International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 16(2024) pp. 1-9
https://doi.org/10.70917/ijcisim-2024-6008
Article

Trustworthy Driver Compensation Automation: Human-in-the-Loop Reconciliation of Telematics, Trip, and Payroll Data

Prasanna Sundaravaradan

Application Development Lead, Anveta Inc

Abstract: Compensation for commercial drivers can be based on data from various telematics sources, duty status, trip execution, payroll policies and rules, location policies, and manual adjustments. A complete manual reconciliation is labor intensive and subject to quality issues, while a complete autonomous adjustment processing can introduce wage, compliance and explainability risks. This paper discusses a human-in-the-loop architecture for trustworthy driver-compensation automation. Allows the ingestion of source events without modifying the original records; provides a way to align telematics and trip timelines; brings compensation rules to the forefront; computes proposed adjustments; classifies by confidence and exception categories; routes uncertain cases for approved review. All suggested and accepted changes remain in the source evidence, calculation path, change rule version, change by reviewer and downstream payroll status. The architecture is analysed with respect to typical examples with breaks, unmatched trips, missing duty events, overlapping activities, delayed source records, and changes in policy, and an evaluation methodology is suggested based on the consistency of the calculations, the accuracy of the exceptions, the reduction of the records for manual review, the processing time, the auditability and the ability of the system to reject incomplete cases safely. The contribution moves beyond efficiency and incorporates all of the elements of explainability, source provenance, human authority, and employee-impact controls as integral parts of the architecture, not mere add-ons.

Keywords: driver compensation; human-in-the-loop automation; telematics reconciliation; payroll systems; explainable automation; algorithmic management; trust in automation; data provenance; exception routing; audit trail; wage compliance; trip-data alignment

1. Introduction

The compensation of commercial drivers relies more and more on the information that is stored on telematics platforms, electronically stored duty status logs, trip-execution systems, payroll rule engines, jurisdictional pay policies and manually submitted requests for adjustments. Each of these systems is usually managed by a separate vendor, with varying dates and frequency of updates and different conventions for the way a driver, a shift, or a trip segment is identified. If compensation relies on the reconciliation of a number of such sources, the two types of failure modes occur again. Manual reconciliation is a slow and auditing after-the-fact process, inconsistent between the payroll analysts, and slow. Purely automatic reconciliation, where the algorithm determines and publishes the pay adjustments without any human oversight, has wage, compliance and explainability risks, as compensation errors impact a worker's livelihood and are under the regulatory control of most jurisdictions.

The concept of a reliable Human-in-the-Loop system for the reconciliation of telematics, trip and payroll data with a view to the automation of driver compensation is addressed in this paper. The architecture is designed to be conservative: Automation is employed to align the data, apply explicit rules, and generate a proposed adjustment and its supporting evidence, and ultimate decision-making power over uncertain or high impact cases resides with a human, authorized reviewer. The goal is to eliminate human judgment from compensation decisions, not to replace it;


Copyright: © 2024 by the authors

rather, to make compensation decisions more consistent, informed, and less cumbersome with reconciliation details of low value.

The rest of the paper is organized as follows. In Section II, the proposed design is placed in the context of the current literature on trust in automation, algorithm aversion, explainable AI, algorithmic management of gig work, data provenance, and human-in-the-loop machine learning. In Section III, the data landscape is characterized and the type of reconciliation exceptions that call such a human in the loop design [1]. The proposed architecture is provided in Section IV. Section V provides details about score-based confidence and exception routing. The auditability and provenance model is described in Section VI. The concept of human authority and safeguards affecting employees is addressed in Section VII. In Section VIII, an assessment approach suitable for the architecture is discussed. Implications and limitations are covered in Section IX, and conclusions are provided in Section X.

