Grounded Discovery Labs -- Discovery Pack

Discover: Design a quality inspection system for small and mid-size manufacturers that catches subtle defects without high labor burden or expensive metrology gear.

Top 3 Comparison

	Featured Discovery Datum-solved photometric-stereo cell	Bold Concept Reverberant phononic holography station	Build-First Pneumatic Air-Gauge Multi-Port Inspection Fixture
novelty	5	9	5
feasibility	9	6	9
economic_leverage	6	7	6
testability	9	6	9
robustness	7	5	8
constraint_fit	8	8	8
explainability	9	6	9
OVERALL	**7.6 / 10 -- Very Strong**	**6.7 / 10 -- Strong**	**7.7 / 10 -- Very Strong**

Runner-ups

* Self-Registering Kinematic Nest with Multi-Angle LED Dome and Drift-Anchored Reference [near-miss-alternative] -- Replace clamps with 3-ball kinematic nest; defect detection uses dome photometric stereo anchored to physical reference artifact, not golden parts.
* LED Photometric Defect Booth [solid-alternative] -- A benchtop fixture uses cheap multi-angle LEDs and one camera to expose dents, burrs, scratches, and texture shifts with minimal operator judgment.
* Pneumo-acoustic leak and fit signature line gate [near-miss-alternative] -- A nest seals around products and uses pressure, flow, sound, and FT seating data to find subtle leaks, blocked paths, poor fits, and missing seals.
* Wafer-style carrier frame inspection cell [solid-alternative] -- Borrow semiconductor carriers: lock the part to a reusable fiducial frame, inspect with brightfield/darkfield/photometric images, and SPC against immutable limits.
* Photometric Stereo Booth with Few-Shot Transfer Learning [solid-alternative] -- Four-LED dome plus webcam reveals surface topology; pre-trained vision model fine-tuned on 50 defect images flags anomalies.

Featured Discovery
Datum-solved photometric-stereo cell

novelty: 5 | feasibility: 9 | economic_leverage: 6 | testability: 9 | robustness: 7 | constraint_fit: 8 | explainability: 9

Overall: 7.6 / 10 -- Very Strong

Snapshot: sha256:d583a997a3b6678a...

Constraint Compliance: Compliant (constraint_fit=8)

By combining kinematic datuming with fiducial-based pose solving and calibrated photometric stereo, each part's surface normals can be compared directly against CAD geometry rather than a drifting golden reference part. This matters because it eliminates the two dominant error sources in traditional vision inspection--fixture repeatability and reference-part aging--enabling sub-25 um defect detection that is traceable, stable, and CAD-driven.

