Journal of Computer Science and Technology Studies
ISSN: 2709-104X
DOI: 10.32996/jcsts
Journal Homepage: www.al-kindipublisher.com/index.php/jcsts

JCSTS
AL-KINDI CENTER FOR RESEARCH AND DEVELOPMENT

| RESEARCH ARTICLE

Breaking Down AI-Powered Case Categorization in Customer Support

Vaibhav Fanindra Mahajan
UNIVERSITY AT BUFFALO, USA
Corresponding Author: Vaibhav Fanindra Mahajan, **E-mail**: vaibhavfmahajan@gmail.com

| ABSTRACT

This article explores the transformative impact of AI-powered case categorization in modern customer support environments. Traditional manual categorization processes create a significant administrative burden for support agents, reducing their capacity for substantive problem resolution while introducing inconsistencies that undermine service quality and analytics. AI-powered categorization systems address these limitations through sophisticated machine learning models, natural language processing capabilities, and continuous learning mechanisms that improve over time. The implementation of these systems in platforms like Salesforce's Einstein Case Classification demonstrates how careful attention to evaluation metrics, threshold configuration, and integration with workflow systems can maximize operational benefits. Beyond efficiency gains, AI categorization delivers improved consistency, enhanced analytical capabilities, optimized resource allocation, and significant return on investment. The article examines both current implementations and emerging directions, including multimodal analysis, personalized categorization, predictive support modeling, generative response capabilities, and causal analysis that promise to further revolutionize customer support operations.

| KEYWORDS

Customer Support Automation, Machine Learning Classification, Natural Language Processing, Administrative Burden Reduction, Predictive Analytics

| ARTICLE INFORMATION

ACCEPTED: 12 April 2025 **PUBLISHED:** 29 April 2025 **DOI:** 10.32996/jcsts.2025.7.3.5

Introduction

In today's fast-paced customer support environment, AI-powered case categorization has emerged as a transformative technology that is reshaping how businesses handle customer inquiries. This technical deep dive explores the mechanics, implementation, and benefits of AI-driven categorization systems in modern support ecosystems.

The Evolution from Manual to Automated Categorization

Traditional customer support systems relied heavily on manual tagging and classification of incoming tickets—a process fraught with inconsistencies, human error, and significant time investment. Support agents would spend valuable time determining the appropriate category for each case before actual problem-solving could begin. Administrative burden theory provides a framework for understanding how these manual classification processes create cognitive loads for support agents, diminishing their capacity for substantive case resolution. Research published in the Journal of Public Administration Research and Theory demonstrates that administrative tasks like ticket categorization represent a significant "time tax" on support personnel, creating opportunity costs that directly impact service quality. The research highlights that customer-facing roles often suffer from what researchers term "administrative exclusion," wherein the complexity of classification systems leads to inconsistent application of categories, which subsequently affects downstream analytics and service delivery. The cognitive costs of continually switching between classification tasks and problem-solving create friction in the support workflow that extends average resolution times.

Copyright: © 2025 the Author(s). This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC-BY) 4.0 license (https://creativecommons.org/licenses/by/4.0/). Published by Al-Kindi Centre for Research and Development, London, United Kingdom.

Aspect	Manual	AI-Powered
Process	Agent manually assigns categories	Algorithms automatically assign categories
Time Investment	High	Low
Consistency	Low - varies by agent	High - consistent rules application
Accuracy	Variable - subject to human error	Consistent - improves over time
Impact on Agents	Creates administrative burden	Reduces administrative tasks
Data Quality	Inconsistent for analytics	Uniform foundation for insights

Table 1: Manual vs. AI-Powered Categorization [1]

AI-powered categorization represents a paradigm shift in this approach by leveraging sophisticated algorithms that can process natural language in real-time, identify customer intent from unstructured text, automatically assign appropriate tags, categories, and priority levels, and route cases to the most suitable agent or department. Implementation of AI-powered categorization systems demonstrates notable efficiency gains across support organizations. The technology effectively mitigates the administrative exclusion phenomenon by standardizing the application of classification schemas, thereby improving both the frontline experience for agents and the downstream analytical capabilities for organizations.