2. Related Concepts and Design Foundations

2.1 Trust and Reliance in Automated Systems

A distinction between overreliance and underreliance is the level of trust that a human operator has in an automation system that corresponds to the actual reliability of that system.Appropriate reliance is trust in the automation system that matches the actual reliability of the system, while overreliance or underreliance does not. Calibrated trust requires information from the system that enables the operator's ability to determine whether a suggestion is likely to be right. A compensation reconciliation system that shows a proposed adjustment with no obvious evidence behind it will either cause it to be accepted on faith or give rise to general mistrust of the process, neither of which is conducive to accurate results or to the interests of workers. The literature lays the foundation for the routing strategy used in Section V, which is confidence-based, where the level of trust an agent can place on the output of the system depends on the system's confidence.

A related issue of interest is algorithm aversion, in which users who notice a minor, visible mistake by an automated system becomes even more skeptical of that system, and may even give up on it entirely, despite the system being more accurate overall. Payroll is a situation in which a mistake will impact on a driver's pay, which, if the mistake is not known or explained, can have a negative impact on the entire system workforce as a whole [2][3]. This calls for default conservative: uncertain cases should not be automatically posted for review, and it will be rare for them to be in error, and when it is, it will be because someone has made an explicit decision, not because of any hidden calculation.

2.2 Explainability and Human-AI Complementarity

Interpretable machine learning has demonstrated that explanations are valuable at the local, instance level in assisting a reviewer to make a decision about whether to believe a specific prediction, and that an explanation with a corresponding confidence signal leads to greater complementary team performance—compared to a human reviewer plus a prediction without an explanation. The results justify two design decisions in the proposed architecture: (a) to include a traceable calculation path instead of a simple numerical output, and (b) to write explanation to help the reviewer in making the decisions, not in trying to design the most complete documentation [4].

2.3 Algorithmic Management in Logistics and Gig Work

One area of scholarship focuses on algorithmic management, the application of software systems to manage, assess, and punish labor, especially within the ride-hailing, delivery, and long-haul trucking settings. This literature confirms that algorithmic control of scheduling, routing and pay can indeed be a contested space where workers' autonomy is undermined, and, in some formulations, where they may be subject to algorithmic domination due to the absence of possibility to complain or appeal decisions that impact their livelihood. The legal analysis of app-based work places algorithmic management as an emerging field for employment relations and human resource management research, and suggests that automated systems for managing pay deserve different types of oversight mechanisms than traditional managerial oversight. Here are the findings that are relevant to an architecture that makes pay adjustments:

No matter how technically correct it is, it is possible for such an architecture to perpetuate or go so far as to mirror the dynamics of domination that this literature warns about because it is not contestable nor does it have a clear path for human recourse.

2.4 Data Provenance and Human-in-the-Loop Machine Learning

For accountabilty, debugging, and even verification, it is essential that the sources which generate a derived value, the transformations that create it, and the versions involved in these transformations are known. Human-in-the-loop ML also classifies how human judgment can be incorporated into automated pipelines, such as verification loops in which a system suggests an output that requires a human to validate before implementation [5]. In Section IV, the architecture is designed around a verification-loop pattern that is strengthened by an explicit provenance record, such that every adjustment (whether automatic or human) is accompanied by a lasting connection to the evidence and logic that led to it.

2.5 Timeline Alignment and Anomaly Detection

To reconcile telematics/trip data, streams of events need to be correlated from separate systems with imperfect clocks and inconsistent identifiers. Map-matching research, which seeks to align noisy positional traces with an underlying reference structure, provides design clues for handling gaps, drift and ambiguous correspondences instead of requiring a clean one-to-one mapping between sources, is closely related. Anomaly detection in event logs from process-aware information systems also helps to identify an event log that has a different sequence, e.g., a duty-status sequence without an event or a sequence with overlapping activities [6]. There is another analogue, derived from the methodology used for collecting statistical data: Here, two sources of data are often partially overlapping and must be merged into a single estimate, with discrepancies between the sources clearly identified and resolved without being averaged over. Last, research on eLDs in long-haul trucking reveals how this technology impacts drivers' work environment and sleep schedules, highlighting that the source data that inform a compensation system is built upon instruments that have their own context of operation, rather than a neutral ground truth.

