International Research Journal of Engineering and Technology (IRJET) e-ISSN: 2395-0056
Volume: 12 Issue: 02 | Feb 2025 www.irjet.net p-ISSN: 2395-0072

Data Quality Guardrails for Artificial Intelligence Applications

Tapan Parekh

Data Engineer at Amazon, New York, NY.

Abstract – *Artificial Intelligence (AI) applications require high-quality data for success because poor data quality creates biased models and inaccurate predictions and causes operational inefficiencies. The research documents the essential function of data quality guardrails in AI systems while introducing a detailed framework to control data integrity, consistency, and fairness. The paper investigates how defective data affects AI decision-making processes while demonstrating instances where incomplete or biased datasets led to operational breakdowns. The paper proposes basic data quality principles, including accuracy, completeness, consistency, timeliness, and validity, together with recommended practices for data governance and other processes to effectively handle these challenges. The paper evaluates leading industry solutions and frameworks that enable scalable and compliant AI decision-making processes. The study highlights that AI models remain trustworthy across different sectors through ongoing development and with high quality data.*

Key Words: **Data Quality, Data Integrity, Data Governance, Bias, Data Cleansing for AI Models, ETL and AI Data Pipelines, Master Data Management for AI**

1.INTRODUCTION

The rapid progression of artificial intelligence (AI) and machine learning (ML) technologies has transformed the manufacturing, entertainment, healthcare, and finance sectors. These technologies enable automation and improved decision-making, which leads to higher operational efficiency. The performance of AI and ML projects relies heavily on maintaining high-quality data standards.

AI systems depend on data to generate insights and automated decisions while enhancing user experiences. When data quality is inadequate, it leads to distorted models and inaccurate forecasts and produces ineffective business strategies. The dependability and efficiency of AI-driven solutions require strong data quality guidelines to be established.

2. IMPACT OF POOR DATA QUALITY ON AI MODELS

AI models are only as good as the data they are trained on. If the data is inconsistent, incomplete, or biased, the resulting AI application may produce unreliable outcomes.

Here are several examples of AI failures caused by poor data quality:

2.1 AI-Powered Resume Screening System – Hiring Bias

A new AI tool was designed to handle resume screening and candidate selection automatically. The system repeatedly granted advantages to individuals from specific demographic groups while denying opportunities to other equally or more qualified candidates. The problem emerged because the training data originated from previous hiring choices that preserved biases regarding gender, ethnicity and educational background. The AI system perpetuated existing discriminatory patterns which produced biased hiring results.

2.2 Healthcare Risk Prediction – Unequal Treatment

Medical experts developed a machine learning tool to identify patients who needed extra medical attention. The system unfairly allocated resources to some groups while ignoring others with equivalent medical needs. The AI system's determination of patient needs through healthcare spending created bias because historically disadvantaged groups exhibited less recorded healthcare spending. The AI system incorrectly assessed patient condition levels because incomplete or biased training data led to inequitable healthcare decisions.

2.3 Facial Recognition Misidentification – Racial and Gender Bias

The facial recognition system used by security forces experienced difficulties in correctly identifying people from different backgrounds. The system showed significantly higher error rates for darker-skinned individuals and women but maintained strong performance with lighter-skinned male faces. The AI training dataset contained inadequate examples of diverse ethnicities and genders which led to the problem. The system often misidentified individuals which led to widespread worry about prejudice within AI-based surveillance and security tools.

2.4 AI Chatbot – Toxic and Offensive Responses

An AI chatbot designed to engage in human-like conversations became problematic when it started generating inappropriate, offensive, and misleading content.

The system was trained on vast amounts of internet text without adequate filtering, allowing it to learn from unmoderated sources that contained biased, false, or harmful information. Without proper safeguards, the chatbot amplified these patterns, demonstrating the risks of training AI on unchecked, poor-quality data.

2.5 Predictive Policing – Reinforcing Bias in Crime Forecasting

The predictive policing AI system was designed to detect areas with high crime rates so that law enforcement resources can be allocated more effectively. The system based on biased historical crime data flagged some neighborhoods more than others despite similar crime rates. The outcome of this system's implementation resulted in excessive police presence in targeted neighborhoods which strengthened pre-existing prejudices and established a cycle that continued unequal policing practices. Incomplete and historically biased crime data used to train the AI led to the emergence of this issue.

