

Human Correction Costs in Automated Hearing Reports

Abstract

A central question for occupational hearing-risk assessment is not simply whether an output is accurate, but whether it survives scrutiny of human oversight. The article develops a research agenda through critical comparison of published evidence. The primary literature contributes a bounded result: the accessible publication describes 7139 initially recruited noise-exposed subjects and longitudinal and external evaluation; baseline 3-kHz and 3-6-kHz measures contribute prognostic information. Oversight is a workflow, not the mere presence of a person. Reviewers need time, authority, and evidence to change an output. Record interventions, reviewer time, and unresolved cases. Compare assisted and unassisted outcomes without crediting the model for decisions that required unreported manual repair. The proposed comparison concerns review queues, correction logs, occupational physicians; it would test whether the specific failure scenario is detected without erasing legitimate uncertainty. Implications are restricted to the evidence roles relevant to occupational hearing-risk assessment, leaving broader efficacy and deployment claims unresolved.

Framing the scientific question

A narrow reading of human correction costs in automated hearing reports reveals a broader problem of evidence design. The question is what an informed user can conclude from review queues when occupational physicians remains uncertain. Consider the possibility of manual repair disappearing from efficiency accounting. This is a proposed stress scenario, not an observed event reported by the cited studies. Its purpose is to expose which evidence contract would be violated even if the final output looked plausible. For noise-exposed workers and clinical follow-up, the relevant claim must identify its population or artifact class, its time window, and the action it is supposed to support. A failure at any of these levels may persist even after a local technical improvement. The analysis therefore treats published methods as bounded evidence and asks how their limitations interact with this particular decision. An operational analogy is useful only if it keeps those limitations visible.

Oversight is a workflow, not the mere presence of a person. Reviewers need time, authority, and evidence to change an output. Hidden correction labor can make an autonomous system look both accurate and efficient while transferring its costs elsewhere. In this article, the relevant contrast is review queues, correction logs, occupational physicians. The proposed counterexample of manual repair disappearing from efficiency accounting puts that contrast under pressure. Record interventions, reviewer time, and unresolved cases. Compare assisted and unassisted outcomes without crediting the model for decisions that required unreported manual repair. A favorable finding would show that the intended interpretation of occupational physicians remains defensible when review queues is perturbed under the specified conditions. An unfavorable finding should narrow the claim rather than be discarded as an inconvenient case.

The scenario of manual repair disappearing from efficiency accounting is designed to challenge, not illustrate success. It supplies a negative case for human oversight that should be judged independently of how fluently the system explains it. The input record should retain enough information to decide which part of that scenario violates the intended contract. That includes the original source type, measurement or artifact status, and the reason the proposed interpretation exceeds its support. The evaluator would compare the system's response with a qualified judgment made without seeing that response. Disagreement should be classified rather than erased by rewriting the prompt. Hidden correction labor can make an autonomous

system look both accurate and efficient while transferring its costs elsewhere. A useful negative result would show that the interface cannot preserve this distinction under realistic use; a useful positive result would demonstrate preservation on independent cases rather than merely on the example used to develop the rule. The primary source remains relevant through its particular contribution: The design distinguishes association with existing impairment from prediction of later impairment. That contribution should be retained even if the integrated workflow fails. Conversely, successful integration would not remove the source's limitations or turn a proposed extension into an already observed result. This asymmetric interpretation protects both the literature and the user from an unjustified all-or-nothing verdict.

From assumptions to observable checks

Replication and transparent workflows are prerequisites for reliable biomedical claims (National Academies 2019). Reproducing a calculation and independently replicating a finding are different achievements. Recoverable data and code allow one to inspect the first; independent observations and design choices inform the second. A review should avoid calling several reanalyses of one dataset separate replications. Clear reporting of analysis changes also makes a failed replication informative instead of allowing it to disappear as an implementation dispute. This is methodological support for examining human oversight, not a claim that the source studied occupational hearing-risk assessment. A separate application-specific test is required at the interface described in this review.

FAIR principles support reusable cohorts, signals, and analysis metadata (Wilkinson et al. 2016). Reusable evidence needs more than a downloadable file. Identifiers, metadata, provenance, and meaningful access conditions help another researcher determine what an artifact contains and whether it can answer a new question. Reuse does not mean universal permission or universal scientific comparability. The information needed to judge those boundaries should travel with the artifact rather than remain in private workflow knowledge. An implementation should expose this limitation to reviewers of noise-exposed workers and clinical follow-up. The result and its evidence qualification should remain recoverable even when the interface condenses them into a short report.

