





FeatureEnVi: Visual Analytics for Feature Engineering Using Stepwise Selection and Semi-Automatic Extraction Approaches

Angelos Chatzimpampas , *Student Member, IEEE*, Rafael M. Martins , *Member, IEEE CS*, Kostiantyn Kucher , *Member, IEEE CS*, and Andreas Kerren , *Senior Member, IEEE*

Abstract—The machine learning (ML) life cycle involves a series of iterative steps, from the effective gathering and preparation of the data—including complex feature engineering processes—to the presentation and improvement of results, with various algorithms to choose from in every step. Feature engineering in particular can be very beneficial for ML, leading to numerous improvements such as boosting the predictive results, decreasing computational times, reducing excessive noise, and increasing the transparency behind the decisions taken during the training. Despite that, while several visual analytics tools exist to monitor and control the different stages of the ML life cycle (especially those related to data and algorithms), feature engineering support remains inadequate. In this paper, we present FeatureEnVi, a visual analytics system specifically designed to assist with the feature engineering process. Our proposed system helps users to choose the most important feature, to transform the original features into powerful alternatives, and to experiment with different feature generation combinations. Additionally, data space slicing allows users to explore the impact of features on both local and global scales. FeatureEnVi utilizes multiple automatic feature selection techniques; furthermore, it visually guides users with statistical evidence about the influence of each feature (or subsets of features). The final outcome is the extraction of heavily engineered features, evaluated by multiple validation metrics. The usefulness and applicability of FeatureEnVi are demonstrated with two use cases and a case study. We also report feedback from interviews with two ML experts and a visualization researcher who assessed the effectiveness of our system.

Index Terms—Feature selection, feature extraction, feature engineering, machine learning, visual analytics, visualization

1 INTRODUCTION

IN machine learning (ML), *classification* is a type of supervised learning where the primary goal is to predict the dependent variable—also known as the target or class label—of every data instance (e.g., rows in a table) given independent features of the data (e.g., columns in a table). **Feature engineering** is the process of converting raw data into a set of features that better expresses the underlying problem, resulting in powerful ML models with enhanced predictive performance based on validation metrics [1]. In practice, most of the time spent in ML is in preparing this set of features, which should be as concise as possible while retaining vital information about the data set [2], [3]. If complications occur in this step and remain undetected, they can spoil the later phases of the ML pipeline, according to the classic “garbage in, garbage out” principle. Another important reason why feature engineering is essential in real-world problems is that it increases the transparency and trustworthiness of the data and, in consequence, the ML process in general [4]. This matter has been drawing attention recently with, for example, the new European General Data Protection Regulation (GDPR) instructions [5]. Furthermore,

domain experts are increasingly requesting clear evidence in order to trust in ML [6].

In general, feature engineering can be subdivided into four major processes: (a) *feature ideation*, (b) *feature generation*, (c) *feature transformation*, and (d) *feature selection* [7], [8]. Feature ideation is the process of coming up with entirely new features from the “raw” data. It is heavily subjective for most applications, except for text data [8], [9], for instance. As we work with numerical values and tabular data, the focus of our analytical approach presented in this paper is on the last three categories, which we describe next.

Feature generation is the action of creating more useful features from the combination of already existing ones. Typical interactions during this process are associated with the experimentation of: (1) *addition*, (2) *subtraction*, (3) *multiplication*, and (4) *division* between correlated features (all these operations are supported by our approach) [10]. ML experts and practitioners usually mix various features in a trial-and-error manner until the new features meet their expectations [2], [11]. The alternative is to use automatic feature generation techniques, but they are computationally expensive and may take a long time to produce an outcome [11], [12], [13]. Another threat in the validity of automatic feature generation is that the optimization is performed according to a single measurement, such as information gain, for example [2]. Additionally, features that look irrelevant in isolation may be informative in combination. Thus, one question that remains open is: (**RQ1**) for a given data set, which features should we compare, and how should we combine them to generate a new feature that boosts performance?

- Angelos Chatzimpampas and Rafael M. Martins are with the Department of Computer Science and Media Technology, Linnaeus University, 35195 Växjö, Sweden.
E-mail: {angelos.chatzimpampas, rafael.martins}@lnu.se.
- Kostiantyn Kucher and Andreas Kerren are with the Department of Computer Science and Media Technology, Linnaeus University, 35195 Växjö, Sweden and the Department of Science and Technology, Linköping University, 60233 Norrköping, Sweden.
E-mail: {kostiantyn.kucher, andreas.kerren}@{lnu, liu}.se.

Manuscript received August 29, 2021; revised December 09, 2021.