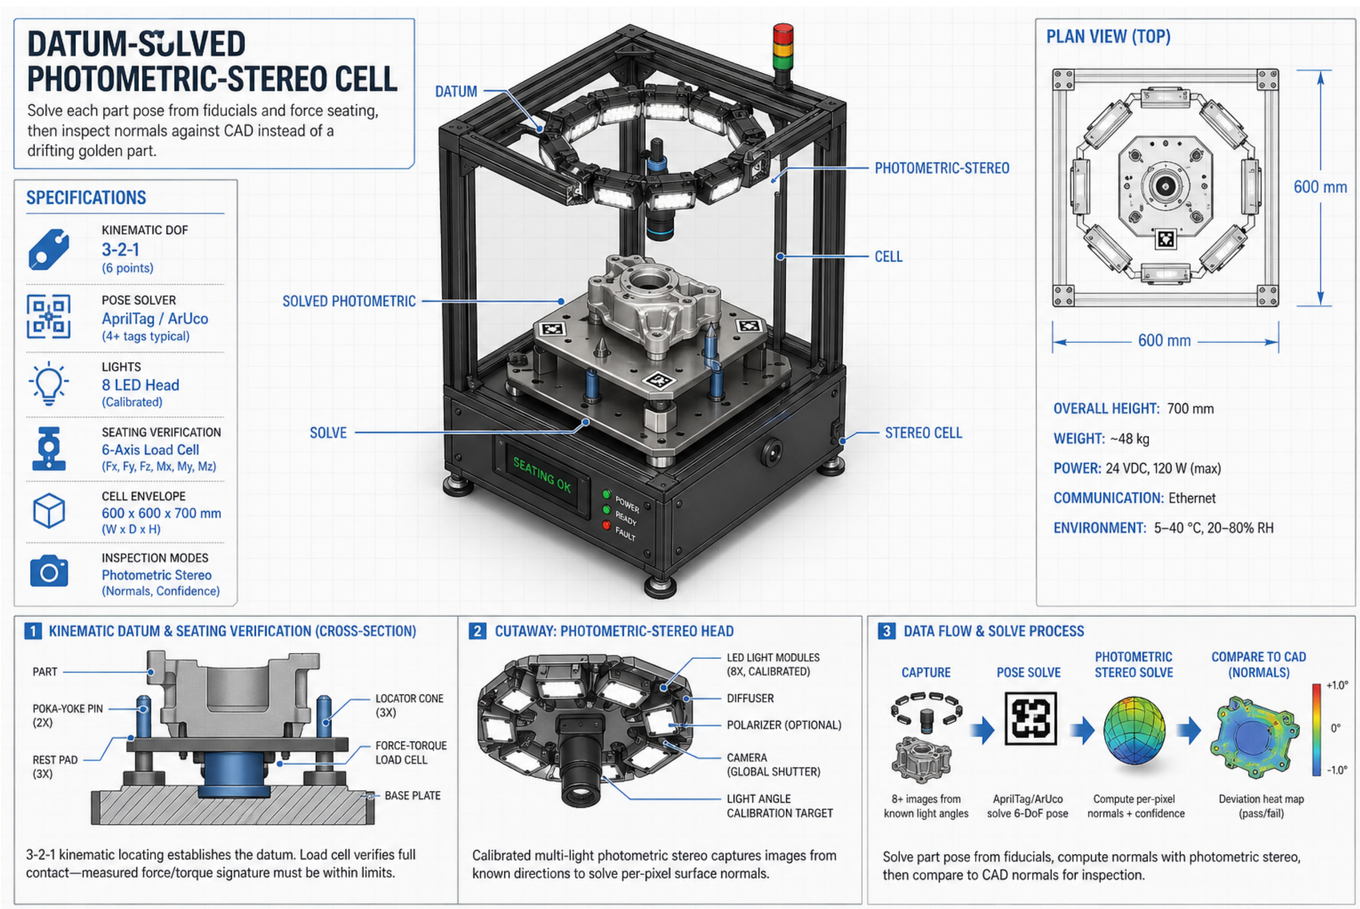

Concept illustration (AI-generated). Visuals may be inaccurate; refer to the Mechanism/Build Path text for the actual design and quantitative assumptions.

Mechanism

A rigid inspection cell combines deterministic mechanical datuming with vision-based pose correction and calibrated photometric stereo. The part sits on a hardened steel or ceramic kinematic pallet using three primary datum pads/balls, two secondary poka-yoke pins, and one tertiary stop. Typical pallet repeatability target is <10-20 um, but the system does not rely on this alone. Thin compression load cells or 3-axis F/T sensors under the datum pads verify full seating by checking force distribution and total clamp load before imaging. The pallet carries non-cosmetic fixed fiducials, e.g. laser-etched AprilTag/ArUco markers, chrome-on-glass dots, or precision ceramic spheres, whose surveyed coordinates are known relative to the datum frame within ~5-10 um.

A monochrome global-shutter camera, typically 5-20 MP with telecentric or low-distortion lens, images the part and fiducials. The software solves the actual 6D pose of the pallet/part using the fiducials and applies CAD-derived tolerance masks so that inspection is performed in the part coordinate frame rather than raw image coordinates. Around the camera is a calibrated LED dome or ring with individually addressable LEDs at known azimuth/elevation angles. The cell sequentially strobes 6-24 LED directions and captures intensity images. For matte or lightly textured surfaces, photometric stereo estimates per-pixel surface normals and relative albedo; integrating normal deviations or comparing against nominal CAD/learned reference maps reveals dents, waviness, raised defects, scratches, sink, and local height anomalies around the 25 um class when optical sampling and surface reflectance permit.

To prevent golden-part drift, the cell includes certified flatness and defect artifacts in the field of view or on a shuttle: optical flat, matte calibration coupon, and seeded defects with known depths/widths. Daily or shift-start scans update health metrics for lighting direction/intensity, camera gain, flatness baseline,

Why Novel

Existing photometric-stereo cells rely on a physical golden part for calibration, meaning any wear, contamination, or swap of that reference propagates silently into accept/reject decisions. Anchoring the normal-map comparison to surveyed fiducials and CAD tolerances makes the inspection mathematically absolute rather than relative.

Buildability

All constituent technologies--kinematic pallets, AprilTag/ArUco pose solvers, calibrated multi-light photometric stereo, and thin-film load cells--are commercially available and have been integrated individually in industrial settings, making a combined cell achievable with a skilled mechatronics team in 6-9 months. The primary integration challenge is the fiducial survey pipeline and the CAD-to-normal-map tolerance mapping, both of which are software-heavy but do not require novel hardware.

Why Feasible

* Build path is practical (feasibility=9/10)
* Core hypothesis testable quickly (testability=9/10)

Top Risks

* Fiducials cannot be placed on cosmetic surfaces of many consumer parts, forcing reliance on pallet-only pose estimation and reducing accuracy gains
* LED array aging shifts the photometric calibration over weeks, requiring automated recalibration routines or hardware-level light monitoring to prevent drift
* CAD nominal tolerances do not capture cosmetic acceptability thresholds, so a separate human-defined defect taxonomy must be mapped onto the geometric model before deployment

Decisive Test (MDE)