Causal inference requires a clearly specified intervention, time order, and exchangeability assumptions (Hernán and Robins 2020). A prediction can be accurate without identifying what an intervention would change. A useful causal analysis begins by defining the hypothetical comparison, eligibility, and time zero. Its assumptions concern the process generating exposure and observation, not merely the flexibility of a fitted model. In this review, causal vocabulary is appropriate only when that comparison can be specified and defended. Within occupational hearing-risk assessment, this distinction is relevant to human oversight. The proposed failure scenario is manual repair disappearing from efficiency accounting; its evaluation must preserve the original source's methodological scope.

Epidemiologic design makes confounding, selection, and measurement error part of interpretation (Rothman, Greenland, and Lash 2008). Confounding, selection, and measurement error are distinct problems and need not move an estimate in the same direction. An analysis should describe them before deciding which adjustment is adequate. A larger dataset can reduce random error while leaving systematic error almost untouched. Precision should therefore be interpreted alongside the quality of the design and the eligibility of its population. This is methodological support for examining human oversight, not a claim that the source studied occupational hearing-risk assessment. A separate application-specific test is required at the interface described in this review.

False-discovery control is central when many molecular features are screened (Benjamini and Hochberg 1995). Multiplicity arises from the family of questions actually searched. Controlling false discoveries within

a named family does not automatically cover later tissue filters, alternative endpoints, or narrative choices. The proposed audit should retain exploratory branches so the family is visible. Independent confirmation then addresses a different question from whether a discovery threshold was correctly applied during screening. This is methodological support for examining human oversight, not a claim that the source studied occupational hearing-risk assessment. A separate application-specific test is required at the interface described in this review.

Medical machine learning must fit data provenance, care pathways, and prospective evaluation (Rajkomar, Dean, and Kohane 2019). The clinical relevance of a model depends on how its output changes a care process. Development data can contain information that becomes available only after the intended decision. A review must identify that availability boundary before comparing algorithms. Prospective use also changes who is measured and who is treated, so evaluation of a frozen historical dataset cannot establish the entire future workflow. Within occupational hearing-risk assessment, this distinction is relevant to human oversight. The proposed failure scenario is manual repair disappearing from efficiency accounting; its evaluation must preserve the original source's methodological scope.

Responsible health AI connects prediction with workflow, monitoring, and patient impact (Wiens et al. 2019). The point at which a model reaches practice is not the end of evaluation. New populations, changing measurements, and feedback can alter failure modes. A responsible deployment needs a route for recognizing deterioration and a person authorized to respond. Monitoring should be linked to feasible corrective actions; an alert that cannot change the workflow provides accountability only in a very limited sense. Within occupational hearing-risk assessment, this distinction is relevant to human oversight. The proposed failure scenario is manual repair disappearing from efficiency accounting; its evaluation must preserve the original source's methodological scope.

Policy distillation connects compression with sequential decision behavior (Rusu et al. 2015). Compressing a policy concerns sequential behavior rather than one isolated prediction. A small local mismatch can compound across actions, while a locally similar policy can explore differently. Evaluation should retain trajectories and consequential state transitions. The review treats policy transfer as a potential operational tool whose task-level contract still needs checking in the destination environment. An implementation should expose this limitation to reviewers of noise-exposed workers and clinical follow-up. The result and its evidence qualification should remain recoverable even when the interface condenses them into a short report.

Proximal policy optimization offers a widely used objective for stable policy updates (Schulman et al. 2017). An optimization procedure can make policy updates tractable without deciding whether the reward represents the desired behavior. Stable training and legitimate task success are separate properties. A review should inspect reward boundaries, information leakage, and held-out failure modes. A larger rewarded score is weak evidence if the policy can improve it by selecting easier tasks or hiding unresolved outputs. This is methodological support for examining human oversight, not a claim that the source studied occupational hearing-risk assessment. A separate application-specific test is required at the interface described in this review.

The information bottleneck formalizes compression that retains task-relevant information (Tishby, Pereira, and Bialek 1999). Compression is useful when it retains the information needed for the intended task. That task must be specified before one can judge what information is expendable. A bottleneck can discard nuisance variation but can also discard rare distinctions important to another user. The analysis should compare retained task information with the explicit downstream requirement rather than treating brevity as self-evidently beneficial. The connection to human correction costs in automated hearing reports is

a constraint on interpretation. It would be inappropriate to borrow the source's methodological authority as direct confirmation of the integrated domain claim.

