

VIAL – A Unified Process for Visual-Interactive Labeling

Jürgen Bernard^{1,4}, Matthias Zeppelzauer², Michael Sedlmair³, and Wolfgang Aigner^{2,5}

¹TU Darmstadt, Germany

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

³Jacobs University, Bremen, Germany

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

⁵TU Wien, Austria

Abstract

The assignment of labels to data instances is a fundamental prerequisite for many machine learning tasks. Moreover, labeling is a frequently applied process in visual-interactive analysis approaches and visual analytics. However, the strategies for creating labels usually differ between these two fields. This raises the question whether synergies between the different approaches can be attained. 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 that unifies both approaches. We describe the six major steps of the process and discuss their specific challenges. Additionally, we present two heterogeneous usage scenarios from the novel VIAL perspective, one on metric distance learning and one on object detection in videos. Finally, we discuss general challenges to VIAL and point out necessary work for the realization of future VIAL approaches.

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 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 tool set for the automatic indexing, organization, and categorization of huge amounts of data. On the other hand, 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, the complementary strengths are often not fully exploited.

Building upon approaches investigating combinations of ML with VIS in general, 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 (i.e., objects, elements, or samples). Labels can be used to find functions f that either map instances to labels, i.e. $f(x) = y$, or define relations between instances, i.e., $f(x_1, x_2) = y$. Labels in this context can be of different type, such as categorical labels [HNH*12] in classification tasks, numerical labels [BSB*15] in regression tasks, relevance scores [SSJK16] in retrieval tasks, as well as labels that represent a relation between two instances (e.g. for learning similarity and distance measures between objects) [BRS*17].

A fundamental difference between ML and VIS approaches is

the way this goal is achieved. ML most often operates fully automatically, i.e., instances with previously defined labels are fed into a supervised ML algorithm which in turn learns the function f from the data. Once trained, the algorithm can be applied for the labeling of new data. As such ML methods are predominantly model-centric. In turn, the VIS perspective emphasizes the knowledge generation of the user, e.g., by visual-interactive labeling interfaces. As such, the VIS perspective is most often user-centric.

In the presence of unlabeled data, incremental ML methods provide iterative solutions to learning ML algorithms, and to refine them continuously [LHW17]. Active Learning (AL) [Set09] is a special type of incremental supervised ML where the user is integrated into the learning process to guide the training. As such, a connecting factor between AL and VIS is the user-in-the-loop principle. In AL an algorithmic model pro-actively asks the user for feedback about unlabeled data to improve a learning model. Thereby, the user is queried for labels of instances the learner is unsure about and which will potentially improve the quality of the model. Typical outputs of the AL process are datasets enriched with labels as well as models learned in the process. Compared to VIS, however, in AL the active role of the user remains rather marginal. In most cases, AL neither incorporates the user's abilities to *select* meaningful instances (e.g., reflecting visual patterns in data), nor does the AL process foster the user's knowledge generation process.

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 and further support 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 numerical interestingness scores to train regression models [BSB*15] or user-defined class labels to train classifiers [HNH*12]. More complex labeling techniques allow, e.g., the manipulation of spatial proximity of data instances [BLBC12] to make statements about (similarity) relations between instances. In contrast to ML, VIS approaches seem to prefer user-centered over model-based criteria. Potential drawbacks of VIS approaches relate to a) the usefulness of purely user-selected instances for labeling to build accurate, robust, and generalizable algorithmic models, and to b) the emphasis on knowledge generation while the enrichment of data with labels is often neglected.

We assume that the model-centered AL and the user-centered VIS perspectives have complementary and unexploited synergies for labeling tasks. Building upon and extending notions of 'interactive learning' presented in pioneer approaches combining AL and VIS [SG10, HNH*12], we investigate the strengths of both, and propose an abstracted and unified process in a Visual Analytics (VA) context that we refer to as *Visual-Interactive Labeling (VIAL)*. Our line of approach complies with established process models in visualization 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]. However, these models often fall short in visual interfaces as well as knowledge generation support [SSS*14].

Most related approaches either indicate the combination of model-based and user-based labeling [HNH*12, SSZ*16] or propose methodologies or concepts for the interactive propagation of user feedback for labeling tasks [MFNP13, BSR*14]. Although these 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 a conceptual cross-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.

We present related work in Section 2. In Section 3, we introduce the VIAL process, followed by two usage scenarios in Section 4 building upon the VIAL process. We discuss limitations and potential future work in Section 5, and conclude with Section 6.

This article is an extension of a previous workshop paper on the same topic [BZSA17], and provides a theoretical framework for our empirical studies in this area [BHZ*17]. The article at hand has been fundamentally reshaped in many ways including a substantial extension of related work, the integration of usage scenarios to better illustrate the application of the model, the extension, and more in-depth descriptions of the model's steps including research challenges and possible solutions. Moreover, it includes more details in most sections of the paper, particularly the discussion, limitations, and future work.

2. Background & Related Work

In the ML and VIS literature, we find different and partly complementary approaches to data labeling. A major contribution of the VIAL process is to join both strategies to develop a broader and more general perspective. In the following, we describe the labeling process in ML (AL) and VIS and point out differences.

2.1. Labeling in Machine Learning

Labeling of datasets is an important pre-requisite for the training of supervised ML models. Recent developments have shown that labeled datasets are not only important to enable the objective comparison of ML approaches but are further necessary to successfully train today's complex classifiers, such as deep neural networks (DNNs) [LBH15]. To reduce the effort of labeling large datasets (of potentially millions of instances) different strategies have been developed (e.g., captcha-based information collection [ZK10], web-based annotation systems [RTMF08], paid micro tasks on platforms like mechanical turk [BKG11], and game-based approaches [VAD04, ROR11]).

Pre-Labeling In supervised ML the dataset is usually generated and annotated completely in an offline process before classifier training starts. To differentiate this type of labeling process from other more interactive and on-line labeling processes, we refer to this process as *pre-labeling*. Pre-labeling is an expensive and time-consuming process that should - in the best case - be performed by domain experts. In the absence of a sufficiently large number of experts, webcrawling represents a relatively inexpensive and promising alternative. This, however may introduce problems such as inconsistent labelings across several annotators, unreliable and wrong labels. Thus, special filtering methods need to be applied to obtain robust labels. A drawback of pre-labeling is that once annotated labels cannot be refined during the learning process, e.g., according to the interpretation of the data by a user.

Incremental Learning and Active Learning An alternative to pre-labeling is the use of incremental and on-line learning methods that enable the continuous refinement of a previously learned model as new labeled data becomes available [LHW17]. AL is special type of incremental ML that explicitly incorporates user knowledge into the learning process [Set09]. In AL, an algorithmic model pro-actively asks the user (referred to as the oracle) for

feedback during the learning process (usually in the form of labels) to iteratively improve the learned model. Since user interactions are time-consuming and thus expensive, AL aims at minimizing 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. The general AL workflow is illustrated in Figure 1. To find the potentially most useful unlabeled instances for learning, different strategies for candidate selection have been introduced which are discussed in detail in a number of surveys on AL [Set09, Ols09, TVC*11, WH11]. These strategies can be partitioned into five groups: (i) uncertainty sampling, (ii) query by committee, (ii) error reduction schemes, (iv) relevance-based selection, and (v) purely data-centered strategies.

Active Learning Strategies *Uncertainty sampling* aims at finding the instances that the learner is most uncertain or unsure about. Widely used strategies search for those instances near decision boundaries of margin-based classifiers [WKBD06] (large-margin based AL) [TVC*11]) or for instances with high entropy of class probabilities [SC08]. *Query by Committee (QBC)* [SOS92] strategies measure the uncertainty of a committee of classifiers. Instances are considered interesting when the committee disagrees with respect to their labeling [Mam98]. *Error reduction* strategies select those instances which may change the underlying classification model most. Selection criteria are expected model change [SCR08], risk reduction [QHR*09], or variance reduction [HJL06] *Relevance-based* strategies [VPS*02] focus on instances which have the highest probability to be relevant for a certain class. This strategy fosters the identification of positive examples for a class instead of difficult to decide instances which is, e.g., useful for ranking search results [WH11]. *Data-driven* strategies work independent of the learning model. Many techniques work with a density-based selection criterion which is a promising strategy for initiating an AL process in the case when no labels are available at all (cold start problem) [API11]. In contrast, diversity-based criteria foster the selection of dissimilar instances for labeling to increase the information gain for the learner [DRH06].

Benefits and Limitations of Active Learning Incremental learning methods like AL are especially useful in cases where large portions of the data are unlabeled, where manual labeling is expensive, and in cases of on-line learning where new unlabeled data needs to be processed continuously. One drawback of AL is that strategies are primarily driven by the learned model – users are not considered in the *identification and selection* of instances, but only in the labeling itself. Hence, the selection of instances is neither based on expert knowledge, nor humans’ ability to quickly identify patterns such as clusters and outliers. This leads to two limitations. First, the integration of the user needs, intentions, and tasks is impeded and second, the AL process does hardly support the generation of user knowledge. The VIAL process is more general and extends AL towards the user-based candidate selection by the integration VIS techniques [SG10, HNH*12].

2.2. Labeling in Interactive Visualization

In VIS, many labeling approaches assign users an active role, using visual interfaces as the means for candidate *identification and selection*. We assess a strong emphasis on knowledge generation.

However, in contrast to AL, VIS has gained little attention to the creation of labeled data. We structure related works in process models and techniques for visual-interactive labeling.

Process Models for Interactive Visualization We draw connections to both general as well as task-specific process models. The VIAL process is derived from general process models for VIS [CMS99, vW05] and VA [KAF*08, CG16], resembling the abstract data and interaction flow, as well as user-based knowledge generation [SSS*14]. However, abstract frameworks often do not reflect the characteristics of specific tasks [SSZ*16] such as labeling. The VIAL process extends these models for labeling tasks. In particular, with the VIAL process the output of the process is refined considerably. With *user knowledge*, *labeled data*, and *trained models* the VIAL process explicitly defines three types of output.

