MLVis 2026: Machine Learning Methods in Visualisation for Big Data (2026)
D. Archambault, I. Nabney and J. Peltonen (Editors)

Parameter Space Analysis through Guided Visual Interpolations

Benedikt Kantz [1], Peter Waldert [1], Stefan Lengauer [1], Clemens Staudinger [2], Stefan Schuster [2], and Tobias Schreck [1]

[1]Graz University of Technology, Austria
[2]voestalpine Stahl GmbH, Linz, Austria

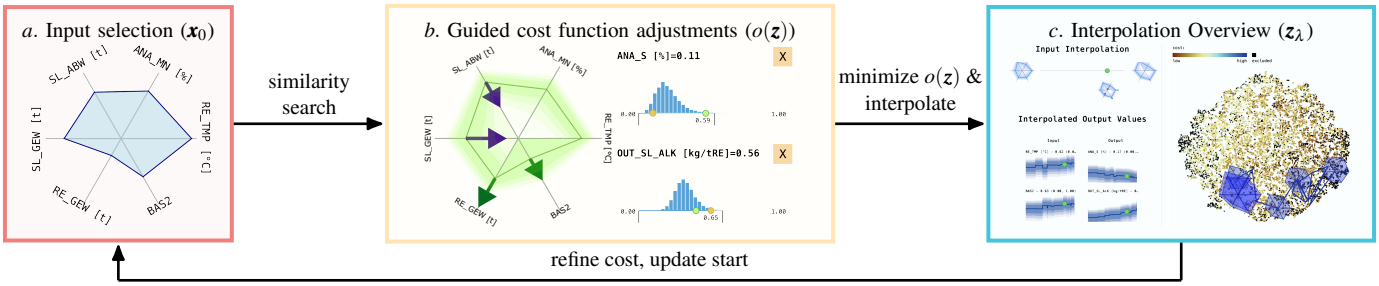

Figure 1: *Our proposed exploration workflow within the ParamInter tool. The input selection (a.) allows the users to select their initial parameters. Next, the guided cost function adjustments (b.) foster the understanding of the dependencies between input parameters and output parameters, while allowing data selection and setup of a cost function. Once the intended cost function has been set up, the interpolation overview (c.) shows a possible path between the start and intended end composition in the data and embedding space. The colors of the embedding space represent the value of the cost function, showing regions with low costs directly.*

Abstract

We propose Parameter Space Analysis through Guided Visual Interpolations (ParamInter), a novel tool for high dimensional input parameter space analysis by making interpolation towards optimal parameter sets explorable using guided analytics. The interpolation is accompanied by both small multiples in linked views, and utilizes t-Distributed Stochastic Neighbor Embedding (t-SNE) representations to show an interpolation overview. ParamInter uses a guided exploration loop focusing on the interpolation towards user-specified target parameters from many output parameters. The exploration process is additionally guided through eXplainable Artificial Intelligence (XAI)-based effect suggestions throughout our tool. ParamInter, compared to prior work, focuses on the integration of state-of-the art effect-based XAI and Uncertainty Quantification (UCQ) approaches for guidance, and introduces an interpolation towards the optimal solution through interpolation between the initial parameter setting and the optimal setting. We also add an interpretability layer for dimensionality-reduced data by displaying our novel interpolation towards the optimum, enhanced by small multiples of the input parameters on top. We demonstrate the direct applicability of our tool on a real-world use case for a blast furnace optimisation process, where a multi-objective problem is solved through modeling and visualisation.

CCS Concepts

• ***Human-centered computing*** → ***Scientific visualization; Visual analytics; Interactive systems and tools;*** • ***Computing methodologies*** → ***Model development and analysis;***