3. Data Landscape and Reconciliation Challenges

3.1 Source Systems

In driver compensation reconciliation, five types of source data are typically used, all of which can be summarized in Table I: the data of telematics events is generated by onboard devices and can be used to capture duty status changes, vehicle motion and location. Trip-execution records are produced by dispatch/trip management systems and contain information about the planned and completed segments of a trip. Payroll rule definitions include base rates, differential pay, break rules and represent the compensation policy in effect [7]. These jurisdictional and location policy data set the rules for pay for a certain portion of work. Manual adjustment submissions occur when a driver or supervisor has a case that they would like to submit for manual adjustment that is not part of the automated pipeline process.

TABLE I SOURCE DATA CATEGORIES

Source System	Origin	Typical Content
Telematics Events	Onboard vehicle / ELD devices	Duty-status transitions, vehicle motion, location traces
Trip-Execution Records	Dispatch / transportation-management system	Planned and completed trip segments, stop records
Payroll Rule Definitions	Payroll policy engine	Base rates, differentials, break rules, effective dates

Location & Policy Data	Jurisdictional reference tables	Applicable wage-and-hour rules by location and date
Manual Adjustment Submissions	Driver / supervisor request channel	Correction requests outside the automated pipeline

3.2 Recurring Exception Categories

In each reconciliation, there are six types of exceptions that occur and motivate the routing design considered in Section V, which are listed in Table II: unlogged or unmatched trips, in which a trip record is not matched to a telematics trace; missing duty status events, in which a transition is expected but not recorded; overlapping activities, in which two systems report the same activity at the same time but record conflicting duty status; qualifying breaks, in which the duration and placement of a break determine whether it is paid or unpaid under the current policy; delayed source records, in which the source system is reporting events after the end of the reconciliation window; and policy changes, in which the compensation rule changes in the middle of the pay period. Each reflects a different approach to the resolution and a different level of acceptance for automated handling [8].

TABLE II EXCEPTION CATEGORY TAXONOMY

Exception Category	Description	Typical Resolution Path
Unmatched Trip	Trip record has no corresponding telematics trace	Route to reviewer with source excerpts
Missing Duty Event	Expected status transition is absent from log	Hold pending correction or manual entry
Overlapping Activity	Concurrent, inconsistent duty states reported	Flag for expedited review
Qualifying Break	Duration or placement affects paid status	Rule-engine evaluation; edge cases reviewed
Delayed Source Record	Event arrives after reconciliation window	Re-run alignment on receipt; interim hold
Mid-Period Policy Change	Compensation rule changes within a pay period	Apply rule versioning; flag boundary events

4. Proposed Architecture

The proposed architecture is a pipeline with an explicit branch point between the automated and human processing, as shown in Fig. 1. The source events are provided from all five input categories to a non-destructive ingestion layer which does not modify the original record, and all downstream computations refer to it, not overwrite it [9]. Ingested events pass through a timeline alignment stage which continues to use tolerance windows and identifier matching based on the principles of map-matching to create a reconciled timeline along with an explicit list of unmatched or ambiguous time intervals, as shown in Fig. 2.

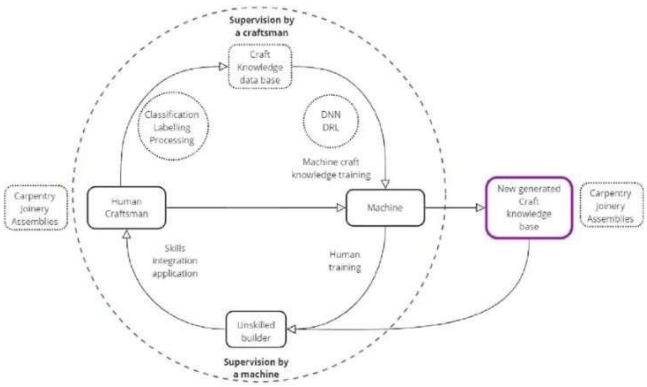


Fig. 1. Human-in-the-loop reconciliation architecture, from source ingestion through confidence-based routing to payroll integration (ResearchGate,2023)