Recently, more specific process models have been proposed for different labeling tasks. Höferlin et al. [HNH*12] explicitly focus on interactive classification techniques, facilitated with AL strategies. The term ‘Interactive Learning’ is proposed, resembling the rationale to keep the user in the loop. We build upon the techniques employed in the pioneering approaches, and represent abstractions in our conceptual VIAL process. Bernard et al. propose a process where users play an active role in selecting data instances and assigning labels [BSB*15]. From this work we take the idea to support data-centered, model-centered, and user-centered criteria for label suggestion. A similarity-modeling approach for mixed data presents techniques for the interpretation of user feedback and discusses pitfalls for the design of labeling approaches [BSR*14]. Mamani et al. propose a visualization-assisted methodology for interacting with instances to transform feature spaces [MFNP13]. Finally, previous work included an empirical user study to compare the performance of either AL or VIS [BHZ*17]. Overarching insights were the complementing strengths of both approaches, suggesting the combination of both in a unified process. Although existing methodological approaches are inspiring, they are specific towards a data type, application goal, user group, or target variable. The VIAL process abstracts from concrete approaches. In addition, VIAL obtains a data-, model-, and user-centric perspective with three outputs: data, models, and knowledge.

Visual-Interactive Labeling Techniques We differentiate VIS techniques for labeling tasks by the type of the label to be defined by the user, such as, categorical labels, numerical labels, as well as labels that represent a relation between two instances. The type of the data to be categorized differs from video [HNH*12], text [HKBE12], time-oriented data content [SSBJ16, BDV*17], to two-dimensional data represented in scatterplots [BKSS14, SA15]. In many cases that accept categorical labels, classification models are learned in parallel or as a downstream task. Categorical labeling tasks differ in the number of groups to be assigned. Binary tasks accept yes/no or relevant/irrelevant statements while multi-labeled tasks consist of three or more classes. Numerical labeling tasks form a second type of labels. Users are enabled to assign fine grained interestingness or relevance scores, e.g., in the context of relevance feedback [SB97, SSJK16]. As an alternative, numerical labels can be used to characterize patient well-being with regression models [BSB*15]. Finally, we shed light on labeling approaches that allow the assignment of relations between instances,

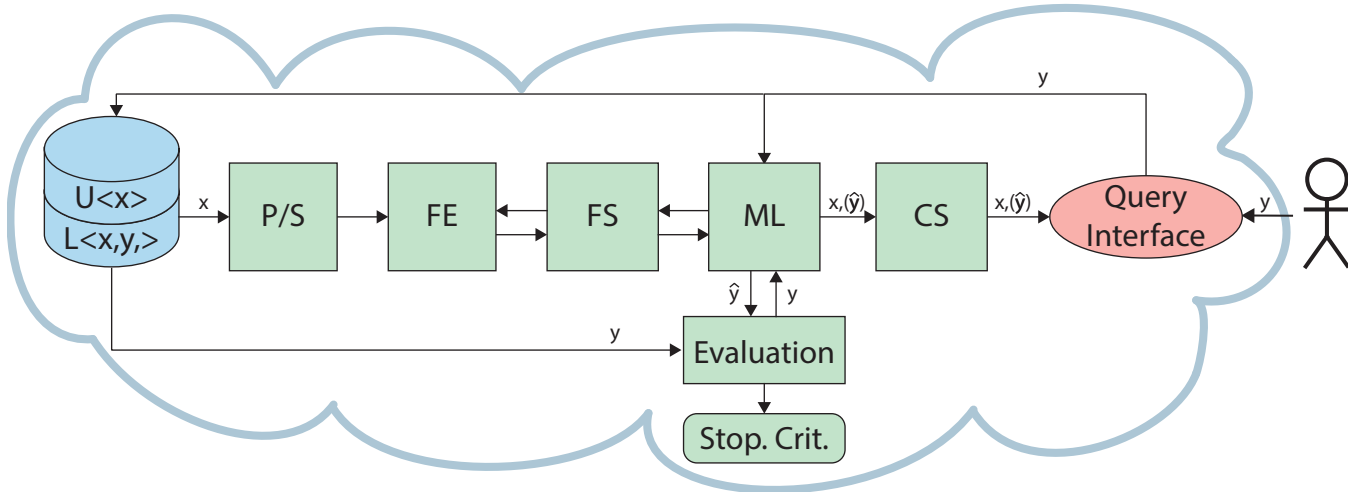


Figure 1: The abstracted AL process. A data source contains unlabeled (U) and labeled (L) data. Data is first preprocessed (e.g., normalized or segmented, “P/S”). Next, feature extraction “FE” and feature selection “FS” are applied to obtain a useful data representation for ML. A learning model “ML” is trained and evaluated iteratively. Candidate suggestion “CS” strategies query new labels y from the oracle which are used to iteratively adapt the ML until a stopping criterion is met [Ols09]. Users provide labels but are not involved in CS. In addition, knowledge generation by users is not part of the model.

e.g., in the sense of similarity scores. Using the example of similarity scores, related approaches aim to learn meaningful similarity (or more generally distance) functions by metric learning [YJ06] that reflect the relations expressed by the user [BRS*17]. Interaction designs such as the re-allocation of instances in a 2D manifold are used to enable users express complex relations between instances in the visual space [BLBC12, MFNP13, BSR*14].

3. The Visual-Interactive Labeling Process

Based on a review of related work in AL and VIS, we propose the VIAL process. We unify the main building blocks into an iterative process consisting of six steps shown in Figure 2. The VIAL process is special in its focus on exploratory instance selection and labeling tasks, as well as its emphasis on three output types, i.e., labeled data, trained learning models, and gained domain knowledge. All three outputs can be of equal importance, depending on the preference in a particular application scenario. As such the VIAL process provides a novel perspective on labeling tasks that can lead to novel innovative labeling approaches, as well as interfaces and tools for instance selection and labeling.

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 benefits and challenges that may emerge when the strengths of AL and VIS are combined in a unified process. In addition, we shed light on pioneer approaches that already address some of these challenges to ease the design of future VIAL solutions. We want to point out that the VIAL process is extensible towards more general problems in ML and VIS, such as visualization of complex parameter spaces originating, e.g., from deep neural networks, model visualization, progressive visual analytics, and user-centered design aspects. In section (Section 5), we discuss future directions of VIAL towards these general aspects.

3.1. Preprocessing and Feature Extraction

Preprocessing is a fundamental step in almost every data analysis approach that needs to be handled with care [KHP*11]. We refer to preprocessing as the cascade of operations that needs to be applied to ensure the usability of the input data by the models applied in later stages of the process. Preprocessing includes, e.g., the identification of erroneous (measurement) data, filtering noise and outliers, and dealing with missing data. We combine the preprocessing step with the mapping of real-world objects into more abstract representations (usually referred to as feature extraction). Existing labeling approaches either directly adopt semantically interpretable attributes of data instances as features (e.g., the age of a person, the GDP of a country), or apply complex derived descriptors [BYRN99] yielding abstract and often high-dimensional feature spaces, e.g., histograms of images or learned basis functions derived from the input signals. Processing and feature extraction are necessary to map the raw input data into a common (potentially high-dimensional) feature space on top of which a model (e.g., a classifier) can be trained. This abstraction is essential to enable (interactive) labeling approaches.

Challenges The analysis of raw data is associated with a series of challenges. In a numerical data set, a single outlying value that is erroneously represented with a multiple of 1.000 can cause serious problems for downstream models if not detected and cleansed appropriately. In general, many of these problems are associated with quality considerations [KHP*11]. Data preprocessing is required to derive secondary data, as a means to yield structured and curated data which is more appropriate for algorithmic models [SAAF17]. Using time series data as an example, taxonomies of dirty data [GGAM12] can foster the awareness for relevant quality challenges. Accordingly, these taxonomies can be used as a guideline to achieve cleansed data. Visual-interactive approaches




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). 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.

for preprocessing raw data [BRG*12] can support this step, e.g., in close collaboration of designers and domain experts. In a design study including similarity search for a retrieval system, the authors report on a cascade of ten preprocessing steps until data quality was achieved and the users’ information need was accomplished in a meaningful way [BDF*15]. Desirable goals for effective data preprocessing include achieving meaningful data representations, comparing raw data with model results, choosing appropriate model parameters, guaranteeing data quality, generalizability of preprocessing workflows, as well as involving users in the preprocessing step [Ber15].

A challenging design consideration in the feature extraction process 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] (third output in the VIAL process). The visualization of semantically interpretable features may be particularly beneficial for non-experts. Semantically interpretable features allow for the intuitive exploration of the feature space, ease the identification of important features for the learning success, may support knowledge generation, and may steer the learning into the intended direction [JSS*14, HNH*12]. The visualization of non-semantic features, however, such as Fourier or Cosine transform coefficients (of, e.g., images), is difficult to grasp even for experts. One possible drawback of visible features in a labeling approach is self-biasing. In a recent collaboration with medical experts, we observed how visible features attracted the attention of a physician in an unfavorable manner. As a result, the model to be learned solely reflected the characteristics of these particular features rather than the intrinsic structures of the data set reflected by all features [BSB*15].

Designing and learning useful feature representations that are representative and discriminative for the underlying data is a challenging task in general [MZB10, BCV13]. In the context of VIAL, high-dimensional feature representations such as those obtained from deep neural networks [LBH15] pose special challenges and are currently subject to intensive research [ZF14, WSW*18]. The integration of deep learning methods in the VIAL process will highly depend on future progress made in the visualization of such techniques as well as on the interplay of ML and VIS in general [SZS*16].

Due to the iterative nature, cases may exist where not the

model state but also learned feature representations may change in labeling process. Changing features may, however, confuse the user, which can be addressed with the VIS perspective. 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 [HNH*12]. This requires, however, that the features are semantically interpretable.

3.2. Learning Model

The choice of learning models primarily depends on the data and the labeling task at hand. Classifiers [HNH*12] support labeling tasks with binary or categorical labels. Regression models [MPI3] can be used to learn numerical target variables. In addition, the VIAL process includes other types of user-defined labels, such as similarity relations between two or more instances (i.e. metric learning) [BHS13, EFN12a]. In the VIAL process the learning model is directly coupled with visual interfaces facilitating analytic reasoning and model refinement [SSZ*16, ERT*17]. The VIAL process is iterative by nature. Thus it requires learning models that are instantly re-trainable (ideally in real-time) and that adapt their internal parameters to changes in the training data. Good candidates for VIAL are supervised incremental and on-line learning methods [LHW17].

The trained models represent a primary output of the VIAL process, building the basis for downstream applications.

Challenges Many learning models (i.e., their internal state) are difficult to visualize [Gle16, SSZ*17]. Traditional machine learning models such as, e.g., support vector machines, are defined in high-dimensional spaces with complex (potentially non-linear) decision boundaries. Neural networks challenge visualization with their millions of different internal parameters, their complex connections, and the absence of an explicitly given decision model [WSW*18].

