

A Unified Process for Visual-Interactive Labeling

Jürgen Bernard^{1,4}, Matthias Zeppelzauer², Michael Sedlmair³, and Wolfgang Aigner²

¹TU Darmstadt, Germany

²St. Pölten University of Applied Sciences, St. Pölten, Austria

³Universität Wien, Vienna, Austria

⁴Fraunhofer Institute for Computer Graphics Research, IGD, Darmstadt, Germany

Abstract

Assigning labels to data instances is a prerequisite for many machine learning tasks. Similarly, labeling is applied in visual-interactive analysis approaches. However, the strategies for creating labels often differ in the two fields. In this paper, we study the process of labeling data instances with the user in the loop, from both the machine learning and visual-interactive perspective. Based on a review of differences and commonalities, we propose the 'Visual-Interactive Labeling' (VIAL) process, conflating the strengths of both. We describe the six major steps of the process and highlight their related challenges.

CCS Concepts

•Human-centered computing → Visual analytics; Information visualization; •Theory of computation → Active learning;

1. Introduction

A central topic in data science is the understanding of data instances and the discovery of knowledge from data. Research has addressed this issue from different perspectives. On the one hand, machine learning (ML) provides a rich toolset for the automatic indexing, organization, and categorization of huge amounts of data. On the other hand, information visualization (VIS) aims at the organization and presentation of data as well as knowledge discovery in a visual-interactive way. While both disciplines have their respective strengths for data analysis, they have an even stronger potential when they are combined in visual analytics (VA) approaches [SSZ*16, ERT*17]. Still, however, the complementary strengths are often not fully exploited.

Building upon approaches investigating combinations of ML with VIS in general [SSZ*16, ERT*17], this work explicitly addresses the common goal of labeling tasks. We refer to labeling as the assignment of labels y to given input instances x (objects, elements, or samples), e.g., to find functions f that map instances to labels, i.e. $f(x) = y$. A fundamental difference between ML and VIS approaches is the way these goal is achieved. ML most often operates fully automatically and is thus predominantly model-centric. In turn, the user-centric VIS perspective emphasizes the information need of the user. Both perspectives are complementary and of high importance for real-world problems.

In ML, *active learning* (AL) strategies have been introduced to incorporate user knowledge. In AL an algorithmic model proactively asks the user (referred to as the oracle) for feedback (e.g., labels) to improve the learning model [Set09]. Since user interactions are time-consuming and thus expensive, AL aims at mini-

mizing the amount of required user interactions by querying only that information that will improve the accuracy of the given model in a best possible way. Popular classes of strategies include uncertainty sampling [CM05], measuring the agreement of a committee of sub-models [SOS92, TVC*11], quantifying the expected model change, reducing the model error, or assessing the output variance [WH11, Set09]. One drawback of model-based AL strategies is that users only play a marginal role in the identification and selection of instances to be labeled. Hence, the selection of instances is neither based on expert knowledge, nor on the human ability to identify patterns.

In the VIS community labeling is an important task as well. Many approaches accept feedback from users for data instances of interest as input to learn the users' information need. Important tasks supported by visual-interactive interfaces are the analysis of model results, the identification and selection of instances, as well as labeling per se. Example labeling interfaces accept user-defined numerical interestingness or similarity scores [BSB*15, BRS*17], categorical labels used for classification tasks [HKBE12, BDV*17], or labels to assign subjective relevance information [BKSS14, SSJK16]. More complex labeling techniques allow, e.g., the manipulation of spatial proximity relations [BLBC12, BSR*14]. In contrast to ML, VIS approaches seem to prefer user-centered over model-based criteria.

We assume that the model-centered AL and the user-centered VIS perspectives have complementarity and unexploited strengths for labeling tasks. Building upon and extending notions of 'interactive learning' presented in pioneer approaches combining AL and VIL [SG10, HNH*12], we investigate the strength of both, and propose an abstracted and unified process in a VA context




Figure 1: The abstracted AL process. A data source contains unlabeled (U) and labeled (L) data. Preprocessing, segmentation (P/S), feature extraction (FE), and feature selection (FS) are upstream steps in the process. A learning model (ML) is trained and evaluated as a black-box approach. Candidate suggestion (CS) strategies query new labels y from the oracle which are used to iteratively adapt the classifier until a stopping criterion is met.

that we refer to as *Visual-Interactive Labeling (VIAL)*. Our line of approach complies with established process models in IV and VA [CMS99, vW05, KAF*08, CG16], resembling the abstract data and interaction flow, as well as user-based knowledge generation [SSS*14]. While these models offer a high degree of abstraction, we extend and substantiate these general process towards labeling tasks. Process models and surveys in AL exist as well [Set09, Ols09, TVC*11, WH11, PG14, HRC14], see Figure 1 for a generalized workflow. However, these models often fall short in visual interfaces as well as knowledge generation support [SSS*14].

