A Novel Multilayer Task-Centric and Data Quality Framework for Autonomous Driving

Yuhan Zhou, Haihua Chen, and Kewei Sha, *University of North Texas, Denton, TX, 76203, USA*

The next-generation autonomous vehicles (AVs), embedded with frequent real-time decision-making, will rely heavily on massive multisource and multimodal data. In real-world settings, the data quality (DQ) of different sources and modalities varies due to unexpected environmental factors or sensor issues. However, both researchers and practitioners in the AV field overwhelmingly concentrate on models/algorithms while undervaluing the DQ. To improve an AV's functionality, efficiency, and trustworthiness, this article proposes a novel task-centric and data quality vase framework, which consists of five layers: data, DQ, task, application, and goal. The proposed framework aims to map DQ with task requirements and performance goals. This paper opens up a range of critical but unexplored challenges at the intersection of DQ, task orchestration, and performance-oriented system development in AVs. It is expected to guide the AV community toward building more adaptive, explainable, and resilient AVs that respond intelligently to dynamic environments and heterogeneous data streams.

The recent deployment of RoboTaxi by Waymo, Tesla, and Mobileye demonstrates the transformative potential of autonomous driving (AD) to improve transportation safety, reduce traffic congestion, and enhance mobility for all.[1] However, the widespread adoption of autonomous vehicles (AVs) still faces significant challenges, particularly in ensuring real-time, accurate, and reliable decision-making in a dynamic environment that involves long-tail risks and high uncertainty. These challenges are compounded by constrained computational and storage resources in AVs.

AVs make intelligent decisions using machine learning (ML) models trained on multisource, multimodal data collected from diverse sensors. The efficiency and effectiveness of AV decisions depend not only on the capacity of ML models but also critically on the quality of data.[2] While significant progress has been made in model development and benchmark-driven evaluation, current AV frameworks often overlook the critical role of data quality (DQ) in determining system performance, such as accuracy, robustness, and efficiency.

Existing data quality frameworks[3] are mostly designed for traditional, homogeneous datasets and fall short in addressing the following complex demands of AV systems. First, AVs must handle multisource, multimodal data streams while simultaneously supporting a wide range of applications, each of which may impose different quality requirements on shared or unique data streams. Second, cost constraints resources in AVs, making it essential to carefully balance data quality with task performance. Third, AVs operate under stringent real-time constraints, where even a correct decision is ineffective if it arrives too late. Thus, a mechanism to evaluate the data quality and select an appropriate subset of data that satisfies all the previous requirements is vital.

While prior research has explored dataset characterization,[4] curation,[5] generation,[6] and testing,[7] to the best of our knowledge, there remains a lack of a unified framework that systematically addresses data quality in the AV context. To bridge this gap, we take a first step to propose a task-centric, multisource, and

Digital Object Identifier 10.1109/MIC.2026.3665121
Date of publication 23 February 2026; date of current version 15 May 2026.

multimodal data quality framework tailored for AVs. This five-layer framework uses decision-making tasks as the central link between data quality metrics and AV performance objectives, offering a novel perspective for evaluating and improving data quality in real-time, safety-critical scenarios.

The contributions of this article are twofold. We first propose a five-layer vase framework with tightly coupled layers—data, data quality, task, application, and goal. Our system treats driving tasks as first-class entities, orchestrates data pipelines in real-time, and evaluates decisions through multidimensional data quality indicators. We not only highlight these critical yet overlooked factors but also advocate for recognizing the relationships among them for iterative improvement. Building on this vision, we put forward open discussions and future directions. We outline a few critical research questions and directions under the multilayer task-centric and data quality framework. These open research questions will inspire and guide the community to explore more efficient and reliable techniques for AVs.

A NOVEL TASK-CENTRIC AND DATA QUALITY VASE FRAMEWORK FOR AVs

Traditional frameworks present the following research gaps:

1) They mainly emphasize model-centric architectural designs and overlook data quality despite its significance in assuring the performance, efficiency, and reliability of AVs.
2) While the quality of multisource and multimodal data in AVs varies under dynamic environments, existing DQ frameworks are not well-suited to handle these complexities, highlighting the need for a novel DQ framework tailored to the unique challenges of AV systems.
3) Current strategies not only overlook DQ evaluation but also treat all the AV tasks with similar DQ requirements. Consequently, the lack of DQ evaluation and task-centric awareness prevents AV systems from feeding quality diagnostics back into the loop and iterating on the performance.