Core Technologies Driving AI Case Categorization
Machine Learning Models

The foundation of effective case categorization lies in properly trained machine learning models. These systems typically employ supervised learning algorithms trained on historically categorized support tickets, text classification models that recognize patterns in customer communications, and transfer learning approaches that adapt pre-trained language models to specific business contexts.
Recent longitudinal research on AI implementation trajectories reveals important insights about model performance over time. Studies examining progression patterns demonstrate that domain adaptation techniques substantially improve categorization outcomes when applied to customer support contexts. Temporal analysis of model performance indicates that there are critical thresholds of training data required before supervised learning approaches begin to outperform rule-based systems consistently across support environments. The research suggests that sequence modeling approaches that incorporate temporal features of customer interactions provide superior predictive power compared to static classification models, particularly for complex product ecosystems where customer issues evolve over product lifecycles.

Technology	Key Components	Primary Benefits
Machine Learning	Supervised learning, Transfer learning	Pattern recognition, Adaptation to business context
NLP	Entity recognition, Contextual understanding	Intent identification, Technical language comprehension
Learning Mechanisms	Feedback loops, Model retraining	Ongoing accuracy improvement, Adaptation to new issues

Table 2: Core Technologies Driving AI Categorization [3]

Natural Language Processing (NLP)

NLP capabilities form the cornerstone of understanding customer intent by breaking down text into meaningful tokens and entities, identifying key phrases that signal specific issues, contextualizing ambiguous language, and recognizing technical terminology specific to products or services.

Research published in Expert Systems with Applications has established that contextual embedding models demonstrate superior performance in domain-specific language understanding compared to traditional vector space models. The cognitive science behind these improvements relates to how transformer architectures process relationships between entities in ways that parallel human language comprehension. Studies demonstrate that attention mechanisms are particularly effective for customer support applications because they can isolate relevant context from lengthy customer descriptions that often contain tangential information. Semantic similarity measures derived from these models have proven especially valuable for identifying novel support issues that don't fit existing categorization schemas, enabling support systems to flag potential new product issues for engineering teams.

Continuous Learning Mechanisms

What makes modern AI categorization systems particularly effective is their ability to improve over time through feedback loops incorporating agent corrections, performance monitoring and model retraining, and adaptation to evolving product features and customer concerns.

Enterprise implementations of large language model architectures require specific adaptation strategies to maintain performance over time. Research examining enterprise data governance models highlights that continuous learning systems must balance model plasticity with stability to avoid catastrophic forgetting of existing category definitions while adapting to emerging support issues. Studies of enterprise LLM implementations demonstrate that human-in-the-loop verification systems produce superior long-term accuracy compared to fully automated retraining pipelines. The research indicates that selective retraining strategies that focus computational resources on boundary cases and misclassifications provide more efficient improvement trajectories than comprehensive retraining approaches that consume significant computational resources.

Real-World Implementation Impact

The business impact of AI-powered case categorization extends beyond technical metrics to transformative effects on service delivery models. Longitudinal studies of enterprise AI deployments reveal that categorization improvements create compound effects throughout the support ecosystem. By reducing the cognitive load associated with case classification, agents experience lower burnout rates and improved job satisfaction scores, which research correlates with higher customer satisfaction metrics. The downstream effects on knowledge management become increasingly significant as accurate categorization enables more precise identification of knowledge gaps and training needs within support organizations.

Furthermore, comprehensive analyses of enterprise support operations validate that well-implemented AI categorization systems deliver substantial returns on investment over multi-year horizons. The efficiency gains accumulate through multiple mechanisms: direct time savings in classification tasks, improved routing accuracy reducing transfer rates, enhanced knowledge discovery through accurate categorization, and improved capacity planning enabled by more reliable support analytics. Research in the Journal of Public Administration Research and Theory suggests that removing administrative burden through automation creates a virtuous cycle that allows organizations to focus more resources on substantive problem resolution rather than classification overhead.

Implementation in Salesforce Environments

Salesforce's Einstein Case Classification exemplifies how AI categorization can be integrated into existing customer relationship management systems. The implementation of classification systems requires careful consideration of evaluation metrics to ensure optimal performance. Research published on classification model evaluation metrics demonstrates that while accuracy is commonly reported as a primary metric, it often provides an incomplete picture of model performance, particularly in contexts with imbalanced class distributions such as customer support environments. The research emphasizes that precision, recall, and F1 scores provide more nuanced insights into classification effectiveness, especially for minority case types that may represent critical customer issues despite their relative infrequency. Studies suggest implementing comprehensive metric frameworks that include confusion matrices to visualize classification errors and ROC (Receiver Operating Characteristic) curves to evaluate performance across various threshold settings, enabling organizations to optimize their implementation based on specific business priorities and risk tolerances.