2.6 AI-Based Loan Approval – Discriminatory Lending Decisions

The AI credit scoring system implemented to streamline loan approvals disproportionately rejected applications from specific demographic groups. The AI training involved historical lending data that exhibited systemic biases concerning financial accessibility and credit history records. The loan system wrongfully rejected qualified applicants based on their association with historically disadvantaged groups. The use of representative and unbiased data sets during training is essential for developing fair financial AI models.

2.7 Automated Content Moderation – Incorrect Flagging of Posts

A content moderation system controlled by AI was introduced for social media platforms to identify and eliminate damaging or unsuitable posts. The system often misclassified harmless content as problematic but missed actual harmful material. The problem originated from incorrect labels in training data and the system's dependency on keyword filtering instead of contextual analysis. The AI moderation system incorrectly suppressed valid conversations but missed harmful content, demonstrating the dangers of poor data quality in these systems.

AI systems experience significant negative impacts from poor data quality which affects their fairness and reliability along with how accurate their results are. AI models trained with incomplete or biased data generate flawed predictions that perpetuate existing disparities. The flawed decision-making process from AI systems can lead to decreased trustworthiness and the emergence of ethical and legal risks. AI models trained on misleading or unfiltered data may generate content that is harmful and offensive or inaccurate which can spread misinformation and intensify bias. The deployment of AI systems becomes riskier and the AI benefits decrease when poor data quality generates errors that result in inefficiencies, resource mis-allocation and unintended harm.

3. KEY GUARDRAILS FOR MAINTAINING HIGH-QUALITY DATA IN AL APPLICATIONS

Organizations should implement strong data quality controls to enable AI models to function with accuracy and reliability. Effective measures preserve data integrity and accuracy while keeping it usable which supports fair and consistent AI decision-making that can be acted upon. The following section outlines essential data quality dimensions together with their importance for AI-based applications.

3.1 Precision – Ensuring Data Reflects Reality

AI models require precise data because inaccuracies in data representation can create flawed decisions and incorrect predictions. AI systems require data that accurately portrays real-world entities alongside their corresponding events and behaviors to ensure proper functioning. AI-driven insights face inconsistencies when human entry errors mix with outdated records and unreliable data sources to produce imprecise data. Organizations need to establish rigorous data validation processes together with automated error detection tools and ongoing data monitoring to maintain data precision.

Example: Manufacturing equipment sensor data inaccuracies prevent AI systems from identifying early signs of failure which leads to unexpected breakdowns and expensive operational stoppages.

3.2 Wholeness – Ensuring Comprehensive Datasets

AI models need fully comprehensive datasets to effectively capture all pertinent information for training and inference. AI systems fail to detect important patterns when faced with incomplete data or missing values which results in poor performance and unreliable insights. Organizations must actively detect missing data points and use suitable imputation methods while establishing systematic data collection protocols to achieve complete coverage of essential attributes.

Example: Healthcare AI applications require complete patient records to deliver correct treatment recommendations because incomplete data can endanger patient safety and care quality.

3.3 Uniformity – Harmonizing Data Across Sources

Maintaining consistent data across various systems requires strict data uniformity. Inconsistent data elements like conflicting records or duplicated entries create unreliable AI outputs and cause decision-making confusion. Organizations need to deploy data standardization guidelines together with data reconciliation methods and centralized governance approaches to address these problems.

Example: An AI-driven fraud detection system risks failing to identify suspicious activities accurately when customer data in a financial institution's CRM does not match the transaction database records which raises security concerns.

3.4 Freshness – Keeping Data Up to Date

Real-time decision-making AI applications need current data to operate properly. When data becomes stale or outdated it leads to inaccurate predictions and poor decision-making while also decreasing operational efficiency. To maintain accurate AI model performance organizations, need to establish periodic data refresh schedules along with real-time data collection systems and automatic pipeline maintenance processes.

Example: Utilizing old financial data in stock market prediction models results in flawed investment advice and lost business prospects.

3.5 Validity – Adhering to Data Format and Constraints

For AI applications to function effectively data must adhere to established formats and constraints. Model training faces disruption and prediction accuracy declines when data contains incorrect types, out-of-range values, or formatting issues. Organizations need to adopt validation rules alongside schema enforcement and data cleansing techniques to keep data valid.

Example: An e-commerce AI recommendation system will struggle to produce price-based recommendations when product prices are entered as text rather than numeric values.