Motivated by addressing the aforementioned gaps and inspired by the five-layer architecture of the Internet of Things,[8] we propose a multilayer task-centric and data quality vase framework in Figure 1. The vase framework is a five-layer architecture: data, data quality, task, application, and goal. It emphasizes that multisource and multimodal data quality must be tailored to the specific goals of each task and its broader application. This framework can guide task-centric data selection and quality evaluation to improve robustness, efficiency, and reliability for AD.

Five-Layer Framework Breakdown

Data Layer

Data in AVs presents two main features: multisource and multimodal, as shown in Figure 2. At the multisource level, AVs gather information from various sensors, such as lidar, cameras, and radar, as well as other feeds like information collected from connected autonomous vehicles (CAVs), user behavior data, weather reports, and infrastructure updates. Since one AV can be equipped with more than one camera, the front, back, and side cameras are considered different data sources. This blend of inputs can validate data and reduce the impact of errors or noises from individual sensors, which enables a more comprehensive view of the driving environment and a more accurate and robust performance.

Modalities are the corresponding data types generated by the sources and widely discussed in data fusion techniques. Multimodal in AVs refers to the diverse forms and types in which data are captured and represented, including raw sensor readings, images, textual data, video, and audio.

Combining multisource and multimodal data can enhance the performance and reliability in AVs compared to a single data source or modality. Zhang et al. improved vehicle localization by integrating vehicle headlights and taillights detected from multiple cameras.[9] Sanchez et al. found that single-source lidar detectors generalize poorly when AVs enter unfamiliar zones, and multisource lidar training can improve domain generalization by over 10% and source-to-source transfer accuracy by 5.3% over any individual dataset.[10] Sensor complementarity is another reason why multisource and multimodal data are essential. Under night, fog, rain, or snow, camera performance degrades, and fusing radar or lidar restores robustness, increasing detection score by up to 29%.[11] In dense urban scenarios, cameras often miss vehicles occluded by large objects, while lidar and radar reveal them.[12] When lidar points are sparse for tiny objects, a camera could help recover them by visual cues.

Data Quality Layer

While multisource and multimodal data enhance AV perception and decision-making, they also introduce data quality challenges with varying accuracy and noise. Conflicting information from various

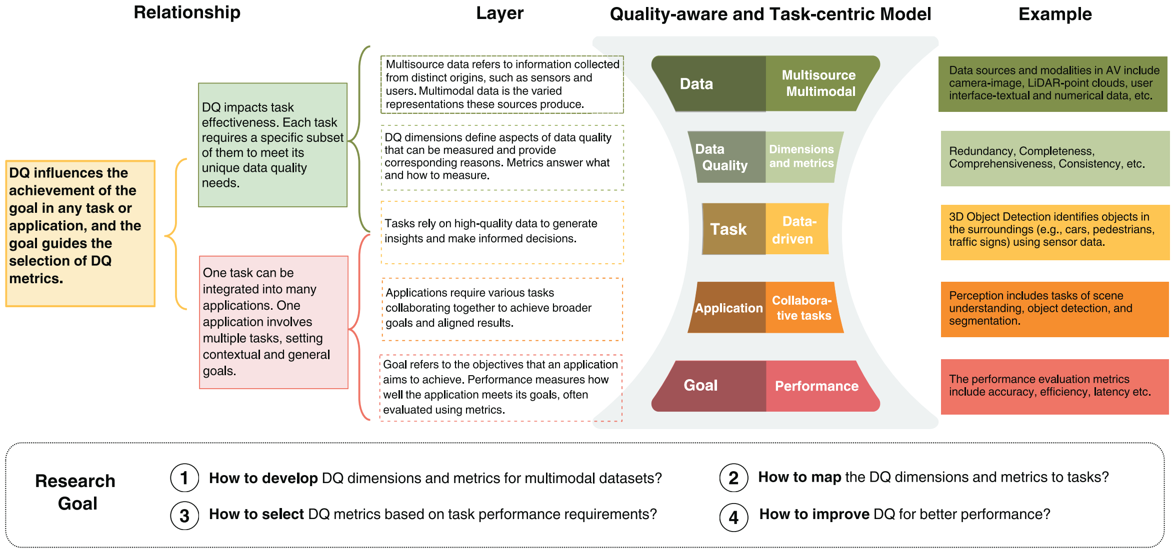

FIGURE 1. The multilayer task-centric and data quality vase framework. It integrates five layers, with corresponding examples, explanations, and relationships.

sources and modalities could impact data consistency and correctness. Fusing these data may cause calibration incorrectness, redundant information, and imbalanced issues. To address these, a structured data quality layer can contribute to the evaluation and improvement of data. The following introduces key data quality dimensions designed for multisource and multimodal data in AVs to emphasize the quality-aware focus of this framework. Table 1 presents the main and common DQ dimensions, which could be expanded or tailored based on different scenarios.