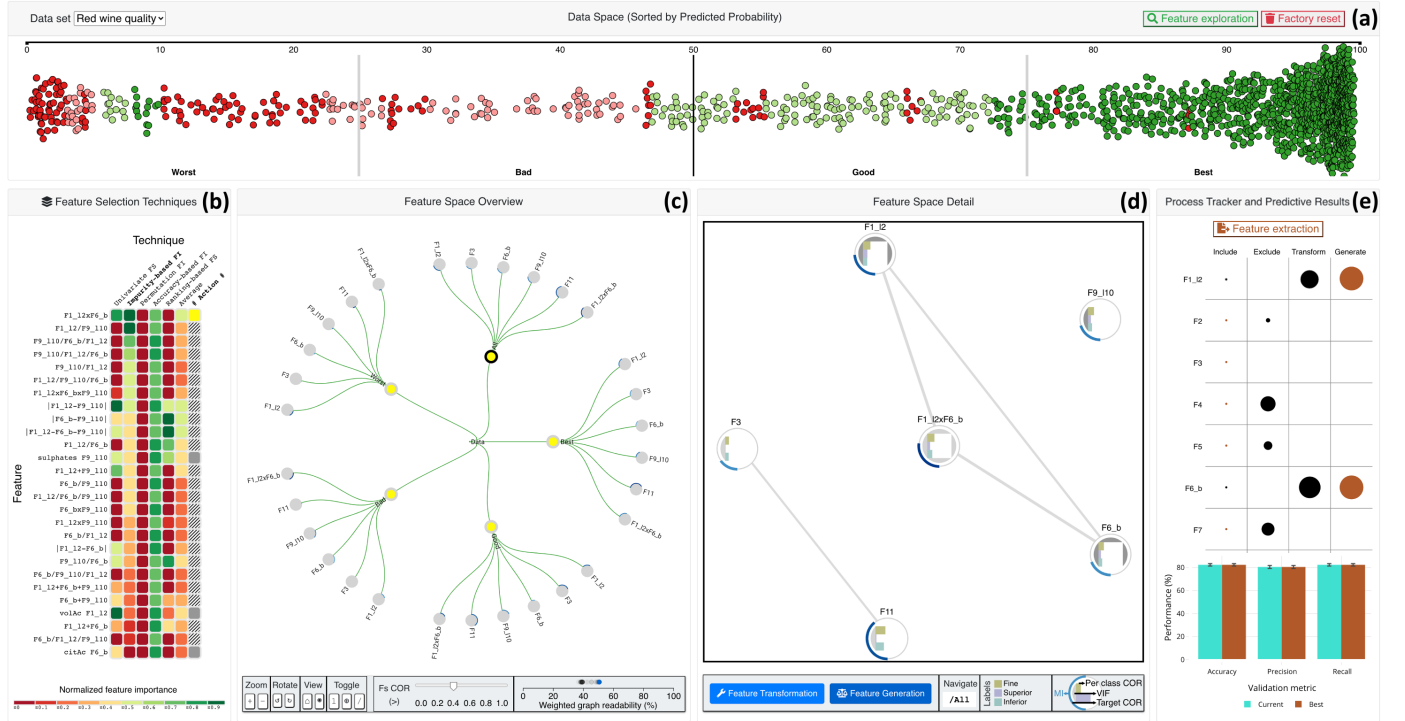


Fig. 1. Selecting important features, transforming them, and generating new features with FeatureEnVi: (a) the horizontal beeswarm plot for manually slicing the data space (which is sorted by predicted probabilities) and continuously checking the migration of data instances throughout the process; (b) the table heatmap view for the selection of features according to feature importances calculated from automatic techniques; (c) the radial tree providing an overview of the features with statistical measures for the different groups of instances, as set by the user-defined data slices; (d) the graph visualization for the detailed exploration of features, their transformation, and comparison between two or three features for feature generation purposes; and (e) the punchcard for tracking the steps of the process and the grouped bar chart for comparing the current vs. the best predictive performance based on three validation metrics.

Feature transformation usually denotes less sophisticated modifications over the features [14]. Some of the standard transformations also supported by our approach are: (1) *rounding*, (2) *binning*, (3) *scaling*, (4) *logarithmic transformations*, (5) *exponential transformations*, and (6) *power functions*. In this scenario, ML experts and practitioners usually perform exploratory data analysis (EDA), sometimes utilizing visualization of the data distribution to understand which transformation they should apply [15]. Another option is to use automated feature transformation approaches such as a transformation graph [1], which can, however, result in overfitting [2]. Unfortunately, such methods were only employed in regression and reinforcement learning problems [1]. This directs us to an additional open question: **(RQ2)** which features should we transform, and how can we understand their impact on the final outcome when using a specific data set?

