton and avalanche ionization when the intensity is high enough, even if the photon energy is below the bandgap [52]. For nanosecond diamond machining, once a graphitized surface layer forms, later pulses couple more strongly because the surface is no longer optically equivalent to pristine diamond [44].

**Beam Quality, Focus, and Depth of Field.** Deep holes are not limited only by absorption. Beam quality controls how tightly and symmetrically the beam can be focused. A near-Gaussian low- $M^2$  beam gives a smaller, cleaner waist, but the Rayleigh range may be short for a high-NA focus. A lower-NA focus gives more depth of field but larger feature size. In high-aspect-ratio drilling, the focus position, plume, debris, wall reflections, and changing entrance geometry all modify the energy delivered to the bottom of the hole. This is why the best deep-hole process is usually not simple percussion drilling at one fixed focus.

**Deep-Hole Quality Strategies.** Several process strategies are used to make deeper, rounder, lower-taper holes. Percussion drilling is fast, but it tends to give more taper and poorer wall quality. Trepanning moves a smaller beam around the desired hole contour, so the hole diameter is set by the beam path rather than only by the focal spot. Helical drilling extends this idea by advancing the ablation front in many circular or spiral passes, which improves circularity, wall control, and aspect-ratio capability. Ultrafast micromachining reviews identify trepanning, helical drilling, and polarization control as major methods for improving high-precision holes [53]. Polarization is especially important in deep microdrilling because wall reflections can create asymmetric intensity distributions. Nolte et al. showed that rotating the polarization during ultrashort-pulse drilling, sometimes called polarization trepanning, can improve hole quality [54]. Additional methods include circular polarization, focus stepping or dynamic focusing, drilling from both sides when geometry allows, assist gas or liquid for debris removal, and using a low-damage finishing pass after a faster roughing pass. Literature on liquid-assisted rear-side ablation, femtosecond Bessel beams, and burst-mode femtosecond ablation also shows how debris removal, extended focal depth, and controlled heat accumulation can improve difficult microdrilling operations. In the rotor context, these methods are best treated as process-quality concepts that could supplement the main 532 nm nanosecond workflow rather than replace it.

The process choice depends on the immediate machining goal. Nanosecond UV or green pulses are favored for fast rough removal and piercing because high pulse energy and thermal graphitization give high removal rate, but the risk is HAZ formation, graphitic residue, cracking, taper, and redeposition. Femtosecond UV, green, or near-IR processing is better for clean surface finishing and damage cleanup because nonlinear absorption and short energy-deposition time reduce heat diffusion, although throughput is lower. Picosecond or femtosecond trepanning and helical drilling are attractive for high-quality microholes because they control the contour and reduce recast or HAZ, but they require precise focus and path control. Deep, low-taper holes require helical drilling, trepanning with focus stepping, polarization control, and debris evacuation to maintain energy delivery and wall symmetry as depth increases. Femtosecond tightly focused beams are also useful for subsurface modification or transparent-material processing through multiphoton or avalanche ionization, but crack control and heat accumulation remain limiting issues at high repetition rate.

For diamond rotors, the process implication is therefore a hybrid one. Nanosecond UV or green processing is usually the most efficient roughing method when the priority is removal rate, especially after a graphitized layer or defect-assisted absorption has started the process. Picosecond or femtosecond processing is more efficient when the priority is low damage, dimensional finishing, or surface cleanup. Deep holes require an optical drilling strategy such as trepanning, helical drilling, polarization control, and focus management; wavelength alone is not enough to predict success. The process narrative in this chapter remains centered on 532 nm nanosecond machining because that is the wavelength and pulse-duration regime used in the principal diamond-rotor fabrication papers and in the process evidence figures.

### 2.4.3 Laser and Motion Architectures

The laser-system comparison is therefore organized around the 532 nm nanosecond process as the benchmark rotor-manufacturing workflow. Other laser platforms are useful for rough tests, finishing, cleanup, or supplemental process development, but they are secondary to the published diamond-rotor machining sequence.

**Oxford 532 nm Nanosecond Laser with Aerotech Stage.** The Oxford system combines a 532 nm nanosecond Nd:YAG-family source with mechanical interpolation. In the rotor process, the diamond is moved under a fixed beam rather than relying only on a galvo scanner. The high-aspect-ratio diamond machining paper reports an Oxford Lasers A-series system with a Q-switched 532 nm diode-pumped Nd:YAG source, a 13–14  $\mu\text{m}$  spot, approximately 20 ns pulse duration, and 5 kHz repetition rate [39]. This is the benchmark system because it is tied directly to the published high-aspect-ratio diamond machining and rotor fabrication papers [2, 39].

**UV 355 nm and 1064 nm MOPA Galvo-Scanning Lasers.** Desktop-industrial galvo systems can scan quickly and are useful for surface patterning, rough dicing, and exploratory ablation. A 355 nm DPSS source has better defect-assisted absorption than infrared. A 1064 nm MOPA fiber laser can deliver high average power and tunable pulses, but diamond coupling is weak until an absorptive layer is present. Galvo systems are limited by depth of focus and beam clipping in deep bores, which can create V-shaped taper.

**Waterjet-Guided Laser.** Waterjet-guided laser machining uses a laminar water jet as an optical waveguide. The water column can guide the beam over long depth, cool the cutting zone, and remove debris. This is attractive for hard materials and low-taper cuts, but nozzle diameter, jet stability, minimum feature size, and integration with sub-millimeter rotor geometry remain constraints.

**PHAROS Femtosecond Platform with Precision Stage.** The femtosecond platform shifts the process toward multiphoton and plasma-mediated modification. Its advantage is low-damage finishing, smaller features, and reduced graphitization. Its challenge is throughput and deep-hole evacuation. In the thesis process chain, femtosecond machining is best positioned as a precision finishing or surface-cleanup tool unless full-depth rotor bores with the required geometry are demonstrated.

Taken together, the 532 nm nanosecond Oxford-stage workflow remains the benchmark deep-bore method because it has the strongest direct rotor evidence and demonstrated high-aspect-ratio holes. UV and MOPA galvo tools are better suited to rough tests, dicing, and surface processing: 355 nm improves initiation, while 1064 nm can supply high-power roughing once an absorptive surface exists, but depth-of-focus limits and beam clipping can create taper. Waterjet-guided machining is attractive for straight, low-taper cuts in larger features because cooling and debris evacuation are built into the process, but nozzle size and jet integration are difficult at sub-millimeter rotor scale. Femtosecond platforms are strongest for finish machining

and low-damage cleanup, with the tradeoff of lower bulk-removal throughput and unresolved deep-hole debris/plasma limits.

## 2.5 Lathe, Air-Bearing, and Clamping System Design

Rotor machining is not only an optics problem. Even a well-chosen laser cannot make a functional rotor if the mechanical stage introduces runout, if the workpiece is clamped unevenly, or if the ID and OD are referenced to different axes. At 0.7 mm scale, small eccentricities produce mass imbalance and bearing instability. Fixture evolution is therefore a major part of the rotor-manufacturing argument.

### 2.5.1 First Generation: Rotary-Assisted Stage and Kinematic Flexure Mount

The first-generation process began by testing conventional machine-tool spindles. Sherline-style or similar spindle trials showed radial runout on the order of 7.5–15  $\mu\text{m}$  depending on seating, which was too large for aerodynamic rotor stability. This drove the transition to precision optical rotary stages. The Thorlabs DDR25 was considered but rejected because axial wobble at realistic stickout translated into excessive runout. The larger Thorlabs DDR100, with brushless direct drive and lower axial wobble, became the basis of the early diamond laser lathe.

Because diamond can cleave along  $\{111\}$ , the workpiece could not be treated like ordinary metal stock. The early system used a flexural clamp integrated into a kinematic mount. A laser displacement sensor measured the workpiece at the base and tip, while X/Y and pitch/yaw adjustments were iterated to minimize radial displacement and angular wobble. This manual tramping process made runout visible and correctable.




Figure 2.15: First-generation rotary-assisted stage used for early diamond rotor laser machining. The figure documents the initial rotary-stage concept that preceded the OD-first and centerless-ID machining workflow. Adapted from patent US20210146475A1 [55].

A major advantage of the first-generation mount was reversibility. The diamond and clamp could be detached, flipped 180 degrees, and remounted so the bore could be ablated from both sides, reducing taper. The limitation was software and datum integration. The rotary stage controller was not fully integrated with the Oxford/Aerotech motion environment, and manual flipping could break the concentricity datum.

## 2.5.2 Second Generation: Split OD and Centerless ID Stages

The second-generation architecture split the process into dedicated OD and ID operations. This is the key correction to the chapter logic. The process is not simply “hold the diamond and drill a hole.” First, the outside of the diamond is shaped with high precision. Then the outside surface becomes the datum for the centerless ID operation.

For OD shaping, the second-generation stage designed and built by David Preiss uses a custom high-precision lathe with graphite air-bearing pads, a DC drive motor, belt drive, precision-ground spindle, adjustable radial preload, magnetic axial preload, pitch/yaw adjustment, a quadrature encoder, and PID-controlled speed regulation [2]. The diamond stock is mounted in a custom four-jaw chuck. The OD is machined to approximately 700  $\mu\text{m}$  with a target tolerance of about  $\pm 3 \mu\text{m}$ , then inspected before ID machining.




Figure 2.16: Second-generation outer-diameter air-bearing stage. The workpiece-defined rotational axis reduces transferred spindle runout during OD shaping. Custom-built high-precision air bearing lathe used for machining the OD of the diamond rotors. (a) A wire frame diagram showing the overall assembly, highlighting the DC drive motor, drive belt, precision ground spindle, and the custom 4-jaw chuck. (b) A sectional 3D CAD view reveals internal components, including the graphite air bearing pads that provide low-friction, high-precision rotation, the magnetic axial preload mechanism, and the positioning of the diamond stock within the chuck. (c) A photograph shows the physical assembled lathe. Adapted from Schaffer et al. [2].

For ID machining, the centerless fixture supports the rotor between follower bearings while a drive wheel or O-ring rotates it. The workpiece is not re-centered by a conventional external chuck; instead, the already-machined OD defines the rotational datum for the bore. This is why the centerless fixture improves concentricity. In Schaffer et al., the workflow includes rough ID machining from both ends at high power and 5 kHz, then finishing at reduced power and higher repetition rate to reach an ID around 470–500  $\mu\text{m}$  [2]. The paper reports that unidirectional radial passes helped reduce internal cracking and lobed fracture compared with less controlled approaches.

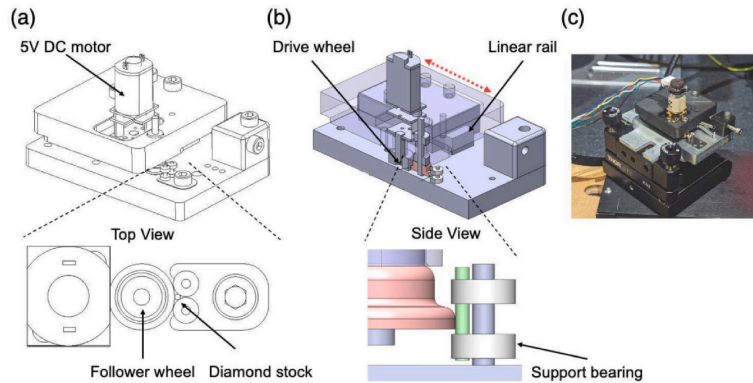


Figure 2.17: Second-generation centerless inner-diameter stage. (a) Wire-frame and top-view layout showing the drive wheel, follower wheels, and diamond stock. (b) Cross-sectional CAD view showing the drive wheel on a linear rail and support bearings that guide the rotor. (c) Assembled fixture on its positioning stage. The already-machined OD acts as the datum for bore machining, improving ID/OD concentricity. Adapted from Schaffer et al. [2].

The machining sequence for high-frequency diamond rotors follows the "Diamond rotors for high magic angle spinning frequencies" process closely because it is the most important demonstrated reference for this class of hardware. The sequence is:

1. Start with HPHT diamond logs, typically [100]-oriented stock, rather than generic CVD stock.
2. Machine the outer diameter first on the custom graphite air-bearing lathe.
3. Measure OD during or after machining with a Keyence telecentric measurement system, targeting approximately 700  $\mu\text{m}$  with  $\mu\text{m}$ -level tolerance.
4. Part off the piece and inspect it by micro-CT before committing to ID machining.
5. Machine the ID in the centerless fixture so the OD defines the rotational datum.
6. Rough from both ends, then finish with lower power and higher repetition rate to control the bore.
7. Clean graphitic residue chemically and thermally, including acid cleaning and high-temperature heat treatment where appropriate.
8. Verify final geometry with micro-CT, SEM, and gage pins, then proceed to packing and spin testing.

This workflow shifts the emphasis from laser parameters alone to fixture and datum control: OD first, CT gate, centerless ID, cleanup, CT/SEM verification, and only then NMR-relevant testing.

### 2.5.3 Through-Bore Air-Bearing Flip-Spindle Micro-Lathe

The through-bore air-bearing flip-spindle micro-lathe is the next architecture for reducing datum loss during diamond rotor machining. The air-bearing spindle supports the rotating axis on a gas film rather than on rolling contact, reducing spindle error, stick-slip, and bearing vibration. A precision collet or compliant micro-chuck holds the diamond with more uniform radial loading than point-contact clamping, while the translation stages control the laser path relative to the rotating workpiece.



Figure 2.18: Through-bore air-bearing flip-spindle micro-lathe used for registered two-sided rotor machining.

High-aspect-ratio machining benefits from accessing the workpiece from both ends, but the important requirement is to reverse the machining direction without releasing the datum. In the preferred concept, the through-bore spindle allows the diamond stock to pass through the rotary axis, and the air-bearing spindle is mounted on a motorized trunnion or B-axis so the part can be rotated 180 degrees while the workpiece coordinate system remains registered to the machine. The laser then machines from the opposite end without removing the diamond from the collet. This keeps the OD datum, bore axis, and machining coordinate system tied to the same registered setup rather than rebuilding concentricity after manual re-clamping.

For fixed-beam systems such as the 532 nm Oxford platform and the 1035 nm PHAROS platform, a low-profile trunnion could flip the workpiece under the objective while keeping the central axis locked in the motion coordinate system. As an alternative, a Swiss-style synchronous transfer could grab the exposed end of the diamond, release the main spindle, and pass the workpiece through without manual handling. This would keep the flipping concept, but it moves the risk from manual re-clamping to the transfer accuracy between two controlled spindles.

## 2.6 Characterization Methods for Single-Crystal Diamond

The characterization workflow has two branches. First, the incoming diamond substrate must be screened for orientation, inclusions, dislocations, point defects, electromagnetic loss, and fracture behavior. Second, the machined rotor must be inspected for geometry and damage: taper, OD/ID concentricity, surface graphitization, residual strain, chips, cracks, cap fit, and spin stability.

### 2.6.1 Incoming Material Characterization

**Manufacturer Defect Mapping and Sectioning.** Rough synthetic diamonds are often mapped before sectioning. Systems such as Sarine Galaxy-type scanners can use index-matching fluid to reduce surface scatter-

ing and map internal inclusions, voids, cracks, or metallic flux catalysts. The manufacturer can then choose cutting planes that avoid major internal defects before delivering rotor blanks.

**Macroscopic Crystal Orientation Verification.** Because orientation controls cleavage risk, the bulk orientation of each blank is a critical incoming-material parameter. The working process retains supplier-provided orientation documentation and, when needed, uses an appropriate macroscopic crystal-orientation measurement before machining. In the absence of direct local orientation data, this step is treated as a proposed stock-qualification gate rather than a completed project measurement.

**Chemical Composition and Raman Profiling.** Confocal Raman spectroscopy can identify the first-order diamond peak near  $1332\text{ cm}^{-1}$ , detect  $\text{sp}^2$  carbon through D and G bands near  $1350$  and  $1580\text{ cm}^{-1}$ , and evaluate residual stress through peak shifts. In the rotor workflow, Raman is both an incoming-screening tool and a post-machining damage tool. Raman spectra of doped diamond can also change through defect-related bands, peak broadening, Fano-type line-shape changes in heavily boron-doped material, and changes in second-order diamond scattering, so the interpretation should be tied to a known material history rather than treated as a standalone chemistry measurement [56–58].

The Bruker Raman comparison below shows that the GW and E6 diamond stocks are spectrally different after Savitzky–Golay smoothing, asymmetric least-squares baseline correction, and normalization to the integrated first-order diamond band. The GW sample has lower relative intensity in the second-order diamond band near  $2600\text{ cm}^{-1}$  than E6, and it also shows a weak band near  $614\text{ cm}^{-1}$  that is not apparent in E6. The  $2600\text{ cm}^{-1}$  feature is consistent with second-order diamond scattering, while the  $614\text{ cm}^{-1}$  band is consistent with a defect- or dopant-related Raman contribution. Because boron-doped diamond can show low-wavenumber Raman features and modified diamond line shapes, a boron-related contribution is a plausible interpretation for the GW spectrum [56–58]. For rotor development, the important point is that GW and E6 should be treated as distinct feedstocks during machining trials because their Raman response indicates different impurity, defect, or growth histories.

**TEM and FIB/TEM.** Transmission electron microscopy can reveal dislocation density, line defects, and lattice damage at nanometer or atomic scale. FIB preparation is required to create electron-transparent lamellae. For rotor development, TEM is most valuable when it can compare HPHT and CVD defect structures or inspect the ablation interface after machining.

**Electromagnetic Transparency by THz-TDS or VNA.** For DNP-relevant rotors, the diamond body should be screened for low dielectric loss and high transmission over microwave to THz frequencies. THz time-domain spectroscopy and vector-network-analysis methods over approximately 200 GHz to 1 THz are possible screening routes. The THz absorption measurement below directly supports the DNP-material criterion by comparing sapphire, single-crystal diamond, and zirconia over approximately 0.1–1.5 THz. Lower absorption coefficient means less field attenuation and less dielectric heating. Diamond remains close to sapphire and far below zirconia over most of the measured band, which supports diamond rather than zirconia when microwave/THz delivery and cryogenic heat management matter.

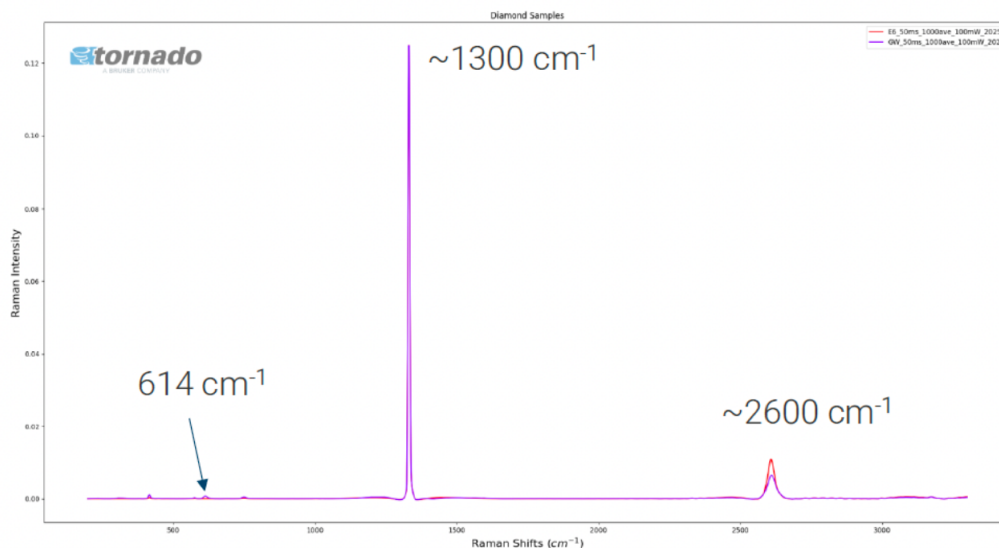


Figure 2.19: Bruker Raman comparison of GW and E6 diamond samples. The spectra were smoothed, baseline corrected, and normalized to the first-order diamond band. Relative to E6, the GW sample shows lower intensity in the second-order diamond band near  $2600\text{ cm}^{-1}$  and an additional weak feature near  $614\text{ cm}^{-1}$ , consistent with a different impurity or defect population in the two diamond feedstocks.

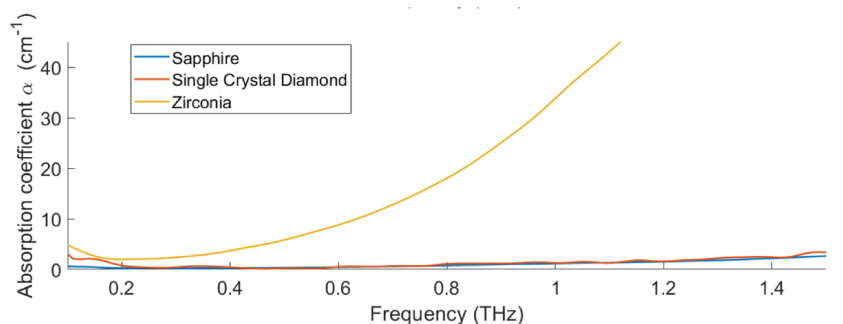


Figure 2.20: THz-TDS absorption-coefficient comparison for sapphire, single-crystal diamond, and zirconia. Diamond exhibits low absorption comparable to sapphire and substantially lower absorption than zirconia across most of the measured THz range, supporting its use as a DNP-compatible MAS rotor material. Measurements were performed with the Nelson laboratory in Chemistry.

**Fracture Toughness and Mechanical Limits.** Direct macroscopic fracture testing can destroy small diamond blanks. Nanoindentation with cube-corner tips or FIB-milled microcantilevers can estimate local fracture toughness and bending strength. These methods are treated as future or supporting material-characterization tools unless direct project data are available. The fracture-mechanics comparison from the first diamond-rotor paper connects characterized material properties back to rotor speed. It does not simply rank hardness. It estimates relative angular-frequency capability from fracture toughness, density, and Poisson's ratio, normalized to yttria-stabilized zirconia (YSZ). In that model, type Ib HPHT diamond and nitrogen-doped CVD diamond outperform YSZ because their lower density and higher fracture toughness allow higher rotor speed before crack-driven failure becomes limiting. This supports diamond selection while preserving the correct caution: actual rotors can fail below the modeled limit if laser damage, chips, cap imbalance, packing defects, or gas-bearing instability create larger flaws than assumed [1].

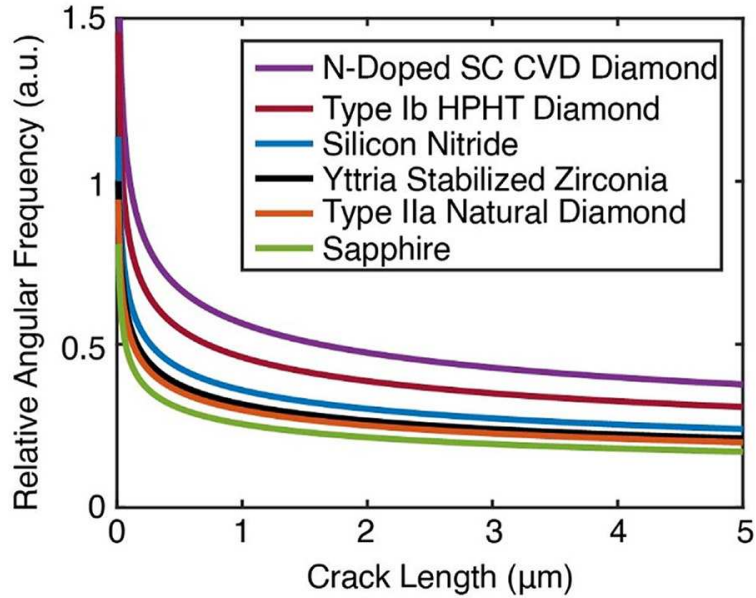


Figure 2.21: Fracture-mechanics-derived relative angular frequencies for candidate MAS rotor materials, normalized to YSZ [1].

## 2.6.2 Machining Feedback Loop

**Micro-CT.** Micro-computed tomography is the most important nondestructive tool for a machined rotor because the bore is inaccessible to ordinary probes. Schaffer et al. used micro-CT after OD machining and again for final rotor geometry, with voxel-scale inspection sufficient to identify chips, cracks, bore taper, and internal damage [2]. Micro-CT should be a pass/fail gate before investing further time in a candidate rotor. The reconstruction in Fig. 2.22 shows why this inspection is useful: internal bore continuity, sidewall condition, taper, chipping, and debris can be evaluated without cutting the rotor open.

**SEM and Electron Backscatter Diffraction (EBSD).** Scanning electron microscopy provides high-depth-of-field imaging of entry and exit holes, chipping, roughness, and surface morphology. EBSD adds local crystallographic orientation information inside the SEM. For diamond, EBSD can help distinguish exposed  $\{100\}$  and  $\{111\}$  planes, map local strain, and identify twin boundaries or orientation changes. This is especially useful when a failure surface suggests cleavage. SEM also provides direct dimensional evidence at the end face: Fig. 2.23 shows an approximately 699  $\mu\text{m}$  OD and 471  $\mu\text{m}$  ID, which is the scale required for compatibility with 0.7 mm-class MAS hardware.




Figure 2.22: Micro-CT reconstruction of a laser-machined diamond rotor section. The three-dimensional reconstruction makes the bore geometry and internal surface condition visible without destructive sectioning, allowing chips, taper, discontinuities, and residual debris to be identified before additional processing or spin testing.




Figure 2.23: SEM end-view of a laser-machined diamond rotor. The measurement overlay shows an outer diameter of approximately 699  $\mu\text{m}$  and an inner diameter of approximately 471  $\mu\text{m}$ , while the end-face morphology provides a direct check for bore shape, edge chipping, and surface damage after machining.

**Scanning Confocal Raman Depth Profiling.** Scanning confocal Raman depth profiling is a future HAZ diagnostic. Raman depth scans can identify graphitic or amorphous carbon layers, quantify stress through shifts of the  $1332\text{ cm}^{-1}$  peak, and estimate how deeply laser damage extends into the sidewall. This is a strong proposed feedback tool because it connects a laser parameter directly to a material-damage depth.

**Cathodoluminescence and Cross-Sectional FIB/TEM.** Cathodoluminescence can map defect emission and strain fields under electron-beam excitation. Cross-sectional FIB/TEM can directly image the boundary between pristine  $\text{sp}^3$  diamond and laser-modified  $\text{sp}^2$  or amorphous carbon. These are future high-resolution methods for understanding failure mechanisms, not routine production gates.

The characterization loop is therefore staged rather than interchangeable. Vendor defect mapping and bulk orientation records screen the blank before cutting. Raman, THz-TDS, and VNA measurements evaluate chemical state, residual strain, graphitic damage, and electromagnetic compatibility. TEM or FIB/TEM is reserved for material comparison and failure analysis when a destructive nanometer-scale view is justified. Micro-CT is the nondestructive process gate for bore continuity, chips, cracks, taper, and concentricity. SEM and EBSD provide the surface morphology and local crystallographic context needed to interpret edge chipping, cleavage, and roughness. Final spin testing remains the system-level acceptance gate because a rotor that passes microscopy can still fail through cap fit, imbalance, packing, or gas-bearing instability.

## 2.7 Preliminary Results: NMR Spectra from Diamond Rotors

The most important proof of the diamond rotor concept is spectroscopic rather than cosmetic or purely mechanical. A rotor must produce useful MAS-NMR spectra in a standard high-field probe without introducing unacceptable background, instability, or RF incompatibility. The first diamond-rotor paper met this requirement by recording proton-detected  $^{13}\text{C}$  and  $^{15}\text{N}$  MAS spectra from a 0.7 mm HPHT diamond rotor at 18.8 T, corresponding to an 800 MHz  $^1\text{H}$  field [1]. In that experiment, the rotor was packed with a mixture containing uniformly  $^{13}\text{C}$ ,  $^{15}\text{N}$ -labeled N-formyl-methionyl-leucyl-phenylalanine and zirconia powder, installed in a three-channel HCN 0.7 mm Bruker MAS probe, and spun at  $\omega_r/2\pi = 111\text{ kHz}$ . The two-dimensional (H)CH and (H)NH correlation spectra in Fig. 2.24 demonstrate that the machined diamond rotor, Vespel closure hardware, gas-bearing operation, cross-polarization transfers, and decoupling could operate together under ultrafast MAS conditions. This is the threshold result that changes the diamond tube from a mechanical prototype into a functional spectroscopic sample container.

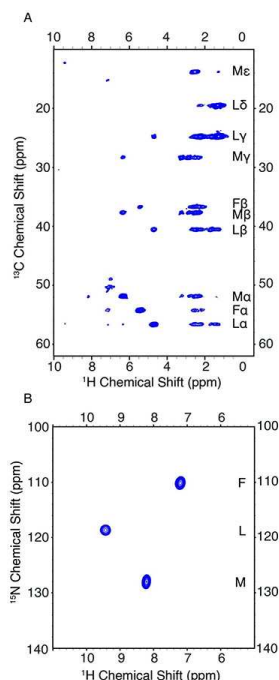


Figure 2.24:  $^1\text{H}$ -detected spectra of  $\text{U-}^{15}\text{N},^{13}\text{C-N-f-MLF-OH}$  recorded using a diamond rotor. (A) (H)CH correlation spectrum. (B) (H)NH correlation spectrum [1].

The later high-frequency diamond-rotor paper extends the evidence from a small model peptide to a biomolecular sample directly relevant to structural biology [2]. After the improved OD-first and centerless-ID machining process, one rotor was repacked with approximately 0.5 mg of uniformly  $^{15}\text{N},^{13}\text{C}$ -labeled  $\text{A}\beta_{1-40}$  and spun at  $\omega_r/2\pi = 100$  kHz in an 800 MHz 0.7 mm HCN MAS probe. The reported data include a one-dimensional  $^{13}\text{C}$  CPMAS spectrum, a  $^{13}\text{C}$ -detected two-dimensional  $^{13}\text{C}$ - $^{13}\text{C}$  RFDR spectrum, and a projection from a three-dimensional  $^1\text{H}$ -detected hCCH/RFDR experiment. We found that the spectra were free of natural-abundance  $^{13}\text{C}$  background from the diamond rotor under the CPMAS conditions, because polarization was sourced from  $^1\text{H}$  spins in the sample and the  $^{13}\text{C}$  spins in the diamond lattice relax on a much longer timescale than those in the protein. The same rotor was reported to spin stably for three weeks, allowing multiple experiments to be collected.

Together, these spectra establish the current functional result for diamond rotors. Diamond rotors can be made sufficiently compatible with commercial 0.7 mm MAS hardware to acquire proton-detected and carbon-detected spectra at 100–111 kHz. The rotor material does not automatically create a dominant  $^{13}\text{C}$  background in the reported biomolecular CPMAS experiments, which is essential because natural-abundance carbon in the diamond wall would otherwise be a serious concern. The improved machining process also matters spectroscopically: lower ID/OD eccentricity and stable packing translate into longer usable experiments, not only better metrology.

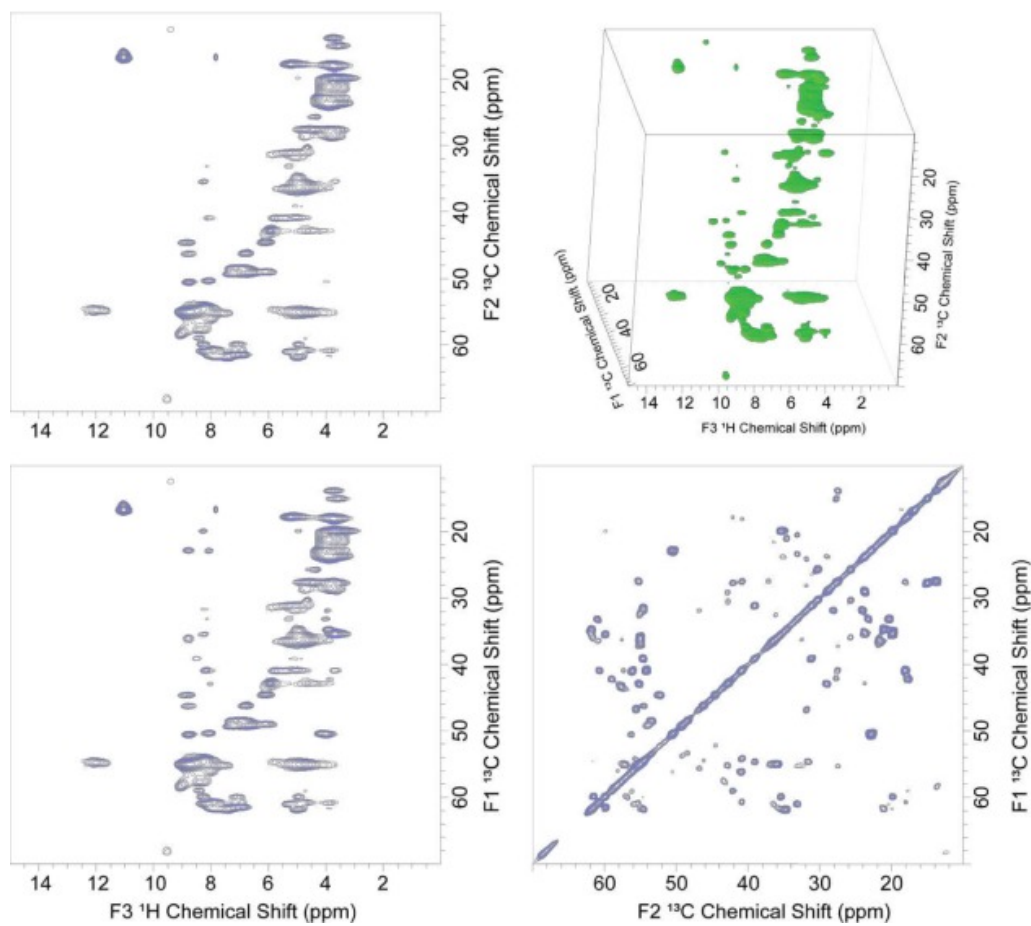


Figure 2.25: Three-dimensional  $^1\text{H}$ -detected RFDR (hCCH) spectrum of  $\text{U-}^{13}\text{C}, ^{15}\text{N-A}\beta_{1-40}$  recorded using an improved diamond rotor spinning at  $\omega_r/2\pi = 100$  kHz at 277 K in an 800 MHz spectrometer [2].

## 2.8 Supporting Architecture

The diamond rotor does not operate alone. It must be compatible with stators, gas bearings, drive tips, end caps, RF coils, packing tools, and the drive gas. At sub-millimeter scale, those supporting parts can become the limiting technology even if the rotor body survives.

### 2.8.1 Stator Material Selection and Manufacturing

The stator is the structural housing of the MAS probe. It integrates bearing interfaces, drive gas nozzles, and RF coil infrastructure while sitting at the magic angle in a strong magnetic field. The material must be non-magnetic, low-loss, thermally stable, and mechanically precise.

Traditional stators use Macor, YSZ, alumina, or related technical ceramics. Macor is machinable because its mica-like microstructure arrests cracks, but it is not the strongest option. YSZ and alumina can offer high strength and hardness but require grinding, ultrasonic machining, or diamond/CBN tooling. Additive manufacturing routes such as SLA ceramic-loaded resins can create complex internal geometry, but debinding and sintering shrinkage can move micron-scale nozzles out of tolerance. Existing printed-stator literature shows that 3D printing is useful for rapid iteration, but high-frequency MAS still requires careful material and tolerance control [59–61].

For the support materials, the useful design question is not the coefficient alone but the actual clearance or press-fit change. The linear thermal expansion coefficient is a normalized strain per temperature interval, so it must be multiplied by the feature size before it becomes an actual dimensional change:

$$\Delta L = \alpha L_0 \Delta T, \quad (2.4)$$

so the table below converts supplier-reported or literature-reported expansion coefficients into the diameter change of a 0.700 mm feature. Positive values in the table are the magnitude of shrinkage during cooling; the same values would be expansion during heating over the same temperature interval. These numbers are first-order estimates because  $\alpha$  is temperature dependent and polymer behavior at cryogenic temperature can be nonlinear. Final cap, stator, and bearing fits still require direct cryogenic fit testing.

Printed stator prototypes remain useful because they allow the bearing pocket, gas manifolds, coil clearance, and assembly interfaces to be checked before committing to ceramic machining. The stator prototype below shows a Formlabs Ceramic Resin part with the major stator features resolved. The result is close enough to demonstrate the feasibility of the geometry, but post-print processing is still required to remove support-contact artifacts, clean the surface, and improve the dimensional accuracy of the bearing and gas-delivery features. Ceramic-loaded SLA materials can capture the bulk stator geometry, but the sintering step introduces shrinkage and distortion that are difficult to tolerate at sub-millimeter MAS scale. For early flow and fit checks, sinter-free glass-filled photopolymer prototypes made from Formlabs Rigid 10K gave decent geometry, which makes them useful for iteration even though they should not be treated as final high-speed bearing materials.

Table 2.5: Estimated dimensional change for a 0.700 mm support or closure feature. The  $\alpha$  column reports the linear thermal expansion coefficient in  $\mu\text{m m}^{-1} \text{K}^{-1}$ ; the next columns translate that normalized material value into actual micrometer-scale diameter change [62–73].

Material	CTE used	$\Delta D$ for 0.700 mm per 100 K	Estimated $\Delta D$ , 293 K to 100 K	Design implication for diamond rotor hardware
Single-crystal diamond	$\sim 1.0$ at 300 K; lower below room temperature	$\sim 0.07 \mu\text{m}$	$< 0.14 \mu\text{m}$	Best thermal match to a diamond rotor body; brittle and orientation-sensitive, so machining damage remains the dominant risk
Vespel SP-1 polyimide	50–54	3.5–3.8 $\mu\text{m}$	6.8–7.3 $\mu\text{m}$	Proven NMR-compatible cap material, but it shrinks tens of times more than diamond and can change cap preload at cryogenic temperature
PCTFE / Kel-F	70–126, supplier dependent	4.9–8.8 $\mu\text{m}$	9.5–17.0 $\mu\text{m}$	Cryogenic-compatible fluoropolymer for seals/fixtures, but the contraction mismatch is too large for an uncontrolled rigid fit to diamond
Torlon 4203 PAI	$\sim 31$	2.2 $\mu\text{m}$	4.2 $\mu\text{m}$	Strong precision-machinable polymer with lower expansion than many plastics; still requires NMR-background and cryogenic-fit validation
UpNano UpPhoto / UpDraft resin	40–80*	2.8–5.6 $\mu\text{m}$	5.4–10.8 $\mu\text{m}$	Useful for printed drive-tip geometry trials; the value is a proxy for related two-photon-polymerized acrylate resins because UpNano does not publish CTE
Formlabs Rigid 10K	41	2.9 $\mu\text{m}$	5.5 $\mu\text{m}$	Sinter-free glass-filled photopolymer for fit checks; not a final high-speed bearing material
Macor glass ceramic	8–9.3	0.56–0.65 $\mu\text{m}$	1.1–1.3 $\mu\text{m}$	Machinable ceramic with moderate mismatch; useful for stator prototypes and non-contact structures
Alumina ceramic	8.2	0.57 $\mu\text{m}$	1.1 $\mu\text{m}$	Strong, low-loss stator candidate; ceramic finishing and feature tolerances are the limiting issues
YSZ / zirconia ceramic	$\sim 10.5$	0.74 $\mu\text{m}$	1.4 $\mu\text{m}$	Mature rotor/stator ceramic baseline, but denser and less favorable than diamond for DNP-oriented rotor bodies

\*Approximate design value only. The UpNano resin data sheet reports mechanical properties for UpPhoto but does not report CTE [67]; the 40–80  $\mu\text{m m}^{-1} \text{K}^{-1}$  range is borrowed from published two-photon-polymerized IP-Dip/acrylate measurements [68, 69] and should be replaced by direct measurement before cryogenic rotor use.




Figure 2.26: Formlabs Ceramic Resin stator prototype for MAS hardware development. The print captures the overall stator envelope and gas-feature layout, while visible support-contact artifacts and surface roughness indicate that post-print machining, finishing, and dimensional verification are required before the geometry can be used as a precision bearing component.

## 2.8.2 Aerostatic Bearing Evolution

MAS rotors are supported by gas films rather than solid contact. Conventional MAS probes use discrete ceramic orifices to create the gas bearing. At 0.7 mm scale, small variations in orifice diameter, inlet angle, or azimuthal position can create asymmetric pressure fields and radial runout. Projection micro-stereolithography can print very small channels, but polymer bearing surfaces may not survive rotor touchdown, and ceramic bearing surfaces require a separate strategy for permeability, finishing, and dimensional control.