Traditional DQ models are inadequate for evaluating multisource and multimodal data in AVs, which introduces unique challenges that necessitate a more rigorous data quality evaluation framework. DQ dimensions of multisource and multimodal data in AVs are similar but not identical. Both types require correctness, consistency, relevance, and timeliness to ensure reliable perception and decision-making. At the same time, the definition of the same DQ dimension may vary for multisource and multimodal data. For example, in multimodal data, completeness refers to capturing and integrating all necessary data from different sensors, whereas in multisource data, it emphasizes collecting data that covers the required spatial areas and time intervals. Redundancy can arise from duplicate detections by multiple sensors on the same vehicle or from overlapping reports across AVs and infrastructure nodes. Managing redundancy effectively prevents unnecessary data processing while preserving critical information for robust task performance. Redundant multimodal data occurs when multiple sensors provide overlapping information, which can be beneficial for fault tolerance and safety, but inefficient if excessive.[7]

Task Layer

Data quality evaluation and improvement should follow the principle of "fit for purpose."[14] There are two reasons for the task-centric layer design:

1) The specific task requirements should guide the selection of data quality dimensions. Each task also involves various conditions, such as weather, scenarios, and light, impacting its own set of critical quality attributes.
2) In real-world scenarios, gathering perfect data to meet all quality dimensions is challenging. Therefore, it is essential to select the most critical dimensions that fit specific tasks and contexts in AVs.

Targeted data quality improvements can directly enhance the performance of each module in the closed-loop system. Tasks such as object detection, segmentation, and tracking require investigation and comparison of the detailed characteristics of the dataset.[4] For instance, tasks that involve real-time decision making, such as obstacle detection, may prioritize

TABLE 1. Data quality dimensions, explanations, and their impacts in AV tasks.

DQ dimension	Explanation	AV task	Impact
Completeness	The extent to which all required data elements are present for a task. For labeling, it ensures all valid objects in a frame are labeled.[2]	Object detection and tracking	Missing annotations and bounding boxes reduce recall and increase bias.
Consistency	The extent to which information provided by sensors and modalities is noncontradictory.	Vehicle recognition and labeling	Cross-modal divergence can trigger a high-fault flag in an information-fusion monitor.[12]
	Data remains noncontradictory over time without discrepancies due to latency, environmental changes, or synchronization issues.	Semantic perception	The representation and labels for the same object viewed at different points in time should be the same.
Correctness	The extent of data values or labels reflects real-world entities, positions, and attributes.[2]	Object detection under adverse weather	Raindrops and snow on windshields or camera lenses impair visibility.
Noise level	The degree to which a sensor stream is contaminated by random or systematic distortions that obscure the true signal.	Object detection for night-time urban driving	High-ISO noise and motion blur in camera images inflate the false-negative rate.
Redundancy	The presence of overlapping information does not increase task performance.[7]	Scene understanding	Consecutive lidar data are highly redundant due to the high sampling frequency and the massive surrounding environment information besides the objects of interest.
Relevance	The degree to which the data contributes useful information to the current task.	Trajectory prediction and planning	Irrelevant or weakly related inputs can cause unsafe actions, such as reacting to distant objects first.
Timeliness	Data is delivered and processed before a deadline.	Prediction and path planning	If the trajectory prediction of localized objects misses its deadline, the planner may make outdated motion hypotheses, leading to unsafe conditions.[13]

timeliness and correctness over others. In contrast, tasks like post-event analysis might place higher value on data completeness. From the task-centric perspective, the cross-modal redundancy requirements may vary for different performance goals. For driving on a highway, the buildings as background are redundant data; For localizing and parking, the building data are not as redundant because this task requires more details on the buildings.

Application Layer

Applications are organically aggregated tasks. One task can be mapped to multiple applications, and one application can involve many goals. For example, object detection is a fundamental task across most AD applications. Recognizing the overall context and subtasks can help refine data quality requirements and ensure performance.

