From Black-Box to Collaborative: Position Paper on Human-Guided Trace Clustering

U. Schlegel[1,2] and G. M. Tavares[1,2] and D. Schuster[1,2,3] and T. Seidl[1,2]

[1]LMU Munich, Germany [2]Munich Center for Machine Learning (MCML), Germany [3]University of Mannheim, Germany

Abstract

Trace clustering techniques partition event logs into coherent groups of process variants but fundamentally lack objective ground-truth labels against which cluster quality can be meaningfully assessed. Current evaluation practices rely on proxy metrics, such as model fitness and silhouette scores, that fail to address the inherently task-dependent and user-specific nature of meaningful process variant groupings. This position paper argues that the field urgently needs human-guided frameworks combining visual analytics and active learning to systematically incorporate domain expertise into trace clustering. We propose that coordinated visual interfaces with strategic query mechanisms, soliciting constraints, prototype labels, and goal-directed feedback, represent the missing systematic paradigm for embedding subjective expertise into algorithmic clustering. Widespread adoption of such frameworks would transform trace clustering from black-box automation into collaborative intelligence, enabling scalable, interpretable trace clustering aligned with real organizational analysis needs.

1. Introduction

Process mining has become a powerful analytical discipline, helping organizations extract insights from event logs across manufacturing, healthcare, and beyond. As process complexity increases, automated clustering techniques have become essential, and researchers have proposed numerous approaches based on distances, profiles, patterns, and graphs [SGVdA08, LNJ23, DKDWvB17,PDD26,dMGG*08]. Yet a critical question remains largely unaddressed: How do we know if the clusters are good? [vLWG12]

Unlike traditional machine learning, process mining lacks labeled ground truth [SGVdA08]. The "correct" grouping of process variants is fundamentally subjective: a compliance officer, a performance analyst, and an organizational researcher examining the same log will each require entirely different clustering criteria. We argue that this is not a methodological flaw. It is an intrinsic feature of analyzing organizational behavior that needs to be exploited. Current evaluation practices sidestep this reality by relying on proxy metrics such as internal clustering scores and model fitness, which assume a universal notion of "good" clustering [GTS*24]. These measures diverge from real analysis needs and reduce analysts to passive consumers of black-box automation.

This paper argues that the field of trace clustering must move beyond such proxies toward human-guided frameworks that treat the absence of ground truth as an opportunity to embed domain expertise into clustering. In fact, the overall challenge of balancing automation and human involvement has been previously identified by Beerepoot et al. [BDR*23], where they discuss the importance of augmenting process mining with domain knowledge. We propose that visual analytics, integrated with active learning, provides the missing architecture: iterative cycles in which analysts steer clustering through interactive visualizations and targeted feedback, and that we base our position on previous positions [Gsc15, MDCSW25]. Such systems would transform trace clustering from an opaque into a dynamic, goal-aligned process.

2. Current State: Algorithmic Solutions to a Human Problem

The process mining community has addressed clustering's inherent subjectivity primarily through algorithmic sophistication, including refined distance metrics [LNJ23] and automatic parameter tuning [DWVBVB13,GTS*24]. These advances are valuable, but they fundamentally sidestep the core issue: clustering quality cannot be defined independently of user intent [vLWG12]. Internal validation metrics, such as silhouette scores or Davies-Bouldin indices, measure mathematical properties like compactness and separation but require continuous feature spaces that often do not align with process mining data [AM15], and also cannot capture whether clusters align with domain-relevant patterns. External validation, meanwhile, requires ground truth labels that simply do not exist in exploratory process analysis [GTS*24]. Current approaches typically present clustering results as static outputs, leaving analysts to manually inspect clusters using statistics and sample traces, a cognitively demanding process with few refinement mechanisms beyond re-running algorithms with different parameters [DWVBVB13] and only sometimes matching pairs [NFFP22]. What is needed is not better algorithms, but better integration between algorithmic capabilities and human expertise, similar to interactive process [SRBF20] and interactive trace clustering [NPF*21].

3. Our Vision: Human-Guided Trace Clustering

We propose reframing trace clustering as a human-guided task rather than a fully automated analysis. This shift acknowledges that ground truth emerges from domain experts rather than solely from algorithms. Visual analytics provides the interface; active learning provides the intelligence; and explainable AI provides the transparency necessary for analysts to *understand, diagnose, and refine* clustering outcomes similar to visual analytics concepts [SSSE19].

3.1. Core Principles

Our approach rests on three core principles. First, *ground truth is constructed, not discovered*: user feedback serves as the authoritative definition of cluster quality, with expert judgments creating context-specific ground truth. Second, *labeling effort must be strategic*: active learning identifies the most informative traces for annotation, making human effort tractable even across large logs and requiring user guidance. Third, *visualization drives understanding and refinement*: coordinated views expose temporal patterns, resource utilization, and control flow structures, enabling users to understand why clusters emerge and how to improve them.