Another important issue is to select a suitable termination criterion for learning and labeling to avoid unnecessary labeling effort. In a recent user experiment, we observed that many participants asked when the labeling will be finished [BHZ*17]. Visualization of the learning status may be one means to inform users about the progress and to avoid potential frustration. Examples are linecharts

of the accuracy of classifiers, or the representation of the variety of instances in the data set that have already been used for labeling [BSB*15, BHZ*17]. From a modeling perspective various alternative termination criteria can be used to inform the user in a visual way. 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 VA criterion is measuring quality aspects that help analysts to validate model convergence [Gle16]. From a ML perspective, a useful strategy for a termination criterion is to detect overfitting of the classifier. Overfitting means that the classifier adapts too much to the training data and thereby loses its ability to generalize to new data. Suitable measures help to avoid overfitting [SHK*14].

Another challenge is the choice of the learning model itself. A series of VA approaches addresses this problem, e.g., by the visual comparison of competitive feature sets and algorithmic models [MP13, KPB14]. In addition to the general problem of algorithm choice, we shed light on the challenge associated with complex learning functions. In fact, categorical labels are usually needed for supervised classifiers, and numerical labels often for regression-based approaches. However, the design space for learning models increases in situations where relations between instances $r(x_1, x_2)$ or distance (or similarity) functions $d(x_1, x_2)$ have to be modeled. An example is metric learning where distance relationships between instances are learned from the data directly [BHS13]. Research has shown that visualization is able to support this process [BLBC12, EFN12a].

3.3. Result Visualization

The visualization of model results corresponds to the VIS perspective on the labeling process. According to the VIS principles, result visualization is particularly suited to amplify knowledge generation [CMS99, SSS*14] as one output of the labeling process. Beneficial examples for the VIAL process include visual representations of classification results [SG10], regression models [MP13], or results of learned distance functions [BLBC12]. 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 [SB-VLK09, 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 We draw the connection to general challenges in the visual representation of high-dimensional data. Visual-interactive interfaces supporting overview and detail visualizations are one option to tackle this issue. Dimensionality reduction [SZS*16] and data aggregation techniques [EF10] help to condense the data, for the price of individual challenges. Examples for individual challenges include the applicability, quality, or uncertainty of algorithms in connection with their parameters. Using dimensionality reduction as an example, it cannot be guaranteed that intrinsic structures of the data is retained after the reduction and that new

structural patterns are introduced that are not present in the original space.

A particular design challenge for labeling approaches is whether and how predicted labels should be visualized. A recent user experiment showed that users change their strategies in selecting instances considerably when additional information about predicted models are depicted [BHZ*17]. Showing unlabeled instances with labels predicted by the model supports the comparison of current situation with targeted label information, i.e., users have a means to improve the learning model in a visual-interactive way. Users can assess whether the learning model is able to explain the already labeled portion of the data. However, showing predicted labels may also cause biases. In contrast, representing instances without their predicted labels helps to focus on data-intrinsic properties such as structural information (e.g., patterns, clusters, or outliers) [BHZ*17]. In this connection, visual data exploration might support instance selection and labeling (see Figure 2).

A central role in VIS is the direct manipulation of learning models and the analysis of respective outcomes. To facilitate this goal, VA approaches such as parameter space analysis support [SHB*14], or techniques for the visual comparison [GAW*11] of different model outputs can be leveraged.

Another class of challenges addresses the visualization of the model state itself, e.g., for the identification of dependencies between hyper parameters defining the model and input variables. While the standard machine learning process does not require visualizations of the model state, VIAL can benefit from the awareness of knowledge between the built model, the data input and possibly the semantics of a given analysis scenario. The identification of shortcomings in the training (e.g., bias towards a class, misleading features, overfitting) is another beneficial but challenging goal to be addressed with model visualization. As an example, in practice the visualization of decision trees is one way to support both the awareness of dependencies between data and model, as well as the selection of meaningful labeling candidates [vdEvW11].

3.4. Candidate Suggestion

Automated candidate suggestion (as in AL) and the visualization of model results (from VIS, see Section 3.3) represent two complementary alternatives for the *identification and selection* of labeling candidates. From an AL perspective, users are queried in a model-centered way that is assumed to improve the model accuracy most [Set09]. In turn, in the VIS perspective the user is typically assigned an active role in the candidate selection process, e.g., by leveraging the gained knowledge about patterns in the data. With the VIAL process, we seek to join both perspectives and propose to either include AL-based strategies as guidance concepts in visual interfaces, or to leverage visual-interactive interfaces for the analysis, verification and steering of AL strategies.

Pioneer VIAL implementations [SG10, HNH*12, BSB*15] provide both options, and give an idea of the potential of combined candidate suggestion and selection strategies. However, approaches that actually combine both strategies to form a generic and unified candidate selection strategy are still missing. As such, research in hybrid approaches remains an open topic and a promising direction of future research.

Challenges A major challenge in candidate suggestion stems from the AL process, i.e., the selection of the most beneficial candidates for labeling. A rich set of techniques for candidate suggestion exists with differences in, e.g., the query strategy [Set09] or the computation costs [Set12] (capability for interactive execution [MPG*14]). 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. From a VIAL perspective the suitability of AL strategies for different user-, data-, and task characteristics remains an open issue.

One central aspect is the cold start problem in AL when no labeled instances are available at all. User-based strategies have shown to outperform AL in early phases of the labeling process, however, questions arise about break-even points and the trade-off between both strategies in the course of the remaining labeling process [BHZ*17]. In addition to the latter, pioneer user experiments started comparing AL with the user-based selection of candidates with respect to different factors such as multi-instance labeling, annotation time, quality measures, and selection strategies [MBS*11, Set11, SG10, LAdS12, BHZ*17]. It remains subject to future work to use the gained knowledge and design hybrid approaches for combined candidate selection.

Considering the need for very large labeled data sets a downstream challenge is the efficient inference of gathered label information to yet unlabeled instances and the exploitation of unlabeled data for the learning. This goal can be addressed by semi-supervised learning which is a branch of supervised learning that tries to improve performance by exploiting information from unlabeled instances [CSZ06]. In VIAL, semi-supervised learning could be integrated in two ways: (i) to learn more robust models by taking unlabeled instances into account and (ii) by inferring labels from labeled to unlabeled instances. Both aspects help to reduce the required user interactions in the VIAL process.

3.5. Labeling Interface

The goal of the labeling interface is to accept labels y from the user which can either be assigned to instances x directly or be used to characterize relations between instances. This data-centered output of the VIAL process can be used to enrich data sets, e.g., for the creation of ground truth data. Every time a user labels an instance in the labeling interface, the labeling loop can be triggered, possibly leading to an improved learning model (see Figure 2). This iterative user-in-the-loop 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 labeling in a meaningful way.

Challenges One challenge is the visual mapping of labeling candidates, i.e., the visual representation of instances in the visual space [CMS99]. In some cases the visual representation of instances is straightforward, e.g., for handwritten digits or other image data. However, other data types build upon complex data characteristics that are more difficult to visualize. Examples from literature include visual representations of unknown patient histories [BSB*15], abstracted features [KPB14], relations between clusters and metadata [BRS*12], or poses in human motion capture data sequences [BDV*17]. In fact, 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 [BSR*14] or images of soccer players [BRS*17]. In other cases where users already have an intrinsic knowledge of the labeling alphabet, e.g., object classes like cats and dogs, visual identifiers are also suitable. A recent VA approach using the example of personal image collections is PICTuReVis, allowing visual-interactive labeling (and classification) of personal image data [vdCvW17]. In more complex situations where, e.g., multimodal data or unknown instances (e.g., of a new class) have to be identified and labeled, visual identifiers become insufficient. More abstract representations and visualizations need to be designed to support decisions. Special interaction designs may be required to explore the instances in detail, e.g., in case of 3D objects or complex data objects with many attributes.

Another challenge relates to rather exploratory situations where users are not aware of all inherent phenomena of a data set. Labeling only small portions of instances users have knowledge about may remain uncovered areas of the data space. As a result, learning models may hardly reflect all characteristics of the data set. The visualization of uncertainty information of the learning model can be one means to amplify the users' ability to select meaningful candidates from a model perspective. A recent user experiment showed that visualizing intermediate class distributions predicted by a classifier supported the participants in the selection of useful instances for model refinement [BHZ*17].

Finally, the interaction design raises challenges in complex learning situations where labels are less distinct and exhibit complex semantics. Categorical and numerical labels can easily be submitted with straightforward text boxes, sliders, or a set of predefined labeling options. Although this works well for many cases more complex learning situations may exceed the limits of baseline interaction designs. For example, applications where models learn relations between more than two instances require more sophisticated interaction designs. One class of approaches where this problem was addressed uses 2D layouts where instances can be arranged spatially (spatialization [EFN12a]) to obtain user feedback. These examples exploit the user feedback to learn similarity relations of the data for dimensionality reduction, clustering, or visual-interactive similarity search [EHM*11, EFN12b, BLBC12, BSR*14, WN17].

3.6. 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., when users create topologies with multiple instances and the feedback interpreter exploits the relations between multiple instances [BLBC12, EHM*11]. Similarly, *implicit* user feedback is more difficult to interpret where user behavior is observed and conclusions are drawn from behavior without explicit queries. In general, meaningful feedback interpretation can facilitate the data-centered and model-centered outputs of the VIAL process.

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 text boxes and slider controls used for feedback submission. Two example interactions are 1) spatial re-arrangements of instances in 2D as described in Section 3.5 [BLBC12] and 2) ranking items in list-based interfaces [WDC*18]. To address this challenge, we draw the connection to *observation-level interaction*, a sensemaking technique based on semantic exploration of data [EHM*11, WN17]. The technique explicitly addresses the question how users can interactively express their reasoning on observations about instances instead of directly modifying parameters of analytical models. For user feedback based on re-arrangements of instances in 2D at least three different mechanisms exist [BSR*14], i.e., the absolute positions of instances, the relative positions (distances), as well as the vicinity of nearest neighbors. The authors postulate that the labeling interface should be designed in a way that it informs users about how their feedback is interpreted by the system. If the feedback concept is based on ranked items [WDC*18] pointwise, pairwise, or listwise approaches can be used to interpret the user feedback and improve the learning model [Liu09].

Challenges from the second perspective are related, but take the discussion deeper into human computer interaction. Mental models [Cra43] of users communicated through visual-interactive interfaces [LS10] open large spaces for interpretation and thus may deviate 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 to capture user feedback and to derive labels continuously. Another challenge associated with implicit feedback is the trust of labels generated from such processes. While user-generated labels usually form ground truth data for learning tasks, implicit feedback however may require additional uncertainty modeling.