Feature selection is about choosing a subset of features from the pool of features available by that time. Feature selection methods can be generally divided into four high-level categories: (1) *filter methods*, (2) *wrapper methods*, (3) *embedded methods*, and (4) *hybrid methods* [16], [17], [18]. Our feature selection strategy belongs to the last category, as we incorporate techniques from all the other categories. Also, instead of appending features progressively (called *forward selection*) or considering all features and then discarding some (known as *backward elimination*), we choose a *stepwise selection approach*. Therefore, we start with all features, but we can add or remove any number of features at different stages. For feature selection, the vast majority of the existing literature is on automatic feature selection techniques.

They may, however, lack transparency without the assistance of visualizations [16], [17], [19], [20], [21]. Furthermore, there is an opportunity to select features from a candidate set, which can be time-consuming if this set is large [22], [23], [24]. Even though a series of analytical tools and systems have been developed to address such challenges with the use of visualizations [25], [26], [27], a remaining open question is: **(RQ3)** which feature selection technique should we follow when they present diverging results, and how can we verify their effectiveness for particular problems?

The complex nature of feature engineering, occasionally declared as “black art” [2], [28], motivated us to concentrate our effort on addressing the three research questions mentioned above. In this paper, we present a visual analytics (VA) system, called *FeatureEnVi* (**Feature Engineering Visualization**, as seen in Fig. 1), that utilizes stepwise selection and semi-automatic feature extraction approaches for the feature engineering process of a state-of-the-art ensemble learning algorithm known as *XGBoost* [29]. FeatureEnVi uses validation metrics for multi-class classification problems, and includes three core iterative feature engineering phases that are intertwined and monitored by automatic techniques and statistical measures (see the details in Section 4). To provide answers to the three research questions **(RQ1-RQ3)**, FeatureEnVi adopts the following **workflow** (cf. Fig. 2 present in Section 4): (i) the initialization of the feature exploration after the user has split the data into four distinct slices sorted by predicted probabilities, (ii) the feature selection subprocess with five automated techniques computed for the choice of features, (iii) the feature transformation that enables the

user to convert features to impactful alternatives while utilizing statistical measures, (iv) the feature generation for the creation of new features from the combination of already existing features, and (v) the tracking of feature changes and analysis of the previous phases using visualizations, while the performance of the current solution is compared to the best result achieved so far. Overall, our contributions consist of the following:

- the definition of a unified visual analytic workflow for feature engineering, fusing stepwise selection and semi-automatic extraction approaches;
- a prototype implementation of the proposed workflow, in the form of our VA system, called FeatureEnVi, that consists of a novel combination of multiple coordinated views to support the complete process of engineering a data set's features;
- a proof-of-concept of the applicability of our proposed system with two use cases and a case study, using real-world data, that verifies the usefulness of our choice to deploy multiple statistical measures and include the human in-between automated methods; and
- the discussion of the methodology and results of interview sessions with ML experts and a visualization researcher, presenting encouraging results.

The rest of this paper is organized as follows. In Section 2, we review automatic feature generation methods, then we continue with feature transformations, and finally, automated and visually-assisted selection of subsets of features. Afterwards, in Section 3, we describe the analytical tasks and design goals for applying VA for feature engineering, and we highlight the necessity for multiple statistical measures. Section 4 presents the system's functionalities and simultaneously describes the first use case to remove unnecessary features, transform influential features to improve the performance, and generate new features for a further increase of the predictive score using physicochemical data; it also includes the second case, in which we demonstrate the applicability and usefulness of FeatureEnVi with a real-world data set focusing on vehicle recognition. In Section 5, we describe a case study carried out with an ML expert, focusing on improving classification for the molecules biodegradation problem. Next, in Section 6, we analyze the feedback on our VA system acquired from the expert interview sessions, including a discussion of limitations that lead us to potential future directions for FeatureEnVi. Subsequently, in Section 7, we discuss alternative design choices for our visual representations as well as additional limitations and opportunities for future work. Finally, Section 8 concludes our paper.

2 RELATED WORK

Several VA systems have been developed to explore and select subsets of features with the help of visualization. Finding which features to transform, and how, together with generating new features from different combinations, are some of the core phases that lack attention from the InfoVis/VA communities. This section reviews prior work on feature generation, feature transformation, and feature selection techniques for both automated and visually-assisted feature engineering. We also discuss the differences of such tools compared to FeatureEnVi to emphasize our system's novelty. To the best of our knowledge, there is no literature explaining the use of VA for the complete feature engineering process (as specified in Section 1) along with the partitioning of the data space based on the extracted predicted probabilities.

TABLE 1
Capabilities Comparison of FeatureEnVi With Other Tools

Tool	Feature Engineering			
	Ideation	Generation	Transformation	Selection
FeatureEnVi		✓	✓	✓
Mühlbacher and Piringer [15]			✓	✓
FeatureInsight [8]	✓			(✓)
Flock [9]	✓			(✓)
Liu et al. [30]		(✓)		(✓)
Other works in Section 2.3				✓

The tick in parentheses (✓) indicates partial support of a capability.