3.6 Integrity – Maintaining Data Relationships and Structure

Data integrity maintains consistent relationships between various data entities. AI-driven models produce inaccurate results when datasets contain duplicated records or misaligned data entries. Creating a strong data model that includes clear relationships between data elements along with primary and foreign keys and integrity constraints through referential checks ensures that AI datasets remain both structured and dependable.

Example: The efficiency of AI-powered customer support systems decreases when incorrect linking between purchase history and customer profiles results in irrelevant AI responses.

4. IMPLEMENTING DATA QUALITY GUARDRAILS FOR AI

Data Quality Guardrails for AI require the establishment of organizational processes and technical solutions that maintain the integrity of data throughout model development and deployment. Guardrails detect incomplete data and inconsistent schema as well as bias and drift to prevent these problems from damaging AI insights. This section provides a detailed sequential method for setting up effective Data Quality Guardrails:

4.1 Establish Clear Data Quality Standards & Metrics

Identify essential data integrity metrics that fulfill technical standards and adhere to industry best practices. Connect your metrics with your organization's strategic goals by associating them with concrete AI applications and results for fraud reduction and customer retention enhancement to ensure data quality improvements generate measurable business advantages.

4.2 Implement a Robust Data Governance Framework

Create a structured governance framework that assigns specific roles and responsibilities to data owners, data stewards and governance committees to manage data quality policies. Develop standardized policies and procedural frameworks for data ingestion and storage activities while setting clear usage guidelines that include privacy, security and compliance measures. Organizations should employ a centralized platform with tools like Collibra or Alation to track and document key data asset lineage and definitions to ensure traceability and organization-wide consistency.

4.3 Integrate Automated Data Profiling & Cleaning

Data profiling tools including Talend Data Quality, Informatica Data Quality, and Trifacta help you spot anomalies, find missing values and detect duplicate records. Establish automated workflows for data cleansing and standardization to handle invalid addresses and incorrect phone formats either in real time or through batch processing. Use open-source libraries like Great Expectations and Deequ to create and enforce data validation rules that maintain consistency across column ranges and data types as well as other essential quality parameters.

4.4 Use Data Observability & Monitoring Platforms

Use platforms such as Monte Carlo, Soda.io, or Bigeye to keep track of data pipelines and detect missing values along with schema drift and latency issues. Set up automated notifications to detect anomalies or threshold violations which enable teams to address issues promptly. Set up Service-Level Agreements (SLAs) for essential data streams alongside developing an incident response strategy to address data quality events quickly.

4.5 Well-designed data model & data logging

To prevent data integrity issues organizations, need to implement well-structured data models that use referential integrity and unique constraints to support consistent data relationships. The schema or application layer's embedded normalization and business rules ensure standardized data capture and usage. A strong ETL pipeline extracts data from sources then transforms it according to integrity standards before loading it into its final destination while addressing anomalies, duplicates, and formatting throughout each process step. The comprehensive methodology generates precise and dependable data which leads to better results in artificial intelligence and analytical processes.

4.6 Incorporate Bias Detection & Expandability

Continuously evaluate fairness metrics throughout model inputs and outputs to identify systematic biases that manifest as unequal false positive and negative rates across diverse groups. Utilize AI fairness tools such as IBM AI Fairness 360, Microsoft Fairlearn, and Google What-If Tool to find and reduce biases in your data and models. Use explainable AI methods like SHAP and LIME to assess model decisions and confirm that data quality problems do not subtly alter results.

4.7 Implement Version Control & Data Lineage

Create versioned datasets with separate data snapshots for every training session to maintain reproducibility and provide transparent traceability of model outcomes. Utilize data lineage tracking tools such as Apache Atlas and DataHub to visualize data movement throughout the pipeline which includes stages of ingestion and transformation before reaching consumption. Build automated documentation systems for creating lineage diagrams to enable teams to swiftly locate error sources and observe data drift.

4.8 Establish Continuous Improvement & Feedback Loops

Regular reviews by data experts or domain specialists should include examining flagged issues and validating anomalies to improve automated rules. Use iterative development to continuously improve cleansing and validation processes in response to changes in data sources. Guardrails require periodic performance evaluations to determine their success in avoiding data quality issues while identifying areas for expanded protection.

