

Ensuring Reliability in Consumer LiDAR-Based Spatial Compliance: A Scan Quality Validation Framework for Regulated Residential Care Environments

Chaitanya Appikonda

Independent Researcher

Preprint. Intended for submission to a peer-reviewed journal. Empirical validation is ongoing.

April 7, 2026

Keywords: LiDAR; SLAM; building compliance; HCBS; scan quality validation; spatial measurement; accessibility; point cloud quality; hybrid verification; regulated residential care

Abstract

Consumer-grade LiDAR sensors, particularly those embedded in mobile devices, have demonstrated sufficient accuracy for many building measurement applications under controlled conditions. However, deployment in regulatory compliance contexts introduces a reliability problem that accuracy specifications alone cannot address: when operated by untrained or semi-trained personnel across heterogeneous residential environments, measurement uncertainty expands substantially and non-uniformly. This paper characterizes the distribution-of-quality problem in consumer LiDAR deployment for Home and Community-Based Services (HCBS) facility compliance under 42 CFR 441.301 and the ADA 2010 Standards for Accessible Design, and presents the Scan Quality Validation Layer (SQVL)—a seven-component pre-processing framework designed to make scan-derived compliance findings defensible in regulatory and legal contexts. The SQVL evaluates point density, surface coverage completeness, SLAM drift, level transition integrity, doorway capture quality, lighting artifacts, and occlusion gaps before any geometry extraction or compliance rule evaluation occurs. Findings are classified by an accuracy band taxonomy that stratifies compliance rules according to their spatial measurement requirements, determining whether consumer LiDAR is sufficient, insufficient, or conditionally sufficient for each rule category. A mandatory hybrid verification workflow is triggered for all precision-dimension rules and all threshold-adjacent measurements. The proposed framework draws design principles from quality assurance practices in autonomous vehicle perception systems and medical imaging device qualification. Proposed thresholds are design estimates grounded in instrument specifications and the SLAM literature; empirical calibration across HCBS facility environments is the subject of a forthcoming validation study. This work defines a methodological foundation for spatially defensible HCBS compliance assessment using accessible scanning hardware.

1 Introduction

The Home and Community-Based Services (HCBS) Settings Rule, codified at 42 CFR 441.301, establishes substantive physical environment requirements for Medicaid-funded residential care settings. Taken together with ADA 2010 Standards for Accessible Design, CMS Conditions of Participation (42 CFR 483.70, 483.73), and NFPA 101 Life Safety Code (2021 edition), the regulatory framework governing HCBS residential facilities imposes more than 125 spatially verifiable requirements—conditions whose compliance status depends on measurable physical attributes of the built environment: corridor widths, door clearances, ramp slopes, grab bar placement, egress path continuity, and room areas, among others.

Despite the inherently spatial nature of these requirements, the prevailing method of compliance assessment remains the paper-based or digital checklist completed by an on-site surveyor. This approach produces compliance documentation that is observer-dependent, non-reproducible, and not anchored to a persistent spatial record of the facility [Gajewski et al., 2006]. When a facility is cited for a deficiency, neither the facility operator nor the regulatory agency can typically produce the measurement that produced the finding. When the same facility is re-surveyed after remediation, neither party can establish a geometric baseline to confirm that remediation occurred.

The consequences of measurement error in this context are not limited to administrative burden. Physical environment deficiencies are among the categories that contribute to civil monetary penalty (CMP) enforcement under the CMS survey and certification system. Applying the severity-weighted attribution methodology of Antonova and Zimmerman [Antonova and Zimmerman, 2012], spatial environment deficiencies contribute proportionally to approximately \$9.5M in annual CMP enforcement nationally, representing a median severity-weighted share of approximately 8.8% of total survey scores at penalized facilities. This figure reflects proportional attribution at the survey level—CMS enforcement actions are issued against surveys rather than individual citations—and should be interpreted as an indicator of economic exposure, not a causal claim that spatial

deficiencies independently drive enforcement.

The emergence of consumer-grade LiDAR sensors—notably the LiDAR scanner integrated into Apple iPhone Pro models since 2020—has made mobile spatial capture accessible at a price point compatible with routine compliance workflows. Published accuracy benchmarks for iPhone LiDAR in indoor environments report point-to-surface errors ranging from sub-millimeter under static conditions to approximately 10 mm in dynamic acquisition [Teo and Yang, 2023], with broader indoor mapping studies reporting mean errors in the range of 27–52 mm depending on application and environment [Catharia et al., 2023, Hou et al., 2024]. These figures appear sufficient for many compliance measurements under favorable conditions. This apparent sufficiency has led to proposals for mobile LiDAR-based building inspection and scan-to-BIM workflows [Federal Highway Administration, 2024, Aş Çemrek et al., 2025].

This paper argues that the framing of consumer LiDAR for compliance assessment as primarily an accuracy problem is incorrect. The operative challenge is reliability: the extent to which measurement quality is consistent and predictable across the full distribution of real-world deployment conditions, including variability in operator training, facility geometry, lighting conditions, and surface texture. A system that achieves ± 20 mm under controlled conditions but ± 55 mm under field conditions with an untrained operator—and that does not distinguish these cases—cannot produce defensible compliance findings.

The contribution of this paper is threefold. First, it characterizes the distribution-of-quality problem specific to consumer LiDAR deployed by non-specialist operators in HCBS residential environments. Second, it presents the Scan Quality Validation Layer (SQVL), a seven-component pre-processing framework that asserts per-scan quality verdicts before any compliance evaluation proceeds. Third, it introduces an accuracy band taxonomy that classifies HCBS compliance rules according to their spatial measurement requirements and maps each class to the appropriate combination of scanning hardware and verification workflow.

The remainder of this paper is organized as follows. Section 2 reviews relevant literature.