The reconciled timeline is then passed to a compensation rule engine that executes on the explicit rules which are version-controlled, and thus the rule active when the event occurs will determine the calculation, not the rule active when the time comes to process the compensation. The result of the rule engine is fed into an adjustment calculation stage which generates an adjustment calculation path and generates a proposed compensation adjustment. At the completion and consistency of evidence stage, the evidence presented to support the proposed adjustment is assessed for its confidence level and classified as an exception (see Section IV), which then takes it to an automated processing stage and logs it for audit at a later date, or holds it and routes it to an authorized human reviewer before any adjustment is processed in payroll [10]. Both paths lead to the same stage in payroll processing, where the approved adjustment is posted along with the complete audit trail.

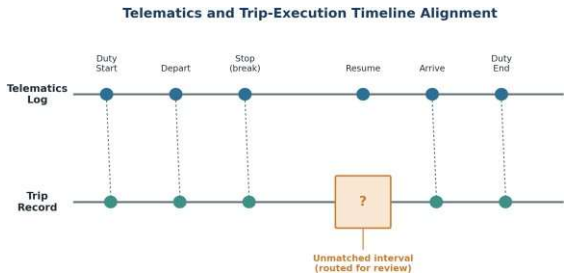


Fig. 2. Alignment of telematics and trip-execution timelines; an interval without a corresponding trip event is flagged and routed rather than estimated.

5. Confidence Scoring and Exception Routing

Three factors are used to determine a composite confidence score for each proposed adjustment: the degree of agreement between independent evaluations of the rules, the completeness of the source evidence available for the interval being adjusted, and the detection or absence of an exception condition defined in Table II. The composite score divides cases into three segments (see Fig. 3). In high-confidence cases, where there is full agreement in the rules and source evidence is complete, they are automatically processed and recorded for audit rather than held for review, as designed, in order to reduce low value manual processing without taking any of the human oversight from consequential cases [11][12]. Medium-confidence cases are identified because only part of the evidence is available, or there is a minor rule conflict, and are flagged for expedited review, where the case is reviewed by a case reviewer who will be provided with a system-generated summary of the case, not the raw source data. Low-confidence cases are retained and sent to an authorized reviewer with all of the evidence.

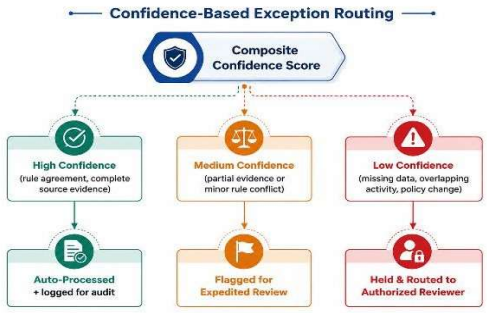


Fig. 3. Confidence-based exception routing from a composite score to auto-processing or authorized human review.

This tiering mirrors the trust-calibration and algorithm-aversion reflections discussed in Section II, in a sense; it is more "conservative" than a "pure" accuracy-maximizing classifier would be, by design, because the "cost" of an error made on a visible signal ("no action" in a compensation context) is asymmetric with the "cost" of a review that is not necessary to make. The level between tiers is not a static value but instead is a versioned, auditable configuration value that an organization can change as they get more confident in the system, and in the quality of the underlying data, over time.

6. Auditability and Provenance Design

All proposed and approved adjustments have a durable log of their origin showing the source evidence referenced from the ingestion layer, the chain of calculations for the adjustment (including any rules logic and intermediate calculations), the version of the compensation rule that was in effect at the time of the underlying event, the action taken by the reviewer on the adjustment, and the downstream status of the payroll adjustment [13]. This is a record that exemplifies the provenance principle that the product of derivation, produced as a value, should be traceable to the inputs and transformations that led to its creation, but here it is extended to incorporate the human decision making as a first-class element of the derivation record, not as an external annotation.

