Kernel Density Estimation for Characterizing Population-Level CT Dose Variability with ACR DIR Benchmarks

Mike Weaver, PhD, Azmul Siddique, MS, Gary Y. Ge, PhD and Jie Zhang, PhD
Department of Radiology, University of Kentucky, Lexington, KY 40513

2026 JULY 19–22 VANCOUVER, BC
JOINT AAPM | COMP MEETING

ABSTRACT

Purpose:
Population-level CT dose monitoring typically relies on summary statistics and benchmark comparisons, which can obscure distributional features such as multimodality and disproportionate upper tails. This preliminary feasibility study evaluated the utility of kernel density estimation (KDE) to characterize CT dose distributions using three years of ACR Dose Index Registry (DIR) data from a single facility and to identify exam categories warranting deeper quality assurance review.

Methods:
A retrospective analysis was performed using three years of adult CT facility-level DIR data. Examinations were mapped to the ten adult exam categories emphasized in the Q3 2025 DIR reporting framework. For each category and year, CTDIvol and DLP distributions were summarized using conventional statistics (median and selected percentiles) and compared with DIR diagnostic reference levels when available. Univariate KDE was applied to visualize distribution shape, assess multimodality, and evaluate upper-tail behavior. Feasibility metrics included successful category mapping, adequacy of sample size for stable density estimation, and the ability of KDE visualizations to identify categories with non-unimodal or long-tailed distributions.

Results:
A total of 191,830 adult CT examinations across three years were included after category mapping. KDE produced stable density estimates for 9 of 10 categories meeting a prespecified minimum sample size threshold (≥100 studies). Compared with percentile summaries alone, KDE identified 7 categories demonstrating multimodal dose distributions and 2 categories exhibiting persistent upper-tail inflation relative to diagnostic reference levels, enabling rapid visual prioritization of exams for further review.

Conclusion:
Using three years of facility-level ACR DIR data, KDE was feasible for visualizing CT dose distributions and provided complementary insight beyond conventional summary metrics by revealing multimodality and upper-tail behavior. These findings support KDE as a practical screening tool for CT dose surveillance and targeted QA triage prior to more detailed protocol or scanner-level investigation.

CONTACT

Mike Weaver, PhD
University of Kentucky
Lexington, KY 40513
mike.weaver@uky.edu
859-218-0404

INTRODUCTION

Population-level CT dose surveillance commonly relies on summary statistics, percentile values, and diagnostic reference level comparisons.[1–3] These metrics are useful for benchmarking but compress each exam category into a small number of values. As a result, important distributional features such as multimodality, protocol inconsistency, scanner effects, or high-dose subgroups may be missed.

Kernel density estimation provides a complementary approach by displaying the full dose distribution rather than selected percentiles alone.[4] When KDE curves are stratified by year and overlaid with ACR DIR reference levels, they can show whether a local exam category is stable, shifted relative to national benchmarks, multimodal, or characterized by an extended upper tail.[1,2]

The objective of this preliminary feasibility study was to evaluate whether KDE visualization of standard ACR DIR export data could identify adult CT exam categories warranting deeper quality assurance review. The goal was not to replace conventional DIR benchmarking, but to determine whether distribution-level visualization adds actionable information for CT dose surveillance.

METHODS AND MATERIALS

This retrospective quality assurance analysis used three years of adult CT facility-level ACR DIR data from a single institution.[1] Examinations were mapped to ten adult ACR DIR exam categories. Study-level CTDIvol and DLP values were extracted and grouped by exam category and annual reporting period. Records with incomplete dose information, pediatric examinations, or studies that could not be mapped to a target category were excluded from KDE visualization.

For each exam category and year, study counts, medians, and selected percentiles were calculated. Local dose distributions were compared with ACR DIR reference levels when available.[1,2] Categories were required to meet a prespecified minimum sample size threshold of at least 100 studies to be included in the final visualization.

Univariate Gaussian KDE was applied separately within each exam category and year.[4] Bandwidth selection used data-driven cross-validation to balance oversmoothing against excessive sensitivity to individual observations. Density curves were displayed as ridge plots, with vertical markers indicating the 25th, 50th, and 75th percentile ACR DIR benchmark values.[1,2]

Primary feasibility endpoints included successful category mapping, adequate sample size for stable density estimation, and the ability of KDE to reveal distributional structure not apparent from percentile summaries alone. Findings of interest included multimodality, persistent right-sided tail extension, and displacement of local distributions relative to DIR benchmarks.

RESULTS

A total of 191,830 adult CT examinations were included after category mapping. Of the ten mapped adult DIR categories, nine met the minimum sample size threshold and were included in KDE ridge plot visualization. The excluded category did not provide sufficient sample size or stable distributional behavior for meaningful interpretation.

Across included categories, KDE demonstrated generally stable distribution shapes across the three annual periods. However, the density plots revealed structure not captured by medians or percentile summaries alone. Seven categories demonstrated multimodal dose distributions, suggesting multiple subpopulations or protocol behaviors within the same mapped exam category.

Two categories demonstrated persistent upper-tail inflation relative to the 75th percentile DIR benchmark.[1,2] This pattern suggests that a subset of examinations had substantially higher dose indices than the central portion of the local distribution. KDE helped distinguish whether benchmark elevation reflected a broad distributional shift or a smaller high-dose subgroup.