Section 3 characterizes the distribution-of-quality problem. Section 4 presents the SQVL architecture. Section 5 introduces the accuracy band classification. Section 6 describes the hybrid verification workflow. Section 7 positions the SQVL relative to QA frameworks in adjacent domains. Section 8 discusses limitations and future work. Section 9 concludes.

2 Background and Related Work

2.1 LiDAR-Based Building Measurement

Terrestrial laser scanning (TLS) for building documentation has an established literature dating to the early 2000s [Boehler et al., 2003]. TLS instruments achieve sub-millimeter accuracy under ideal conditions and have been deployed for as-built documentation, heritage preservation, and structural assessment. Comparative studies of TLS against handheld mobile laser scanning systems have confirmed that while TLS retains accuracy advantages, handheld systems offer substantially reduced setup time and operational complexity [Russhakim et al., 2019]. However, TLS instruments require specialized operation, station-by-station acquisition, and post-processing workflows that are incompatible with routine compliance inspection timescales and budgets.

Structured-light and time-of-flight depth sensors at consumer price points—including Microsoft Kinect-class devices and the Intel RealSense family—have been studied for indoor mapping [Henry et al., 2012]. The integration of solid-state LiDAR into mobile devices (Apple LiDAR Scanner, introduced with iPhone 12 Pro in 2020) represents a qualitative shift in accessibility, enabling point cloud capture as a workflow step rather than a specialized operation.

Recent empirical evaluations of iPhone LiDAR for indoor measurement have reported point-to-surface errors of less than 1 mm under static acquisition and approximately 10 mm under dynamic acquisition [Teo and Yang, 2023]. Broader indoor mapping studies report mean errors in the range of 27–52 mm depending on scan conditions and application [Catharia et al., 2023, Hou et al., 2024]. Area measurement errors have been reported

in the range of 1–4% for typical room geometries [Teo and Yang, 2023]. These figures are consistent with accessibility to Band B and Band C compliance rules as defined in this paper (Section 5), provided that measurement uncertainty is properly characterized.

However, published evaluations uniformly report performance under controlled or near-controlled conditions: professionally operated scans, adequate lighting, textured surfaces, and standard room geometries. The emergence of pocket LiDAR for construction inspection and as-built documentation [Federal Highway Administration, 2024] and scan-to-BIM workflows based on mobile point cloud acquisition [Aş Çemrek et al., 2025] has extended the application domain, but these workflows assume professional operators and standard building geometries. Performance under HCBS field conditions—occupied residential spaces, institutional surfaces, variable lighting, irregular geometries derived from converted single-family housing—has not been systematically characterized. This gap motivates the distribution-of-quality analysis in Section 3.

2.2 SLAM for Indoor Mapping

Simultaneous Localization and Mapping (SLAM) is the algorithmic foundation of mobile LiDAR capture on consumer devices [Cadena et al., 2016]. Consumer LiDAR applications, including Matterport, Polycam, and the native Apple RoomPlan API, employ variants of point cloud registration SLAM. SLAM-based capture introduces a distinctive error mode not present in station-based TLS: drift, the progressive accumulation of pose estimation error over a scanning trajectory. Drift magnitude is a function of trajectory length, surface texture quality, loop closure frequency, and scan speed [Zhou et al., 2021].

Loop closure detection—the identification of previously visited locations—is the primary mechanism for drift correction in SLAM systems [Xu et al., 2024]. In regular, rectilinear geometries, loop closures are frequent and drift remains bounded. In irregular geometries—split-level floor plans, non-orthogonal walls, large open areas—loop closures are sparse and drift can accumulate to levels that corrupt compliance-critical measurements [Zhou et al., 2021]. This is directly relevant to HCBS facilities, which are predom-

inantly converted residential properties with non-standard geometries.

2.3 Compliance Assessment in HCBS Settings

The HCBS Settings Rule, finalized in 2014 and subject to ongoing CMS guidance, establishes requirements in two categories: person-centered rights protections and physical environment standards [Centers for Medicare & Medicaid Services, 2014]. The physical environment provisions include requirements for accessible path of travel, privacy, private storage, and environmental safety that cannot be evaluated without reference to the facility’s spatial characteristics.

Current practice for HCBS compliance assessment relies primarily on surveyor observation supported by standardized instruments such as CMS Form 2567 and state-specific equivalents [Centers for Medicare & Medicaid Services, 2024a,b]. Surveyors are not required to use measurement instruments for the majority of compliance determinations, and inter-rater reliability of observation-based compliance assessment has been identified as a concern in adjacent regulatory contexts [Gajewski et al., 2006].

No published framework exists for spatially anchored, measurement-based HCBS compliance assessment at the time of this writing. This gap represents the primary motivation for the proposed system.

2.4 Point Cloud Quality Frameworks in Building Inspection

The closest existing work is that of Lo et al. [Lo et al., 2025], who propose a composite quality framework integrating accuracy, density, and completeness metrics for indoor LiDAR-based building inspection, with application to corridor and doorway geometries. While the quality dimensions overlap with several SQVL checks, the frameworks diverge architecturally and in purpose. Lo et al. treat quality evaluation as a precursor to point cloud enhancement—the output is an improved dataset. The SQVL treats quality evaluation as a compliance gate—the output is either a blocked pipeline or a zone-level uncertainty map that propagates into individual finding confidence. This distinction is not

incidental: in a regulatory compliance context, a finding derived from a degraded scan that passes without flagging creates false legal certainty, which is a qualitatively different failure mode from imprecise documentation. The SQVL’s architecture is a response to this requirement, not an extension of prior building inspection quality frameworks.

More broadly, point cloud quality assessment has an established literature in airborne and terrestrial LiDAR applications, addressing density, coverage, and completeness as individual metrics [Stanley and Laefer, 2021, Aryan et al., 2021]. The SQVL draws on these established metrics as component methods; the contribution rests on their integration as a mandatory compliance gate calibrated to the specific failure mode consequences of regulatory assessment—a combination that has no precedent in the building inspection quality literature.