Mode	Description	Key Advantages
Recommendation	Suggests categories for agent confirmation	Preserves human judgment, Builds agent trust
Auto-Apply	Automatically assigns categories when confident	Maximum efficiency, Faster routing
Hybrid	Auto-categorization for high-confidence cases with human oversight for edge cases	Balances efficiency and accuracy, Adaptable to case complexity

Table 3: Salesforce Einstein Implementation Modes [4]

Historical data mining forms the foundation of effective implementation, with the quality of training data directly influencing model performance. Research on sustainable systems implementation highlights the importance of data governance frameworks in establishing reliable training datasets. The research emphasizes that historical case data often suffers from inconsistent labeling practices, which can perpetuate biases and errors if not addressed during the preprocessing phase. Studies recommend implementing data quality assessment protocols that evaluate completeness, consistency, and accuracy of historical classifications before model training begins. This approach aligns with sustainable implementation practices that prioritize long-term system viability over rapid deployment timelines.

Configuration of prediction confidence thresholds represents a critical implementation decision point that directly influences both automation rates and error profiles. Research on classification thresholds demonstrates that optimal threshold settings vary significantly based on case type, with routine inquiries tolerating lower confidence requirements compared to complex or high-risk issues. The research emphasizes the value of implementing dynamic thresholds that adapt based on case characteristics rather than applying uniform confidence requirements across all categories, which can lead to suboptimal automation rates for well-understood case types while still allowing excessive errors for complex scenarios.

Integration with workflow rules and escalation paths amplifies the impact of accurate categorization by ensuring that classification decisions trigger appropriate downstream actions. Research on sustainable service systems emphasizes that classification represents only the initial stage in an integrated service delivery framework, with its value largely determined by how effectively these classifications inform subsequent workflow decisions. Studies demonstrate that organizations implementing end-to-end integration between classification systems and downstream workflows realize substantially greater benefits compared to those using categorization in isolation, particularly regarding time-to-resolution and customer satisfaction metrics.

Dashboard monitoring represents an essential feedback mechanism for continuous improvement, with sustainable implementation models emphasizing the importance of establishing robust monitoring frameworks. Research on classification evaluation metrics highlights the importance of tracking performance across multiple dimensions, including category-specific accuracy, temporal stability of predictions, and alignment between model confidence and actual performance. Studies recommend implementing comprehensive dashboards that enable both aggregate performance monitoring and detailed analysis of specific error patterns to guide targeted improvement initiatives.

The system can operate in two primary modes that represent different balances between automation and control. Recommendation mode suggests categories for agent confirmation, preserving human judgment while reducing cognitive load through decision support. Auto-apply mode directly assigns categories when confidence exceeds predefined thresholds, increasing efficiency but introducing potential automation risks. Research on sustainable service systems suggests that the optimal approach often involves a hybrid implementation that applies auto-categorization for high-confidence predictions while maintaining human oversight for edge cases and evolving issue types, with thresholds calibrated based on ongoing performance monitoring.

Technical Benefits and Operational Improvements

The technical advantages of AI-powered case categorization extend beyond simple automation to fundamental improvements in service delivery capabilities and operational efficiency. Research on classification model evaluation demonstrates that consistent application of well-defined classification rules represents a significant advancement over human-only processes, which typically exhibit substantial variation between individuals and across time periods. Studies highlight that AI systems apply classification criteria with higher consistency than human agents, resulting in standardized categorization regardless of agent experience level or workload pressures.

By applying the same criteria across all incoming cases, AI categorization ensures reduced miscategorization rates compared to manual methods. Research on classification metrics emphasizes that reduction in false positives and false negatives represents a critical performance indicator for support operations, with miscategorization leading to inefficient resource allocation and extended resolution timelines. Studies examining classification performance across service operations highlight that accuracy improvements are particularly significant for complex product ecosystems with extensive taxonomy requirements that exceed human cognitive capacity for consistent application.