BMF projection micro-stereolithography was evaluated because it can resolve the small bearing inner diameter and radial feed passages needed for sub-millimeter rotor support. The printed-bearing photograph below shows an air-bearing component with the surrounding stator prototype geometry, while the CT image shows a commercial Bruker ceramic air-bearing geometry used as a dimensional reference for nitrogen-spun zirconia rotors. The CT section has an ID of approximately 0.746 mm. The bearing uses seven radial feed holes; the outer pore entrance is approximately 0.5 mm, the inner inlet is approximately 0.15 mm, and the transition between them is a conical feature with an included angle near 110 degrees. A circumferential outer groove distributes gas to the feed holes. The bearing disc is approximately 1 mm thick, the groove depth is about 0.4 mm, and the rotor-facing land is narrowed to roughly 0.5 mm by 90 degree chamfers on both sides. These features are not arbitrary: the groove distributes supply gas, the converging feed passages act as restrictors, the chamfers reduce sharp-edge interference with the rotor, and the narrow land defines the gas-film region that controls stiffness, flow rate, and touchdown tolerance.

The project is also evaluating porous-medium air bearings as an alternative to discrete-orifice bearings. A porous restrictor distributes gas through many small interconnected pores, forming a more spatially uniform pressure field at the bearing surface than a small number of jets. Porous-bearing literature explains the attraction of this geometry: many distributed restrictors can improve pressure uniformity, stiffness, damping, and running accuracy compared with discrete orifices, while reducing sensitivity to a single blocked or



Figure 2.27: BMF-printed air-bearing and stator prototype components used to evaluate sub-millimeter bearing geometry and feed-channel resolution.




Figure 2.28: Micro-CT cross-section of a Bruker ceramic air-bearing reference geometry, showing the central ID and radial feed passages used to support a sub-millimeter MAS rotor.

oversized feed hole [74, 75]. The design risk is pneumatic hammer and unstable compressible-film behavior; Yoshimoto et al. showed that a lower-permeability surface-restricted layer can suppress pneumatic hammer while maintaining useful stiffness [76]. This suggests that a sub-millimeter MAS bearing should not simply maximize porosity. It should tune permeability and possibly use a restricted surface layer.

For the present scale, an initial porous-ceramic target is a fine open-pore structure on the order of 2–5  $\mu\text{m}$  with about 12% open porosity, then experimentally adjust permeability after lift-off and spin testing. That target is intentionally lower in porosity than the 20–35% range often cited for larger aerostatic pads working at 5–15  $\mu\text{m}$  design gaps, because the MAS bearing has a very small diameter and must avoid excessive through-flow and weak pressure recovery [74]. Other porous-ceramic studies show that porosity, pore size, and permeability must be designed together: porous SiC membranes with 25.8% porosity and 20.55  $\mu\text{m}$  pore size produced high stiffness in a larger aerostatic bearing, while porous-alumina work emphasizes permeability near  $10^{-14}$ – $10^{-16}$   $\text{m}^2$  and porosity below about 30% for practical ultra-precision bearings [77, 78]. For a diamond MAS rotor, the porous bearing would therefore be evaluated by four coupled outputs rather than by pore size alone: lift-off pressure, recess/source pressure ratio, flow rate and cooling, and rotor runout or spin stability under nitrogen and helium.

### 2.8.3 End Caps and Micro-Impellers

End caps seal the sample and act as micro-impellers that couple the drive gas to rotor motion. Vespel is widely used because it has favorable NMR background and mechanical properties. However, machining a complex 0.7 mm-scale impeller from Vespel is difficult. High-precision CNC platforms such as KERN machines, hydrostatic guideways, and sub- $\mu\text{m}$  encoders can help, but tool geometry limits remain severe.

Two-photon polymerization can print true three-dimensional microstructures with sub- $\mu\text{m}$  feature resolution because polymerization occurs only at the femtosecond laser focal voxel. This makes it attractive for complex drive tips and micro-impellers, as shown by the printed drive-tip prototypes below. The drive-tip printing work was carried out with Nick Wiesner and Samuel Strymis. The UpNano NanoOne platform is well matched to this scale: MIT.nano lists minimum horizontal features of at least 170 nm, minimum vertical features of at least 550 nm, and high writing speed for NanoOne-class printing [79]. A representative starting recipe for UpNano drive-tip prototypes is to print a coarse support body first and reserve a fine print profile for the turbine fins and cap-sealing features. For UpPhoto or UpDraft, the manufacturer-recommended fine-mode windows include 20x printing with approximately 0.3  $\mu\text{m}$  line distance, 0.9–1.0  $\mu\text{m}$  layer slicing, and 300 mm/s infill speed, or 40x printing with approximately 0.15–0.2  $\mu\text{m}$  line distance, 0.45–0.5  $\mu\text{m}$  layer

slicing, and 150 mm/s infill speed. Development can start with two to three IPA baths of at least 10 minutes per bath, with longer development for small internal channels [80]. The exposure window should still be screened locally because fin thickness, resin age, objective choice, and support strategy affect whether the fins survive development and handling.

Preliminary tests on a 1.3 mm rotor assembly showed that the printed drive-tip concept could spin to 60 kHz without issue in that larger rotor format. This result supports continued development of printed micro-impellers for rapid prototyping, but it does not eliminate the material mismatch: polymer caps have different thermal expansion, strength, and DNP compatibility than diamond. A long-term method that is still being explored is to machine the drive tip and cap from a single piece of diamond so the spinning assembly has matched thermal expansion, high stiffness, and a cleaner material interface.

Direct diamond drive tips have also been tested as a matched-material alternative to polymer or Vespel drive components. In a prototype by Zach Fredin, the drive tip was fabricated as a two-part diamond assembly and bonded under a micro-probe station using a thin epoxy layer. Roughing was performed at a feed rate of 0.1 mm/s, and finishing was performed on the inner three contours at 0.0125 mm/s. The process required approximately 12 hours per fin. This prototype shows that diamond drive-tip fabrication is technically plausible, but the method still requires further work on toolpath efficiency, fin-edge quality, adhesive-joint reliability, balance, and survivability under gas-drive loading before it can be treated as a production rotor-closure method.




Figure 2.29: Rotor-closure and gas-drive prototypes. (A) Two-part diamond drive-tip prototype fabricated by Zach Fredin, showing turbine-fin geometry, central post, and laser-machined surface texture; the machining marks and assembled construction motivate improved finishing and single-piece diamond closure fabrication. (B) Micro-printed drive-tip and impeller-style prototypes developed with Nick Wiesner and Samuel Strymis for sub-millimeter rotor closure development; additive microfabrication can generate turbine-like features for rapid iteration, but final use still depends on balance, cap fit, NMR background, thermal expansion, and spin survivability.

#### 2.8.4 Aerostatic-Bearing Modeling

The gas film can be modeled using the isothermal compressible Reynolds equation under the assumption that film thickness  $h$  is much smaller than bearing radius  $R$ . A simplified form is:

$$\nabla \cdot (h^3 p \nabla p) = 6\mu U \frac{\partial(ph)}{\partial x} + 12\mu \frac{\partial(ph)}{\partial t}, \quad (2.5)$$

where  $\mu$  is gas viscosity and  $U = \omega R$  is rotor surface velocity. The bearing compressibility number can be written as:

$$\Lambda = \frac{6\mu\omega R^2}{p_a c^2}, \quad (2.6)$$

where  $p_a$  is ambient pressure and  $c$  is radial clearance [81]. As MAS frequency increases,  $\omega$  drives the bearing into a more compressible regime, and stability becomes sensitive to clearance, nozzle geometry, and gas choice.

### 2.8.5 Thermal-Viscous Dynamics in Helium

Helium is attractive for ultrafast MAS because its speed of sound is much higher than that of air or nitrogen. As calculated in Table 2.2, a 0.7 mm rotor at 150 kHz has a surface speed of approximately 330 m/s, which is near Mach 1 in air but only about Mach 0.33 in helium at room temperature. The practical advantage is not that the rotor becomes stable at the speed of sound. The advantage is that helium keeps the same mechanical surface speed in a lower-Mach-number gas-flow regime, reducing sonic choking, compressibility losses, and acoustic instability in the drive and bearing flow [13]. Helium operation is nevertheless not a simple gas substitution. The gas viscosity, density, compressibility, heat capacity, source pressure, and bearing clearance all change the pressure field and the viscous power dissipated in the gas film. Aerostatic-bearing modeling for MAS hardware shows that bearing stiffness and stability are strongly coupled to clearance, nozzle/restrictor geometry, and compressible flow through the bearing feed [81]. The shear stress in a gas film can be approximated as:

$$\tau = \mu \frac{U}{c} = \mu \frac{\omega R}{c}, \quad (2.7)$$

and the frictional power loss scales approximately as:

$$P_{\text{loss}} = \tau A U = \frac{\mu \omega^2 R^2 A}{c}. \quad (2.8)$$

A preliminary flow model by Nick Wiesner at the MIT Francis Bitter Magnet Laboratory was used to evaluate this clearance problem for nitrogen and helium. Gauge pins from approximately 0.64 to 0.74 mm were used as rotor-diameter proxies inside the bearing. Larger gauge-pin diameter corresponds to smaller radial clearance, while a smaller gauge pin corresponds to a larger rotor-bearing gap. The CFD result in Fig. 2.30A shows the local mechanism: gas accelerates through the bearing restrictor, turns into the thin bearing film, and redistributes around the rotor-support region. High local velocity at the restrictor and in the turning region implies that bearing flow is controlled by both restrictor geometry and the available film clearance.

The pressure-ratio calculation in Fig. 2.30B converts the same clearance problem into a bearing-performance metric. As gauge-pin diameter increases from 0.64 to 0.74 mm, the effective clearance decreases and the recess/source pressure ratio rises. Helium gives a higher recess/source pressure ratio than nitrogen under comparable pressure cases. In practical terms, smaller clearance can increase pressure recovery and bearing stiffness, but it also increases viscous shear, heating, and sensitivity to dimensional error. Larger clearance increases gas through-flow and can reduce shear loss, but if the clearance becomes too large the bearing loses stiffness and the rotor can become less stable. This is the clearance tradeoff that links rotor OD, bearing ID, gas selection, and spin efficiency.

Because heating scales with  $\omega^2$ , pushing to very high speed increases thermal load sharply. For a fixed bearing ID near 745  $\mu\text{m}$ , a 700  $\mu\text{m}$  rotor gives a radial clearance of approximately 22.5  $\mu\text{m}$ , while a 690  $\mu\text{m}$  rotor gives approximately 27.5  $\mu\text{m}$ . That 5  $\mu\text{m}$  increase in radial clearance is small geometrically but large aerodynamically: from the simplified shear-loss expression above, the factor  $R^2/c$  is reduced by about 20% when the rotor OD changes from 700 to 690  $\mu\text{m}$  at the same spin frequency. The compressibility number,

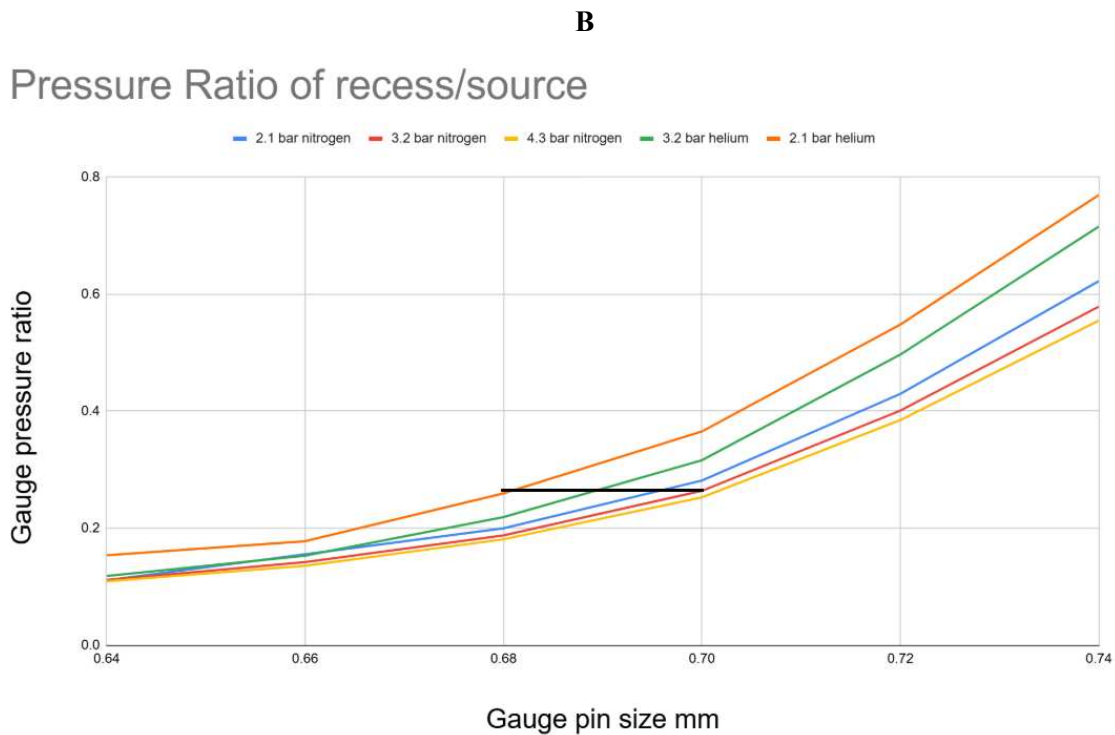
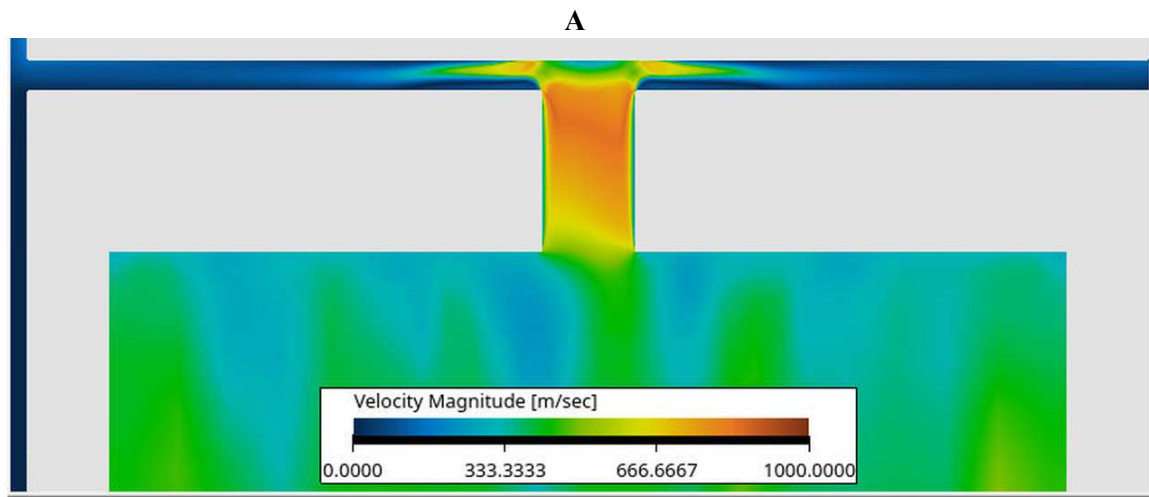


Figure 2.30: Helium-bearing flow simulations from preliminary modeling by Nick Wiesner at the MIT Francis Bitter Magnet Laboratory. (A) CFD velocity field for gas injection into the sub-millimeter rotor-bearing clearance, showing gas acceleration through the restrictor and turning into the rotor-bearing film. (B) Recess/source pressure ratio as a function of gauge-pin size for nitrogen and helium bearing-flow cases; gauge-pin size is used as a rotor-diameter proxy in a fixed bearing geometry, so increasing gauge-pin size reduces clearance and changes pressure recovery. Together, the panels show that rotor-bearing gap, restrictor geometry, and gas choice must be optimized as a coupled aerodynamic system.

which scales approximately as  $R^2/c^2$ , is reduced by about 35%. Machining helium rotors closer to 690  $\mu\text{m}$  is therefore a rational design adjustment when the goal is to increase bearing flow and reduce viscous gas-film losses. The cost is reduced pressure recovery and potentially lower bearing stiffness, so the final clearance has to be chosen from the coupled gas-flow, thermal, and spin-stability behavior rather than from rotor diameter alone [13, 81].

## Chapter 3

# Precision Motion and the Fabrication Ecosystem

### 3.1 Why Accessible Sub- $\mu\text{m}$ Motion Matters

Precision motion is the mechanical layer that turns a toolpath into a manufactured feature and a scan command into a registered measurement. In this thesis it links three classes of linear capability: accessible piezoelectric fine motion, long-travel high-speed servo motion for laser fabrication, and synchronized motion/process control for machining and characterization. These are related, but they are not interchangeable. A piezo stage solves the problem of inexpensive local sub- $\mu\text{m}$  or nanometer-scale positioning over a short range. A ball-screw or direct-drive linear stage solves the problem of moving a payload over millimeters to hundreds of millimeters while maintaining geometry, speed, and process synchronization.

Cheap, accessible sub- $\mu\text{m}$  motion is important because many laboratory instruments fail at the same bottleneck: the optical, mechanical, or material process is capable in principle, but the lab cannot position the workpiece or sensor repeatably enough to exploit it. Commercial nanopositioning systems are powerful, but their cost and closed architecture can make early-stage instrument development slow, fragile, and difficult to modify. A locally built piezo stage changes that experimental structure. It gives the lab a fine-motion module that can be repaired, instrumented, adapted, and treated as part of the experiment rather than as a sealed purchased component. That does not remove the need for calibration. It makes calibration more meaningful because the mechanism, electronics, and controller are visible.

For laser micromachining and optical metrology, linear axis error becomes feature-placement error. Angular error becomes Abbe-amplified tool-center error. Servo following error becomes pulse-placement error. Fixture compliance becomes a lost datum after re-clamping. For Raman imaging, confocal microscopy, and autocorrelation measurements, the same motion stack determines whether spectra, images, and time-delay measurements can be registered to the physical part. The precision-motion layer is crucial part of the fabrication method.

### 3.2 Motion Architectures Used in the Fabrication Ecosystem

The chapter separates the motion systems by linear-motion role and physical drive mechanism. Short-stroke piezo flexures use a piezo stack or bender against elastic guidance, giving nanometer-to-micrometer motion, high stiffness, high small-signal bandwidth, no backlash, and no sliding guide wear; their limits are stroke, high-voltage electronics, hysteresis, creep, preload, flexure stress, and resonance. Long-travel stepping piezo stages use stick-slip, inchworm, walking, or ultrasonic frictional accumulation to move over millimeters to centimeters, but their motion is stepwise and usually slower than servo long-travel stages; friction, surface

condition, preload, wear, humidity, payload, and velocity ripple become part of the motion budget [82, 83]. Ball-screw servo stages convert rotary motor motion into long linear travel through a screw, nut, and rolling guide system. They are mature, stiff, high-payload, and cost effective, but lead error, backlash or reversal behavior, friction, screw windup, thermal growth, and transmission compliance separate command increment from tool-point accuracy [84, 85]. Direct-drive mechanical-bearing stages remove the screw and use a linear motor, linear encoder, and crossed-roller or precision guide bearings; they are better for smooth scanning and high acceleration, but motor heat, magnetic forces, cable drag, bearing cage behavior, encoder alignment, and servo tuning become central [86, 87]. Direct-drive air-bearing stages remove rolling contact entirely and can give excellent velocity stability and geometric averaging, but they require clean dry air, filtration, pressure interlocks, a flat support structure, contamination control, and crash protection [88].

A motion axis should be specified at the coordinate where the experiment happens. Catalog resolution, motor encoder counts, and controller interpolation describe only parts of the chain. The manufactured feature responds to the total tool-center-point error:

$$\mathbf{e}_{tcp}(t) = \mathbf{e}_{lin}(t) + \mathbf{e}_{servo}(t) + \mathbf{e}_{geo}(t) + \mathbf{e}_{thermal}(t) + \mathbf{e}_{fixture}(t) + \mathbf{e}_{process}(t), \quad (3.1)$$

where  $\mathbf{e}_{lin}$  includes calibrated linear position error,  $\mathbf{e}_{servo}$  includes following error during motion,  $\mathbf{e}_{geo}$  includes straightness and angular errors,  $\mathbf{e}_{thermal}$  includes stage and scale drift,  $\mathbf{e}_{fixture}$  includes clamping and compliance, and  $\mathbf{e}_{process}$  includes forces, debris, ablation recoil, plasma shielding, vibration, and measurement registration effects. ISO 230-2 provides formal definitions and testing procedures for positioning accuracy and repeatability of numerically controlled axes [89]. Nanopositioning and nanomeasuring reviews make the same broader point: useful precision depends on actuator, guide mechanism, sensor, structure, controller, and environment rather than on command increment alone [90, 91].

Abbe error shows why the point of measurement matters. If an encoder measures an axis at an offset  $L$  from the tool point, angular error  $\theta$  produces an additional linear error:

$$e_A \approx L\theta. \quad (3.2)$$

This is why a high-resolution encoder can coexist with micron-scale tool error if the workpoint is far above the scale and the carriage pitches or yaws during travel.

### 3.3 Piezoelectric Stages for Accessible Fine Motion

#### 3.3.1 Operating Principle and Design Envelope

Piezoelectric nanopositioners exploit the inverse piezoelectric effect: an applied electric field produces strain in a piezoelectric material. In the linear strain-charge form, the constitutive relation can be written as [92]

$$\mathbf{S} = \mathbf{s}^E \mathbf{T} + \mathbf{d}^t \mathbf{E}, \quad (3.3)$$

$$\mathbf{D} = \mathbf{d} \mathbf{T} + \epsilon^T \mathbf{E}, \quad (3.4)$$

where  $\mathbf{S}$  is strain,  $\mathbf{T}$  is stress,  $\mathbf{E}$  is electric field,  $\mathbf{D}$  is electric displacement,  $\mathbf{s}^E$  is elastic compliance at constant electric field,  $\mathbf{d}$  is the piezoelectric coefficient matrix, and  $\epsilon^T$  is dielectric permittivity at constant stress. In practical PZT stack actuators, the longitudinal coefficient  $d_{33}$  is used for axial extension and the transverse coefficient  $d_{31}$  is used for bending or lateral deformation. A simplified unloaded multilayer stack estimate is

$$\Delta L_{free} \approx Nd_{33}V, \quad (3.5)$$

where  $N$  is the number of active layers and  $V$  is the voltage per layer. The blocked-force estimate follows from the actuator stiffness  $k_p$ :

$$F_{block} \approx k_p \Delta L_{free}. \quad (3.6)$$

Piezo stacks can be stiff, compact, fast, and electrically controllable, but their raw strain is small. A useful laboratory stage must therefore choose one of four strategies: accept a short fine range, amplify the motion elastically, accumulate steps through friction or clamping, or combine a coarse stage with a fine piezo stage.




Figure 3.1: Multilayer piezo stack concept from the laboratory slide deck. Alternating ceramic and electrode layers convert applied electric field into small axial strain; many active layers increase usable displacement at a practical drive voltage.

### 3.3.2 Review of Piezo Stage Types

The piezo-stage review separates actuator geometry from stage architecture. Stack, tube, and bending actuators describe how the piezo element deforms. Flexure-amplified stages describe how a platform is guided and amplified. Stepping stages describe how many small piezo motions accumulate into long travel. These are not interchangeable design choices. A stack actuator uses the axial  $d_{33}$  response of many active layers, giving high stiffness, high force, fast small-signal response, and compact packaging, but only short raw stroke, preload sensitivity, hysteresis, and creep [92, 93]. Tube scanners use bending and extension of a cylindrical piezo with segmented electrodes; they are compact multi-axis scanners for SPM-style instruments, but long tubes bring cross-axis coupling, creep, and low structural modes [93, 94]. Flexure-amplified stages use elastic hinges and levers to transform the piezo input into platform motion; they are backlash-free, clean, and wear-free in the guide, but amplification trades stroke against output stiffness, parasitic motion, stress concentration, and resonance [94, 95]. Stick-slip, inchworm, and ultrasonic stages obtain longer travel by frictional accumulation or clamping sequences. They can be compact and low cost, but their step size, velocity ripple, wear, preload dependence, humidity sensitivity, and surface condition must be treated as part of the motion system [82, 83].

Short-range solid-state stages use stack, bending, or tube actuators directly. Stack actuators are strong and stiff, but they must normally be preloaded because piezo ceramics tolerate compression better than tension.

Tube actuators are compact multi-axis scanners, but their cross-axis coupling and low structural modes can limit accuracy. Bending actuators produce larger tip motion than stacks, but at lower stiffness and force.

Flexure-amplified stages place the piezo actuator against a moving platform constrained by elastic hinges. Motion occurs through elastic deformation rather than sliding contact, so flexures can eliminate backlash, rolling-element friction, lubricant contamination, and wear in the guide [93, 94]. The tradeoff is stress, fatigue, parasitic motion, and the onset of low-frequency structural modes. An amplified mechanism should be evaluated by its displacement gain,

$$A = \frac{\Delta x_{out}}{\Delta x_{in}}, \quad (3.7)$$

and by the stiffness penalty that comes with amplification. A useful scaling rule from energy conservation is

$$k_{out} \approx \frac{k_{in}}{A^2}. \quad (3.8)$$

Stepping piezo stages accumulate many small motions to obtain long travel. In stick–slip motion, an asymmetric waveform drives a slow stick phase and a fast slip phase. During the slow phase, static friction carries the slider; during the rapid return, inertia prevents the slider from following the actuator completely. In inchworm motion, clamps and an extender execute a hold-extend-transfer-retract cycle. Ultrasonic piezo motors use resonant vibration to generate elliptical or traveling-wave contact motion. These methods make long travel possible, but friction, preload, surface condition, humidity, wear, and payload become part of the motion budget [82].

### 3.3.3 Hysteresis, Creep, and Position Sensing

Piezoelectric stages must be treated as nonlinear mechatronic systems. Hysteresis means that displacement depends on voltage history. Creep means that position drifts after a nominally fixed command. Structural modes limit usable bandwidth. Sensor noise and drift limit closed-loop stability. Nanopositioning control surveys therefore treat modeling, sensing, and feedback as central parts of the stage rather than optional refinements [93, 94].

A common phenomenological model for rate-independent hysteresis is the Bouc–Wen form. One simplified actuator representation is

$$x(t) = k_v V(t) - h(t), \quad (3.9)$$

$$\dot{h} = \alpha \dot{V} - \beta |\dot{V}| |h|^{n-1} h - \gamma \dot{V} |h|^n, \quad (3.10)$$

where  $V$  is drive voltage,  $x$  is displacement,  $h$  is an internal hysteretic state, and  $\alpha$ ,  $\beta$ ,  $\gamma$ , and  $n$  are fitted parameters. Bouc–Wen variants are useful because they can be identified and inverted for compensation, but rate dependence and asymmetry often require extended models or feedback [96, 97]. If a capacitive sensor, interferometer, optical encoder, strain gauge, or magnetic sensor is used, its own noise, linearity, bandwidth, and thermal drift must be reported [91].

## 3.4 In-House Low-Cost Piezo Stage

### 3.4.1 Design Intent

The in-house piezo stage is the access-oriented fine-motion contribution of this chapter. This build was led by Bowen Zeng, with help from Max Aalto and Jake Read. Its purpose is not to replace every commercial nanopositioner. Its purpose is to create a locally controlled sub- $\mu\text{m}$  positioning module that can

be built, repaired, modified, instrumented, and understood inside the laboratory. The low-cost open-source XYZ nanopositioner reported by Liao et al. shows why this direction is technically credible: a magnet-preloaded stick–slip architecture achieved centimeter-scale horizontal travel, millimeter-scale vertical travel, micrometer-scale high-resolution scanning ranges, and a 12 kg horizontal payload demonstration under the published test conditions [83]. The assembled in-house stage is characterized separately because preload, surface condition, electronics, sensors, and assembly tolerances can change the measured result.




Figure 3.2: Front CAD view of the in-house low-cost piezo micro-motion stage. The stage was designed by Bowen Zeng, with work from Max Aalto and Jake Read, and was inspired by the low-cost open-source XYZ nanopositioner architecture reported by Liao et al. [83]. The assembly combines piezo stack actuation, miniature ball-slide guidance, magnetic encoder boards, magnetic scale tape, and a compact fixture stack.


### 3.4.2 Mechanism, Flexure, and Material Considerations

The mechanical design of an in-house stage is evaluated by travel, resolution, stiffness, preload stability, cross-axis coupling, moving mass, resonance, and ease of fabrication. Flexure stages and stick–slip stages solve different problems. Flexures are best when backlash-free short-range motion and clean guidance dominate. Stick–slip mechanisms are attractive when long travel and low component cost matter more than perfectly deterministic open-loop step size.


The useful flexure variables are direct and unforgiving. A thinner hinge neck increases compliance and travel, but consumes stress margin, fatigue life, and off-axis stiffness; it needs FEA and a measured stiffness check. Longer flexure beams lower actuation force and increase travel, but reduce resonant frequency and enlarge parasitic motion, so modal response and frequency response matter as much as static stroke. A higher amplification ratio increases output motion for a given piezo stroke, but reduces output stiffness

approximately with the square of the ratio and lowers payload bandwidth. Parallel-kinematic layouts improve axis decoupling and platform stiffness, but they can become overconstrained if machining and assembly tolerances are not controlled. Low moving mass improves settling and bandwidth only if the fixture remains stiff. The position sensor should be placed as close as practical to the workpoint because Abbe error and compliance between sensor and payload cannot be corrected by a finer command increment.

Flexure material choice is a mechanical design decision rather than a convenience. Flexures require high elastic strain capacity, predictable modulus, fatigue resistance, machinability, dimensional stability, and compatibility with the process environment. Material properties are reported by exact alloy, heat treatment, and stock condition because generic values for steel, stainless steel, aluminum, or titanium can vary by heat treatment, cold work, grain direction, and surface finish.



(a) Flexure detail in CAD.



(b) Waterjet-cut flexure trial.

Figure 3.3: Flexure-stage fabrication evidence from the laboratory slide deck. The waterjet trial was unexpectedly clean for the geometry: the flexure used approximately 1 mm internal radii and 0.5 mm wall sections through 1/4 in (6.35 mm) stock, dimensions that would normally push abrasive-waterjet cutting toward a coupon-validation regime.

The waterjet result in Figure 3.3 was a useful surprise for the laboratory. The flexure looked visually closer to a wire-EDM trial than expected, even though it was produced on a much faster and easier OMAX 5-axis abrasive-waterjet workflow. This matters because flexure prototypes need many geometry iterations. A waterjet can rough or even trial-cut compliant mechanisms quickly without wire threading, submerged dielectric setup, or the conductivity requirement of wire EDM. The caveat is that this trial is near the edge of what should be assumed without testing: OMAX lists fine abrasive-waterjet kerf widths of 0.015 in (0.38 mm) for the 7/15 MAXJET5 nozzle and 0.020 in (0.5 mm) for the Mini-Jet nozzle, while taper, jet lag, material thickness, speed, pierce strategy, and arc geometry all change the final edge shape [98, 99]. A 0.5 mm web in 6.35 mm-thick stock is therefore only about one fine-waterjet kerf wide and has an aspect ratio of roughly 13:1 through the plate. The correct lesson is not that waterjet replaces wire EDM for final fatigue-critical flexures, but that waterjet may be good enough for fast design exploration if the hinge roots, taper, striation marks, abrasive embedding, and wall width are inspected after cutting.

The comparison with wire EDM is still important. The lab also has access to a Sodick wire EDM running 0.010 in wire, which is a 0.254 mm electrode diameter. A Sodick AP450L-class fine-wire machine is specified for 0.002–0.010 in wire diameters, linear motor drives, glass-scale feedback, and high-precision micromachining [100]. For a 0.010 in wire, the geometric lower bound on an inside radius is set by the wire radius plus spark overcut, and practical design rules normally leave extra radius and clearance for flushing, wire vibration, skim cuts, and material thickness. Even with those allowances, the characteristic wire-EDM corner scale is much smaller than a waterjet kerf. Wire EDM is therefore the better final process for small

conductive flexure hinges when surface integrity and dimensional control dominate; the OMAX waterjet is the faster screening process when the question is whether a geometry, stiffness, and assembly concept is worth carrying forward.

The material decision and the fabrication decision should be kept separate. A material table asks what elastic, thermal, fatigue, corrosion, and interface behavior the part should have. A fabrication table asks how the selected material and geometry can actually be cut, inspected, and finished. Mixing these questions hides important tradeoffs: for example, aluminum may be attractive for the stage body even if the piezo contact hardware should be steel, while waterjet may be useful for rapid geometry trials even when wire EDM remains the better final process for fatigue-critical conductive flexures.

For compliant precision parts, 4140 or a similar alloy steel gives high stiffness and high strength after controlled heat treatment, but needs corrosion planning and exact documentation of alloy, heat treatment, yield strength, surface finish, and fatigue assumptions. 316 stainless is laboratory friendly and corrosion resistant, but annealed strength, work hardening, and magnetic response must be checked before it is treated as a precision flexure material. Aluminum is excellent for the stage body, prototypes, datum plates, and low-cost machining, but it has lower modulus, lower fatigue margin, higher thermal expansion, and soft contact surfaces; piezo contact and preload interfaces should therefore use steel inserts or steel hardware. Titanium offers high specific strength and corrosion resistance, but its lower modulus, galling risk, and machining difficulty only make sense if mass or corrosion requirements justify the fabrication penalty.

For fabrication, wire EDM, abrasive waterjet, and CNC finishing occupy different places in the workflow. Wire EDM is the clean final route for intricate conductive flexures because the 0.010 in (0.254 mm) wire scale is far smaller than a fine waterjet kerf, but it adds setup time, wire threading, conductivity requirements, recast-layer risk, microcrack risk, and surface-finish validation [100–102]. Abrasive waterjet is fast, cold, and material-flexible; in this project it produced the surprising 1 mm-radius, 0.5 mm-wall, 6.35 mm-thick flexure trial, but the final part still needs kerf, taper, striation, abrasive-embedding, and hinge-root inspection [98, 99, 103, 104]. CNC milling or grinding is best kept for body plates, datum faces, threaded interfaces, preload blocks, burr removal, and post-process cleanup where tool access and surface finish can be controlled.

### 3.4.3 Bill of Materials

The in-house XY piezo stage was designed around low-cost, commercially available motion components, compact sensing hardware, and a small embedded controller. The bill of materials in Table 3.1 records the costed implementation used to evaluate the accessibility of the design, including a separate fabrication allowance for printed or machined parts.

Table 3.1: Bill of materials for the in-house XY piezo stage.

Component function	Specification	Qty.	Cost
Piezoelectric actuators	Thorlabs PC4WM co-fired piezo actuators, 150 V, 9.5 $\mu$ m nominal travel	2	\$158.36
Mechanical guides	THK LS852 / Misumi miniature linear ball slides	2	\$198.12
Piezo driver amplifier	PDu150 three-channel, 150 V output, low-noise piezo driver	1	\$401.00
Position sensing	Magneto 4 Click AS5311 Hall-effect encoder modules	2	\$72.00

Component function	Specification	Qty.	Cost
Metrology scale	Multipole magnetic strip or tape compatible with the AS5311 pole-pitch requirement	1	\$8.99
Controller logic	ESP32/RP2040-class microcontroller board	1	\$48.40
Signal generation	FG-100 DDS function generator, 1 Hz–500 kHz operating range	1	\$29.22
Interface and fasteners	4N25 optocouplers, transistors, M2/M2.5 screws, small hardware	Lot	\$41.55
Stage body fabrication	Printed prototype body or machined aluminum body; steel actuator-contact and preload hardware at the piezo interface	Lot	\$100–\$300 allowance
Prototype total	XY prototype component total including fabrication allowance	–	\$1,057.64– \$1,257.64

The mechanical body is deliberately treated as part of the bill of materials because it controls preload alignment, contact stiffness, and repairability. A printed body is acceptable for rapid layout checks and early assembly trials, but a machined aluminum body is the better thesis-build direction because it improves stiffness, thread quality, dimensional stability, and thermal/mechanical robustness. The parts that directly contact or preload the piezo actuator should be steel rather than printed polymer or soft aluminum, so the ceramic stack is loaded through hard, repeatable, wear-resistant interfaces.

### 3.4.4 Motion-Control Architecture and Measured Repeatability

The assembled in-house piezo stage has already demonstrated repeatability below 1  $\mu\text{m}$  in laboratory testing. This is a meaningful result for the thesis because it shows that the low-cost architecture has crossed from a mechanical prototype into a usable sub- $\mu\text{m}$  positioning platform. The claim should be stated as repeatability rather than absolute accuracy: repeatability describes how closely the stage returns to the same measured position under a specified command sequence, while accuracy also requires an independently calibrated position reference, a coordinate error map, and uncertainty analysis. The final archived test record should therefore include axis, payload, command waveform, approach direction, number of cycles, target spacing, encoder mode, filtering, environmental conditions, and the definition of the reported repeatability statistic.

The peer-reviewed open-source XYZ nanopositioner reported by Liao et al. provides the closest control-architecture reference for the local stick–slip concept [83]. Its motion control is deliberately simple. A piezo stack drives a magnetically preloaded slider on a miniature linear slide. A triangular command waveform uses the piezo stack as a short-stroke scanner and gives a small continuous fine-motion range. A sawtooth command waveform creates the stick–slip sequence: the slow ramp moves the slider through static friction, while the fast return causes slip because slider inertia prevents full reverse motion. Liao et al. report approximately 3.7  $\mu\text{m}$  high-resolution scanning range, 12 mm horizontal stepping travel, 6 mm vertical stepping travel, and up to 12 kg horizontal payload under their test conditions [83]. Their low-cost stepping controller uses an Arduino Mega, a DAC board, and voltage/current amplification. The Arduino generates 8-bit digital signals for the DAC board, the DAC produces sawtooth waveforms, the documented default stepping frequency is 5 Hz, and a TDA2050 audio power amplifier drives the piezo stack up to a reported 35 V and 5 A peak-current envelope for the low-cost controller configuration [83]. The important thesis takeaway is that centimeter-scale travel does not require a complicated servo architecture if the mechanism is allowed to

move in frictional increments, but the resulting step size and return behavior must be characterized because friction, preload, surface condition, and waveform shape become part of the motion loop.

The antonysigma piezo-stage PID-board project is useful as a second reference because it addresses the next layer: closed-loop piezo-flexure focus control rather than open-loop stick-slip stepping [105]. The project keeps the high-voltage piezo amplifier as a separate analog power stage, while a custom board generates a conditioned low-voltage command and reads a differential quadrature encoder. Its documented target is a 0  $\mu\text{m}$ –250  $\mu\text{m}$  Z-axis travel range, 5  $\mu\text{m}$  command steps, and 50 ms closed-loop settling for focus scanning. Several design choices transfer directly to the local stage: the DAC output should be filtered before the piezo amplifier to suppress zero-order-hold steps and amplifier excitation; the encoder signal should be converted and decoded close to the controller; the measured position should be filtered before feedback; and large setpoint changes should pass through a slew-rate limiter to prevent integral windup and mechanical ringing.

A practical firmware-level motion loop for the local stage can therefore be expressed as

$$r_s[k] = r_s[k - 1] + \text{clip}(r[k] - r_s[k - 1], -v_{\max}T_s, v_{\max}T_s), \quad (3.11)$$

$$e[k] = r_s[k] - \hat{y}[k], \quad (3.12)$$

$$u[k] = \text{sat}(C(z)e[k]), \quad (3.13)$$

where  $r[k]$  is the requested position,  $r_s[k]$  is the slew-limited setpoint,  $v_{\max}$  is the allowed command velocity,  $T_s$  is the controller sample time,  $\hat{y}[k]$  is the filtered encoder position,  $e[k]$  is following error,  $C(z)$  is the digital controller, and  $u[k]$  is the bounded low-voltage command sent to the piezo driver. The PID-board project reports mapping a continuous controller into the z-domain with a bilinear transform and, after physical tuning, using a discrete form

$$C(z) = G \frac{z^{-1} + 10z^{-2} - 20z^{-3}}{1 - z^{-1}}, \quad (3.14)$$

where  $G$  is the proportional gain adjusted on the physical system [105]. This exact transfer function should not be copied blindly into the thesis stage because resonance, encoder noise, amplifier bandwidth, friction, and payload differ. Its value is architectural: it shows that a low-cost controller can combine setpoint shaping, encoder feedback, digital compensation, and analog signal conditioning to make a piezo stage settle repeatably instead of merely applying voltage.