2.5 Quality Assurance in Sensor-Based Measurement Systems

The problem of qualifying measurement data before downstream use is well-established in two adjacent technical domains: autonomous vehicle perception and medical imaging device qualification. These domains share with the compliance scanning context the fundamental challenge of deploying sensors in uncontrolled environments where measurement failure modes are non-obvious to the end user and where erroneous outputs carry consequential downstream risk.

In autonomous vehicle perception, the deployment of LiDAR, radar, and camera sensors in diverse environmental conditions has produced extensive literature on sensor performance characterization and failure mode classification [Heinzler et al., 2019]. The concept of operational design domain (ODD) specification—formally bounding the conditions under which a sensor or system is validated to perform [SAE International, 2021]—is directly analogous to the threshold-adjacency analysis proposed in Section 4.

In medical imaging, device qualification frameworks including NEMA standards, ACR phantom protocols, and FDA 510(k) substantial equivalence requirements establish that device accuracy must be demonstrated at the boundary of clinical actionability, not merely

under nominal conditions [Adjeiwaah et al., 2020]. This principle—that the hardest cases are at decision boundaries—directly motivates the SQVL’s threshold-adjacency analysis.

3 The Distribution-of-Quality Problem

3.1 Accuracy vs. Reliability in Compliance Contexts

Published accuracy benchmarks for consumer LiDAR characterize central tendency performance—mean absolute error, root mean square error—across a sample of measurements under controlled conditions. This characterization is appropriate for applications where errors are independent, random, and symmetric, and where no single measurement is dispositive.

Compliance assessment does not satisfy these conditions. In a compliance context, a single measurement—a corridor width, a door clear opening, a ramp slope—may determine whether a facility receives a deficiency citation. The distribution of possible errors for that single measurement, under the actual deployment conditions of that facility and that operator, is the relevant statistical object. The mean of a well-controlled laboratory study is not.

This paper proposes that the operative metric for consumer LiDAR in compliance applications is not mean accuracy but distributional reliability: the probability that a measurement falls within an acceptable error bound, evaluated across the full range of field conditions and operators. This reframing from accuracy to reliability is the organizing principle of the SQVL design.

3.2 Operator Variance

Consumer LiDAR for indoor capture is marketed as requiring minimal training. In practice, scan quality is sensitive to operator behavior in ways that are not visible to the operator in real time: scan speed, path selection, proximity to surfaces, handling of doorways, and behavior in low-texture zones all materially affect point density, surface coverage, and drift accumulation.

Empirical evaluation of operator variance in consumer depth sensor applications has demonstrated substantial performance differences across experience levels [Aryan et al., 2021]. For the purposes of characterizing the distribution-of-quality problem, this paper distinguishes three operator classes:

- **Expert operator:** Trained scanning professional with documented experience in SLAM-based capture and post-processing review. Achieves conditions near published accuracy benchmarks.
- **Trained operator:** Non-specialist staff who have completed a structured scanning protocol and conducted at least three supervised scans. Achieves intermediate performance, with degradation in irregular geometries.
- **Naive operator:** No prior scanning experience, scanning after brief instruction. Exhibits the full range of performance degradation described below.

Under expert operation, iPhone LiDAR achieves approximately ± 20 mm RMS error for linear dimensions in typical indoor environments [Teo and Yang, 2023]. Under naive operation in HCBS residential environments—characterized by irregular geometries, mixed lighting, and frequent furniture occlusion—the authors estimate RMS error expands to approximately ± 55 mm based on extrapolation from operator variance studies in comparable sensor deployments. This estimate is a design assumption motivating the SQVL framework and is subject to empirical validation in a forthcoming pilot study. The expansion of approximately $2.75\times$ renders numerous compliance-critical measurements unreliable without additional quality controls.

3.3 Environmental Variance in HCBS Facilities

HCBS group homes are predominantly converted single-family residential properties. This creates a built environment substantially different from the commercial office, retail, or new-construction educational facilities that represent the modal use case in published indoor scanning studies.

Key environmental factors affecting scan quality in HCBS facilities include:

Irregular geometry. Converted residential properties frequently feature non-orthogonal walls, alcoves, split-level floor plans, and interior modifications that create concave geometries and partial occlusions. These features reduce loop closure frequency in SLAM systems, increasing drift accumulation [Zhou et al., 2021].

Occupied spaces. HCBS facilities are occupied during the majority of feasible scanning windows. Residents, staff, and furnishings create dynamic occlusions that produce point cloud gaps in compliance-critical zones, including accessibility clearances and egress paths.

Mixed and adverse lighting. Consumer LiDAR systems supplement point cloud registration with visual-inertial odometry or photogrammetric methods. Low-light conditions, direct sunlight through windows, and reflective surfaces degrade visual tracking and introduce artifact voids in the point cloud.

Low-texture surfaces. White painted walls and ceilings—common in institutional residential environments—provide minimal texture for visual feature matching. In long, uniform corridors, this creates extended low-texture zones with elevated drift risk.

3.4 The Consequence: Undetected Measurement Failure

Without an explicit quality validation step, a compliance assessment system operating on a degraded scan produces three classes of error:

1. **False compliance:** A measurement that violates a regulatory threshold appears compliant due to scan error inflating the measured dimension. The facility receives a clean finding that would not survive physical verification.
2. **False deficiency:** A measurement that meets a regulatory threshold appears non-compliant due to scan error deflating the measured dimension. The facility receives a citation that is not warranted, generating remediation cost and legal exposure.
3. **Undetected threshold adjacency:** A measurement near a compliance threshold is reported with apparent precision that does not account for hardware uncertainty. The compliance determination is indeterminate but is presented as definitive.

All three error classes are consequential in the HCBS compliance context: false compliance

exposes residents to safety risk, false deficiency generates operational and legal burden for facility operators, and undetected threshold adjacency creates regulatory and legal liability for any party relying on the compliance report. The SQVL is designed to detect and block all three error classes before compliance evaluation proceeds.