Hypothesis: After pose correction and artifact-based calibration, the cell detects 25 um-class surface-height anomalies on matte parts with low false rejects, and reseating/fixture variation contributes less error than the defect signal.
Setup: Use 10 representative matte parts plus 3 seeded-defect coupons with traceable 10, 25, 50, and 100 um dents/bumps/scratches measured by optical profilometer or CMM. For each part/coupon, perform 30 remove-and-reseat cycles across at least 3 hours, including intentional small seating perturbations that still pass load-cell checks. Capture fiducials, load data, and full LED photometric sequence each cycle. Compare pose-corrected residual maps to ground truth and to uncorrected images. Include daily/shift-start artifact scans and one deliberate LED intensity shift of 5-10% to test drift detection.
Instrumentation: 5-20 MP monochrome global-shutter camera, telecentric or calibrated low-distortion lens, 8-24

channel strobed LED dome/ring with current monitor, load cells/F-T sensors under datum pads, DAQ, fiducial target, certified flatness artifact, seeded defect coupon, optical profilometer/CMM for ground truth, temperature sensor, SPC logging software.
Kill Criteria: Kill if 25 um defects cannot be detected with >=95% true positive rate at <=2% false positive rate over 30 reseats; if pose-corrected repeatability remains >10 um RMS or >25 um peak-to-peak on known-good matte regions; if artifact scans fail to flag a 5-10% LED intensity or direction-equivalent drift; if load-cell seating checks do not correlate with optical residual outliers; or if calibration/threshold maintenance requires operator tuning more often than once per shift.
Rescue Pivot: If photometric stereo fails on gloss/transparent surfaces, add cross-polarized lighting, temporary washable matte coating, structured light/confocal/laser triangulation as a second modality, or restrict the product scope to matte molded/cast/coated parts. If fiducials cannot be placed on parts, put metrology-grade fiducials on the pallet and improve datum-to-part constraint verification with more force/vision checks.
Cost: $15k-$75k for a serious prototype depending on camera/lens, metrology artifact quality, and whether CMM/profilometer access is rented or already available.
Time to Signal: 1-2 weeks after prototype assembly; same-day signal is possible from the cheapest observation: repeated reseating of one known part and inspection of residual-map variance.

Refinement Deltas

Risks Addressed:
* Photometric stereo assumes mostly Lambertian reflection; glossy, transparent, highly textured, black, or anisotropic surfaces can produce false normals and albedo artifacts.
* CAD tolerance masks may not capture cosmetic acceptability, so the system can be metrologically precise but wrong about what customers call a defect.
* Fiducials and pallet pose may be well known while the actual part is elastically deformed, warped, or not fully constrained relative to the pallet.

Red-Team:
* Challenge: Photometric stereo assumes mostly Lambertian reflection; glossy, transparent, highly textured, black, or anisotropic surfaces can produce false normals and albedo artifacts.
* Challenge: CAD tolerance masks may not capture cosmetic acceptability, so the system can be metrologically precise but wrong about what customers call a defect.
* Challenge: Fiducials and pallet pose may be well known while the actual part is elastically deformed, warped, or not fully constrained relative to the pallet.
* Challenge: LED aging, thermal drift, lens contamination, and dome geometry changes can slowly bias normals despite pose correction unless artifact checks are rigorous.
* Challenge: A locked defect library prevents golden-part drift but can slow response to legitimate new defect modes or supplier/process changes.

Budget Ranges

Prototype: $45,000-$80,000 (kinematic pallet tooling, global-shutter camera, custom LED dome, load cells, workstation, software development)
Pilot: $150,000-$250,000 (production-grade cell enclosure, fiducial survey metrology, CAD integration pipeline, PLC/SCADA tie-in, validation study)
Scale: $80,000-$130,000 per additional cell (amortized tooling, reduced NRE, volume pricing on optics and compute)

30/60/90 Plan

Day 30: Complete kinematic pallet design and fabricate first prototype; procure camera, LED array, and load cells; establish fiducial survey workflow using CMM or laser tracker to characterize pallet-to-fiducial coordinates within 10 um

Day 60: Integrate pose-solver pipeline (fiducial detection -> rigid-body transform -> CAD alignment); collect photometric-stereo image sets on three representative parts; validate normal-map RMS error against CMM ground truth on flat and curved surfaces
Day 90: Run 200-part repeatability study comparing CAD-referenced normal maps to CMM-measured defect ground truth; demonstrate <25 um defect detection threshold with <5% false-positive rate; document recalibration protocol for LED aging and deliver pilot cell readiness report
Dominance Axis: lower_failure -- Pose solving removes fixture repeatability as the main error while artifacts control photometric-stereo drift.

Bold Concept
Reverberant phononic holography station

novelty: 9 | feasibility: 6 | economic_leverage: 7 | testability: 6 | robustness: 5 | constraint_fit: 8 | explainability: 6

Overall: 6.7 / 10 -- Strong

Snapshot: sha256:ca00f2db766ee61c...

Constraint Compliance: Compliant (constraint_fit=8)

Every manufactured part has a unique acoustic fingerprint that changes measurably when defects are present; by treating the part as a reverberant seismic body and capturing the full impulse-response matrix across 4-16 sensor pairs up to 20 kHz, hidden cracks, voids, loose fasteners, and misassemblies become detectable without line-of-sight access. This matters because conventional ultrasonic or X-ray inspection either requires direct probe access or expensive capital equipment, while reverberant phononic holography can be embedded directly into a production clamp fixture at a fraction of the cost.