There are different categorizations of autonomous driving applications. A widely adopted application pipeline is a serial design of perception, prediction, planning, and controlling. Perception is the detection and interpretation of objects and their surroundings using sensors. Prediction anticipates future movements of detected objects based on current data. Planning refers to deciding an optimal trajectory by the environmental context and the predicted behavior of objects. Controlling executes these planned vehicle operations, such as steering, acceleration, and braking. There are also end-to-end approaches integrating the pipelines.

Goal Layer

The goal layer is the highest level within the vase framework, providing the overall objectives that applications and tasks aim to achieve, including efficiency,

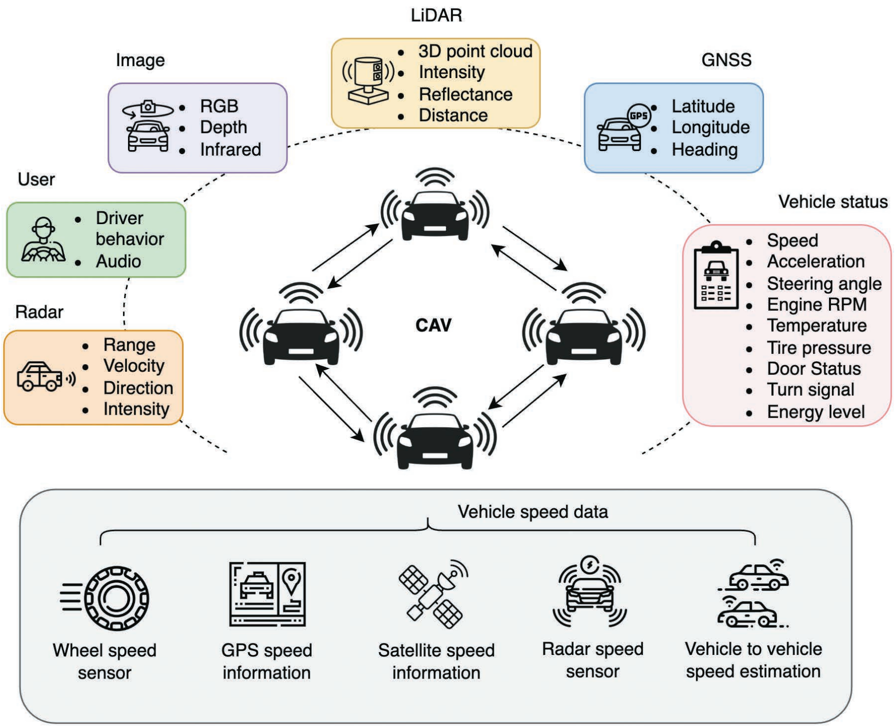

FIGURE 2. Illustration of multisource and multimodal data in autonomous driving.

latency, and accuracy. Without this layer, data quality and task optimizations could not get feedback and improve iteratively. Taking pedestrian prediction as an example, the Goal layer defines high-level performances as sufficient accuracy and low response latency. Flagged sensor data informs quality issues in data curation and labeling, and engineers could refine the model to restore speed or precision.

High efficiency can enhance the driving experience and reduce operational costs and environmental impact in AVs. Latency evaluates if the model's response time and processing speed satisfy the real-time demands of autonomous driving across data collection, computation, consumption, and communication. High latency can compromise real-time decision-making capabilities, resulting in delayed reactions to dynamic road conditions and potentially increasing safety risks. High accuracy typically demands high computational power, leading to increased energy consumption and reduced efficiency.

Relationships Between Layers

The relationships between layers are important for the following reasons. First, there are tradeoffs in AV design. Higher-resolution sensors can capture smaller or distant objects more reliably, but also increase data processing costs. In addition, enhancing one DQ dimension may come at the cost of another. These tradeoffs must be carefully managed to align with overall goals. Second, feedback loops in these layers provide better traceability and interpretability, answering questions such as "For object detection at night, what are the DQ requirements to achieve an 80% precision?" Finally, a relational view supports modular iteration and deployment. New sensors, learning algorithms, or application modes can be integrated quickly when their relationships in a task are clear. In summary, analyzing the interlayer relationships can enhance a coherent, verifiable, and adaptive AD system. The aforementioned reasons also clarify that tasks, data, and quality are not isolated elements, but are part of an overall framework.