4.9 Plan for Scalability & Compliance

Develop a scalable architecture that functions in cloud-based, on-premises, or hybrid environments to manage growing data volumes and complexity. Ensure your compliance activities meet relevant privacy laws such as GDPR and CCPA along with industry-specific regulations like HIPAA for healthcare and PCI for finance through the integration of appropriate guardrails. Establish audit and reporting functions through consistent log maintenance and dashboard creation which delivers transparency of compliance status and data quality advancements to auditors and stakeholders.

The implementation of Data Quality Guardrails in AI requires an extensive combination of processes along with policies and tools that uphold reliable and precise data while ensuring ethical standards throughout model development and deployment. Defining precise data quality standards forms the foundation of data quality management with the help of strong governance frameworks together with automated profiling and active monitoring. Data integrity remains protected by well-designed data models and efficient ETL pipelines while bias detection and explainability promote fairness in AI results. Through version control systems combined with continuous improvement protocols and strict compliance standards the organization maintains data quality at large volumes which strengthens both trust and transparency in AI-based decision-making.

Organizations can establish effective Data Quality Guardrails for AI through the combination of organizational best practices such as data governance and clear roles with technical safeguards including automated validation and observability. A comprehensive methodology enables data accuracy, consistency, and reliability throughout the entire AI lifecycle stages including ingestion, processing, model training, deployment and continuous monitoring.4.1 Establish Clear Data Quality Standards & Metrics.

5. INDUSTRY SOLUTIONS FOR MANAGING DATA QUALITY IN AI APPLICATIONS

Enterprise data quality management solutions enable organizations to maintain high-quality and reliable data that remains free from bias for AI application needs. The solutions cover data profiling, validation, governance, observability, and monitoring aspects.

5.1 Enterprise Data Quality AND Cleansing Solutions

The platforms conduct profiling and cleaning of data and standardize it before validating its quality for AI training purposes.

- IBM InfoSphere QualityStage – Data cleansing, deduplication, and standardization for enterprise data.
- Talend Data Quality – AI-powered profiling, validation, and automatic data cleansing.
- Informatica Data Quality – Real-time monitoring, anomaly detection, and governance.
- Ataccama ONE – AI-driven data governance and quality control.
- Trifacta (by Alteryx) – Self-service data preparation and anomaly detection.

5.2 Data Observability & Monitoring Platforms

These tools continuously track data pipelines to detect schema drift, missing values, and performance issues.

- Monte Carlo – Monitors data reliability and pipeline health.
- Great Expectations – Open-source framework for automated data validation.
- Bigeye – ML-powered data quality monitoring.
- Soda.io – Observability and rule-based monitoring for AI data pipelines.
- Datafold – Data validation and reconciliation for AI workflows.

5.3 Master Data Management (MDM) & Data Governance Solutions

MDM platforms help enforce data integrity, lineage tracking, and compliance for AI applications.

- Collibra Data Governance – Enterprise-wide data governance and lineage tracking.
- Alation – Metadata-driven data cataloging and quality tracking.
- SAP Master Data Governance – Ensures high-quality, unified datasets for AI/ML models.
- Profisee MDM – AI-powered master data management platform.

5.4 Bias Detection & Explainability Frameworks

These tools ensure that AI models remain fair, transparent, and explainable.

- IBM AI Fairness 360 – AI bias detection and mitigation toolkit.
- Microsoft Fairlearn – Fairness-aware model auditing.
- Google What-If Tool – Interactive fairness analysis for AI models.
- SHAP & LIME – Model explainability techniques to interpret AI predictions.

5.5 Version Control & Data Lineage Tracking

These solutions help manage data versions, transformations, and lineage to maintain reproducibility in AI workflows.

- Apache Atlas – Open-source metadata and lineage tracking for big data pipelines.
- DataHub – Data lineage and governance platform for AI applications.
- LakeFS – Version control system for large-scale AI data lakes.
- DVC (Data Version Control) – Git-like versioning system for ML datasets.

5.6 ETL & Data Preparation Solutions for AI

Efficient Extract, Transform, Load (ETL) pipelines ensure AI models receive structured, high-quality data.