4. Usage Scenarios

In Section 3, we presented the six phases of the VIAL process using relevant pioneer approaches as examples [SG10, HNH*12, HKBE12, BSB*15, BRS*17, BDV*17, WN17, BHZ*17]. In this section, we discuss two real-world VIAL approaches in detail which represent heterogeneous and complementary examples. This line of approach is inspired by Sacha et al. [SSS*14] using, e.g., Jigsaw [SGL08] as an explicit example. Our rationale is to demonstrate the range of possible applications, to shed light on solutions that actually address the described challenges, and to validate our process model and the interplay between AL and VIS.

To achieve high variability, the selected use cases differ in at least one aspect in each of the six phases on the VIAL process, i.e., (1) semantical vs. non semantical features, (2) label types and learning models, (3) visualizations of the model state, (4) AL model used for candidate suggestion, (5) the type of visual representation of instances, and (6) feedback interpretation and propagation. The first approach applies VIAL for the generation of categorical labels in the context of object detection in videos. The second approach applies VIAL for learning similarities between complex objects (soccer players). Both approaches are concrete instances of the VIAL process and address all the six steps. Note, however, that not every

VIAL step is exploited with the full complexity the general VIAL process allows. The gap that remains between the concrete examples and the general VIAL process are good indicators for future work.

4.1. Visual-Interactive Labeling for Video Classification

4.1.1. Introduction

In visual classification and recognition, classifiers are usually pre-trained on large amounts of annotated video data and are then applied to new data (e.g., face detection in surveillance applications). Pre-training is possible when the target object is known a priori and when sufficient training data for this object has been collected and annotated. Pre-training is, however, not possible for ad-hoc queries defined by the user because in general no pre-defined classifiers and labeled datasets exist a priori for such queries. Höferlin et al. present an interactive learning-based approach for the detection of user-specified objects in a video that allows the user to actively select instances for labeling and thereby to steer the learning process [HNH*12]. Additionally, the trained model is visualized by VA methods to provide feedback and to increase the trust of the user in the model. The authors demonstrate their method for the task of person identification in videos. An illustration of the approach is depicted in Figure 3. This work is an early example which resembles all steps of the VIAL process to a certain degree. For this reason, we select this work as a first usage scenario and discuss it from the perspective of the VIAL process.

4.1.2. The VIAL Process

(1) Preprocessing and Feature Extraction

The data employed by Höferlin et al. are surveillance videos from a multi camera tracking dataset. According to the authors no special pre-processing is performed on the videos and they are taken as they come from the dataset. The authors addressed the learning of representations by employing rectangle features which gained great popularity with the strongly influencing work of Viola and Jones on face detection and which are applicable to arbitrary objects [VJ04]. Rectangle features represent image filters composed of black and white areas. They act as basic edge and corner filters and thus can be directly interpreted visually by humans (see Figure 3(b) for examples). Thus, they are well-suited for a visualization of the internals of the learning model which is built on top of them in the course of the VIAL process.

(2) Learning Model The major challenge in model learning is to identify which combination of filters are most characteristic for a certain type of object. To find a suitable representation, the learning model selects rectangle features dynamically and trains a cascade of classifiers from these features. Each input instance is fed through the entire cascade. Only if it passes all classifiers in the cascade with positive prediction it is declared as a positive example. Otherwise it is rejected. This asymmetric scheme is necessary to account for the fact that there are by far more negative instances than positive instances in typical visual detection tasks.

For model learning, the authors employ a slightly modified version of the boosting algorithm proposed by [VJ04]. Adaptions to the training algorithm where necessary to enable fast iterative re-training of the algorithm (on-line Adaboost) which is particularly

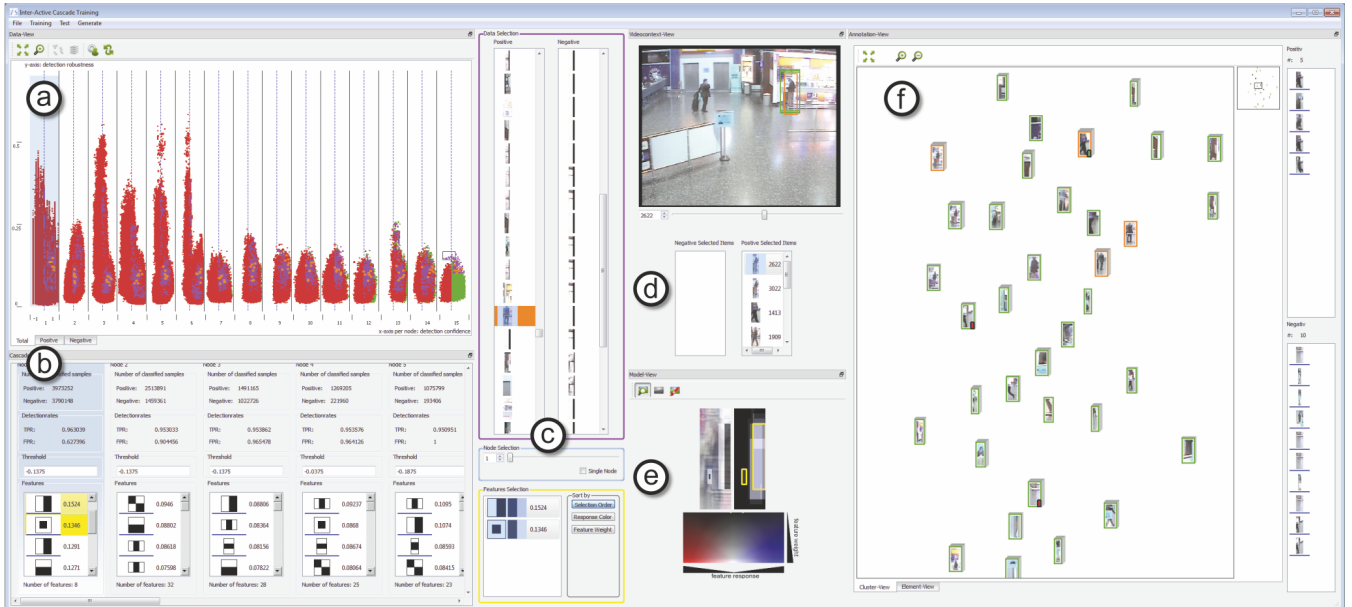


Figure 3: A screenshot of the VIAL approach by Höferlin et al. [H^{NH}*12]. The interface is split into several interlinked areas. (a) the scatterplots for each strong classifier of the cascade; (b) detail view of each strong classifier including the respective features selected by the classifier; (c) selected instances by the user (from the scatterplots); (d) visualization of selected instances in the original video frame (context); (e) visualization of rectangle features on selected instances; (f) a 2D view for instance labeling. Reprint with permission of IEEE.

important in VIAL to provide instant feedback to the user. Thereby, the authors solved the first challenge related to the learning model (see Section 3.2). The second challenge is the specification of a stopping criterion. In the approach the point in time when training should be terminated is defined by the user. To support the user in this decision, performance metrics such as true positive rate (TPR) and false positive rate (FPR) are provided (see Figure 3(b)). Although this is a sound approach it has the drawback that the user needs a certain knowledge of the performance metrics and some basic principles of ML. The use of visually interpretable features, their visualization, and the possibility to provide relevance feedback on them to guide the learning process (see below) represent basic concepts of the VIAL process.

(3) Result Visualization

Challenges addressed in the context of result visualization are the visual representation of the model state, its learned representation, and designing visual techniques that enable the user to guide the learning process. Höferlin et al. design several visualizations to show the internal workings of the machine learner. First, for each classifier in the cascade a scatterplot is shown that displays the decision boundary and the positively and negatively predicted instances (left and right of the boundary, see Figure 3(a)). From these scatter plots single or even multiple instances can be selected. The selected items are visualized in Figure 3(c). Note that this visualization does not provide the true label of already annotated instances. Adding this information could further improve the guidance in the selection and exploration process. In a second view (Figure 3(b)), the selected rectangle features (image filters) are displayed for each strong classifier together with their importance. This view helps the

user to inspect which types of image structures (e.g., vertical lines, horizontal lines, corners) the classifier relies on and if these patterns also reflect what the user expects to be important structures for the searched for objects (inclusion of domain knowledge). To further verify what the classifiers learn, the user can overlay the rectangular features over selected instances, see Figure 3(e) and provide relevance feedback to the learner.

(4) Candidate Suggestion

The major VIAL challenge that Höferlin et al. address in candidate suggestion step is the automated selection of the most important instances to be labeled. For this purpose [H^{NH}*12] implement an AL strategy based on uncertainty as an alternative to user-driven candidate selection. The authors compute the uncertainty of the learner for each instance and present the most uncertain instances to the user for annotation [GB06, VSF08]. This process is initiated only on demand by the user and is thus an optional step.

In the concrete example the selected instances are directly shown to the user (see Figure 3(f)). An interesting extension following the idea of VIAL would be to highlight these instances also in the result visualization, i.e., in the scatterplots of each classifier. Furthermore, visualizing the most uncertain instances in the (projected) feature space or relative to the decision boundary would allow the identification of ambiguous features which may lead to wrong decisions. The latter aspects indicate that the scope of the VIAL process goes beyond the implementation of in the usage scenario. VIAL aims at a stronger integration of AL and VIS, e.g., by the combination of AL-based and user-based strategies for candidate selection. This has hardly been investigated so far and is a challenging topic for future research [BHZ*17].

(5) Labeling Interface

The primary VIAL challenge addressed in the context of the labeling interface is the visual mapping of candidate instances and their arrangement. The authors visualize the instances for labeling in 2D in a similarity-preserving way (see Figure 3(f)) to support the exploration of the space. Similar objects are grouped into clusters and can be labeled simultaneously which saves time. Thumbnails are used for the visual mapping, which is natural as the instances represent images. For each thumbnail the predicted class is shown as well as the true class if instances have already been labeled. To provide contextual information, selected instances are further shown in the context of the video frame they stem from, see Figure 3(d).

The interface enables the quick tagging of instances. The approach could be extended with an encoding of uncertain instances where the true label is different from the predicted one. Such items are particularly interesting as they may indicate that the learning model is not able to capture the complexity of the data (e.g., too simple decision boundary, underfitting). An important problem emphasized in the VIAL process and solved by the particular labeling interface is the cold start problem, i.e., the initialization of the learner when no labeled instances exist at all. In the concrete example, the user is simply asked to query positive and negative instances (used to build an initial model).