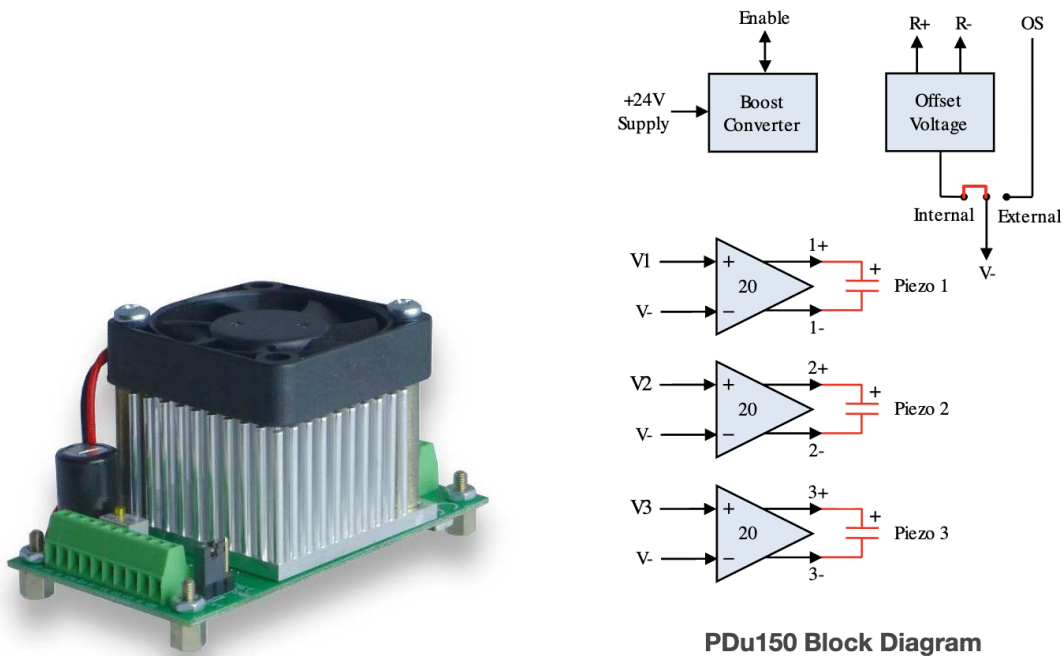
For the in-house XY stage, the immediate control objective is therefore not only smaller commanded step size. It is reproducible return-to-position behavior. The sub-1  $\mu\text{m}$  repeatability data show that the current hardware is already in the useful regime for microscope alignment, fixture nudging, and local process adjustment. The next control milestone is to pair this result with an explicit error histogram, approach-direction comparison, drift trace, and frequency-response record so that repeatability, hysteresis, creep, and bandwidth are not conflated.

### 3.4.5 Key Commercial Components

The prototype uses four key purchased subsystems: a high-voltage piezo driver, multilayer piezo stack actuators, magnetic encoder modules with a multipole scale, and miniature linear ball slides. These components are sized by voltage, capacitance, bandwidth, sensor resolution, and guide behavior rather than by nominal displacement alone. Each component solves a specific accessibility problem, but each also introduces limits that must be measured on the assembled stage.

## Piezo Driver

The PDu150 driver, shown in Figure 3.4, is a useful purchased amplifier because it avoids the need to design a custom high-voltage, low-noise analog output stage during the first prototype cycle. It provides three channels with a  $-30\text{ V}$  to  $150\text{ V}$  output range, up to  $100\text{ mA}$  peak current per channel,  $80\text{ kHz}$  power bandwidth at  $150\text{ V}$  peak-to-peak, and  $180\text{ kHz}$  signal bandwidth [106]. The block diagram is important because it shows the actual motion chain: a low-voltage command input is amplified, an internal boost converter produces the high-voltage rail from a  $24\text{ V}$  supply, and an offset-voltage stage defines the unipolar or bipolar operating range. For a laboratory prototype, the main advantages are compact packaging, three synchronized channels, reduced high-voltage design risk, and compatibility with simple function-generator or DAQ commands.



(a) PDu150 driver module.

(b) Driver internal block diagram.

Figure 3.4: PiezoDrive PDu150 driver used to command the prototype piezo stacks. The module turns low-voltage command signals into high-voltage piezo drive outputs while providing the boost conversion and offset-voltage structure needed for practical bench operation [106].

For a capacitive piezo load, the amplifier current required for a sinusoidal command is approximately

$$I_{pk} \approx 2\pi f C V_{pk}, \quad (3.15)$$


where  $f$  is command frequency,  $C$  is actuator capacitance, and  $V_{pk}$  is voltage amplitude. This equation is why high-voltage range alone is not sufficient: fast motion requires enough current to charge and discharge the piezo capacitance.

## Piezo Stack Actuator

The selected Thorlabs PC4WM actuator, shown in Figure 3.5, is a compact co-fired multilayer piezo stack. The advantage of this actuator class is that it gives useful force, high stiffness, and a simple two-wire drive

interface in a small package, while still being inexpensive enough for a low-cost stage. The PC4WM should be treated as a 150 V, 9.5  $\mu\text{m}$  free-stroke actuator, with Thorlabs listing 9.5  $\mu\text{m} \pm 15\%$  displacement at 150 V, 3.5 mm by 4.5 mm by 10.0 mm stack dimensions, 115 kHz  $\pm 10\%$  no-load resonant frequency, recommended preload below 125 N, 180 nF  $\pm 15\%$  capacitance, and 250 N blocking force at 150 V [107]. The inner working is the same multilayer mechanism introduced in Figure 3.1: many thin active ceramic layers produce axial strain at a practical drive voltage. The main design cost is short stroke, hysteresis, creep, preload sensitivity, and the need to avoid tensile stress in the ceramic.

## Co-Fired Piezo Actuator, 150 V, 9.5 $\mu\text{m}$ Travel



(a) Thorlabs PC4WM co-fired stack actuator.


PC4WM <sup>a</sup>	
Drive Voltage Range	0 to 150 V
Displacement (Free Stroke) at 150 V <sup>b</sup>	9.5 $\mu\text{m} \pm 15\%$
Hysteresis	$\leq 15\%$ (See Graph on Next Page)
Load for Maximum Displacement <sup>c</sup>	100 N (22.5 lbs)
Recommended Preload	<100 N (22.5 lbs)
Blocking Force at 150 V	250 N (55 lbs)
Resonant Frequency	115 kHz $\pm 10\%$ (No Load)
Impedance at Resonant Frequency	700 m $\Omega$
Anti-Resonant Frequency	165 kHz $\pm 10\%$ (No Load)
Dissipation Factor <sup>d</sup>	<2.0%
Capacitance <sup>e</sup>	180.0 nF $\pm 15\%$
Operating Temperature	-25 to 110 °C
Curie Temperature	230 °C
Dimensions	Width 1: 3.5 mm + 0/-0.5 mm Width 2: 4.5 mm + 0/-0.5 mm Length: 10.0 mm $\pm 5 \mu\text{m}$

(b) PC4WM specification excerpt used in the lab slide deck.


Figure 3.5: Thorlabs PC4WM piezo stack actuator selected for the low-cost stage. The actuator gives compact fine motion, but the usable stage performance depends on preload, amplifier current, guide friction, flexure stiffness, and closed-loop sensing [107].

## Magnetic Encoder and Scale

The AS5311 magnetic encoder module, shown in Figure 3.6, is the low-cost position-sensing element in the prototype. The sensor reads the field pattern from a multipole magnetic strip or ring, so the measurement is contactless and mechanically simple to package near the moving stage. This is a major advantage over a contact probe or potentiometric sensor because it adds little friction and can be mounted around the compact slide structure. The limitation is that the sensor and scale are not an interferometer. The datasheet describes a 1.0 mm pole-length magnetic strip or ring, a typical 0.3 mm sensing gap, 488 nm absolute step size over a 2.0 mm pole pair, and 1.95  $\mu\text{m}$  incremental output step size [108]. Sub- $\mu\text{m}$  claims for the assembled stage therefore require direct calibration of the selected encoder mode, magnetic scale, analog/digital interface, interpolation, mounting gap, and noise.



(a) AS5311 encoder board.



(b) Magnetic scale working principle.


Figure 3.6: AS5311 magnetic encoder and multipole-scale concept used for compact feedback in the in-house piezo stage. The scale converts linear travel into a repeating magnetic field pattern, while the encoder reads the field without adding sliding contact to the motion axis [108].

### Miniature Linear Slide

The THK LS852 linear ball slide, shown in Figure 3.7, supplies the compact guideway for the prototype. Its advantage is pragmatic: it provides a purchased, miniature, low-friction linear guide so the first prototype can focus on actuator, sensing, preload, and control rather than on manufacturing a perfect monolithic guide. The LS852 size has a nominal 25 mm stroke, 8 mm height, 14.2 mm width, 53.7 mm base length, and published dynamic/static load ratings of 68.6 N/118 N [109]. The cost is that the guide introduces rolling-contact friction, guide straightness, preload, and cage behavior into what would otherwise be a purely solid-state piezo actuator problem.



(a) THK LS852 slide.



(b) Slide dimensional diagram.

Figure 3.7: THK LS852 miniature linear ball slide used as the prototype guideway. This component makes the stage mechanically accessible and easy to assemble, but the final motion behavior must include rolling-element friction, straightness error, and cage behavior [109].

### 3.4.6 Characterization Data Products

The stage is evaluated by measured behavior rather than by cost alone. A low-cost stage does not need to outperform commercial nanopositioners to be useful; it needs to make a defined experimental capability accessible and characterize its limits honestly. The core data products are step-size distribution, scan linearity, hysteresis, creep, drift, payload sensitivity, cross-axis coupling, bidirectional repeatability, and frequency response. The axis calibration map records position error versus command so repeatable scale error can become compensation data. The reversal map records approach-direction error, which predicts bidirectional raster mismatch. Straightness, pitch, yaw, roll, and Abbe offsets translate carriage geometry into tool-center error. The frequency-response record stores resonant peaks, notch filters, bandwidth, and settling behavior so tuning can be repeated instead of guessed. A thermal drift log records temperature, encoder-scale behavior, and position drift over the job time. The final validation artifact should be a measured cut, scan, or alignment task that connects stage metrics to a physical outcome.

## 3.5 Long-Travel High-Speed Linear Motion for Laser Systems

### 3.5.1 Why Long-Travel High-Speed Stages Are a Separate Class

Long-travel high-speed linear stages solve a different problem from the piezo stage. A piezo stack can have high intrinsic bandwidth because its active element is solid state and its moving mass can be small, but the useful stroke is normally in the  $\mu\text{m}$  range. A stepping piezo can move over millimeters or centimeters, but it achieves that travel by frictional or clamping cycles, so the motion is discrete, preload-dependent, and slower than a servo axis. A laser machining stage must instead move a payload over tens to hundreds of millimeters, preserve geometry, maintain smooth contouring velocity, coordinate  $Z$  focus, and trigger process events at the correct spatial locations. The design concerns are bearing geometry, encoder placement, cable forces, thermal drift, servo bandwidth, controller configuration, safety limits, and the way following error maps into feature error.

### 3.5.2 Ball-Screw Stages and the Oxford Laser System

A ball screw converts rotary motion into linear motion through a screw, nut, and recirculating balls. The ideal kinematic relation is

$$x = \frac{p}{2\pi}\theta, \quad (3.16)$$

where  $p$  is screw lead and  $\theta$  is motor angle. With an encoder of  $N$  counts per revolution and interpolation factor  $R$ , the ideal command increment is

$$\Delta x_{cmd} = \frac{p}{NR}. \quad (3.17)$$

This command increment is not the same as axis accuracy. Lead error, backlash, preload, nut stiffness, guide stiffness, motor coupling, friction, thermal expansion, and torsional windup all appear between the motor and the carriage. A minimal two-inertia feed-drive model captures the compliance:

$$\delta = r\theta - x, \quad r = \frac{p}{2\pi}, \quad (3.18)$$

$$F_s = k_s\delta + c_s\dot{\delta}, \quad (3.19)$$

$$J_m\ddot{\theta} = T_m - rF_s, \quad (3.20)$$

$$m_l\ddot{x} = F_s - F_f - F_d. \quad (3.21)$$

Here  $J_m$  is motor inertia,  $m_l$  is moving load mass,  $T_m$  is motor torque,  $k_s$  and  $c_s$  are transmission stiffness and damping,  $F_f$  is friction, and  $F_d$  is disturbance force. This model explains why a motor encoder can look settled while the load is still deflected or oscillating. Reviews of ball-screw feed drives emphasize that recirculating rolling elements, guideway joints, preload, and thermal effects require both static and dynamic modeling for high-precision applications [85].

In the Oxford laser micro-machining system, the PRO190SL ball-screw stage is paired with Aerotech Automation 3200 control and an Ndrive drive. Its role is robust, contamination-tolerant, long-travel positioning for nanosecond laser machining. Aerotech's current PRO190SL data sheet lists travel options from 100 mm to 600 mm, standard accuracy from  $\pm 6 \mu\text{m}$  to  $\pm 15 \mu\text{m}$  over that range, calibrated accuracy from  $\pm 1 \mu\text{m}$  to  $\pm 2 \mu\text{m}$ ,  $\pm 1 \mu\text{m}$  bidirectional repeatability, 300 mm/s maximum speed, 60 kg horizontal load capacity, 50 kg vertical/axial load capacity, and a brushless rotary servomotor drive system with rotary encoder feedback [110]. The installed Oxford stage should be documented by its exact travel, encoder, motor, foldback, and calibration options because these options change the achievable tool-center performance.

The PRO115SL is a smaller member of the same ball-screw/mechanical-bearing design family and is a useful compact reference for any in-house screw-driven rail stage. Aerotech describes the PRO115SL as a precision-ground ballscrew stage with a long-life recirculating linear guide bearing system, sealing against particulates, a brushless rotary servomotor, and rotary encoder feedback options of a 1000-line  $1 V_{pp}$  encoder or a 2500-line digital RS422 encoder [111]. The current data sheet lists 50 mm–400 mm standard travel options, 300 mm/s maximum speed, calibrated accuracy from  $\pm 1 \mu\text{m}$  to  $\pm 3 \mu\text{m}$  over that range,  $\pm 1 \mu\text{m}$  bidirectional repeatability, 40 kg horizontal load capacity, and ideal minimum incremental motion of  $0.1 \mu\text{m}$  with the amplified sine encoder or  $1.0 \mu\text{m}$  with the digital encoder [111]. This is an important distinction: the sine encoder increases the electrical subdivision of the rotary motor measurement, but the ball screw, nut preload, guide rail, and load-side compliance still set the carriage-level accuracy.

The important design lesson is that ball-screw stages trade cost, payload, and contamination robustness against transmission physics. If building a comparable axis in-house, the screw lead, nut preload, bearing block stiffness, coupling stiffness, motor inertia, load inertia, guideway straightness, encoder location, thermal path, and way-cover friction must be specified together. A high-resolution motor encoder does not eliminate screw lead error or load-side deflection; the performance claim must be made at the carriage or tool point.

### 3.5.3 Direct-Drive Mechanical-Bearing Stages and Compact Auxiliary Axes

Direct-drive linear motors remove the screw and convert current directly into force:

$$F = K_f i, \quad (3.22)$$

where  $K_f$  is force constant and  $i$  is phase-current-equivalent command. A simplified load equation is

$$m\ddot{x} + c\dot{x} + kx = K_f i - F_d. \quad (3.23)$$

Because the transmission is removed, backlash and screw windup disappear. A linear encoder can be placed near the payload, reducing measurement offset. This makes linear motors attractive for high-speed contouring, scanning, and synchronized laser processing. The tradeoffs are heat, magnetic attraction forces in iron-core designs, cable forces, cost, and servo tuning sensitivity [84, 86]. The current laboratory scope in this chapter focuses on screw-driven and crossed-roller/direct-drive systems.

For long-travel high-speed scanning, an ALS130/ALS130H-class direct-drive linear stage is a useful reference architecture. Aerotech's ALS130H data sheet describes a mechanical-bearing, direct-drive linear motor stage with a high-accuracy noncontact glass-scale encoder, anti-cage-creep crossed-roller bearings,

0.5 nm resolution with Aerotech controls, bidirectional repeatability to 75 nm, standard straightness and flatness below  $\pm 2 \mu\text{m}$  over the travel, 300 mm/s maximum velocity, and 1g peak no-load acceleration [87]. It also notes that neither the bearing system nor the drive system uses recirculating elements, which is why this architecture is attractive for low-ripple scanning. The design lesson is not that the catalog resolution should be treated as the machining accuracy. It is that a direct-drive mechanical-bearing stage moves the dominant error sources away from screw reversal and transmission compliance and toward encoder quality, bearing straightness, motor heat, cable force, payload placement, and servo tuning.

Compact screw-driven axes are still useful for vertical, focus, or auxiliary motion when high payload stiffness, covers, holding behavior, and packaging matter more than continuous high-speed contouring. The ATS100 illustrates this tradeoff. It is a screw-driven mechanical-bearing stage protected by integral bellows; Aerotech lists travel lengths up to 200 mm, sub- $\mu\text{m}$  positioning accuracy, bidirectional repeatability to 0.3  $\mu\text{m}$ , a precision-ground ball screw, linear motion guide bearings, brushless rotary motors with integral encoders, and vacuum-compatible options [112]. In a build, this type of axis should be evaluated by focus-coordinate repeatability, vertical stiffness, brake or holding behavior, contamination protection, cable routing, and the Abbe offset between the axis measurement line and the optical work point.

The ANT95L stage is used in the autocorrelator build, where the motion requirement is not cutting force or payload but stable, repeatable optical delay. The current ANT95L data sheet describes a compact direct-drive linear motor stage with high-precision crossed-roller bearings, 25 mm, 50 mm, and 100 mm travel classes,  $\pm 75 \text{ nm}$  bidirectional repeatability, PLUS accuracy down to  $\pm 175 \text{ nm}$  on shorter-travel high-performance configurations, 500 mm/s maximum speed for the  $1 V_{pp}$  encoder configurations, and in-position stability below 0.5 nm for the high-performance encoder option [113]. Its feedback options are also unusually instructive: the standard amplified-sine encoder option is specified as  $1 V_{pp}$  with a 20  $\mu\text{m}$  signal period, while the high-performance amplified-sine option is  $1 V_{pp}$  with a 4  $\mu\text{m}$  signal period [113]. The exact installed encoder option should therefore be recorded before claiming 0.5 nm, 1 nm, or 5 nm command increments.

The laboratory does not currently have an air-bearing linear stage in this chapter, but the ABL1000 is worth keeping as a design benchmark because it isolates what changes when rolling contact is removed. Aerotech's ABL1000 family uses a fully preloaded noncontact air bearing, noncontact linear motor, and noncontact feedback device; the data sheet lists 0.5 nm resolution with the higher-performance encoder option,  $\pm 50 \text{ nm}$  repeatability, 300 mm/s maximum speed, 1g no-load acceleration, and straightness/flatness values down to  $\pm 0.25 \mu\text{m}$  on shorter travel configurations [88]. The same data sheet also shows the practical cost of this architecture: the air supply must be clean and dry, filtered to 0.25  $\mu\text{m}$  or better, and protected by pressure interlocks. This is the right comparison for bearing noise and velocity stability, but it is not a current in-house axis.

### 3.5.4 Aerotech Stage Specifications as Design Targets

The commercial systems are not copied as black boxes. They are useful because their published specifications reveal the design variables an in-house long-travel stage would have to satisfy: travel, speed, acceleration, payload, encoder type, guide type, straightness, angular error, thermal stability, and calibration procedure. Figure 3.8 anchors three physical reference architectures: an ALS130H direct-drive crossed-roller stage, a PRO115SL protected ball-screw stage, and an ANT95L compact nanopositioning direct-drive stage. The ABL1000 air-bearing family remains a comparison benchmark only; it is not a current in-house axis.




Figure 3.8: Commercial Aerotech reference stages used to anchor the long-travel stage comparison. The images show the physical distinction between a compact direct-drive crossed-roller stage, a protected ball-screw mechanical-bearing stage, and a compact nanopositioning direct-drive stage [87, 111, 113].

The dry design targets are the useful part. PRO115SL-class screw stages show the compact ball-screw envelope: 50 mm–400 mm travel, 300 mm/s maximum speed,  $\pm 1 \mu\text{m}$  bidirectional repeatability, calibrated accuracy from  $\pm 1 \mu\text{m}$  to  $\pm 3 \mu\text{m}$ , 40 kg horizontal load capacity, and rotary servomotor feedback options of 1000-line  $1 V_{pp}$  sine or 2500-line digital encoder [111]. ALS130H-class stages show the direct-drive mechanical-bearing target: 25 mm–150 mm travel classes, 0.5 nm controller resolution, 75 nm bidirectional repeatability, 300 mm/s velocity, 1g no-load acceleration, crossed-roller bearings, and noncontact glass-scale feedback [87]. AT5100-style auxiliary axes show why compact screw stages still matter for Z or focus motion: up to 200 mm travel, sub- $\mu\text{m}$  positioning, 0.3  $\mu\text{m}$  bidirectional repeatability, precision-ground screw drive, linear guides, bellows, and brushless rotary motor options [112]. ANT95L-class optical-delay axes show the nanopositioning target: 25 mm–100 mm travel classes,  $\pm 75 \text{ nm}$  bidirectional repeatability, 500 mm/s speed for  $1 V_{pp}$  encoder configurations, and 20  $\mu\text{m}$  or 4  $\mu\text{m}$  encoder signal-period options [113]. ABL1000 air-bearing specifications show what improves when rolling contact is removed: noncontact bearing, motor, and feedback; 0.5 nm resolution with the high-performance encoder option;  $\pm 50 \text{ nm}$  repeatability; and 300 mm/s speed, but only with clean dry air, 0.25  $\mu\text{m}$  filtration, pressure interlocks, and crash protection [88]. The design lesson is that an in-house long-travel axis must specify the whole chain, not only travel: motor force constant, moving mass, screw or linear-motor selection, guideway class, encoder signal period, interpolation, cable force, thermal path, safety limits, tuning files, and calibration method.

### 3.5.5 Encoder, Bearing, and Servo Architecture Inside Commercial Stages

The exact encoder readhead model used inside a commercial Aerotech stage is not consistently published in the public stage data sheets. The supportable statement for the thesis is therefore narrower than saying “Aerotech uses a specific Renishaw model.” The PRO115SL documentation identifies rotary motor encoder options, the ANT95L data sheet identifies 20  $\mu\text{m}$  and 4  $\mu\text{m}$   $1 V_{pp}$  linear encoder signal periods, and the ALS130H data sheet identifies a high-accuracy noncontact glass-scale linear encoder, but the public ALS130H and ANT95L documents do not name the installed Renishaw readhead part number [87, 111, 113]. The correct final-thesis procedure is to record the installed stage option string, readhead label, Aerotech order sheet, or service documentation before naming a part number. Renishaw documentation is still useful here because it explains the class of encoder architecture that makes these stage specifications possible: a micron-pitch scale, an optical readhead, sine/cosine signal generation, interpolation, signal conditioning, and a high-speed digital feedback interface [114–116].

The important point is that nanometer feedback does not require a scale physically marked every nanometer. A linear optical scale can have a 20  $\mu\text{m}$  or 4  $\mu\text{m}$  signal period, while the readhead and drive electronics estimate the phase within each period. In an ideal quadrature model,

$$V_s(x) = G_s \sin\left(\frac{2\pi x}{p}\right) + o_s + n_s, \quad (3.24)$$

$$V_c(x) = G_c \cos\left(\frac{2\pi x}{p} + \phi_e\right) + o_c + n_c, \quad (3.25)$$

$$\hat{x} = p\left(N + \frac{1}{2\pi} \text{atan2}(V_s, V_c)\right), \quad (3.26)$$

$$\Delta x_{\text{ideal}} = \frac{p}{M}. \quad (3.27)$$

Here  $p$  is signal period,  $N$  is the integer fringe count,  $M$  is interpolation factor,  $G_s$  and  $G_c$  are channel gains,  $o_s$  and  $o_c$  are offsets,  $n_s$  and  $n_c$  are noise terms, and  $\phi_e$  represents quadrature phase error. Renishaw describes open incremental encoders that filter the scale image into sine and cosine signals, use dynamic signal conditioning to control gain, balance, and offset, and interpolate those signals for fine position feedback [114]. For example, Renishaw’s TONiC family is specified as a compact noncontact incremental encoder system with speeds up to 10 m/s, resolutions down to 1 nm with the Ti interface, 20  $\mu\text{m}$  scale options, very low jitter, and low sub-divisional error [115]. Aerotech’s XC4 drive gives the matching controller-side picture: square-wave encoder feedback can be accepted at up to 40 million counts/s, and optional sine-wave encoder multiplication can multiply sine encoders up to 16,384 times [117]. If a 20  $\mu\text{m}$  signal period is multiplied by 16,384, the ideal count size is about 1.22 nm; if a 4  $\mu\text{m}$  signal period is multiplied by 16,384, it is about 0.244 nm. These are ideal digital subdivisions, not guaranteed absolute accuracy. Scale calibration, interpolation nonlinearity, readhead alignment, analog noise, thermal expansion, Abbe error, and servo dynamics still determine the tool-point result.

The bearing system is the second half of the same story. A cheap linear rail can move accurately enough for many laboratory tools, but high-end scanning stages are built around guide errors that are small and repeatable enough for the encoder and servo to matter. Recirculating guide bearings introduce ball-passage vibration, friction ripple, preload sensitivity, and noise from ball return paths. NSK’s high-accuracy HA/HS linear-guide literature explicitly connects ball-passage vibration to slide posture changes, table straightness, surface-finish defects, and high-precision table motion error; NSK reports that its high-accuracy series reduces this vibration to one-third of conventional models in its test comparison [118]. SKF’s LLT profile rail documentation gives the complementary design variables: four ball rows at a 45° contact angle, sealed carriages, integrated lubrication, and multiple preload and accuracy classes to trade stiffness, precision, friction, and assembly tolerance [119]. The public Aerotech documents do not identify the exact rail manufacturer or rail grade used in PRO115SL/PRO190SL stages, so the useful thesis claim is the design criterion: a self-built precision mechanical-bearing axis should specify rail accuracy class, preload class, straightness, mounting flatness, lubrication, sealing, carriage spacing, and measured ball-passage or bearing-noise signatures instead of simply buying a “linear rail.”

Crossed-roller direct-drive stages address a different part of the bearing-noise problem. The ALS130H uses anti-cage-creep crossed-roller bearings and explicitly notes that the stage avoids recirculating elements in both the bearing and drive systems, which supports smoother low-speed scanning than a recirculating screw/guide architecture [87]. The ANT95L uses crossed-roller bearings in a smaller nanopositioning package [113]. Air-bearing stages push this further by removing mechanical rolling contact entirely, which reduces wear and rolling-element noise, but only by adding clean gas supply, filtration, pressure monitoring, and very careful crash and contamination control [88]. The clean design comparison is therefore: ball-screw stages are robust and economical; crossed-roller direct-drive stages remove screw transmission error and reduce ripple; air-bearing direct-drive stages remove rolling contact but add pneumatic infrastructure and higher integration risk.

Servo architecture closes the loop around the mechanics and encoder. A practical precision axis uses a

fast current loop in the drive, velocity and position feedback around the stage, and feedforward terms from the commanded trajectory. A compact representation is

$$e(t) = r(t) - y(t), \quad (3.28)$$

$$u_c(t) = C(s)e(t), \quad (3.29)$$

$$u_{ff}(t) = K_v \dot{r}(t) + K_a \ddot{r}(t) + K_j \dddot{r}(t) + K_s r^{(4)}(t), \quad (3.30)$$

$$u_{total}(t) = u_c(t) + u_{ff}(t). \quad (3.31)$$

The feedback controller rejects disturbances and stabilizes the plant; the feedforward terms reduce tracking error during predictable moves without increasing feedback gain. Aerotech's Automation1 tuning documentation makes the same division operationally: measure the frequency response, tune servo gains and filters for stability margins, then compare plant and inverse-feedforward responses to adjust feedforward gain and timing [120]. Its FeedforwardAdvance parameter explicitly advances velocity, acceleration, jerk, and snap feedforward contributions in time to compensate plant latency, and its feedforward filter setup enables second-order filters in the feedforward path [121, 122]. This matters for laser stages because high acceleration, corner blending, or nonsmooth trajectories can excite mechanical resonances even when static positioning looks excellent.

Bode plots are the experimental language for this tuning. If  $P(j\omega)$  is the plant from current command to measured position and  $C(j\omega)$  is the feedback controller, the open-loop response is

$$L(j\omega) = C(j\omega)P(j\omega), \quad (3.32)$$

$$S(j\omega) = \frac{1}{1 + L(j\omega)}, \quad (3.33)$$

$$T(j\omega) = \frac{L(j\omega)}{1 + L(j\omega)}. \quad (3.34)$$

The magnitude plot shows resonances, anti-resonances, loop crossover, and disturbance rejection. The phase plot shows how much stability margin remains before the loop becomes oscillatory. Aerotech's frequency-response module directly exposes open-loop, plant, sensitivity, closed-loop, feedforward, and related responses; it also recommends using phase margin, gain margin, and sensitivity peak as stability metrics, with notch filtering used only when a measured resonance limits the loop [120]. For thesis validation, the Bode plot should not be treated as a decorative control figure. It should be paired with following-error logs and a fabricated or measured artifact, because the real question is whether the tuned axis produces better spatial pulse placement, smoother scanning, lower delay jitter, or more repeatable feature geometry.

The practical motion-system tuning discussion in this section also benefited substantially from guidance from Michael Selepec at Aerotech. His help was especially important in connecting the controller, drive, encoder, frequency-response, and following-error information into one coupled motion system rather than treating them as isolated specifications.

### 3.5.6 Drive and Controller Architecture: Advantages and Drawbacks

The stage mechanics only become a machine when they are paired with an amplifier, feedback interface, real-time controller, safety wiring, process I/O, and a tuning workflow. The controller must do more than send independent axis commands. It generates coordinated trajectories, manages interpolation and electronic gearing, applies compensation, supervises limits and faults, records data, and synchronizes process events such as laser triggers. Aerotech's A3200 documentation describes this class of system as a software-based

machine controller that coordinates up to 32 axes, supports servo, stepper, voice-coil, galvo, piezo, and other drive types, and provides interpolation, high-speed I/O, camming, IFOV, and position-synchronized output functions [123]. For an in-house system, the relevant lesson is that the controller architecture is part of the motion design, not a replaceable user interface.

In a PC-based controller, the industrial PC becomes the machine controller rather than only the operator terminal. The user sees Windows, the programming environment, data files, CAD/CAM tools, diagnostics, and network access. The time-critical part is handled by a real-time layer that runs beside Windows and owns deterministic scheduling, interrupt handling, motion updates, and drive communication. Aerotech's A3200 product literature identifies INtime as the real-time operating system used by the controller, and Aerotech's PC configuration guide emphasizes that PC selection, BIOS settings, power management, device drivers, and Windows configuration affect real-time control performance [123, 124]. TenAsys describes INtime as a real-time operating environment that runs real-time applications alongside Windows while providing deterministic execution for time-critical tasks [125]. The practical meaning is straightforward: Windows may host the interface and file workflow, but the machine must reserve a deterministic execution path for servo timing and I/O.

This architecture has important advantages. It reduces the number of separate controller boxes, allows the motion program, process recipe, HMI, logging, and CAD/CAM workflow to live on one machine, and makes it easier to coordinate many axes with a common trajectory model. It also supports rich diagnostics: following error, current command, encoder position, limit state, temperature, and process triggers can be recorded against the same job timeline. For laser processing, PC-based control is especially useful because toolpath generation, PSO setup, hatch strategy, skywriting transitions, and IFOV deployment can be tied directly to the same controller environment rather than passed through a loose stack of separate tools [126, 127]. This is why the software architecture matters when designing a motion platform: it determines whether the machine can close the loop between geometry, trajectory, trigger timing, and measured outcome.

The drawbacks are just as real. A PC-based controller makes the PC a metrology-critical and safety-critical component. CPU power states, chipset behavior, BIOS settings, Windows updates, graphics or network drivers, thermal throttling, background services, and unsupported hardware can all degrade determinism or make faults difficult to reproduce. The system is also more vendor-ecosystem dependent: drive firmware, controller software version, fieldbus hardware, licenses, encoder interfaces, STO wiring, and API compatibility become part of the machine configuration. Troubleshooting becomes layered because a following-error fault may originate in mechanics, encoder alignment, drive tuning, current limits, real-time scheduling, communication, or process-side I/O. A self-built stage must therefore treat the controller PC as controlled hardware: document the image, BIOS, driver versions, controller version, drive firmware, real-time configuration, and tuning files with the same care as the stage drawing.

The drive still determines the electrical and timing envelope. The Oxford system's Ndrive CP class of A3200 network drive is specified with a 20 V–340 V DC motor bus, 10 A, 20 A, or 30 A peak-current options, corresponding 5 A, 10 A, and 10 A continuous-output classes, 20 kHz PWM switching, 12.5 MHz PSO output frequency, and 160 ns single-axis PSO firing latency [128]. Newer compact Aerotech drive families show the same design pressures in different packages. The XC2 is a compact PWM drive with a 100 V DC bus, up to 10 A peak current, 20 kHz current and servo loops, data collection up to 320 kHz, drive-array memory, and optional multi-axis PSO and I/O expansion [129]. The XC4 extends the single-axis envelope with bus options up to 340 V DC, 10 A, 20 A, or 30 A peak-current versions, 20 kHz current and servo loop updates, 40 million counts/s square-wave encoder input capability, optional sine-encoder interpolation, and PSO capability [117]. These numbers matter because they constrain motor force, acceleration, high-speed encoder reading, trigger timing, and data capture. Selecting a drive by axis name alone is not enough; the drive must be sized from force constant, moving mass, desired acceleration, back-EMF, bus voltage, thermal foldback, encoder rate, safety I/O, and process-trigger latency.

The final disadvantage is conceptual: real-time software does not make a weak mechanical axis precise.

A high-performance controller can expose problems rather than solve them. Cable drag still adds position-dependent force. A poorly mounted scale still produces interpolation error. A direct-drive axis still heats the payload path. A screw stage still has reversal and compliance. The purpose of the control architecture is to make these effects measurable and tunable; it cannot erase the need for stiffness, alignment, thermal management, and independent calibration.

### 3.5.7 Assembly and Tuning Considerations

Precision stage installation begins with mechanics, not software. The mounting surface must support the stage without twist; the payload must be placed near the intended load path; cable carriers must not inject lateral force; limit switches and homes must be unambiguous; and the metrology loop must measure the coordinate that matters. Thermal expansion and air currents should be treated as part of the machine, especially when scans or cuts last long enough for drift to accumulate.

The mechanical checklist is compact but strict. Base flatness and bolt preload control rail distortion, carriage pitch/yaw, bearing preload variation, and position-dependent friction; bolts should be torqued progressively and straightness should be mapped after installation. Payload placement controls bandwidth and Abbe error, so the moving mass should sit near the carriage load center and the workpoint height should be recorded. Encoder scale and readhead alignment control cyclic error, interpolation noise, false velocity ripple, and in-position stability; gap, yaw, pitch, signal quality, and slow bidirectional motion should be verified before aggressive servo tuning. Cable routing can dominate a good stage if service loops inject side load, heat, or drag, so full-travel following error must be checked with the real cable carrier. Thermal paths should be logged through warmup, because motor heat, scale growth, and optical stack drift change the coordinate frame. Home and limit logic must be tested as a recovery procedure, not only as an installation step. Contamination control must protect guides, scales, optics, and process fixtures before routine cutting or scanning.

Control tuning usually follows a cascaded structure: current loop, velocity loop, and position loop. The current loop determines force production, the velocity loop shapes damping and dynamic response, and the position loop determines tracking and settling. Bode plots expose resonance, anti-resonance, phase margin, and gain margin. If a lightly damped mode with natural frequency  $\omega_n$  limits the loop gain, a common continuous-time notch filter is

$$N(s) = \frac{s^2 + 2\zeta_z\omega_n s + \omega_n^2}{s^2 + 2\zeta_p\omega_n s + \omega_n^2}, \quad \zeta_z < \zeta_p. \quad (3.35)$$

Notch tuning should be validated against measured following error, settling time, and feature quality rather than only against a cleaner Bode plot [130].

### 3.5.8 Motion Error and Fabrication Outcomes

The largest feature errors often come from the full motion stack: stage geometry, encoder location, controller tuning, payload, fixture stack, Z focus, and process synchronization. Linear positioning error becomes feature placement error, raster offset, or misplaced alignment datums, and it should be checked with interferometer or encoder calibration, feature metrology, and ISO-style repeated-position tests [89]. Straightness, pitch, yaw, and roll create taper, nonparallel walls, focus shift across travel, and Abbe-amplified tool-center error; they require straightness and angular maps, not only a single repeatability number. Backlash or reversal error appears as cusps at direction changes, bidirectional raster mismatch, and circular-path lobing. Dynamic following error produces pulse placement error, contour error, and roughness changes during acceleration, so controller logs must be compared against test cuts at realistic feed rates. Fixture compliance and re-clamping error move the datum after loading, flipping, or tightening; they need load/unload repeatability studies or

coordinate registration artifacts. Thermal drift changes focus, scale length, and feature size over long jobs. Sensor noise and quantization become closed-loop jitter or unstable scan registration, so in-position stability and noise spectra should be reported with the fabricated result [91].

### 3.5.9 Position-Synchronized Output

Laser micromachining couples motion to energy deposition. If pulses are fired by a clock while the stage accelerates or slows, the spatial pulse pitch changes:

$$\Delta s_{time}(t) = \frac{v(t)}{f_{rep}}, \quad (3.36)$$

where  $v(t)$  is instantaneous path speed and  $f_{rep}$  is laser repetition rate. During deceleration,  $\Delta s_{time}$  shrinks, increasing pulse overlap and local thermal load. During acceleration, the pitch grows, reducing local dose. Position-synchronized output instead triggers a pulse when measured distance reaches the next spatial interval:

$$\|\mathbf{y}(t_k) - \mathbf{y}(t_{k-1})\| \geq \Delta s_{PSO}. \quad (3.37)$$

Aerotech's documentation describes PSO as a feature that controls drive output in response to real-time position feedback from one or more axes, including fixed-distance firing, array-based masking, windowing, and laser/process triggering modes [127, 131]. PSO is therefore the bridge between motion control and process physics: it converts encoder distance into laser timing.

## 3.6 Chapter Synthesis

This chapter began with the practical problem that accessible sub- $\mu\text{m}$  motion changes what people can build and debug. The in-house piezo stage is the local answer to that problem: it is not presented as a replacement for commercial nanopositioners, but as an inspectable fine-motion module whose actuator, guide, encoder, controller, material choices, and fabrication route can be modified inside the lab. The measured sub-1  $\mu\text{m}$  repeatability claim is therefore important because it shows that the design has crossed from a component collection into a usable positioning tool, while still requiring explicit calibration of repeatability, drift, hysteresis, and bandwidth.

The long-travel stage discussion completes the same argument at a larger scale. Ball-screw, crossed-roller direct-drive, and air-bearing architectures all solve different versions of the motion problem, and their specifications only become meaningful when tied to encoder location, bearing noise, servo tuning, thermal drift, process synchronization, and fabricated feature quality. Whether the system is low-cost piezo motion or a commercial linear stage, understanding the manufacturing of the components and their tuning is what makes the platform useful. Encoder interpolation, Bode-plot tuning, and position-synchronized output all serve the same purpose: making the machine's motion error visible enough that it can be improved, calibrated, and used to produce reliable results that were not feasible before.

## Chapter 4

# Laser Micromachining Systems

### Chapter Summary

This chapter reviews state-of-the-art laser micromachining systems and records the two laser systems being developed in house. It is not a rewrite of the diamond rotor fabrication process. Diamond remains an important motivating material because it places strict demands on taper, heat damage, debris removal, sidewall quality, and metrology, but the organizing question of the chapter is broader: how advanced femtosecond and waterjet-guided laser micromachining platforms are designed, characterized, monitored, and evaluated.

The chapter is organized around two system classes. The first is a femtosecond micromachining platform built around an industrial ultrafast source, custom beam delivery, precision motion, pulse diagnostics, and process-development tools. The second is a waterjet-guided laser (WJGL) platform built in house from a pulsed laser source, hydraulic subsystem, coupling optics, water/nozzle head, and staged instrumentation. In both cases, the research contribution is not simply access to a laser. The contribution is the conversion of sources, optics, mechanics, fluidics, and diagnostics into repeatable micromachining systems.

## 4.1 Femtosecond Laser Micromachining Systems

### 4.1.1 Introduction to Femtosecond Laser

Femtosecond laser micromachining includes drilling, cutting, grooving, surface texturing, internal modification, and direct writing at length scales where the beam, material response, motion system, and diagnostic stack must be designed together. Compared with nanosecond processing, femtosecond pulses deposit energy before substantial lattice heating and bulk thermal diffusion. This can reduce heat-affected zones and improve edge quality, but it does not remove every damage mechanism. Plasma shielding, incubation, nonlinear propagation, focus walk-off, redeposition, and low ablation efficiency can still limit the process window [132–134].

For this reason, a useful femtosecond-platform review compares more than pulse duration. The relevant questions are how optical energy couples into the material, how the beam evolves at the workpiece, how pulses are synchronized to scanner or stage motion, how heat and debris are managed, how delivered pulse energy and duration are verified, and how final quality is measured by microscopy, roughness, phase/chemical characterization, and repeatability. Claims about ultrafast interaction physics are supported by peer-reviewed literature; claims about the in-house system are supported by hardware records, optical layouts, alignment data, and measured beam or feature data.