1. Introduction

The discovery of optimal or advantageous data points within high-volume, high-dimensional datasets can be formulated either mathematically as minimisation problems or via iterative filtering. Our problem setting is motivated by a typical input-output model, where continuous input parameters $\boldsymbol{x} \in \boldsymbol{X}$ drive an unknown process with multiple outputs $\mathbf{f}(\boldsymbol{x})$, yielding continuous $\boldsymbol{y} \in \boldsymbol{Y}$ output vectors, i.e. a basic input-output model [SHB*14]. These can be optimised under a specific cost function $o(\boldsymbol{y})$, i.e. $\arg\min_{\boldsymbol{x}} o(\mathbf{f}(\boldsymbol{x}))$ to obtain

the best parameter set. Approaches within visualisation, and visual analytics specifically, often employ parallel coordinates or turn towards the Pareto front [CSP*19; YWR03; CMMK20; CAS*13]. We, however, turn in our visual parameter space analysis towards optimisation of the parameter through interpolations, as well as uncertainties and sensitivity analysis within the space [SHB*14]. Our tool, ParamInter, through these aspects, fosters the understanding of the progression towards the optimal solution by providing an interface to smoothly interpolate between a chosen start configuration and an optimisation target. We additionally employ XAI as a guidance approach to provide a local hint of relevant input parameters interactively the cost function composition. ParamInter additionally, integrates UCQ into the visualisations, providing an indication of the confidence within the data itself, and supporting the decision-making process [SPC*24; MBZK24].

Our approach centers on the use of glyphs, namely radar charts, as a mediator for setting the initial input parameter search and for displaying input parameter sets. We also utilize the space within the glyphs for our XAI and UCQ displays whenever appropriate, guiding the users in their analysis. The radar charts also show the interpolation while the user scrubs through it interactively, strengthening the interpretation of the interpolation. The glyphs are supported by small multiples of histograms for brushing and composition of the optimisation function, and line charts with uncertainty displays. Our interpolation is, finally, also drawn over t-SNE plots, as shown in Figure 1 *c*.. Each t-SNE plot is concerned with one input- or output group, to aid with a clearer understanding of the relations. The interpolation and visualisations update on each change to the cost function $o(\boldsymbol{z})$ and filter $s(\boldsymbol{Z})$, enhancing understanding of the data composition and optimal parameter regions, including guidance as shown in Figure 1 *b*..

The general setting of having arbitrary continuous input-output parameters allows us to apply our tool to various optimisation and general problems [KH13]. We specifically conducted an informal expert session for a metallurgy setting to gauge the applicability of our approach. The novel tool, finally, provides **guided creation of cost functions** through interactive elements, UCQ, XAI, **explorable interpolation** towards optimal solution. The code is archived on Zenodo [Kan26].

2. Related Works

Prior tools and explorations have already tackled optimisation problems in conjunction with parameter space analysis, either from a multi-objective optimality perspective with many possible solutions (Pareto front), using parallel coordinates as the core visualisation paradigm, or by integrating UCQ and XAI into the decision-making process.

Pareto Front: Works within the high-dimensional optimisation field focus on variation within the Pareto front and parallel coordinates [CP07; GFED17; CSP*19; CMMK20; GWR] to visualise multiple solutions jointly. Prior work to apply these principles within visualisation explore joint cost-optimisation problems in various settings [CAS*13]. The tool enables the discovery of choices with the same costs, exploring the trade-offs between various cost-optimal points. Further, works also investigate overlaying markers for the underlying dimensionality-reduced space to explore their distributions [RVW24]. Existing visual analytic tools, furthermore, already use parallel coordinates to display Pareto-optimal solutions, visualising them in an integrated fashion [CSP*19; CMMK20].

XAI and Uncertainty visualisation: Integrating UCQ into visualisation has been shown to be advantageous [MBZK24; BAZ*21; CSKV24], with certain visual parameters being more effective compared to others – especially full probability density functions (pdfs) over Confidence Intervals (CIs) [SPC*24]. All cited works, nevertheless, make the case that integrating uncertainty into user interfaces often improve the quality of decision-making and user trust. Additionally, XAI is already used in some applications to foster transparency within tools, while there are still open challenges like deployment and evidence for long term benefit within applications [GA25; KZW25]

However, most prior works focus on specific, singular solutions with no sensitivity analysis, hindering the comparison of instances. Hence, we introduce a novel **interpolation scheme** to search for optimal data points across the **input and output spaces**, supported by UCQ and XAI to guide the user.

