Foreword to Special Section: Highlights from EuroVA 2024

The International Workshop on Visual Analytics (EuroVA) has become the premier European forum for the exchange of ideas, methods, and applications in Visual Analytics. Each year, EuroVA brings together researchers and practitioners from academia and industry to discuss the latest advances, challenges, and opportunities in the field. In its 16th installment, held in conjunction with the EuroVis International Conference in Odense, Denmark on May 27th 2024, the EuroVA workshop continued this tradition, showcasing a vibrant cross-section of the community's work. This special section of Computers & Graphics presents selected highlights from EuroVA 2024. The articles included here reflect the diversity and maturity of the field, as well as its ongoing evolution in response to new data challenges and application domains.

Bridging automated process analysis with visual-interactive exploration, Palma and Angelini [1] introduce a Visual Analytics system *IMPAVID* to support compliance assessment of incident management processes (IMP) in organizations. Its novelty lies in integrating process mining with a parameterizable cost model for compliance, allowing security assessors not only to detect and analyze deviations from reference standards but also to interactively explore and configure the parameters that drive compliance scoring. By leveraging machine learning to guide parameter selection and providing visual, data-driven feedback through coordinated views, *IMPAVID* empowers users to perform both high-level overviews and fine-grained what-if analyses of compliance scenarios and thus enables more transparent and context-aware decision-making in security and compliance domains.

A fundamentally new way of interactive data labeling is introduced by Matt et al. [2]. Their approach of class-centric labeling enables users to iteratively focus on one class at a time, efficiently identifying and labeling multiple relevant instances in batch. This outperforms traditional instance-centric approaches that struggle to scale with large datasets and many classes as they assign classes to individual instances. The class-centric perspective reduces cognitive load by transforming the labeling task into a series of simpler, binary decisions, and leverages property measures and interactive visualizations to guide users through the labeling process. By addressing the challenges of scalability and cognitive overload in large, class-rich datasets, the class-centric approach provides a new, more effective paradigm for human-in-the-loop data annotation, which very deservedly earned the authors the EuroVA Best Paper Award 2024.

The problem of interpretable inverse projections of dimensionality reduction techniques is of great importance for supporting interactive exploration and explanation of high-dimensional data in Visual Analytics. Blumberg et al. [3] propose a novel and interpretable inverse projection technique *MultiInv* that leverages multilateration—a geometric method traditionally used in geopositioning—to invert distance-preserving dimensionality reduction techniques such as MDS, PCA, Sammon's mapping, and CCA. Unlike existing inverse projection methods that often rely on complex, black-box machine learning models, *MultiInv* demonstrates competitive accuracy under favorable projection conditions and superior smoothness in gradient and decision maps, making it particularly suitable for user-driven visual analysis tasks.

Another way of making high-dimensional data accessible is investigated by Blum et al. [4]. Their approach enables the visual-interactive exploration of large tabular datasets through the simultaneous creation and navigation of both item and attribute hierarchies. Its novelty lies in supporting user-controlled, multi-level hierarchical aggregation based on both semantic metadata and statistical similarity for scalable analysis across different levels of data granularity. In combination with dual hierarchy support, flexible aggregation criteria, and a tailored multiple coordinated view setup, this abstraction-driven, user-in-the-loop hierarchy construction empowers users to uncover patterns, outliers, and relationships in high-dimensional tabular data in complex analytical workflows.

Pérez-Messina et al. [5] address the mismatch between the ephemeral nature of user interactions on one side and their importance for steering visual-interactive analyses on the other side. They argue that in order to trace analytical workflows, support collaboration, and enhance the transfer of expert knowledge across tasks, users, and sessions, interactions in Visual Analytics should be persistent. To that end, they outline a unified theoretical framework that distinguishes persistent interaction from traditional, transient interactions by emphasizing three defining qualities: persistence (artifacts endure beyond a single session), user subjectivity (artifacts encapsulate personal judgments), and data attachment (artifacts are meaningfully linked to specific data). By conceptualizing persistent interaction as a core design principle—rather than a byproduct or secondary metadata—they highlight how the intentional capture and reuse of user-generated artifacts can support knowledge transfer, generalizability, and guidance in Visual Analytics systems.

The work by Han et al. [6] lies at the intersection of the aforementioned topics of dimensionality reduction and user-driven interactive analysis. Unlike traditional dimensionality reduction methods, which often obscure the

contribution of individual features due to complex, often non-linear combinations, their method *DimenFix* allows users to explicitly preserve the values of a chosen feature (either ordinal or nominal, including external metadata like labels) along a selected axis of the embedding. This novel approach enables hybrid visualizations that combine the strengths of scatter plots for the preserved feature and dimensionality reduction techniques for the combined features. This has substantial benefits for the analyst as *DimenFix* enables the creation of task-driven visualizations that support analytical goals such as feature-focused exploration, class separation, and the integration of external information, all while retaining the benefits of dimensionality reduction.