Figure 3 presents how the framework's layers interact in loops. Consider an object detection and tracking task within a perception application. Current practice seeks robustness by adding heavier models, stacking sensors, or applying data augmentation strategies.[15] These measures address only parts of the DQ issues and rarely explain how each change influences downstream performance, making the system less adaptable to task requirements.

Under our framework, the workflow changes. Assume the task is performed in a scene with good visibility. During data selection, the pipeline may keep

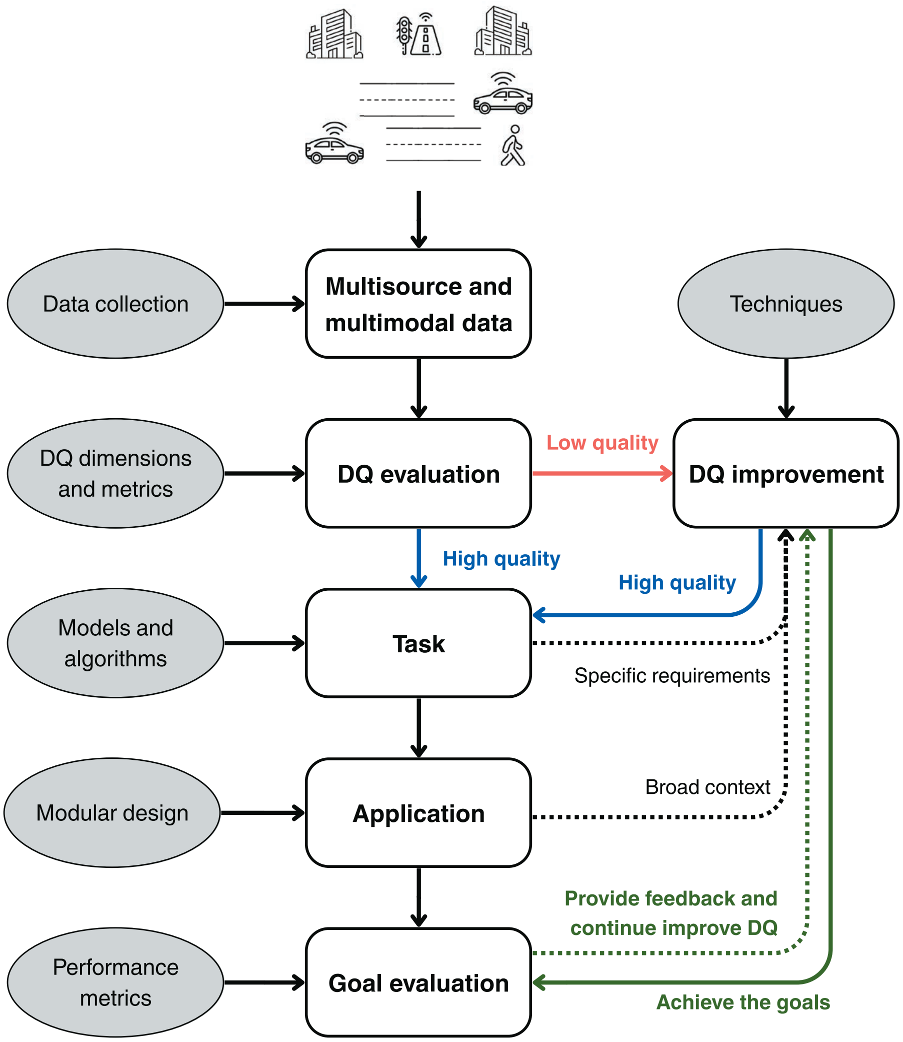

FIGURE 3. Illustration of the relationships between layers in the vase framework. The red text and lines indicate low-quality data that requires improvement, and the blue ones represent the high-quality data that can fulfill the task goals. The loops between DQ improvement and task highlight that this data quality framework is task-centric.

only the image modality and discard lidar streams. The subsequent DQ evaluation stage chooses redundancy as its primary dimension because other dimensions, such as correctness, are satisfied under this setting. The evaluation quantifies overlap among camera views and retains only nonredundant information for training. Next, performance metrics validate the data and model. If the reduced redundancy degrades detection accuracy, the loop returns to DQ improvement and relaxes the redundancy thresholds. The flow repeats until the data satisfies the task-specific quality. This process shows how the continuous feedback turns data, data quality, and task performance into mutually reinforcing levers, driving the system toward ever-higher robustness and efficiency.