3. Interpolation through the Parameters

Our interpolation approach within ParamInter smoothly transitions from a starting point chosen from the input parameters, $\boldsymbol{x}_0 \in \boldsymbol{X}$, to an optimal point, $\boldsymbol{x}_1 \in \boldsymbol{X}$. We start from a selection

$$\boldsymbol{x}_{\text{sel}} \in \left\{ (x^0,\dots,x^l)^T \;\middle|\; \min_{\bar{\boldsymbol{x}}\in\boldsymbol{X},\bar{\boldsymbol{x}}=\{\bar{x}^0,\dots,\bar{x}^k\}} \bar{x}^i \le x^i \le \max_{\bar{\boldsymbol{x}}\in\boldsymbol{X},\bar{\boldsymbol{x}}=\{\bar{x}^0,\dots,\bar{x}^k\}} \bar{x}^i \right\}.$$

The selected input parameters are used to search for the k closest samples $\boldsymbol{x}_{\text{knn}} \in \boldsymbol{X}$ from the provided input dataset $\boldsymbol{x} \in \boldsymbol{X}$ using k-Nearest Neighbors, from which the start sample $\boldsymbol{z}_0 \in \boldsymbol{Z}$ is chosen, with $\boldsymbol{x} = (z^0,\dots,z^l)^T$ and $\boldsymbol{y} = (z^{l+1},\dots,z^{l+1+j})^T$, i.e. $\boldsymbol{z} \in \boldsymbol{Z}$ is the combination of input and output parameters. This combination allows us to include both input and output variables of a system into our cost and filter functions. l and j are the number of input and output parameters, respectively. The other end of the interpolation, i.e. $\boldsymbol{z}_1$, is determined by a user-specified composite cost function of the form $o(\boldsymbol{z}) = \|\mathbf{d}(\boldsymbol{z})\|$ with

$$\mathbf{d}(\boldsymbol{z}) = (z - z^{*,h}, \dots \mid h \in H)^T,$$

with H being the set of chosen target parameters indices, and $z^{*,h}$ being the chosen target value corresponding to the index. The minimum of this cost function is chosen as the target for interpolation by evaluating the cost $o(\boldsymbol{z})$ for each value within the dataset $\boldsymbol{Z}$ after applying a user-specific filter operation on the data set $\boldsymbol{Z}_s = s(\boldsymbol{Z})$. Our tool supports bounding each variable for now, as seen in the next section. This leads to the optimum $\boldsymbol{z}_1 = \arg\min_{\boldsymbol{z}\in\boldsymbol{Z}_s} o(\boldsymbol{z})$.

Once the final data sample $\boldsymbol{z}_1$ has been chosen, we interpolate between them in both input and output space. First, the input space is interpolated using the convex combination

$$\boldsymbol{x}^*_\lambda = \lambda \boldsymbol{x}_0 + (1-\lambda)\boldsymbol{x}_1,$$

over a discretized $\lambda \in (0,1)$. We then feed each sample into a learned model Light Gradient Boosting Machine (LightGBM) ensemble $\boldsymbol{f}^*$ [ZSH17] to get all outputs $\boldsymbol{y}^*_\lambda = \boldsymbol{f}^*(\boldsymbol{x}^*_\lambda)$, and stacking $\boldsymbol{x}^*$

and $\boldsymbol{y}^*$ into $\boldsymbol{z}^*$. These samples are then used to find the closest real data point within the filtered set with a composite cost, i.e.

$$\boldsymbol{z}_\lambda = \arg\min_{z \in \boldsymbol{Z}_s} \left\| \boldsymbol{z} - \boldsymbol{z}^*_\lambda \right\| + \left| o(\boldsymbol{z}) - o(\boldsymbol{z}^*_\lambda) \right| .$$

Only the resulting nearest data point is then used for visualisations and interpolation, keeping the shown data within the set.

4. Visual Representations

The introduced interpolation principle is incorporated into the proposed *visual* analytics flow to yield our ParamInter. We will show the progression by introducing our visual elements one-by-one.