Most related approaches indicate the combination of model-based and user-based labeling. Seifert and Granitzer [SG10] as well as Höferlin et al. [HNH*12] present visual-interactive classification techniques, both with an emphasis on AL-support. We build upon the techniques employed in these pioneer VIL approaches, abstract primary steps for a conceptualization of the VIAL process, and additionally shed light on challenges occurring in the process. Bernard et al. propose a regression-based process where users play an active role in assigning numerical labels [BSB*15]. From this work, we take away the idea to support data-centered, model-centered, and user-centered criteria for label suggestion. In [BSR*14] a labeling approach is presented that models distance functions for mixed data. Inspiring for our approach is the series of pitfalls for the design of this specific type of labeling approach, which we will adopt towards labeling in general. Finally, Mamani et al. propose a visualization-assisted methodology for interacting with instances to transform feature spaces [MFNP13].

Although the latter approaches are inspiring, they are specific towards a data type, employed technique, application goal, user group, or target variable y . In contrast, the rationale of our unified VIAL process is to abstract from concrete approaches and to propose a general and conceptual labeling workflow. Furthermore, one aspect of the labeling process remains largely uncharted—the three types of output: labeled data, trained models, and gained knowledge. VIAL, on the contrary, obtains a data-, model-, and user-centric perspective with three outputs: data, models, and knowledge.

In this work, we contribute an abstract conceptual trans-disciplinary process that combines the AL and the VIS perspective. We explain the six crucial steps of the VIAL process, point out

their interplay, and describe how AL and VIS can contribute to the respective step. In addition, we discuss the major design and development challenges in every step from both the AL and the VIS perspective. Future approaches may benefit from the VIAL process in two ways. First, we provide an integrated view of AL and VIS in a VA setting that may inspire novel innovative approaches that go beyond the borders of the individual disciplines. Second, the outlined challenges help to overcome inherent hurdles in the VIAL process and to make informed design decisions.

2. The Visual-Interactive Labeling Process

Based on a review of related works in AL and VIS, we propose the VIAL process. We unify the main building blocks to an iterative process consisting of six steps shown in Figure 2. The VIAL process is special in its detail for exploration and labeling tasks, as well as its emphasis on three output types, i.e., labeled data, trained learning models, and gained knowledge. In the following, we describe each of the six steps in detail. For each step, we present the particular challenges from the ML and VIS perspectives together with additional challenges that may emerge when the strengths of AL and VIS are combined in a unified process.

2.1. Preprocessing and Feature Extraction

Preprocessing is a fundamental step in almost every data analysis approach that needs to be handled with care. We combine the preprocessing step with the *mapping* of real-world objects into more abstract representations (features). Existing labeling approaches either directly adopt semantically interpretable attributes of data instances (e.g., the GDP of a country [BSR*14]) or apply complex descriptors [BYRN99] yielding abstract feature spaces.

Challenges A challenging design consideration is whether internal feature representations should be visible to the user. From a VIS perspective transparent feature spaces can be beneficial for the knowledge generation process [KPB14, Gle16]. The visualization of semantically interpretable features may be particularly beneficial for non-experts. Non-semantic features, however, such as Fourier or Cosine transform coefficients (of e.g., images) are difficult to grasp even for experts. One possible drawback of visible features for the VIAL process is self-biasing [BSR*14]. Users being aware of individual features may be tempted to trim new labels with respect to feature values instead of respective instances. Another challenge regards novel methodologies for feature extraction, such as deep learning [LBH15] and sparse representations [WMM*10] that learn abstract representations directly from the data. These representations may adapt during the iterative training phase. Not only the selected feature subsets, but now also the features themselves may change, which may confuse the user. To address this complexity, VIAL approaches may, e.g., provide visual representations showing the evolution of the features, or support the interactive adaption of features [KPB14]. Finally, in VIAL the visualization of the features themselves could further be used as an indicator for training progress, evaluation, and success.