Uniform data quality for downstream analytics represents a fundamental improvement that extends beyond immediate operational benefits. Research on sustainable systems implementation demonstrates that reliable classification creates a foundation for increasingly sophisticated analytical capabilities by ensuring that historical data accurately reflects actual case characteristics. Studies emphasize that inconsistent classification represents one of the primary barriers to effective analytics implementation, with data quality issues propagating through analytical processes and undermining confidence in resulting insights.

Resource allocation becomes more efficient through intelligent routing based on agent expertise and availability. Research on predictive analytics highlights that routing optimization represents a natural extension of accurate classification, with case characteristics informing assignment decisions based on agent skill profiles and current workload. Studies demonstrate that classification-informed routing significantly reduces resolution time and escalation requirements by matching cases to appropriate expertise from the initial contact, particularly for specialized technical issues that benefit from targeted knowledge.

Prioritization based on issue severity and business impact improves service level agreement compliance and customer satisfaction. Research on predictive analytics emphasizes that effective prioritization relies on accurate classification combined with business rules that reflect organizational priorities and customer impact assessments. Studies highlight that classification systems enable more nuanced prioritization frameworks that consider both explicit case characteristics and derived features such as predicted resolution complexity or customer sentiment, resulting in more effective queue management.

Load balancing across support teams improves agent utilization rates and operational efficiency. Research on sustainable service systems demonstrates that classification-informed workload distribution enables more effective capacity planning by providing detailed visibility into case mix and complexity patterns. Studies highlight that balanced workload distribution represents a critical factor in agent satisfaction and retention, with classification systems enabling more equitable assignment practices based on objective case characteristics rather than arbitrary distribution rules.

Integration with Analytics

Well-categorized cases enable sophisticated analytics that identify emerging issue clusters with significantly improved accuracy. Research on predictive analytics trends demonstrates that pattern recognition capabilities depend fundamentally on data quality and consistency, with classification systems providing the structured data necessary for effective cluster analysis. Studies highlight that accurately categorized historical data enables the identification of emerging trends before they become widespread problems, allowing proactive intervention that reduces overall impact and resolution costs.

Prediction of support volume patterns becomes substantially more accurate with properly classified historical data. Research on predictive analytics emphasizes that forecasting models rely on consistent historical patterns to generate reliable projections, with classification consistency representing a critical input quality factor. Studies demonstrate that forecasting precision improves significantly when based on consistently categorized historical data, enabling more accurate capacity planning and resource allocation decisions that reduce both overstaffing costs and service risks.

Product development priorities informed by accurately categorized support data lead to more targeted improvement initiatives. Research on sustainable systems highlights that product enhancement decisions increasingly rely on customer feedback signals, with support data representing a particularly valuable source of improvement opportunities. Studies demonstrate that classification systems enable more effective prioritization by accurately quantifying issue frequency and impact across product components, resulting in development roadmaps that more effectively address customer pain points.

Measuring Effectiveness and ROI

Organizations implementing AI categorization should establish comprehensive measurement frameworks to track system effectiveness and return on investment. Research on classification metrics emphasizes the importance of multidimensional evaluation approaches that consider both technical performance and business outcomes. Studies recommend implementing balanced scorecard approaches that combine operational metrics such as time-to-resolution with customer experience indicators and financial measures to provide holistic performance assessment.

Reduction in average time-to-resolution represents a primary performance indicator that directly impacts both operational efficiency and customer satisfaction. Research on sustainable service systems demonstrates that resolution time improvements typically result from a combination of factors, including more accurate initial routing, improved context availability, and reduced rework requirements due to appropriate categorization. Studies emphasize that time savings typically vary by case complexity, with the most significant improvements observed for issues that previously experienced high misrouting rates.

Decrease in escalation or reassignment rates indicates improved first-level resolution capability and routing accuracy. Research on classification metrics highlights that escalation reduction represents a particularly valuable outcome that reduces both operational costs and customer friction by minimizing handoffs between agents. Studies demonstrate that classification-informed routing significantly reduces the need for subsequent transfers by matching cases to appropriate expertise from the initial contact.