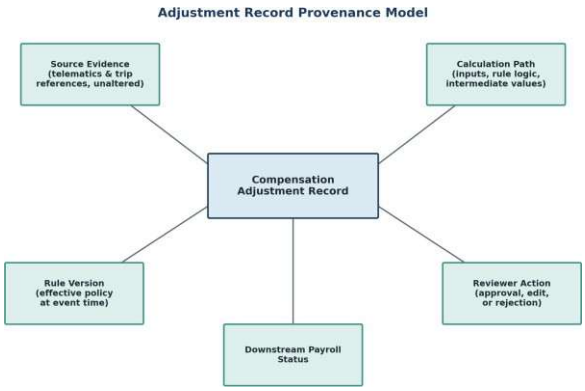


Fig. 4. Provenance model linking a compensation adjustment record to its evidentiary and decision history.

The fields that are captured in the audit trail record are summarized in Table III. The ingestion layer does not alter source records, so an auditor always knows what the underlying telematics and trip data was before an adjustment is made, regardless of whether he or she later receives a correction to the source data.

TABLE III AUDIT TRAIL RECORD FIELDS

Field	Description
Source Evidence Reference	Pointer to unaltered telematics and trip records

Calculation Path	Rule logic and intermediate values used to derive the adjustment
Rule Version	Compensation rule in effect at the time of the underlying event
Confidence Tier	Assigned routing classification (high / medium / low)
Reviewer Action	Approval, edit, or rejection, with reviewer identity and timestamp
Payroll Status	Downstream posting state of the adjustment

7. Human Authority and Employee-Impact Safeguards

The architecture allows for the presence of human review authority as a structural requirement, not as a last resort, as the algorithmic-management literature mentioned in Section II does. The following three safeguards are the result of this treatment. Safe rejection: If the ingest or alignment stage cannot form an adequate supported timeline for an interval, it should not suggest an adjustment, but let manual handling take over completely. Second, contestability: a driver or supervisor should have the ability to cause a re-review of an auto-processed adjustment and Table I contains the manual-adjustment input channel just for that purpose in order to make the manual route out of the automated pipeline. Third, reviewer accountability: reviewer assignments and the timing of all non-automated decisions are included in the provenance record in Section VI, so that human reviewers are not only accountable for their actions, but they have a documented audit trail of their actions [14][15].

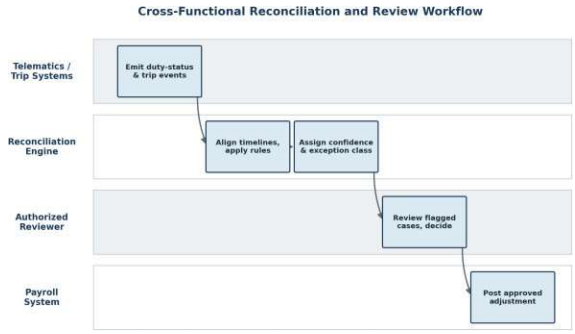


Fig. 5. Cross-functional workflow showing how case responsibility moves between source systems, the reconciliation engine, an authorized reviewer, and payroll.

The cross-functional workflow for the implementation of these safeguards is illustrated by Fig. 5, which shows the interaction between the source systems, the reconciliation engine, an authorized reviewer and the payroll system as the information passes through each step. The explicit handoff points in this workflow are designed to bring to the fore, and thus to make accountable, the separation between algorithmic and human decisions about pay, a concern that otherwise might be left unmarked.