3.2. Conceptual Architecture

We envision three tightly coupled components.

Clustering Engine. An initial clustering establishes baseline groupings using established techniques (e.g., hierarchical clustering on trace edit distances [LNJ23]). This serves as a starting point for refinement, not a final answer. The engine continuously incorporates user feedback, dynamically adjusting cluster assignments as constraints accumulate.

Visual Analytics Workspace. Multiple linked views provide complementary perspectives on process variants. A cluster overview displays sizes, cohesion, and inter-cluster relationships. Process models may provide visual summaries of clusters. Trace detail views highlight intra- and inter-cluster similarities, while a performance dashboard exposes operational metrics such as cycle time and resource utilization. Critically, explainability overlays surface *why* the model assigned traces to particular clusters, exposing the dominant features and decision boundaries driving each grouping. This supports a three-stage analyst workflow similar to Spinner et al. [SSSE19]: *understand* the current clustering structure, *diagnose* misalignments with domain goals, and *refine* assignments through direct interaction. Uncertainty visualization further guides attention to cluster boundaries, where user input yields the greatest improvement in model performance.

Active Learning Module. This component selects which traces to present for labeling, treating user time as a scarce resource and employing query strategies such as uncertainty sampling near cluster boundaries, query-by-committee for traces where multiple models disagree, and diversity sampling to ensure broad variant coverage, all within a human-guided active learning setup where the model shows visualizations for labeling but takes only minimal initiative, reducing mixed-initiative guidance [SJB*21].

3.3. The Interaction Loop

The system operates through iterative refinement cycles: (1) *initialization* via standard clustering algorithms; (2) *exploration* through linked visual views; (3) *querying* via active learning; (4) *feedback* from the user, optionally with rationale; (5) *refinement* as constraints are incorporated and uncertainty recalculated; and (6) *repetition* until convergence or diminishing returns. Each cycle sharpens the model's alignment with the user's specific goals, embodying the principle that meaningful ground truth emerges through interaction rather than as a prerequisite.

4. Research Challenges

Realizing this vision requires addressing challenges across visual analytics, machine learning, and process mining. **Visualization Design:** Designing effective encodings for high-dimensional process traces, model differences, and uncertainty, without overwhelming users or sacrificing detail, remains an open problem [TOJC23]. **Active Learning Strategies:** It is unclear which query strategies best suit this setting: whether uncertainty or diversity sampling is more effective [DWVBVB13], whether pair, triplet, or single-trace formats are preferable [AM15], and how to handle inconsistent feedback as analytical understanding evolves [ZZW23]. **Clustering Integration:** Incorporating user feedback via constraint-based clustering, metric learning, or ensemble weighting presents significant technical challenges, particularly when real-time responsiveness is required. **Evaluation:** Without ground truth, measuring success is non-trivial [vLWG12]. Rigorous evaluation will require a combination of user studies, synthetic benchmarks, consistency metrics, and downstream task assessments. **Cognitive Load:** Practical deployment raises questions about session query limits, how to meaningfully contextualize feedback requests, and when to open to free exploration.

5. Implications and Directions

Human-guided clustering has the potential to catalyze important shifts in process mining practice. Users can actively steer clustering toward specific goals rather than passively observing outcomes. Making algorithmic reasoning visible builds the trust essential for adoption in high-stakes settings, while lowering technical barriers democratizes access for domain experts. The feedback captured during interactive sessions represents reusable knowledge, encoding expertise that can inform future analyses [SAD*26]. Looking ahead, integration with LLMs offers a compelling direction, enabling analysts to guide clustering through language [KBSvdA24].

6. Conclusion

The absence of ground truth in trace clustering is not a methodological flaw but a reflection of organizational reality. Different analysts, pursuing different goals, legitimately require different groupings of the same event log, and no algorithm can resolve this subjectivity on their behalf. We argue that visual analytics and active learning together provide a principled architecture for addressing this challenge, transforming clustering from an opaque preprocessing step into a goal-aligned discovery process. Significant challenges remain, but the potential rewards are substantial: clusters that are meaningful rather than merely mathematically compact, and insights grounded in domain expertise rather than proxy metrics. We call on the community to direct efforts toward systems that amplify human judgment rather than replace it.

Proceedings published by Eurographics - The European Association for Computer Graphics.