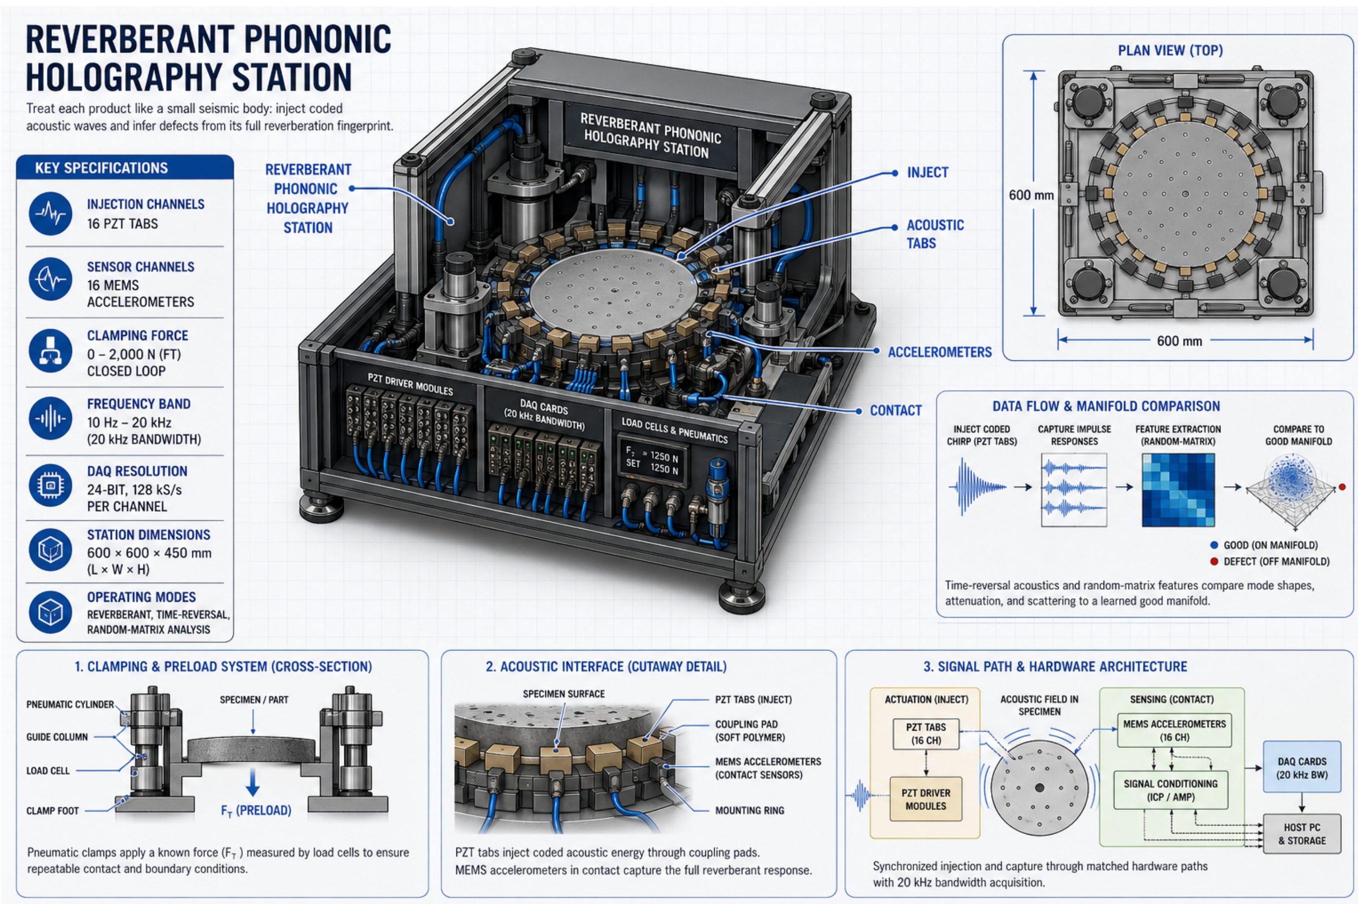

Concept illustration (AI-generated). Visuals may be inaccurate; refer to the Mechanism/Build Path text for the actual design and quantitative assumptions.

Mechanism

A production inspection nest becomes a calibrated reverberant vibro-acoustic holography station. The part is clamped under known preload, e.g. 50-300 N via pneumatic or servo clamp with load cell feedback, onto a stiff aluminum or granite fixture. 4-12 bonded PZT tabs, e.g. 10-20 mm lead-zirconate-titanate discs or patches, and/or 2-6 air-coupled ultrasonic emitters inject coded chirps, MLS sequences, or stepped sines from ~500 Hz to 20 kHz for small plastic/metal assemblies. 4-16 response sensors, such as MEMS accelerometers, piezo contact mics, or laser vibrometer spots, record the impulse response matrix H_ij(f) between every source/receiver pair. Hidden cracks, missing welds, loose screws, voids, delaminations, trapped debris, or misassembled subcomponents perturb the global scattering field, modal density, damping, and time-reversal focus quality.