Additionally, a summary of the tools discussed in this section and comparison of their capabilities with FeatureEnVi are presented in Table 1. A tick indicates that the tool has the corresponding capabilities, while a tick in parentheses means the tool offers implicit support (i.e., it could be done manually, in an ad hoc manner, but is not explicitly supported). The table does not include works that do not contain a concrete visualization tool as their research contribution, as in the work by Khurana et al. [1], for instance. This table thus highlights the similarities/dissimilarities of FeatureEnVi against previously developed VA tools. From Table 1, it is evident that no other VA tool offers all functionalities of our system and is directly comparable to FeatureEnVi.

2.1 Feature Generation

As mentioned in Section 1, many studies confirm the importance of high-quality features in ML (e.g., the works by Domingos [2] or Yang and Pedersen [22]). Practitioners often spend a substantial amount of time experimenting with custom combinations of features or utilizing algorithmic feature generation [2], [11]. In several application domains, algorithms have been useful for creating high-level features from data elements related to speech [13], real-time sensors [31], and general ML problems [12]. However, there is no single metric that is consistently suitable across all data sets. Therefore, certain levels of involvement and manual effort are still required from analysts, and they receive little to no guidance during this complicated procedure.

A use case present in a visual diagnosis tool revealed that feature generation involving the combination of two features is capable of a slight increase in performance [30]. The authors tested the same mathematical operations as in our system (i.e., addition, subtraction, multiplication, and division), but the generation was performed manually by the analysts. Also, the decision for this action was based solely upon the similarity in those features' distributions [30]. In FeatureEnVi, determining which features to match during feature generation is achieved by analyzing linear and nonlinear relations present in the data. For the former, one of the most well-known approaches is *Pearson's correlation coefficient* between features and with the target variable [32], [33]. For the latter, *mutual information* is used in our VA system (also used by May et al. [26], for instance). Features are added to capture the missing information and improve the classifier's performance [34]. The magnitude of correlation with the dependent variable and in-between features is key to such decisions [35], [36]. However, the aforementioned VA tools work with regression problems and only support feature selection.

2.2 Feature Transformation

Automatic feature transformation has been examined within the ML community with positive results in reinforcement learning. In the work by Khurana et al. [1], the authors conduct a performance-driven exploration of a transformation graph which systematically enumerates the space of given options. A single “best” measurement is not possible, however, since the options might conflict with each other. In contrast, FeatureEnVi focuses on classification problems and presents users with various statistics about each feature in four different slices of the data space (the ones considered in Section 2.1, along with *variance influence factor* and *per class correlation*). This is similar to ExplainExplore [37] for classification and HyperMoVal [38] for regression. The authors of these tools work with the probabilities of instances to belong in different classes or clusters depending on the features. In our case, we take into account the ground truth values of the target variable and compute the probability (0% to 100%) of the ML model to classify each instance correctly.

A VA system for regression analysis has been proposed by Mühlbacher and Piringer [15]. The system is more similar to our work as it also incorporates feature transformation in its design (specifically logarithmic, exponential, and power functions). The main difference between this work and ours is our focus on classification rather than regression models. Moreover, they compared only the distribution of data values per feature to choose which transformation is better. In FeatureEnVi, we use multiple statistical measures to understand the relation of features with their transformed counterparts.

2.3 Feature Selection

There is a rather large body of existing work on automatic feature selection techniques [16], [17], [19]. However, one limitation is that features can be redundant if there is a strong correlation among them, and the correlation coefficient is unable to characterize nonlinear relationships. Thus, this is a problem where the feature selection techniques struggle to find a solution because of multiple parameters they have to optimize simultaneously. Guyon and Elisseeff [16] performed a survey including an extensive description of automatic feature selection pitfalls. The authors stress the general problem of finding the smallest possible subset of features for a given data set. They suggest that an automated method cannot be expected to find the best feature subset in all cases by itself. Other methods that face the same challenge are wrappers that use regression or classification models to find an ideal feature subset by iteratively including or excluding features. The combination of learning models (e.g., SVM [39]) and wrapper methods (e.g., RFE [40]) is a commonly used approach for automatic feature selection [20], [21]. Also, metric-based ranking followed by the selection of the k best features [23], [24] and more complex metrics—such as those used in genetic algorithms [41]—have been examined in the past. But they suffer from the same issues as described before. In our VA system, we implement several alternative feature selection techniques belonging to different types, and we allow users to decide if their aggregation is ideal or they want to focus on one of them.