2.2. Learning Model

The choice of learning models primarily depends on the data and the labeling task. Classifiers [HNH*12], regression models

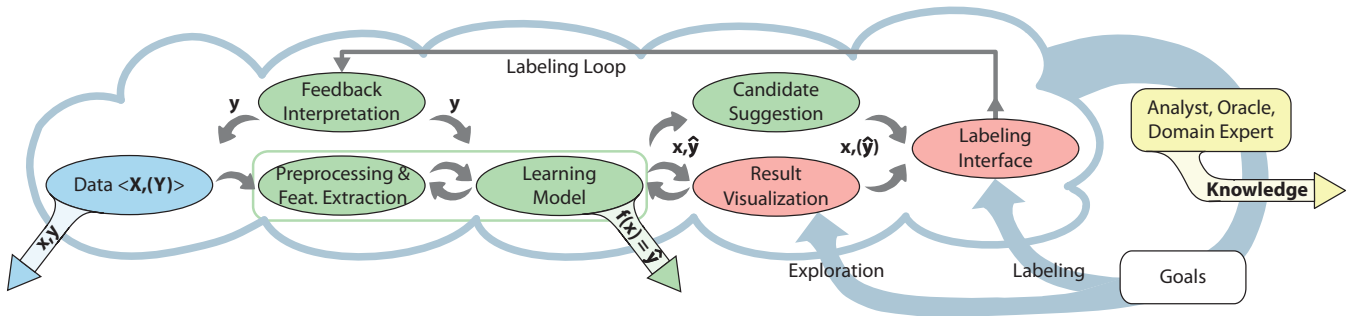


Figure 2: The VIAL process. Four algorithmic models (green) and two primary visual interfaces (red) are assembled to an iterative labeling process. To resemble the special characteristics of the AL and the VIS perspective, the VIAL process contains a branch (from “Learning Model” to “candidate suggestion” and “result visualization”, since both are complementary). At a glance the VIAL process can be applied for data exploration and labeling tasks. The output of the VIAL process is threefold: labeled data, learned models, and gained knowledge.

[BSB*15], or more complex ensembles thereof may be appropriate. In the VIAL process the learning model is directly coupled with visual interfaces facilitating analytic reasoning and model refinement [SSZ*16, ERT*17]. The trained models represent a primary output of the VIAL process, building the basis for downstream applications.

Challenges The VIAL process is iterative by nature which raises at least two challenges. First, learning models need to be instantly re-trainable, ideally in real-time. Second, result visualizations need to be sensitive to model changes. Thus, learning models are required that can iteratively adapt their internal parameters to changes training data, such as decision trees [vdEvW11] or neural networks [SB-VLK09]. Many learning models are, however, difficult to visualize [Gle16]. Another important issue is to select a suitable termination criterion for learning and labeling. Due to the limited capacity of most classifiers [KM97, Vap13] the learning progress converges at some point in time. Termination criteria can be both intrinsic (e.g., model change) or extrinsic (e.g., classification accuracy) [Set09]. A traditional visual analytics approach is measuring quality aspects that help analysts to validate labeling or model convergence [Gle16].

2.3. Result Visualization

Result visualization corresponds to the VIS perspective on the labeling process. We identify three primary benefits for the VIAL process. First, result visualization can facilitate exploration tasks supporting hypotheses and insight generation about the data as well as the knowledge generation process [SSS*14]. Second, tightly coupled learning models and result visualizations enable user-centered model refinement [SBVLK09, vdEvW11]. Third, result visualization allows users to select meaningful candidates for labeling and thus, serves as a complement to model-based AL heuristics for the suggestion of candidates [SG10, HNH*12].

Challenges In general, result visualization poses challenges in the representation of high-dimensional data. Visual-interactive interfaces supporting overview and detail visualizations are one option to tackle this issue. Dimension reduction [SZS*16] and data aggregation techniques [EF10] help to condense the data, for the price of individual challenges such as the applicability, quality, or uncertainty of algorithms in connection with their parameters.