Improvement in first-contact resolution percentages directly enhances customer experience while reducing operational costs. Research on predictive analytics emphasizes that first-contact resolution represents a critical efficiency and satisfaction driver, with each additional contact significantly reducing customer satisfaction while increasing service delivery costs. Studies highlight that classification systems improve first-contact resolution by ensuring cases reach appropriately skilled agents with relevant knowledge resources, particularly when integrated with knowledge recommendation systems that leverage classification data.

Agent time saved on administrative tasks represents a significant productivity enhancement that enables increased focus on value-added customer engagement. Research on sustainable service systems demonstrates that administrative burden reduction produces compound benefits by simultaneously increasing available capacity and improving agent satisfaction through more meaningful work allocation. Studies quantify administrative time savings as a direct productivity benefit while emphasizing the additional value of improved agent experience and reduced turnover.

Return on investment timelines vary based on implementation approach and organizational context. Research on sustainable systems implementation emphasizes that ROI calculation should incorporate both direct cost savings from automation and indirect benefits such as improved customer retention and enhanced analytical capabilities. Studies highlight that comprehensive change management programs significantly accelerate ROI achievement by ensuring effective adoption and utilization of new capabilities.

Future Directions in AI Case Categorization

The evolution of this technology continues with promising developments in several key areas that extend classification capabilities beyond current implementations. Research on classification metrics highlights that multimodal analysis incorporating images and attachments represents a significant advancement for technical support applications where visual evidence provides critical diagnostic information. Studies demonstrate that integrating visual and textual analysis enables more accurate classification of complex technical issues, particularly hardware failures and environmental factors that manifest visually.

Personalized categorization based on customer history and preferences is emerging as a significant enhancement to traditional classification approaches. Research on AI in marketing personalization demonstrates that contextual models incorporating customer-specific factors can substantially improve classification accuracy for returning customers with established interaction patterns. The research emphasizes that personalization extends beyond simple demographic segments to incorporate behavioral patterns, preference history, and relationship characteristics that influence both problem presentation and optimal resolution approaches. Studies highlight that personalized categorization represents a critical evolution that bridges standardized classification with individualized service delivery, particularly for complex products with diverse usage patterns.

Capability	Description	Impact
Multimodal Analysis	Integration of text and image analysis	Better diagnosis of visual/hardware issues
Personalized Categorization	Incorporating customer history into classification	Improved accuracy for returning customers
Predictive Support	Anticipating issues before tickets are created	Proactive resolution, reduced ticket volume
Generative AI	Automated responses based on classification	Faster initial responses for common issues
Causal Analysis	Identifying root causes rather than symptoms	Higher first-time fix rates

Table 4: Future Directions [7, 8]

Predictive support modeling that anticipates issues before formal tickets are created represents a transformative capability emerging from mature classification systems. Research on predictive analytics trends demonstrates that accurate historical classification enables the development of predictive models that identify patterns preceding specific support issues, allowing proactive intervention before customers experience significant impact. Studies highlight that predictive capabilities demonstrate particularly high value in environments with telemetry data that can be correlated with support patterns to identify emerging issues before they affect multiple customers.

Integration with generative AI for automated response generation leverages accurate classification to enable more precise and relevant automated communications. Research on predictive analytics highlights that response generation represents a natural extension of classification capabilities, with category-specific response templates evolving toward increasingly sophisticated generative models that incorporate case-specific context. Studies demonstrate that classification-informed response systems achieve substantially higher accuracy compared to generic approaches, particularly when developed through a phased approach that begins with common inquiries before addressing more complex scenarios.

Advanced systems are increasingly focusing on causal analysis capabilities that move beyond correlation to identify root causes of support issues. Research on classification metrics emphasizes that causal understanding represents the next frontier in support analytics, with classification systems providing the structured data necessary for identifying relationship patterns between symptoms and underlying causes. Studies highlight that causal analysis significantly improves first-time fix rates by addressing fundamental issues rather than symptoms, establishing a foundation for increasingly proactive support models that resolve underlying problems rather than repetitively addressing their manifestations.

Conclusion

AI-powered case categorization represents a fundamental advancement in customer support technology that addresses longstanding operational challenges while creating new strategic capabilities. By automating the labor-intensive and error-prone process of ticket classification, these systems simultaneously reduce administrative burden on agents and improve data consistency for downstream processes. The most successful implementations balance automation with human oversight, applying machine intelligence where it excels while preserving human judgment for complex edge cases.