### 4.1.2 State of the Art and Cold-Ablation Physics

Ultrafast laser processing matured from laboratory demonstrations of clean ablation into an industrial manufacturing technology as compact chirped-pulse amplifiers, Yb-based sources, harmonic conversion, beam scanners, high-precision stages, and process control improved. Sugioka and Cheng describe ultrafast lasers as practical tools for surface micromachining, micro/nanostructuring, nanoablation, and three-dimensional internal processing in transparent materials [133]. Malinauskas et al. extend this picture from science to industry, emphasizing fabrication from below  $1\ \mu\text{m}$  to  $\mu\text{m}$ -scale dimensions, three-dimensional structuring in transparent materials, and throughput as the field moves from laboratory processing toward manufacturing [135].

The common description of femtosecond machining as “cold ablation” is useful but incomplete. Femtosecond excitation can confine energy deposition and reduce thermal diffusion during the pulse, yet the final material state still depends on fluence, repetition rate, overlap, burst timing, focusing geometry, and post-pulse relaxation. Figure 4.1 summarizes this contrast: longer pulses allow electron-to-lattice heat transfer during irradiation, producing a larger heat-affected region and debris, while femtosecond pulses confine the early energy deposition and can leave a cleaner ablation region. Lin and Hong also compare nanosecond and femtosecond holes in a  $100\ \mu\text{m}$  steel foil at  $780\ \text{nm}$  [136]. Modern platforms therefore use pulse duration, wavelength, repetition rate, burst mode, beam shaping, spatiotemporal focusing, adaptive optics, and closed-loop monitoring as coupled degrees of freedom [137, 138].

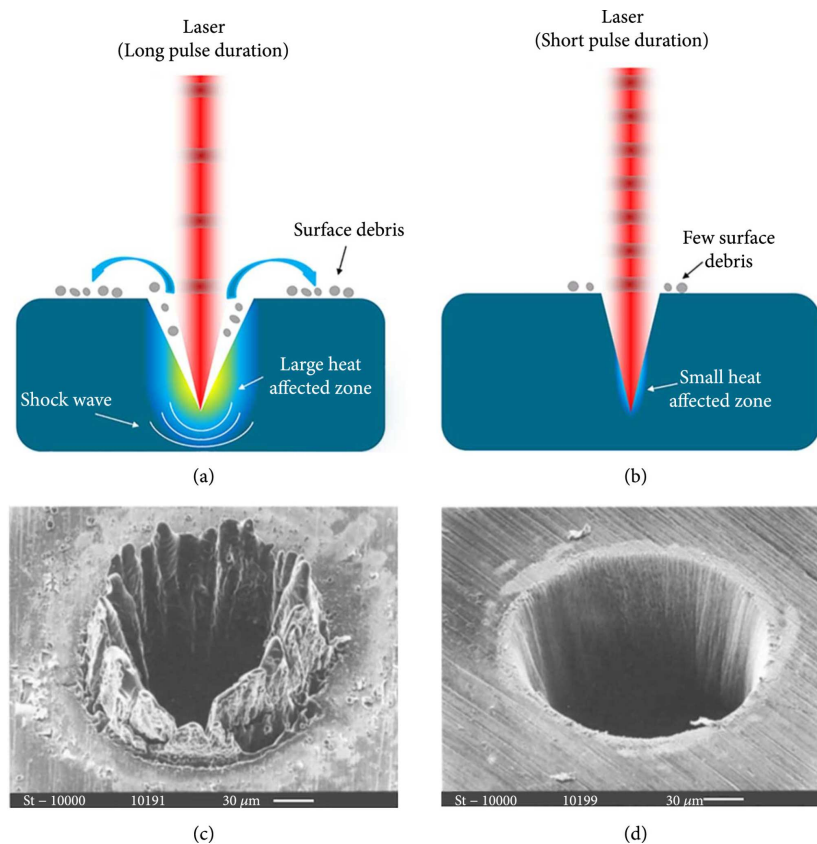


Figure 4.1: Cold-ablation comparison from Lin and Hong, illustrating long-pulse thermal damage and femtosecond-pulse localized removal [136].

For transparent and wide-bandgap materials, femtosecond pulses can drive nonlinear absorption and free-carrier generation even when the unmodified material has weak linear absorption at the laser wavelength.

For metals and semiconductors, ultrafast energy deposition is often modeled with two-temperature or related nonequilibrium frameworks. For dielectrics and transparent materials, photoionization, avalanche ionization, transient Drude response, incubation, and morphology evolution become central. No single model is sufficient across all materials and regimes. A practical micromachining model is hierarchical: optical modeling for the beam, carrier/thermal modeling for the material response, motion modeling for pulse placement, and empirical calibration for feature size, taper, roughness, and failure mode [134, 139].

### 4.1.3 In-House Femtosecond Platform Integration

The in-house femtosecond system is treated as a platform build rather than as a commercial source placed on an optical table. The objective is to turn an existing industrial ultrafast source into a repeatable manufacturing instrument with measured pulse delivery, modular beam shaping, synchronized motion, and diagnostic access. A complete platform record reports the installed source configuration, available harmonics, pulse duration at the workpiece, repetition-rate limits, pulse-energy range by wavelength, burst capability, polarization control, average-power stability, and safety interlocks.



Figure 4.2: In-house femtosecond micromachining platform, showing the ultrafast source, optical table, beam-delivery optics, vertical optical breadboard, objective/imaging path, motion stage, chiller, and supporting electronics.

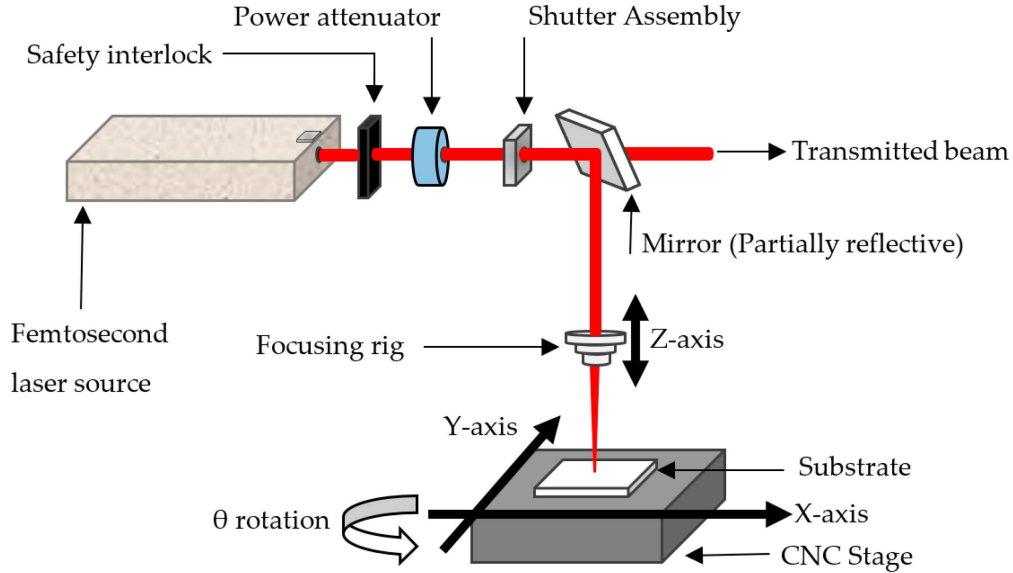


Figure 4.3: Simplified femtosecond laser micromachining setup from Owusu-Ansah and Dalton for normal-incidence machining of a 0.5 mm tungsten sheet, showing the source, shutter, attenuator, safety interlock, steering optics, focusing lens, and CNC stage [140].

Figures 4.2 and 4.3 define the hardware map used for the in-house build. The platform consists of five assemblies: the ultrafast source and harmonic-output region; the beam-conditioning and safety path, including the shutter or interlocked beam block, attenuator, power pickoff, irises, beam expansion, and steering mirrors; the process head, including the final focusing optic, micro-vision path, vertical optical breadboard, beam blocks, and local shielding; the motion and workholding assembly; and the utility layer, including the chiller, electronics, interlocks, and supporting controls.

Alignment starts from the 1034 nm path, which is used as the downstream optical datum. Reference irises define the table height and the vertical breadboard path. The selected wavelength is steered through those irises, centered through the final focusing optic, and checked against the camera field of view and stage coordinates. A usable alignment record includes delivered average power or pulse energy at the workpiece, beam diameter before the focusing optic, focus height, camera-to-beam registration, and a small coupon mark or equivalent focus check.

The micro-vision and focusing optics are built as an infinity-corrected camera arm with two practical objective states. The permanent receiver is a FLIR Flea3 FL3-U3-120S3C-C USB 3.0 color CMOS camera with a C-mount interface. The camera is treated as the fixed image plane of the system once the receiver arm is calibrated, rather than as the routine focusing adjustment. A Thorlabs AC254-200-A 200 mm achromatic doublet serves as the tube lens for visible imaging. It is mounted in the 30 mm cage system with the flat side facing the camera and the nominal tube-lens distance set by the 200 mm focal length, so collimated light from the objective is focused onto the FLIR sensor [141, 142].

The first objective state is the finder configuration. This uses a Mitutoyo M Plan Apo 2x long-working-distance, infinity-corrected apochromatic objective. With the 200 mm tube lens and the objective's 100 mm focal length, the imaging magnification is 2x. The 34.0 mm working distance gives enough clearance to find the sample, navigate the workpiece, and establish a safe absolute z-height before machining. This objective is therefore used for white-light setup, sample approach, and low-risk navigation rather than for high-fluence laser delivery [143].

The second objective state is the machining configuration. This uses a Thorlabs ASL2520 uncoated UV fused-silica aspheric lens with 20.0 mm effective focal length and  $NA = 0.65$ . In the same 200 mm

tube-lens imaging geometry, the equivalent magnification is approximately 10x. The current mechanical setup gives tight clearance, with a local working clearance of approximately 14.3 mm, so this configuration is less forgiving mechanically than the finder objective. It is nevertheless the practical high-NA machining path because the single uncoated fused-silica element avoids the cemented interfaces and visible-optimized coatings that can fail under high-power ultraviolet or green femtosecond operation. In the present build, this lens has been used for approximately 15 W ultraviolet/green service; quantitative machining results are tied to delivered average power, pulse energy, spot-size measurement, and lens-condition inspection rather than nominal laser-head power alone. In use, this optic focuses the expanded, approximately 12 mm laser beam toward the machining plane. Using the simple diffraction estimate  $r \approx 0.61\lambda/\text{NA}$ , the ideal first-minimum radius is approximately  $0.32 \mu\text{m}$  at 343 nm,  $0.48 \mu\text{m}$  at 515 nm, and  $0.97 \mu\text{m}$  at 1034 nm. These values are optical estimates, not measured feature sizes; aberration, beam quality, pulse-front distortion, focus error, material threshold, and pulse overlap still determine the real machined feature [144–146].

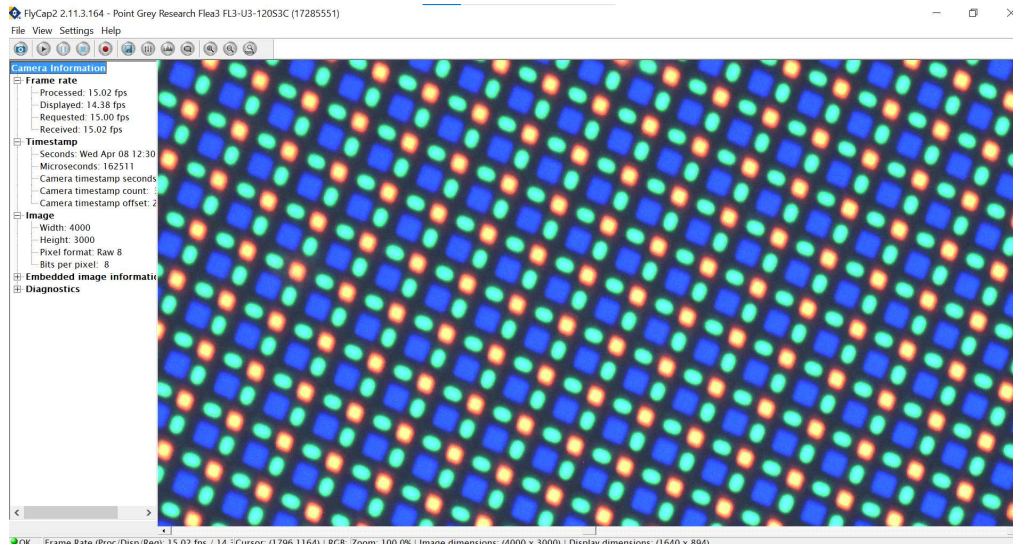


Figure 4.4: iPhone 15 OLED display pixels imaged through the in-house femtosecond laser micro-vision path. The specified 2556 by 1179 display at 460 ppi gives an approximately  $55 \mu\text{m}$  display-pixel pitch; the separated RGB subpixel structure demonstrates  $\mu\text{m}$ -scale inspection capability rather than machining calibration [147].

Figure 4.4 is an alignment and inspection result, not a final machining-resolution proof. The pitch estimate comes from the display specification,  $25,400 \mu\text{m}$  per inch divided by 460 pixels per inch, and refers to the full display-pixel spacing rather than the width of each colored subpixel. Figure 4.4 shows that the present vision path can resolve  $\mu\text{m}$ -scale display substructures and can therefore support focus finding, camera-to-stage registration, and visual inspection during process development. True cutting or feature-placement claims below  $1 \mu\text{m}$  require a stricter vision and metrology stack: the camera sampling at the workpiece must satisfy the intended inspection tolerance, the imaging path must be calibrated to the motion axes, the focus position must be measured or controlled within the useful depth of focus, and final features must be confirmed by SEM, AFM, confocal microscopy, or calibrated optical profilometry rather than only by the live process camera.

Common high-precision laser-processing vision architectures fall into several practical families. The simplest is a coaxial or near-coaxial microscope path that shares the focusing objective through a beamsplitter or dichroic, with illumination and filters chosen so plasma emission or machining light cannot saturate the camera. This is useful for setup and registration, but the camera image alone is usually not enough for focus control below  $1 \mu\text{m}$ . More advanced systems use active autofocus signals: diffractive beam samplers

and CCD/CMOS spot geometry for real-time focal-position sensing, machine-vision analysis of the laser focus pattern, low-coherence interferometry using the femtosecond beam itself, or dynamic z-scanning with a variable-focus element to relax the surface-flatness requirement during machining [148–152]. For this platform, the immediate development standard is to treat the live microscope image as a  $\mu\text{m}$ -scale observation tool and to add calibrated autofocus or interferometric focus tracking before claiming machining control below 1  $\mu\text{m}$ .

After alignment, coupon tests are recorded with the parameters needed to reproduce the machine state: wavelength, repetition rate, delivered average power or pulse energy, attenuator setting, focusing optic, beam diameter before the focusing optic, focus offset or z-height, scan speed, hatch spacing, pass count, polarization if controlled, sample material, and post-process measurements. This ties trial-to-trial comparisons to the optical alignment and motion state rather than only to the laser setting.

#### 4.1.4 Multi-Wavelength Machining Setup

The in-house femtosecond platform uses the fundamental, second-harmonic, and third-harmonic outputs as selectable process wavelengths rather than as simultaneous beams. The downstream machining axis is first aligned to the fundamental port near 1034 nm, which becomes the mechanical datum for beam height, iris locations, the camera reference, focusing optics, and stage coordinates. This datum is important because peer-reviewed diamond micromachining work has shown that the same CVD diamond material responds differently at 1030 nm, 515 nm, and 343 nm; wavelength choice must therefore be treated as a process variable, not as a cosmetic change in the source setting [153].




Figure 4.5: Three-output wavelength-routing region of the in-house femtosecond platform, showing the infrared, green, and ultraviolet output geometry and the fixed/kinematic mirror mounts used to route one selected wavelength at a time onto the common downstream machining axis.

The green path near 515 nm exits close to the fundamental port, so the current practical solution is a removable Thorlabs kinematic mirror-mount pair placed in front of the green output. The two mirrors fold the green beam onto the already established fundamental downstream axis. When this mirror set is installed,

it physically intercepts the original fundamental route, so the platform is operated in a wavelength-switched configuration: the shared machining path is common, but the active wavelength is selected mechanically. This is a deliberately simple architecture. It avoids rebuilding the beam-delivery rail for every harmonic, preserves the same workpiece coordinate system, and gives tip/tilt correction through the kinematic mount while keeping the number of inserted optics small.

The ultraviolet path near 343 nm is brought into the same common-axis scheme with a fixed folding mirror that directs the UV output toward the switching mirror geometry. In practice, the user swaps between the IR-aligned open path, the green kinematic routing set, and the UV routing set rather than trying to overlap all three wavelengths at the same time. The design should therefore be described as wavelength routing or beam-path selection, not coherent beam combining. After each wavelength swap, the acceptance check is beam height through the reference irises, beam centering at the steering mirrors, pointing through the focusing optic, camera-to-beam registration, delivered average power and pulse energy, focal-plane position, spot size, and a spectrum or pulse-duration health check when the optical path has changed.

The difficult part of this integration is not only steering the beams. Mirror routing is mostly geometric, but focusing femtosecond pulses at 343 nm, 515 nm, and 1034 nm through practical optics introduces wavelength-dependent focal shift, coating limits, transmission limits, laser-damage concerns, group-delay dispersion, and spatiotemporal distortion. Bor showed that lenses and lens systems can delay the pulse front relative to the phase front by picoseconds, and Kempe and Rudolph showed that chromatic and spherical aberration jointly affect the temporal distribution of focused ultrashort pulses [145, 146]. More recent autocorrelation z-scan work demonstrates why the focused femtosecond pulse should be evaluated as both a spatial and temporal object near the focal region rather than only as a static spot [154]. For the in-house platform, this means a single shared delivery axis is reasonable, but a single shared final lens is acceptable only after wavelength-specific validation. If one optic cannot preserve transmission, focus quality, pulse duration, and damage margin across all three wavelengths, the practical build should use wavelength-specific focusing optics or a repeatable lens-carrier swap with separate calibration records.

#### **4.1.5 Beam Shaping and Non-Diffractive Bessel Methods**

Beam shaping is a critical technique for moving beyond simple Gaussian focusing. A Gaussian focus has a small waist and limited depth of focus. In deep holes, the entrance aperture can clip the beam, the focal volume changes as the hole deepens, and debris or plasma can shield the bottom. Non-diffractive and Bessel-like beams extend the axial interaction length and can self-reconstruct after partial obstruction. This makes them attractive for high-aspect-ratio structuring, although side lobes, roughness, taper, and material-specific damage must be measured rather than assumed.

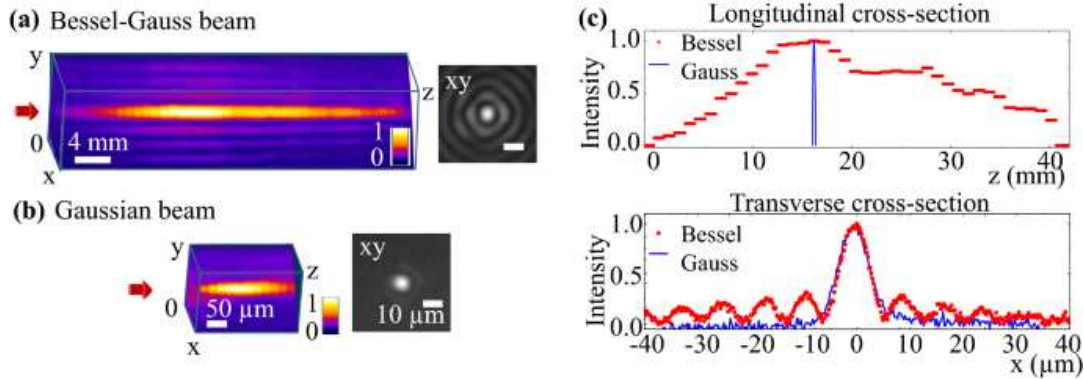


Figure 4.6: Non-diffractive beam comparison from Nguyen et al.: (a) Bessel-Gauss beam with  $3^\circ$  conical half-angle,  $11\ \mu\text{m}$  core diameter, and  $43\ \text{mm}$  Bessel length; (b) Gaussian beam with similar core diameter and  $0.24\ \text{mm}$  confocal range; (c) longitudinal and transverse intensity sections [155].

Figure 4.6 is important because it separates the useful Bessel idea from a vague claim of “better focusing.” Panel (a) shows that a Bessel-Gauss beam can maintain a narrow central core over a much longer axial distance than a Gaussian waist, while panel (b) shows how quickly a conventional Gaussian focus leaves its confocal region. Panel (c) also shows the tradeoff: the Bessel-like field has side lobes, so the method can extend the machining depth but may also expose surrounding material to additional energy. Nguyen et al. provide the clearest visual reference for the beam-shaping concept, but the study is focused on metals. The related claim that Bessel beams work especially well in glass is supported by glass-specific work such as Vetter et al., who studied high-aspect-ratio structuring of glass with ultrafast Bessel beams [156]. For the in-house system, a Bessel module would be a future or optional beam-shaping path: it could extend focus length and improve deep-feature access, but it must be validated locally by feature depth, taper, sidewall quality, side-lobe damage, and material-state measurements.

#### 4.1.6 Burst Mode and Wavelength Agility

Burst-mode processing is another important ultrafast-system variable. Closely spaced packets of pulses can change ablation efficiency, heat accumulation, melt dynamics, and shielding. Kerse et al. demonstrated ablation-cooled material removal with ultrafast bursts, showing that intra-burst timing can increase material removal efficiency relative to isolated pulses [157]. For the in-house platform, burst mode is treated as a controlled experimental variable rather than a generic improvement. A useful build record reports pulse energy, repetition rate, intra-burst spacing, number of pulses per burst, scan speed, overlap, and post-process roughness or phase data.

Wavelength selection is part of the same development layer. Common Yb-based harmonic wavelengths near  $1030\ \text{nm}$ ,  $515\ \text{nm}$ , and  $343\ \text{nm}$  give different tradeoffs in spot size, available pulse energy, optics durability, nonlinear propagation, and coupling to the material. Infrared operation can provide higher available pulse energy and simpler optics, green operation can reduce diffraction-limited spot size and improve practical coupling in many materials, and ultraviolet operation can support finer finishing where pulse energy and optics lifetime remain acceptable. Wavelength is therefore a controlled experimental variable rather than a parameter with one universally best value.

#### 4.1.7 Femtosecond Source Characterization: Spectrum, Bandwidth, and Dispersion

Femtosecond pulse characterization is a measurement chain because no normal photodiode or oscilloscope can directly follow a  $100\ \text{fs}$  pulse. The pulse is therefore inferred from quantities that electronics can mea-

sure after optical conversion: spectrum, nonlinear autocorrelation, frequency-resolved gated spectra, and dispersion-dependent nonlinear spectra. Peer-reviewed pulse-metrology literature treats ultrashort-pulse measurement as an indirect phase-retrieval problem; intensity autocorrelation and spectrum are useful but incomplete, while FROG-, ptychographic-, SPIDER-, MIIPS-, and d-scan-type methods are used when the time-dependent field is required [158–160].

The optical spectrum is the first and most important daily check because spectral bandwidth sets the shortest possible pulse duration. A short pulse requires a broad optical bandwidth: if the measured spectrum narrows, the pulse cannot remain equally short unless the pulse shape changes in a compensating and usually undesirable way. For a narrow spectrum centered at wavelength  $\lambda_0$ , the approximate frequency bandwidth is

$$\Delta\nu \simeq \frac{c \Delta\lambda}{\lambda_0^2}.$$

The transform-limited duration is then bounded by the time-bandwidth product. For intensity full-width at half-maximum (FWHM), a Gaussian pulse has  $\Delta\nu\Delta\tau \approx 0.441$ , while a  $\text{sech}^2$  pulse has  $\Delta\nu\Delta\tau \approx 0.315$ . These numbers are not a pulse-duration measurement by themselves. They are a lower bound obtained only if the spectral phase is flat, or at least properly compensated. In the in-house platform, the spectrum is therefore recorded at the source and after the installed beam path whenever optics, harmonics, compressor settings, or high-power operating conditions change.

The reason spectrum is so powerful is that it separates two questions that are often confused. The first question is whether the pulse has enough bandwidth to be short. The second question is whether the bandwidth is phased correctly so that those colors arrive together at the workpiece. A broad spectrum with large chirp can be much longer than its transform limit; a narrow spectrum cannot produce an extremely short pulse even with perfect compression. This is why spectrum plus autocorrelation is a useful health check, while spectrum plus FROG or d-scan is a field characterization.

The missing quantity in a spectrum-only measurement is spectral phase. The spectral phase can be expanded about the carrier angular frequency  $\omega_0$ :

$$\phi(\omega) = \phi_0 + \phi_1(\omega - \omega_0) + \frac{1}{2}\phi_2(\omega - \omega_0)^2 + \frac{1}{6}\phi_3(\omega - \omega_0)^3 + \dots$$

The constant phase  $\phi_0$  is usually irrelevant for micromachining. The linear term  $\phi_1$  is the group delay. The quadratic term  $\phi_2$  is group-delay dispersion (GDD), normally reported in  $\text{fs}^2$ ; it gives a linear chirp that makes different spectral components arrive at different times. The cubic term  $\phi_3$  is third-order dispersion (TOD), normally reported in  $\text{fs}^3$ ; it can create asymmetric pulses, pedestals, or satellite structure even when the FWHM looks reasonable. Higher-order terms and spatiotemporal couplings such as spatial chirp and pulse-front tilt become important when broad bandwidth beams pass through objectives, tilted windows, gratings, prisms, or imperfect harmonic-conversion optics.

For machining, dispersion is not a cosmetic laser parameter. At fixed pulse energy, the approximate peak power scales as  $E_p/\tau$ ; if the pulse stretches from 300 fs to 1 ps after a transmissive optics stack, peak power falls by about a factor of three before any material interaction changes are considered. That changes nonlinear absorption, plasma formation, ablation threshold, redeposition, and heat accumulation. Therefore, a serious process recipe specifies the delivered pulse energy and the pulse duration or diagnostic state at the workpiece, not only the laser-head setting.

The practical characterization workflow consists of measuring average power and repetition rate to obtain pulse energy, measuring spectrum to estimate the transform-limit lower bound and watch for bandwidth loss, using an autocorrelator for daily pulse-duration health, and using FROG/XFROG or d-scan when optical-path dispersion, compressor tuning, harmonic generation, or process sensitivity makes the actual pulse field important. This chain creates a defensible link from source setting to delivered fluence, peak power, and nonlinear coupling at the material.

#### 4.1.8 Autocorrelation for Daily Pulse-Duration Checks

A second-harmonic autocorrelator is an appropriate first pulse-duration instrument for routine femtosecond-platform operation because it is relatively compact, robust, and fast enough for alignment checks. The instrument splits a pulse into two replicas, delays one replica, overlaps the replicas in a nonlinear crystal, and records the second-harmonic signal versus delay. In scanning instruments the delay is produced mechanically; in single-shot non-collinear instruments the delay can be encoded spatially across a camera image [161].

What the autocorrelator returns is an intensity autocorrelation trace, not the pulse itself. A pulse duration is obtained by deconvolving the autocorrelation width with an assumed pulse shape. The deconvolution factor is different for Gaussian,  $\text{sech}^2$ , Lorentzian, and structured pulses, so a single autocorrelation width is not overinterpreted. The autocorrelator also cannot recover spectral phase. A chirped pulse, a structured pulse, or a pulse with a pedestal can produce an autocorrelation that appears acceptable while the workpiece sees a lower peak intensity or delayed energy. This limitation motivates reporting autocorrelation together with spectrum and, when needed, a field-retrieval diagnostic [158, 159].

Ramaiah et al. give a useful implementation example because the apparatus is simple enough to map onto a research micromachining platform. A beam splitter creates two paths whose lengths must be matched for temporal overlap in the KDP crystal. One path is made variable with a retro-mirror pair on a translation stage, while the other is folded through fixed mirrors. The two replicas are line-focused by a 300 mm focal-length cylindrical lens into a type-I KDP crystal; the setup uses a  $4^\circ$  crossover angle, a 20 mm by 20 mm by 6 mm KDP crystal, CCD recording, an aperture to reject the fundamental beams and unwanted second-harmonic light, KG-3 filters to suppress scattered fundamental radiation, and neutral-density filters to keep the camera signal in range [161]. This is the level of detail needed in the in-house diagnostic record: the pulse-width number is only meaningful if the measured optical path, nonlinear crystal, detector, filtering, and assumed pulse shape are reported together.

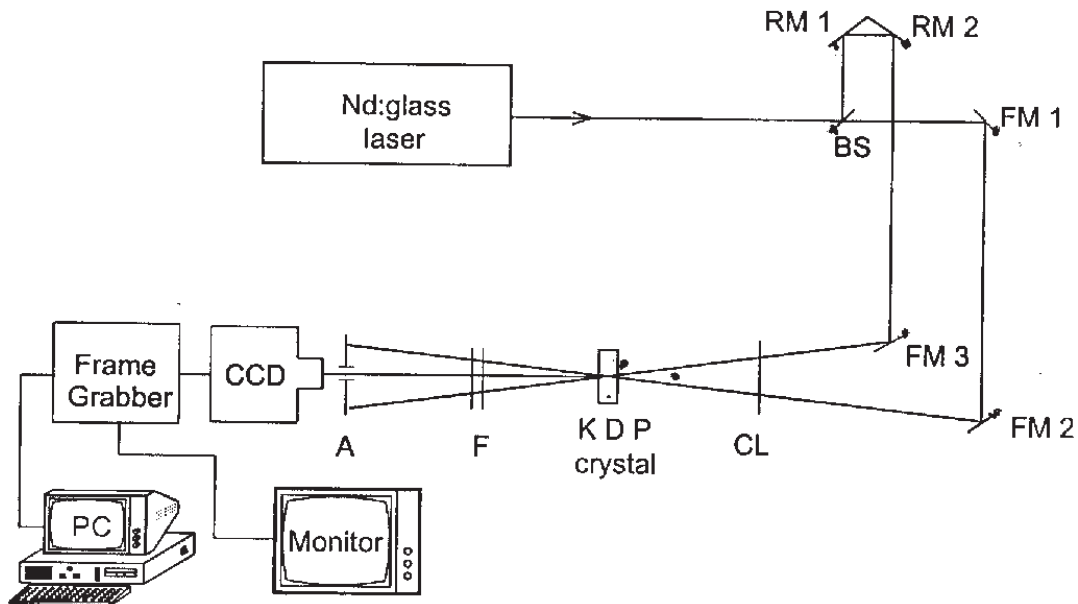


Figure 4.7: Experimental set-up of single-shot autocorrelator. RM: retro mirror, FM: folding mirror, BS: beam splitter, CL: cylindrical lens, F: filters, A: aperture. From Ramaiah et al. [161].

For the in-house femtosecond system, Figure 4.7 is useful because it shows the practical measurement chain rather than only the nonlinear-optics principle: beam splitting, path-length control, beam overlap in

the nonlinear crystal, filtering, aperture rejection, CCD acquisition, and computer readout. The autocorrelator functions as a pulse-duration and drift-monitoring tool. A complete characterization record includes the autocorrelation FWHM, assumed pulse shape, deconvolved pulse duration, measured optical spectrum, estimated transform limit, time-bandwidth product, wavelength, repetition rate, pulse energy, and optical path being measured. If the autocorrelation width increases while the spectrum stays constant, dispersion or compressor drift is likely. If the spectrum narrows, the transform limit itself has changed. If autocorrelation and spectrum disagree strongly, the pulse may have chirp, pedestal energy, or multiple-pulse structure, and the next diagnostic is a field-retrieval method such as FROG, XFROG, or d-scan.

Pulse-train stability is also treated carefully. Multi-shot measurements average many pulses. Jafari, Khosravi, and Trebino show that unstable pulse trains can create misleading diagnostic traces, and they discuss how discrepancies between measured and retrieved FROG traces can be used as an instability indicator [162]. This matters for machining because the material integrates many pulses along a toolpath. A stable average power reading does not prove that the individual pulses have stable shape, timing, or peak intensity.

#### **4.1.9 FROG/XFROG for Pulse-Field Retrieval**

Frequency-resolved optical gating (FROG) is the reference method for measuring the intensity and phase of an ultrashort pulse without assuming a pulse shape.

FROG spectrally resolves the nonlinear signal generated by two pulse replicas as a function of delay. The measurement is a two-dimensional trace of signal intensity versus frequency and delay. Trebino and Kane showed that this trace can be inverted as a two-dimensional phase retrieval problem to obtain the pulse intensity and phase [163]. Kane and Trebino also described practical FROG characterization of arbitrary femtosecond pulses [164], and the later Review of Scientific Instruments article gives the broader measurement framework and beam-geometry comparisons [158].

FROG answers the question the autocorrelator cannot answer: whether the measured bandwidth is actually compressed into a short pulse at the target. A useful FROG report includes the measured trace, retrieved trace, retrieval error, temporal intensity, temporal phase, spectrum, spectral phase, pulse duration, and time-bandwidth product. If the measured and retrieved traces disagree, the diagnostic is not treated as a clean pulse measurement; the disagreement may indicate algorithm stagnation, pulse instability, calibration error, or an optical artifact [160, 162].

XFROG uses a known reference pulse to characterize an unknown test pulse. It is especially useful when the test pulse is weak, spectrally shifted, temporally stretched, or generated by a process that cannot easily be self-referenced. Nakano et al. demonstrate cross-correlation, SHG, and self-diffraction FROG approaches for reliable femtosecond-pulse characterization [165]. Figure 4.8 shows why this class of diagnostic is more complex than an autocorrelator: the setup must manage reference-pulse routing, test-pulse routing, nonlinear mixing, calibrated delay, spectral collection, and a retrieval algorithm. For the in-house platform, FROG/XFROG is a future or shared-lab diagnostic: not required for every machining day, but important when a new harmonic path, compressor condition, beam-shaping module, or nonlinear delivery path is introduced.

#### **4.1.10 d-scan for Dispersion-Aware Pulse Diagnostics**

Dispersion-scan, or d-scan, characterizes femtosecond pulses by scanning a known or retrieved amount of dispersion while recording the spectrum of a nonlinear signal, commonly SHG. The resulting trace encodes the spectral phase because different dispersion settings compress or stretch the pulse differently. Miranda et al. demonstrated d-scan for broadband few-cycle pulses and discussed its sensitivity to noise, bandwidth, and pulse complexity [166]. Alonso et al. later extended the method with a self-calibrating approach that can use an arbitrary compressor and retrieve both the pulse and the compressor dispersion [167].

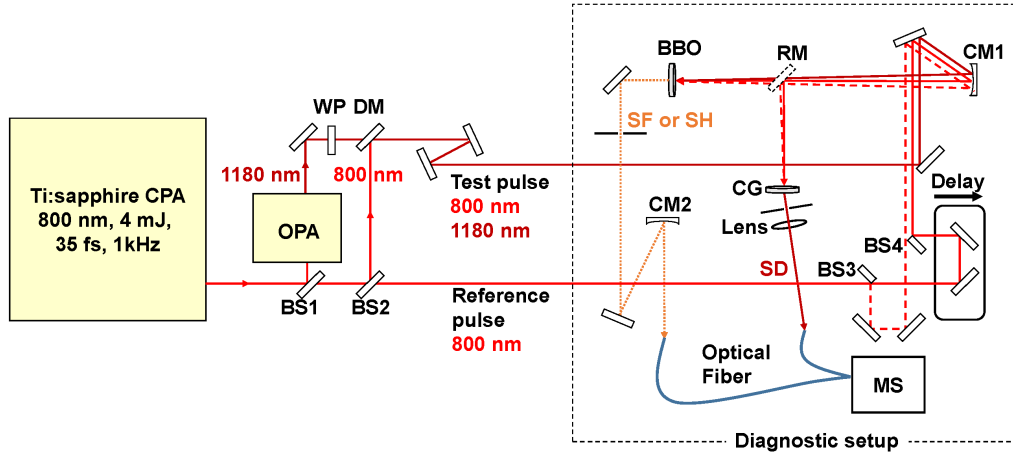


Figure 4.8: FROG/XFROG diagnostic setup from Nakano et al., showing reference- and test-pulse routing, half-wave plate, dichroic mirror, beam splitters, concave mirrors, BBO crystal, removable mirror, delay line, fiber coupling, and multichannel spectrometer for cross-correlation, SHG-FROG, and self-diffraction FROG measurements [165].

For an industrial CPA ultrafast source, d-scan is especially relevant because the laser already contains a compressor and the installed micromachining path adds additional dispersion. Alonso et al. note that for many practical systems the variable compressor dispersion can be modeled with GDD and TOD, while the pulse spectral phase itself can remain more complex [167]. Figure 4.9 shows this idea experimentally: prism insertion is scanned in a double-pass compressor, the pulse is focused onto a thin BBO crystal, the residual fundamental is filtered, and the SHG spectrum is recorded over many compressor settings. This is exactly the kind of information needed when a machining result changes after a lens, window, harmonic path, beam expander, or compressor setting is changed.

A d-scan characterization report includes the measured fundamental spectrum, nonlinear trace, retrieved spectral phase, retrieved temporal pulse, GDD/TOD estimate or compressor calibration, residual retrieval error, and the physical location of the measurement. The location matters: a pulse measured at the laser head is not automatically the pulse delivered after attenuators, harmonic optics, lenses, windows, and focusing optics. For process development, the operating standard is “pulse at the process plane” whenever the diagnostic can be practically inserted there.

#### 4.1.11 Live Femtosecond Process Monitoring and Adaptive Control

Process monitoring is separated from pulse diagnostics. Pulse diagnostics measure the delivered laser pulse; process monitoring observes the material interaction during machining. This distinction is essential for live feedback because a stable pulse does not guarantee stable ablation. Focus drift, debris shielding, plasma shielding, surface curvature, incubation, changing absorptivity, motion following error, and local material heterogeneity can all change the feature being written. Monitoring reviews and ultrashort-pulse manufacturing studies identify imaging, reflected or scattered light, plasma emission, acoustic signals, thermal radiation, vibration, and synchronized machine-state logs as the main in-situ channels for anomaly detection and feedback control [138, 168].

The most direct closed-loop demonstrations use a real-time proxy for ablation state rather than waiting for post-process microscopy. Chang et al. used laser-triggered plasma as the feedback signal for closed-loop ultrafast laser milling, showing that plasma emission can act as an observable linked to the state of the milling process [169]. Chen et al. developed a prediction-correction scheme for femtosecond-laser microchannel milling in which the measured process signal is used to correct the machining trajectory or process condition




Figure 4.9: Self-calibrating d-scan setup from Alonso et al., using prism-compressor insertion as the scan variable, an off-axis parabola, thin BBO crystal, blue filtering of the residual fundamental, and spectrometer acquisition of the SHG trace over 51 prism positions [167].

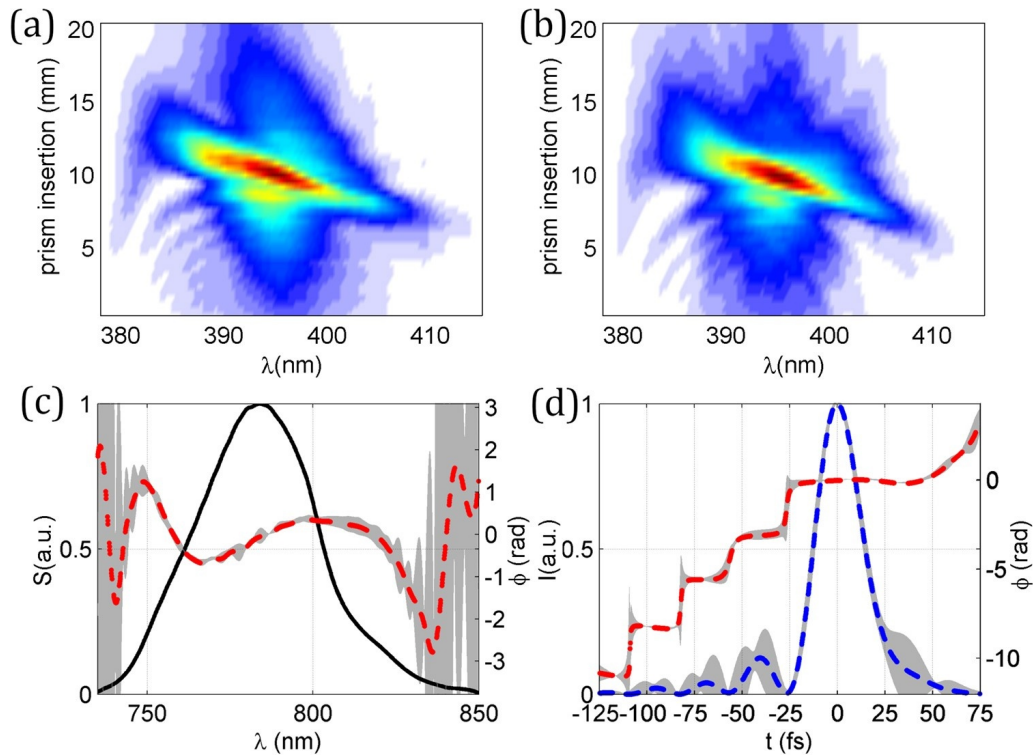


Figure 4.10: Measured and retrieved d-scan results from Alonso et al., including d-scan traces, pulse spectrum, retrieved spectral phase, temporal intensity, temporal phase, measurement uncertainty, and retrieved compressor GDD/TOD per unit insertion [167].

during fabrication [170]. More recent work combines theoretical modeling and machine learning to regulate tapered microchannels in silica glass, reinforcing the broader point that high-quality femtosecond machining is a data-driven control problem rather than a fixed recipe [171].