(6) Feedback Interpretation

The last step to complete the VIAL loop is feedback interpretation. In the approach, feedback interpretation is rather straightforward as (i) user feedback is explicit and (ii) feedback consists of predefined mutually exclusive categories (i.e., person or no person). Thus, the provided label is simply passed to the machine learner as provided by the user. We want to note, however, that even in this simple example there is a certain room for interpretation and a certain fuzziness. Should, e.g., an instance which shows a person only partially (e.g., one half or less) be considered a positive instance? And where to draw the border between positive and negative instance, i.e., how much of a person must be at least visible to consider an instance positive? There are in most labeling tasks such ambiguities which need to be treated in some way. VIAL provides a means to make such ambiguities explicit to the user and to raise awareness towards such issues.

4.1.3. Application Example

Höferlin et al. demonstrate the capabilities of their system on the task of person identification in videos. The VIAL process starts with completely unlabeled videos. In an initial step the user annotates persons in the videos (positive examples) by drawing rectangles. Next, negative examples are generated in the same way. After an initial training on these examples interactive labeling is performed to improve the result. It shows that after a few cycles a similar performance level can be reached than with pure AL (uncertainty sampling) but, however, with much less training examples. This indicates the potential of combining user-interaction and active learning in the VIAL process.

The approach generates all three outputs of the VIAL process. Labeled data is obtained in the form of positively and negatively labeled image regions. The learned model is achieved together with

a representation that can be re-used for the detection of the target object. The gained knowledge of the user is twofold. First the user gains insights into the inner workings of the machine learner. Second, the user gets knowledge about the application domain, its challenges (difficult cases), ambiguities (e.g. patterns which are easily confused with the target object), and misleading image structures which leads to a better understanding of the task.

4.2. Visual-Interactive Labeling for Similarity Modeling

4.2.1. Introduction

The second usage scenario describes a similarity learning approach applied for soccer players [BRS*17] (see Figure 4). Experts as well as non-experts can create a personal similarity function for soccer players in a visual-interactive way. In general, such similarity functions can subsequently be used for downstream algorithms to address retrieval, clustering, or other ML tasks. In this approach, soccer players are characterized by a series of attributes gathered from Wikipedia; all attributes are semantically interpretable. A visual-interactive interface allows the definition of similarity relations between pairs of soccer players. A similarity modeling component interprets the user feedback and builds a model that can be applied on the entire data set. A retrieval interface supports the visual validation of (intermediate) similarity models for soccer players. Overall, the approach resembles all six steps of the VIAL process.

4.2.2. The VIAL Process

(1) Preprocessing and Feature Extraction The data set of the approach consists of (European) soccer players represented by a set of attributes (name, position on the field, goals, national games, size, etc.). The feature vector used in the approach consists of primary attributes extracted from Wikipedia (DBpedia) and secondary attributes derived from primary attributes. As such, all features are semantically interpretable, even by non-experts. The approach addresses the challenge to deal with mixed data, i.e., the features can either be numerical, categorical, or binary. Challenges in the data preprocessing process were the acquisition of data, as well as dealing with missing values. The authors' design foresees an explicit visualization of existing features (attributes) in order to facilitate the visual comparison of player characteristics in detail, see Figure 4 (left). This design decision comes with the cost of self-biasing caused by a focus on individual attributes (see Section 3.1). Simple representations make the approach applicable even for non-experts, e.g., a straightforward glyph design depicts the position of players on a soccer field.

(2) Learning Model The function to be learned is $f(x_1, x_2) = y$ where y is a numerical value between 0 (dissimilar) and 1 (very similar), and x are the players. A two-step approach first correlates every feature of the data set against the user labels and then pairwise distances for all pairs of players in the data set can be calculated. Every attribute is correlated individually with different correlation measures for numerical, categorical, and binary features. Individual treatment of features is fast (allows rebuilding the model in real-time) and results are easily interpretable, i.e., a weight is calculated for every feature. However, the approach does not address dependencies of attributes between each other. Different distance measures for numerical, categorical, and binary features are

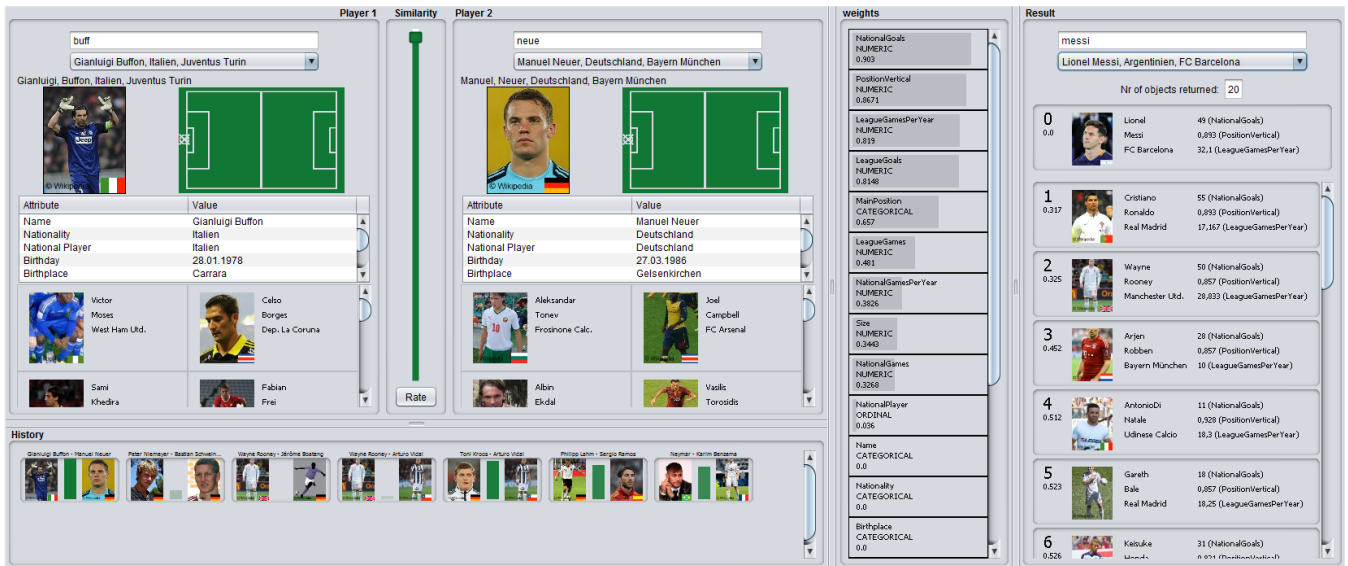


Figure 4: Visual-interactively labeling the similarity between pairs of soccer players [BRS* 17]. An active learning model suggests other players to be labeled next. The learned similarity function can be validated in a retrieval component. Most recently, the user has defined a 100% similarity between the two goal keepers Gianluigi “Gigi” Buffon and Manuel Neuer. At the bottom previously labeled players can be seen including their similarity scores (bars in green-to-gray). The retrieval algorithm was applied with Lionel Messi, according to the user-defined similarity function most similar soccers are Cristiano Ronaldo, Wayne Rooney, and Arjen Robben.

applied, until a global similarity value is calculated, successively. Differing weights of features indicate model changes and build the basis for an intrinsic and comprehensible termination criterion. The model-centered output of the approach is an algorithmic measure that reflects user-defined similarity.

(3) Result Visualization In line with the VIAL principles, the approach provides visual interfaces for both model visualization and model result validation. A straightforward list-based interface with barcharts depicts the weights of individual features, see Figure 4 (halfway right). In this way, users can analyze the fitness of every feature to the given similarity feedback. A retrieval interface applies the similarity function, and represents its results visually, see Figure 4 (right). Similarity scores for every retrieved player give an indication for the similarity relations of retrieval results with respect to a queried player. As such, users have a means to interact with the data, explore the data set, and gain knowledge, e.g., about previously unknown players. The visual retrieval result can also be used for the selection (and labeling) of new pairs of players. With the retrieval component, the authors omit the challenge to represent the data of the high-dimensional feature space in an abstract way. One drawback of the strategy is a missing overview of the similarity relations of the entire data set.

(4) Candidate Suggestion The approach combines the selection of candidates with AL and VIS, as recommended by VIAL process. The AL strategy builds upon a data-centered criterion, i.e., it exploits structure of the data [Set12]: suggested candidates have highest distances in the vector space to all labeled instances. A set of players is suggested and depicted in the list-based AL-interface. VIS principles are implemented in three different ways. First, a textual search allows querying players known to the users. Second, the

history can be used to re-select (an re-label) players, see Figure 4 (bottom left). Third, players retrieved with the retrieval component of the tool can be selected for labeling.

(5) Labeling Interface According to the VIAL process, the approach builds upon a highly iterative labeling strategy. The labeling interface visualizes two players, a slider control in between allows the assignment of a similarity score, see Figure 4 (left). A tabular interface depicts details about attributes (features) of every player as a means to grasp the characteristics of queried instances. The glyph with the position on the soccer field allows direct comparison of this (probably important) attribute. Images of soccer players serve as visual identifiers of the players. At least experienced users are able to lookup known players without difficulties.

(6) Feedback Interpretation With the slider control users can explicitly define the similarity score between two players. However, the complexity of the approach is the similarity relation in itself ($f(x_1, x_2) = y$). The authors address this problem with correlation measures used to relate every feature to the given similarity feedback. Features with high weights match the similarity characteristics of the set of labeled players. Pairs of players in combination with assigned similarity scores define the data-centered output of the approach.

4.2.3. Application Example

In Figure 4 the approach can be seen in practice. The history of labeled pairs of players at the bottom reveals that 7 pairs have been labeled yet. Most recently, two of the best goal keepers in the world were labeled with a 100% similarity (green bar in the labeling interface). The weights of the similarity model (list-based bar visualization center-right) indicate that the notion of similarity of the user is

mainly based on national goals, the (vertical) position on the field, as well as the numbers of league games and goals per year. According to the learned similarity function, the soccer player most similar to the queried player Lionel Messi is Cristiano Ronaldo.

5. 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, but also for challenges and existing solutions. In addition, we presented explicit and holistic reflections of two application examples in the context of VIAL. Future work includes implementations of the VIAL process in application examples with the aim for a broad range of real-world settings. In this connection, collaborations with experts seems to be a promising approach, allowing the application of principles from user-centered design and design studies. Another line of future work involves points of reference such as deep learning or progressive visual analytics to be included in the VIAL process more explicitly. Finally, future work includes evaluations of the challenges described in the six steps.

5.1. Deep Learning

Deep learning has gained a lot of attention recently. We discuss possible connection points to the VIAL process. Neural networks in general are suited for VIAL approaches as they can be trained and refined iteratively. However, a number of limitations may hamper the use of deep networks in the VIAL process.