Various visualization techniques have been proposed for the task of feature selection, including correlation matrices [42], [43], radial visualizations [44], [45], [46], scatterplots [47], scatterplot matrices [48], feature ranking [49], [50], [51], [52], [53], [54], [55], [56], feature clustering [57], and dimensionality reduction

(DR) [53], [58], [59]. The category of techniques more related to our work is *feature ranking*, since we use automatic feature selection techniques to rank the importance of the different features. For example, a VA tool called INFUSE [50] was designed to aid users in understanding how features are being ranked by the automated feature selection techniques. It presents an aggregated view of results produced by automatic techniques, assisting the user in learning how these work and compare their results with multiple algorithms. Similarly, Klemm et al. [60] propose an approach that performs regression analysis exhaustively between independent features and the target class. These approaches take into account the user’s ability to identify patterns from analyzing the data (e.g., with the colors in a heatmap representation) or choose the feature subset by some quantitative metric. A few other VA systems have leveraged a balanced blending between automatic and visual feature selection techniques. RegressionExplorer [61] is one example for examining logistic regression models. Additionally, the exploration of linear relationships among features was studied by Barlowe et al. [62]. FeatureEnVi offers rather similar characteristics with the tools analyzed above. However, we combine several automatic feature selection techniques and statistical heuristics in cohesive visualizations for evaluating feature selection and feature extraction concurrently.

Visual support for the task of feature subset selection requires displaying information on different levels of granularity; highly detailed views are not optimal because they do not scale well with many features. For instance, the tool by Hohman et al. [63] facilitates the visual comparison of feature distributions for a high number of features. It visually exposes the divergence of distributions in terms of training and testing splits, predictive performance (correct vs. incorrect predictions), and multiple data versions. Overall, this tool is less related to the context of our work because it focuses on the challenge of data iteration, in which analysts continuously add and/or remove instances over iterations. Bernard et al. [64] present a tool that forms a connection panorama between features and between their bins. However, their VA tool is limited to categorical data. Another feature selection approach is SmartStripes by May et al. [25]. It supports the investigation of interdependencies between different features and entity subsets. Their main idea is to find features with the strongest correlation with instance partitions of the data set. Alternatively, our approach groups instances based on the predicted probability of an ML algorithm. It also indicates features that played a vital role in distinguishing instances from the worst to the best predicted.

3 ANALYTICAL TASKS AND DESIGN GOALS

This section specifies the fundamental analytical tasks (**T1–T5**) that users should be capable of when performing feature engineering while they obtain assistance and guidance from a VA system. Afterwards, we report the corresponding design goals (**G1–G5**) that led to the development of our proposed FeatureEnVi system.

3.1 Analytical Tasks for Feature Engineering

We derived the analytical tasks described in this section from the in-depth analysis of the related work in Section 2. The three analytical tasks from Krause et al. [50], the three experts who expressed their requirements in Zhao et al. [32], and the user tasks acquired through expert interviews from Collaris and van Wijk [37] also inspired our work. Additionally, our own recently conducted surveys [65], [66] and experiences with prototyping VA

tools for ML [67], [68], [69] played a significant role in the final formulation of the following analytical tasks.

T1: Inspect the impact of features in both local and global perspectives. In data sets with many features, it is not trivial to acknowledge each feature's contribution in the final prediction. Notably, a user needs a global mechanism to examine the relationships between those features (if there are any). Before users decide about the engineering strategy behind every feature, another factor that they must consider is the impact of features in representative groups of instances [37]. Finding such groups should be a goal for VA systems (cf. **G1**).

T2: Contrast various feature selection techniques. The results produced by automatic feature selection techniques can be remarkably inconsistent. A user should be able to compare alternative feature selection techniques according to their rankings for all features. The interviewees of Krause et al. [50] stated that they would like to get answers to questions such as: "Which features are consistently ranked highly?", "How much do the techniques differ in their ranking?", and "Are there features that have a high rank with some techniques and a low rank with some others?" The visual feedback from VA systems should provide users the ability to respond to those inquiries (**G2**).

T3: Analyze the effect of diverse feature transformations. When some features' distributions are skewed left or right, logarithmic or exponential transformations respectively could further boost each feature's contribution. Moreover, particular ML algorithms (e.g., SVM [39] and XGBoost [29]) might expect a few preprocessing steps, for instance, scaling the data values with z-score normalization [70]. A user should be competent in judging the influence of feature transformations before applying them [37]. As explained in **G3**, visual clues are mandatory for assessing every transformation in practice.

T4: Combine features to generate new features. It is common in feature engineering to combine features in many ways by using mathematical operators. Feature generation produces features with higher levels of contribution because they pack additional information. Here, users need means to select a few features and a way to evaluate the new combination. A typical question connected to **T2** (about feature subset selection) and this task was made by the participants of Krause et al.'s interview session [50]: "How does the performance of the model increase or decrease if I remove or add these features?" The same could apply for newly-generated features (that we are unaware of) when compared against the original pool of features (see **G4**).