For the in-house femtosecond platform, the appropriate near-term architecture is a two-loop control hierarchy. The fast loop records process-plane signals during the toolpath: pulse trigger, stage or scanner position, average power, reflected or scattered light, plasma brightness or spectrum, and coaxial/off-axis images. That loop can detect loss of focus, loss of coupling, debris shielding, abnormal plasma intensity, or sudden changes in ablation state. The slower loop uses microscopy, profilometry, Raman spectroscopy, and cross-section data to update the process model between runs. For diamond and other carbon materials, Raman spectroscopy remains essential because feature geometry alone does not reveal graphitization, amorphous carbon, residual strain, or sub-surface damage [44, 172].

The feedback channels are therefore written into the process record rather than left as a disconnected instrument list. Coaxial and off-axis imaging track spot position, plume shape, debris accumulation, edge state, and large focus changes; these signals can trigger a pause, refocus, region rejection, or a revised scan order. Plasma or plume emission tracks ablation intensity, shielding, and changes in material coupling; this can support changes in pulse energy, scan speed, repetition rate, or pass count after calibration against real feature data. Reflected and scattered light are useful for detecting focus drift, entrance-hole clipping, beam clipping, or loss of process-plane coupling. Power logs, trigger logs, and stage-position logs separate laser drift from material-response changes by recording delivered average power, pulse-energy estimates, pulse timing, and synchronization. Post-process Raman, profilometry, microscopy, and cross-sections close the loop by connecting live optical proxies to graphitization, strain, roughness, depth, taper, and final geometry.

An adaptive fs-laser machining experiment should therefore be reported as a control system. The input variables are pulse energy, repetition rate, burst setting, wavelength, polarization, focus offset, scan speed, hatch spacing, pass count, and trajectory. The measured state is a synchronized vector of optical, emission, motion, and post-process signals. The control objective is not only to maximize removal rate; it is to maintain the desired feature depth, width, sidewall quality, heat-affected-zone limit, and material phase. This is a major research area for the platform because the same feedback logic can be transferred from coupon tests to diamond rotor bores, transparent-material channels, and waterjet-guided coupling studies.

## **4.2 Waterjet-Guided Laser Micromachining Systems**

### **4.2.1 Brief History and Current Direction**

Waterjet-guided laser machining began as an attempt to solve a very practical limitation of dry laser cutting: a tightly focused beam gives high intensity, but it also has a short depth of focus, a thermally loaded interaction zone, debris redeposition, and sensitivity to surface height. The key idea was to couple the laser into a fine laminar water jet so that the water-air interface guides the beam by total internal reflection while the same jet cools the cut and clears debris. Reviews trace the modern WJGL process to Bernold Richerzhagen's work at EPFL in the early 1990s, followed by commercial development through Synova's Laser MicroJet platform in the late 1990s [173, 174].

The early patent record is useful because it already frames WJGL as an opto-fluidic machine rather than as a laser with water added afterward. Richerzhagen's liquid-guided laser patent describes focusing a laser into a nozzle-fed liquid jet and explicitly discusses thermal lensing in the water above the nozzle as a source of focus shift and potential nozzle damage [175]. That observation remains central today: the beam, chamber flow, window, water absorption, nozzle, pressure stability, and workpiece response must be designed together.

The field has now moved from proving that a laser can be guided in water to controlling how much useful energy reaches the material and what feature quality results. Current WJGL work includes optical-field

regulation for taper and deep-processing control, dual-beam coupling to increase delivered power, gas-assist modules for deep blind features, real-time beam/process measurement, response-surface or multi-objective optimization, and diamond-specific machining studies [176–180]. The in-house WJGL system in this thesis is positioned inside that modern transition: the first goal is a repeatable research platform with measured coupling, jet stability, and feature quality; the longer-term goal is controlled WJGL micromachining for hard, brittle, and heat-sensitive materials.

#### 4.2.2 Working Principle

Waterjet-guided laser micromachining couples a pulsed laser into a laminar water jet so that the jet acts as a liquid optical waveguide by total internal reflection. The result is not simply water-assisted laser machining. The water jet provides an extended, nearly cylindrical energy-delivery region, cools the interaction zone, removes debris, and increases working distance relative to a conventional dry focused beam [173, 174].

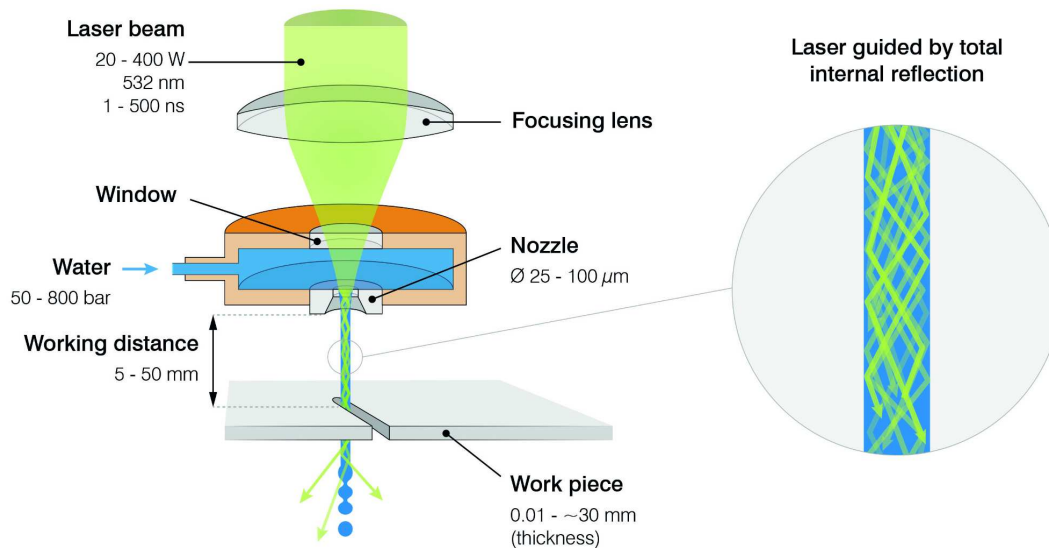


Figure 4.11: Basic working principle of waterjet-guided laser machining: a laser is focused into a pressurized water jet, total internal reflection guides the beam along the jet, and water cooling/debris removal occur at the workpiece.

The WJGL head is a coupled opto-fluidic system. Stable performance depends on laser wavelength, pulse duration, focusing geometry, protection-window loading, water absorption, chamber flow, nozzle geometry, jet coherence, standoff, assist gas, sweep gas, drainage, and workpiece response. Richerzhagen’s early liquid-guided laser patent is important because it identified thermal lensing in water above the nozzle as a root cause of focus shift and nozzle damage, making clear that water flow and optical coupling must be co-designed rather than treated separately [175].

#### 4.2.3 System Architecture

A practical WJGL system contains the laser source, beam conditioning optics, focusing unit, protective window, coupling chamber, water supply, pump, filtration, pressure regulation, nozzle, motion system, workholding, catcher/drain, and process observation. Figure 4.12 is useful because it separates these functions into laser/optical, water-supply, air-supply, coupling, imaging, and motion/workpiece modules while still showing that they meet at the coupling head. Each subsystem can limit performance. A dirty or damaged window reduces coupling. A worn nozzle changes jet diameter and stability. Bubbles or particles disrupt

the jet. Poor drainage causes splash and secondary damage. Incorrect pressure can break the jet or reduce guidance. Poor motion synchronization causes taper or kerf variation.

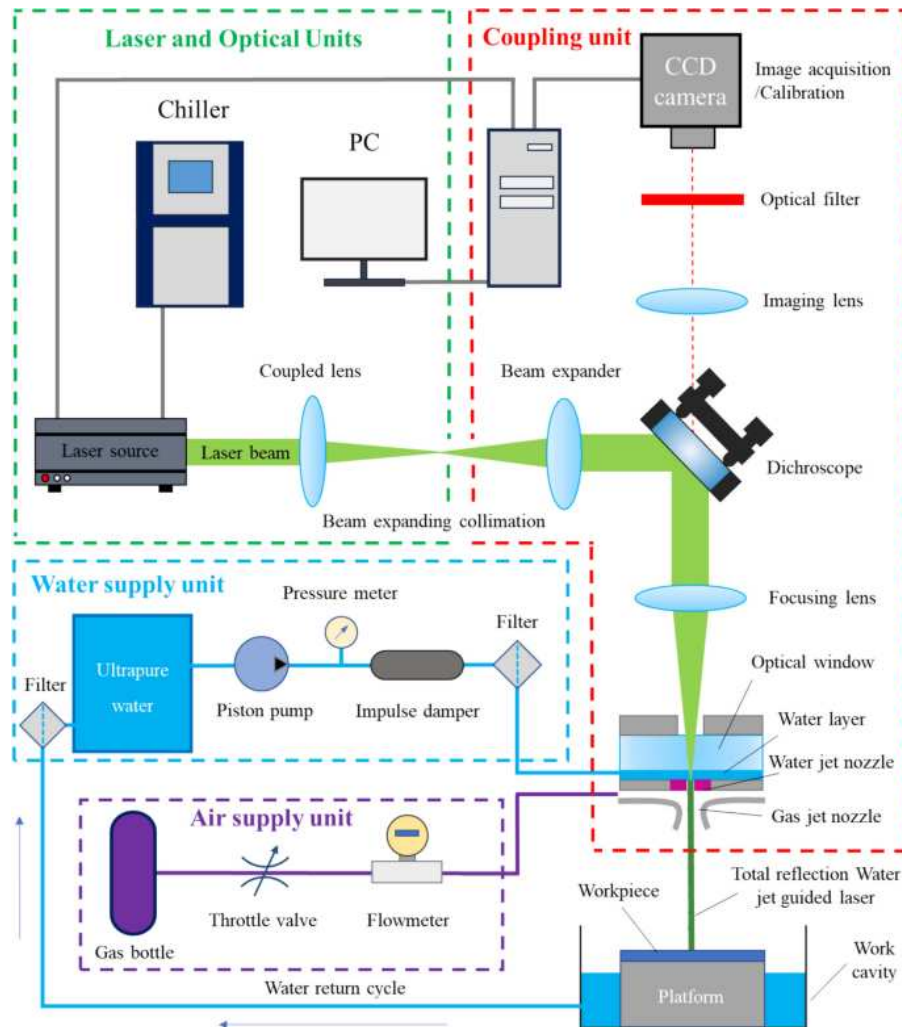


Figure 4.12: Waterjet-guided laser system architecture from He et al., identifying the laser/optical unit, coupling unit, water-supply unit, air-supply unit, work platform, CCD observation path, beam expansion/collimation optics, dichroic element, focusing lens, optical window, nozzle, gas jet, water layer, and total-reflection guided path [173].

#### 4.2.4 External Industry Benchmark

Commercial Laser MicroJet and related WJGL systems show why this process class is attractive: a guided water jet can produce long working distance, debris removal, cooling, and low-taper cutting in hard or heat-sensitive materials. The Avonisy demonstration is retained as an external industry benchmark rather than as in-house work.

#### 4.2.5 In-House WJGL Platform Integration

The in-house WJGL system is a research-grade opto-fluidic machine build. The platform combines an existing pulsed laser source, beam-conditioning optics, focusing optics, a protection-window and water-chamber




Figure 4.13: Avonisys finished-material demonstration showing 0.25 mm-pitch features in 18 mm aluminum; the black material on the left is silicon.

head, a high-pressure pump, pressure regulation, filtration, nozzle hardware, camera access, drainage/catcher hardware, and motion/workholding. This section does not claim that the system already matches a commercial green WJGL engine; it documents the optical, hydraulic, mechanical, and sensing stack being built to study and eventually control WJGL micromachining.

The WJGL source is evaluated by delivered pulse energy and coupling survivability, not only by nominal average power. All three in-house generations were built around the available 60 W, 1064 nm MOPA laser source. This source is flexible and programmable, but it creates a harder coupling problem than a mature commercial green WJGL system. Green wavelengths near 515 nm to 532 nm are common in WJGL because water absorption is much lower and coupling is better established; near-infrared operation increases the burden on absorption, thermal lensing, protection-window loading, and component protection [174, 177].

The hydraulic system is treated as pressure headroom plus regulation. The current in-house water system uses an air-powered hydraulic pump and has been operated at approximately 23,000 psi water pressure. This pressure level makes filtration, pressure control, pulsation damping, and near-head pressure sensing as important as maximum pressure. The nozzle is also not just a consumable aperture. Nozzle material, inlet geometry, orifice diameter, cavitation behavior, coaxial gas, sweep gas, drainage, and catcher design all influence jet stability and process repeatability. In the current third-generation head, assist air is supplied at approximately 80 psi and the lens-cleaning air is supplied at approximately 5 psi, so gas flow is treated as a separately regulated machine variable rather than an unspecified shop-air connection.

The in-house development path is therefore divided into four coupled builds. The optical build protects the source while delivering a controlled focus into the water chamber; this requires a clean window, repeatable lens-to-nozzle distance, lateral alignment access, and a way to measure coupled or transmitted power. The hydraulic build produces a stable, filtered, bubble-free jet at the nozzle, with pressure measured close enough to the head that pump pulsation and line losses are visible. The mechanical build makes the head serviceable, because nozzle inspection, window cleaning, lens changes, and alignment iterations happen repeatedly. The process-observation build allows jet imaging, workpiece viewing, and post-process inspection so that coupling failures are not mistaken for material limits.

The design record preserves the WJGL hardware iterations as engineering evidence. Figure 4.14 is

valuable because the first generation proved physical access and revealed packaging problems. Figure 4.15 is valuable because the second generation brought the focusing lens, water chamber, and nozzle into a more coherent stack. Figure 4.16 is valuable because the third generation moves toward a compact, serviceable research head with room for sensing and future gas-assist modules. Figures 4.14–4.16 are therefore placed with the corresponding generation rather than grouped as a visual appendix: each figure documents what changed, what limitation drove the change, and what measurement will decide whether the new version is better.

#### **4.2.6 First-Generation In-House WJGL Design**

The first-generation design established the initial hardware stack: laser delivery from the 60 W MOPA source to a vertical head, basic workpiece access below the nozzle, a fixture path for the jet, and a mechanical structure that could hold the head relative to the workpiece. This generation was designed and built by Max Aalto, Vineet Padia, and John Zhang during Neil Gershenfeld's How to Make Something That Makes (almost) Anything class. The design was intentionally close to a catalog-optics prototype. It relied heavily on Thorlabs-style commercial mounts and tube/cage hardware so that the optical axis, water chamber, and mechanical support could be assembled quickly and modified without waiting for fully custom machined parts. The optical concept used an aspheric focusing lens and a 12 mm optical window between the dry optical path and the water chamber. That large window was useful for early packaging, but it was not ideal for a compact WJGL head because it increased the pressure-loaded optical area, made the head bulkier, and created a larger element to seal, clean, and protect from contamination or damage. The value of the first generation was therefore not final performance; it made the coupling, sealing, plumbing, drainage, and access problems visible.

The first-generation lens/head CAD design recorded the first attempt to align the focusing path, head body, nozzle, and gas/water interfaces as a single assembly rather than as separate bench components. It also marks the beginning of the transition from commercial optomechanical parts toward purpose-built WJGL hardware.




Figure 4.14: First-generation in-house WJGL design and hardware: 1064 nm MOPA laser, tip-tilt mount, XY stage, CMOS camera path, camera focusing lens, dichroic element, aspheric focusing optic, and initial laser-water coupling unit with protective lens, water chamber, and nozzle. Designed and built by Max Aalto, Vineet Padia, and John Zhang.

#### 4.2.7 Second-Generation In-House WJGL Design

The second-generation design moved toward a more integrated machine layout and represents the intermediate stage between a catalog-part prototype and a custom head. This version was designed and developed by Max Aalto. Compared with the first generation, this version treated the optical input, focusing optics, water chamber, window, and nozzle as a more compact stack. It still used the same 60 W MOPA laser source, but the head geometry began to move away from general-purpose Thorlabs-style assemblies toward machined parts that constrained the optical and hydraulic axes more directly.

The second-generation head and focusing-lens design is the clearest record of this intermediate coupling stack. This version explored an axicon-lens element together with an aspheric element used as part of the window or coupling interface. The motivation was to improve the optical field delivered into the water/nozzle region while reducing the bulk associated with the first-generation flat-window stack. The limitation was mechanical and operational robustness: using shaped optical elements as part of the window/interface makes sealing, cleaning, alignment tolerance, pressure loading, and replacement more difficult than a head built around a separate flat pressure window and a protected focusing optic. The second generation therefore clarified the need for a more custom, serviceable head rather than serving as a final design.




Figure 4.15: Second-generation in-house WJGL design: 1064 nm MOPA source, XY stage, high-magnification imaging path, ring light, prism mirror, and revised head/lens section with axicon lens, focusing ring, aspheric lens, O-ring seal, and 50  $\mu\text{m}$  nozzle. Designed and developed by Max Aalto.

#### 4.2.8 Third-Generation In-House WJGL Design

The third-generation design is the first version organized as a practical research machine rather than a fast optomechanical prototype. This version was worked on with Max Aalto, Leo Cheng, and Abigail Suk. It moves to custom mounts and a custom head architecture while retaining the same 60 W MOPA source. The optical stack returns to an aspheric focusing lens, but the pressure/window interface is reduced to a 4 mm flat optical window. This separates focusing from pressure sealing more cleanly than the second-generation shaped-optic interface, reduces the pressure-loaded area relative to the first-generation 12 mm window, and makes window replacement and cleaning more manageable. The operating architecture also separates the fluid and gas functions: the air-powered hydraulic pump supplies the water line at approximately 23,000 psi, assist air is supplied at approximately 80 psi, and lens-cleaning air is supplied at approximately 5 psi. These pressures are treated as machine-state variables because they affect jet stability, optical-window cleanliness, splash behavior, and process repeatability. The third-generation composite in Figure 4.16 records the head CAD, installed vertical head, optical/camera access, plumbing, nozzle region, workpiece/catcher area, and granite-gantry machine concept in a single figure.




Figure 4.16: Third-generation in-house WJGL design: compact custom head CAD, assembled vertical head in the machine frame, and granite bridge/base machine concept with workpiece stage and catcher region. Developed with Max Aalto, Leo Cheng, and Abigail Suk.

#### 4.2.9 Dual-Beam Coupling as a Future WJGL Upgrade

Dual-beam coupling is a state-of-the-art WJGL direction, not an in-house result in the current system. The basic idea is to combine beams before coupling into the water jet so that more optical energy can be delivered into the liquid waveguide while controlling polarization and coupling geometry. Figure 4.17 shows one implementation in which two 532 nm solid-state laser channels are combined with polarization optics and time-domain pulse control before focusing into the micro water jet [177]. This matters for deep processing because the limiting variable is often not nominal source power but power successfully coupled into the stable jet.

In the in-house system, this idea would require a future combining stage before the focusing lens, polarization control or beam-combining optics, a power monitor before the head, and a second power or scatter-based diagnostic after coupling into the jet. It is best added after the single-beam system has stable jet formation and a repeatable coupling metric. Otherwise, added source power would make the failure mode harder to interpret.

#### 4.2.10 Dual-Beam Coupling Performance Metrics

Figure 4.18 is useful because it shows the kinds of machine metrics that the in-house WJGL system eventually reports: coupled power, transmitted power in the jet, pulse energy, and response to optical alignment parameters. Figure 4.18 includes beam deflection at the water-air interface, the optical path inside the jet, coupling efficiency versus jet diameter and processing distance, coupled power versus coupled jet length, and pulse energy versus coupled jet length. Wang et al. report up to 90.4% coupling efficiency at 10 kHz, corresponding to 36.01 W average coupled power; those values are their result, not a transferable guarantee for the in-house system [177].

For local development, the important lesson is the measurement style. A WJGL power claim identifies incident laser power, loss before the head, power coupled into the jet, transmitted or process-side power where measurable, jet diameter, pressure, nozzle condition, and alignment state. A cut-quality claim is then




Figure 4.17: Dual-beam WJGL coupling setup from Wang et al., combining two 532 nm solid-state laser channels with polarization optics and time-domain pulse control before focusing into the micro water jet [177].

ties to these machine metrics. This prevents a false comparison between a well-coupled low-power condition and a poorly coupled high-power condition.

#### 4.2.11 Optical-Field Regulation

Optical-field regulation is important because WJGL cuts can still develop taper and nonuniform energy delivery even when the beam is guided by water. Zhang et al. show an approach for reshaping or regulating the optical field inside the water jet to reduce taper and improve deep-processing capability [176]. Figure 4.19 compares centered and eccentric beam conditions, the corresponding cross-section energy density, and the groove-depth response. For the in-house platform, this suggests a future path in which the coupling optics are not only aligned for maximum power but designed to control the radial and axial energy distribution in the jet.

This would be incorporated as an optical upgrade after baseline coupling is stable. The practical route would be to compare a standard focused beam with a shaped or regulated field at the same jet diameter and pressure, then measure kerf taper, sidewall roughness, groove depth, and debris behavior. If the regulated field improves taper without sacrificing coupling stability, it becomes a candidate for deep-slot and high-aspect-ratio processing.

#### 4.2.12 Gas Assistance for WJGL Head Development

Gas assistance in WJGL is different from assist gas in dry laser cutting. In a dry laser cutter, gas mainly removes melt, changes chemistry, or protects the optic. In WJGL, the water jet is also the optical waveguide, so any gas must preserve the stable water-air interface that guides the laser. A useful gas-assist design can shield the jet from reverse splash, extend the coherent jet length, or sweep stagnant water and debris from a

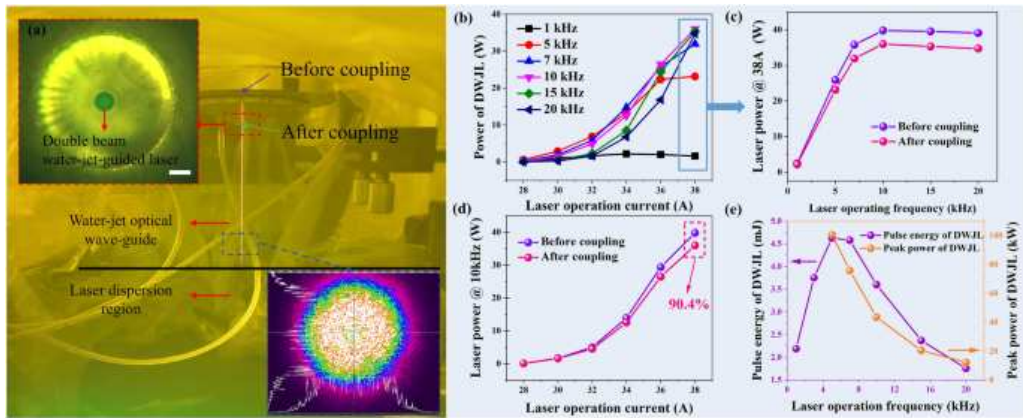


Figure 4.18: Dual-beam WJGL coupling-performance measurements from Wang et al., showing micro-jet transmission behavior, coupling efficiency versus jet diameter and processing distance, coupled power versus coupled jet length, and pulse energy versus coupled jet length [177].

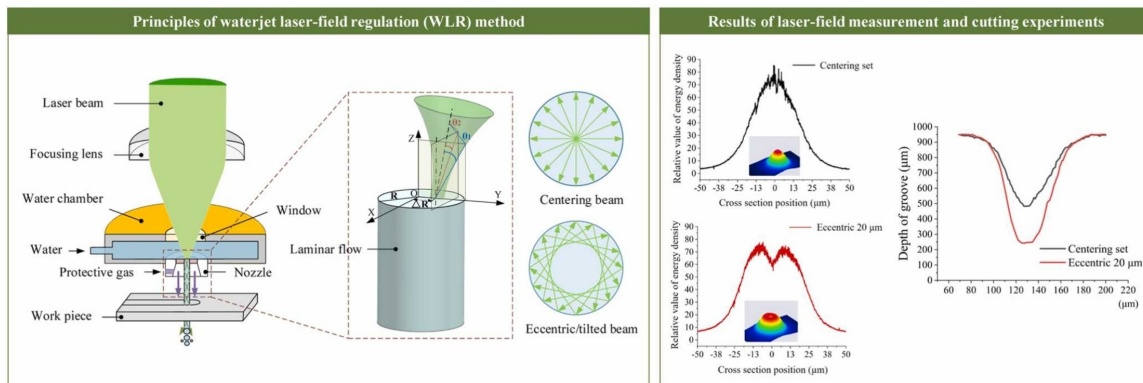


Figure 4.19: Optical-field regulation method for WJGL from Zhang et al., comparing centered and eccentric beam conditions, cross-section energy-density distribution, and groove-depth response for taper reduction in deep processing [176].

deep blind hole. A poor gas-assist design can shear the jet, inject bubbles, disturb total internal reflection, and reduce coupling.

The “double coaxial” language in Cao et al. refers to a nested flow concept rather than simply the number of visible ports in Figure 4.20. Their baseline system already used a water jet with coaxial helium assistance. The proposed upgrade added an outer coaxial argon jet around that existing water-plus-helium structure. The helium flow is used as an inner protective envelope around the water jet; previous work cited by Cao et al. reports that a coaxial helium envelope can increase stable water-jet length and machining depth. Helium is attractive in that role because it is inert, low density, and can form a fast protective sheath with relatively low optical absorption. The outer argon is used differently: it is a denser sweep/shielding gas intended to push away reverse water and deposited water in deep blind holes. In Cao et al., this reduced the penetration time for a 0.7 mm diameter, 4 mm deep hole from 410 s to 173 s, and the paper reports 0.7 mm holes with an aspect ratio of 14.6:1 in a superalloy thermal-barrier-coating specimen [178]. A related coaxial-annulus argon study also frames gas assistance as a way to improve WJGL performance on hard-to-process materials, including ceramic-matrix composites and CVD diamond [181].

For the in-house WJGL system, helium or argon should not be treated as a simple add-on to the existing water line. Gas must not be mixed into the water supply or nozzle cavity, because bubbles would break the optical waveguide and destabilize the jet. If gas assistance is used locally, it should be a separate annular head module outside the water jet, with its own regulator, flow meter, pressure gauge, sealing, exhaust/drainage path, and visual jet-stability check. Helium also has practical disadvantages: it is expensive, leak-prone, and requires dedicated plumbing to justify the benefit. The current in-house priority is therefore water-only baseline coupling and jet stability. Gas assistance is a future module to test only after the baseline head can repeatedly form a clean jet, couple measurable optical power, and machine simple validation features.

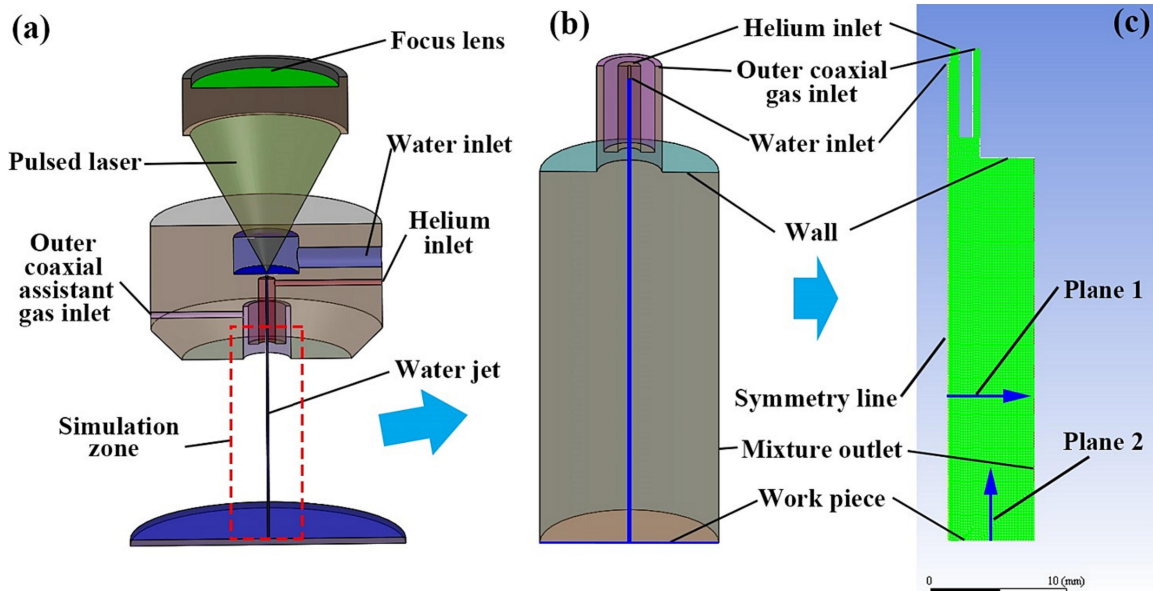


Figure 4.20: Gas-assisted WJGL reference from Cao et al., showing the processing head, simulation geometry, focusing lens, pulsed laser, water inlet, helium inlet, outer coaxial gas inlet, water jet, gas nozzle, workpiece, mixture outlet, and deep-hole flow model [178].

#### 4.2.13 Diamond-Oriented WJGL Direction

The diamond-oriented WJGL papers point toward future platform development after the basic in-house machine is stable. Wang et al. model WJGL diamond machining as a multi-objective optimization problem

rather than a single-depth maximization problem. Cutting depth, kerf width, cutting speed, focal-plane offset, pulse frequency, pulse width, and surface roughness are coupled response variables, and their improved NSGA-III workflow shows the kind of dataset structure needed for local process development [179]. Figure 4.21 is especially useful as a checklist because it connects the WJGL processing system, water-light coupling schematic, coupling-power measurement, diamond morphology measurement, and optimization workflow. The reported equipment stack includes a 100 W, 532 nm laser, 50 bar to 500 bar ultrapure-water system, coupling controller, motion control, compressed-air supply, helium supply, and microscopy for diamond feature measurement [179]. This direction treats WJGL diamond machining as a controlled system-identification problem.

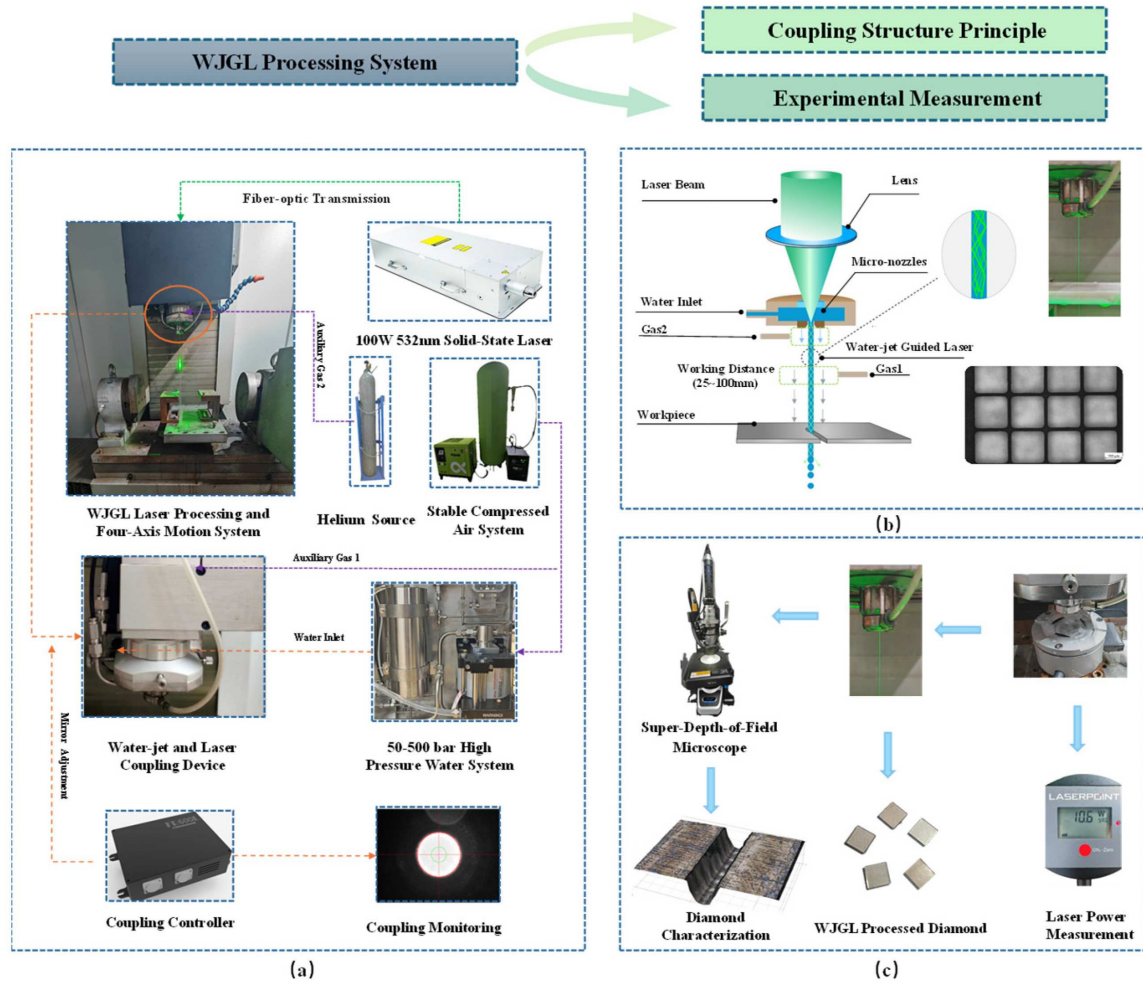


Figure 4.21: WJGL diamond machining system, metrology, and optimization workflow from Wang et al., linking the processing system, water-light coupling, coupling-power measurement, diamond morphology measurement, and NSGA-III process optimization [179].

Chen et al. study energy-dependent WJGL processing of single-crystal diamond and show why power cannot be treated as a monotonic improvement variable. Figure 4.22 shows average laser power decreasing with underwater depth for initial powers from 8 W to 20 W, then compares effective jet length at 10 W and 20 W; the bright scattering spots at higher power indicate water-jet disturbance [180]. Their results frame diamond WJGL machining as a transition from threshold-limited removal to higher-energy grooves where removal rate, aspect ratio, kerf width, water-jet disturbance, and surface roughness are optimized together [180]. For the in-house system, this paper motivates recording machine settings, jet state, coupled power,

groove or kerf geometry, surface roughness, and material-state data for every diamond coupon instead of only reporting whether a cut completed.




Figure 4.22: Energy-dependent WJGL groove response in single-crystal diamond from Chen et al., showing underwater power attenuation for initial powers of 8 W to 20 W, effective jet length at 10 W and 20 W, and high-power water-jet disturbance [180].

#### 4.2.14 WJGL Validation Workflow

WJGL performance is evaluated by machine metrics and feature metrics. Machine metrics include coupling efficiency, transmitted power in the jet, jet diameter, breakup length, nozzle lifetime, pressure stability, window temperature or damage, alignment tolerance, and pressure ripple. Feature metrics include kerf width, depth, taper, sidewall roughness, entrance/exit chipping, redeposition, heat-affected zone, and repeatability. For difficult materials, microscopy and 3D profiling are supplemented by chemical or microstructural characterization, because reduced thermal damage is not the same as zero material modification [182].

The recommended in-house validation sequence is staged. First, prove a stable water jet without laser: pressure stability, jet diameter, breakup length, splash behavior, drainage, and catcher performance. Second, prove optical coupling at low power: coupling image, transmitted power, window temperature or damage inspection, nozzle inspection, and tolerance to focus and lateral offsets. Third, run inert or low-risk materials to map kerf width, depth, taper, redeposition, and repeatability. Fourth, move to difficult materials only after machine-level metrics are stable. This sequence keeps source, hydraulic, and coupling failures from being misinterpreted as material-process limits.

## Chapter 5

# Characterization and Spectroscopy Tools

### 5.1 Chapter Overview

This chapter describes two linked spectroscopy instrument-building efforts: a single-laser 532 nm confocal Raman imaging platform and a low-cost personal multirole 3D-printed modular spectrometer project. The chapter is intentionally centered on instrumentation. The main contribution is the local construction, validation logic, and capability expansion of spectroscopy tools.

The two projects operate at different levels of performance and accessibility. The 532 nm confocal Raman platform is a research-facing instrument intended for local spectra, optical imaging, point scans, raster maps, and spatially registered process feedback once its calibration, resolution, confocal rejection, SNR, and repeatability are measured. The personal spectrometer is an access-facing instrument: it exposes the slit, grating, optics, detector, calibration, and software steps that commercial instruments usually hide. It is useful for education, replication, and selected measurements after validation, but it is not presented as a substitute for a calibrated confocal Raman microscope.

### 5.2 Motivation for Open Spectroscopy

The motivation for building open Raman and open spectrometer systems is not only cost reduction. Cost matters, because a commercial confocal Raman microscope can be financially inaccessible to small laboratories, teaching groups, community colleges, high schools, and researchers in resource-limited settings. But the deeper motivation is scientific agency: when the optical path, electronics, software, calibration files, and bill of materials are visible, the instrument becomes something that can be understood, repaired, modified, and taught rather than simply operated.

Open scientific hardware changes the social role of instrumentation. Pearce argues that free and open-source hardware can lower the cost of research equipment and make customized instruments more accessible to scientists who cannot justify or afford low-volume commercial systems [183]. Baden et al. make a related case for open labware: shared designs, local fabrication, and inexpensive electronics can help laboratories and schools build useful tools faster, especially where supply chains and budgets are limiting factors [184]. This matters for spectroscopy because light-based measurements are broadly useful: they can support chemical education, materials identification, quality control, environmental sensing, cultural-heritage analysis, biological sample screening, and fabrication feedback.

An open spectrometer can let a student see how a slit, grating, detector, and calibration curve turn light into data. An open Raman system can let a researcher see how laser wavelength, Rayleigh rejection, confocal aperture, detector noise, and signal processing shape a spectrum. That visibility changes the user from a button-pusher into an instrument builder. In practical terms, this can help people troubleshoot local prob-

lems, reproduce measurements, adapt a system to unusual samples, and share improvements back to a wider community.

The humanitarian value should be stated carefully. A low-cost open Raman system is not automatically a medical diagnostic device, a regulatory instrument, or a substitute for a calibrated commercial microscope. Its value is that it lowers the threshold for learning, experimentation, and local capability. When more people can build and understand measurement tools, more people can participate in science, identify materials in their own environment, teach spectroscopy with real hardware, and adapt instruments to problems that a commercial vendor may never target. In that sense, open spectroscopy is part of a larger movement toward more distributed, repairable, and locally owned scientific infrastructure.

## 5.3 Single-Laser 532 nm Confocal Raman Platform

### 5.3.1 Design Target

Raman spectroscopy measures inelastically scattered light, so it can identify molecular vibrations, crystal phases, strain-sensitive peak shifts, and unwanted background or fluorescence depending on sample class. Confocal Raman microscopy adds spatial selectivity by focusing excitation into a small volume and rejecting out-of-focus light with a pinhole, fiber core, or slit. Modern Raman imaging reviews emphasize that this combination can support chemical imaging from micrometer-scale maps down toward specialized near-field and enhanced modalities, but only when the optics, calibration, and data processing are treated as a coupled system [185, 186].

The local system was built as a single-excitation-wavelength platform. Earlier design discussions considered both visible and near-infrared excitation, but the completed Raman path described here is organized around 532 nm. That distinction matters because wavelength is not a cosmetic choice; it defines Raman efficiency, fluorescence risk, sample heating, optical filtering, detector response, diffraction-limited spot size, and laser-safety practice.

### 5.3.2 Laser Wavelength Selection

The laser wavelength for a Raman microscope is selected by balancing signal strength, background, optical resolution, sample compatibility, and component availability. For spontaneous Raman scattering, all else equal, the scattering efficiency increases strongly at shorter excitation wavelengths and is commonly discussed with an approximate  $1/\lambda^4$  dependence. This favors visible excitation when fluorescence and absorption are manageable [187]. A shorter wavelength also improves the diffraction-limited focus because the lateral spot size scales with wavelength for a fixed objective numerical aperture.

532 nm was therefore a practical choice for the first local confocal Raman build. It is compatible with common diode-pumped solid-state lasers, visible optics, dichroic beamsplitters, long-pass or notch filters, CCD-class detectors, and alignment cameras. It also provides a tighter focus than 785 nm excitation under the same objective. The tradeoff is that 532 nm can increase fluorescence, photothermal loading, and photochemical risk in absorbing or fluorescent samples. A 785 nm path remains a useful future option for polymers, biological samples, dark powders, and unknown materials, but it should be described as a future or comparative capability unless it is built, interlocked, aligned, calibrated, and tested to the same standard as the 532 nm path.