OPEN RESEARCH QUESTIONS

While our proposed task-centric and DQ framework offers a foundational approach, it also opens up a range of critical but unexplored challenges that merit further investigation. This section outlines key open research questions at the intersection of DQ, task orchestration, and performance-oriented system development in AVs.

How to develop a taxonomy of task-specific DQ dimensions for multisource and multimodal data in AVs? A formal taxonomy would help standardize how quality is defined for images, lidar, radar, GPS, and so on. In addition, how does each DQ type contribute to specific tasks? For example, how do spatial resolution, temporal consistency, and labeling accuracy impact object tracking, while pixel-wise accuracy, label distribution balance, and occlusion influence semantic segmentation?

How to design DQ metrics that are task-adaptive and dynamic across AV modules? Different tasks, such as object detection, lane keeping, trajectory planning, and emergency braking in AVs, may require different sources and modalities of high-quality data. A single static DQ metric may not be sufficient. This question will investigate task-conditioned DQ metrics that dynamically adjust based on the goal and current environmental context. For example, in low-light environments, the DQ metric could weigh lidar data higher than image data for perception. The objection detection task in AVs may prioritize visual clarity and spatial resolution, while the trajectory planning task may prioritize temporal continuity.

What is the optimal tradeoff between data volume, DQ, and system latency in AVs? As datasets for AVs continue to expand, the cost of storing, processing, and computing for multisource and multimodal data, especially videos and images, will escalate rapidly. Exploiting methods for subsampling large datasets and evaluating the most relevant DQ dimensions to a set of downstream multimodal vision and language tasks are critical. This question explores how to mathematically model and optimize tradeoffs between data quality, task performance, and latency.

How can the AV systems detect and respond to dynamic DQ changes in real time? Sensor degradation due to dirt, fog, rain, occlusion, or hardware failure may compromise data streams. Investigating real-time DQ monitoring mechanisms and autonomous failover strategies will enhance the robustness, adaptability, and safety of AVs. For example, if the front camera is occluded, how can we design an algorithm to guide the system to reweight lidar and side cameras for decision-making within a short time?

How can DQ-aware mechanisms be applied to enhance explainability in AVs? Explainable artificial intelligence (XAI) in AVs aims to understand why a model made a particular decision, which is very important for public trust, regulatory compliance, and system safety. However, current XAI in AVs focuses on the model internals, such as feature attributions, attention weights, or decision trees, while ignoring that DQ could be an equally critical source of error. This question will bridge DQ with XAI by understanding the relationship between the quality of input data and the explanation made in AVs.

How can a task-centric DQ evaluation framework be used to support synthetic data generation and augmentation in AVs? Data synthesis and data augmentation using LLMs have been widely used in AVs. However, one of the key challenges is to evaluate and select high-quality data for training. The proposed framework can be instrumental in determining what synthetic data to generate, evaluating the quality of synthetic samples from the perspective of downstream tasks, and helping determine which synthetic data should be integrated into training datasets.

How can we integrate feedback from human users into a task-centric DQ framework to enhance data-driven decision-making in AVs? Semi-autonomous driving systems, such as those offering driver-assist features, involve continuous human–machine collaboration. In such systems, the vehicle may perform most tasks, but human drivers still oversee, intervene, or take over under certain conditions. It opens up an interdisciplinary avenue involving human–computer interaction, cognitive modeling, and data engineering.

FUTURE DIRECTIONS

Future directions aim to guide the community toward building more adaptive, explainable, and resilient AVs that respond intelligently to dynamic environments and heterogeneous data streams.

DQ Metrics for AV Benchmarks and Task-Centric Evaluation

Guided by the vase framework, future work will focus on defining, modeling, and measuring DQ metrics using established AD benchmarks. The proposed framework should also be extended across multiple data sources and sensor modalities, including camera images, lidar cloud points, and radar. In addition, it is essential to investigate how different AV tasks prioritize DQ under varying operational conditions, considering driving scenarios and environmental factors. These conditions not only affect multisource and multimodal inputs and the standards for DQ metrics, but also shape the relationships among each layer. As the five layers form a "many-to-many" mapping relationship, further research could conduct extensive studies of "multi dimension-one task-multi goal," "multi dimension-multi task-multi goal," and so on.