4 The Scan Quality Validation Layer (SQVL)

4.1 Design Principles

The SQVL is designed as a first-class pipeline component that executes before any geometry extraction, BIM generation, or compliance rule evaluation. This placement is deliberate: downstream processing should never receive scan data whose quality is unknown or insufficient. The SQVL produces a per-scan quality verdict—PASS, BORDERLINE, or REJECT—and, for BORDERLINE and PASS scans, a zone-level quality map that propagates uncertainty into downstream compliance findings.

The seven SQVL checks are independent. Each check evaluates a distinct failure mode and can produce a BORDERLINE or REJECT verdict independent of the others. A REJECT verdict from any single check blocks all downstream processing. This independence is a design requirement: quality failures in real-world scans are not correlated, and a system that required multiple simultaneous failures to trigger REJECT would allow many consequential scan degradations to pass undetected.

Each SQVL check addresses a distinct failure mode identifiable in the prior literature. The contribution of the SQVL is not any individual check but the integration of these checks as an interdependent, mandatory gate calibrated to the specific error consequences of the compliance context. A drift-free scan with incomplete doorway coverage fails at Check 5. A dense scan with severe SLAM drift fails at Check 3. Neither failure mode is detectable by the other’s check. The combination is necessary because the failure modes are independent—a system designed to catch only the most common failure will pass scans that fail on a less common but equally consequential dimension. This composite

architecture has no precedent in building inspection quality literature, where quality assessment has historically been applied to individual metrics rather than integrated as a mandatory pipeline gate.

4.2 Check 1: Point Density Per Zone

Purpose: Verify that each spatial zone of the facility has been captured with sufficient point density to support geometry extraction at the required accuracy level.

Method: The scan volume is partitioned into cubic zones at a resolution appropriate to the smallest compliance-critical element in that zone. Point counts per zone are computed and compared against density thresholds.

Thresholds:

- PASS: $> 500 \text{ pts/m}^3$
- BORDERLINE: $200\text{--}500 \text{ pts/m}^3$
- FAIL (triggers REJECT): $< 200 \text{ pts/m}^3$

Rationale: Point density is a necessary but not sufficient condition for geometric accuracy. Zones below the PASS threshold are flagged for hybrid verification regardless of whether a compliance-critical measurement falls within them. This check is particularly sensitive to standoff distance errors (operator scanning from too far from a surface) and occlusion. The thresholds stated here are design estimates subject to empirical calibration.

4.3 Check 2: Surface Coverage Completeness

Purpose: Verify that all wall surfaces, floor, and ceiling in each zone have been captured with sufficient coverage to support extraction of compliance-critical dimensions.

Method: Surfaces are segmented from the point cloud using plane-fitting methods. Coverage is evaluated as the fraction of each surface area for which point cloud density exceeds a minimum threshold.

Thresholds:

- PASS: $\geq 85\%$ coverage per wall segment
- BORDERLINE: 70–85% coverage
- FAIL (triggers REJECT): $< 70\%$ coverage in any zone containing a compliance-critical element

Rationale: Partial wall capture is the most common source of width measurement error in SLAM-based indoor scanning. A wall segment with a gap produces a reconstructed surface displaced from the true surface position, systematically biasing any dimension measured to or from that surface. The 85% PASS threshold is a design estimate corresponding to an estimated maximum width error of ± 15 mm for standard corridor widths, subject to empirical validation.

4.4 Check 3: SLAM Drift Detection via Loop Closure Analysis

Purpose: Detect accumulation of SLAM pose estimation error (drift) that would corrupt geometric measurements across the facility.

Method: Loop closures are identified in the scan trajectory as points where the scanner returns to a previously scanned location. The residual misalignment at each loop closure is computed. The root mean square of loop closure residuals is the drift score.

Thresholds:

- PASS: ≤ 15 mm RMS loop closure residual
- BORDERLINE: 15–30 mm RMS
- REJECT: > 30 mm RMS

Rationale: Drift is the dominant accuracy determinant for whole-floor SLAM captures [Zhou et al., 2021, Xu et al., 2024]. A REJECT threshold of 30 mm RMS corresponds approximately to the hardware accuracy limit of the Leica DISTO D510 laser distance meter (the reference instrument used for hybrid verification), meaning that scan drift above this threshold degrades measurement quality to below the reference instrument's

accuracy—an incoherent state for a compliance system.

In scans with fewer than three loop closures, the SQVL automatically escalates to BORDERLINE regardless of residual magnitude, as the RMS estimate is insufficiently stable for a reliable verdict.

4.5 Check 4: Level Transition Integrity

Purpose: For multi-story or split-level facilities, verify that floor-to-floor scan registration is within bounds compatible with accurate egress and accessibility analysis.

Method: At each stairway, elevator, or level transition, the SQVL compares point cloud geometry of shared transition elements between floor-level scans. Misalignment is measured as the 3D offset between matched features.

Thresholds:

- PASS: ≤ 20 mm misalignment at all level transitions
- FAIL (triggers REJECT): > 20 mm at any level transition

Rationale: Level transition misalignment propagates into all accessibility and egress measurements that cross floor levels. The 20 mm threshold corresponds to the maximum allowable error for ADA ramp slope determination at the steepest permissible slope (1:12) across a 2-meter ramp span.

4.6 Check 5: Doorway Capture Quality

Purpose: Verify that doorway geometries—the highest-density compliance measurement points in HCBS facilities—have been captured with sufficient fidelity to support clear opening width determination.

Method: Door jamb positions are identified via vertical edge detection. Point counts at each jamb are evaluated. Both jambs must be present and above density threshold.