The selection criteria used for the local system were:

- strong Raman scattering from materials that tolerate visible excitation;
- compact and affordable laser, filter, dichroic, objective, and detector options;

- visible beam alignment during optical construction;
- compatibility with confocal imaging and stage-scanned maps;
- acceptable laser-safety controls for a benchtop prototype.

In practical terms, the laser was selected as an instrument component rather than as a wavelength label alone. The source must have a narrow enough linewidth that it does not blur Raman features, good beam quality so the objective can form a small and stable confocal focus, and power control stable enough that repeated spectra and maps are not dominated by source drift. Power supply stability is part of the optical design: the system must provide enough sample-plane power for SNR while avoiding fluorescence growth, thermal shifts, and photodamage.

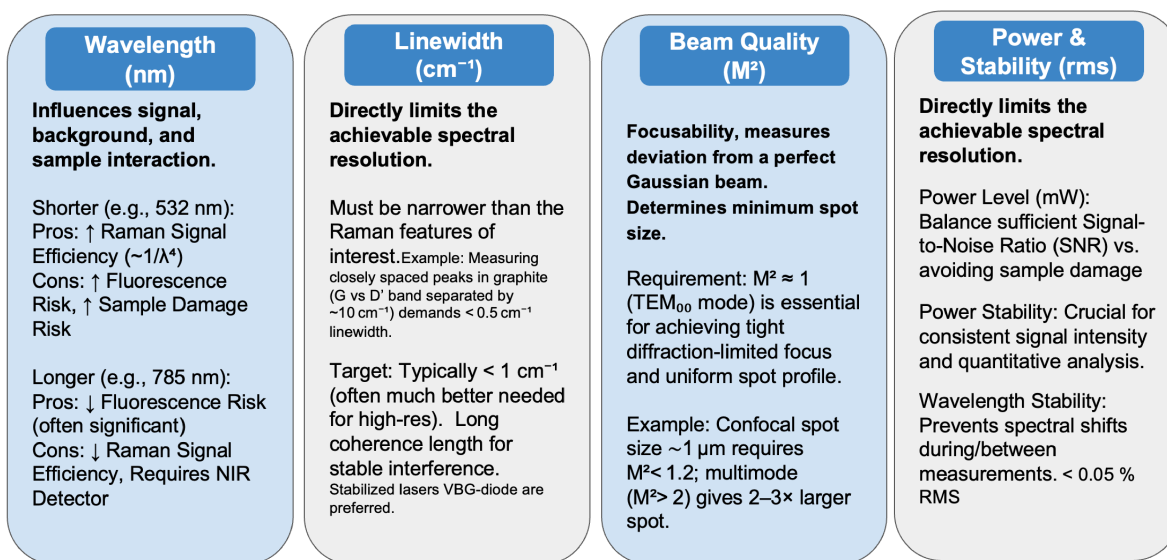


Figure 5.1: Laser and power-supply selection criteria for the local Raman platform. Wavelength sets the Raman-efficiency/background tradeoff, linewidth limits usable spectral resolution, beam quality controls confocal focus size, and power stability determines whether repeated spectra can be compared quantitatively.

### 5.3.3 Optical Architecture

A robust in-house confocal Raman layout can be understood as seven coupled subsystems. The excitation subsystem provides the 532 nm laser, cleanup filter, attenuator, shutter, power monitoring, beam steering, and beam dump. The microscope subsystem sends the beam through a dichroic or beamsplitter into the objective and collects the backscattered signal. The imaging subsystem splits a small fraction of the return path to a camera for white-light inspection and region selection. The filtering subsystem removes Rayleigh light with an edge or notch filter. The confocal subsystem places a pinhole, fiber core, or slit at an image plane so the spectrograph receives mostly light from the focal region. The spectrometer subsystem disperses the light and records it on a cooled detector. The motion subsystem provides x-y mapping and z focus control.

Díaz Barrero et al. provide the closest peer-reviewed reference point for this style of system. Their off-the-shelf confocal Raman microscope used 532 nm excitation, fiber coupling to the spectrometer, a wide-field imaging branch, motorized xyz translation, interchangeable pinholes, and micrometer-scale mapping with a 10x objective [188]. The local platform follows the same general instrument logic while using a build-specific mechanical and optical layout.




Figure 5.2: Component-level layout used to explain the local 532 nm confocal Raman path. The design separates the laser delivery path, microscope objective, dichroic or beam-splitting region, imaging branch, confocal aperture, Rayleigh-rejection filter, fiber coupling, and Raman spectrometer so each subsystem can be aligned and validated independently.




Figure 5.3: CAD rendering of the local confocal Raman architecture, showing the excitation source, microscope optics, collection path, and compact spectrometer connection.



Figure 5.4: Bench implementation of the local Raman optical path during development. The system uses cage-mounted optics so the excitation, imaging, filtering, and spectrometer-coupling paths can be modified and realigned during iteration.

### 5.3.4 Component Selection Criteria

The component choices are organized around a simple principle: every optical element should either increase useful Raman signal, reject unwanted light, preserve spatial registration, or make the instrument easier to align and reproduce. This is especially important for a local build because a commercial Raman microscope hides many of these choices inside factory-aligned modules.

Table 5.1: Design criteria for the major Raman microscope components.

Subsystem	Selection criterion	Reason in the local build
532 nm laser source	Narrow linewidth, good beam quality, stable power, controllable output, and compatible safety hardware	Maximizes usable Raman scattering while preserving a small confocal focus and repeatable spectra
Fiber delivery and collimator	Repeatable coupling, stable collimated beam, modular FC or SMA-style interfaces where useful	Decouples the laser head from the microscope path and makes the bench layout easier to rebuild
Microscope objective	Moderate magnification, adequate numerical aperture, usable working distance, and affordable replacement path	The first build uses a 10x objective because it is forgiving for alignment and sample handling while still supporting micrometer-scale mapping
Dichroic or beam splitter	Wavelength-matched separation of excitation and collected Raman/imaging light	The later design moves away from a simple 10:90 cube toward wavelength-selective splitting to improve collection efficiency and imaging quality
Camera branch	Visible field of view aligned to the laser focus and stage coordinates	Lets spectra and maps be tied to physical sample regions rather than to anonymous stage positions
Rayleigh-rejection filter	Edge or notch filter matched to the 532 nm laser with high optical density at the laser line	Protects the detector from elastic scatter and makes weak Stokes-shifted Raman bands measurable
Confocal aperture or fiber	Pinhole, fiber core, or slit placed at an image plane with x-y adjustment	Rejects out-of-focus light and defines the spatial sampling volume before the spectrometer
Spectrometer and detector	Grating, slit/fiber width, focal length, detector quantum efficiency, dark noise, and cooling selected for the desired Raman-shift range and resolution	Determines whether the instrument can separate relevant bands and whether weak signals can be measured without excessive smoothing
Mechanical frame	Cage rails, translation mounts, and optical breadboard mounting	Keeps the system reconfigurable while retaining enough rigidity for repeated alignment and mapping

### 5.3.5 Spectrometer and Detector Selection

The spectrometer is the measurement endpoint of the Raman path, so its specifications directly control data accuracy. A detector with attractive marketing sensitivity is not enough if the grating, slit, focal length, thermal stability, and calibration workflow do not match the Raman bands of interest. During development, the author spent considerable time attempting to revive unreliable used spectrometer hardware, but those efforts did not produce a reproducible calibrated data stream suitable for thesis-quality Raman measurements. That experience made the spectrometer subsystem a primary design problem rather than a peripheral accessory.

For 532 nm excitation, a silicon CCD or CMOS detector is a practical detector class because the relevant Stokes-shifted Raman wavelengths fall in the visible to near-infrared response range of silicon. Active thermoelectric cooling should be considered whenever long integration times or weak Raman signals are expected, because reduced dark current improves signal-to-noise ratio and lowers the risk that baseline subtraction hides real weak peaks. A narrow entrance slit or fiber image, often in the 10  $\mu\text{m}$  to 50  $\mu\text{m}$  class for research spectrographs, improves spectral resolution but reduces throughput. The grating, slit, focal length, detector pixel size, and array format therefore must be chosen as one coupled system rather than as independent catalog parts. These tradeoffs motivated the decision to build and document a local spectrometer path rather than depend on opaque or unstable second-hand hardware.

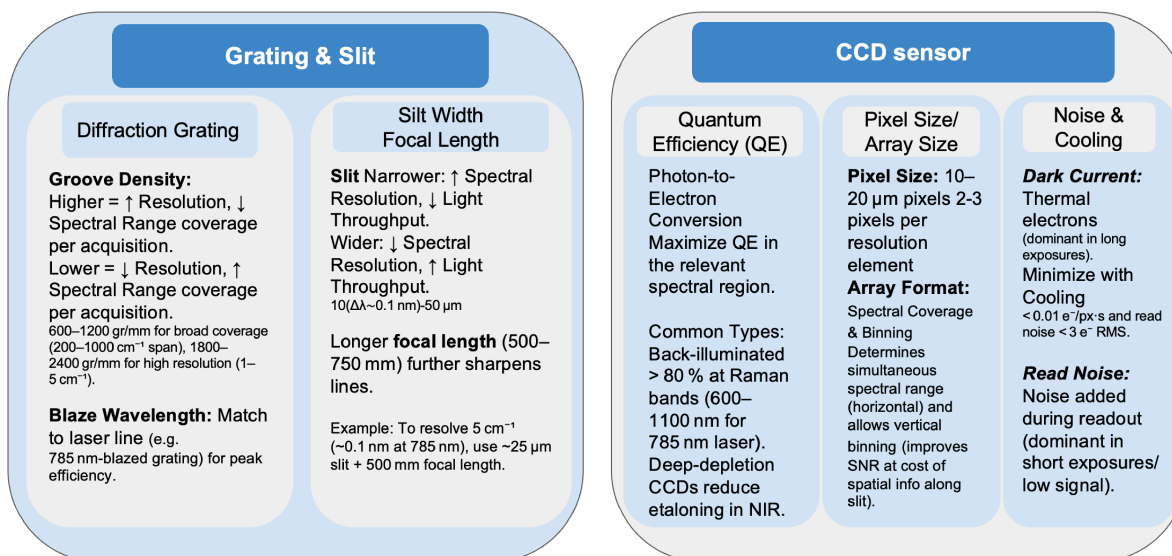


Figure 5.5: Spectrometer selection criteria for the Raman platform. Grating groove density and blaze wavelength set the resolution/coverage/efficiency tradeoff; slit width and focal length set line sharpness and throughput; detector quantum efficiency, pixel sampling, cooling, dark current, and read noise determine how weak Raman bands can be measured.

### 5.3.6 Raman Prototype Bill of Materials

The working bill of materials below comes from the original Raman system and is used here as the costed build lineage for the local instrument. The optical mounts, lenses, fiber hardware, objective, pinhole, and spectrometer entries are directly representative of the build strategy. Wavelength-specific entries in the original system used 785 nm labels; in the completed single-laser thesis system these must be replaced by the corresponding 532 nm laser, dichroic, cleanup, and edge-filter parts. Prices are included only as reference values from the 2025/2026 build period and should be refreshed against current quotes before any final

purchasing or cost claim.

Table 5.2: Costed Raman prototype BOM from the local original system. Prices are reference values from the 2025/2026 build period.

Item or model	Function in the Raman build	Qty.	Source cost
KC1T	Kinematic mounting hardware	1	\$110.18
AD8F	Fiber mount	1	\$33.56
F220FC	Fiber collimation or fiber optic focus package	1	\$181.00
CM1-DCH	Dichroic filter mount	1	\$192.44
Wavelength-matched dichroic beam splitter	Separates excitation from the collected Raman and imaging paths	1	\$660.00
CP33	Objective mount	1	\$19.09
MY10X-803	10x microscope objective	1	\$160.19
CCM1-4ER	Beam-splitter cube mount	1	\$155.62
BS026-10:90	Non-polarizing beamsplitter cube used in early imaging/collection tests	1	\$250.65
AC254-030-A	Achromatic doublet focusing lenses	2	\$96.22 each
SM1ZA	Z-axis translation mount	1	\$220.20
Wavelength-matched Raman edge filter	Rejects Rayleigh light before the confocal/spectrometer path	1	\$745.00
AC254-075-A	Achromatic doublet focusing lens	1	\$89.10
P75K	75 $\mu\text{m}$ pinhole for confocal spatial filtering	1	\$78.62
CXY1A	X-y translation mount for aperture or fiber alignment	1	\$201.47
M65L01	Optical fiber	1	\$141.37
SM1SMA	Fiber mount	1	\$33.61
532 nm DPSS laser source	Raman excitation source, substituted for the earlier wavelength-specific BOM line	1	\$995.00
Compact spectrometer	Disperses and records the filtered Raman signal	1	\$1,060.00
Total recorded prototype cost	Optomechanics, optics, excitation source, and compact spectrometer	–	\$5,519.54

### 5.3.7 Reference Off-the-Shelf Confocal Raman BOM

Díaz Barrero et al. are useful here because they do not only describe the concept of a custom confocal Raman microscope; they publish the kinds of components that make such a system reproducible. Their design used a 30 mm cage-based architecture, a 532 nm excitation path, a wide-field imaging branch, a confocal pinhole arm, fiber coupling to the spectrometer, and motorized xyz sample motion [188].

Table 5.3: Peer-reviewed off-the-shelf CRM component reference from Diaz Barrero et al.

Subsystem	Representative parts reported in the paper	Design lesson for the local build
Laser coupling	F260FC-A fixed-focus fiber collimator and single-mode fiber delivery	Fiber coupling can make the excitation path modular and helps isolate the laser source from the microscope body
Beam steering and cleanup	KC1-S/M steering mounts and LL01-532-12.5 laser-line filter	The laser path should be cleaned and adjustable before it reaches the objective
Excitation/collection separator	Di02-R532-25x36 single-edge dichroic beam splitter	A wavelength-selective dichroic is preferred over a generic splitter when Raman collection efficiency matters
Objective and sample focus	10x infinity-corrected objective used for the reported measurements	A moderate objective is a practical first configuration and can still support micrometer-scale maps
Power monitoring	SM1PD1A silicon photodiode power monitor	Power at the sample should be tracked rather than assumed from the laser controller
Wide-field imaging branch	BS025 10:90 beamsplitter, AC254-150-A-ML achromat, and CS165CU/M color CMOS camera	Imaging and spectroscopy should remain registered so spectra can be interpreted spatially
Confocal detection arm	LP03-532RU-25 edge filter, AC254-150-A-ML focusing lens, pinholes, and ST1XY-S/M x-y translator	Confocal rejection depends on both the aperture size and precise aperture placement
Spectrometer coupling	Custom circular-to-slit 48-fiber bundle into the spectrometer	Fiber geometry can be used to bridge the microscope image plane and the spectrograph entrance geometry
Motion system	Motorized xyz sample translation assembly	Raman maps require motion metadata, step-size validation, and focus control, not just point spectra

### 5.3.8 Build Workflow

The system was assembled as a sequence of testable modules rather than as one monolithic optical table. A practical build order is:

1. establish laser safety: enclosure, eyewear, beam height, beam dumps, shutter state, interlocks where possible, and a written operating procedure;
2. collimate and clean the 532 nm beam, then verify beam quality and power range before it reaches the

microscope;

3. align the dichroic/objective path with a low-power target and verify that the focused spot stays centered through objective changes or focus movement;
4. align the white-light camera branch so the camera image, stage coordinates, and laser focus refer to the same physical region;
5. insert Rayleigh rejection optics and verify that the spectrograph is not saturated by elastic scatter;
6. place and center the confocal aperture or fiber at the image plane using a strong Raman target before testing weak samples;
7. calibrate the spectrograph wavelength or Raman-shift axis with a known standard;
8. measure spatial resolution, confocal rejection, SNR, and repeatability before making material claims.

This workflow keeps failures local. If the Raman signal is weak, the likely causes can be traced to laser power, focus, objective NA, sample response, filter orientation, pinhole alignment, fiber coupling, spectrograph throughput, detector settings, or background subtraction.

### **5.3.9 Scan Control and Data Products**

A confocal Raman map is only useful if each spectrum knows where and how it was acquired. The control stack should log stage coordinate, objective, laser power, grating, slit or pinhole, exposure, accumulations, detector temperature, calibration file, background file, and software version. The map should also preserve the relationship between the white-light image and the stage coordinate frame. Stop-and-settle raster scanning is slow, but it is the safest first mode for a custom system because the stage position, focus, and detector exposure are easier to audit than in continuous scanning.

Confocal Raman microscopy is easy to overinterpret. Everall's review of confocal Raman artifacts is a useful warning: pinhole size, aberration, refractive-index mismatch, focus drift, detector saturation, and poor depth calibration can all produce convincing but misleading depth profiles or low-SNR maps [189]. For the local instrument, representative raw spectra, processed spectra, calibration residuals, and repeat maps are therefore more convincing than polished heat maps alone.

### **5.3.10 Spectrum Processing, Noise, and SNR**

During development, spectra from the OpenRAMAN project were used as external processing examples for propanol/isopropanol-type Raman data, not as proof of local instrument performance [190]. Their value in this chapter is that they make the data-processing path visible: raw spectra contain detector offset, baseline curvature, possible cosmic-ray spikes, and random noise before peaks are assigned.

The signal-to-noise ratio should be reported numerically rather than implied by a visually smooth line. In this workflow, signal can be defined as the height or integrated area of a diagnostic Raman band after baseline correction, and noise can be defined as the root-mean-square residual in a nearby feature-free spectral window. Dark subtraction and repeated accumulations reduce fixed detector offset and random noise. Isolated spikes are removed by median or outlier rejection across neighboring pixels or repeated accumulations. Broad fluorescence or instrument background is corrected with an asymmetric penalized least-squares baseline model, while Savitzky–Golay smoothing is used only for display or derivative estimates when its window is smaller than the peak spacing that must be preserved [191, 192]. Peak positions and SNR values should always be traced back to the unsmoothed, calibrated spectrum.

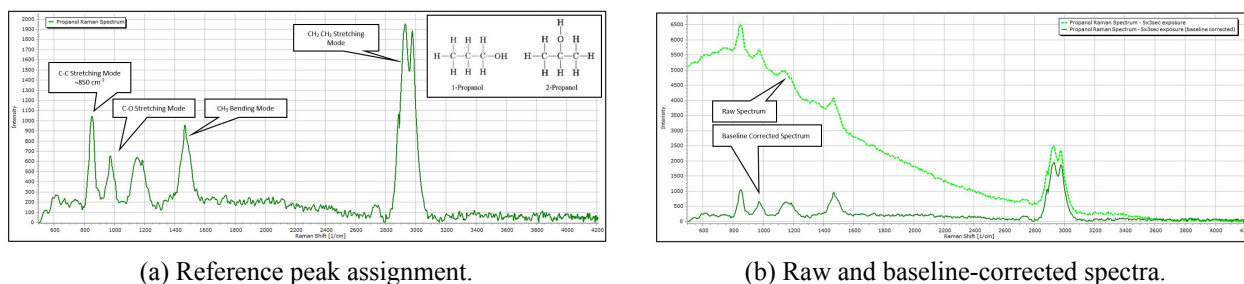


Figure 5.6: OpenRAMAN project spectra used as processing benchmarks during system development. The key quality metric is the SNR of assigned Raman bands after dark subtraction, spike removal, and baseline correction, with smoothing limited to display so that narrow vibrational features are not artificially broadened or shifted [190].

### 5.3.11 Calibration and Validation

Calibration is treated as a routine measurement, not a one-time setup step. ASTM E1840 provides a standards anchor for Raman shift calibration practice, and a silicon reference near 520.7/cm is a common day-to-day check [193]. Each Raman dataset should record the laser wavelength, objective, grating, slit or pinhole, integration time, accumulations, laser power at the sample, detector temperature, stage coordinates, calibration standard, and processing software version.

The minimum validation package for the local 532 nm confocal Raman system is therefore straightforward. Raman-shift calibration should report the standard used, the expected line position, the measured line position, the fit residual, and any correction applied; otherwise peak-shift claims can be calibration artifacts. Spectral resolution should be measured with a narrow-line source or a known standard under the actual grating and slit/fiber settings used for data acquisition. Lateral resolution should be measured with an edge, bead, patterned target, or small Raman-active feature scanned with a stated step size. Axial response should be measured by a z scan through a thin Raman-active layer or reflective interface so the confocal sectioning depth is known. Sample-plane power should be measured directly and repeated spectra should be acquired at increasing exposure to detect photodamage, thermal drift, or sample modification. SNR and repeatability should come from repeated spectra at the same point and repeated small maps of the same feature. Finally, the white-light image, laser focus, stage coordinates, and Raman map must be registered so a plotted map can be tied back to the physical sample.

## 5.4 Published Custom and Open Raman Systems

The useful comparison class for the local build is not only commercial Raman microscopy. Published custom systems and open Raman projects show how much can be built with standard optics, careful alignment, and honest performance reporting.

Díaz Barrero et al. provide the most direct peer-reviewed custom CRM reference for this chapter: a 532 nm confocal Raman microscope built mostly from off-the-shelf optomechanics, with a wide-field imaging arm, interchangeable pinholes, motorized xyz sample motion, and fiber coupling to a spectrometer [188]. The important lesson is not that a custom system automatically matches a commercial microscope; it is that a custom system can be credible when the optical layout, component choices, alignment procedure, spatial resolution, and test measurements are all reported.

AlShehab et al. show a second custom Raman path in which the construction of the microscope and its scientific use in plasmon-enhanced Raman measurements are presented together [194]. That model is useful because it links hardware construction to a sample question, but it also shows why calibration and sample

data must be included before the instrument can carry scientific claims. Mohr et al. demonstrate a lower-cost 532 nm Raman spectrometer for undergraduate and graduate experiments [195]; its design ambition is below confocal imaging, but it is stronger than a simple classroom spectroscope because it treats Raman excitation, filtering, spectrometer coupling, and teaching use as one coherent instrument.

OpenRAMAN is the main open-hardware Raman comparison point because it documents starter and performance versions, mechanical files, optics, enclosure choices, software, and community-facing build logic [190, 196]. AutoOpenRaman extends that direction by adding automation, scanning, and a Python GUI around an OpenRAMAN-style system [197]. These projects are valuable because they expose the whole stack: excitation, filtering, spectrometer design, mechanical packaging, acquisition software, data export, and operating limitations. They also make clear that open Raman is not just a cheap copy of a commercial tool; it is a documentation and learning ecosystem that can be modified, repaired, audited, and improved.

## 5.5 Commercial Raman Benchmarks and Robotic Integration

Commercial Raman systems set the benchmark for a mature instrument ecosystem: multiple lasers, automated calibration, confocal optics, high-throughput spectrographs, mapping stages, environmental control options, focus tracking, software pipelines, and service support. Official manufacturer specifications are useful for benchmarking, but they remain vendor claims rather than independent validation.

Renishaw's inVia family is a mature confocal Raman benchmark because it combines automated alignment/calibration options, multiple laser choices, high spectral resolution options, mapping stages, and Qontor focus tracking [198]. HORIBA's LabRAM Odyssey represents a similar high-end research platform with multi-laser and multi-grating configurations, fast imaging options, ultralow-frequency modules, and particle-analysis workflows [199]. Oxford Instruments' WITec systems emphasize high-speed confocal Raman imaging, modular microscopy integration, and future paths toward AFM or near-field coupling [200]. Thermo Scientific's DXR3xi shows the turnkey imaging side of the landscape: automated mapping, autofocus, confocal depth discrimination, and software-driven analysis for routine sample workflows [201]. Bruker SENTERRA II and Nanophoton/RAMANtouch show two additional commercial directions: automated microscope Raman on one side and high-speed line-illumination Raman imaging on the other [202, 203]. These systems make clear that the commercial state of the art is not only a laser and spectrometer; it is a tightly integrated optical, mechanical, calibration, software, and service stack.

Handheld and robotic Raman systems solve a different problem. Metrohm MIRA/XTR and IBEX are field and standoff tools for identification in settings where a sample may be dangerous, inaccessible, or operationally difficult to bring to a microscope [204, 205]. Robotic Raman integration is therefore best treated as a future extension of the spectroscopy ecosystem rather than as a replacement for the benchtop microscope. A mobile platform moves the technical problem away from diffraction-limited confocal imaging and toward pose control, standoff distance, laser safety, target selection, and confidence reporting. MIT News documented a 2026 hands-on Raman spectroscopy workshop in which a handheld Raman device was mounted to a robot dog to demonstrate remote material identification in dangerous environments such as crime scenes or toxic industrial sites [206]. That example captures the practical value of robotic Raman: it trades microscope-grade spatial resolution for access, distance, and operator safety.



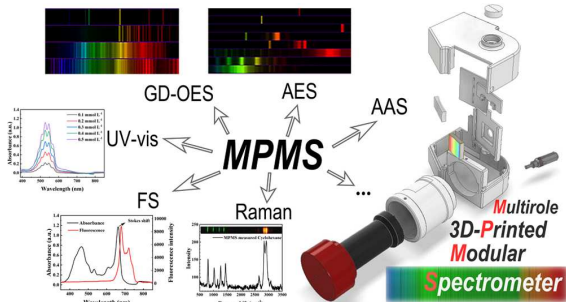
Figure 5.7: Robotic Raman integration concept for remote or hazardous-site sensing, reproduced from an MIT News report on a 2026 Raman spectroscopy workshop [206]. In this chapter it serves as a future-facing comparison to the local benchtop confocal Raman platform rather than as a completed local capability.

## 5.6 Low-Cost Personal Spectrometer


### 5.6.1 Collaborative Baseline

The personal spectrometer project follows a different purpose from the confocal Raman microscope. It makes spectroscopy buildable and understandable at low cost by exposing the entrance slit, collimation, grating dispersion, detector response, calibration, and software processing steps. The strongest peer-reviewed baseline for this part of the chapter is Huanyu Ren’s multirole 3D-printed modular spectrometer project, which we replicated in the lab and are now improving with more rigid mechanical design, better optical mounts, cleaner module interfaces, and a path toward more repeatable alignment. The published MPMS design supports UV–visible absorption, fluorescence, flame emission and absorption, glow discharge, and Raman demonstrations, with a reported approximate cost of \$600–\$700 and a practical working range of about 370 nm to 1000 nm under the described optical and camera constraints [207].

The thesis contribution is therefore not only replication; it is the movement from a teaching demonstrator toward a more mechanically disciplined open spectrometer platform. The failed attempt to extract reliable measurements from used spectrometer hardware is part of this motivation: it showed that an inexpensive instrument is only useful when the slit geometry, detector state, cooling, calibration, and raw-data workflow are visible and repeatable. Building the spectrometer locally makes those variables inspectable instead of leaving the quality of the Raman or fluorescence data dependent on undocumented hardware behavior.




(a) Multirole MPMS concept.



(b) Representative optical configurations.

Figure 5.8: Published MPMS images showing the baseline open spectrometer concept that we replicated and are improving in the lab [207].



(a) Ray-trace model of the CD-grating configuration.



(b) Printed monochromator and installed optics.

Figure 5.9: MPMS monochromator design, including the ray-traced geometry and the corresponding 3D-printed optical body. These figures support the collaborative low-cost spectrometer section: we replicated this concept in the lab and are improving the mechanical stiffness, mount repeatability, and alignment interfaces [207].

Table 5.4: MPMS reference BOM and module set.

Subsystem	Reported or replicated component choice	Role in the low-cost spectrometer
Printed chassis	Modular 3D-printed monochromator and plug-in experiment modules	Makes the optical geometry visible, low cost, and mechanically reproducible enough for teaching
Entrance slit	Razor-blade slit mounted at the monochromator entrance	Defines resolution and throughput while keeping the slit inexpensive and replaceable
Dispersive element	CD grating, DVD grating, or blazed grating; reported examples include CD at 625 lines/mm and DVD at 1350 lines/mm	Allows low-cost comparison between wavelength range, dispersion, and efficiency
Collimation optic	Concave mirror with reported focal length of 70 mm	Collimates light from the slit before grating dispersion in the compact folded path
Focusing/detector lens	CCTV lens in the reported 5 mm to 50 mm class	Focuses the dispersed spectrum onto an array detector while remaining inexpensive and adjustable
Detector	Astronomical CMOS camera with adjustable integration, gain, raw 16-bit output, and usable response toward the NIR	Gives better raw-data access and longer exposure control than a normal phone or webcam camera
UV-vis and fluorescence modules	Multipurpose light source, cuvette holder, broadband tungsten-halogen source, and LED excitation options	Supports absorption and fluorescence teaching experiments with the same monochromator body
Atomic spectroscopy modules	Flame atomizer, ultrasonic nebulizer, low-pressure sodium lamp, Tesla-coil glow discharge fixtures, and gas discharge tubes	Demonstrates emission and absorption spectra without a commercial atomic spectrometer
Raman module	532 nm or 660 nm laser, long-pass filters near 550 nm or 670 nm, dark chamber, and cuvette scattering geometry	Enables qualitative Raman demonstrations while keeping the optical path simpler than a backscattering confocal microscope
Calibration and software	Hg-Ar or fluorescent-lamp lines, Fraunhofer lines where appropriate, raw image processing, and wavelength calibration	Converts camera pixels into wavelength or Raman-shift axes and exposes calibration error to the user

Subsystem	Reported or replicated component choice	Role in the low-cost spectrometer
Reported cost	Approximately \$600–\$700 for the MPMS materials and modules	Establishes the access-oriented cost target for the personal spectrometer effort

## 5.6.2 Design Choices

A standard grating spectrometer has an entrance slit, a collimating optic, a dispersive grating, a focusing optic, a detector or camera, an enclosure, and a wavelength calibration workflow. Once those subsystems are visible, students and builders can directly see why slit width changes resolution and throughput, why grating angle changes wavelength range, why stray light matters, and why a calibration curve is more than a cosmetic plot.

The major design choices are:

- entrance slit width, which trades spectral resolution against throughput;
- grating groove density and blaze wavelength, which set dispersion and efficiency;
- focal length and detector size, which set wavelength range and sampling;
- detector type, which controls sensitivity, noise, exposure range, and raw-data access;
- enclosure and baffling, which control stray light and mechanical stability;
- calibration source, which determines how well pixel position can be converted to wavelength;
- software, which must handle dark frames, calibration curves, response correction, and exportable raw data.

## 5.6.3 Calibration and Validation

The personal spectrometer is only useful if its limitations are measured. A two-line calibration can work for a narrow educational demonstration, but a broader visible-range instrument usually needs a polynomial calibration because camera lenses, grating geometry, and off-axis imaging introduce curvature. Mercury, argon, neon, compact fluorescent, or other line sources can be used if the line identities are known. Fraunhofer lines can help in solar spectroscopy, but they should be used carefully because the sunlight path, exposure, and line identification matter.

For Raman demonstrations, the wavelength axis must be converted to Raman shift:

$$\Delta\tilde{\nu} = 10^7 \left( \frac{1}{\lambda_0} - \frac{1}{\lambda_s} \right),$$

where  $\lambda_0$  and  $\lambda_s$  are in nanometers and  $\Delta\tilde{\nu}$  is in  $1/\text{cm}$ . This conversion makes calibration error more serious because a small wavelength error can become a meaningful Raman-shift error.

The validation ladder for the personal spectrometer starts with assembly evidence: a labeled photograph, slit/grating/detector geometry, BOM, enclosure notes, and a record of any module changes. Wavelength calibration then requires a line-source spectrum, identified reference lines, a calibration fit, and a residual plot; without the residual plot, the spectrum may look convincing while still carrying unacceptable wavelength

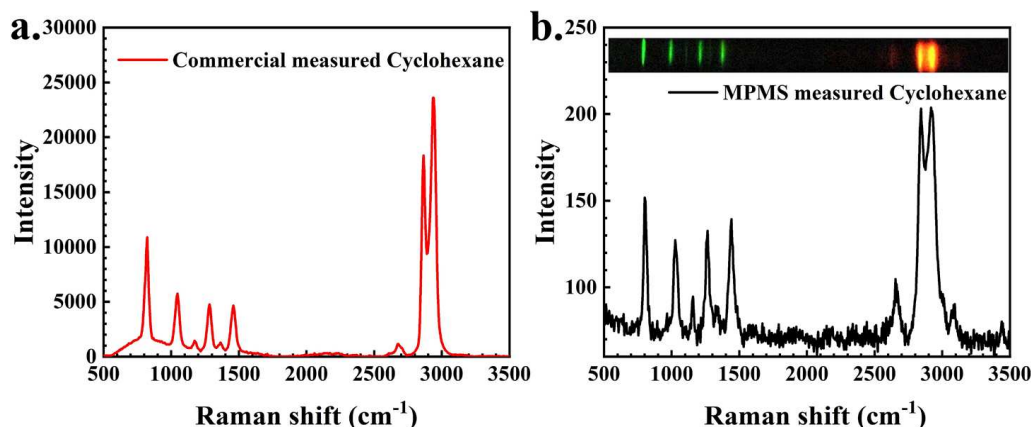


Figure 5.10: Cyclohexane Raman comparison from the MPMS project, showing how the modular spectrometer was evaluated against a commercial measurement while remaining an educational/access instrument [207].

error. Resolution should be measured with a narrow line or line-pair separation at the actual slit width and grating used for the measurement. Intensity behavior should be checked with dark frames, blank/reference spectra, exposure linearity, and the saturation limit of the detector. Stray-light testing should include a covered-source measurement, a blocked-beam measurement, and an off-band leakage check under normal room-light conditions. Repeatability should be measured by removing and replacing the sample or module and then reacquiring the calibration or reference spectrum. Safety evidence is also part of validation when the module uses lasers, flame, high voltage, or chemical samples, because an accessible instrument is not useful if it cannot be operated responsibly.

## 5.7 Commercial Spectrometer Benchmarks

Commercial spectrometers span a wide range: compact OEM modules, cooled high-sensitivity array spectrometers, high-speed OCT spectrometers, astronomy spectrographs, and complete UV–visible instruments. For a personal spectrometer chapter, the most useful comparison is not the best possible spectrometer, but the feature set that mature commercial designs use to solve the problems that self-built instruments expose.

The 2025/2026 commercial landscape is useful because each product class solves a different instrument problem. Ocean Insight QE Pro and related HR-class array spectrometers emphasize cooled, high-sensitivity modular detection for fluorescence, low-light spectroscopy, and Raman-adjacent measurements, so they are a benchmark for detector cooling, dynamic range, stray-light control, and integrated software [208]. Avantes AvaSpec-HERO represents the mature compact-lab category: broad wavelength configuration options, high signal-to-noise, high dynamic range, and cooled detection in a packaged system [209]. Ibsen Photonics FREEDOM shows the OEM design path, where the optical train is miniaturized into reproducible grating modules with application-specific wavelength and resolution choices [210].

At the other size extreme, Hamamatsu’s C12880MA MEMS micro-spectrometer integrates optics and detector into a fingertip-scale visible/NIR module, making it attractive for embedded sensing even though its resolution and optical flexibility are limited compared with larger spectrographs [211]. Wasatch Photonics Cobra-S is a reminder that some state-of-the-art spectrometers are optimized for speed and system integration, especially OCT, rather than for low-cost teaching use [212]. Shelyak Sol’Ex/Star’Ex and eShel occupy a particularly relevant middle ground: they connect self-build astronomy spectroscopy, calibration discipline, and high-resolution pro-am practice, showing that serious non-industrial spectroscopy can be built around careful documentation rather than sealed black-box hardware [213, 214].

## 5.8 Open and Self-Build Spectrometer Landscape

The open and self-build spectrometer landscape is broad. Some projects are true open-source hardware, some are published teaching instruments, and some are well-documented hobbyist or pro-am systems. The categories need to be labeled honestly because open documentation, peer-reviewed validation, and commercial support are different forms of evidence.

Several examples are especially useful for positioning the personal spectrometer work:

- **MPMS modular spectrometer:** Ren's peer-reviewed 3D-printed modular spectrometer is the main technical reference for the personal spectrometer direction because it combines absorption, fluorescence, flame emission/absorption, glow-discharge spectroscopy, and Raman demonstrations in one documented low-cost platform [207]. Its value is not only low cost; it exposes slit geometry, grating choice, detector response, calibration, and experiment modules as teachable subsystems.
- **PhysicsOpenLab webcam spectrometer:** the PhysicsOpenLab project is useful as practical documentation for very low-cost spectrometer construction using a webcam-class detector [215]. It belongs in the educational-access landscape, but it should be treated as project documentation rather than as peer-reviewed validation.
- **Shelyak Sol'Ex/Star'Ex:** the Shelyak ecosystem shows that non-industrial spectroscopy can be technically serious when the optical layout, calibration procedure, and user documentation are mature [213]. It is especially relevant as a model for build documentation, repeatable alignment, and community improvement around astronomy spectroscopy.
- **Open-source fluorescence spectrometers:** peer-reviewed open fluorescence instruments demonstrate how files, cost, calibration, and educational use can be published together rather than hidden behind a finished black box [216]. This is the documentation standard that makes an accessible instrument scientifically useful.
- **Low-cost teaching spectrophotometers:** the Viking spectrophotometer and related low-cost teaching instruments show that inexpensive optical systems can still report calibration, resolution limits, and validation data when the measurement goal is clearly scoped [217].
- **Smartphone and 3D-printable spectrometers:** smartphone-adaptor and 3D-printable spectrophotometer papers define the bottom end of the access ladder [218, 219]. They are valuable for teaching and outreach, but their calibration stability, stray-light rejection, detector linearity, and spectral resolution must be stated explicitly before they are used for quantitative claims.

## 5.9 Instrument-Building Continuum

The confocal Raman setup and the personal spectrometer are a continuum, not competitors. The personal spectrometer teaches the optical grammar: slit, collimation, dispersion, detector response, calibration, stray light, and raw data. The confocal Raman microscope turns that grammar into a research instrument by adding laser safety, microscope focusing, Rayleigh rejection, confocal spatial filtering, precision motion, and map registration. Commercial systems set the integrated performance benchmark, while open-source and self-build projects define the documentation, repairability, and cost-transparency benchmark.

## 5.10 Detailed Build and Validation Protocols

### 5.10.1 Raman Build Record

The Raman build record links spectra and maps to the physical instrument state that produced them. This record makes later comparison possible when optics, stages, filters, or software are changed.

For the laser source, the record should state wavelength, model, output range, linewidth if known, shutter and attenuator state, cleanup filter, beam dumps, and eyewear rating, followed by measured sample-plane power as a function of control setting. The beam-delivery record should identify mirrors, cage rails, dichroic, objective, focusing geometry, beam diameter, and the alignment target used to center the focus and return path. The imaging branch should record illumination, beamsplitter, tube lens, camera model, pixel size, magnification, and field of view, with a calibration target or grid image that converts camera pixels to micrometers. The Raman filtering record should give the edge or notch filter model, orientation, cutoff, and blocking specification when available, then show that Rayleigh light does not saturate the detector under the chosen acquisition settings. The confocal record should identify the pinhole, fiber core, or slit dimension, the aperture mount, and the adjustment method, with an axial-response or thin-layer z-scan used as evidence. The spectrograph and detector record should include grating, slit or fiber input, detector model, cooling state, gain, readout mode, exposure range, calibration spectrum, spectral resolution, dark noise, and cosmic-ray handling. The motion record should give coarse and fine stage axes, travel, step size, controller, software link, repeatability test, map registration, and settling-time estimate. Finally, the software record should preserve the acquisition script or package, file format, metadata fields, processing pipeline, example raw file, metadata export, and version number.

### 5.10.2 Raman Alignment Checklist

The alignment checklist is run whenever the optical path is rebuilt, a filter is swapped, the objective changes, or the instrument has been idle long enough for drift to matter.

The operating checklist is short enough to keep in prose. First, confirm eyewear, enclosure state, shutter state, warning sign, beam dumps, and the lowest practical laser-power setting; the pass condition is that the laser can be enabled without exposed uncontrolled beam paths. Second, check collimation and beam centering through the excitation path; the beam should remain centered on the steering apertures and objective back aperture. Third, focus onto a benign alignment target and record sample-plane power; the focus should be visible, stable, and below the chosen power limit. Fourth, align the white-light camera to the same region as the laser focus; the camera crosshair or marked pixel should correspond to the laser focus within the stated registration error. Fifth, insert Raman filters and verify Rayleigh suppression; the detector should not saturate under the target and exposure settings. Sixth, maximize signal through the pinhole, fiber, or slit with a strong Raman target; the signal should rise smoothly with aperture alignment and fall when deliberately misaligned. Seventh, acquire a calibration spectrum and fit the reference peak or lines; the measured offset must remain within the acceptance limit for that dataset. Eighth, run a short repeat spectrum and a small repeat map; the spectra and maps should agree within expected noise and registration limits before any sample claim is made.