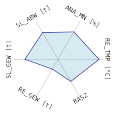

The selection process for $\boldsymbol{x}_{\text{sel}}$ is driven by a user-editable radar chart, as shown in the small glyph to the left, where the rough input parameter settings can be configured, with the k nearest possible actual data points as selectable choices for starting points $\boldsymbol{z}_0$. We use radar charts as the central glyph throughout our interface, as they allow easy editing within the space they require, despite their flaws [ALBR16]. The radar charts are also used, whenever appropriate, to convey uncertainty within the data, as we display it directly within the retrieved chart, allowing direct comparison between dimensions and harnessing the glyph space more effectively.

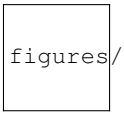

The optimisation function $o(\boldsymbol{z})$ can be tuned by adjusting variables from the input and output columns using an interactable histogram shown on the left, once the initial input parameter set has been chosen. Selecting a parameter adds it to the set H, which allows the user, through the histogram tool, to both define the target points $z^{*,h}$ by dragging the target point in yellow, and crop the data region for later interpolation, effectively defining $s(\boldsymbol{Z})$.

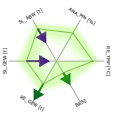

We also display a sensitivity indicator on the radar chart whenever the user hovers over an output parameter. This indicator is based on a learned LightGBM [ZSH17] ensemble on the individual parameter y^i, i.e. $y^i = f^{*,i}(\boldsymbol{x})$, i being the output dimension. We calculate the sensitivity $\boldsymbol{w}^i = \Phi_{f^{*,i}}(\boldsymbol{x})$ using the XAI method SmoothGrad, effectively approximating the gradient $\nabla_{\boldsymbol{x}} f^{*i}(\boldsymbol{x})$. This method has been shown to be quite robust in the face of noise, especially when combined with a tree-based Machine Learning method and densely sampled input space. It has also been tested in an applied industrial setting, similar to our use cases demonstrated in this paper [KSF*24]. The adjustments of the output parameters is used to provide the users an estimate for the interpolation calculation by providing data-driven effect estimation. The uncertainty estimates are also shown as shaded areas, providing an estimate of the data variability at this point. We use an uncertainty-aware Variational Autoencoder with a Gaussian process as the decoder [TSMF23]. Prior research has shown that a complete view of the pdf is advantageous in single-dimensional settings for more accurate decision-making [SPC*24].

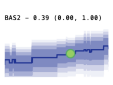

The changes in each output variable y^{*i}_λ relative to the interpolated variables are, furthermore, shown as small multiples, enriched with uncertainty estimates for all dimensions, complementing the visualisations within the radar charts. Users can hover over the full range to get the value for the data point at that specific progression, enabling them to "scrub" through the interpolation, related to how videos are usually scrubbed through [SHH15].

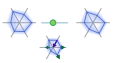

ParamInter, furthermore, provides a view of the interpolation between the input parameters, with explanations and uncertainties shown, again on top of the radar charts and updated as the user scrubs through the interpolation. This interface is linked to all small charts and shows the values for the specific data point simultaneously during scrubbing.

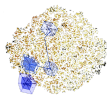

Finally, the interpolations are overlaid on a t-SNE-reduced [vdMH08] scatter plot of the various parameter groups, e.g. input, output, and combined using fast GPU-accelerated libraries [RPN20]. These provide users with an indication of the extent of the interpolation across the entire data space. We chose a nonlinear interpolation to reveal possible hidden close relations in the data, which a linear embedding like Principal Component Analysis [MR93] can not. Any linear projection also produces simple linear paths for our interpolation scheme, i.e. a linearly projected $\hat{\boldsymbol{z}}_\lambda$ will stay linear as

$$\begin{aligned}\hat{\boldsymbol{z}}_\lambda &= \mathbf{P}\boldsymbol{z}_\lambda = \mathbf{P}(\lambda \boldsymbol{z}_0 + (1-\lambda)\boldsymbol{z}_1) = \lambda \mathbf{P}\boldsymbol{z}_0 + (1-\lambda)\mathbf{P}\boldsymbol{z}_1 \;= \\ &= \lambda \hat{\boldsymbol{z}}_0 + (1-\lambda)\hat{\boldsymbol{z}}_1 .\end{aligned}$$

Much more interesting insights of the high-dimensional space can therefore be achieved using non-linear layouts, which has been shown to be advantageous when grouping and imposing structure [EHA*22]. This grouping, however reduces the semantic meaning of our interpolation lines, which has to kept in mind during analysis [CKJ*25]. The t-SNE plot is enriched by a linked display of both the cost function $o(\boldsymbol{z})$ and filter $s(\boldsymbol{z})$ applied to each data point, while compositing the function. We use color mappings [CSH20] to present costs in an accessible, contrast-rich view, even in dense or noisy regions. Users can still refer to the scrubbing view to get actual data values for their exploration.