Thresholds:

- PASS: ≥ 50 points per jamb, both jambs captured
- FAIL (triggers REJECT for that doorway): < 50 points at either jamb, or either jamb not detected

Rationale: Door clear opening width is one of the most frequently cited deficiencies in HCBS and ADA surveys. The 32-inch (813 mm) minimum clear opening under ADA 2010 §404.2.3 and the 36-inch (914 mm) requirement for accessible route doors create decision thresholds where moderate measurement error (± 25 mm) produces non-trivial false classification rates. The 50-point threshold is a design estimate subject to empirical calibration.

4.7 Check 6: Lighting Artifact Detection

Purpose: Identify point cloud regions corrupted by lighting conditions that degrade LiDAR signal quality.

Method: Saturation voids are detected as contiguous spatial gaps in the point cloud that align with window positions and sun angle at scan time, with void boundaries exhibiting characteristic high-density rings at the illumination falloff boundary. Low-texture SLAM zones are detected by cross-referencing point cloud density with scan-concurrent RGB imagery; regions where RGB texture entropy is below threshold and point density is below PASS threshold are flagged.

Verdict: Saturation voids and low-texture SLAM zones in compliance-critical areas (doorways, corridors, accessibility clearance zones) trigger BORDERLINE for those zones; in non-critical areas, they are recorded in the quality map without affecting the global verdict.

Rationale: Lighting artifacts are invisible to the operator during scanning and are not correctable in post-processing. Identifying their extent and location before compliance evaluation prevents scan-derived measurements in artifact zones from being presented as reliable.

4.8 Check 7: Occlusion Gap Mapping

Purpose: Identify gaps in the point cloud attributable to occlusion (furniture, residents, equipment) in compliance-critical zones, and distinguish these from gaps attributable to incomplete scanning.

Method: Occluded gaps are characterized by abrupt point cloud boundaries with foreground object surfaces present. Scanning-incomplete gaps are characterized by open boundaries without foreground surfaces. For compliance-critical zones, gaps exceeding 0.5 m in maximum extent are flagged.

Verdict: Flagged gaps produce mandatory hybrid verification requirements for any compliance measurement whose spatial extent overlaps the gap region.

Rationale: Occupied HCBS facilities present persistent occlusion challenges that cannot be fully resolved through scanning protocol optimization. The 0.5 m threshold is set below the minimum clearance requirement in ADA 2010 §403.5.1 (60-inch passing clearance, approximately 1.5 m), ensuring that gaps sufficient to conceal a compliance violation are always flagged.

4.9 Threshold-Adjacency Analysis

Following completion of the seven SQL checks, all measurements extracted from PASS or BORDERLINE scans undergo threshold-adjacency analysis. This analysis applies the principle of conformance testing under measurement uncertainty—equivalent to guard band methodology in ISO 14253-1—to the spatial compliance context. For each compliance rule with a numerical threshold, the extracted measurement is compared against that threshold. If the absolute difference between the measured value and the regulatory threshold falls within the hardware accuracy margin of the scanning instrument, the finding is classified as threshold-adjacent.

Formally: if $|m - t| \leq \varepsilon$, where m is the measured value, t is the regulatory threshold, and ε is the hardware accuracy margin ($\varepsilon \approx 25$ mm for iPhone LiDAR in typical field conditions, a design estimate subject to empirical refinement), then the finding is flagged

as THRESHOLD-ADJACENT and mandatory hybrid verification is triggered.

This analysis is applied after SQLV verdicts, not during, because threshold-adjacency is a property of the measurement-to-rule relationship, not of scan quality per se. A high-quality scan can still produce threshold-adjacent measurements, and those measurements require the same escalation as scan-quality-driven uncertainties.

5 Accuracy Band Classification

5.1 Overview

Not all compliance rules impose equivalent measurement requirements. A rule that requires a fire extinguisher to be present is topologically verifiable—the relevant question is binary. A rule that requires a toilet to be positioned within 18 inches of a side wall requires sub-inch accuracy at a specific location. Treating all rules as imposing equivalent scanning requirements either over-specifies simple rules or under-specifies precision rules.

The accuracy band classification system stratifies compliance rules into five categories based on their spatial measurement requirements, mapping each category to the minimum hardware and workflow necessary for defensible assessment. The band proportion estimates below are preliminary figures based on the authors' analysis of HCBS and ADA regulatory text; systematic validation through a full regulatory coding process is the subject of a companion paper in preparation.

5.2 Band A — Topological Rules

Definition: Rules whose compliance status depends on the presence or absence of a spatial feature, not its dimensions.

Examples: Emergency lighting present in egress path; exit signage at egress doors; handrails present on both sides of stairways.

Estimated proportion: 35–40% of HCBS rules.

Scanning requirement: Consumer LiDAR is sufficient. Feature detection at the point cloud or RGB image level is adequate.

5.3 Band B — Area and Volume Rules

Definition: Rules whose compliance status depends on the area or volume of a space.

Examples: Bedroom minimum net floor area; required storage area per resident; minimum common area.

Estimated proportion: 15–20%.

Scanning requirement: Consumer LiDAR is sufficient unless measured area is within 5% of the regulatory threshold. iPhone LiDAR area measurement error of 1–4% [Teo and Yang, 2023] creates a zone of indeterminacy at threshold proximity that requires hybrid verification.

5.4 Band C — Linear Dimension Rules

Definition: Rules whose compliance status depends on a single linear measurement: width, height, depth, or clearance.

Examples: Minimum door clear opening (32 inches); minimum corridor clear width (36 inches or 60 inches for accessible routes); minimum turning radius clearance (60-inch diameter).

Estimated proportion: 30–35%.

Scanning requirement: Consumer LiDAR is sufficient for clear compliance (measured dimension exceeds threshold by more than the hardware accuracy margin). Consumer LiDAR is insufficient for threshold-adjacent measurements. For threshold-adjacent Band C measurements, the Leica DISTO D510 is the reference instrument.

5.5 Band D — Precision Dimension Rules

Definition: Rules requiring measurement at sub-inch accuracy at specific geometric relationships: offsets, slopes, heights above finished floor at a specified location.