First, training and refinement usually takes considerable time which makes real-time training hardly feasible today. The capability of algorithms for interactive execution [MPG*14] is a challenge for many ML approaches in general, so it is for the VIAL process. Second, large numbers of labeled data are necessary for initial training which strongly conflicts with the cold start problem that is present in many situations. This requirement does not meet the scope of the VIAL process, rather deep learning may be a well-suited downstream process when training data are available as a result of the VIAL process. In addition, a solution for the latter problem may be *transfer learning* where existing models trained for similar tasks are fine-tuned incrementally instead of training from scratch from the given data [PY10, Ben12, YCBL14]. This approach saves training time considerably and could be a first step towards the integration of deep learning into VIAL. Transfer learning requires that a compatible model already exists, such as the model-based output of the VIAL process. Third, a big and still open challenge in the context of VIAL is how to visualize the *representations* learned by the network and the *model* itself. Basic visualization techniques for representations learned from images have been introduced recently [SAB*17, YCN*15, ZF14, WSW*18]. However, the challenge remains of how to visualize thousands of such learned representations and their hierarchical relations. For more abstract and complex data (e.g., multimodal data) which has complex relations and is difficult to interpret visualization becomes significantly more challenging [RFFT17]. We thus conclude that an integration of deep learning into VIAL is at the current state of research

hardly feasible, however, research that tries to mitigate current limitations is currently ongoing.

5.2. Visualization of Active Learning Models/Strategies

Recent experiments have shown that candidate suggestion by a machine learner and candidate selection by the user may differ significantly and that user-based selection can outperform AL [BHZ*17]. One strength of VIAL is the combination of both strategies. This combination opens up a novel design space for combined selection and suggestion strategies. From the perspective of AL the major novelty lies in the visualization of active learning strategies. Such a visualization should, on the one hand, enable to validate a given strategies against the intuition of the user. On the other hand, it makes the learning process more transparent for the user and fosters a deeper understanding of the model and its actual “needs” (in terms of labeled instances). Furthermore, it bears a potential to improve existing or develop novel AL strategies by learning from the user and his or her selection strategies.

The type of AL strategy strongly influences the visualization approach. Purely data-driven approaches, such as density-based selection [WKBD06] can be visualized more easily because density is a concept well perceivable by humans. For relevance-based approaches [VPS*02], visual identifiers or glyphs (in case of more abstract data) may be good measures to visualize suggested instances and to let the user verify their actual relevance for a given category. For strategies that are driven by model-specific properties, such as error reduction schemes [QHR*09] and uncertainty sampling [SC08] but also for query-by-committee [SOS92] the visualization becomes more complex, because internals of the models need to be visualized to explain to the user why a certain instance is at the current point of learning the most important one for the learner. The development of visualization techniques for this purpose is a challenging direction for future work.

5.3. Progressive Visual Analytics

The VIAL process intrinsically relies on the smooth combination of iterative user input and computational models. With the growing complexity of the computational components, this approach needs to balance a critical tradeoff though. On the one hand, the user expects a fluent interaction with immediate feedback from the labeling system and the underlying computational models. Delays as small as seconds can substantially interfere with and interrupt the user’s workflow [CRM91]. Thus, fast and fluent responses from the system are crucial. On the other hand, however, increasing complexity of models often comes with increasing computational requirements, that is, longer run times. The user would thus need to wait for results longer.

To deal with this tradeoff, researchers have suggested to work on progressive visual analytics solutions [SPG14]. The basic idea is, instead of waiting for the final results of computational components, to quickly compute intermediate results, offer these to the user, and refine them over time [MPG*14, CPI3, TKBH17]. An interesting strategy, for instance, might be to train a classifier first on a subset of the original input data only (to keep response times short). Once trained, data from the remaining set can be used to refine the model. Restricting the amount of data used for a single training

pass could lead to a reduction of disturbing delays and thereby to a smoother user experience. Such approaches will be crucial for successful implementations of the VIAL process, specifically if the underlying models come with a high computational complexity.

5.4. User-Centered Design and Evaluation Strategies

The current state of pioneer VIAL approaches demonstrates the applicability of VIAL for various combinations of data sets, application fields, label types, and learning models. Most existing approaches have an emphasis on techniques, as well as on data and control workflows, proving the applicability of VIAL. One aspect with remaining potentials is designing VIAL approaches for specific user groups. Following the design study principle [SMM12], future approaches may emphasize methodology for the characterization of users and their application problem, an iterative design phase, and a careful reflection of lessons learned in the design process. Similarly, more general principles from user-centered design can be borrowed to facilitate both usefulness and usability.

Similarly, the evaluation of VIAL approaches is still in an early state. We observe that several existing works already validate learning models, as well as the model conversion process [BSR*14, SBKK14, BSB*14, BSB*15, BRS*17]. However, many other types of evaluation are possible. One strategy for future work may be the validation of all six steps proposed in the VIAL process. This principle helps to validate that individual modules provided in the interactive and iterative workflow work together seamlessly. Another promising direction for evaluation strategies is the assessment of strategies for candidate selection either performed by AL, by users, or by both. A recent experimental study identified that the strategies for the selection of candidates even differs for various users [BHZ*17]. This opens new evaluation strategies for the assessment of various user-based strategies, as well as in the corporation of analytical guidance.

6. 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. In the same way, the VIAL process mitigates the missing knowledge generation in AL as well as shortcomings in data creation observed in VIS. Based on a review of the AL process as well as pioneer VIS approaches for labeling, we identified six core steps assembled to the VIAL process. For every step, we described both the AL and the VIS perspective, discussed respective challenges, and outlined possible solutions. In two inspiring examples, we indicated the applicability of the VIAL process also for other application domains. Finally, we discussed four core aspects in the broader sphere of VIAL and indicated connecting points referred to as future work.

In summary, VIAL provides a novel perspective on labeling by bringing together two complementing approaches with a common goal. The VIAL process opens up a space for the design future labeling approaches that has not been explored so far. Especially the combined candidate selection and suggestion – which is the point where AL and VIS meet – represents a challenging direction of research which bears the potential for significant improvements of

existing labeling strategies. VIAL may lead to faster model convergence and more generalizable models in shorter time. In this connection, we amplified the potentials of combining the strengths of both humans and algorithmic models in a unified process, and outlined possible future work addressing this direction.

Acknowledgments

This work is an extended version of a previous EuroVA paper [BZSA17] entitled “A Unified Process for Visual-Interactive Labeling”. This work was supported by the Deutsche Forschungsgemeinschaft (DFG), Project No. I 2850, Lead Agency Verfahren (DACH) “Visual Segmentation and Labeling of Multivariate Time Series (VISSECT)”, the Austrian Research Promotion Agency (FFG), Project No. 7179681, No. 7189193, the Austrian Ministry for Transport, Innovation and Technology under the initiative “ICT of the future” via the project “VALiD” (Project No. 845598), the Austrian Research Fund (FWF) via the projects “KAVA-Time” (Project No. P25489-N23) and “VisOnFire” (Project No. P27975-NBL), as well as by the Lower Austrian Research and Education Company and the Provincial Government of Lower Austria (NFB), Department of Science and Research via the project “IntelIGait” (Project No. LSC14-005).

References

- [AP11] ATTENBERG J., PROVOST F.: Inactive learning?: Difficulties employing active learning in practice. *SIGKDD Explor. Newsl.* 12, 2 (Mar. 2011), 36–41. doi:10.1145/1964897.1964906. 3
- [BCV13] BENGIO Y., COURVILLE A., VINCENT P.: Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence* 35, 8 (2013), 1798–1828. 5
- [BDF*15] BERNARD J., DABERKOW D., FELLNER D., FISCHER K., KOEPLER O., KOHLHAMMER J., RUNNWERTH M., RUPPERT T., SCHRECK T., SENS I.: Visinfo: a digital library system for time series research data based on exploratory search—a user-centered design approach. *International Journal on Digital Libraries (IJoDL), Springer* 16, 1 (2015), 37–59. doi:10.1007/s00799-014-0134-y. 5
- [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. 3, 7, 8
- [Ben12] BENGIO Y.: Deep learning of representations for unsupervised and transfer learning. In *ICML Workshop on Unsupervised and Transfer Learning* (2012), pp. 17–36. 12
- [Ber15] BERNARD J.: *Exploratory search in time-oriented primary data*. Dissertation, phd, Technische Universität Darmstadt, Graphisch-Interaktive Systeme (GRIS), Darmstadt, Germany, 2015. URL: <http://tuprints.ulb.tu-darmstadt.de/5173/>. 5
- [BHS13] BELLET A., HABRARD A., SEBBAN M.: A survey on metric learning for feature vectors and structured data. *CoRR abs/1306.6709* (2013). URL: <http://arxiv.org/abs/1306.6709>, arXiv:1306.6709. 5, 6
- [BHZ*17] BERNARD J., HUTTER M., ZEPELZAUER M., FELLNER D., SEDLMAIR M.: Comparing visual-interactive labeling with active learning: An experimental study. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 24, 1 (2017). accepted for TVCG. doi:10.1109/TVCG.2017.2744818. 2, 3, 5, 6, 7, 8, 9, 12, 13
- [BKG11] BUHRMESTER M., KWANG T., GOSLING S. D.: Amazon’s mechanical turk: A new source of inexpensive, yet high-quality, data? *Perspectives on psychological science* 6, 1 (2011), 3–5. 2
- [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. 8