T5: Evaluate the results of the feature engineering process. At any stage of the feature engineering process (**T2–T4**), a user should be able to observe the fluctuations in performance with the use of standard validation metrics (e.g., *accuracy*, *precision*, and *recall*) [32]. Also, users could possibly want to refer to the history of the actions they performed to identify the key spots that might have corresponded to improvements in the outcome. VA systems must thus be able to provide ways of monitoring the performance, as described in **G5**.

3.2 Design Goals for FeatureEnVi

To satisfy the previously defined analytical tasks, we have determined five design goals to be addressed by our tool/system. The implementation of these design goals is illustrated in Section 4.

G1: Division of data space into slices based on predicted probabilities, for transparent local feature contribution. Our

goal is to assist in the search for distinctive features that might contribute more to instances that are harder or easier to classify (**T1**). By splitting the data space into quadrants, we aim to assure that users' interference with the engineering of specific features does not cause problems in key parts of the data space. The tool should begin by training an initial model with the original features of each data set, which will be a starting point (i.e., "state zero") for users to compare future interactions and take further decisions.

G2: Deployment of different types of feature selection techniques to support stepwise selection. There are several different techniques for computing feature importance that produce diverse outcomes per feature. The tool should facilitate the visual comparison of alternative feature selection techniques for each feature (**T2**). Another key point is that users should have the ability to include and exclude features during the entire exploration phase.

G3: Application of alternative feature transformations according to feedback received from statistical measures. In continuation of the preceding goal, the tool should provide sufficient visual guidance to users to choose between diverse feature transformations (**T3**). Statistical measures such as *target correlation* and *mutual information* shared between features, along with *per class correlation*, are necessary to evaluate the features' influences in the result. Also, the tool should use *variance influence factor* and *in-between features' correlation* for identifying colinearity issues. When checking how to modify features, users should be able to estimate the impact of such transformations.

G4: Generation of new features and comparison with the original features. With the same statistical evidence as defined in **G3**, users should get visually informed about strongly correlated features that perform the same for each class. Next, the tool can use automatic feature selection techniques to compare the new features with the original ones, using the same methods as in **G2**. Finally, the tool should let users select the proper mathematical operation according to their prior experience and the visual feedback (**T4**).

G5: Reassessment of the instances' predicted probabilities and performance, computed with appropriate validation metrics. In the end, users' interactions should be tracked in order to preserve a history of modifications in the features, and the performance should be monitored with validation metrics (**T5**). At all stages, the tool should highlight the movement of the instances from one data slice to the other due to the new prediction probability values of every instance. The best case is for all instances to relocate from the left to the right-most side.

4 FEATUREENVI: SYSTEM OVERVIEW AND APPLICATION

Following the analytical tasks and our derived design goals, we have developed FeatureEnVi, an interactive web-based VA system that allows users to utilize multiple metrics and automatic methods in order to engineer useful features.

The system consists of five main interactive visualization panels (Fig. 1): (a) data space (\rightarrow **G1**), (b) feature selection techniques (\rightarrow **G2**), (c) feature space overview, (d) feature space detail (\rightarrow **G3** and **G4**), and (e) process tracker and predictive results (\rightarrow **G5**). We propose the following **workflow** for the integrated use of these panels (cf. Fig. 2): (i) choose four suitable data space slices, which are then used for evaluating the impact of each feature on particular groups of instances (Fig. 1(a)); (ii) in the exploration phase, choose subsets of features using

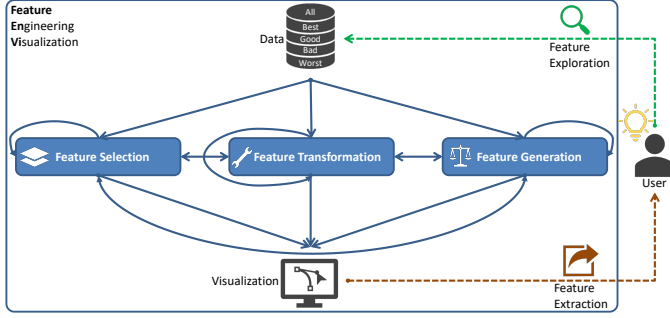


Fig. 2. The FeatureEnVi workflow begins with partitioning the data set according to the prediction probabilities of instances. The data is passed to three different feature engineering processes (selection, transformation, and generation) which are executed iteratively, under the control of the user through the interface.

diverse automatic feature selection techniques (see Fig. 1(b)); (iii) in the overview phase, explore the impact of features in both the individual slices and the entire data space, with various computed metrics (Fig. 1(c)); (iv) during the detailed examination phase, check the different transformations of the features with statistical measures and compare the combinations of two or three features that result in newly-generated features (cf. Fig. 1(d)); and (v) contrast the performances of the *best* predictive performance found so far vs. the *current* result according to three validation metrics in Fig. 1(e). All these processes are iterative and could happen in any other order. The final outcome is the generated knowledge acquired from the extracted features. Note that the typical workflow, which is followed in the two use cases and the case study in Section 5, is linear to the design of the views of FeatureEnVi (i.e., from **G1** to **G5**). The individual panels and the workflow are discussed in more detail below.