This flow enables the users to iteratively refine their parameters, alternating between input and output spaces, while being alerted to sensitivities with respect to the outputs. ParamInter, therefore fosters the discovery of advantageous parameters by showing the intermediate steps between a target output sample and initial input parameters, while also providing iterative means to exploration if the parameter set is not yet satisfactory through small multiples and guidance by low-dimensional representation. The user is, however, also free to choose any points in the low-dimensional space by hovering over them to reveal the input parameters and set them as a new start point.

5. Case Study: Blast Furnace

We demonstrate ParamInter on a use case co-developed with experts from the steelmaking industry, a blast furnace process optimisation. Our specific use case concerns the output of bases via the slag to prevent instabilities in the blast furnace [GCL20], while also keeping sulfur output via the hot metal minimal advantageous for downstream processes. The problem has six input variables comprising different parameters for the blast furnace operations, i.e. temperature, slag and iron output, manganese concentration, and

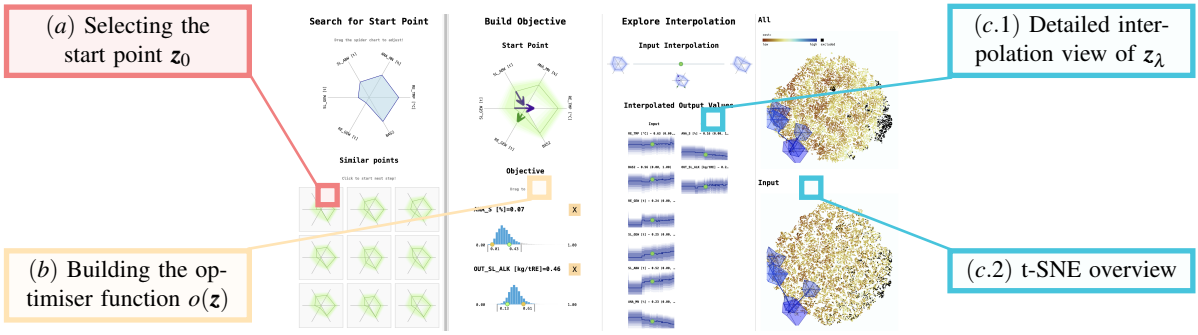

Figure 2: *Overview over the full interface of ParamInter.* $(a), (b), (c)$ *show the stages of the exploration flow within for the loaded dataset and use case; the blast furnace optimisation. We note that the dataset is normalized to* $[0, 1]$. *The expert selects the start using the radar chart as the input method in* (a)*, then sets up the cost function* $o(\boldsymbol{z})$ *and filter* $s(\boldsymbol{Z})$ *in* (b)*, and can finally observe the interpolation over the values in* $(c.1)$ *and across the t-SNE in* $(c.2)$.

basicity. These parameters need to be adjusted for these two output parameters. We gathered about 16,000 samples for this analysis. The specific optimisation problem was raised by industry experts and requires the discovery of parameters using an informed approach based on a repository of prior process data. We conducted an open-ended expert interview series, where we first demonstrated the tool, then gathered feedback. These were conducted with our two co-authors from the voestalpine Stahl GmbH. Their feedback was then incorporated into the tool, presented again with the updates; and finally provided to the experts for their down stream use. The tool is now deployed for their own internal use. We want to outline the analysis process the experts would use to come to a decision for their parameters next.