Source-specific contributions

The relevant question in Prediction of Risk of Hearing Loss by Industry Noise from Cross-Sectional and Longitudinal Data is whether baseline hearing measurements can predict subsequent speech-frequency hearing loss in noise-exposed industrial workers (Xiao Yu et al. 2026d). Its methodological basis is sex-stratified cross-sectional modeling, longitudinal generalized estimating equations, and Cox-model nomograms using baseline high-frequency audiometry. The accessible evidence is specific: the accessible publication describes 7139 initially recruited noise-exposed subjects and longitudinal and external evaluation; baseline 3-kHz and 3-6-kHz measures contribute prognostic information. The design distinguishes association with existing impairment from prediction of later impairment. Limited follow-up, a smaller longitudinal sample, occupational selection, and untested prospective clinical impact constrain deployment claims. The source supplies one part of the comparison involving review queues, correction logs, occupational physicians. Its evidence should retain its original unit and scope; the present synthesis proposes a connection without claiming that the paper directly tested this integrated workflow.

The relevant question in BoostAPR: Boosting Automated Program Repair via Execution-Grounded Reinforcement Learning with Dual Reward Models is how reinforcement learning for program repair can overcome sparse execution feedback and coarse sequence-level credit (Yuanhao Li et al. 2026a). Its methodological basis is a three-stage pipeline combining execution-verified supervised traces, sequence- and line-level reward models, and PPO with reward redistribution to critical edit regions. The accessible evidence is specific: evaluation spans SWE-bench Verified, Defects4J transfer, HumanEval-Java, and QuixBugs, with reported gains in both repository repair and cross-language generalization. Line-level credit assignment makes test execution informative at the granularity where patches actually change behavior. Benchmark success does not by itself establish maintainability, specification fidelity, or robustness to incomplete and flaky tests. This source enters the present review as a methodological comparison for evidence processing, not as evidence about disease biology, patient outcomes, or clinical efficacy. The transferable question is whether review queues remains correctly interpreted after a computational choice involving occupational physicians. Its task-specific validation cannot be borrowed to validate the clinical claim.

Adaptive Quantization Strategies for Robust ML Inference under Distribution Shift asks how an inference system should revise its quantization policy when deployment data no longer resembles calibration data (Yinghao. “Adaptive Quantization Strategies for Robust ML Inference under Distribution Shift.” *Proceedings of the 2026 5th International Conference on Cyber Security Sang 2026). Its methodological basis is an adaptive quantization strategy that treats numerical precision as a deployment control rather than a fixed compression choice. The accessible evidence is specific: the study is framed around robustness under distribution shift and therefore makes the stability of low-precision inference, not compression alone, the central outcome. It connects efficient inference with the broader problem of shift-aware reliability. Its conclusions should be tested across architectures, bit widths, hardware kernels, and shifts that were not represented during calibration. This source enters the present review as a methodological comparison for evidence processing, not as evidence about disease biology, patient outcomes, or clinical efficacy. The transferable question is whether review queues remains correctly interpreted after a computational choice involving occupational physicians. Its task-specific validation cannot be borrowed to validate the clinical claim.

The sources should be compared through the roles they play, not merged into a single performance claim. Prediction of Risk of Hearing Loss by Industry Noise from Cross-Sectional and Longitudinal Data

addresses whether baseline hearing measurements can predict subsequent speech-frequency hearing loss in noise-exposed industrial workers (Xiao Yu et al. 2026d). By contrast, Adaptive Quantization Strategies for Robust ML Inference under Distribution Shift addresses how an inference system should revise its quantization policy when deployment data no longer resembles calibration data (Yinghao. “Adaptive Quantization Strategies for Robust ML Inference under Distribution Shift.” *Proceedings of the 2026 5th International Conference on Cyber Security Sang 2026). That difference matters for human oversight: the evidence concerning review queues is not automatically evidence concerning occupational physicians. A successful software operation, a plausible mechanistic account, and a validated outcome are different achievements. Their integration may improve traceability, but it cannot create an unmeasured link. A coherent review should identify where the sources complement one another and where their scopes remain incompatible. Here the proposed failure scenario is manual repair disappearing from efficiency accounting. The analytical value of that scenario is that it can occur without requiring any cited paper to be wrong: an otherwise valid local finding can be misused at a later interface. Preventing that misuse requires evidence about the interface itself, not simply a larger collection of local results.

An agenda for independent evaluation

The question left open by the literature is whether the intended distinction survives independent checking. Record interventions, reviewer time, and unresolved cases. Compare assisted and unassisted outcomes without crediting the model for decisions that required unreported manual repair. These procedures are proposed, not completed. The case set would include the scenario of manual repair disappearing from efficiency accounting, with its validity judged independently of the system's response. The relevant comparison axes are review queues, correction logs, occupational physicians.