Examples: Grab bar height (33–36 inches AFF); grab bar centerline distance from toilet centerline (42 inches minimum); accessible ramp slope ($\leq 1:12$); handrail height (34–38 inches AFF).

Estimated proportion: 10–15%.

Scanning requirement: Consumer LiDAR is never sufficient for Band D rules. The iPhone LiDAR accuracy margin (± 20 – 55 mm depending on conditions) is larger than the permitted tolerance range for many Band D rules. The Leica DISTO D510, with stated accuracy of ± 1.0 mm, is mandatory for all Band D measurements. Band D rules always trigger hybrid verification.

5.6 Band E — Documentary Rules

Definition: Rules whose compliance status depends on documentary evidence rather than spatial measurement.

Examples: HVAC maintenance records; fire suppression system inspection logs; emergency evacuation drill documentation.

Estimated proportion: 10–15%.

Scanning requirement: None. Band E rules are not spatial and are outside the scope of the SQVL.

6 Hybrid Verification Workflow

6.1 Mandatory Trigger Conditions

Hybrid verification—on-site physical measurement using a calibrated reference instrument—is mandatory under four conditions:

1. **All Band D measurements.** Consumer LiDAR accuracy is insufficient for precision dimension rules without exception.
2. **All threshold-adjacent Band C measurements.** Any measurement falling within the hardware accuracy margin of a regulatory threshold requires physical confirmation.
3. **All measurements in SQVL low-coverage zones.** Elements in zones that received BORDERLINE verdicts on Check 2 (surface coverage), Check 6 (lighting artifacts), or Check 7 (occlusion gap) require physical verification.
4. **All HIGH severity findings derived solely from scan data.** Findings at the highest severity level supported only by scan-derived measurements require physical confirmation before being reportable.

6.2 Reference Instrument

The reference instrument for hybrid verification is the Leica DISTO D510, a class II laser distance meter with stated accuracy of ± 1.0 mm and range of 0.05–200 m. The DISTO D510 is used for Band D measurement collection, Band C threshold-adjacency resolution, and level-surface slope measurement (integrated inclinometer, accuracy $\pm 0.2^\circ$).

The DISTO D510 is specified rather than a tape measure because its accuracy is documented, its calibration is traceable, and its measurement records are digitally exportable, enabling integration into the compliance evidence chain.

6.3 Compliance Score Blocking

The compliance score for a facility assessment is blocked until all mandatory hybrid verification items are resolved. This is a hard constraint, not a soft warning. An unresolved mandatory hybrid item represents an unknown that cannot be bounded within acceptable error limits; producing a compliance score in this state would present false confidence to the facility operator and any relying authority.

While the compliance score is blocked, a partial report is available containing: the spatial

risk map (zones with compliance concerns); the egress model (corridor connectivity, exit path analysis, occupant load calculation); and the ADA clearance visualization (plan-view representation of measured clearances). These outputs do not require resolution of hybrid items because they do not make binary compliance determinations.

6.4 Evidence Chain Requirements

For a compliance finding to be reportable—positive or negative—the following evidence chain must be complete:

- **Scan data:** The raw point cloud, with timestamp, device identifier, and SQLV verdict.
- **Extraction record:** The measurement extracted from the point cloud, with zone identifier, surface fitting parameters, and threshold-adjacency flag.
- **Hybrid measurement record (if applicable):** The physical measurement, with instrument identifier, operator identifier, measurement timestamp, and measurement method.
- **Rule evaluation record:** The specific regulatory provision evaluated, the version of the encoded rule, the measurement(s) applied, and the compliance determination.

This evidence chain is designed to support both internal quality assurance and regulatory or legal review. The requirement for an evidence chain at the finding level, rather than at the report level, is a deliberate design choice intended to enable challenge of individual findings without invalidating the full report.

7 Comparison to QA Frameworks in Adjacent Domains

7.1 Design Parallels

The SQLV embodies several design principles that are well-established in quality assurance frameworks for sensor-based decision systems in adjacent domains. This section

situates the SQVL within that literature to support its design rationale and identify opportunities for methodological borrowing.

7.2 Autonomous Vehicle Perception

The deployment of LiDAR, radar, and camera sensors for autonomous vehicle (AV) perception presents a structural analogue to consumer LiDAR for compliance assessment: sensors with characterized mean performance are deployed in uncontrolled conditions where specific failure modes are consequential and not visible to the end user.

Operational design domain (ODD) specification formally bounds the conditions under which a perception system is validated to perform [SAE International, 2021]. The SQVL’s REJECT threshold system is functionally equivalent to ODD boundary detection: a scan that fails one or more checks has exited the validated operating domain and its outputs cannot be trusted.

Failure mode classification in AV perception distinguishes false positives, false negatives, and missed detections as distinct failure types with distinct risk profiles [Heinzler et al., 2019]. The three compliance error classes defined in Section 3—false compliance, false deficiency, and undetected threshold adjacency—are the analogous taxonomy for compliance sensing.

Sensor qualification under degraded conditions—particularly for LiDAR performance in adverse weather [Heinzler et al., 2019]—has produced characterization methods for the systematic expansion of error distributions under adverse conditions that directly inform the operator variance analysis in Section 3.

7.3 Medical Imaging Device Qualification

Medical imaging QA frameworks, including ACR phantom protocols and FDA performance testing guidance, are built on a principle directly applicable to the SQVL: accuracy must be demonstrated at the boundary of clinical actionability, not merely under nominal conditions [Adjeiwaah et al., 2020].

In mammographic screening, image quality validation is performed specifically at lesion sizes near the detection threshold, because that is where classification errors occur [Huda et al., 2002]. The SQVL’s threshold-adjacency analysis embodies the same logic: measurement quality is most consequential at the boundary between compliance and deficiency, and that boundary is precisely where the SQVL requires the highest-quality evidence.