When applying ParamInter to this data to arrive at decisions about operating procedures– as in Figure 2 – the metallurgy experts can start by selecting the appropriate or current start point $\boldsymbol{z}_0$ of the blast furnace, as seen in part (a). This adjustment is performed by dragging the values in the radar chart to the desired position and then selecting one of the proposed data points. Next, the experts can adapt the optimiser function $o(\boldsymbol{z})$ within (b) by interacting with the histograms. The expert chooses the sulfur output (`ANA_S`) and base output (`OUT_SL_ALK`), then clips away the extreme tails of the histogram as they represent extreme operating regions not usually visited. Finally, the target values are set using the circle within the histogram, setting everything up for the interpolation. The t-SNE output in $(c.2)$ is already updated, indicating that specific regions are omitted and showing the most advantageous areas right away. The radar charts additionally reveal, through their change towards a bigger area, that the parameters will most likely increase. While the expert is aware of the influence of the process parameters on each of the targets individually, it is a delicate task to control the blast furnace process in a way that both are optimal simultaneously. The full effect can only be understood by investigating the interpolation in $(c.1)$ – it directly shows in which direction the parameters have to move, once the interpolation progresses towards the target – both using XAI and the line charts containing the uncertainties for each interpolated point. The case at hand requires a rise in almost all input parameters to increase the base content and decrease the sulfur content in the output, as already indicated in the t-SNE overview. The specifics of the progression can be investigated as well, as the expert can hover over a specific part of the interpolation to retrieve the relevant output value, while also getting a visual indication of the distributional uncertainty at that point.

The interpolated radar chart at the top provides additional intuition into the sensitivity of the parameters, offering further insights into which direction might be advantageous to optimise in the next step. These insights would therefore allow the experts to tune the process parameters according to the final input variables, or use another operating point on the path to the optimum that might be less optimal, but offer a lower uncertainty, i.e. display narrower bands in the interpolation view, or be more stable with regard to the outputs, which can be identified through the XAI markers. The experts also remarked on some limitations of the interface, especially in regard to the radar plots – as the precise configuration of inputs is difficult using the current interface, and would require improved facilities.

6. Discussion & Conclusion

Our tool, ParamInter, provides an iterative, visual, and guided optimisation tool that enables experts to discover input-output relations of parameters and achieve ideal compositions through interpolation. The optimisation process uses guidance to select input parameters, build an optimisation function, and interpolate towards the optimum. ParamInter is already deployed for the blast-furnace use case, strengthening our case for real-world applicability.

Improvements to the visual guidance could include integration of the sensitivity results to steer the user not only visually but also to incorporate gradient estimation into the interpolation algorithm. Further research could also examine the decision-making capabilities of the UCQ and XAI integrations, especially in our industrial use case.

Acknowledgements

This work is partially supported by the HEREDITARY Project, as part of the European Union's Horizon Europe research and innovation programme under grant agreement No GA 101137074.

References

[ALBR16] ALBO, YAEL, LANIR, JOEL, BAK, PETER, and RAFAELI, SHEIZAF. "Off the Radar: Comparative Evaluation of Radial Visualization Solutions for Composite Indicators". *IEEE Transactions on Visualization and Computer Graphics* 22.1 (Jan. 2016), 569–578. ISSN: 1077-2626. DOI: 10.1109/tvcg.2015.2467322 3.

[BAZ*21] BHATT, UMANG, ANTORÁN, JAVIER, ZHANG, YUNFENG, et al. "Uncertainty as a Form of Transparency: Measuring, Communicating, and Using Uncertainty". *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. AIES '21. Virtual Event, USA: Association for Computing Machinery, 2021, 401–413. DOI: 10.1145/3461702.3462571 2.

[CAS*13] CHEN, SHAHAR, AMID, DAVID, SHIR, OFER M., et al. "Self-organizing maps for multi-objective pareto frontiers". *2013 IEEE Pacific Visualization Symposium (PacificVis)*. IEEE, Feb. 2013, 153–160. DOI: 10.1109/pacificvis.2013.6596140 2.

[CKJ*25] CASHMAN, DYLAN, KELLER, MARK, JEON, HYEON, et al. "A Critical Analysis of the Usage of Dimensionality Reduction in Four Domains". *IEEE Transactions on Visualization and Computer Graphics* 31.10 (Oct. 2025), 9405–9423. ISSN: 2160-9306. DOI: 10.1109/tvcg.2025.3567989 3.