A particular design challenge for labeling approaches is whether and how predicted labels should also be visualized as they may cause biases. Patient well-being may serve as an example where physicians may be affected by trained models from other experts [BSB*15]. The improvement of learning models by direct manipulation is associated with more general visual analytics challenges [KAF*08, SSS*14]. For that purpose the VIAL process can be facilitated with parameter space analysis support [SHB*14], or techniques for the visual comparison [GAW*11] of different model outputs.

2.4. Candidate Suggestion

Automated candidate suggestion (as in AL) and the visualization of model results (from VIS), cf. Section 2.3 represent two complementary alternatives for the identification of labeling candidates. From an AL perspective, users are queried in a model-centered way to improve the model accuracy [Set09]. In turn, in the VIS perspective the user is typically assigned an active role in the candidate selection process. The VIAL process joins both perspectives and proposes to either include AL-based guidance concepts included in visual interfaces, or visual-interactive interfaces for the analysis and steering of AL strategies.

Challenges A major challenge in the candidate suggestion comes with the AL process, i.e. the selection of AL heuristics. A rich set of techniques for candidate suggestion exists [Set09]. The applicability of individual AL heuristics depends on the data, the types of labels, and the ML model [WH11, Set09], as well as on the interplay of model-based and user-based candidate selection. The VIAL process proposes the joint suggestion and selection of candidates performed by the user (VIS) and the model (AL). Pioneer VIAL implementations [SG10, HNH*12, BSB*15] indicate the potential of combined candidate suggestion and selection strategies. However, such hybrid approaches remain an open topic and a promising direction of future research. Considering the need for very large labeled data sets, e.g., used for deep learning [LBH15, Sch15] or sparse coding [WMM*10], a downstream challenge is the generalization of gathered label information for yet unlabeled instances. Selections of most representative objects (centroids) or multiple instances at once are two promising approaches for future work.

2.5. Labeling Interface

The goal of the labeling interface is to create pairs of instances x and labels y . Every time a user labels an instance, the labeling loop can be triggered, possibly leading to an improved learning model. This iterative approach is supported from both the AL and the VIS perspective and is resembled in the VIAL process. Particularly the VIS perspective requires meaningful visualization and interaction designs to support the labeling process in a meaningful way.

Challenges One challenge is the visual mapping of labeling candidates. In order to submit qualified feedback, users must be able to grasp the characteristics of queried instances. In case users already know individual instances, visual identifiers can be used, e.g., national flags for countries or images of soccer players [BSR*14, BRS*17]. In other cases users already have an intrinsic knowledge of the labeling alphabet, e.g., object classes like cats and dogs. If complex instances (e.g., multimodal data) or unknown instances (e.g., of a new class) are to be identified and labeled, detailed information needs to be visualized to support decisions, possibly in combination with special interaction designs. Examples include visual representations of unknown patient histories [BSB*15], abstracted features [KPB14], or relations between clusters and metadata [BRS*12]. Another class of challenges relates to the candidate suggestion of the AL process. Candidate suggestion has the primary goal to improve learning and to reduce labeling effort. Thus, the users' information need is not explicitly captured. Additionally, users may only be able to label small portions of instances they have knowledge about. Meaningful visualization designs may be one strategy to address these problems. Finally, the interaction design raises challenges in complex learning situations where labels are less distinct and exhibit complex semantics as in similarity learning [BLBC12, BSR*14].

2.6. Feedback Interpretation

An often neglected question is how to interpret complex user feedback and pass it to the learning model [SZS*16]. We assume that the difficulty to interpret feedback is related to the complexity of user interaction. For simple labeling tasks such as selecting a category, feedback interpretation may be straightforward. For more complex tasks the situation becomes more challenging, e.g., for relations between multiple instances [BLBC12, BSR*14], or *implicit* user feedback where user behavior is observed without explicit queries.

Challenges We elaborate challenges in feedback interpretation from two perspectives: the *concrete interaction* and the *abstract user intent*. The first perspective arises in more complex interaction paradigms that go beyond simple labeling tasks. An example is user interaction in terms of spatial re-arrangements of instances in 2D [BLBC12] which can be interpreted in at least three ways [BSR*14]. Challenges from the second perspective are related, but take the discussion deeper into human computer interaction. Mental models of users communicated through visual-interactive interfaces open large spaces for interpretation and thus may deviate significantly from the measured feedback. Implicit feedback falls into this category [Nor02], as well as data from sensor devices such as eye tracking [BKR*14] which may be addressed in future VIAL approaches.