Quality-Aware Data Selection

Developing an integrated and practical data quality evaluation workflow based on this framework is essential. This overarching approach should specify the methodologies to systematically and iteratively evaluate, improve, and monitor data quality in AVs and other important multimodal domains. To select high-quality data and improve low-quality data, future studies could adopt active learning to selectively retain, discard, or request new data based on evolving task demands and data quality profiles.

Hardware-Aware DQ Optimization on the Edge

Future work should move beyond software-only fixes and pursue sensor–pipeline co-design. This calls for joint optimization models that balance sensor settings, edge-compute schedules, data processing pipelines, and quality-aware task adaptation.

DQ-Aware Simulation and Digital Twins for Safety Evaluation

Current simulators vary in weather and traffic but assume "perfect" data. The next step is to simulate the effects of varying data quality in synthetic environments, such as injecting controlled quality degradation of sensor noise and time-sync drift. Embedding these capabilities in digital-twin platforms would accelerate investigation of the feedback loop, robust testing, certification, and fault injection studies.

Reinforcement Learning (RL) Agents for DQ-Task–Goal Alignment

RL agents could help dynamically adjust data selection, sensor configurations, and computation strategies to optimize task outcomes. By optimizing a reward that aligns DQ metrics with task goals, an RL agent can navigate complex tradeoffs.

Multimodal Models and Agentic AI in AVs

Multimodal foundation models trained on large-scale and heterogeneous datasets without rigorous DQ evaluation may face cross-modal inconsistency, redundancy, object class imbalance issues, and others. Moreover, agentic driving systems rely on iterative perception–reasoning–action loops. DQ issues in one modality or one source can propagate across planning and decision layers. Overlooking how multimodal misalignment, annotation ambiguity, and source imbalance can constrain model robustness, generalization, and safety.

Human-in-the-loop Interfaces for DQ-Aware Debugging and Governance

Data engineers and regulators need to capture how data quality propagates into decisions. Visual analytics dashboards could plot quality metrics alongside task goals, enable data quality-task relationship checks, facilitating simulations and debugging under altered quality thresholds. Advances in this area will facilitate DQ standard design, foster transparency, and build public trust in DQ-aware autonomous systems.

CONCLUSION

AV systems intake and output data from diverse sources and modalities, and they need a holistic DQ framework that tailors quality dimensions to each task's unique demands and ensures data meets required standards at each stage. This underscores the necessity of a framework that aligns data quality with task requirements and overall objectives, bridging data-level quality metrics and application-level goals. Motivated by this, we put forward a five-layered framework and emphasize that the data work in this field should be task-centric and quality-aware. The relationships between the five layers could further guide experiment design. This again emphasizes the idea that dataset choice should be aligned with the task requirements in a structured and goaloriented way. At last, the research goals and open questions point out the following main directions of our study.

From a practical perspective, we outline the steps for operationalizing data quality across multisource and multimodal AV systems. First, the design of DQ metrics should consider the source and modality features. Second, DQ evaluation should be task-centric. Recognizing that different tasks require varying standards for specific DQ issues. Third, the framework should support condition-aware monitoring with adaptive DQ thresholds. Finally, feedback loops can be incorporated by linking each layer, from targeted data collection to the task goals, enabling continuous DQ improvement throughout the AV systems.

ACKNOWLEDGMENTS

The authors would like to thank Dr. Junhua Ding from the Department of Data Science, and Drs. Heng Fan and Yunhe Feng from the Department of Computer Science and Engineering at the University of North Texas for their precious feedback and suggestions. This work was supported in part by the National Science Foundation (NSF) under Award 2505686.

REFERENCES

1. D. Omeiza, H. Webb, M. Jirotka, and L. Kunze, "Explanations in autonomous driving: A survey," *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 8, pp. 10,142–10,162, Aug. 2022, doi: 10.1109/TITS.2021.3122865.