All visual encodings designed for the panels of FeatureEnVi are summarized in Table 2. On the right-hand side, we can observe the optimal states for the available statistical measures. However, in reality, many of the statistical measures will be contradictory to each other, and human decisions are essential on such occasions. To the best of our knowledge, little empirical evidence exists for choosing a particular measure over others. In general, target correlation and mutual information (both related to the influence between features and the dependent variable) may be good candidates for identifying important features [71]. After these first indicators, the remaining collinearity measures can be useful too. Although it can be claimed that a high level of correlation between features is from 0.9 and above [72], no precise rules exist for judging the significance of the collinearity (see our discussion on the variance influence factor in Section 4.4). To accommodate for these challenges, a variety of measures are simultaneously visualized in our tool. In detail, for the *Data Space* panel, the predicted probabilities should become dark or light green. The more important a feature, it is more likely to impact the outcome of the ML model, thus, in the *Feature Selection Techniques* panel, the normalized importance should be mostly green or close to green colors. The *Feature Space Overview* panel contains a subset of the measures present in the *Feature Space Detail* panel. For the former, target correlation (*COR*) should be—as high as possible—depicted with a full circle bar. The same applies to mutual information (*MI*), but the indication here is the dark or darker blue color. For the difference between features correlation,

TABLE 2
Summary of the Various Measurements and their Visual Mappings for the Different Panels Supported by FeatureEnVi

Panel	Measurement	Visual Mapping				
Data Space	Predicted Probability	Worst	Bad	Good	Best	
Feature Selection Techniques	Normalized Feature Importance	0%	25%	50%	75%	100%
Feature Space Overview	Target Correlation (<i>COR</i>) [*]	0%	25%	50%	75%	100%
	Mutual Information (<i>MI</i>) [†]	Light Blue	Lighter Blue	Blue	Darker Blue	Dark Blue
	Difference Between Features Correlation [‡]	Increasing Feature Correlation		Reducing Feature Correlation		
Feature Space Detail	Target Correlation (<i>COR</i>) [*]	0%	25%	50%	75%	100%
	Mutual Information (<i>MI</i>) [†]	Light Blue	Lighter Blue	Blue	Darker Blue	Dark Blue
	Variance Influence Factor (<i>VIF</i>)	...>10	10>...>5	5>...>2.5	2.5>...	
	Per Class Correlation (<i>COR</i>)	0%	25%	50%	75%	100%
	Between Features (<i>Fs</i>) Correlation	1.0	0.75	0.50	0.25	0.0
	Difference Between Features Correlation [‡]	0.0	0.25	0.50	0.75	1.0
		Reducing Feature Correlation		Increasing Feature Correlation		
Process Tracker and Predictive Results	Step of the Exploration	1	2	3	4	5
	Score/Performance	Current		Best		

This table is hierarchically organized according to the five main panels that follow our system's workflow. From left to right, we present the worst to the best possible outcome for all measurements.

^{*}, [†], and [‡]: These measurements are shared across the Feature Space Overview and the Feature Space Detail panels, using similar visual mappings.

the optimal is to reduce, hence, green is the expected outcome. The same encodings also apply for the latter panel with the addition of the variance influence factor (*VIF*), per class correlation (*COR*), and between features correlation (*Fs*). The per class correlation is the only of the three measurements that should be substantially high, which is illustrated with long bars for the horizontal bar chart. The remaining should both be decreasing, as a result, the gray color visible in this panel should be diminishing—as much as possible. Finally for the *Process Tracker and Predictive Results* panel, moving forward to the next exploration step is mapped to bigger in size circles with brown color used for the best setup of features explored so far.

The workflow of FeatureEnVi is model-agnostic within the scope of classification problems, meaning that it could be paired with any classification algorithm. In this current version, the implementation uses a state-of-the-art ensemble learning method called XGBoost [29]. This choice was made intentionally because some algorithms (e.g., SVM [39] and XGBoost [29]) are susceptible to specific types of transformations (e.g., scaling). To make our approach even more future-proof, we train this ML algorithm with the *Bayesian Optimization* package [73]. For this section, we validate our results with cross-validation using 8-folds, and we scan the hyperparameter space for 25 iterations, choosing the model with the best accuracy. The hyperparameters we experimented with (and their intervals) are: *number of trees* (5–200), *learning rate* (0.0–0.3), *maximum depth of a tree* (6–12),

