

Stereo Reconstruction and Classification of Ionizing Particle Tracks in a Diffusion Cloud Chamber

Andrew Simonson

Golisano College of Computing and Information Sciences

Rochester Institute of Technology

Rochester, NY, USA

abs1907@rit.edu

Abstract—Diffusion cloud chambers make ionizing radiation directly visible as condensation trails. Classifying the resulting tracks by particle type has historically relied on a human observer eyeballing track length and width, a practice that dates to the earliest bubble- and cloud-chamber photography. A single camera makes this worse, since a track angled out of the imaging plane is foreshortened in its two-dimensional projection, biasing any length-based classification with only estimations available for depth. This project builds a diffusion cloud chamber imaged by three synchronized, calibrated cameras, reconstructs particle tracks in three dimensions via triangulation, and classifies particle type (alpha, beta, cosmic-ray muon) using three approaches on the same features: classical length/width thresholds, a supervised machine-learning baseline, and a Bayesian classifier that propagates measurement uncertainty into a calibrated posterior probability. The same pipeline applies statistical process control to session-level data to flag anomalous experimental conditions; the spike-in protocol’s deliberate shifts also serve as a built-in validation of that monitoring layer. Data collection follows a controlled spike-in protocol: a background-only baseline session, followed by sessions with known and easily obtainable alpha and beta sources placed at fixed positions. This gives weakly labeled ground truth without requiring any licensed material. The central empirical question has two parts: whether the constructed apparatus achieves geometrically valid 3D reconstruction, confirmed directly against an object of known length instead of assumed from the general historical track record of stereo photography, and, given that, whether three-dimensional reconstruction measurably corrects the foreshortening bias inherent to single-camera classification, and by how much.

Index Terms—cloud chamber, ionizing radiation, particle track classification, multi-view stereo, Bayesian inference, uncertainty quantification

I. INTRODUCTION

The diffusion cloud chamber, first developed by C.T.R. Wilson in the early twentieth century, was the instrument through which several fundamental particles were first observed. Unlike most detectors, which infer particle interactions from an indirect electronic signal, it remains one of the few devices that makes those interactions directly visible to the human eye [1]. For decades, particle identification from chamber photographs was done by a trained observer classifying tracks by their visual signature where alpha particles leave short, thick, straight trails compared to beta particles’ (energetic electrons) long, thin, wandering ones and cosmic-ray muons’ long, straight, minimally scattered tracks that pass through with little deflection. Multi-camera stereo photography

of bubble and cloud chambers was standard practice in mid-twentieth-century particle physics, so a track’s true three-dimensional trajectory could be reconstructed instead of read off a single flat photograph [2].

Low-cost, open-hardware particle detectors, most notably the CosmicWatch desktop muon detector developed for physics outreach, have made this kind of instrumentation newly accessible outside a laboratory setting [3]. This project sits at the intersection of that trend and modern computer vision and Bayesian statistics. It modernizes the historical stereo-photography approach to cloud chamber analysis: consumer cameras and automated image processing take the place of a trained physicist manually plotting stereo pairs by hand.

The remainder of this proposal is organized as follows. Section II defines the problem and its significance. Section III surveys related work in cloud chamber instrumentation, particle-based imaging, and the computer-vision and Bayesian methods this project draws on. Section IV describes the proposed apparatus, data collection protocol, processing pipeline, and evaluation plan, along with anticipated challenges and concrete deliverables.

II. PROBLEM

A. Problem Definition

A cloud chamber viewed by a single camera records only a two-dimensional projection of an inherently three-dimensional particle track; a track angled away from that camera’s imaging plane appears shorter than it actually is, which systematically biases any classification method that depends on track length. Existing classification of chamber tracks, whether performed by a human observer or a simple image-processing heuristic, produces a hard particle-type label with no quantification of confidence, which is precisely where it is least reliable: on foreshortened or otherwise ambiguous tracks. This project addresses both limitations by reconstructing tracks in three dimensions from multiple synchronized, calibrated cameras, and by replacing hard-label classification with a Bayesian model that outputs a calibrated posterior probability over particle type.