3. Discussion and Future Work

In this work, we carved out the benefits of joint approaches using AL and VIS for labeling data instances. While we focused on the conceptual baseline, the quantification of success of the VIAL process remains future work. When describing the six core steps of the VIAL process, we went for a broad overview of techniques instead of one primary application example. Future work includes two strategies: first, implementations of the VIAL process in application examples and second, a more explicit and holistic reflection of single application examples. The latter is inspired by the Sacha et al. [SSS*14] using, e.g., Jigsaw [SGL08] as an explicit example. Finally, future work includes evaluations of the challenges described in the six steps.

4. Conclusion

We presented the VIAL process that adopts and extends the process model from AL and VIS and thereby combines the strengths of model-centered active learning with user-centered visual-interactive labeling. Overall we identified six core steps of the process. For every step, we described both the AL and the VIS perspective, discussed respective challenges, and outlined inspiring examples and open topics.

References

- [BDV*17] BERNARD J., DOBERMANN E., VÖGELE A., KRÜGER B., KOHLHAMMER J., FELLNER D.: Visual-interactive semi-supervised labeling of human motion capture data. In *Visualization and Data Analysis (VDA)* (2017). doi:<https://doi.org/10.2352/ISSN.2470-1173.2017.1.VDA-387>. 1
- [BKR*14] BLASCHECK T., KURZHALS K., RASCHKE M., BURCH M., WEISKOPF D., ERTL T.: State-of-the-Art of Visualization for Eye Tracking Data. In *EuroVis (STAR)* (2014), Eurograph. doi:[10.2312/eurovisstar.20141173.4](https://doi.org/10.2312/eurovisstar.20141173.4)
- [BKSS14] BEHRISCH M., KORKMAZ F., SHAO L., SCHRECK T.: Feedback-driven interactive exploration of large multidimensional data supported by visual classifier. In *IEEE Visual Analytics Science and Technology (VAST)* (2014), pp. 43–52. 1
- [BLBC12] BROWN E. T., LIU J., BRODLEY C. E., CHANG R.: Disfunction: Learning distance functions interactively. In *IEEE VAST* (2012), IEEE, pp. 83–92. 1, 4
- [BRS*12] BERNARD J., RUPPERT T., SCHERER M., SCHRECK T., KOHLHAMMER J.: Guided discovery of interesting relationships between time series clusters and metadata properties. In *Knowledge Management and Knowledge Technologies (i-KNOW)* (2012), ACM, pp. 22:1–22:8. doi:[10.1145/2362456.2362485.4](https://doi.org/10.1145/2362456.2362485.4)
- [BRS*17] BERNARD J., RITTER C., SESSLER D., ZEPPELZAUER M., KOHLHAMMER J., FELLNER D.: Visual-interactive similarity search for complex objects by example of soccer player analysis. In *Proc. of IVAPP, VISIGRAPP* (2017), pp. 75–87. doi:[10.5220/0006116400750087.1.4](https://doi.org/10.5220/0006116400750087.1.4)
- [BSB*15] BERNARD J., SESSLER D., BANNACH A., MAY T., KOHLHAMMER J.: A visual active learning system for the assessment of patient well-being in prostate cancer research. In *VIS Workshop on Visual Analytics in Healthcare* (2015), ACM, pp. 1–8. doi:[10.1145/2836034.2836035.1.2.3.4](https://doi.org/10.1145/2836034.2836035.1.2.3.4)
- [BSR*14] BERNARD J., SESSLER D., RUPPERT T., DAVEY J., KUIJPER A., KOHLHAMMER J.: User-based visual-interactive similarity definition for mixed data objects-concept and first implementation. *Journal of WSCG 22* (2014). 1, 2, 4