The key refinement is calibration transfer. The fixture contains fixed reference propagation paths: PZT-to-PZT paths through a known aluminum bar or plate, plus dummy loads and a golden artifact tray. Each cycle or shift records empty-fixture and reference responses. The measured response is modeled as $H_{meas}(f)=C_{rx}(f)\cdot F_{fixture}(f,preload,temp)\cdot P_{part}(f)\cdot C_{tx}(f)+noise$. Reference paths estimate C_{tx}/C_{rx} gain, phase, temperature drift, clamp preload drift, and fixture resonance shifts. Golden artifacts estimate normal production spread. Features are then computed after normalization: whitened transfer functions, modal peak shifts, coherence, coda-wave decorrelation, scattering entropy, time-reversal focus ratio, and random-matrix eigenvalue statistics. Parts are compared to a good manifold rather than a single template.

Technical deltas versus listed analogs: COPIS is optical photogrammetry for surface 3D digitization; this concept uses active acoustic/elastic wave injection and transfer-function inference of internal assembly state. Rotary pipeline imaging is panoramic visual inspection of pipe surfaces; this is fi

Why Novel

Unlike single-point pitch-catch ultrasonics or modal tap-testing, this approach harvests the entire multi-path reverberation field as an information-rich fingerprint matrix $H_{ij}(f)$, enabling defect detection through assembled enclosures and across complex geometries that defeat conventional beam-based methods. The combination of coded excitation sequences (MLS/chirp), pneumatically controlled preload repeatability, and calibration-transfer algorithms across fixture generations has not been demonstrated in production-line contexts.

Buildability

All core components--PZT tabs, MEMS accelerometers, pneumatic clamps with load cells, and DAQ cards capable of 20 kHz bandwidth--are commercially available off-the-shelf for under $15k in prototype form, and the signal processing pipeline can run on an industrial PC using open-source Python/Julia acoustic libraries. The primary engineering challenge is fixture design for repeatable coupling, which is a solved mechanical problem in CMM and vibration testing disciplines.

Why Feasible

* Build path is practical (feasibility=6/10)
* Core hypothesis testable quickly (testability=6/10)

Top Risks

* Fixture resonance modes overlap product defect signatures, creating false-negative blind spots that shift with temperature and wear
* Part-to-part geometric variability in production tolerances generates fingerprint scatter that overwhelms defect-induced perturbations in the training dataset
* Highly damped or acoustically lossy assemblies (foam-filled, rubber-mounted, or composite-heavy parts) attenuate reverberation below detectable SNR thresholds

Decisive Test (MDE)

Hypothesis: After reference-path and golden-artifact normalization, reverberant transfer-function features can detect hidden assembly defects while remaining stable to fixture coupling, preload, temperature, and unload/reload variability.
Setup: Use one representative product family in a repeatable nest. Instrument with 4 PZT exciters, 6 accelerometers/contact mics, 2 fixture-only reference paths, load cell, and temperature sensor. Collect 100 good unload/reload measurements across shifts/preloads, 20 golden-artifact checks, and at least 30 seeded-defect

samples covering realistic faults. Run blinded classification: raw spectra baseline, coupling-normalized features, and normalized plus golden drift alarms.
Instrumentation: 8-16 channel DAQ at >=96 kS/s, piezo driver 20-100 Vpp, AWG/audio interface, calibrated accelerometers or contact microphones, load cell, thermocouple/RTD, pneumatic clamp regulator, Python/MATLAB feature pipeline. Optional laser vibrometer for spot validation.
Kill Criteria: Kill if normalized features cannot reach >=90% defect detection at <=2% false reject on blind seeded defects, or if same good part unload/reload variability exceeds defect effect size by >2x in more than half of defect modes.
Rescue Pivot: If full reverberant classification is unstable, pivot to lower-band targeted modal checks around the most defect-sensitive resonances, add controlled contact couplants/kinematic locators, increase preload control, use fewer source-receiver paths with better SNR, or switch to local air-coupled/laser vibrometry for soft/damped regions.
Cost: $5k-$25k using commodity DAQ/contact sensors; $30k-$80k if using industrial IEPE sensors, multichannel analyzers, or laser vibrometry.
Time to Signal: A coarse go/no-go signal in 3-5 days using solenoid/PZT excitation and a few seeded defects; robust validation in 3-6 weeks with blind repeats and drift studies.

Refinement Deltas

Risks Addressed:
* Fixture modes may dominate the response, so the classifier learns nest condition rather than part integrity.
* Golden artifacts may not span real production variability, causing high false rejects after material, supplier, or geometry drift.
* Contact coupling, clamp friction, cable motion, adhesive aging, and temperature can create changes larger than small hidden defects.