B. Significance of Problem

The immediate beneficiaries are physics education and outreach: an automated, uncertainty-aware classification pipeline

turns a classic hands-on demonstration into a quantitative instrument, without a licensed radioactive source or laboratory-grade equipment. More broadly, pairing an inexpensive, self-built sensing instrument with rigorous statistical treatment of its measurement uncertainty is directly transferable to other low-cost sensing problems where the instrument, not the modeling, is the limiting factor; with better modeling, a less precise instrument can still become a viable measurement tool. Such an expansion of viability reduces costs, improving quantity and accessibility of instruments. Section IV’s process-monitoring layer extends that pattern by turning a physical risk this proposal already anticipates, an unstable low-cost instrument, into a quantitatively monitored condition. This project’s claims are modest, falsifiable, and build on each other. First, that the constructed three-camera rig achieves 3D reconstruction accurate enough to recover the true geometry of a known reference object: stereo triangulation itself is not new, but this specific build’s accuracy is. Second, once that holds, that three-dimensional reconstruction measurably reduces the length-foreshortening bias a single 2D camera view introduces into particle-type classification. This proposal reports both: geometric validation against known ground truth, then the 2D-versus-3D classification comparison it makes possible.

III. PRIOR WORK

Wilson’s original cloud chamber work established the instrument and first demonstrated that ionizing particle tracks could be made visible via vapor condensation [1]. Stereo photography of bubble and cloud chambers, using two or three cameras at known, calibrated positions, was the standard approach to recovering a track’s true three-dimensional geometry throughout the mid-twentieth century, and remains the conceptual basis for the multi-camera approach in this project [2]. More recently, the CosmicWatch project demonstrated that a rigorous, quantitatively useful particle detector can be built from off-the-shelf components at hobbyist cost, motivating the broader premise that meaningful particle physics measurements do not require institutional laboratory equipment [3]. Cosmic-ray muon radiography, which uses ambient muon scattering and attenuation to image the interior of large, dense objects without an artificial radiation source, is a related example of particle-based imaging that has been applied to cargo scanning, volcanology, and structural inspection [4]. It is methodologically adjacent to this project’s use of ambient and spiked particle flux as a measurement signal.

On the technical side, this project’s multi-camera triangulation depends on standard camera calibration methodology. Zhang’s flexible calibration technique is the widely adopted approach [5], implemented here using the OpenCV computer vision library [6]. The Bayesian classification approach follows standard Bayesian pattern-recognition methodology, treating particle-type identification as inference over a probabilistic model instead of a deterministic decision rule [7]. This project also monitors the data-collection process itself for anomalous conditions, using standard statistical process control methodology: Shewhart and cumulative-sum control charts

computed on session-level aggregate statistics, the standard toolkit for distinguishing an unexpected process shift from ordinary variation in an ongoing measurement process [8].

The multi-camera “freeze a moment, view it from anywhere” concept in this proposal takes direct inspiration from World Labs’ Atlas, a generative multi-view world model that takes a small number of camera views with known poses and uses a learned diffusion-based model to synthesize novel viewpoints and freeze-frame reframings of a captured scene [9]. Atlas itself is not used in this project because it utilizes a generative approach trained for photorealistic video which can create plausible but misleading data, especially against thin, faint vapor trails in a semi-opaque mist. That distinction matters when the goal is to quantify measurement uncertainty, not to produce a visually convincing reconstruction. Classical calibrated triangulation is used here because its error sources (calibration residual, correspondence error) are known and can be propagated into the Bayesian classifier in a principled way; a black-box generative prior offers no equivalent.

IV. PROPOSED METHODOLOGY

A. Plan

The project proceeds in four stages: apparatus construction, data collection, processing, and modeling/evaluation.

Apparatus construction. A diffusion cloud chamber is built from a cooled black-felt floor, likely either dry ice or a Peltier stack, an isopropanol-soaked felt strip as the vapor source, and a grazing-angle LED light sheet for track visibility. The choice between the two cooling methods is made during early construction based on achieved temperature and run duration, informed by standard cloud-chamber operating practice [12] and, if Peltier cooling is used, published module-selection and thermal-performance testing for that approach [11]. Three cameras (Raspberry Pi Camera Modules, hardware-triggered over GPIO for frame-level synchronization) are mounted at fixed, non-collinear positions around the chamber so that triangulation is well conditioned, as shown in Fig. 1. Recent dynamic multi-view work is suggestive, not a direct validation: 4C4D reconstructs full dynamic 3D scenes from as few as four portable cameras [10], though it targets opaque, texture-rich subjects, while a chamber track is thin and low-contrast (a harder correspondence problem, discussed in Section IV-B). Three cameras is a geometrically well-conditioned choice for triangulating a single curve, not a claim that this problem is as easy as 4C4D’s; reconstructing a diffuse target like the surrounding vapor, not attempted here, is a further step removed still. Each camera is calibrated individually and jointly using a checkerboard target photographed inside the chamber’s viewing volume, following Zhang’s method [5] as implemented in OpenCV [6]. Calibration alone does not confirm the pipeline reconstructs real geometry correctly, so a rigid object of known length is triangulated at several positions and orientations and its reconstructed length compared against the true value. This is the apparatus-viability check from Section II-B, run before any particle-track data is treated as usable.