Additionally, medical imaging QA frameworks distinguish device performance from system performance: a correctly calibrated device operated incorrectly produces incorrect results. The SQVL’s explicit treatment of operator variance and the hybrid verification workflow’s requirement for documented operator identity reflect the same system-level perspective.

7.4 Structural Health Monitoring

Structural health monitoring (SHM) systems that deploy sensor networks in civil infrastructure share the SQVL’s challenge of producing measurements for consequential decisions in environments where sensor failure modes are not directly observable [Farrar and Worden, 2007]. SHM practice has developed statistical control chart methods for detecting sensor degradation and data quality anomalies, including the use of multivariate control limits on sensor outputs as a proxy for measurement reliability. The SQVL’s seven-check architecture can be understood as a domain-specific instance of multivariate quality control, with checks calibrated to the failure modes specific to SLAM-based indoor LiDAR.

8 Discussion

8.1 Limitations

The SQVL architecture presented in this paper is a designed framework rather than an empirically validated system. The thresholds proposed—point density (500 pts/m³), surface coverage (85%), and drift REJECT threshold (30 mm RMS)—are grounded in

first-principles reasoning from published instrument specifications and SLAM literature, but have not been validated against ground-truth measurements in HCBS facility environments. These values are design estimates and should be understood as such. Empirical validation against physical measurement ground truth across a sample of HCBS facilities, with representative operator variance, is the subject of planned future work. The contribution of this paper is the framework architecture and the compliance-driven design rationale, not the specific threshold values.

The operator variance estimates (± 20 mm expert, ± 55 mm naive) are extrapolated from published studies in related scanning contexts rather than from direct measurement in HCBS environments. These estimates should be treated as motivating the need for the SQVL rather than as precise empirical claims.

The accuracy band proportion estimates (Band A = 35–40%, etc.) are based on the authors’ analysis of HCBS and ADA regulatory text and have not been independently validated through a systematic regulatory coding process. A companion paper in preparation conducts that systematic analysis.

8.2 Regulatory and Legal Implications

The SQVL is designed with the requirements of a defensible compliance determination in mind. “Defensible” in this context means: producible as evidence in an administrative proceeding (state agency survey, CMS audit) or legal challenge without the producing party needing to make claims about measurement accuracy that go beyond what the evidence chain supports.

The field has no established standard for what constitutes a defensible spatial compliance determination. The SQVL framework represents a proposed answer to that question—one that is conservative (requiring hybrid verification wherever uncertainty is material) and transparent (producing a documented quality verdict for every scan). The authors anticipate that the specific thresholds will require refinement through regulatory dialogue and empirical validation.

8.3 Future Work

A planned validation study will conduct a structured comparison of scan-derived measurements against Leica DISTO ground truth across a sample of HCBS facilities, with scans conducted by operators at each experience level (expert, trained, naive). That study will:

- Empirically characterize the operator variance distribution across rule categories.
- Test the SQVL’s REJECT and BORDERLINE verdicts against ground-truth measurement error—specifically, whether a REJECT-triggering scan produces measurements outside acceptable error bounds.
- Characterize the threshold-adjacency rate in real HCBS buildings (what fraction of compliance determinations fall in the indeterminate zone?).
- Calibrate SQVL thresholds based on empirical findings.

9 Conclusion

Consumer LiDAR presents a genuine opportunity to make spatially anchored compliance assessment accessible in HCBS residential settings. The opportunity is real but conditioned: it is conditioned on the existence of a quality validation framework that makes the difference between defensible and indefensible scan-derived measurements explicit, and that routes measurements to appropriate verification workflows based on their classification.

This paper has characterized the distribution-of-quality problem—the gap between controlled-condition accuracy benchmarks and field performance across real operators and real HCBS environments—and has presented the Scan Quality Validation Layer as a design response. The SQVL’s seven-check architecture, accuracy band classification, hybrid verification workflow, and evidence chain requirements are proposed as the minimum infrastructure for scan-derived compliance findings to be treated as reliable by regulatory and legal bodies.

The deeper contribution is methodological: the paper proposes that reliability, not accu-

racy, is the operative metric for compliance sensing, and that the relevant statistical object is the distribution of measurement quality across operators and conditions rather than its central tendency. This reframing—from accuracy to reliability, from instrument specification to system performance—is the basis on which any spatially grounded compliance system must be designed.

References

- Byron J. Gajewski, Sarah Thompson, Nancy Dunton, Annette Becker, and Marcia Wrona. Inter-rater reliability of nursing home surveys: a Bayesian latent class approach. *Statistics in Medicine*, 25(2):325–344, 2006. doi: 10.1002/sim.2224.
- Evgeniya Antonova and David Zimmerman. Scope and severity index: a metric for quantifying nursing home survey deficiency number, scope, and severity adjusted for the state-related measurement bias. *Journal of the American Medical Directors Association*, 13(2):188.e7–188.e12, 2012. doi: 10.1016/j.jamda.2011.04.021. SSI weights confirmed: A=5, B=10, C=15, D=20, E=30, F=40, G=35, H=50, I=65, J=55, K=75, L=100.
- T.-A. Teo and C.-C. Yang. Evaluating the accuracy and quality of an iPad Pro’s built-in LiDAR for 3D indoor mapping. *Developments in the Built Environment*, 14:100169, 2023. doi: 10.1016/j.dibe.2023.100169.
- O. Catharia, F. Richard, H. Vignoles, P. Véron, A. Aoussat, and F. Segonds. Smartphone LiDAR data: a case study for numerisation of indoor buildings in railway stations. *Sensors*, 23(4):1967, 2023. doi: 10.3390/s23041967.
- Jiwei Hou, Patrick Hübner, Jakob Schmidt, and Dorota Iwaszczuk. Indoor mapping with entertainment devices: evaluating the impact of different mapping strategies for Microsoft HoloLens 2 and Apple iPhone 14 Pro. *Sensors*, 24(4):1062, 2024. doi: 10.3390/s24041062. iPhone 14 Pro Gaussian mean distance 2.7–5.2 cm against TLS ground truth in indoor corridor.