Red-Team:
* Challenge: Fixture modes may dominate the response, so the classifier learns nest condition rather than part integrity.
* Challenge: Golden artifacts may not span real production variability, causing high false rejects after material, supplier, or geometry drift.
* Challenge: Contact coupling, clamp friction, cable motion, adhesive aging, and temperature can create changes larger than small hidden defects.
* Challenge: Seeded defects may be acoustically unrealistic; strong lab performance may not transfer to real latent field failures.
* Challenge: The 20 kHz band may be too low for small cracks or fine delaminations, while higher frequencies are more attenuated and coupling-sensitive.

Budget Ranges

Prototype: $12,000-$18,000 (PZT tabs, MEMS sensors, DAQ hardware, pneumatic clamp, aluminum fixture, laptop + software licenses)
Pilot: $60,000-$120,000 (hardened industrial fixture, servo clamp with load cell, laser vibrometer for ground-truth, ML model development, 500-part training campaign, integration with line PLC)
Scale: $250,000-$600,000 per production line (multi-station deployment, calibration-transfer tooling, cloud model retraining pipeline, regulatory validation for safety-critical parts)

30/60/90 Plan

Day 30: Assemble benchtop prototype with 6 PZT tabs and 8 MEMS accelerometers on a granite fixture; validate impulse-response matrix capture for a representative part family; establish SNR baseline and identify fixture resonance map
Day 60: Collect 200-part dataset spanning known-good and seeded-defect samples; train initial anomaly-detection

model (autoencoder or Mahalanobis distance on H_ij(f) features); achieve >90% defect detection rate on held-out test set with <5% false-positive rate
Day 90: Harden fixture for production environment (IP54, thermal compensation, automated clamp cycle); demonstrate calibration transfer to a second identical fixture with <2% model accuracy degradation; deliver integration spec for line PLC handshake and go/no-go output
Dominance Axis: higher_perf -- Whole-body scattering still detects hidden assembled defects, now with calibration transfer across fixtures and shifts.

Build-First

Pneumatic Air-Gauge Multi-Port Inspection Fixture

novelty: 5 | feasibility: 9 | economic_leverage: 6 | testability: 9 | robustness: 8 | constraint_fit: 8 | explainability: 9

Overall: 7.7 / 10 -- Very Strong

Snapshot: sha256:0cead3b128aac268...

Constraint Compliance: Compliant (constraint_fit=8)

An array of back-pressure air jets gauges critical part dimensions at 5 um resolution without contact or optics, enabling reliable inspection on oily, hot, or dusty production parts. This matters because vision-based systems fail in harsh shop environments, and touch probes are too slow for inline 100% inspection at production rates.