Data collection. Data collection follows a controlled spike-in experiment, not passive monitoring: a background-only baseline session (capturing ambient cosmic-ray muons and residual radon-daughter alphas) is recorded first, followed by sessions with a known source placed at a fixed position and distance from the chamber. Alpha sources must be surface-exposed to escape self-absorption; an open lantern mantle or the bare foil from a smoke detector are suitable, unlicensed choices. Beta sources (a potassium chloride salt substitute, or uranium glass) tolerate more bulk material, since beta particles penetrate farther. Every source used is legal to possess and requires no license, satisfying the constraint that this project not depend on regulated radioactive material. All data is self-generated by the constructed apparatus; no external dataset is used. The spike-in protocol supplies approximate ground-truth labels, the dominant particle type expected in a given session, against which the classifier’s output is compared.

Processing. Recorded video is processed as shown in Fig. 2: per-camera background subtraction isolates candidate tracks, epipolar-constrained matching identifies the same track across camera views, and triangulation using the calibrated camera geometry reconstructs each track in three dimensions. Track length, width (a proxy for ionization density), curvature, and orientation are extracted from the reconstructed 3D curve. The same feature extraction is also performed on each individual camera’s 2D projection alone, which is required for the comparison described below.

Modeling and evaluation. Three tiers of trail-identification approach are compared on the same extracted feature set, so the added statistical machinery at each tier can be judged against what it actually buys over the simpler alternative: classical rule-based thresholds (replicating the historical manual approach), a standard supervised classifier (logistic regression or a random forest, the off-the-shelf baseline a practitioner would reach for first), and a Bayesian classifier that propagates measurement uncertainty from calibration and triangulation error into the features and produces a posterior probability over particle type instead of a hard label. An optional unsupervised clustering pass (e.g. a Gaussian mixture model) is also run as an exploratory check on whether particle types separate without labels. Classification directly from raw track imagery (e.g. a convolutional network) is deliberately out of scope: it would trade away the interpretable features the uncertainty propagation depends on, and is noted as future work. Once the apparatus-viability check confirms the pipeline recovers real geometry, the central evaluation compares tier performance on single-camera 2D features versus triangulated 3D features from the same events, testing whether 3D reconstruction corrects the foreshortening bias from Section II-A.

Post-detection process monitoring. Beyond classifying individual tracks, session-level aggregate statistics (event rate, per-track feature distributions, and calibration reprojection error) are monitored across the full data-collection period using statistical process control [8]. Control limits are established from the background-only baseline sessions, which define the chamber’s normal operating variation, and every subsequent

session is checked against those limits. This is intended as an anomaly-detection layer and to set a baseline for the spike-in protocol, which introduces a deliberate, known shift in event rate and feature distribution. It also functions as a built-in validation of the monitoring system itself: a correctly functioning control-chart pipeline should flag spike-in sessions as out-of-control for the expected reason, cleanly distinguishing that case from an unexplained anomaly of actual concern.

Section IV-C lists which deliverables are expected at each project checkpoint.

B. Challenges or Barriers

The largest technical risk is cross-camera correspondence: cloud chamber tracks are thin, faint, and transient, a harder matching problem than the rigid, textured objects most stereo-vision methods assume. This is mitigated by validating correspondence on a small, manually verified set of events before relying on it at scale, with a semi-manual fallback for a smaller high-quality dataset if automated matching underperforms. Chamber stability is a second risk: diffusion cloud chambers are sensitive to temperature gradients, vibration, and drafts, so consistent track visibility requires real commissioning time, budgeted before any session is treated as final. Event rate is a third consideration, since background cosmic-ray and ambient alpha events are infrequent; baseline sessions are budgeted in hours, not minutes, and the spike-in protocol is meant to raise the event rate for the particle type of interest. Camera synchronization is verified explicitly, a shared visible trigger confirmed to land in the same frame index across all three streams, instead of assumed from the hardware specification alone. The broadened analysis scope (the ML baseline and process-monitoring layer) is low-risk to add, since both run on the same extracted features and session-level data the core pipeline already produces: they add analysis code, not new instrumentation risk. Finally, the apparatus itself may take longer to commission than planned; if the chamber is not producing stable, visible tracks by the first checkpoint, the fallback is to shift toward the classical-baseline and correspondence-validation work using whatever background-only data is available, preserving a complete but smaller-scope comparison instead of delaying the project on hardware readiness.