Fast algorithms to support interactivity are crucial for Visual Analytics. Without them, the computations in the backend of a Visual Analytics system would not be able to keep up with the fast-paced nature of interactive visual exploration the analyst performs on the frontend. Wang et al. [7] make an important contribution in this direction by presenting a highly efficient technique for generating image-based visual explanations of machine learning classification models, such as decision maps and classifier maps. Traditional methods for creating these visualizations—used to reveal decision boundaries, confidence zones, and other model behaviors—are computationally expensive, often prohibiting their use in interactive Visual Analytics workflows. With the proposed approach running in near-linear time, the detailed, high-resolution classifier visualizations it creates can be used in interactive, human-in-the-loop workflows—for example, to support model interpretation, debugging, and active learning at speeds previously unattainable.

Finally, Antweiler and Fuchs [8] use Visual Analytics to address trust issues in the adoption of machine learning in clinical decision support systems. They introduce the novel concept of trust quality gates: structured checkpoints that define, assess, and provide evidence for trust-related claims at each step of the development and deployment process. These gates are operationalized through Visual Analytics methods, enabling multidisciplinary teams to collaboratively evaluate and build trust in machine learning models. This contribution has far-reaching implications and applications as it not only enhances transparency and communication among stakeholders but also helps to align with regulatory requirements in healthcare. For its practical applicability and importance to the domain of health IT, the paper received the EuroVA 2024 Honorable Mention award.

Together, the papers featured in this special section exemplify the breadth and high quality of research conducted in the field of Visual Analytics in general and presented at EuroVA 2024 in particular. They span foundational algorithmic advances, innovative interaction paradigms, novel systems and domain-driven applications, each addressing critical challenges at the forefront of the discipline. Importantly, these contributions are not only of academic significance; they provide concrete solutions and conceptual frameworks that are poised to shape the daily practice of Visual Analytics, empowering practitioners and researchers alike to tackle increasingly complex data-driven problems with confidence, transparency, and impact.

Hans-Jörg Schulz
Marco Angelini

References

[1] Alessandro Palma and Marco Angelini: “*IMPAVID: Enhancing incident management process compliance assessment with visual analytics*”, Computers & Graphics, article no. 104243, 2025

[2] Matthias Matt, Jana Sedlakova, Jürgen Bernard, Matthias Zeppelzauer, Manuela Waldner: “*Scalable Class-Centric Visual Interactive Labeling*”, Computers & Graphics, article no. 104240, 2025

[3] Daniela Blumberg, Yu Wang, Alexandru Telea, Daniel A. Keim, Frederik L. Dennig: “*MultiInv: Inverting multidimensional scaling projections and computing decision maps by multilateration*”, Computers & Graphics, article no. 104234, 2025

[4] Michael Blum, Jonas Blum, Madhav Sachdeva, Jürgen Bernard: “*HAXplorer: Interactive visual exploration of hierarchical item and attribute spaces*”, Computers & Graphics, article no. 104233, 2025

[5] Ignacio Pérez-Messina, Davide Ceneda, Victor Schetinger, Silvia Miksch: “Persistent interaction: A conceptualization of user-generated artefacts in Visual Analytics”, Computers & Graphics, article no. 104232, 2025

[6] Zixuan Han, Diede van der Hoorn, Thomas Höllt, Qiaodan Luo, Leonardo Christino, Evangelos Milios, Fernando V. Paulovich: “*DimenFix: A novel meta-strategy to preserve user-defined data values on dimensionality*

reduction layouts", Computers & Graphics, article no. 104231, 2025

[7] Yu Wang, Cristian Grosu, Alexandru Telea: "*Computing fast and accurate maps for explaining classification models*", Computers & Graphics, article no. 104230, 2025

[8] Dario Antweiler, Georg Fuchs: "*Trust at every step: Embedding trust quality gates into the visual data exploration loop for machine learning-based clinical decision support systems*", Computers & Graphics, article no. 104212, 2025

Hans-Jörg Schulz received his doctorate degree in Computer Science in 2010 from University of Rostock, Germany. Since 2018, he is an associate professor in the Department of Computer Science at Aarhus University in Aarhus, Denmark. His research interests include information visualization, visual analytics, and human-data interaction. Over the course of his career, he has contributed visual analytics methods and tool in application domains like food science, climate science, and system biology. More about him can be found at http://hjschulz.net Contact: hjschulz@cs.au.dk.

Marco Angelini is an associate professor in computer science with Link University of Rome since 2023. He is the Director of the xAIber Research Laboratory focused on advancing Human-centered AI research, component of the HCI research group at Sapienza University of Rome, and member of the AI Commission of the Presidency of the Council of the Italian Government. His main research interests include Visual Analytics, Human-Centered AI, Cybersecurity. More about him can be found at https://sites.google.com/dis.uniroma1.it/angelini. Contact: m.angelini@unilink.it