References

[AM15] APPICE A., MALERBA D.: A co-training strategy for multiple view clustering in process mining. *IEEE transactions on services computing* (2015). 1, 2

[BDR*23] BEEREPOOT I., DI CICCIO C., REIJERS H. A., RINDERLE-MA S., BANDARA W., BURATTIN A., CALVANESE D., CHEN T., COHEN I., DEPAIRE B., DI FEDERICO G., DUMAS M., VAN DUN C., FEHRER T., FISCHER D. A., GAL A., INDULSKA M., ISAHAGIAN V., KLINKMÜLLER C., KRATSCH W., LEOPOLD H., VAN LOOY A., LOPEZ H., LUKUMBUZYA S., MENDLING J., MEYERS L., MODER L., MONTALI M., MUTHUSAMY V., REICHERT M., RIZK Y., ROSEMANN M., RÖGLINGER M., SADIQ S., SEIGER R., SLAATS T., SIMKUS M., SOMEH I. A., WEBER B., WEBER I., WESKE M., ZERBATO F.: The biggest business process management problems to solve before we die. *Computers in Industry* (2023). 1

[DKDWvB17] DE KONINCK P., DE WEERDT J., VANDEN BROUCKE S. K. L. M.: Explaining clusterings of process instances. *Data Mining and Knowledge Discovery* (2017). 1

[dMGG*08] DE MEDEIROS A. K. A., GUZZO A., GRECO G., VAN DER AALST W. M. P., WEIJTERS A. J. M. M., VAN DONGEN B. F., SACCÀ D.: Process mining based on clustering: A quest for precision. In *Business Process Management Workshops* (2008). 1

[DWVBVB13] DE WEERDT J., VANDEN BROUCKE S., VANTHIENEN J., BAESENS B.: Active trace clustering for improved process discovery. *IEEE Transactions on Knowledge and Data Engineering* (2013). 1, 2

[Gsc15] GSCHWANDTNER T.: Visual analytics meets process mining: challenges and opportunities. In *International Symposium on Data-Driven Process Discovery and Analysis* (2015). 1

[GTS*24] GRIGORE I. M., TAVARES G. M., SILVA M. C. D., CERAVOLO P., BARBON JUNIOR S.: Automated trace clustering pipeline synthesis in process mining. *Information* (2024). 1

[KBSvdA24] KOURANI H., BERTI A., SCHUSTER D., VAN DER AALST W. M. P.: Process modeling with large language models. In *International Conference on Business Process Modeling, Development and Support* (2024). 2

[LNJ23] LAKSHMI NARAYANA N., JAGADISHWARI V.: Trace clustering techniques for process mining. In *International Conference on Advances in Electrical, Computing, Communication and Sustainable Technologies (ICAECT)* (2023). 1, 2

[MDCSW25] MIKSCH S., DI CICCIO C., SOFFER P., WEBER B.: Visual analytics meets process mining: Challenges and opportunities. *IEEE Computer Graphics and Applications* (2025). 1

[NFFP22] NEUBAUER T. R., FERNANDES A. G. L., FANTINATO M., PERES S. M.: Interactive trace clustering to enhance incident completion time prediction in process mining. In *WCCI Workshop on Computational Intelligence for Process Mining (CI4PM)* (2022). 1

[NPF*21] NEUBAUER T. R., PERES S. M., FANTINATO M., LU X., REIJERS H. A.: Interactive clustering: a scoping review. *Artificial Intelligence Review* (2021). 1

[PDD26] PEEPERKORN J., DE SMEDT J., DE WEERDT J.: Model-driven stochastic trace clustering. *Information Systems* (2026). 1

[SAD*26] SCHUSTER D., AIGNER W., DI FRANCESCOMARINO C., TURKAY C., ZERBATO F.: Kava-pm: Knowledge-assisted visual process mining. *Information Systems* (2026). 2

[SGVdA08] SONG M., GÜNTHER C. W., VAN DER AALST W. M. P.: Trace clustering in process mining. In *International conference on business process management (BPM)* (2008). 1

[SJB*21] SPERRLE F., JEITLER A., BERNARD J., KEIM D., EL-ASSADY M.: Co-adaptive visual data analysis and guidance processes. *Computers & Graphics* (2021). 2

[SRBF20] SCHUHMANN S., REHSE J.-R., BAUMANN S., FETTKE P.: Interactive process clustering with t-sne. In *CEUR Workshop Proceedings* (2020). 1

[SSSE19] SPINNER T., SCHLEGEL U., SCHÄFER H., EL-ASSADY M.: explAIner: A Visual Analytics Framework for Interactive and Explainable Machine Learning. *IEEE Transactions on Visualization and Computer Graphics* (2019). 2

[TOJC23] TAVARES G. M., OYAMADA R. S., JUNIOR S. B., CERAVOLO P.: Trace encoding in process mining: A survey and benchmarking. *Engineering Applications of Artificial Intelligence* (2023). 2

[vLWG12] VON LUXBURG U., WILLIAMSON R. C., GUYON I.: Clustering: Science or art? In *ICML Workshop on Unsupervised and Transfer Learning* (2012). 1, 2

[ZZW23] ZIMMERMANN L., ZERBATO F., WEBER B.: What makes life for process mining analysts difficult? a reflection of challenges. *Software and Systems Modeling* (2023). 2