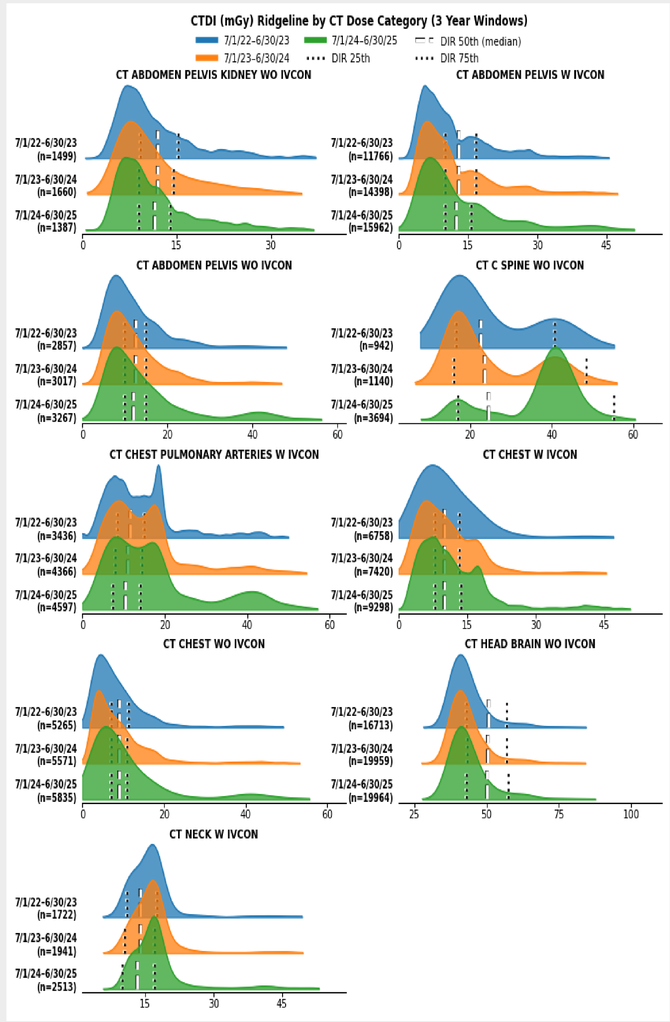

Figure 1. Figure 1: KDE ridge plots for nine adult ACR DIR categories with ≥100 examinations. For each category, density curves for three consecutive years are overlaid to illustrate temporal stability and distributional shape. Vertical lines indicate the 25th, 50th, and 75th ACR DIR reference levels. Multimodality and extended upper tails relative to reference levels are visible in selected categories.

DISCUSSION

KDE provides a direct way to inspect the full CTDIvol or DLP distribution for each DIR exam category, rather than relying only on medians and selected percentiles. This is useful because a single elevated percentile does not indicate whether the issue is a broad shift in the entire category or a smaller subset of high-dose examinations.

In this dataset, several categories showed multimodal distributions or persistent right-sided tails. For CT dose review, these features are the useful findings. A secondary peak may indicate mixed protocol usage, scanner-specific behavior, different clinical indications grouped within the same DIR category, or patient-size effects. A long upper tail identifies the region of the distribution where outlier studies can be selected for case-level review.

The intended use of KDE in this workflow is therefore to guide follow-up analysis. Categories with stable, compact distributions may require no immediate review. Categories with distinct secondary peaks or extended upper tails can be prioritized for stratification by scanner, protocol name, acquisition parameters, scan length, and patient habitus. The highest-dose studies within these regions can then be reviewed directly to determine whether the observed distribution reflects appropriate clinical variation or correctable protocol/process variation.

In this sense, KDE complements DIR benchmarking. DIR percentiles indicate how the local distribution compares with national reference data, while KDE helps identify where within the local distribution the relevant cases are located.

CONCLUSION

KDE visualization of standard ACR DIR data was feasible for adult CT dose surveillance at a single institution.[1,4] Compared with conventional percentile summaries alone, KDE provided additional information about distribution shape, multimodality, and upper-tail behavior.

This approach may be useful as a QA triage tool. Categories with stable unimodal distributions may require routine monitoring, while categories with multimodality or persistent upper-tail inflation can be prioritized for scanner-, protocol-, and patient-level review. KDE should therefore be viewed as complementary to DIR benchmarking: it helps identify where to look next, not as a replacement for clinical judgment or detailed protocol analysis.[1–3]

REFERENCES

1. American College of Radiology. ACR Dose Index Registry.
2. Kanal KM, Butler PF, Sengupta D, Bhargavan-Chatfield M, Coombs LP, Morin RL. U.S. Diagnostic Reference Levels and Achievable Doses for 10 Adult CT Examinations. Radiology. 2017;284(1):120–133.
3. American College of Radiology, American Association of Physicists in Medicine, Society for Pediatric Radiology. ACR–AAPM–SPR Practice Parameter for Diagnostic Reference Levels and Achievable Doses in Medical X-Ray Imaging.
4. Silverman BW. Density Estimation for Statistics and Data Analysis. Chapman & Hall; 1986.