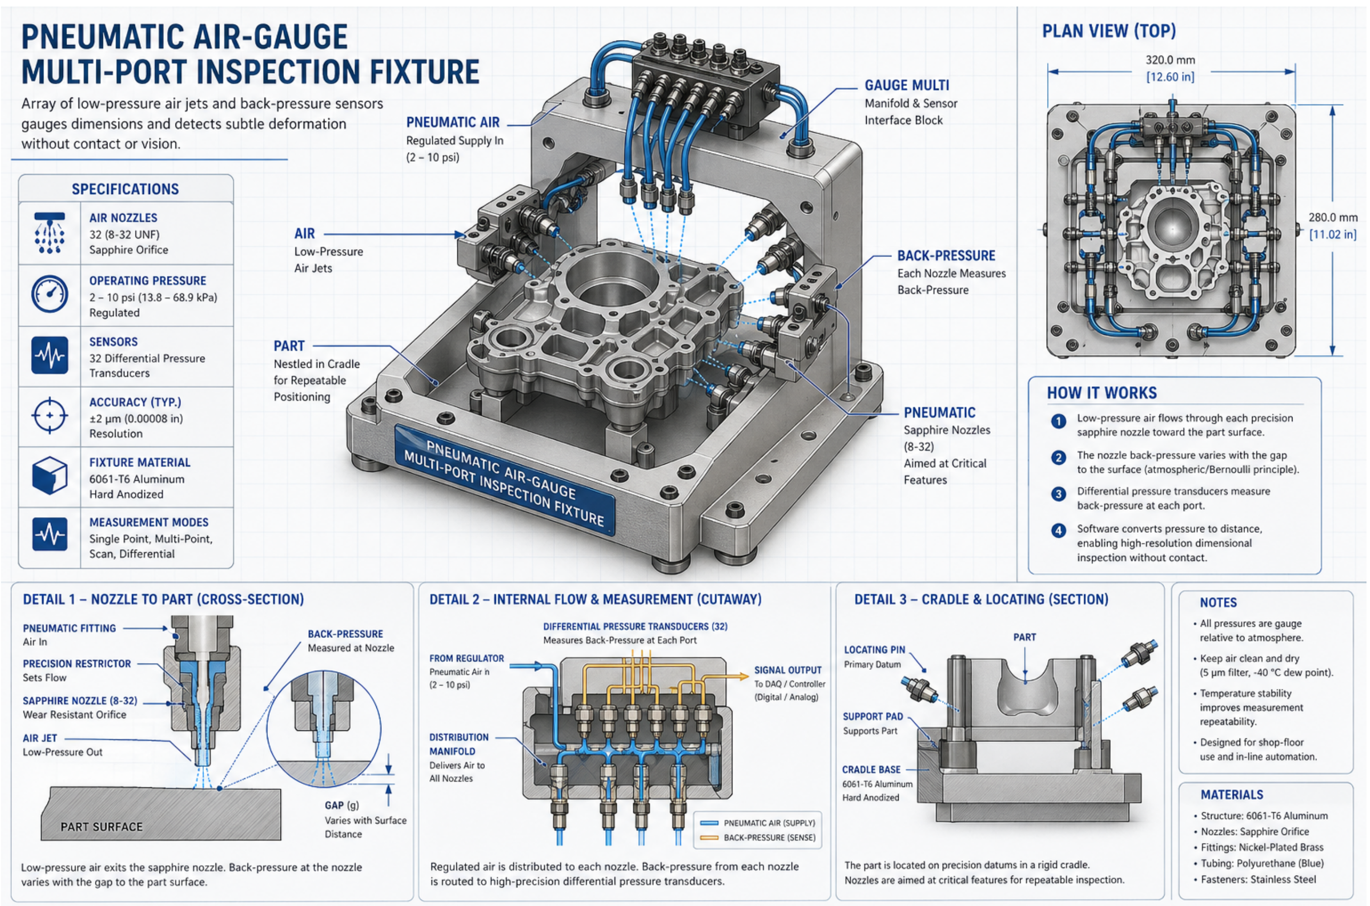

Concept illustration (AI-generated). Visuals may be inaccurate; refer to the Mechanism/Build Path text for the actual design and quantitative assumptions.

Mechanism

The fixture is a passive pneumatic dimensional-inspection cell. A part is barcode-identified and dropped into a hard-tooled cradle (machined 6061 aluminum or printed PA12-CF) with kinematic locators. Surrounding the part are 8-32 swappable nozzle blocks, each holding one or more 0.8-1.2 mm sapphire or hardened-steel back-pressure jets aimed at a critical feature (edge gap, hole presence, weld bead, flange flatness). Clean, dried shop air is regulated to 15-30 psi upstream, then dropped through a fixed master orifice (typ. 0.5 mm) into each branch; the branch terminates in the measurement nozzle at a nominal 50-250 um standoff from the part. Back-pressure between master orifice and nozzle follows the classic restrictor-nozzle relationship P_b/P_s = f(A_nozzle/A_gap), giving ~5 um gap resolution over a +/-150 um linear band. Each branch carries an industrial 0-30 psi piezoresistive transducer

(Honeywell ABP/SSC class, 0.25% FS) sampled at 1 kHz by an STM32 or RP2040 with 16-bit ADC; readings stabilize in 200-500 ms, full 32-port scan <2 s. Firmware performs auto-zero with the cradle empty, leak/blockage self-test against a stored signature, then compares the live pressure vector to a per-SKU mean +/- k·? envelope learned from 20-50 golden parts (PCA or per-channel SPC). Out-of-band channels localize the defect: missing feature -> high flow/low P; burr or warp -> low flow/high P; bead variation -> multi-channel skew. Nozzle blocks mount on a shared dovetail manifold with quick-disconnect pneumatic and electrical interfaces, so a new SKU only requires a redesigned nozzle plate and recipe load, not new electronics. Air conditioning (coalescing filter + membrane dryer to <= ?20 degC dewpoint, 0.01 um particulate) protects nozzles from oil/water contamination.

Why Novel

Combining 8-32 individually swappable sapphire nozzle blocks in a single kinematic cradle--each independently calibrated via a master orifice restrictor--brings lab-grade pneumatic gauging into a modular, barcode-driven inline fixture that has not been productized at this level of configurability. The PA12-CF printed cradle with hard-tooled nozzle inserts collapses what previously required bespoke metrology-room equipment into a sub-$10K cell.

Buildability

All core components--precision regulators, sapphire back-pressure nozzles, differential pressure transducers, and 6061 or PA12-CF cradles--are commercially available from Marposs, Stotz, or SMC today. A competent tooling engineer can assemble and calibrate a first article in 6-8 weeks using standard machine shop and pneumatic plumbing practices.