8. Evaluation Methodology

When considering an architecture of this type, one must consider dimensions other than just the accuracy of the computations since the design objective is to have trustworthy automation, not just automation. Table IV summarizes six evaluation dimensions: calculation consistency (how many times doing the calculation with the same source data produces the same proposed adjustment); exception precision (how many cases that are routed or auto-processed are correctly classified as an exception); manual-review reduction (how much less work does a reviewer have to do when processing a case, compared to a fully manual baseline); processing time (how much time does it take to take a case from the point of ingestion to the point of a posted adjustment); auditability (how complete is the provenance record for a sample case that was adjusted); and safe rejection (how many times, when the evidence is too incomplete, does the system correctly decline to propose an adjustment instead of producing an unsupported estimate).

TABLE IV PROPOSED EVALUATION DIMENSIONS

Dimension	What It Measures
Calculation Consistency	Repeatability of proposed adjustments over identical inputs
Exception Precision	Correctness of the confidence-tier routing decision
Manual-Review Reduction	Change in reviewer workload versus a fully manual baseline
Processing Time	Latency from source-event ingestion to a posted adjustment
Auditability	Completeness and reconstructability of the provenance record
Safe Rejection Rate	Frequency of correctly declining to propose an adjustment

A representative test suite for this evaluation should include the exception categories identified in Table II, namely unmatched trips, missing duty events, overlapping activities, qualifying breaks, delayed source records, and mid-period policy changes, so that the architecture's behavior on precisely the cases it was designed to route can be examined directly rather than inferred from aggregate performance alone.

9. Discussion

Here, the architecture is chosen to be deliberately between full automation and full manual reconciliation, and there are trade-offs. An appropriate balance between the risk of errors that might be visible and foster trust erosion and reviewer workload, is probably organization- and jurisdiction-specific, rather than universal, especially when wage and hour regulation limit the ways in which compensation decisions can be made and/or reviewed. The cost of the provenance-centric design is storage and engineering overhead compared to a system that discards the intermediate state of calculation, but is seen here as a legitimate cost of being explainable and being able to be traced back to its origin, in accordance with the idea that explainability and provenance are legitimate architectural requirements for automation which affects pay, rather than reporting features added at the end.

Another restriction is the source data quality. The reliability of the timeline will be as good as the telematics and trip-execution systems that are driving it; and prior studies of electronic logging devices suggest a wider environment of device failures, connectivity issues and driver workarounds [16]. However, the safe-rejection safeguard in Section VII is not enough to cure data-quality problems; it simply means that they don't silently propagate to pay, and they are a design issue of the source-system that is being reconciled and is beyond the scope of the reconciliation architecture.

10. Conclusion and Future Work

The architecture proposed in this paper has been a human-in-the-loop approach to reconciling telematics, trip and payroll data for the purpose of ensuring accurate driver compensation automation. The architecture does not modify the source events, provides telematics and trip event alignment, contains explicit versioned compensation rules and passes along proposed events to automated or human event processing based on a composite confidence score, whilst ensuring a full provenance for all events along the way, irrespective of processing path. The design addresses a concern from the trust-in-automation, algorithm-aversion, and algorithmic-management literatures (Section II): explainability, source provenance, human review authority, employee-impact safeguards are fundamental requirements, not optional extras.

The next steps involve empirical testing of the architecture with the test-case categories and dimensions described in Section VIII, research into how routing thresholds should evolve as more confidence is gained in a particular data source over time, and further integration of the manual-adjustment channel with the contestability safeguard so that driver-driven review requests can be conveyed and tracked in the same way as system-driven exceptions are conveyed and tracked. Another direction proposed by the literature on provenance and reconciliation reviewed here is to extend the provenance model, to support cross-organization audits, where applicable, when drivers travel between carriers in a single pay period.