The true value of AI categorization emerges when viewed as a foundational capability rather than an isolated feature. Organizations realizing the greatest benefits integrate classification systems into comprehensive service delivery frameworks that leverage accurate categorization to inform routing decisions, prioritization schemes, knowledge recommendations, and analytical insights. This integrated approach transforms support operations from reactive problem resolution to proactive issue prevention through increasingly sophisticated analytical capabilities.

As these systems continue to evolve, the integration of multimodal analysis, personalized categorization approaches, predictive modeling, and causal analysis will further extend their capabilities. The progression from basic classification to anticipatory support represents a fundamental shift in service delivery paradigms, enabling organizations to address customer needs before they manifest as formal support requests. This evolution toward increasingly proactive and personalized support, grounded in accurate classification and enhanced by advanced analytics, promises to redefine customer support excellence in the coming years.

Organizations implementing these systems should establish comprehensive measurement frameworks that capture both immediate operational improvements and longer-term strategic benefits. By balancing short-term efficiency gains with sustained investment in system refinement and capability extension, support leaders can position their organizations to deliver exceptional service experiences while continuously improving operational performance. The journey from manual to AI-powered categorization ultimately represents not just a technological upgrade but a fundamental reimagining of how customer support delivers value in the digital age.

Funding: This research received no external funding.
Conflicts of Interest: The authors declare no conflict of interest.
Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers.

References

[1] Carolyn Heinrich, "The Bite of Administrative Burden: A Theoretical and Empirical Investigation," December 2015, Journal of Public Administration Research and Theory 26(3):muv034, DOI:10.1093/jopart/muv034, Available: https://www.researchgate.net/publication/286509699_The_Bite_of_Administrative_Burden_A_Theoretical_and_Empirical_Investigation

[2] Cheonsu Jeong, "A Study on the Implementation of Generative AI Services Using an Enterprise Data-Based LLM Application Architecture," September 2023, DOI:10.48550/arXiv.2309.01105, Available: https://www.researchgate.net/publication/373685686_A_Study_on_the_Implementation_of_Generative_AI_Services_Using_an_Enterprise_Data-Based_LLM_Application_Architecture

[3] Oluwaseyi Kolawole Oladele, "AI for Longitudinal Analysis: Tracking Disease Progression Over Time." Research Gate, March 2025, Available: https://www.researchgate.net/publication/389909343_AI_for_Longitudinal_Analysis_Tracking_Disease_Progression_Over_Time

[4] Prajval Sudhir, Varun Deshakulkarni Suresh, "Comparative study of various approaches, applications and classifiers for sentiment analysis," Global Transitions Proceedings, Volume 2, Issue 2, November 2021, Pages 205-211, Available: https://www.sciencedirect.com/science/article/pii/S2666285X21000327

[5] Qiang Wu, "Optimization of AI-driven communication systems for green hospitals in sustainable cities," Sustainable Cities and Society, Volume 72, September 2021, 103050, Available: https://www.sciencedirect.com/science/article/abs/pii/S2210670721003346

[6] Sodiq Odetunde Babatunde, et al, "THE ROLE OF AI IN MARKETING PERSONALIZATION: A THEORETICAL EXPLORATION OF CONSUMER ENGAGEMENT STRATEGIES," March 2024, International Journal of Management & Entrepreneurship Research 6(3):936-949,

DOI:10.51594/ijmer.v6i3.964, Available: https://www.researchgate.net/publication/379393342_THE_ROLE_OF_AI_IN_MARKETING_PERSONALIZATION_A_THEORETICAL_EXPLORATION_OF_CONSUMER_ENGAGEMENT_STRATEGIES

[7] Vaibhav Kumar, M. L., "Predictive Analytics: A Review of Trends and Techniques," July 2018, International Journal of Computer Applications 182(1):31-37, DOI:10.5120/ijca2018917434, Available: https://www.researchgate.net/publication/326435728_Predictive_Analytics_A_Review_of_Trends_and_Techniques

[8] Zeljko Vujovic, "Classification Model Evaluation Metrics," July 2021, International Journal of Advanced Computer Science and Applications Volume 12(Issue 6):599-606, DOI:10.14569/IJACSA.2021.0120670, Available: https://www.researchgate.net/publication/352902406_Classification_Model_Evaluation_Metrics