- Fivetran – Automated ETL pipelines with real-time syncing.
- dbt (Data Build Tool) – Analytics engineering for AI data transformation.
- Matillion – ETL solutions optimized for AI/ML workloads.
- Google Dataflow – Real-time data transformation pipelines for AI models

Selecting the appropriate data quality management approach for AI applications requires consideration of multiple essential factors. Organizations that manage extensive AI models with large and complex datasets need effective ETL and Master Data Management solutions to maintain data quality and structured processing. Real-time monitoring and observability platforms enhance AI-driven automation and decision-making which makes the choice between real-time and batch processing crucial. Industries bound by stringent compliance regulations in fields like healthcare and finance need to use data governance solutions such as Collibra and SAP MDM to meet regulatory requirements. Organizations aiming to address bias detection and fairness should utilize AI fairness frameworks like IBM AI Fairness 360 and Fairlearn because these tools help lessen model bias risks and promote ethical AI operations.

6. IMPACT OF ENSURING DATA QUALITY GUARDRAILS IN AI APPLICATIONS

AI applications that implement data quality guardrails achieve better accuracy while enhancing fairness, reliability, and scalability of AI solutions. Organizations can develop AI systems that deliver trustworthy performance and regulatory compliance through effective data governance combined with validation and monitoring procedures. The presence of poor data quality produces biased results and operational issues while diminishing trust and hence necessitates strong AI data integrity measures.

Data quality guardrails produce significant benefits by enhancing model accuracy and performance. AI models achieve effective training through high-quality data which minimizes errors while improving predictive capabilities. Higher data quality enables dependable decisions in areas including healthcare diagnostics and fraud detection as well as personalized recommendations. A medical AI system that receives accurate patient data training performs more precise disease detection and reduces misdiagnosis risks.

The reduction of bias and promotion of fair AI decisions represent another major impact. Discriminatory patterns emerge from AI models trained using biased or incomplete datasets which create unfair results in hiring practices as well as lending and law enforcement operations. Organizations can establish fair and ethically sound AI systems through the implementation of bias detection frameworks alongside the use of diverse and representative datasets. AI-driven recruitment platforms can objectively assess applicants while preventing discrimination based on gender or race which leads to fair hiring practices.

Organizations can depend on AI-generated insights because data quality guardrails strengthen both data integrity and trustworthiness. AI systems produce trustworthy outcomes when their outputs match real-world situations backed by verifiable information resulting in increased user confidence and reduced misinformation. Accurate and real-time transaction data processing in AI models is essential for preventing false alerts and detecting fraudulent activities during financial fraud detection.

Organizations are prompted to adopt data quality measures because they need to comply with regulations such as GDPR, HIPAA, and CCPA. AI systems within highly regulated industries need to follow rigorous data privacy and security guidelines. Through data governance frameworks businesses protect themselves against legal penalties and reputation harm that result from violations in AI compliance. Healthcare AI requires strong patient privacy safeguards to provide top-tier diagnostic information.

High-quality data operationally reduces risks and minimizes AI failure-associated costs. Inaccurate AI predictions caused by low-quality data lead to necessary manual fixes and additional costs that generate inefficiencies. Organizations that take active steps to keep their data clean reduce rework requirements which leads to increased productivity and better business results. The implementation of AI technology in supply chain management combined with validated and real-time inventory information helps avoid stock calculation mistakes which leads to decreased financial losses.

AI solutions for real-time applications including predictive maintenance and automated fraud detection depend on accurate and clean data to function properly. AI systems generate accurate predictions through the use of continuously refreshed data inputs. AI models for stock market trend analysis that use outdated financial data can lead to investment mistakes and financial losses. AI applications achieve optimal responsiveness and effectiveness through ongoing data observability and integrity maintenance.

Trustworthiness in AI systems emerges as a significant advantage when high data quality standards are sustained. Organizations and consumers will adopt AI systems more readily if they operate on trustworthy and impartial data that maintains transparency. Customer service AI systems require chatbots and virtual assistants to deliver precise and context-aware responses. When chatbots receive training from well-selected high-quality data collections they boost customer experiences and drive higher engagement and satisfaction levels.

Establishing data quality guardrails provides significant scalability benefits. When AI models use structured and high-quality datasets, they achieve efficient scalability across various sectors including finance and healthcare as well as retail and manufacturing. Organizations that use standardized data quality practices gain the ability to modify AI solutions for different applications which helps speed up innovation and boost operational performance. An e-commerce AI recommendation system that accesses structured customer data delivers precise product recommendations which increase both sales and customer loyalty.