### 5.10.3 Raman Acquisition Template

Every Raman dataset used as thesis evidence has a companion record with the following fields:

- sample ID, sample description, preparation state, operator, date, and purpose of scan;
- laser wavelength, measured sample-plane power, objective, numerical aperture if known, grating, slit/pinhole/fiber setting, detector temperature, gain, exposure, and accumulations;

- calibration standard, expected line position, measured line position, residual or correction, and calibration file;
- white-light image file, stage coordinate origin, scan size, step size, z position, focus strategy, scan direction, dwell time, and settle time;
- raw spectrum file, dark/background file, cosmic-ray removal method, baseline method, smoothing if any, and peak-fit model if used;
- plotted outputs and the exact script or software version used to produce them.

This record is intentionally heavier than a quick lab note. A custom instrument becomes credible when the data can be regenerated and questioned later.

#### **5.10.4 Personal Spectrometer Build and Validation Sequence**

The personal spectrometer build record is folded into the validation sequence so the document emphasizes measured evidence rather than a separate checklist. Before reporting spectra, the build should identify the slit material and width, collimating optic, grating groove density and angle, focusing optic, detector model, enclosure or baffles, calibration source, and software path from raw image to spectrum. Each item should be tied to evidence: slit width to resolution and throughput, grating geometry to wavelength range, detector settings to dark frame and saturation behavior, enclosure design to stray-light tests, calibration source to residual plots, and software to reproducible raw-data processing.

A low-cost spectrometer is therefore validated in a sequence that makes failure modes visible:

1. record a dark frame and a covered-detector frame to estimate detector offset and room-light leakage;
2. record a calibration lamp and identify several lines across the intended wavelength range;
3. fit pixel position to wavelength and plot residuals rather than only showing the fitted spectrum;
4. measure a narrow line or line pair to estimate spectral resolution for the chosen slit and grating;
5. repeat the calibration after moving or reopening the enclosure to estimate mechanical repeatability;
6. measure a known sample or lamp and compare the peak positions to a reference spectrum;
7. for intensity work, record blank/reference spectra and avoid saturated pixels;
8. for Raman or laser demonstrations, document laser safety controls and optical filtering before acquiring spectra.

#### **5.10.5 Chapter Summary**

The Raman microscope and the personal spectrometer are related because they sit on the same optical foundation. Both depend on controlled light delivery, defined apertures, wavelength-selective optics, detector response, calibration, dark subtraction, stray-light control, and metadata-rich acquisition. The personal spectrometer exposes these ideas in their simplest form: a slit defines the source image, a grating disperses the spectrum, a detector records intensity as a function of position, and calibration turns pixels into physical wavelengths. The confocal Raman platform builds on the same grammar but adds the requirements of a research instrument: laser safety, microscope focusing, Rayleigh rejection, confocal spatial filtering, motion control, Raman-shift calibration, SNR reporting, and spatial registration.

Together, these projects define a practical path from accessible spectroscopy education to local research capability. The low-cost spectrometer lowers the entry barrier for students, small laboratories, and resource-constrained groups that need to learn how spectra are produced rather than simply read numbers from a closed instrument. The Raman platform shows how the same open-instrument mindset can move toward more demanding measurements when the optical path, build record, bill of materials, alignment procedure, calibration workflow, and processing pipeline are documented. This does not mean that an open or local instrument automatically replaces a commercial system. It means that the knowledge needed to build, repair, adapt, and validate a scientific measurement tool becomes more widely available.

That distinction is important for people and institutions that need locally made, low-cost scientific hardware. In many settings, the limiting factor is not only the price of a commercial spectrometer or Raman microscope; it is the inability to repair the system, change the optical configuration, inspect the software pipeline, replace a failed part quickly, or adapt the instrument to a local sample, classroom, field site, or fabrication process. A documented open spectrometer can support teaching, calibration training, fluorescence and emission demonstrations, and selected validated measurements. A locally built Raman system can support material identification, process feedback, and spatially resolved spectroscopy once its calibration, resolution, confocal response, SNR, repeatability, and safety controls are measured. The broader contribution of this chapter is therefore a spectroscopy-building framework: start with transparent optical subsystems, validate each claim against measured evidence, keep cost and repairability visible, and make the instrument understandable enough that other people can reproduce, improve, and use it for their own scientific problems.

The 532 nm Raman project benefited from technical assistance, build feedback, and experimental discussion with Max Aalto, Koseki J. Kobayashi-Kirschvink, Marcello Tania, and Nikhil Lal.

## Chapter 6

# Synthesis and Future Work

### 6.1 Synthesis

This thesis argues that micro-manufacturing for scientific instrumentation should be treated as a coupled laboratory capability rather than as a single machine or process. The central problem is not only whether a feature can be made once. The harder requirement is to connect material selection, motion control, laser processing, characterization, data management, and iteration tightly enough that a laboratory can design and improve micron-scale parts for its own scientific questions.

The diamond MAS-NMR rotor is the strongest case study because it forces all parts of that loop to meet at one demanding object. The rotor must be made from a material that survives high hoop stress, laser-induced damage, cap insertion, gas-bearing operation, and cryogenic or DNP-relevant experimental conditions. Its geometry is sub-millimeter, but its failure modes are system-level: a rotor can fail because of a crystal defect, a laser-machining flaw, an ID/OD eccentricity error, a cap or drive-tip problem, a bearing-flow instability, or a characterization gap. The rotor therefore makes the thesis concrete. It shows why micro-manufacturing for science is not merely small-scale fabrication, but fabrication with feedback from the intended instrument.

The supporting chapters develop the same argument from different parts of the stack. The precision-motion chapter connects stage architecture, piezo actuation, encoder feedback, and mechanical design to the repeatability needed for fabrication and measurement. The laser-machining chapter separates demonstrated local machine-building from literature-based process directions, and shows why femtosecond and waterjet-guided laser platforms are best understood as complementary tools with different damage, taper, throughput, and feedback requirements. The spectroscopy chapter extends the loop into characterization: Raman microscopy and a custom spectrometer are not side projects, but examples of making the measurement tools needed to verify materials and fabrication outcomes.

### 6.2 Contributions

The first contribution is a system-level framing for micro-manufacturing in scientific laboratories. The thesis links the Center for Bits and Atoms tradition of accessible fabrication to a stricter instrumentation standard: open or in-house tools are useful only when they can produce measured, repeatable, scientifically relevant parts.

The second contribution is the organization of the diamond-rotor workflow as a manufacturing ecosystem. Material choice, HPHT/CVD defect structure, crystallographic orientation, laser wavelength, pulse duration, rotary fixturing, micro-CT, SEM, Raman, THz/VNA measurements, stators, gas bearings, drive tips, and helium operation are treated as coupled design variables rather than isolated topics.

The third contribution is the integration of supporting machine-building projects into that ecosystem. The piezo stage, femtosecond platform, waterjet-guided laser platform, Raman microscope, and custom spectrometer show how a laboratory can begin to build its own fabrication and characterization infrastructure around a difficult scientific part. Some of these systems are still at the design or early validation stage, but they establish the practical roadmap: make the machine, measure the machine, measure the part, and use the result to revise both.

### **6.3 Limitations**

The maturity of the evidence is uneven, and the thesis should preserve that distinction. The diamond-rotor workflow has the strongest connection to published rotor results and local process evidence. The motion, laser, Raman, spectrometer, and automation systems provide important supporting infrastructure, but they require additional quantitative validation before they can support final performance claims.

The most important missing data are not exotic. They are the measurements that make instruments trustworthy: stage repeatability, hysteresis, drift, calibration residuals, delivered pulse energy, coupled waterjet power, beam quality, feature geometry, taper, heat-affected or graphitized material, spectral resolution, signal-to-noise ratio, detector stability, and repeated success/failure statistics. Without those records, a build remains a promising apparatus rather than a validated scientific instrument.

A further limitation is the uneven provenance record for figures and references across the assembled manuscript. Images derived from literature, patents, collaborators, internal project records, or redrawn schematics require clear attribution and should be distinguished from original experimental data. Figures that are only illustrative should be replaced with high-resolution, properly sourced versions when they support technical claims, or removed when they do not. The same standard applies to references: citations should support the exact claim being made, not only the general topic.

### **6.4 Implications**

The broader implication is that many scientific groups need intermediate manufacturing capability. Industrial suppliers are essential, but they are often too slow, expensive, or specialized for early-stage instrument development. Hobby-grade fabrication tools are accessible, but usually lack the precision, materials, metrology, and documentation required for advanced scientific hardware. The opportunity is the space between those extremes: laboratory-scale micro-manufacturing systems that can make and measure parts with enough rigor to accelerate research.

For MAS-NMR, this means a laboratory could iterate rotor materials, bore geometry, stators, caps, gas bearings, and inspection methods without treating each change as a full industrial procurement cycle. For laser processing, it means that femtosecond and waterjet-guided platforms can be developed around local process feedback instead of used as black boxes. For spectroscopy and Raman characterization, it means that measurement tools can be designed to serve the fabrication loop directly.

### **6.5 Roadmap**

The immediate roadmap is to finish the evidence chain for each subsystem. For diamond rotors, the next step is a structured build log linking stock identity, growth method, orientation, machining parameters, CT/SEM/Raman results, cap fit, and spin-test outcome. For precision motion, the next step is repeated calibration under realistic loads and synchronized triggering with the relevant laser or imaging system. For the femtosecond and waterjet-guided laser platforms, the next step is a measured process map that connects

optical settings to feature geometry and damage metrics. For Raman and the custom spectrometer, the next step is calibration against known standards and quantified signal stability.

The longer-term roadmap is an open, inspectable micro-manufacturing workflow for scientific instrumentation. In that workflow, a researcher should be able to define a micron-scale part, fabricate it with documented machine state, inspect it with local metrology, update the design from the measurement, and share enough data that another laboratory can reproduce the process or improve it. That is the scientific version of democratized fabrication: not only access to machines, but access to the measured capability to make parts that expand what an experiment can do.

# Appendix A: Lessons Learned

One of my favorite sections in Will Langford's thesis was his lessons learned section [220]. Inspired by that example, this appendix summarizes several lessons from my time at MIT: practical habits, judgment calls, and research patterns that shaped the work in this thesis.

## The Value of Systematic Learning

Quick literature searches are useful for orientation, but they are insufficient for complex research problems. Books, review articles, and even the table of contents of a well-structured textbook can provide the scaffolding needed to understand a field from its foundational assumptions. When a project crosses optics, mechanics, materials, electronics, and measurement, systematic learning prevents the work from becoming a collection of disconnected tricks.

## Economy of Scale in Design

Custom parts are often slow and expensive to make. Whenever possible, a design should adapt mass-produced commercial off-the-shelf components before committing to fully custom hardware. COTS components are not automatically ideal, but they make a system more affordable, reproducible, repairable, and accessible to everyone else.

## The Buy vs. Build Dilemma

MIT's maker culture encourages building from first principles, and that culture is powerful. However, time is finite. The decision to build should be weighed against the opportunity cost of not using an available tool, module, or service. Choosing what not to build is as important as choosing what to build.

## Ruthless Prioritization

It is not possible to finish every promising task. Neil often notes that only a fraction of a daily to-do list will actually be completed. Research progress depends on isolating the critical path, choosing battles carefully, and letting lower-priority tasks remain unfinished when they do not change the central result.

## The Necessity of Seeking Help

Research is rarely solitary. When a project stalls, asking for help is often the fastest path forward. People in this community are generally willing to help, but they are also busy. Unanswered emails, delayed replies, and scheduling friction should not be interpreted as rejection. Persistence, specificity, and gratitude matter.

## Cultivating Perspective

Research can become all-consuming. Graduate school occupies years of life that should not be reduced entirely to laboratory output. Time outside the lab is not a distraction from good research; it is part of maintaining judgment, curiosity, and well-being over the duration of a difficult project.

## **Balancing Theory and Execution**

Action and reflection must feed one another. Without clear direction, physical effort can become motion without progress. Without execution, theory remains inert. The most useful project rhythm is iterative: model enough to act, build enough to learn, measure enough to correct, and then repeat.

## **Rigorous Documentation**

Documenting work early and often is essential. Photographs, videos, sketches, measurements, filenames, and short notes from unfinished or failed attempts often become the artifacts that best explain a project's trajectory. Documentation is most valuable when it is collected during the work, not reconstructed after the result is already known.

## **Process vs. Results**

It is easy to spend unlimited time optimizing workflows, tooling, and documentation systems. Good process improves reproducibility, but excessive process development can become a way to avoid finishing the work itself. The process should serve the result, not replace it.

# References

- [1] Natalie C. Golota, Zachary P. Fredin, Daniel P. Banks, David Preiss, Salima Bahri, Prashant Patil, William K. Langford, Camron L. Blackburn, Erik Strand, Brian Michael, Blake Dastrup, Keith A. Nelson, Neil Gershfeld, and Robert G. Griffin. Diamond rotors. *Journal of Magnetic Resonance*, 352:107475, 2023. doi: 10.1016/j.jmr.2023.107475. URL <https://doi.org/10.1016/j.jmr.2023.107475>.
- [2] Lauren Schaffer, David Preiss, Ravi Shankar Palani, Nicholas Wiesner, Jiaming Liu, Samuel Strymish, Salima Bahri, Sara Linse, Neil Gershfeld, and Robert G. Griffin. Diamond rotors for high magic angle spinning frequencies. *Journal of Magnetic Resonance*, 379:107909, 2025. doi: 10.1016/j.jmr.2025.107909. URL <https://doi.org/10.1016/j.jmr.2025.107909>.
- [3] Marta Bonaccorsi, Tanguy Le Marchand, and Guido Pintacuda. Protein structural dynamics by magic-angle spinning NMR. *Current Opinion in Structural Biology*, 70:34–43, 2021. doi: 10.1016/j.sbi.2021.02.008. URL <https://doi.org/10.1016/j.sbi.2021.02.008>.
- [4] Yusuke Nishiyama, Guangjin Hou, Vipin Agarwal, Yongchao Su, and Ayyalusamy Ramamoorthy. Ultrafast magic angle spinning solid-state NMR spectroscopy: Advances in methodology and applications. *Chemical Reviews*, 123(3):918–988, 2023. doi: 10.1021/acs.chemrev.2c00197. URL <https://doi.org/10.1021/acs.chemrev.2c00197>.
- [5] E. R. Andrew, A. Bradbury, and R. G. Eades. Nuclear magnetic resonance spectra from a crystal rotated at high speed. *Nature*, 182:1659, 1958. doi: 10.1038/1821659a0. URL <https://doi.org/10.1038/1821659a0>.
- [6] I. J. Lowe. Free induction decays of rotating solids. *Physical Review Letters*, 2:285–287, 1959. doi: 10.1103/PhysRevLett.2.285. URL <https://doi.org/10.1103/PhysRevLett.2.285>.
- [7] Kai Xue, Riddhiman Sarkar, Carina Motz, Sam Asami, Venita Decker, Sebastian Wegner, Zdenek Tosner, and Bernd Reif. Magic-angle spinning frequencies beyond 300 khz are necessary to yield maximum sensitivity in selectively methyl protonated protein samples in solid-state NMR. *The Journal of Physical Chemistry C*, 122(28):16437–16442, 2018. doi: 10.1021/acs.jpcc.8b05600. URL <https://doi.org/10.1021/acs.jpcc.8b05600>.
- [8] Ivan V. Sergeev, Keith Fritzsche, Rivkah Rogawski, and Ann McDermott. Resolution in cryogenic solid state NMR: Challenges and solutions. *Protein Science*, 33(7):e4803, 2024. doi: 10.1002/pro.4803. URL <https://doi.org/10.1002/pro.4803>.
- [9] Adam N. Smith and Joanna R. Long. Dynamic nuclear polarization as an enabling technology for solid state nuclear magnetic resonance spectroscopy. *Analytical Chemistry*, 88(1):122–132, 2016. doi: 10.1021/acs.analchem.5b04376. URL <https://doi.org/10.1021/acs.analchem.5b04376>.

- [10] Ravi Shankar Palani, Michael Mardini, Yifan Quan, Yifu Ouyang, Aditya Mishra, and Robert G. Griffin. Dynamic nuclear polarization with P1 centers in diamond. *The Journal of Physical Chemistry Letters*, 15(46):11504–11509, 2024. doi: 10.1021/acs.jpcllett.4c02612. URL <https://doi.org/10.1021/acs.jpcllett.4c02612>.
- [11] Engineering ToolBox. Speed of sound: Equations, 2026. URL [https://www.engineeringtoolbox.com/speed-sound-d\\_82.html](https://www.engineeringtoolbox.com/speed-sound-d_82.html). Formula and reference values for acoustic velocity in ideal gases; accessed 2026-05-11.
- [12] Michael R. Moldover, Roberto M. Gavioso, James B. Mehl, Laurent Pitre, Michael de Podesta, and Jintao Zhang. Acoustic gas thermometry, 2014. URL <https://www.nist.gov/publications/acoustic-gas-thermometry>. NIST publication page; used for the dilute-gas sound-speed relation.
- [13] Dirk Wilhelm, Armin Porea, and Frank Engelke. Fluid flow dynamics in MAS systems. *Journal of Magnetic Resonance*, 257:51–63, 2015. doi: 10.1016/j.jmr.2015.05.006. URL <https://doi.org/10.1016/j.jmr.2015.05.006>.
- [14] Qi Liang, Chih-Shiue Yan, Yufei Meng, Joseph Lai, Szczesny Krasnicki, Ho-Kwang Mao, and Russell J. Hemley. Enhancing the mechanical properties of single-crystal CVD diamond. *Journal of Physics: Condensed Matter*, 21(36):364215, 2009. doi: 10.1088/0953-8984/21/36/364215. URL <https://doi.org/10.1088/0953-8984/21/36/364215>.
- [15] Gemological Institute of America. Digging into diamond types, 2014. URL <https://4cs.gia.edu/en-us/blog/digging-diamond-types/>. Accessed 2026-05-09.
- [16] Francis P. Bundy. The P,T phase and reaction diagram for elemental carbon, 1979. *Journal of Geophysical Research: Solid Earth*, 85(B12):6930–6936, 1980. doi: 10.1029/JB085iB12p06930. URL <https://doi.org/10.1029/JB085iB12p06930>.
- [17] Francis P. Bundy, William A. Bassett, Michael S. Weathers, Russell J. Hemley, Ho-kwang Mao, and Alexander F. Goncharov. The pressure-temperature phase and transformation diagram for carbon; updated through 1994. *Carbon*, 34(2):141–153, 1996. doi: 10.1016/0008-6223(96)00170-4. URL [https://doi.org/10.1016/0008-6223\(96\)00170-4](https://doi.org/10.1016/0008-6223(96)00170-4).
- [18] Francis P. Bundy. Diamond synthesis with non-conventional catalyst-solvents. *Nature*, 241:116–118, 1973. doi: 10.1038/241116a0. URL <https://doi.org/10.1038/241116a0>.
- [19] Kun Luo, Biao Liu, Jiecheng Yang, Shangzhou Zhang, Qiangwei Li, Zhisheng Wu, Fei Zeng, Ye Tian, Bingbing Liu, Sijie Zhang, Xueyong Zhou, Quan Huang, Yusheng Zhao, Qiaoshi Zeng, Jian Sun, and Zhaodong Liu. Coherent interfaces govern direct transformation from graphite to diamond. *Nature*, 607:486–491, 2022. doi: 10.1038/s41586-022-04863-2. URL <https://doi.org/10.1038/s41586-022-04863-2>.
- [20] Tetsuo Irifune, Ayako Kurio, Shizue Sakamoto, Toru Inoue, and Hitoshi Sumiya. Ultrahard polycrystalline diamond from graphite. *Nature*, 421:599–600, 2003. doi: 10.1038/421599b. URL <https://doi.org/10.1038/421599b>.
- [21] Quan Huang, Dongli Yu, Bo Xu, Wentao Hu, Yanming Ma, Yanbin Wang, Zhisheng Zhao, Bin Wen, Julong He, Zhongyuan Liu, and Yongjun Tian. Nanotwinned diamond with unprecedented hardness and stability. *Nature*, 510:250–253, 2014. doi: 10.1038/nature13381. URL <https://doi.org/10.1038/nature13381>.

- [22] Vadym N. Mochalin, Olga Shenderova, Dean Ho, and Yury Gogotsi. The properties and applications of nanodiamonds. *Nature Nanotechnology*, 7(1):11–23, 2012. doi: 10.1038/nnano.2011.209. URL <https://doi.org/10.1038/nnano.2011.209>.
- [23] Vincent Pichot, Benedikt Risse, Fabien Schnell, Julien Mory, and Denis Spitzer. Understanding ultra-fine nanodiamond formation using nanostructured explosives. *Scientific Reports*, 3:2159, 2013. doi: 10.1038/srep02159. URL <https://doi.org/10.1038/srep02159>.
- [24] Chao Fang, Yuewen Zhang, Weixia Shen, Shishuai Sun, Zhuangfei Zhang, Lisha Xue, and Xiaopeng Jia. Synthesis and characterization of HPHT large single-crystal diamonds under the simultaneous influence of oxygen and hydrogen. *CrystEngComm*, 19(38):5727–5734, 2017. doi: 10.1039/C7CE01349C. URL <https://doi.org/10.1039/C7CE01349C>.
- [25] Shuai Fang, Yongkui Wang, Liangchao Chen, Zhiyun Lu, Zhenghao Cai, Chao Fang, Zhandong Zhao, Hongan Ma, and Xiaopeng Jia. The effect of pressure on synthetic diamond crystals at high temperatures and pressures in an Fe/Ni catalyst system. *CrystEngComm*, 23(6):1406–1414, 2021. doi: 10.1039/D0CE01452D. URL <https://doi.org/10.1039/D0CE01452D>.
- [26] Christoph E. Nebel. CVD diamond: a review on options and reality. *Functional Diamond*, 3(1):1–10, 2023. doi: 10.1080/26941112.2023.2201592. URL <https://doi.org/10.1080/26941112.2023.2201592>.
- [27] Hitoshi Sumiya and Kenji Tamasaku. Large defect-free synthetic type IIa diamond crystals synthesized via high pressure and high temperature. *Japanese Journal of Applied Physics*, 51(9R):090102, 2012. doi: 10.7567/JJAP.51.090102. URL <https://doi.org/10.7567/JJAP.51.090102>.
- [28] Yongkui Wang, Shuai Fang, Liangchao Chen, Chao Fang, Zhandong Zhao, Hongan Ma, and Xiaopeng Jia. Graphite/diamond transformation mechanism under the action of an iron-based catalyst solvent under high pressure and high temperature. *CrystEngComm*, 25:1884–1893, 2023. doi: 10.1039/D2CE01202B. URL <https://doi.org/10.1039/D2CE01202B>.
- [29] Branko Deljanin. Stable enhancements and unstable treatments of coloured diamonds. In *Proceedings of the 2nd Mediterranean Gem and Jewellery Conference*, Valencia, Spain, May 2016. URL [https://www.researchgate.net/publication/308929181\\_Stable\\_Enhancements\\_and\\_Unstable\\_Treatments\\_of\\_Coloured\\_Diamonds](https://www.researchgate.net/publication/308929181_Stable_Enhancements_and_Unstable_Treatments_of_Coloured_Diamonds). Figure source accessed through ResearchGate, 2026-05-15.
- [30] Diamond Land. HPHT diamond synthesis: Technology, process, and applications, 2026. URL <https://www.diamond.land/en/blog/articles/hpht-synthesis/>. Web article; accessed 2026-05-15.
- [31] Devi Shanker Misra. Studies of dislocations in type Ib, type IIa HPHT and CVD single crystal diamonds. *Crystals*, 13(4):657, 2023. doi: 10.3390/cryst13040657. URL <https://doi.org/10.3390/cryst13040657>.
- [32] Richard S. Balmer, J. R. Brandon, S. L. Clewes, H. K. Dhillon, J. M. Dodson, I. Friel, P. N. Inglis, T. D. Madgwick, M. L. Markham, T. P. Mollart, N. Perkins, G. A. Scarsbrook, D. J. Twitchen, A. J. Whitehead, J. J. Wilman, and S. M. Woollard. Chemical vapour deposition synthetic diamond: materials, technology and applications. *Journal of Physics: Condensed Matter*, 21(36):364221, 2009. doi: 10.1088/0953-8984/21/36/364221. URL <https://doi.org/10.1088/0953-8984/21/36/364221>.

- [33] E. V. Bushuev, V. Yu. Yurov, A. P. Bolshakov, V. G. Ralchenko, A. A. Khomich, I. A. Antonova, E. E. Ashkinazi, V. A. Shershulin, V. P. Pashinin, and V. I. Konov. Express in situ measurement of epitaxial CVD diamond film growth kinetics. *Diamond and Related Materials*, 72:61–70, 2017. doi: 10.1016/j.diamond.2016.12.021. URL <https://doi.org/10.1016/j.diamond.2016.12.021>.
- [34] Diamond Land. CVD diamond synthesis: Principles, process, and applications, 2026. URL <https://www.diamond.land/en/blog/articles/cvd-synthesis/>. Web article; accessed 2026-05-15.
- [35] Yufei Meng, Chih-Shiue Yan, Joseph Lai, Szczesny Krasnicki, Jinfu Shu, Ting Yu, Qi Liang, Hockwang Mao, and Russell J. Hemley. Enhanced optical properties of chemical vapor deposited single crystal diamond by low-pressure/high-temperature annealing. *Proceedings of the National Academy of Sciences*, 105(46):17620–17625, 2008. doi: 10.1073/pnas.0808230105. URL <https://doi.org/10.1073/pnas.0808230105>.
- [36] S. J. Charles, J. E. Butler, B. N. Feygelson, M. E. Newton, D. L. Carroll, J. W. Steeds, H. Darwish, C.-S. Yan, H.-K. Mao, and R. J. Hemley. Characterization of nitrogen doped chemical vapour deposited single crystal diamond before and after high pressure, high temperature annealing. *physica status solidi (a)*, 201(11):2473–2485, 2004. doi: 10.1002/pssa.200405175. URL <https://doi.org/10.1002/pssa.200405175>.
- [37] Kazuyuki Ueda, Makoto Kasu, Yoshihiro Yamauchi, Toshiki Makimoto, Michael Schwitters, Daniel J. Twitchen, Geoffrey A. Scarsbrook, and Steven E. Coe. High-pressure and high-temperature annealing effects on CVD homoepitaxial diamond films. *Diamond and Related Materials*, 15(11–12):1789–1791, 2006. doi: 10.1016/j.diamond.2006.07.023. URL <https://doi.org/10.1016/j.diamond.2006.07.023>.
- [38] R. H. Telling, C. J. Pickard, M. C. Payne, and J. E. Field. Theoretical strength and cleavage of diamond. *Physical Review Letters*, 84(22):5160–5163, 2000. doi: 10.1103/PhysRevLett.84.5160. URL <https://doi.org/10.1103/PhysRevLett.84.5160>.
- [39] Natalie C. Golota, David Preiss, Zachary P. Fredin, Prashant Patil, Daniel P. Banks, Salima Bahri, Robert G. Griffin, and Neil Gershenfeld. High aspect ratio diamond nanosecond laser machining. *Applied Physics A*, 129:490, 2023. doi: 10.1007/s00339-023-06755-2. URL <https://doi.org/10.1007/s00339-023-06755-2>.
- [40] Yun Zhao, Juping Tu, Liangxian Chen, Junjun Wei, Jinlong Liu, and Chengming Li. Toughness enhancement of single-crystal diamond by the homoepitaxial growth of periodic nitrogen-doped nanomultilayers. *International Journal of Minerals, Metallurgy and Materials*, 30(4):766–771, 2023. doi: 10.1007/s12613-022-2497-1. URL <https://doi.org/10.1007/s12613-022-2497-1>.
- [41] Bin Wen, Bo Xu, Yanbin Wang, Guoying Gao, Xiang-Feng Zhou, Zhisheng Zhao, and Yongjun Tian. Continuous strengthening in nanotwinned diamond. *npj Computational Materials*, 5:117, 2019. doi: 10.1038/s41524-019-0256-2. URL <https://doi.org/10.1038/s41524-019-0256-2>.
- [42] Yunqi Zhang, Yumei Zhu, and Zhihong Li. Preparation and characterization of carbide particle-toughened Si–B system of high thermostability polycrystalline diamond by HPHT sintering. *Materials*, 16(11):3933, 2023. doi: 10.3390/ma16113933. URL <https://doi.org/10.3390/ma16113933>.

- [43] T. V. Kononenko, V. G. Ralchenko, I. I. Vlasov, S. V. Garnov, and V. I. Konov. Ablation of CVD diamond with nanosecond laser pulses of UV–IR range. *Diamond and Related Materials*, 7(11–12):1623–1627, 1998. doi: 10.1016/S0925-9635(98)00198-8. URL [https://doi.org/10.1016/S0925-9635\(98\)00198-8](https://doi.org/10.1016/S0925-9635(98)00198-8).
- [44] Vitali V. Kononenko. Modification of diamond surface by femtosecond laser pulses. *Photonics*, 10(10):1077, 2023. doi: 10.3390/photonics10101077. URL <https://doi.org/10.3390/photonics10101077>.
- [45] H. R. Phillip and E. A. Taft. Kramers-kronig analysis of reflectance data for diamond. *Physical Review*, 136(5A):A1445–A1448, 1964. doi: 10.1103/PhysRev.136.A1445. URL <https://doi.org/10.1103/PhysRev.136.A1445>.
- [46] Mikhail N. Polyanskiy. Refractiveindex.info database of optical constants. *Scientific Data*, 11:94, 2024. doi: 10.1038/s41597-023-02898-2. URL <https://doi.org/10.1038/s41597-023-02898-2>.
- [47] J. Walker. Optical absorption and luminescence in diamond. *Reports on Progress in Physics*, 42(10):1605–1659, 1979. doi: 10.1088/0034-4885/42/10/001. URL <https://doi.org/10.1088/0034-4885/42/10/001>.
- [48] D. Howell, A. T. Collins, L. C. Loudin, P. L. Diggle, U. F. S. D’Haenens-Johansson, K. V. Smit, A. N. Katrusha, J. E. Butler, and F. Nestola. Automated FTIR mapping of boron distribution in diamond. *Diamond and Related Materials*, 96:207–215, 2019. doi: 10.1016/j.diamond.2019.02.029. URL <https://doi.org/10.1016/j.diamond.2019.02.029>.
- [49] E. Rohrer, C. E. Nebel, M. Stutzmann, A. Flöter, R. Zachai, X. Jiang, and C. P. Klages. Photoconductivity of undoped, nitrogen- and boron-doped CVD- and synthetic diamond. *Diamond and Related Materials*, 7(6):879–883, 1998. doi: 10.1016/S0925-9635(97)00318-X. URL [https://doi.org/10.1016/S0925-9635\(97\)00318-X](https://doi.org/10.1016/S0925-9635(97)00318-X).
- [50] B. N. Chichkov, C. Momma, S. Nolte, F. von Alvensleben, and A. Tünnermann. Femtosecond, picosecond and nanosecond laser ablation of solids. *Applied Physics A*, 63(2):109–115, 1996. doi: 10.1007/BF01567637. URL <https://doi.org/10.1007/BF01567637>.
- [51] Bakhtiar Ali, I. V. Litvinyuk, and Maksym Rybachuk. Femtosecond laser micromachining of diamond: Current research status, applications and challenges. *Carbon*, 179:209–226, 2021. doi: 10.1016/j.carbon.2021.04.025. URL <https://doi.org/10.1016/j.carbon.2021.04.025>.
- [52] Rafael R. Gattass and Eric Mazur. Femtosecond laser micromachining in transparent materials. *Nature Photonics*, 2:219–225, 2008. doi: 10.1038/nphoton.2008.47. URL <https://doi.org/10.1038/nphoton.2008.47>.
- [53] Jian Cheng, Chang Liu, Shao Shang, Dong Liu, Walter Perrie, Geoff Dearden, and Ken Watkins. A review of ultrafast laser materials micromachining. *Optics & Laser Technology*, 46:88–102, 2013. doi: 10.1016/j.optlastec.2012.06.037. URL <https://doi.org/10.1016/j.optlastec.2012.06.037>.
- [54] S. Nolte, C. Momma, G. Kamlage, A. Ostendorf, C. Fallnich, F. von Alvensleben, and H. Welling. Polarization effects in ultrashort-pulse laser drilling. *Applied Physics A*, 68:563–567, 1999. doi: 10.1007/s003390050941. URL <https://doi.org/10.1007/s003390050941>.

- [55] Prashant Patil, Daniel Banks, Salima Bahri, Will Langford, Camron Blackburn, Zach Fredin, Robert Griffin, and Neil Gershenfeld. Diamond rotors for MAS-NMR. U.S. Patent Application Publication US 2021/0146475 A1, 2021. URL <https://patents.google.com/patent/US20210146475A1/en>. Published 2021-05-20; assignee Massachusetts Institute of Technology.
- [56] Steven Praver and Robert J. Nemanich. Raman spectroscopy of diamond and doped diamond. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 362 (1824):2537–2565, 2004. doi: 10.1098/rsta.2004.1451. URL <https://doi.org/10.1098/rsta.2004.1451>.
- [57] M. Bernard, A. Deneuve, and P. Muret. Non-destructive determination of the boron concentration of heavily doped metallic diamond thin films from Raman spectroscopy. *Diamond and Related Materials*, 13(2):282–286, 2004. doi: 10.1016/j.diamond.2003.10.051. URL <https://doi.org/10.1016/j.diamond.2003.10.051>.
- [58] P. W. May, W. J. Ludlow, M. Hannaway, P. J. Heard, J. A. Smith, and K. N. Rosser. Raman and conductivity studies of boron-doped microcrystalline diamond, faceted nanocrystalline diamond and cauliflower diamond films. *Diamond and Related Materials*, 17(2):105–117, 2008. doi: 10.1016/j.diamond.2007.11.005. URL <https://doi.org/10.1016/j.diamond.2007.11.005>.
- [59] Daniel Banks, Brian Michael, Natalie Golota, and Robert G. Griffin. 3D-printed stators and drive caps for magic-angle spinning NMR. *Journal of Magnetic Resonance*, 335:107126, 2022. doi: 10.1016/j.jmr.2021.107126. URL <https://doi.org/10.1016/j.jmr.2021.107126>.
- [60] Ke Xu, Oliver Pecher, Marco Braun, and Jörn Schmedt auf der Günne. Stable magic angle spinning with low-cost 3D-printed parts. *Journal of Magnetic Resonance*, 333:107096, 2021. doi: 10.1016/j.jmr.2021.107096. URL <https://doi.org/10.1016/j.jmr.2021.107096>.
- [61] Daniel Pereira, Mariana Sardo, Ildefonso Marín-Montesinos, and Luís Mafra. One-shot resin 3D-printed stators for low-cost fabrication of magic-angle spinning NMR probeheads. *Analytical Chemistry*, 95(27):10384–10389, 2023. doi: 10.1021/acs.analchem.3c01323. URL <https://doi.org/10.1021/acs.analchem.3c01323>.
- [62] Element Six. CVD Diamond Handbook, 2022. URL [https://e6cvd.com/media/wysiwyg/pdf/Element\\_Six\\_CVD\\_Diamond\\_handbook\\_2022.pdf](https://e6cvd.com/media/wysiwyg/pdf/Element_Six_CVD_Diamond_handbook_2022.pdf). Technical handbook values for diamond thermal expansion and thermal properties.
- [63] DuPont. Vespel SP-1, 2026. URL <https://www.dupont.com/products/vespel-sp-1.html>. Manufacturer property page; accessed 2026-05-11.
- [64] Curbell Plastics. PCTFE Kel-F, Neoflon Cryogenic Applications and Material Properties, 2026. URL <https://www.curbellplastics.com/materials/plastics/pctfe/>. Supplier material-property page; accessed 2026-05-12.
- [65] Plastic Machining Inc. PCTFE (Kel-F) Data Sheet, 2021. URL <https://www.plasticmachininginc.com/pctfe-data-sheet/>. Supplier data sheet; accessed 2026-05-11.
- [66] Torlon.site. Torlon 4203, 2026. URL <https://torlon.site/torlon-grades/torlon-4203/>. Material property summary; accessed 2026-05-11.

- [67] UpNano. UpPhoto Resin Data Sheet, 2024. URL [https://www.upnano.com/wp-content/uploads/2024/04/UpNano\\_Resins\\_Data\\_Sheet\\_UpPhoto.pdf](https://www.upnano.com/wp-content/uploads/2024/04/UpNano_Resins_Data_Sheet_UpPhoto.pdf). Manufacturer data sheet; accessed 2026-05-11.
- [68] Qingda Ji, Johnny Moughames, Xue Chen, Guoyuan Fang, Juan J. Huaroto, Vincent Laude, Jose A. Iglesias Martínez, Gwenn Ulliac, Cédric Clévy, Philippe Lutz, Kanty Rabenoroso, Vincent Guelpa, Arnaud Spangenberg, Jinyang Liang, Anne Mosset, and Muamer Kadic. 4D thermomechanical metamaterials for soft microrobotics. *Communications Materials*, 2:93, 2021. doi: 10.1038/s43246-021-00189-0. URL <https://www.nature.com/articles/s43246-021-00189-0>.
- [69] Jia Qu, Muamer Kadic, Andreas Naber, and Martin Wegener. Micro-structured two-component 3D metamaterials with negative thermal-expansion coefficient from positive constituents. *Scientific Reports*, 7:40643, 2017. doi: 10.1038/srep40643. URL <https://www.nature.com/articles/srep40643>.
- [70] Formlabs. Rigid 10K Resin, 2026. URL <https://formlabs.com/global/products/rigid-10k-resin/>. Manufacturer material property page; accessed 2026-05-11.
- [71] Corning. MACOR Machinable Glass Ceramic for Industrial Applications, 2026. URL <https://www.corning.com/media/worldwide/csm/documents/71759a443535431395eb34ebad091cb.pdf>. Manufacturer data sheet; accessed 2026-05-11.
- [72] CoorsTek. Alumina, 2026. URL <https://www.coorstek.com/en/materials/alumina/>. Manufacturer technical-ceramic property page; accessed 2026-05-11.
- [73] GW Ceramics. Zirconia, 2026. URL <https://gwc ceramics.com/materials/zirconia/>. Technical-ceramic property page; accessed 2026-05-11.
- [74] T. H. Panzera, J. C. C. Rubio, C. R. Bowen, and P. J. Walker. Microstructural design of materials for aerostatic bearings. *Cement and Concrete Composites*, 30(7):649–660, 2008. doi: 10.1016/j.cemconcomp.2007.08.013. URL <https://doi.org/10.1016/j.cemconcomp.2007.08.013>.
- [75] Yandong Gu, Jinwu Cheng, Chaojie Xie, Longyu Li, and Changgeng Zheng. Theoretical and numerical investigations on static characteristics of aerostatic porous journal bearings. *Machines*, 10(3):171, 2022. doi: 10.3390/machines10030171.
- [76] S. Yoshimoto, H. Tozuka, and S. Dambara. Static characteristics of aerostatic porous journal bearings with a surface-restricted layer. *Proceedings of the Institution of Mechanical Engineers, Part J: Journal of Engineering Tribology*, 217(2):125–132, 2003. doi: 10.1243/13506500360603552.
- [77] Xin Xiao, Jianzhou Du, Yu Zhang, Jingyi Yan, Yunping Li, Kongjun Zhu, and Luming Wang. Study on static characteristics of aerostatic bearing based on porous SiC ceramic membranes. *Membranes*, 12(9):898, 2022. doi: 10.3390/membranes12090898. URL <https://doi.org/10.3390/membranes12090898>.
- [78] Jin Zhao, Linying Wang, Xiaojian Mao, Liqiong An, Yutao Liu, Shiwei Wang, Jian Zhang, and Kai Feng. Preparation and properties of porous alumina ceramics for ultra-precision aerostatic bearings. *Ceramics International*, 48(9):13311–13318, 2022. doi: 10.1016/j.ceramint.2022.01.210. URL <https://doi.org/10.1016/j.ceramint.2022.01.210>.
- [79] MIT.nano. NanoOne 3D Printer from UpNano, n.d. URL <https://nanousers.mit.edu/prototyping-facility/nanoone>. Facility equipment specification page; accessed 2026-05-09.