A discriminating protocol would specify a preferred interpretation and at least one credible alternative before choosing measurements. It would include negative controls and observations that could favor the alternative. Reviewers would classify whether disagreement concerns data quality, model structure, or the scientific question itself. The comparison should reduce uncertainty about that distinction rather than add more observations equally compatible with every account.

An occupational study would need decision-time eligibility, subsequent outcome observation, and a realistic follow-up pathway. Model comparison should account for worker turnover and baseline impairment. Calibration in a validation cohort is informative, but improved prevention or referral quality would require evaluation of the action policy and its burdens, not a higher discrimination statistic alone.

Transfer and access limitations

Access limits are material for Adaptive Quantization Strategies for Robust ML Inference under Distribution Shift. Where only bibliographic scope or a limited abstract is available, the discussion does not assign unverified datasets, algorithms, effect sizes, or clinical recommendations. This review is a bounded conceptual synthesis rather than a systematic review with an exhaustive search protocol. It cannot establish the effectiveness of the proposed integrated workflow. Its specific contribution is to articulate how human oversight changes the interpretation of review queues, correction logs, occupational physicians. Technical sources support computational comparisons only; they do not validate biological mechanisms or patient benefit. Reviews and correspondence retain their publication roles and are not counted as new experiments. Proposed stress cases are illustrative, not observed failures attributed to the papers. Future work should therefore test the weakest inference described here, preserve contrary results, and disclose the cost of checking rather than treating additional automation as intrinsically beneficial.

Conclusion

The counterexample of manual repair disappearing from efficiency accounting shows why local success and integrated validity must be evaluated separately. Record interventions, reviewer time, and unresolved cases. Compare assisted and unassisted outcomes without crediting the model for decisions that required unreported manual repair. The central methodological contribution is specific: The design distinguishes association with existing impairment from prediction of later impairment. Extending it to occupational physicians is an additional hypothesis, not a result supplied by the citation. A trustworthy workflow would preserve the evidence source, the assumptions of the transition, and the right to stop when support is insufficient. The proposed evaluation agenda gives those obligations observable tests while withholding claims of completed validation.

References

1. Yu, Xiao, et al. "Prediction of Risk of Hearing Loss by Industry Noise from Cross-Sectional and Longitudinal Data." *Communications Medicine*, vol. 6, 2026, <https://doi.org/10.1038/s43856-026-01463-3>.
2. Li, Yuanhao, et al. "BoostAPR: Boosting Automated Program Repair via Execution-Grounded Reinforcement Learning with Dual Reward Models." *arXiv*, version 4, 18 Sept. 2026, <https://arxiv.org/abs/2605.09134v4>. Preprint.
3. Sang, Yinghao. "Adaptive Quantization Strategies for Robust ML Inference under Distribution Shift." *Proceedings of the 2026 5th International Conference on Cyber Security, Artificial Intelligence and Digital Economy*, ACM, 2026, pp. 362-68, <https://doi.org/10.1145/3819836.3819892>.
4. National Academies of Sciences, Engineering, and Medicine. *Reproducibility and Replicability in Science*. National Academies Press, 2019.
5. Wilkinson, Mark D., et al. "The FAIR Guiding Principles for Scientific Data Management and Stewardship." *Scientific Data*, vol. 3, 2016, article 160018.
6. Hernán, Miguel A., and James M. Robins. *Causal Inference: What If*. Chapman & Hall/CRC, 2020.
7. Rothman, Kenneth J., Sander Greenland, and Timothy L. Lash. *Modern Epidemiology*. 3rd ed., Lippincott Williams & Wilkins, 2008.
8. Benjamini, Yoav, and Yosef Hochberg. "Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing." *Journal of the Royal Statistical Society: Series B*, vol. 57, no. 1, 1995, pp. 289-300.
9. Rajkomar, Alvin, Jeffrey Dean, and Isaac Kohane. "Machine Learning in Medicine." *New England Journal of Medicine*, vol. 380, 2019, pp. 1347-58.
10. Wiens, Jenna, et al. "Do No Harm: A Roadmap for Responsible Machine Learning for Health Care." *Nature Medicine*, vol. 25, 2019, pp. 1337-40.
11. Rusu, Andrei A., et al. "Policy Distillation." *arXiv*, 2015, arXiv:1511.06295.
12. Schulman, John, et al. "Proximal Policy Optimization Algorithms." *arXiv*, 2017, arXiv:1707.06347.
13. Tishby, Naftali, Fernando C. Pereira, and William Bialek. "The Information Bottleneck Method." *Proceedings of the 37th Annual Allerton Conference on Communication, Control, and Computing*, 1999, pp. 368-77.