2. H. Chen, J. Chen, and J. Ding, “Data evaluation and enhancement for quality improvement of machine learning,” *IEEE Trans. Rel.*, vol. 70, no. 2, pp. 831–847, Jun. 2021, doi: 10.1109/TR.2021.3070863.
3. M. Priestley, F. O’donnell, and E. Simperl, “A survey of data quality requirements that matter in ml development pipelines,” *ACM J. Data Inf. Qual.*, vol. 15, no. 2, pp. 1–39, 2023, doi: 10.1145/3592616.
4. M. Liu et al., “A survey on autonomous driving datasets: Statistics, annotation quality, and a future outlook,” *IEEE Trans. Intell. Veh.*, vol. 9, no. 11, pp. 7138–7164, Nov. 2024, doi: 10.1109/TIV.2024.3394735.
5. H. Caesar et al., “nuScenes: A multimodal dataset for autonomous driving,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 11,618–11,628, doi: 10.1109/CVPR42600.2020.01164.
6. L. Chen et al., “Driving with LLMs: Fusing object-level vector modality for explainable autonomous driving,” in *Proc. IEEE Int. Conf. Robot. Automat. (ICRA)*, Piscataway, NJ, USA: IEEE Press, 2024. pp. 14,093–14,100, doi: 10.1109/ICRA57147.2024.10611018.
7. F. Li, C. Fu, D. Sun, H. Marzbani, and M. Hu, “Reducing redundancy in maps without lowering accuracy: A geometric feature fusion approach for simultaneous localization and mapping,” *ISPRS Int. J. Geo-Inf.*, vol. 12, no. 6, 2023, Art. no. 235, doi: 10.3390/ijgi12060235.
8. M. Wu, T.-J. Lu, F.-Y. Ling, J. Sun, and H.-Y. Du, “Research on the architecture of Internet of Things,” in *Proc. 3rd Int. Conf. Adv, Comput. Theory Eng. (ICACTE)*, Piscataway, NJ, USA: IEEE Press, 2010, pp. V5-484–V5-487, doi: 10.1109/ICACTE.2010.5579493.
9. X. Zhang, B. Story, and D. Rajan, “Night time vehicle detection and tracking by fusing vehicle parts from multiple cameras,” *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 7, pp. 8136–8156, Jul. 2022, doi: 10.1109/TITS.2021.3076406.
10. J. Sanchez, J.-E. Deschaud, and F. Goulette, “COLA: COarse-LAbel multi-source LiDAR semantic segmentation for autonomous driving,” *IEEE Trans. Robot.*, vol. 41, pp. 1742–1754, 2025, doi: 10.1109/TRO.2025.3543302.
11. H. Su et al., “Object detection in adverse weather for autonomous vehicles based on sensor fusion and incremental learning,” *IEEE Trans. Instrum. Meas.*, vol. 73, pp. 1–10, 2024, doi: 10.1109/TIM.2024.3472860.
12. W. Hou, W. Li, and P. Li, “Fault diagnosis of the autonomous driving perception system based on information fusion,” *Sensors*, vol. 23, no. 11, 2023, Art. no. 5110, doi: 10.3390/s23115110.
13. T. Wu and W. Shi, “Timeliness in autonomous driving: Hype or reality?” *IEEE Internet Comput.*, vol. 28, no. 5, pp. 75–84, Sep./Oct. 2024, doi: 10.1109/MIC.2024.3441869.
14. R. Y. Wang and D. M. Strong, “Beyond accuracy: What data quality means to data consumers,” *J. Manage. Inf. Syst.*, vol. 12, no. 4, pp. 5–33, 1996, doi: 10.1080/07421222.1996.11518099.
15. C. Zhang and A. Eskandarian, “Image-guided outdoor LiDAR perception quality assessment for autonomous driving,” *IEEE Trans. Intell. Veh.*, vol. 10, no. 9, pp. 4455–4466, Sep. 2025, doi: 10.1109/TIV.2024.3485004.

YUHAN ZHOU is a Ph.D. student of data science at the College of Information, University of North Texas, Denton, TX, 76203, USA. Her research interests include data quality, including frameworks, dimensions, open source tools, and improvement methods. Zhou received her M.S. degree in business analytics from The University of Hong Kong. Contact her at yuhanzhou@my.unt.edu.

HAIHUA CHEN is an assistant professor of data science at the College of Information, University of North Texas (UNT), Denton, TX, 76203, USA. His research interests include data quality, health informatics, applied data science, and data-centric artificial intelligence, with a special focus on improving data quality and developing effective machine learning models for important domains such as education, health care, legal intelligence, and cybersecurity. Chen received his Ph.D. degree in information science from UNT. He is a Member of IEEE. Contact him at haihua.chen@unt.edu.

KEWEI SHA is an associate professor of data science within the College of Information, University of North Texas, Denton, TX, 76203, USA. His research interests include various aspects of data quality, including modeling, management, and enhancement. Contact him at kewei.sha@unt.edu.