- [80] UpNano. HubNano Resins: Material Handling and Recommended Printing Parameters, n.d. URL <https://www.upnano.com/hubnano-resins/>. Manufacturer material-handling and process-parameter guide; accessed 2026-05-09.
- [81] Rajdeep Deb, Andreas Hunkeler, Dirk Wilhelm, Patrick Jenny, and Beat H. Meier. Numerical modeling and design decisions for aerostatic bearings with relatively large nozzle sizes in magic-angle spinning (MAS) systems. *Tribology International*, 175:107855, 2022. doi: 10.1016/j.triboint.2022.107855. URL <https://doi.org/10.1016/j.triboint.2022.107855>.
- [82] Jianping Li, Hu Huang, and Takeshi Morita. Stepping piezoelectric actuators with large working stroke for nano-positioning systems: A review. *Sensors and Actuators A: Physical*, 292:39–51, 2019. doi: 10.1016/j.sna.2019.04.006.
- [83] Hsien-Shun Liao, Christian Werner, Roman Slipets, Peter Emil Larsen, Ing-Shouh Hwang, Tien-Jen Chang, Hans Ulrich Danzebrink, Kuang-Yuh Huang, and En-Te Hwu. Low-cost, open-source XYZ nanopositioner for high-precision analytical applications. *HardwareX*, 11:e00317, 2022. doi: 10.1016/j.ohx.2022.e00317. URL <https://doi.org/10.1016/j.ohx.2022.e00317>.
- [84] Yusuf Altintas, A. Verl, C. Brecher, L. Uriarte, and G. Pritschow. Machine tool feed drives. *CIRP Annals*, 60(2):779–796, 2011. doi: 10.1016/j.cirp.2011.05.010.
- [85] Changyou Li, Mengtao Xu, Wenjun Song, and Hongzhuang Zhang. A review of static and dynamic analysis of ball screw feed drives, recirculating linear guideway, and ball screw. *International Journal of Machine Tools and Manufacture*, 188:104021, 2023. doi: 10.1016/j.ijmachtools.2023.104021.
- [86] G. Pritschow. A comparison of linear and conventional electromechanical drives. *CIRP Annals*, 47(2):541–548, 1998. doi: 10.1016/S0007-8506(07)63241-7.
- [87] Aerotech. ALS130H Series Mechanical Bearing, Direct-Drive Linear Stage Data Sheet, n.d.. URL <https://www.aerotech.com/wp-content/uploads/2020/10/ALS130H.pdf>. Official vendor data sheet; accessed 2026-05-09.
- [88] Aerotech. ABL1000 Series Air-Bearing, Direct-Drive Linear Stage Data Sheet, n.d.. URL <https://www.aerotech.com/wp-content/uploads/2020/10/ABL1000-Data-Sheet-D20250110.pdf>. Official vendor data sheet; accessed 2026-05-09.
- [89] International Organization for Standardization. ISO 230-2:2014 Test code for machine tools – Part 2: Determination of accuracy and repeatability of positioning of numerically controlled axes, 2014. URL <https://www.iso.org/standard/55295.html>. Confirmed 2025; accessed 2026-05-08.
- [90] Eberhard Manske, Gerd Jager, Tino Hausotte, and Roland Fussl. Recent developments and challenges of nanopositioning and nanomeasuring technology. *Measurement Science and Technology*, 23(7):074001, 2012. doi: 10.1088/0957-0233/23/7/074001.
- [91] Andrew J. Fleming. A review of nanometer resolution position sensors: Operation and performance. *Sensors and Actuators A: Physical*, 190:106–126, 2013. doi: 10.1016/j.sna.2012.10.016.
- [92] IEEE. ANSI/IEEE Std 176-1987 IEEE Standard on Piezoelectricity, 1988. Published 1988-01-29; inactive/withdrawn standard; used for notation and the linear constitutive framework.
- [93] Santosh Devasia, Evangelos Eleftheriou, and S. O. Reza Moheimani. A survey of control issues in nanopositioning. *IEEE Transactions on Control Systems Technology*, 15(5):802–823, 2007. doi: 10.1109/TCST.2007.903345.

- [94] Yuen Kuan Yong, S. O. Reza Moheimani, Brian J. Kenton, and Kam K. Leang. Invited review article: High-speed flexure-guided nanopositioning: Mechanical design and control issues. *Review of Scientific Instruments*, 83(12):121101, 2012. doi: 10.1063/1.4765048.
- [95] Qing Yao, J. Dong, and P. M. Ferreira. Design, analysis, fabrication and testing of a parallel-kinematic micropositioning xy stage. *International Journal of Machine Tools and Manufacture*, 47(6):946–961, 2007. doi: 10.1016/j.ijmachtools.2006.07.007.
- [96] Jingnan Cai, Wei Dong, and Ryozi Nagamune. A survey of bouc-wen hysteretic models applied to piezo-actuated mechanical systems: Modeling, identification, and control. *Journal of Intelligent Material Systems and Structures*, 34(16):1843–1863, 2023. doi: 10.1177/1045389X231157361.
- [97] Jinqiang Gan and Xianmin Zhang. Nonlinear hysteresis modeling of piezoelectric actuators using a generalized bouc-wen model. *Micromachines*, 10(3):183, 2019. doi: 10.3390/mi10030183.
- [98] OMAX. OMAX Waterjet FAQs, n.d.. URL <https://promo.omax.com/en/us/omax-waterjet-faqs>. Manufacturer FAQ; accessed 2026-05-09.
- [99] OMAX. Taper Compensation, n.d.. URL <https://www.omax.com/en/us/media-center/tips/taper-compensation>. Manufacturer technical note; accessed 2026-05-09.
- [100] Sodick. AP450L Wire EDM, n.d. URL <https://sodick.com/product/ap450l/>. Manufacturer specification page; accessed 2026-05-09.
- [101] K. H. Ho and S. T. Newman. State of the art electrical discharge machining (edm). *International Journal of Machine Tools and Manufacture*, 43(13):1287–1300, 2003. doi: 10.1016/S0890-6955(03)00162-7.
- [102] Laurentiu Slatineanu, Oana Dodun, Margareta Coteata, Gheorghe Nagit, Irina Besliu Bancescu, and Adelina Hrituc. Wire electrical discharge machining—a review. *Machines*, 8(4):69, 2020. doi: 10.3390/machines8040069.
- [103] Mohamed Hashish. Abrasive waterjet machining. *Materials*, 17(13):3273, 2024. doi: 10.3390/ma17133273.
- [104] OMAX. Accessory Comparison: A-Jet vs. Tilt-A-Jet, n.d.. URL <https://www.omax.com/en/us/media-center/tips/jet-vs-tilt-jet>. Manufacturer technical note; accessed 2026-05-09.
- [105] Antony C. S. Chan. Piezo-flexure nano-positioning stage PID controller, 2026. URL <https://github.com/antonysigma/piezo-stage-pid-board>. Open-source hardware and firmware project; accessed 2026-05-09.
- [106] PiezoDrive. PDU150 – Three Channel Ultra-Low Noise 150V Piezo Driver, n.d. URL <https://www.piezodrive.com/modules/pdu150/>. Manufacturer specifications; accessed 2026-05-09.
- [107] Thorlabs. PC4WM Co-Fired Piezo Actuator, 9.5 Micrometer Max Displacement, n.d. URL <https://www.thorlabs.com/item/PC4WM>. Manufacturer product specifications; accessed 2026-05-14.
- [108] ams OSRAM. AS5311 High Resolution Magnetic Linear Encoder Data Sheet, n.d. URL <https://look.ams-osram.com/m/27ea8328172c5288/original/AS5311-DS000200.pdf>. Manufacturer data sheet; accessed 2026-05-14.

- [109] THK. Linear Ball Slide Model LS, n.d. URL [https://www.thk.com/eu/en/products/other\\_linear\\_motion\\_guides/linear\\_ball\\_slide/ls/ls](https://www.thk.com/eu/en/products/other_linear_motion_guides/linear_ball_slide/ls/ls). Official product page; accessed 2026-05-09.
- [110] Aerotech. PRO190SL Mechanical Bearing, Ball-Screw Stage Data Sheet, n.d.. URL <https://www.aerotech.com/wp-content/uploads/2020/10/PRO190SL-Data-Sheet-D20230328.pdf>. Official vendor data sheet; accessed 2026-05-09.
- [111] Aerotech. PRO115SL Mechanical Bearing, Ball-Screw Stage Data Sheet, n.d.. URL <https://www.aerotech.com/wp-content/uploads/2020/10/PRO115SL-Data-Sheet-D20230227.pdf>. Official vendor data sheet; accessed 2026-05-09.
- [112] Aerotech. ATS100 Mechanical Bearing, Screw-Driven Linear Stage, n.d.. URL <https://www.aerotech.com/product/ats100-mechanical-bearing-screw-driven-linear-stage/>. Official product page; accessed 2026-05-09.
- [113] Aerotech. ANT95L Single-Axis Linear Nanopositioning Stage Data Sheet, n.d.. URL <https://www.aerotech.com/wp-content/uploads/2020/10/ANT95L-Data-Sheet-D20230609.pdf>. Official vendor data sheet; accessed 2026-05-09.
- [114] Renishaw. How Renishaw Optical Encoders Work, n.d.. URL <https://www.renishaw.com/en/how-optical-encoders-work--36979>. Official technical overview; accessed 2026-05-09.
- [115] Renishaw. TONiC Optical Incremental Encoder Series, n.d.. URL <https://www.renishaw.com/en/tonic-optical-incremental-encoder-series--37824>. Official product documentation; accessed 2026-05-09.
- [116] Renishaw. Optical Encoders Scale Range, n.d.. URL <https://www.renishaw.com/en/optical-encoders-scale-range--36031>. Official product documentation; accessed 2026-05-09.
- [117] Aerotech. Automation1 XC4 PWM Servo Drive, n.d.. URL <https://www.aerotech.com/product/automation1-xc4-pwm-digital-drive/>. Official product page; accessed 2026-05-09.
- [118] NSK. NSK Linear Guide High-Accuracy Series: HA/HS Models, n.d. URL <https://www.nsk.com/content/dam/nsk/common/catalogs/ctrGpdf/precision/e3329b.pdf>. Official manufacturer catalog; accessed 2026-05-09.
- [119] SKF Evolution. SKF Launches LLT Profile Rail Guides, n.d. URL <https://evolution.skf.com/skf-launches-llt-profile-rail-guides/>. SKF technical news article; accessed 2026-05-09.
- [120] Aerotech. Automation1 Frequency Response Module, n.d.. URL <https://help.aerotech.com/automation1/Content/Guide-Servo-Tuning/1.6-Module-Frequency-Response.htm>. Official Automation1 servo tuning documentation; accessed 2026-05-09.
- [121] Aerotech. Automation1 FeedforwardAdvance Parameter, n.d.. URL <https://help.aerotech.com/automation1/Content/Parameters/FeedforwardAdvance.htm>. Official Automation1 parameter documentation; accessed 2026-05-09.
- [122] Aerotech. Automation1 FeedforwardFilterSetup Parameter, n.d.. URL <https://help.aerotech.com/automation1/Content/Parameters/FeedforwardFilterSetup.htm>. Official Automation1 parameter documentation; accessed 2026-05-09.

- [123] Aerotech. A3200 Software-Based Machine Controller, n.d.. URL <https://www.aerotech.com/wp-content/uploads/2021/04/A3200-Software-Based-Machine-Controller-2021.pdf>. Official product document; accessed 2026-05-09.
- [124] Aerotech. A3200 PC-Based Controller Selection, Setup, Configuration, and Optimization, n.d.. URL <https://help.aerotech.com/a3200/user-guides/A3200-PC-Based-Controller-Selection-Setup/Content/Online-Guide-Automation1-A3200-PC-Based-Controller-Selection-Setup/PC-Selection-Configuration-Optimization.htm>. Official setup guide; accessed 2026-05-09.
- [125] TenAsys. About INtime, n.d. URL [https://support.tenasys.com/6-2/About\\_Overview](https://support.tenasys.com/6-2/About_Overview). Official vendor documentation; accessed 2026-05-09.
- [126] Aerotech. CADFusion Graphical Motion Programming Software, n.d.. URL <https://www.aerotech.com/product/cadfusion-graphical-motion-programming-software/>. Official product page; accessed 2026-05-09.
- [127] Aerotech. Position Synchronized Output (PSO), n.d.. URL <https://help.aerotech.com/automation1/Content/Position-Synchronized-Output-PSO.htm>. Official Automation1 documentation; accessed 2026-05-08.
- [128] Aerotech. Ndrive CP Hardware Manual, n.d.. URL <https://www.aerotech.com/wp-content/uploads/2020/09/Ndrive-CP.pdf>. Official vendor hardware manual; accessed 2026-05-09.
- [129] Aerotech. Automation1 XC2 Compact PWM Servo Drive, n.d.. URL <https://www.aerotech.com/product/automation1-xc2-pwm-digital-drive/>. Official product page; accessed 2026-05-09.
- [130] Chung-Ching Liu, Meng-Shiun Tsai, Mao-Qi Hong, and Pu-Yang Tang. Development of a novel tuning approach of the notch filter of the servo feed drive system. *Journal of Manufacturing and Materials Processing*, 4(1):21, 2020. doi: 10.3390/jmmp4010021.
- [131] Aerotech. Position Synchronized Output (PSO) – Coordinate Part Position with Process Control, n.d.. URL <https://www.aerotech.com/position-synchronized-output-pso-coordinate-part-position-with-process-control/>. Official application note; accessed 2026-05-08.
- [132] B. N. Chichkov, C. Momma, S. Nolte, F. von Alvensleben, and A. Tünnermann. Femtosecond, picosecond and nanosecond laser ablation of solids. *Applied Physics A*, 63:109–115, 1996. doi: 10.1007/BF01567637. URL <https://doi.org/10.1007/BF01567637>.
- [133] Koji Sugioka and Ya Cheng. Ultrafast lasers—reliable tools for advanced materials processing. *Light: Science & Applications*, 3:e149, 2014. doi: 10.1038/lssa.2014.30. URL <https://doi.org/10.1038/lssa.2014.30>.
- [134] Rafael R. Gattass and Eric Mazur. Femtosecond laser micromachining in transparent materials. *Nature Photonics*, 2:219–225, 2008. doi: 10.1038/nphoton.2008.47. URL <https://doi.org/10.1038/nphoton.2008.47>.
- [135] Mangirdas Malinauskas, Albertas Zukauskas, Satoshi Hasegawa, Yoshio Hayasaki, Vyngantas Mizeikis, Ricardas Buividas, and Saulius Juodkazis. Ultrafast laser processing of materials: From

- science to industry. *Light: Science & Applications*, 5:e16133, 2016. doi: 10.1038/lssa.2016.133. URL <https://doi.org/10.1038/lssa.2016.133>.
- [136] Zhenyuan Lin and Minghui Hong. Femtosecond laser precision engineering: From micron, submicron, to nanoscale. *Ultrafast Science*, 2021:9783514, 2021. doi: 10.34133/2021/9783514. URL <https://doi.org/10.34133/2021/9783514>.
- [137] Heng Guo, Jiawang Xie, Guangzhi He, Dezhi Zhu, Ming Qiao, Jianfeng Yan, Jiachen Yu, Ji-aqun Li, Yuzhi Zhao, Ma Luo, and Haoze Han. A review of ultrafast laser micro/nano fabrication: Material processing, surface/interface controlling, and devices fabrication. *Nano Research*, 17(7):6212–6230, 2024. doi: 10.1007/s12274-024-6644-z. URL <https://doi.org/10.1007/s12274-024-6644-z>.
- [138] Kerim Yildirim, Balasubramanian Nagarajan, Tegoeh Tjahjowidodo, and Sylvie Castagne. Review of in-situ process monitoring for ultra-short pulse laser micromanufacturing. *Journal of Manufacturing Processes*, 133:1126–1159, 2025. doi: 10.1016/j.jmapro.2024.12.011. URL <https://doi.org/10.1016/j.jmapro.2024.12.011>.
- [139] Vasiliki E. Alexopoulou and Angelos P. Markopoulos. A critical assessment regarding two-temperature models. *Archives of Computational Methods in Engineering*, 31:93–123, 2024. doi: 10.1007/s11831-023-09974-1. URL <https://doi.org/10.1007/s11831-023-09974-1>.
- [140] Ebenezer Owusu-Ansah and Colin Dalton. Femtosecond laser machining of an x-ray mask in a 500 micron-thick tungsten sheet. *Micromachines*, 14(11):2071, 2023. doi: 10.3390/mi14112071. URL <https://doi.org/10.3390/mi14112071>.
- [141] FLIR. *Flea3 USB3 Technical Reference*, 2017. URL <https://www.eureca.de/files/pdf/optoelectronics/flir/FL3-U3-Technical-Reference.pdf>. Accessed 2026-05-13.
- [142] Thorlabs. AC254-200-A: f = 200 mm, 1 inch achromatic doublet, 400–700 nm antireflection coating, 2026. URL <https://www.thorlabs.com/thorProduct.cfm?partNumber=AC254-200-A>. Accessed 2026-05-13.
- [143] Mitutoyo. M Plan Apo brightfield infinity-corrected microscope objectives, 2026. URL <https://dev.pim.mitutoyo.com/products/optical-measurement/infinity-corrected-microscope-objectives/m-plan-apo-for-brightfield-observation/>. Accessed 2026-05-13.
- [144] Thorlabs. ASL2520: 25.0 mm uv fused silica aspheric lens, f = 20.0 mm, na = 0.65, uncoated, 2026. URL [https://www.thorlabs.com/newgrouppage9.cfm?objectgroup\\_id=14329&partnumber=ASL2520](https://www.thorlabs.com/newgrouppage9.cfm?objectgroup_id=14329&partnumber=ASL2520). Accessed 2026-05-13.
- [145] Z. Bor. Distortion of femtosecond laser pulses in lenses and lens systems. *Journal of Modern Optics*, 35(12):1907–1918, 1988. doi: 10.1080/713822325. URL <https://doi.org/10.1080/713822325>.
- [146] M. Kempe and W. Rudolph. Impact of chromatic and spherical aberration on the focusing of ultrashort light pulses by lenses. *Optics Letters*, 18(2):137–139, 1993. doi: 10.1364/OL.18.000137. URL <https://doi.org/10.1364/OL.18.000137>.
- [147] Apple. iPhone 15 technical specifications, 2023. URL <https://support.apple.com/en-us/111831>. Accessed 2026-05-13.

- [148] Binh Xuan Cao, Phuong Le Hoang, Sanghoon Ahn, Jeng-o Kim, Heeshin Kang, and Jiwhan Noh. Real-time laser focusing system for high-precision micromachining using diffractive beam sampler and advanced image sensor. *Optics and Lasers in Engineering*, 107:13–20, 2018. doi: 10.1016/j.optlaseng.2018.03.002. URL <https://doi.org/10.1016/j.optlaseng.2018.03.002>.
- [149] Si-Jia Xu, Yan-Zhao Duan, Yan-Hao Yu, Zhen-Nan Tian, and Qi-Dai Chen. Machine vision-based high-precision and robust focus detection for femtosecond laser machining. *Optics Express*, 29(19): 30952–30960, 2021. doi: 10.1364/OE.438537. URL <https://doi.org/10.1364/OE.438537>.
- [150] Marcus Paulo Ruele, Lucas Ramos De Pretto, Wagner de Rossi, Nilson Dias Vieira Jr., and Ricardo Elgul Samad. Focus tracking system for femtosecond laser machining using low coherence interferometry. *Scientific Reports*, 9:4167, 2019. doi: 10.1038/s41598-019-40749-6. URL <https://doi.org/10.1038/s41598-019-40749-6>.
- [151] Ting-Hsuan Chen, Romain Fardel, and Craig B. Arnold. Ultrafast z-scanning for high-efficiency laser micro-machining. *Light: Science & Applications*, 7(4):17181, 2018. doi: 10.1038/lsa.2017.181. URL <https://doi.org/10.1038/lsa.2017.181>.
- [152] Viet Hoang Dinh, Le Phuong Hoang, Yen Nhi Thai Vu, and Binh Xuan Cao. Auto-focus methods in laser systems for use in high precision materials processing: A review. *Optics and Lasers in Engineering*, 167:107625, 2023. doi: 10.1016/j.optlaseng.2023.107625. URL <https://doi.org/10.1016/j.optlaseng.2023.107625>.
- [153] L. K. Nolasco, F. A. Couto, M. B. Andrade, and C. R. Mendonça. Femtosecond laser micro-machining study with multiple wavelengths in CVD diamond. *Diamond and Related Materials*, 131:109589, 2023. doi: 10.1016/j.diamond.2022.109589. URL <https://doi.org/10.1016/j.diamond.2022.109589>.
- [154] P. Castro-Marín, G. Castro-Olvera, J. Garduño-Mejía, M. Rosete-Aguilar, N. C. Bruce, D. T. Reid, and O. G. Rodríguez-Herrera. Autocorrelation z-scan technique for measuring the spatial and temporal distributions of femtosecond pulses around the focal region of lenses. *Optics Express*, 25(13):14473–14482, 2017. doi: 10.1364/OE.25.014473. URL <https://doi.org/10.1364/OE.25.014473>.
- [155] Huu Dat Nguyen, Enrique Moreno, Anton Rudenko, Nicolas Faure, Xxx Sedao, Cyril Mauclair, Jean-Philippe Colombier, and Razvan Stoian. Super-efficient drilling of metals with ultrafast non diffractive laser beams. *Scientific Reports*, 12:2074, 2022. doi: 10.1038/s41598-022-05967-5. URL <https://doi.org/10.1038/s41598-022-05967-5>.
- [156] Christian Vetter, Remo Giust, Luca Furfaro, Cyril Billet, Luc Froehly, and François Courvoisier. High aspect ratio structuring of glass with ultrafast bessel beams. *Materials*, 14(22):6749, 2021. doi: 10.3390/ma14226749. URL <https://doi.org/10.3390/ma14226749>.
- [157] Can Kerse, Hamit Kalaycioglu, Parviz Elahi, Barbaros Cetin, Denizhan K. Kesim, Onder Akcaalan, Seydi Yavas, Mehmet D. Asik, Bulent Oktem, Heinar Hoogland, Ronald Holzwarth, and Fatih Omer Ilday. Ablation-cooled material removal with ultrafast bursts of pulses. *Nature*, 537:84–88, 2016. doi: 10.1038/nature18619. URL <https://doi.org/10.1038/nature18619>.
- [158] Rick Trebino, Kenneth W. DeLong, David N. Fittinghoff, John N. Sweetser, Marco A. Krumbügel, Bruce A. Richman, and Daniel J. Kane. Measuring ultrashort laser pulses in the time-frequency domain using frequency-resolved optical gating. *Review of Scientific Instruments*, 68(9):3277–3295, 1997. doi: 10.1063/1.1148286. URL <https://doi.org/10.1063/1.1148286>.

- [159] Daniel J. Kane and Andrei B. Vakhtin. A review of ptychographic techniques for ultrashort pulse measurement. *Progress in Quantum Electronics*, 81:100364, 2022. doi: 10.1016/j.pquantelec.2021.100364. URL <https://doi.org/10.1016/j.pquantelec.2021.100364>.
- [160] Rick Trebino, Rana Jafari, S. A. Akturk, Pamela Bowlan, Zhe Guang, Peng Zhu, Enrique Escoto, and Günter Steinmeyer. Highly reliable measurement of ultrashort laser pulses. *Journal of Applied Physics*, 128(17):171103, 2020. doi: 10.1063/5.0022552. URL <https://doi.org/10.1063/5.0022552>.
- [161] M. Raghu Ramaiah, A. K. Sharma, Prasad Naik, P. D. Gupta, and R. A. Ganeev. A second-order autocorrelator for single-shot measurement of femtosecond laser pulse durations. *Sadhana*, 26(6): 603–611, 2001. doi: 10.1007/BF02703461. URL <https://doi.org/10.1007/BF02703461>.
- [162] Rana Jafari, Soroush D. Khosravi, and Rick Trebino. Reliable determination of pulse-shape instability in trains of ultrashort laser pulses using frequency-resolved optical gating. *Scientific Reports*, 12:21006, 2022. doi: 10.1038/s41598-022-25193-3. URL <https://doi.org/10.1038/s41598-022-25193-3>.
- [163] Rick Trebino and Daniel J. Kane. Using phase retrieval to measure the intensity and phase of ultrashort pulses: Frequency-resolved optical gating. *Journal of the Optical Society of America A*, 10(5):1101–1111, 1993. doi: 10.1364/JOSAA.10.001101. URL <https://doi.org/10.1364/JOSAA.10.001101>.
- [164] D. J. Kane and R. Trebino. Characterization of arbitrary femtosecond pulses using frequency-resolved optical gating. *IEEE Journal of Quantum Electronics*, 29(2):571–579, 1993. doi: 10.1109/3.199311. URL <https://doi.org/10.1109/3.199311>.
- [165] Yuta Nakano, Yuichiro Kida, Kazuya Motoyoshi, and Totaro Imasaka. Cross-correlation frequency-resolved optical gating for test-pulse characterization using a self-diffraction signal of a reference pulse. *Applied Sciences*, 6(11):315, 2016. doi: 10.3390/app6110315. URL <https://doi.org/10.3390/app6110315>.
- [166] Miguel Miranda, Cord L. Arnold, Thomas Fordell, Francisco Silva, Benjamín Alonso, Rosa Weigand, Anne L’Huillier, and Helder Crespo. Characterization of broadband few-cycle laser pulses with the d-scan technique. *Optics Express*, 20(17):18732–18743, 2012. doi: 10.1364/OE.20.018732. URL <https://doi.org/10.1364/OE.20.018732>.
- [167] Benjamin Alonso, Inigo J. Sola, and Helder Crespo. Self-calibrating d-scan: Measuring ultrashort laser pulses on-target using an arbitrary pulse compressor. *Scientific Reports*, 8:3264, 2018. doi: 10.1038/s41598-018-21701-6. URL <https://doi.org/10.1038/s41598-018-21701-6>.
- [168] Tapio Purtonen, Antti Kalliosaari, and Antti Salminen. Monitoring and adaptive control of laser processes. *Physics Procedia*, 56:1218–1231, 2014. doi: 10.1016/j.phpro.2014.08.038. URL <https://doi.org/10.1016/j.phpro.2014.08.038>.
- [169] T.-L. Chang, L.-T. Chen, C.-C. Chen, D.-J. Lee, and Y.-F. Chen. Closed-loop control in ultrafast laser milling process using laser triggered plasma. *International Journal of Machine Tools and Manufacturing*, 60:35–44, 2012. doi: 10.1016/j.ijmachtools.2012.04.002. URL <https://doi.org/10.1016/j.ijmachtools.2012.04.002>.
- [170] J. X. Chen, X. Zhou, S. Lin, and Y. Tu. A prediction-correction scheme for microchannel milling using femtosecond laser. *Optics and Lasers in Engineering*, 91:115–123, 2017. doi: 10.1016/j.optlaseng.2016.11.012. URL <https://doi.org/10.1016/j.optlaseng.2016.11.012>.

- [171] Kai Liao, Wenjun Wang, Xuesong Mei, Wenwen Tian, Hai Yuan, Mingqiong Wang, and Bozhe Wang. Shape regulation of tapered microchannels in silica glass ablated by femtosecond laser with theoretical modeling and machine learning. *Journal of Intelligent Manufacturing*, 34:2907–2924, 2023. doi: 10.1007/s10845-022-01950-z. URL <https://doi.org/10.1007/s10845-022-01950-z>.
- [172] Tzveta Apostolova, Vasyi Kurylo, and Iaroslav Gnilitzkyi. Ultrafast laser processing of diamond materials: A review. *Frontiers in Physics*, 9:650280, 2021. doi: 10.3389/fphy.2021.650280. URL <https://doi.org/10.3389/fphy.2021.650280>.
- [173] Zhouwei He, Xiaoyan Sun, Haikuan Chen, Ji'an Duan, and Youwang Hu. Water jet-guided laser cutting technology: A review. *The International Journal of Advanced Manufacturing Technology*, 140:3415–3434, 2025. doi: 10.1007/s00170-025-16273-5. URL <https://doi.org/10.1007/s00170-025-16273-5>.
- [174] Yan Liu, Meirong Wei, Tao Zhang, Hongchao Qiao, and Hui Li. Overview on the development and critical issues of water jet guided laser machining technology. *Optics and Laser Technology*, 137:106820, 2021. doi: 10.1016/j.optlastec.2020.106820. URL <https://doi.org/10.1016/j.optlastec.2020.106820>.
- [175] Bernold Richerzhagen. Method and apparatus for machining material with a liquid-guided laser beam. U.S. Patent US5902499A, 1999. URL <https://patents.google.com/patent/US5902499A/en>.
- [176] Guangyi Zhang, Jiayu Wang, Zhongan Chen, Yaowen Wu, Binying Bao, and Wenwu Zhang. An optical field regulation method for waterjet-guided laser: Reducing taper and improving deep-processing capability. *Journal of Materials Processing Technology*, 334:118637, 2024. doi: 10.1016/j.jmatprotec.2024.118637. URL <https://doi.org/10.1016/j.jmatprotec.2024.118637>.
- [177] Ji Wang, Chunguang Chen, Guolong Wang, Kaijie Cheng, and Wenwu Zhang. Efficient coupling of dual beam combined laser into micro water jet for deep processing. *Scientific Reports*, 15:1083, 2025. doi: 10.1038/s41598-025-85552-8. URL <https://doi.org/10.1038/s41598-025-85552-8>.
- [178] Zhihe Cao, Hongchao Qiao, Jibin Zhao, and Dongyu Han. Numerical and experimental study on double coaxial gas assisted water jet guided laser machining of superalloy. *Journal of Manufacturing Processes*, 108:321–334, 2023. doi: 10.1016/j.jmapro.2023.11.004. URL <https://doi.org/10.1016/j.jmapro.2023.11.004>.
- [179] Mengjian Wang, Jianwei Wang, Weizhe Wang, Jinhuan Guan, Haoqing Jiang, and Hongxing Xu. Experiments and numerical optimization of water-jet guided laser diamond machining based on the improved nsga-iii algorithm. *Micromachines*, 17(2):206, 2026. doi: 10.3390/mi17020206. URL <https://doi.org/10.3390/mi17020206>.
- [180] Shilong Chen, Tianye Jin, Zihao Lin, Mingtao Wu, Junquan Huang, Junyun Chen, and Anmin Nie. Energy-dependent machining mechanism and process in water-jet guided laser processing single crystal diamond. *Alexandria Engineering Journal*, 118:681–691, 2025. doi: 10.1016/j.aej.2025.01.115. URL <https://doi.org/10.1016/j.aej.2025.01.115>.
- [181] Yuan Li, Shuiwang Wang, Ye Ding, Bai Cheng, Wanda Xie, and Lijun Yang. Investigation on the coaxial-annulus-argon-assisted water-jet-guided laser machining of hard-to-process materials. *Materials*, 16(16):5569, 2023. doi: 10.3390/ma16165569. URL <https://doi.org/10.3390/ma16165569>.

- [182] Zhirong Liao, Dongdong Xu, Dragos Axinte, Jeremie Diboine, and Anders Wretland. Surface formation mechanism in waterjet guided laser cutting of a ni-based superalloy. *CIRP Annals*, 70(1):155–158, 2021. doi: 10.1016/j.cirp.2021.03.007. URL <https://doi.org/10.1016/j.cirp.2021.03.007>.
- [183] Joshua M. Pearce. Building research equipment with free, open-source hardware. *Science*, 337(6100):1303–1304, 2012. doi: 10.1126/science.1228183. URL <https://doi.org/10.1126/science.1228183>.
- [184] Tom Baden, Andre Maia Chagas, Greg Gage, Timothy Marzullo, Lucia L. Prieto-Godino, and Thomas Euler. Open labware: 3-D printing your own lab equipment. *PLoS Biology*, 13(3):e1002086, 2015. doi: 10.1371/journal.pbio.1002086. URL <https://doi.org/10.1371/journal.pbio.1002086>.
- [185] Lothar Opilik, Thomas Schmid, and Renato Zenobi. Modern Raman imaging: Vibrational spectroscopy on the micrometer and nanometer scales. *Annual Review of Analytical Chemistry*, 6:379–398, 2013. doi: 10.1146/annurev-anchem-062012-092646. URL <https://doi.org/10.1146/annurev-anchem-062012-092646>.
- [186] Stephanie Stewart, Ryan J. Priore, Michael P. Nelson, and Patrick J. Treado. Raman imaging. *Annual Review of Analytical Chemistry*, 5:337–360, 2012. doi: 10.1146/annurev-anchem-062011-143152. URL <https://doi.org/10.1146/annurev-anchem-062011-143152>.
- [187] Ewen Smith and Geoffrey Dent. *Modern Raman Spectroscopy: A Practical Approach*. Wiley, Chichester, 2 edition, 2019. ISBN 9781119440550. URL <https://www.wiley-vch.de/en/areas-interest/natural-sciences/modern-raman-spectroscopy-978-1-119-44055-0>.
- [188] Deseada Díaz Barrero, Genrich Zeller, Magnus Schlösser, Beate Bornschein, and Helmut H. Telle. Versatile confocal Raman imaging microscope built from off-the-shelf opto-mechanical components. *Sensors*, 22(24):10013, 2022. doi: 10.3390/s222410013. URL <https://doi.org/10.3390/s222410013>.
- [189] Neil J. Everall. Confocal Raman microscopy: Common errors and artefacts. *Analyst*, 135(10):2512–2522, 2010. doi: 10.1039/C0AN00371A. URL <https://doi.org/10.1039/C0AN00371A>.
- [190] OpenRAMAN. OpenRAMAN: An open source Raman spectrometer, 2026. URL <https://www.open-raman.org/>. Official project documentation; accessed 2026-05-08.
- [191] Seung-Jin Baek, Aaron Park, Young-Jin Ahn, and Jaebum Choo. Baseline correction using asymmetrically reweighted penalized least squares smoothing. *Analyst*, 140(1):250–257, 2015. doi: 10.1039/C4AN01061B. URL <https://doi.org/10.1039/C4AN01061B>.
- [192] Abraham Savitzky and Marcel J. E. Golay. Smoothing and differentiation of data by simplified least squares procedures. *Analytical Chemistry*, 36(8):1627–1639, 1964. doi: 10.1021/ac60214a047. URL <https://doi.org/10.1021/ac60214a047>.
- [193] ASTM International. ASTM E1840-96(2022): Standard Guide for Raman Shift Standards for Spectrometer Calibration, 2022. URL <https://store.astm.org/e1840-96r22.html>.
- [194] Maryam AlShehab, Saba Siadat Mousavi, Maude Amyot-Bourgeois, Jaspreet Walia, Anthony Olivieri, Behnood Ghamsari, and Pierre Berini. Design and construction of a Raman microscope and characterization of plasmon-enhanced Raman scattering in graphene. *Journal of the Optical Society of America B*, 36(8):F49–F59, 2019. doi: 10.1364/JOSAB.36.000F49. URL <https://doi.org/10.1364/JOSAB.36.000F49>.

- [195] Christian Mohr, Claire L. Spencer, and Michael Hippler. Inexpensive Raman spectrometer for undergraduate and graduate experiments and research. *Journal of Chemical Education*, 87(3):326–330, 2010. doi: 10.1021/ed800081t. URL <https://doi.org/10.1021/ed800081t>.
- [196] OpenRAMAN. OpenRAMAN starter and performance editions, 2026. URL <https://www.open-raman.org/build/>. Official project documentation; accessed 2026-05-08.
- [197] Hardeep Sharma, Marisa K. Bird, Sierra Bruinsma, Maricruz Velazquez-Rivera, An Vu, Ginny Goodwin, Margaret Kim, Leonardo San Emeterio, and Jacqueline Briggs. AutoOpenRaman: An automated, open-source Raman spectrometer, 2025. URL <https://doi.org/10.57844/arcadia-7vbd-n3ry>. Open research resource.
- [198] Renishaw. inVia confocal Raman microscope, 2026. URL <https://www.renishaw.com/en/invia-confocal-raman-microscope--6260>. Official product page; accessed 2026-05-08.
- [199] HORIBA Scientific. LabRAM Odyssey raman microscope, 2026. URL <https://www.horiba.com/int/scientific/products/detail/action/show/Product/labram-odyssey-raman-microscope-1848/>. Official product page; accessed 2026-05-08.
- [200] Oxford Instruments. witec360 raman imaging microscope, 2026. URL <https://raman.oxinst.com/products/witec360/>. Official product page; accessed 2026-05-08.
- [201] Thermo Fisher Scientific. DXR3xi raman imaging microscope, 2026. URL <https://www.thermofisher.com/order/catalog/product/IQLAADGAAGFAHDMVX>. Official product page; accessed 2026-05-08.
- [202] Bruker. SENTERRA II raman microscope, 2026. URL <https://www.bruker.com/en/products-and-solutions/infrared-and-raman/raman-spectrometers/senterra-ii.html>. Official product page; accessed 2026-05-08.
- [203] Bruker. RAMANtouch raman microscope, 2026. URL <https://www.bruker.com/en/products-and-solutions/infrared-and-raman/raman-microscopes/ramantouch.html>. Official product page; accessed 2026-05-08.
- [204] Metrohm. MIRA XTR DS raman spectrometer, 2026. URL <https://www.metrohm.com/en/products/raman-spectroscopy/handheld-raman/mira-xtr-ds.html>. Official product page; accessed 2026-05-08.
- [205] Metrohm. IBEX site assessment robot, 2026. URL <https://www.metrohm.com/en/products/raman-spectroscopy/ibex.html>. Official product page; accessed 2026-05-08.
- [206] Jason Sparapani. Lasers, robots, action: MIT workshop explores Raman spectroscopy, March 2026. URL <https://news.mit.edu/2026/lasers-robots-action-mit-workshop-explores-raman-spectroscopy-0320>. MIT News, published March 20, 2026; accessed 2026-05-14.
- [207] Huanyu Ren, Jiajun Zhang, Liqi Peng, Huihui Li, and Zhi Su. Multirole 3D-printed modular spectrometer for various teaching spectral experiments in the classroom and at home. *Journal of Chemical Education*, 100(11):4307–4317, 2023. doi: 10.1021/acs.jchemed.3c00406. URL <https://doi.org/10.1021/acs.jchemed.3c00406>.

- [208] Ocean Insight. QE Pro spectrometer, 2026. URL <https://www.oceaninsight.com/products/spectrometers/high-sensitivity-spectrometers/qe-pro/>. Official product page; accessed 2026-05-08.
- [209] Avantes. AvaSpec-HERO spectrometer, 2026. URL <https://www.avantes.com/products/spectrometers/sensline/avaspec-hero/>. Official product page; accessed 2026-05-08.
- [210] Ibsen Photonics. FREEDOM spectrometers, 2026. URL <https://ibsen.com/products/freedom-spectrometers/>. Official product page; accessed 2026-05-08.
- [211] Hamamatsu Photonics. C12880MA mems micro-spectrometer, 2026. URL <https://www.hamamatsu.com/us/en/product/optical-sensors/spectrometers/mini-spectrometer/C12880MA.html>. Official product page; accessed 2026-05-08.
- [212] Wasatch Photonics. Cobra-S oct spectrometer, 2026. URL <https://wasatchphotonics.com/product/cobra-s-oct-spectrometer/>. Official product page; accessed 2026-05-08.
- [213] Shelyak Instruments. Sol'Ex and Star'Ex spectrographs, 2026. URL <https://www.shelyak.com/produit/solex-english/?lang=en>. Official product page; accessed 2026-05-08.
- [214] Shelyak Instruments. eShel II high resolution echelle spectrograph, 2026. URL <https://www.shelyak.com/produit/eshel/?lang=en>. Official product page; accessed 2026-05-08.
- [215] PhysicsOpenLab. Webcam DIY spectrometer, 2015. URL <https://physicsopenlab.org/2015/09/30/webcam-diy-spectrometer/>. Project documentation; accessed 2026-05-08.
- [216] Hyejeong Jeong, Suyeon Shin, Jihun Hwang, Yoon-Jin Kim, and Sungyoung Choi. Open-source fluorescence spectrometer for noncontact scientific research and education. *Journal of Chemical Education*, 98(11):3493–3501, 2021. doi: 10.1021/acs.jchemed.1c00560. URL <https://doi.org/10.1021/acs.jchemed.1c00560>.
- [217] Pavel V. Kolesnichenko, Axl Eriksson, Linnea Lindh, Donatas Zigmantas, and Jens Uhlig. Viking spectrophotometer: A home-built, simple, and cost-efficient absorption and fluorescence spectrophotometer for education in chemistry. *Journal of Chemical Education*, 100(3):1128–1137, 2023. doi: 10.1021/acs.jchemed.2c00679. URL <https://doi.org/10.1021/acs.jchemed.2c00679>.
- [218] Ryan Bogucki, Mary Greggila, Paul Mallory, Jiansheng Feng, Kelly Siman, Banafsheh Khakipoor, Hunter King, and Adam W. Smith. A 3D-printable dual beam spectrophotometer with multiplatform smartphone adaptor. *Journal of Chemical Education*, 96(7):1527–1531, 2019. doi: 10.1021/acs.jchemed.8b00870. URL <https://doi.org/10.1021/acs.jchemed.8b00870>.
- [219] Elise K. Grasse, Morgan H. Torcasio, and Adam W. Smith. Teaching UV–Vis spectroscopy with a 3D-printable smartphone spectrophotometer. *Journal of Chemical Education*, 93(1):146–151, 2016. doi: 10.1021/acs.jchemed.5b00654. URL <https://doi.org/10.1021/acs.jchemed.5b00654>.
- [220] William Kai Langford. *Discrete Robotic Construction*. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, September 2019. URL <https://cba.mit.edu/docs/theses/19.09.Langford.pdf>. Doctoral thesis; accessed 2026-05-15.