- Federal Highway Administration. Leveraging pocket LiDAR for construction inspection and digital as-builts — phase 1 [TechBrief]. Technical Report FHWA-HRT-25-017, U.S. Department of Transportation, Federal Highway Administration, 2024. URL <https://rosap.nhtl.bts.gov/view/dot/79631>.
- Handan Aş Çemrek, Seda Sönmez, and Ümit Işıkdağ. Mobile 3D LiDAR-based Scan2BIM process: a comparison of point cloud processing tools. In *Proceedings of eCAADe 2025 — 43rd Conference on Education and Research in Computer Aided Architectural Design in Europe*, Ankara, Turkey, 2025. URL https://papers.cumincad.org/data/works/att/ecaade2025_589.pdf.
- W. Boehler, M. B. Vicent, and A. Marbs. Investigating laser scanner accuracy. In *International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, volume 34, pages 696–701, 2003. XIXth CIPA Symposium, Antalya, Turkey.
- N. A. S. Russhakim, M. F. M. Ariff, Z. Majid, K. M. Idris, and N. Darwin. The suitability of terrestrial laser scanning for building survey and mapping applications. In *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, volume XLII-2/W9, pages 663–668, 2019. doi: 10.5194/isprs-archives-XLII-2-W9-663-2019.
- P. Henry, M. Krainin, E. Herbst, X. Ren, and D. Fox. RGB-D mapping: using Kinect-style depth cameras for dense 3D modeling of indoor environments. *The International Journal of Robotics Research*, 31(5):647–663, 2012. doi: 10.1177/0278364911434148.
- C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J. J. Leonard. Past, present, and future of simultaneous localization and mapping: toward the robust-perception age. *IEEE Transactions on Robotics*, 32(6):1309–1332, 2016. doi: 10.1109/TRO.2016.2624754.
- Lipu Zhou, Daniel Koppel, and Michael Kaess. LiDAR SLAM with plane adjustment for indoor environment. *IEEE Robotics and Automation Letters*, 6(4):7073–7080, 2021. doi: 10.1109/LRA.2021.3073076.

Xu Xu, Lianwu Guan, Junjie Zeng, Yibo Sun, Yanbin Gao, and Qiang Li. When-to-Loop: enhanced loop closure for LiDAR SLAM in urban environments based on SCAN CONTEXT. *Micromachines*, 15(10):1212, 2024. doi: 10.3390/mi15101212.

Centers for Medicare & Medicaid Services. Medicaid program; state plan home and community-based services, 5-year period for waivers, provider payment reassignment, and home and community-based setting requirements for community first choice and home and community-based services (HCBS) waivers. Final Rule. *Federal Register*, 79(11):2948–3039, 2014. URL <https://www.govinfo.gov/content/pkg/FR-2014-01-16/pdf/2014-00487.pdf>. January 16, 2014. To be codified at 42 C.F.R. § 441.301(c)(4).

Centers for Medicare & Medicaid Services. Statement of deficiencies and plan of correction (Form CMS-2567). Centers for Medicare & Medicaid Services, Baltimore, MD, 2024a. URL <https://www.cms.gov/medicare/cms-forms/cms-forms/downloads/cms2567.pdf>.

Centers for Medicare & Medicaid Services. State operations manual: Appendix PP — guidance to surveyors for long term care facilities. Rev. 225. Centers for Medicare & Medicaid Services, Baltimore, MD, 2024b. URL <https://www.cms.gov/medicare/provider-enrollment-and-certification/guidanceforlawsandregulations/downloads/appendix-pp-state-operations-manual.pdf>. Issued August 8, 2024.

T. T. Lo, Y. Zhang, J. Cao, L. Di Sarno, and J. Chen. Towards a LiDAR-based point cloud evaluation and enhancement in building inspection, 2025. SSRN preprint. Full citation to be updated on journal publication.

Matthew H. Stanley and Debra F. Laefer. Metrics for aerial, urban LiDAR point clouds. *ISPRS Journal of Photogrammetry and Remote Sensing*, 175:268–281, 2021. doi: 10.1016/j.isprsjprs.2021.01.010.

Afroz Aryan, Frédéric Bosché, and Pingbo Tang. Planning for terrestrial laser scanning

in construction: a review. *Automation in Construction*, 125:103551, 2021. doi: 10.1016/j.autcon.2021.103551.

Robin Heinzler, Philipp Schindler, Jürgen Seekircher, Werner Ritter, and Wilhelm Stork. Weather influence and classification with automotive LiDAR sensors. In *2019 IEEE Intelligent Vehicles Symposium (IV)*, 2019. doi: 10.1109/IVS.2019.8814205.

SAE International. Taxonomy and definitions for terms related to driving automation systems for on-road motor vehicles. Technical Report J3016_202104, SAE International, Warrendale, PA, 2021. URL https://www.sae.org/standards/content/j3016_202104/.

Mary Adjeiwaah, Anders Garpebring, and Tufve Nyholm. Sensitivity analysis of different quality assurance methods for magnetic resonance imaging in radiotherapy. *Physics and Imaging in Radiation Oncology*, 13:21–27, 2020. doi: 10.1016/j.phro.2020.03.001.

Walter Huda, Anthony M. Sajewicz, Kent M. Ogden, Ernest M. Scalzetti, and David R. Dance. How good is the ACR accreditation phantom for assessing image quality in digital mammography? *Academic Radiology*, 9(7):764–772, 2002. doi: 10.1016/S1076-6332(03)80345-8.

Charles R. Farrar and Keith Worden. An introduction to structural health monitoring. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 365(1851):303–315, 2007. doi: 10.1098/rsta.2006.1928.