C. Project Deliverables

Deliverables are tied to the checkpoint schedule, so progress is demonstrable at each milestone rather than only at the end.

By the first project checkpoint: (1) a working, calibrated three-camera diffusion cloud chamber apparatus; (2) a geometric validation result confirming the reconstruction pipeline’s accuracy against a known-length reference object; (3) an initial background-only dataset, used to establish the process-control baseline described in Section IV-A.

By the second project checkpoint: (4) a complete labeled dataset of background and spike-in recording sessions, with source, position, and duration metadata; (5) an open cross-camera correspondence and triangulation pipeline reconstruct-

ing particle tracks in three dimensions; (6) a classical rule-based baseline classifier; (7) a supervised machine-learning baseline classifier over the same extracted features.

By the final report and presentation: (8) a Bayesian classifier producing calibrated posterior probabilities over particle type, with measurement uncertainty propagated from the reconstruction pipeline; (9) a statistical process-control monitoring layer flagging anomalous data-collection sessions, validated against the known spike-in sessions; (10) an empirical comparison quantifying the effect of 3D reconstruction on classification accuracy and calibration relative to single-camera 2D features; (11) the final project report and presentation summarizing methodology, results, and limitations.

As an optional stretch, attempted only if the above is complete with time remaining, the reconstructed track may be composited into a phone-LiDAR-scanned 3D model of the static apparatus for a freeze-frame presentation visualization; this is presentation polish, not a project dependency.




Fig. 1. Top-down schematic of the stereo cloud chamber apparatus: three calibrated cameras view a cooled (dry ice or Peltier), alcohol-vapor chamber with a fixed spike-in source port and an optional magnet for future curvature-based charge discrimination.




Fig. 2. Processing and analysis pipeline, from synchronized multi-camera video through triangulated 3D feature extraction to classical and Bayesian classification. The 2D single-camera comparison arm feeds the same classifiers to isolate the effect of 3D reconstruction.

- [10] J. Zhou, Z. Yang, L. Han, W. Zhang, K. Shi, S. Xu, and Y.-S. Liu, "4C4D: 4 camera 4D Gaussian splatting," in *Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR)*, 2026.
- [11] S. Al-Naabi and A. Radi, "A Peltier-cooled diffusion cloud chamber for visualizing charged particle tracks," *Int. J. Res. Innov. Appl. Sci.*, vol. 10, no. 11, pp. 1584–1591, 2025.
- [12] CERN S'Cool LAB, "Cloud chamber – DIY manual," International Particle Physics Outreach Group, 2022. [Online]. Available: https://ippog.org/ippog_resource_database/cloud-chamber-diy-manual

V. REFERENCES

- [1] C. T. R. Wilson, "On an expansion apparatus for making visible the tracks of ionising particles in gases and some results obtained by its use," *Proc. R. Soc. Lond. A*, vol. 87, no. 595, pp. 277–292, 1912.
- [2] K. Kleinknecht, *Detectors for Particle Radiation*, 2nd ed. Cambridge, U.K.: Cambridge Univ. Press, 1998.
- [3] S. N. Axani, K. Frankiewicz, and J. M. Conrad, "The CosmicWatch desktop muon detector: A self-contained, pocket sized particle detector," *J. Instrum.*, vol. 13, p. P03019, 2018.
- [4] K. N. Borozdin, G. E. Hogan, C. Morris, W. C. Priedhorsky, A. Saunders, L. J. Schultz, and M. E. Teasdale, "Radiographic imaging with cosmic-ray muons," *Nature*, vol. 422, p. 277, 2003.
- [5] Z. Zhang, "A flexible new technique for camera calibration," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 22, no. 11, pp. 1330–1334, 2000.
- [6] G. Bradski, "The OpenCV library," *Dr. Dobb's Journal of Software Tools*, 2000.
- [7] C. M. Bishop, *Pattern Recognition and Machine Learning*. New York, NY, USA: Springer, 2006.
- [8] D. C. Montgomery, *Introduction to Statistical Quality Control*, 8th ed. Hoboken, NJ, USA: Wiley, 2020.
- [9] World Labs, "Atlas: A world model for spatial intelligence," 2026. [Online]. Available: <https://www.worldlabs.ai/blog/atlas>