- [BYRN99] BAEZA-YATES R. A., RIBEIRO-NETO B.: *Modern Information Retrieval*. Addison-Wesley Longman, 1999. 2
- [CG16] CHEN M., GOLAN A.: What may visualization processes optimize? *IEEE TVCG* 22, 12 (2016), 2619–2632. doi:10.1109/TVCG.2015.2513410. 2
- [CM05] CULOTTA A., MCCALLUM A.: Reducing labeling effort for structured prediction tasks. In *Proc. of 20th Nat. Conf. on AI* (2005), AAAI'05, AAAI Press, pp. 746–751. 1
- [CMS99] CARD S. K., MACKINLAY J. D., SHNEIDERMAN B. (Eds.): *Readings in Information Visualization: Using Vision to Think*. Morgan Kaufmann, San Francisco, CA, USA, 1999. 2
- [EF10] ELMQVIST N., FEKETE J.-D.: Hierarchical aggregation for information visualization: Overview, techniques, and design guidelines. *IEEE TVCG* 16, 3 (2010), 439–454. doi:10.1109/TVCG.2009.84. 3
- [ERT*17] ENDERT A., RIBARSKY W., TURKAY C., WONG B. W., NABNEY I., BLANCO I. D., ROSSI F.: The state of the art in integrating machine learning into visual analytics. *CGF* (2017). doi:10.1111/cgf.13092. 1, 3
- [GAW*11] GLEICHER M., ALBERS D., WALKER R., JUSUFI I., HANSEN C. D., ROBERTS J. C.: Visual comparison for information visualization. *Inf. Vis.* 10, 4 (2011), 289–309. doi:10.1177/1473871611416549. 3
- [Gle16] GLEICHER M.: A framework for considering comprehensibility in modeling. *Big Data* 4, 2 (2016), 75–88. doi:10.1089/big.2016.0007. 2, 3
- [HKBE12] HEIMERL F., KOCH S., BOSCH H., ERTL T.: Visual classifier training for text document retrieval. *IEEE TVCG* 18, 12 (2012), 2839–2848. 1
- [HNN*12] HÖFERLIN B., NETZEL R., HÖFERLIN M., WEISKOPF D., HEIDEMANN G.: Inter-active learning of ad-hoc classifiers for video visual analytics. In *Visual Analytics Science and Technology (VAST)* (2012), IEEE, pp. 23–32. doi:10.1109/VAST.2012.6400492. 1, 2, 3
- [HRC14] HASAN M., ROY-CHOWDHURY A. K.: Incremental activity modeling and recognition in streaming videos. In *Proceedings of IEEE Conf. on Comp. Vis. & Pat. Rec.* (2014). 2
- [KAF*08] KEIM D., ANDRIENKO G., FEKETE J.-D., GÖRG C., KOHLHAMMER J., MELANÇON G.: *Information visualization*. Springer-Verlag, Berlin, Heidelberg, 2008, ch. Visual Analytics: Definition, Process, and Challenges, pp. 154–175. doi:10.1007/978-3-540-70956-5_7. 2, 3
- [KM97] KARPINSKI M., MACINTYRE A.: Polynomial bounds for vc dimension of sigmoidal and general pfaian neural networks. *J. of Comp. & Sys. Sc.* 54, 1 (1997), 169–176. doi:http://dx.doi.org/10.1006/jcss.1997.1477. 3
- [KPB14] KRAUSE J., PERER A., BERTINI E.: Infuse: Interactive feature selection for predictive modeling of high dimensional data. *IEEE Trans on Vis. & Comp. Graph.* 20, 12 (2014), 1614–1623. doi:10.1109/TVCG.2014.2346482. 2, 4
- [LBH15] LECUN Y., BENGIO Y., HINTON G.: Deep learning. *Nature* 521, 7553 (2015), 436–444. 2, 3
- [MFNP13] MAMANI G. M. H., FATORE F. M., NONATO L. G., PAULOVICH F. V.: User-driven feature space transformation. In *Eurographics Conf. on Visualization* (2013), Eurographics, pp. 291–299. doi:10.1111/cgf.12116. 2
- [Nor02] NORMAN D. A.: *The Design of Everyday Things*, reprint paperback ed. Basic Books, New York, 2002. 4
- [Ols09] OLSSON F.: A literature survey of active machine learning in the context of natural language processing. 2
- [PG14] PUTTEMANS S., GOEDEMÉ T.: Optimal object categorization under application specific conditions. In *Proc. of DCVISIGRAPP* (2014), vol. 1, VISIGRAPP, pp. 25–34. 2