[CMMK20] CIBULSKI, LENA, MITTERHOFER, HUBERT, MAY, THORSTEN, and KOHLHAMMER, JÖRN. "PAVED: Pareto Front Visualization for Engineering Design". *Computer Graphics Forum* 39.3 (June 2020), 405–416. DOI: 10.1111/cgf.13990 2.

[CP07] CHINCHULUUN, ALTANNAR and PARDALOS, PANOS M. "A survey of recent developments in multiobjective optimization". *Annals of Operations Research* 154.1 (Oct. 2007), 29–50. DOI: 10.1007/s10479-007-0186-0 2.

[CSH20] CRAMERI, FABIO, SHEPHARD, GRACE E., and HERON, PHILIP J. "The misuse of colour in science communication". *Nature Communications* 11.1 (Oct. 2020). ISSN: 2041-1723. DOI: 10.1038/s41467-020-19160-7 3.

[CSKV24] CRESSWELL, JESSE C., SUI, YI, KUMAR, BHARGAVA, and VOUITSIS, NOËL. *Conformal Prediction Sets Improve Human Decision Making*. 2024. arXiv: 2401.13744 [cs.LG] 2.

[CSP*19] CAJOT, SÉBASTIEN, SCHÜLER, NILS, PETER, MARKUS, et al. "Interactive Optimization With Parallel Coordinates: Exploring Multidimensional Spaces for Decision Support". *Frontiers in ICT* 5 (Jan. 2019). ISSN: 2297-198X. DOI: 10.3389/fict.2018.00032 2.

[EHA*22] ECKELT, KLAUS, HINTERREITER, ANDREAS, ADELBERGER, PATRICK, et al. "Visual Exploration of Relationships and Structure in Low-Dimensional Embeddings". *IEEE Transactions on Visualization and Computer Graphics* 29.7 (Mar. 2022), 3312–3326. DOI: 10.1109/TVCG.2022.3156760 3.

[GA25] GARN, WOLFGANG and AMIRGHASEMI, MEHRDAD. "Transparency of combinatorial optimisations via machine learning and explainable AI". *Annals of Operations Research* 354.1 (Nov. 2025), 427–458. DOI: 10.1007/s10479-025-06684-8 2.

[GCL20] GEERDES, M., CHAIGNEAU, R., and LINGIARDI, O. *Modern Blast Furnace Ironmaking: An Introduction*. 4th ed. IOS Press, 2020, 208–210. ISBN: 1643681222,9781643681221 3.

[GFED17] GOGUELIN, STEVEN, FLYNN, JOSEPH M., ESSINK, WESLEY P., and DHOKIA, VIMAL. "A Data Visualization Dashboard for Exploring the Additive Manufacturing Solution Space". *Procedia CIRP* 60 (2017), 193–198. ISSN: 2212-8271. DOI: 10.1016/j.procir.2017.01.016 2.

[GWR] GUO, ZHENYU, WARD, MATTHEW O., and RUNDENSTEINER, ELKE A. "Nugget Browser: Visual Subgroup Mining and Statistical Significance Discovery in Multivariate Datasets". *2011 15th International Conference on Information Visualisation*. IEEE, 13–15. DOI: 10.1109/IV.2011.21 2.

[Kan26] KANTZ, BENEDIKT. *Dakantz/ParamInter: MLVis 2026*. 2026. DOI: 10.5281/ZENODO.19886990 2.

[KH13] KEHRER, JOHANNES and HAUSER, HELWIG. "Visualization and Visual Analysis of Multifaceted Scientific Data: A Survey". *IEEE Transactions on Visualization and Computer Graphics* 19.3 (Mar. 2013), 495–513. ISSN: 1077-2626. DOI: 10.1109/tvcg.2012.110 2.

[KSF*24] KANTZ, BENEDIKT, STAUDINGER, CLEMENS, FEILMAYR, CHRISTOPH, et al. "Robustness of Explainable Artificial Intelligence in Industrial Process Modelling". *ICML'24 Workshop ML for Life and Material Science: From Theory to Industry Applications*. 2024. URL: https://openreview.net/forum?id=ieAn7KOFBz 3.