- [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. [3](#)
- [BLBC12] BROWN E. T., LIU J., BRODLEY C. E., CHANG R.: Disfunction: Learning distance functions interactively. In *IEEE Visual Analytics Science and Technology (VAST)* (2012), IEEE, pp. 83–92. [2](#), [4](#), [6](#), [7](#), [8](#)
- [BRG*12] BERNARD J., RUPPERT T., GOROLL O., MAY T., KOHLHAMMER J.: Visual-interactive preprocessing of time series data. In *SIGRAD, Swedish Chapter of Eurographics* (2012), vol. 81 of *Linköping Electronic Conference Proceedings*, Linköping University Electronic Press, pp. 39–48. URL: http://www.ep.liu.se/ecp_article/index.en.aspx?issue=081;article=006. [5](#)
- [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](#). [7](#)
- [BRS*17] BERNARD J., RITTER C., SESSLER D., ZEPELZAUER M., KOHLHAMMER J., FELLNER D.: Visual-interactive similarity search for complex objects by example of soccer player analysis. In *IVAPP, VISI-GRAPP* (2017), pp. 75–87. [doi:10.5220/0006116400750087](#). [1](#), [4](#), [7](#), [8](#), [10](#), [11](#), [13](#)
- [BSB*14] BERNARD J., SESSLER D., BERISCH M., HUTTER M., SCHRECK T., KOHLHAMMER J.: Towards a user-defined visual-interactive definition of similarity functions for mixed data. In *IEEE Visual Analytics Science and Technology (Poster Paper)* (2014). [doi:10.1109/VAST.2014.7042503](#). [13](#)
- [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](#), [5](#), [6](#), [7](#), [8](#), [13](#)
- [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). [2](#), [3](#), [4](#), [7](#), [8](#), [13](#)
- [BYRN99] BAEZA-YATES R. A., RIBEIRO-NETO B.: *Modern Information Retrieval*. Addison-Wesley Longman, 1999. [4](#)
- [BZSA17] BERNARD J., ZEPELZAUER M., SEDLMAIR M., AIGNER W.: A Unified Process for Visual-Interactive Labeling. In *EuroVis Workshop on Visual Analytics (EuroVA)* (2017), Sedlmair M., Tominski C., (Eds.), Eurographics. [doi:10.2312/eurova.20171123](#). [2](#), [13](#)
- [CG16] CHEN M., GOLAN A.: What may visualization processes optimize? *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 22, 12 (2016), 2619–2632. [doi:10.1109/TVCG.2015.2513410](#). [2](#), [3](#)
- [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](#), [3](#), [6](#), [7](#)
- [CP13] CHOO J., PARK H.: Customizing computational methods for visual analytics with big data. *IEEE Computer Graphics and Applications (CG&A)* 33, 4 (2013), 22–28. [12](#)
- [Cra43] CRAIK K. (Ed.): *The nature of explanation*. Cambridge University Press, Cambridge, UK, 1943. [8](#)
- [CRM91] CARD S. K., ROBERTSON G. G., MACKINLAY J. D.: The information visualizer, an information workspace. In *SIGCHI Conference on Human Factors in Computing Systems (CHI)* (1991), ACM, pp. 181–186. [doi:10.1145/108844.108874](#). [12](#)
- [CSZ06] CHAPPELLE O., SCHÖLKOPF B., ZIEN A.: Semi-supervised learning. adaptive computation and machine learning series, 2006. [7](#)
- [DRH06] DAGLI C. K., RAJARAM S., HUANG T. S.: Leveraging active learning for relevance feedback using an information theoretic diversity measure. In *Conference on Image and Video Retrieval* (2006), Springer, pp. 123–132. [doi:10.1007/11788034_13](#). [3](#)
- [EF10] ELMQVIST N., FEKETE J.-D.: Hierarchical aggregation for information visualization: Overview, techniques, and design guidelines. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 16, 3 (2010), 439–454. [doi:10.1109/TVCG.2009.84](#). [6](#)
- [EFN12a] ENDERT A., FIAUX P., NORTH C.: Semantic interaction for sensemaking: Inferring analytical reasoning for model steering. *IEEE Transactions on Visualization and Computer Graphics* 18, 12 (Dec 2012), 2879–2888. [doi:10.1109/TVCG.2012.260](#). [5](#), [6](#), [7](#)
- [EFN12b] ENDERT A., FIAUX P., NORTH C.: Semantic interaction for visual text analytics. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2012), CHI '12, ACM, pp. 473–482. [doi:10.1145/2207676.2207741](#). [7](#)
- [EHM*11] ENDERT A., HAN C., MAITI D., HOUSE L., LEMAN S., NORTH C.: Observation-level interaction with statistical models for visual analytics. In *2011 IEEE Conference on Visual Analytics Science and Technology (VAST)* (Oct 2011), pp. 121–130. [doi:10.1109/VAST.2011.6102449](#). [7](#), [8](#)
- [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. *Computer Graphics Forum (CGF)* (2017). [doi:10.1111/cgf.13092](#). [1](#), [5](#)
- [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](#). [6](#)
- [GB06] GRABNER H., BISCHOF H.: On-line boosting and vision. In *Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on* (2006), vol. 1, IEEE, pp. 260–267. [9](#)
- [GGAM12] GSCHWANDTNER T., GÄRTNER J., AIGNER W., MIKSCH S.: *A Taxonomy of Dirty Time-Oriented Data*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012, pp. 58–72. [doi:10.1007/978-3-642-32498-7_5](#). [4](#)
- [Gle16] GLEICHER M.: A framework for considering comprehensibility in modeling. *Big Data* 4, 2 (2016), 75–88. [doi:10.1089/big.2016.0007](#). [5](#), [6](#)
- [HJL06] HOI S. C., JIN R., LYU M. R.: Large-scale text categorization by batch mode active learning. In *World Wide Web* (2006), ACM, pp. 633–642. [doi:10.1145/1135777.1135870](#). [3](#)
- [HKBE12] HEIMERL F., KOCH S., BOSCH H., ERTL T.: Visual classifier training for text document retrieval. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 18, 12 (2012), 2839–2848. [3](#), [8](#)
- [HNH*12] HÖFERLIN B., NETZEL R., HÖFERLIN M., WEISKOPF D., HEIDEMANN G.: Inter-active learning of ad-hoc classifiers for video visual analytics. In *IEEE Visual Analytics Science and Technology (VAST)* (2012), IEEE, pp. 23–32. [doi:10.1109/VAST.2012.6400492](#). [1](#), [2](#), [3](#), [5](#), [6](#), [8](#), [9](#)
- [JSS*14] JANETZKO H., SACHA D., STEIN M., SCHRECK T., KEIM D. A., DEUSSEN O.: Feature-driven visual analytics of soccer data. In *2014 IEEE Conference on Visual Analytics Science and Technology (VAST)* (Oct 2014), pp. 13–22. [doi:10.1109/VAST.2014.7042477](#). [5](#)
- [KAF*08] KEIM D., ANDRIENKO G., FEKETE J.-D., GÖRG C., KOHLHAMMER J., MELANÇON G.: *Visual Analytics: Definition, Process, and Challenges*. Springer Berlin Heidelberg, 2008, pp. 154–175. [doi:10.1007/978-3-540-70956-5_7](#). [2](#), [3](#)
- [KHP*11] KANDEL S., HEER J., PLAISANT C., KENNEDY J., VAN HAM F., RICHE N. H., WEAVER C., LEE B., BRODBECK D., BUONO P.: Research directions in data wrangling: Visualizations and transformations for usable and credible data. *Information Visualization* 10, 4 (2011), 271–288. [doi:10.1177/1473871611415994](#). [4](#)
- [KM97] KARPINSKI M., MACINTYRE A.: Polynomial bounds for vc dimension of sigmoidal and general pfaffian neural networks. *J. of Comp.*