- [SBVLK09] SCHRECK T., BERNARD J., VON LANDESBERGER T., KOHLHAMMER J.: Visual cluster analysis of trajectory data with interactive kohonen maps. *Information Visualization* 8, 1 (2009), 14–29. doi:10.1057/ivs.2008.29. 3
- [Sch15] SCHMIDHUBER J.: Deep learning in neural networks: An overview. *Neural networks* 61 (2015), 85–117. 3
- [Set09] SETTLES B.: *Active Learning Literature Survey*. Tech. Report 1648, Univ. of Wisconsin–Madison, 2009. 1, 2, 3
- [SG10] SEIFERT C., GRANITZER M.: User-based active learning. In *IEEE Int. Conf. on Data Mining Worksh.* (2010), pp. 418–425. doi:10.1109/ICDMW.2010.181. 1, 2, 3
- [SGL08] STASKO J., GÖRG C., LIU Z.: Jigsaw: Supporting investigative analysis through interactive visualization. *Information Visualization* 7, 2 (2008), 118–132. doi:10.1145/1466620.1466622. 4
- [SHB*14] SEDLMAIR M., HEINZL C., BRUCKNER S., PIRINGER H., MÖLLER T.: Visual parameter space analysis: A conceptual framework. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (2014), 2161–2170. doi:10.1109/TVCG.2014.2346321. 3
- [SOS92] SEUNG H. S., OPPER M., SOMPOLINSKY H.: Query by committee. In *Proc. of the 5th Ann. Worksh. on Comput. Learning Theory* (New York, NY, USA, 1992), COLT '92, ACM, pp. 287–294. doi:10.1145/130385.130417. 1
- [SSJK16] SEEBACHER D., STEIN M., JANETZKO H., KEIM D. A.: Patent Retrieval: A Multi-Modal Visual Analytics Approach. In *EuroVis Workshop on Visual Analytics (EuroVA)* (2016), Eurographics, pp. 013–017. 1
- [SSS*14] SACHA D., STOFFEL A., STOFFEL F., KWON B. C., ELLIS G. P., KEIM D. A.: Knowledge generation model for visual analytics. *IEEE Trans. Vis. Comput. Graph.* 20, 12 (2014), 1604–1613. doi:10.1109/TVCG.2014.2346481. 2, 3, 4
- [SSZ*16] SACHA D., SEDLMAIR M., ZHANG L., LEE J. A., WEISKOPF D., NORTH S. C., KEIM D. A.: Human-Centered Machine Learning Through Interactive Visualization: Review and Open Challenges. In *Symp. on Artificial Neural Networks, Computational Intelligence and Machine Learning* (2016). 1, 3
- [SZS*16] SACHA D., ZHANG L., SEDLMAIR M., LEE J. A., PELTONEN J., WEISKOPF D., NORTH S. C., KEIM D. A.: Visual Interaction with Dimensionality Reduction: A Structured Literature Analysis. *IEEE TVCG* 23, 01 (2016), 241–250. doi:10.1109/TVCG.2016.2598495. 3, 4
- [TVC*11] TUIA D., VOLPI M., COPA L., KANEVSKI M., MUNOZ-MARI J.: A survey of active learning algorithms for supervised remote sensing image classification. *IEEE Journal of Selected Topics in Signal Processing* 5, 3 (2011), 606–617. 1, 2
- [Vap13] VAPNIK V.: *The nature of statistical learning theory*. Springer science & business media, 2013. 3
- [vdEvW11] VAN DEN ELZEN S., VAN WIJK J. J.: Baobabview: Interactive construction and analysis of decision trees. In *IEEE Conf. on Vis. Anal. Sc. & Techn.* (2011), pp. 151–160. doi:10.1109/VAST.2011.6102453. 3
- [vW05] VAN WIJK J. J.: The value of visualization. In *VIS 05. IEEE Visualization, 2005.* (2005), pp. 79–86. doi:10.1109/VISUAL.2005.1532781. 2
- [WH11] WANG M., HUA X.-S.: Active learning in multimedia annotation and retrieval: A survey. *ACM Trans. Intell. Syst. Technol.* 2, 2 (2011), 10:1–10:21. 1, 2, 3
- [WMM*10] WRIGHT J., MA Y., MAIRAL J., SAPIRO G., HUANG T. S., YAN S.: Sparse representation for computer vision and pattern recognition. *Proceedings of the IEEE* 98, 6 (2010), 1031–1044. 2, 3