CONCLUSIONS

AI applications rely on their training data quality to function effectively and reliably. AI-driven innovations face limitations because poor data quality leads to biased models, inaccurate predictions and raises ethical issues. Organizations need to establish strong data quality controls through automated validation processes, ongoing monitoring systems, governance frameworks, and fairness evaluations to overcome these challenges.

Businesses can maintain data integrity and compliance while reducing AI decision-making risks through industrial solutions like IBM InfoSphere, Talend, Monte Carlo, and AI

fairness frameworks. Implementing a well-designed ETL pipeline together with real-time data monitoring and bias detection capabilities improves the trustworthiness, precision, and equity of AI systems.

Organizations must maintain high-quality scalable data practices that adhere to ethical standards to build reliable and unbiased AI models that support meaningful responsible innovation as AI technology progresses.

REFERENCES

1. Kelleher, J. D. (2019). Deep Learning: The AI Revolution and Its Data Challenges. MIT Press.
2. O'Neil, C. (2016). Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy. Crown Publishing Group.
3. Sicular, S. (2018). Managing AI Data Pipelines: The Role of Data Quality and Observability. McGraw-Hill.
4. Cappiello, C., & Francalanci, C. (2011). Data Quality in Information Systems and Decision Making. Springer.
5. Provost, F., & Fawcett, T. (2013). Data Science for Business: What You Need to Know about Data Mining and Data-Analytic Thinking. O'Reilly Media.
6. Danks, D., & London, A. J. (2020). Algorithmic Bias in AI: Ethical and Policy Considerations. Cambridge University Press.
7. Scholkopf, B., & Smola, A. J. (2002). Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond. MIT Press
8. IBM InfoSphere QualityStage – IBM. (n.d.). Enterprise Data Quality and Governance Solutions. Retrieved from https://www.ibm.com/products/infosphere-qualitystage
9. Towards Data Science. (2022). Data Quality for AI: The Missing Piece of the Puzzle.
10. Talend Data Quality – Talend. (n.d.). Automated Data Profiling and Validation Tools. Retrieved from https://www.talend.com/products/data-quality
11. Harvard Business Review (HBR). (2021). Why AI Models Fail: The Role of Data Quality in Machine Learning. Highlights how organizations can improve AI performance by maintaining high-quality datasets.
12. Informatica Data Quality – Informatica. (n.d.). AI-Powered Data Cleansing and Monitoring. Retrieved from https://www.informatica.com/products/data-quality.html
13. Great Expectations – Great Expectations. (n.d.). Open-Source Data Validation Framework for AI Applications. Retrieved from https://greatexpectations.io/
14. Bigeye Data Observability – Bigeye. (n.d.). Automated Data Quality Monitoring for AI Pipelines. Retrieved from https://www.bigeye.com/
15. Collibra Data Governance – Collibra. (n.d.). Enterprise Metadata and Lineage Tracking. Retrieved from https://www.collibra.com/
16. Google Cloud Blog. (2022). Improving AI Data Quality with Machine Learning Operations (MLOps). Google Cloud. Explores best practices for managing data pipelines and ensuring high-quality AI data.
17. Alation Data Cataloging – Alation. (n.d.). Metadata Management and Data Discovery for AI Applications. Retrieved from https://www.alation.com/
18. Gebru, T., et al. (2018). Datasheets for Datasets. arXiv preprint arXiv:1803.09010. Proposes a standardized approach to document datasets for AI applications, ensuring transparency and better data quality.
19. Hynes, N., Dao, D., Yan, J., Zhang, D., & Song, D. (2018). A Demonstration of Data Lint: Detecting Anomalies in Dataset Distribution for Machine Learning. Proceedings of the VLDB Endowment, 11(12), 2082-2085
20. AWS Machine Learning Blog. (2021). How to Use AWS Data Wrangler to Ensure Data Quality for AI Workloads. Demonstrates techniques to clean and validate data before AI model training.
21. Barocas, S., Hardt, M., & Narayanan, A. (2019). Fairness and Machine Learning: Limitations and Opportunities. arXiv preprint arXiv:1908.09635. Examines how data quality influences fairness in AI models and provides strategies for mitigating bias
22. Cheng, X., Li, Y., & Jin, H. (2020). Data Quality for Machine Learning: State of the Art and Research Directions. IEEE Access, 8, 75427-75441. Reviews data quality dimensions and their impact on AI/ML models.
23. IBM Data & AI Blog. (2023). AI Data Quality: Why It Matters and How to Improve It.