[KZW25] KUCHER, KOSTIANTYN, ZOHREVANDI, ELMIRA, and WESTIN, CARL A. L. "Towards Visual Analytics for Explainable AI in Industrial Applications". *Analytics* 4.1 (Feb. 2025), 7. ISSN: 2813-2203. DOI: 10.3390/analytics4010007 2.

[MBZK24] MARUSICH, LAURA, BAKDASH, JONATHAN, ZHOU, YAN, and KANTARCIOGLU, MURAT. "Using AI Uncertainty Quantification to Improve Human Decision-Making". *Proceedings of the 41st International Conference on Machine Learning*. Vol. 235. Proceedings of Machine Learning Research. PMLR, July 2024, 34949–34960. URL: https://proceedings.mlr.press/v235/marusich24a.html 2.

[MR93] MAĆKIEWICZ, ANDRZEJ and RATAJCZAK, WALDEMAR. "Principal components analysis (PCA)". *Computers & Geosciences* 19.3 (Mar. 1993), 303–342. ISSN: 0098-3004. DOI: 10.1016/0098-3004(93)90090-r 3.

[RPN20] RASCHKA, SEBASTIAN, PATTERSON, JOSHUA, and NOLET, COREY. *Machine Learning in Python: Main developments and technology trends in data science, machine learning, and artificial intelligence*. 2020. arXiv: 2002.04803 [cs.LG] 3.

[RVW24] RAVAL, SHIVAM, VIÉGAS, FERNANDA, and WATTENBERG, MARTIN. "Hypertrix: An indicatrix for high-dimensional visualizations". *2024 IEEE Visualization and Visual Analytics (VIS)*. IEEE, Oct. 2024, 1–5. DOI: 10.1109/vis55277.2024.00073 2.

[SHB*14] SEDLMAIR, MICHAEL, HEINZL, CHRISTOPH, BRUCKNER, STEFAN, et al. "Visual Parameter Space Analysis: A Conceptual Framework". *IEEE Transactions on Visualization and Computer Graphics* 20.12 (Dec. 2014), 2161–2170. ISSN: 1077-2626. DOI: 10.1109/tvcg.2014.2346321 1, 2.

[SHH15] SCHOEFFMANN, KLAUS, HUDELIST, MARCO A., and HUBER, JOCHEN. "Video Interaction Tools: A Survey of Recent Work". *ACM Comput. Surv.* 48.1 (Sept. 2015). DOI: 10.1145/2808796 3.

[SPC*24] SARMA, ABHRANEEL, PU, XIAOYING, CUI, YUAN, et al. "Odds and Insights: Decision Quality in Exploratory Data Analysis Under Uncertainty". *Proceedings of the CHI Conference on Human Factors in Computing Systems*. CHI '24. ACM, May 2024, 1–14. DOI: 10.1145/3613904.3641995 2, 3.

[TSMF23] TRAN, BA-HIEN, SHAHBABA, BABAK, MANDT, STEPHAN, and FILIPPONE, MAURIZIO. "Fully Bayesian Autoencoders with Latent Sparse Gaussian Processes". *Proceedings of the 40th International Conference on Machine Learning*. Vol. 202. Proceedings of Machine Learning Research. PMLR, 23–29 Jul 2023, 34409–34430 3.

[vdMH08] Van der MAATEN, LAURENS and HINTON, GEOFFREY. "Visualizing Data using t-SNE." *Journal of Machine Learning Research* 9.11 (Nov. 2008), 2579. ISSN: 1532-4435 3.

[YWR03] YANG, JING, WARD, MATTHEW O., and RUNDENSTEINER, ELKE A. "Interactive hierarchical displays: a general framework for visualization and exploration of large multivariate data sets". *Computers & Graphics* 27.2 (Apr. 2003), 265–283. ISSN: 0097-8493. DOI: 10.1016/s0097-8493(02)00283-2 2.

[ZSH17] ZHANG, HUAN, SI, SI, and HSIEH, CHO-JUI. *GPU-acceleration for Large-scale Tree Boosting*. 2017. arXiv: 1706.08359 [stat.ML] 2, 3.