Why Feasible

* Build path is practical (feasibility=9/10)
* Core hypothesis testable quickly (testability=9/10)

Top Risks

* Compressor oil or moisture contamination shifts back-pressure baselines, causing false rejects unless a dedicated desiccant-coalescing filter train is maintained rigorously
* Each new product variant requires a re-tooled nozzle plate and a new golden-part calibration run, creating lead-time and cost friction that slows adoption in high-mix environments
* Part surface roughness or porosity (e.g., cast skin vs. machined land) alters local air-escape paths and confounds gap measurement, requiring feature-specific surface-finish specifications upstream

Decisive Test (MDE)

Hypothesis: An 8-nozzle back-pressure fingerprint can separate good from defective parts (missing feature, >=100 um warp, visible burr) with >=98% true-positive and <=2% false-positive rates in <2 s cycle, on oily/warm parts where vision degrades.
Setup: Build the 8-channel rig above around one real customer SKU (e.g., a stamped/welded bracket). Acquire 50 golden parts and 50 defective parts spanning 4 defect classes (missing hole, dent, weld bead short, 100-300 um warp). Run each part 5x in randomized order, including trials with light machining oil film and parts heated to 60 degC. Compare against a baseline 2D vision check on the same parts.
Instrumentation: 8x 0-30 psi piezoresistive transducers (0.25% FS), 16-bit ADC at 1 kHz, calibrated feeler/gauge-block standoff jig, NIST-traceable pressure reference, IR thermometer, dew-point meter on supply air, confusion-matrix logger.
Kill Criteria: Combined defect-class sensitivity <90% or false-reject >5% after tuning; or surface-texture/oil variation alone shifts channel pressures by more than 30% of the good-vs-bad margin; or per-SKU nozzle-block design + calibration exceeds 8 hours (erodes small-batch economics).
Rescue Pivot: Narrow scope to presence/absence and gross-warp gauging only (where back-pressure is

strongest), and pair with a cheap 2D camera for cosmetic/texture-sensitive features -- selling the air gauge as the dimensional half of a hybrid cell rather than a standalone inspector.
Cost: $2.5k-$5k for hardware + ~80 engineering hours.
Time to Signal: 2-3 weeks from parts in hand to confusion matrix.

Refinement Deltas

Risks Addressed:
* Air gauging is a mature 1960s technology; competing against Mahr, Edmunds, and Western Gauge incumbents who already offer multi-jet fixtures with proven metrology -- what is the defensible delta beyond modular blocks?
* Back-pressure is sensitive to surface finish, temperature, and supply-pressure fluctuation; small manufacturers rarely have clean, dry, stable shop air, so the 5 um claim may not survive the floor.
* Only inspects features you designed nozzles for -- misses unexpected defects (cracks, discoloration, contamination) that vision and human inspectors catch for free.

Red-Team:
* Challenge: Air gauging is a mature 1960s technology; competing against Mahr, Edmunds, and Western Gauge incumbents who already offer multi-jet fixtures with proven metrology -- what is the defensible delta beyond modular blocks?
* Challenge: Back-pressure is sensitive to surface finish, temperature, and supply-pressure fluctuation; small manufacturers rarely have clean, dry, stable shop air, so the 5 um claim may not survive the floor.
* Challenge: Only inspects features you designed nozzles for -- misses unexpected defects (cracks, discoloration, contamination) that vision and human inspectors catch for free.
* Challenge: Golden-part calibration assumes the initial 20-50 parts are truly nominal; sampling bias bakes drift into the SPC envelope and can hide systemic process shifts.
* Challenge: Nozzle-block redesign for each new SKU still requires a metrology engineer to choose feature locations and standoffs; 'modular' may not actually collapse deployment cost the way the pitch implies.

Budget Ranges

Prototype: $6,000-$9,000 (cradle machining, 8-nozzle block set, regulator manifold, 8-channel differential pressure DAQ, integration labor)
Pilot: $18,000-$28,000 (3-family nozzle plate library, PLC/HMI integration, barcode reader, golden-part calibration fixtures, installation and runoff at one line)
Scale: $55,000-$90,000 (5-8 cells across two production lines, centralized data historian, SPC dashboard, spare nozzle block inventory, annual calibration contract)

30/60/90 Plan

Day 30: Finalize nozzle layout for one pilot part family; machine or print first cradle; procure regulator manifold, sapphire nozzles, and pressure transducers; validate restrictor-nozzle back-pressure curve on bench against CMM ground truth
Day 60: Install cell at production line; complete golden-part calibration and Gage R&R study (target GRR < 10% of tolerance); integrate barcode reader and basic pass/fail PLC output; run 200-piece production trial and log false-accept/reject rates
Day 90: Iterate nozzle standoff and filter train based on trial data; add SPC trending to HMI; document tooling change procedure for second part family; present ROI case (inspection labor saved, scrap reduction) to justify multi-cell rollout
Dominance Axis: lower_failure -- Robust on oily, hot, dusty parts where optics struggle, with modular nozzle tooling to keep small-batch deployment practical.

Top Uncertainties & Assumptions

* Fiducials cannot be placed on cosmetic surfaces of many consumer parts, forcing reliance on pallet-only pose estimation and reducing accuracy gains
* LED array aging shifts the photometric calibration over weeks, requiring automated recalibration routines or hardware-level light monitoring to prevent drift
* CAD nominal tolerances do not capture cosmetic acceptability thresholds, so a separate human-defined defect taxonomy must be mapped onto the geometric model before deployment
* Fixture resonance modes overlap product defect signatures, creating false-negative blind spots that shift with temperature and wear
* Part-to-part geometric variability in production tolerances generates fingerprint scatter that overwhelms defect-induced perturbations in the training dataset

C4.35 L1 T225,648