REFERENCES

1. G. Bansal, T. Wu, J. Zhou, R. Fok, B. Nushi, E. Kamar, M. T. Ribeiro, and D. S. Weld, "Does the whole exceed its parts? The effect of AI explanations on complementary team performance," in Proc. 2021 CHI Conf. Human Factors in Computing Systems, 2021, pp. 1–16, doi: 10.1145/3411764.3445717.
2. H. K. Benecha, D. A. Abreu, R. Jennings, and L. J. Young, "Data reconciliation and estimation in an agricultural survey," Statistical Journal of the IAOS, vol. 39, no. 2, pp. 461–472, 2023, doi: 10.3233/SJI-220087.
3. F. Bezerra and J. Wainer, "Algorithms for anomaly detection of traces in logs of process aware information systems," Information Systems, vol. 38, no. 1, pp. 33–44, 2013, doi: 10.1016/j.is.2012.04.004.
4. A. M. Crizzle, J. Malkin, G. A. Zello, R. Toxopeus, P. Bigelow, and M. Shubair, "Impact of electronic logging devices on fatigue and work environment in Canadian long-haul truck drivers," Journal of Transport & Health, vol. 24, Art. no. 101295, 2022, doi: 10.1016/j.jth.2021.101295.
5. B. J. Dietvorst, J. P. Simmons, and C. Massey, "Algorithm aversion: People erroneously avoid algorithms after seeing them err," Journal of Experimental Psychology: General, vol. 144, no. 1, pp. 114–126, 2015, doi: 10.1037/xge0000033.
6. J. Duggan, U. Sherman, R. Carbery, and A. McDonnell, "Algorithmic management and app-work in the gig economy: A research agenda for employment relations and HRM," Human Resource Management Journal, vol. 30, no. 1, pp. 114–132, 2020, doi: 10.1111/1748-8583.12258.
7. M. Hashemi and H. A. Karimi, "A critical review of real-time map-matching algorithms: Current issues and future directions," Computers, Environment and Urban Systems, vol. 48, pp. 153–165, 2014, doi: 10.1016/j.compenvurbsys.2014.07.009.
8. K. C. Kellogg, M. A. Valentine, and A. Christin, "Algorithms at work: The new contested terrain of control," Academy of Management Annals, vol. 14, no. 1, pp. 366–410, 2020, doi: 10.5465/annals.2018.0174.
9. J. D. Lee and K. A. See, "Trust in automation: Designing for appropriate reliance," Human Factors, vol. 46, no. 1, pp. 50–80, 2004, doi: 10.1518/hfes.46.1.50_30392.
10. M. Möhlmann and L. Zalmanson, "Hands on the wheel: Navigating algorithmic management and Uber drivers' autonomy," in Proc. 38th Int. Conf. Information Systems (ICIS 2017), Assoc. for Information Systems, 2017.
11. J. Muldoon and P. Raekstad, "Algorithmic domination in the gig economy," European Journal of Political Theory, vol. 22, no. 4, pp. 587–608, 2023, doi: 10.1177/14748851221082078.
12. M. T. Ribeiro, S. Singh, and C. Guestrin, "'Why should I trust you?': Explaining the predictions of any classifier," in Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining, 2016, pp. 1135–1144, doi: 10.1145/2939672.2939778.
13. W. Samek, G. Montavon, S. Lapuschkin, C. J. Anders, and K.-R. Müller, "Explaining deep neural networks and beyond: A review of methods and applications," Proceedings of the IEEE, vol. 109, no. 3, pp. 247–278, 2021, doi: 10.1109/JPROC.2021.3060483.
14. K. E. Schaefer, J. Y. C. Chen, J. L. Szalma, and P. A. Hancock, "A meta-analysis of factors influencing the development of trust in automation: Implications for understanding autonomy in future systems," Human Factors, vol. 58, no. 3, pp. 377–400, 2016, doi: 10.1177/0018720816634228.
15. Y. L. Simmhan, B. Plale, and D. Gannon, "A survey of data provenance in e-science," ACM SIGMOD Record, vol. 34, no. 3, pp. 31–36, 2005, doi: 10.1145/1084805.1084812.
16. X. Wu, L. Xiao, Y. Sun, J. Zhang, T. Ma, and L. He, "A survey of human-in-the-loop for machine learning," Future Generation Computer Systems, vol. 135, pp. 364–381, 2022, doi: 10.1016/j.future.2022.05.014.