- & *Sys. Sc.* 54, 1 (1997), 169–176. doi:<http://dx.doi.org/10.1006/jcss.1997.1477>. 6
- [KPB14] KRAUSE J., PERER A., BERTINI E.: Infuse: Interactive feature selection for predictive modeling of high dimensional data. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 20, 12 (2014), 1614–1623. doi:[10.1109/TVCG.2014.2346482](https://doi.org/10.1109/TVCG.2014.2346482). 5, 6, 7
- [LAdS12] LEWIS J. M., ACKERMAN M., DE SA V. R.: Human cluster evaluation and formal quality measures: A comparative study. In *Annual Meeting of the Cognitive Science Society (CogSci)* (2012), pp. 1870–1875. 7
- [LBH15] LECUN Y., BENGIO Y., HINTON G.: Deep learning. *Nature* 521, 7553 (2015), 436–444. 2, 5
- [LHW17] LOSING V., HAMMER B., WERSING H.: Incremental on-line learning: A review and comparison of state of the art algorithms. *Neuro-computing* (2017). 1, 2, 5
- [Liu09] LIU T.-Y.: Learning to rank for information retrieval. *Found. Trends Inf. Retr.* 3, 3 (Mar. 2009), 225–331. doi:[10.1561/1500000016](https://doi.org/10.1561/1500000016). 8
- [LS10] LIU Z., STASKO J.: Mental models, visual reasoning and interaction in information visualization: A top-down perspective. *IEEE Transactions on Visualization and Computer Graphics* 16, 6 (2010), 999–1008. doi:[10.1109/TVCG.2010.177](https://doi.org/10.1109/TVCG.2010.177). 8
- [Mam98] MAMITSUKA N. A. H.: Query learning strategies using boosting and bagging. In *International Conference on Machine Learning (ICML)* (1998), vol. 1, Morgan Kaufmann Pub. 3
- [MBS*11] MÖHRMANN J., BERNSTEIN S., SCHLEGEL T., WERNER G., HEIDEMANN G.: Improving the usability of interfaces for the interactive semi-automatic labeling of large image data sets. In *Human Computer Interaction. Design and Development Approaches* (2011), Springer, pp. 618–627. 7
- [MFNP13] MAMANI G. M. H., FATORE F. M., NONATO L. G., PAULOVICH F. V.: User-driven feature space transformation. *Computer Graphics Forum (CGF)* 32, 3 (2013), 291–299. doi:[10.1111/cgf.12116](https://doi.org/10.1111/cgf.12116). 2, 3, 4
- [MP13] MÜHLBACHER T., PIRINGER H.: A partition-based framework for building and validating regression models. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 19, 12 (2013), 1962–1971. doi:[10.1109/TVCG.2013.125](https://doi.org/10.1109/TVCG.2013.125). 5, 6
- [MPG*14] MÜHLBACHER T., PIRINGER H., GRATZL S., SEDLMAIR M., STREIT M.: Opening the black box: Strategies for increased user involvement in existing algorithm implementations. *IEEE transactions on visualization and computer graphics* 20, 12 (2014), 1643–1652. 7, 12
- [MZB10] MITROVIĆ D., ZEPELZAUER M., BREITENEDER C.: Features for content-based audio retrieval. *Advances in computers* 78 (2010), 71–150. 5
- [Nor02] NORMAN D. A.: *The Design of Everyday Things*, reprint paperback ed. Basic Books, New York, 2002. 8
- [Ols09] OLSSON F.: *A literature survey of active machine learning in the context of natural language processing*. Tech. rep., Swedish Institute of Computer Science, 2009. 3, 4
- [PY10] PAN S. J., YANG Q.: A survey on transfer learning. *IEEE Transactions on knowledge and data engineering* 22, 10 (2010), 1345–1359. 12
- [QHR*09] QI G.-J., HUA X.-S., RUI Y., TANG J., ZHANG H.-J.: Two-dimensional multilabel active learning with an efficient online adaptation model for image classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* 31, 10 (2009), 1880–1897. doi:[10.1109/TPAMI.2008.218](https://doi.org/10.1109/TPAMI.2008.218). 3, 12
- [RFFT17] RAUBER P. E., FADEL S. G., FALCAO A. X., TELEA A. C.: Visualizing the hidden activity of artificial neural networks. *IEEE transactions on visualization and computer graphics* 23, 1 (2017), 101–110. 12
- [ROR11] RIEK L. D., O’CONNOR M. F., ROBINSON P.: Guess what? a game for affective annotation of video using crowd sourcing. In *International Conference on Affective Computing and Intelligent Interaction* (2011), Springer, pp. 277–285. 2
- [RTMF08] RUSSELL B. C., TORRALBA A., MURPHY K. P., FREEMAN W. T.: Labelme: a database and web-based tool for image annotation. *International journal of computer vision* 77, 1 (2008), 157–173. 2
- [SA15] SEDLMAIR M., AUPETIT M.: Data-driven evaluation of visual quality measures. *Computer Graphics Forum (CGF)* 34, 3 (2015), 201–210. doi:[10.1111/cgf.12632](https://doi.org/10.1111/cgf.12632). 3
- [SAAF17] SHURKHOVETSKYY G., ANDRIENKO N., ANDRIENKO G., FUCHS G.: Data abstraction for visualizing large time series. *Computer Graphics Forum (CGF)* (2017), n/a–n/a. doi:[10.1111/cgf.13237](https://doi.org/10.1111/cgf.13237). 4
- [SAB*17] SEIFERT C., AAMIR A., BALAGOPALAN A., JAIN D., SHARMA A., GROTTTEL S., GUMHOLD S.: *Visualizations of Deep Neural Networks in Computer Vision: A Survey*. Springer International Publishing, Cham, 2017, pp. 123–144. doi:[10.1007/978-3-319-54024-5_6](https://doi.org/10.1007/978-3-319-54024-5_6). 12
- [SB97] SALTON G., BUCKLEY C.: Improving retrieval performance by relevance feedback. *Readings in Information Retrieval* 24 (1997), 5. doi:[10.1002/\(SICI\)1097-4571\(199906\)41:4<288::AID-ASIS8>3.0.CO;2-H](https://doi.org/10.1002/(SICI)1097-4571(199906)41:4<288::AID-ASIS8>3.0.CO;2-H). 3
- [SBKK14] SESSLER D., BERNARD J., KUIJPER A., KOHLHAMMER J.: Adopting mental similarity notions of categorical data objects to algorithmic similarity functions, 2014. Poster Paper. URL: <http://www.vmv2014.gcc.tu-darmstadt.de/sites/program.html>. 13
- [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](https://doi.org/10.1057/ivs.2008.29). 6
- [SC08] SETTLES B., CRAVEN M.: An analysis of active learning strategies for sequence labeling tasks. In *Empirical Methods in Natural Language Processing* (2008), Computational Linguistics, pp. 1070–1079. 3, 12
- [SCR08] SETTLES B., CRAVEN M., RAY S.: Multiple-instance active learning. In *Advances in neural information processing systems* (2008), pp. 1289–1296. 3
- [Set09] SETTLES B.: *Active Learning Literature Survey*. Tech. Report 1648, Univ. of Wisconsin–Madison, 2009. 1, 2, 3, 6, 7
- [Set11] SETTLES B.: Closing the loop: Fast, interactive semi-supervised annotation with queries on features and instances. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)* (2011), Computational Linguistics, pp. 1467–1478. 7
- [Set12] SETTLES B.: Active learning. *Synthesis Lectures on Artif. Intell. and Machine Learning* 6, 1 (2012), 1–114. 7, 11
- [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](https://doi.org/10.1109/ICDMW.2010.181). 2, 3, 6, 7, 8
- [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](https://doi.org/10.1145/1466620.1466622). 8
- [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 (TVCG)* 20, 12 (2014), 2161–2170. doi:[10.1109/TVCG.2014.2346321](https://doi.org/10.1109/TVCG.2014.2346321). 6
- [SHK*14] SRIVASTAVA N., HINTON G. E., KRIZHEVSKY A., SUTSKEVER I., SALAKHUTDINOV R.: Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research* 15, 1 (2014), 1929–1958. 6
- [SMM12] SEDLMAIR M., MEYER M., MUNZNER T.: Design study methodology: Reflections from the trenches and the stacks. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 18, 12 (2012), 2431–2440. doi:[10.1109/TVCG.2012.213](https://doi.org/10.1109/TVCG.2012.213). 13

- [SOS92] SEUNG H. S., OPPER M., SOMPOLINSKY H.: Query by committee. In *Worksh. on Comput. Learning Theory (COLT)* (New York, NY, USA, 1992), ACM, pp. 287–294. doi:10.1145/130385.130417. 3, 12
- [SPG14] STOLPER C. D., PERER A., GOTZ D.: Progressive visual analytics: User-driven visual exploration of in-progress analytics. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (2014), 1653–1662. 12
- [SSBJ16] SARKAR A., SPOTT M., BLACKWELL A. F., JAMNIK M.: Visual discovery and model-driven explanation of time series patterns. In *Visual Languages and Human-Centric Computing (VL/HCC)* (2016), IEEE, pp. 78–86. doi:10.1109/VLHCC.2016.7739668. 3
- [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, 3
- [SSS*14] SACHA D., STOFFEL A., STOFFEL F., KWON B. C., ELLIS G. P., KEIM D. A.: Knowledge generation model for visual analytics. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 20, 12 (2014), 1604–1613. doi:10.1109/TVCG.2014.2346481. 2, 3, 6, 8
- [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 *Artificial Neural Networks, Computational Intelligence and Machine Learning* (2016). 1, 2, 3, 5
- [SSZ*17] SACHA D., SEDLMAIR M., ZHANG L., LEE J. A., PELTONEN J., WEISKOPF D., NORTH S. C., KEIM D. A.: What you see is what you can change: Human-centered machine learning by interactive visualization. *Neurocomputing* (2017). ISSN = 0925-2312. doi:10.1016/j.neucom.2017.01.105. 5
- [SSZ*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 Transactions on Visualization and Computer Graphics (TVCG)* 23, 01 (2016), 241–250. doi:10.1109/TVCG.2016.2598495. 5, 6, 7
- [TKBH17] TURKAY C., KAYA E., BALCISOY S., HAUSER H.: Designing progressive and interactive analytics processes for high-dimensional data analysis. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 23, 1 (2017), 131–140. 12
- [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. 3
- [VAD04] VON AHN L., DABBISH L.: Labeling images with a computer game. In *Conference on Human factors in computing systems (SIGCHI)* (2004), ACM, pp. 319–326. 2
- [Vap13] VAPNIK V.: *The nature of statistical learning theory*. Springer science & business media, 2013. 6
- [vdCvW17] VAN DER CORPUT P., VAN WIJK J. J.: Comparing personal image collections with picturevis. *Computer Graphics Forum (CGF)* 36, 3 (2017), 295–304. doi:10.1111/cgf.13188. 7
- [vdEvW11] VAN DEN ELZEN S., VAN WIJK J. J.: Baobabview: Interactive construction and analysis of decision trees. In *IEEE Visual Analytics Science and Technology (VAST)* (2011), pp. 151–160. doi:10.1109/VAST.2011.6102453. 6
- [VJ04] VIOLA P., JONES M. J.: Robust real-time face detection. *International journal of computer vision* 57, 2 (2004), 137–154. 8
- [VPS*02] VENDRIG J., PATRAS I., SNOEK C., WORRING M., DEN HARTOG J., RAAIJMAKERS S., VAN REST J., VAN LEEUWEN D. A.: Trec feature extraction by active learning. In *TREC* (2002). 3, 12
- [VSF08] VISENTINI I., SNIDARO L., FORESTI G. L.: On-line boosted cascade for object detection. In *Pattern Recognition, 2008. ICPR 2008. 19th International Conference on* (2008), IEEE, pp. 1–4. 9
- [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, 3
- [WDC*18] WALL E., DAS S., CHAWLA R., KALIDINDI B., BROWN E. T., ENDERT A.: Podium: Ranking data using mixed-initiative visual analytics. *IEEE Trans. Vis. Comput. Graph.* 24, 1 (2018), 288–297. 8
- [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. doi:10.1145/1899412.1899414. 3, 7
- [WKBD06] WU Y., KOZINTSEV I., BOUGUET J.-Y., DULONG C.: Sampling strategies for active learning in personal photo retrieval. In *IEEE International Conference on Multimedia and Expo* (2006), IEEE, pp. 529–532. doi:10.1109/TCME.2006.262442. 3, 12
- [WN17] WENSKOVITCH J., NORTH C.: Observation-level interaction with clustering and dimension reduction algorithms. In *Workshop on Human-In-the-Loop Data Analytics (HILDA)* (2017), ACM, pp. 14:1–14:6. doi:10.1145/3077257.3077259. 7, 8
- [WSW*18] WONGSUPHASAWAT K., SMILKOV D., WEXLER J., WILSON J., MANÉ D., FRITZ D., KRISHNAN D., VIÉGAS F. B., WATTENBERG M.: Visualizing dataflow graphs of deep learning models in tensorflow. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (Jan 2018), 1–12. doi:10.1109/TVCG.2017.2744878. 5, 12
- [YCBL14] YOSINSKI J., CLUNE J., BENGIO Y., LIPSON H.: How transferable are features in deep neural networks? In *Advances in Neural Information Processing Systems* (2014), Ghahramani Z., Welling M., Cortes C., Lawrence N. D., Weinberger K. Q., (Eds.), Curran Associates, Inc., pp. 3320–3328. 12
- [YCN*15] YOSINSKI J., CLUNE J., NGUYEN A., FUCHS T., LIPSON H.: Understanding neural networks through deep visualization. *arXiv preprint arXiv:1506.06579* (2015). 12
- [YJ06] YANG L., JIN R.: Distance metric learning: A comprehensive survey. *Michigan State University* 2, 2 (2006). 4
- [ZF14] ZEILER M. D., FERGUS R.: *Visualizing and Understanding Convolutional Networks*. Springer International Publishing, Cham, 2014, pp. 818–833. doi:10.1007/978-3-319-10590-1_53. 5, 12
- [ZK10] ZHU Q., KEOGH E. J.: Using captchas to index cultural artifacts. In *International Symposium on Advances in Intelligent Data Analysis IX* (2010), Springer, pp. 245–257. 2