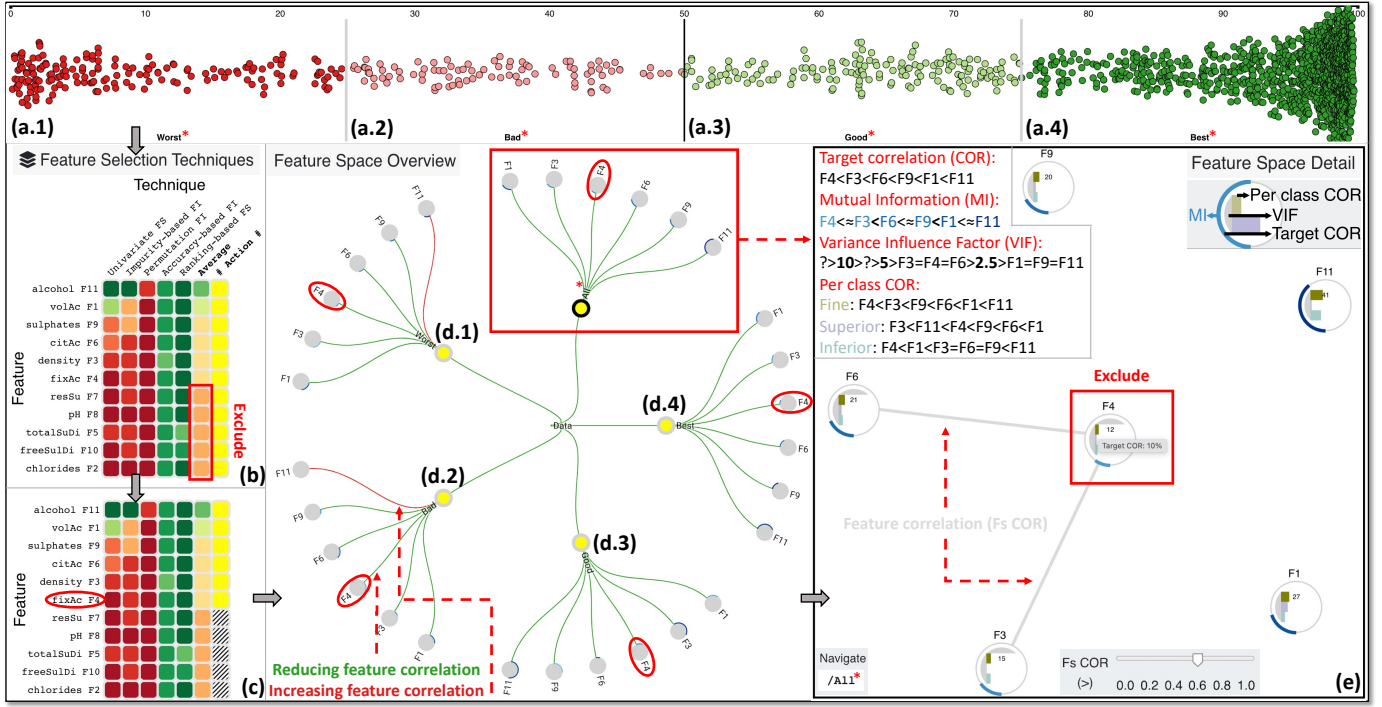


Fig. 3. Exploration of features with FeatureEnVi. The *default* slicing thresholds for the data space separate the instances into four quadrants that represent intervals of 25% predicted probability (see (a.1–a.4)). View (b) presents a table heatmap with five different feature selection techniques and their average value per feature. We exclude the less contributing features, as shown in the duplicated view (c). In the radial tree, the paths from (d.1) to (d.4) are the features for the groups formed at (a.1–a.4), respectively, while the features’ impact for the entire data set is shown in the red box. The whole data space is displayed with even more details in the graph visualization in (e), where additional metrics’ results are reported. A summary of the meaning of the visual encodings for these metrics is visible in the top-left corner in (e). More details about these views are described in the text.

subsample ratio of the training instances (0.8–1.0), and *fraction of columns to be subsampled* (0.8–1.0). All those values can be easily adjusted within the code.

In the following subsections, we explain the system by using a running example with the *red wine quality* physicochemical data set [74] obtained from the UCI ML repository [75]. The data set represents a very imbalanced multi-class classification problem and consists of 11 numerical features and 1,599 instances. Consequently, we mapped the six quality categories of wine for the dependent variable to three new classes (similar to Laughter and Omari [76]) to alleviate this problem. The *fine* class contains 1,319 instances and comprises qualities of 5 and 6. The *superior* quality consists of 217 instances and categories 7 and 8. Finally, the *inferior* class has 63 instances, and it was created by merging qualities 3 and 4.

4.1 Data Space

In FeatureEnVi, data instances are sorted according to the predicted probability of belonging to the ground truth class, as shown in Fig. 1(a). The initial step before the exploration of features is to pre-train the XGBoost [29] on the original pool of features, and then divide the data space into four groups automatically (i.e., *Worst*, *Bad*, *Good*, and *Best*). The vertical black line is the stable threshold anchored precisely at 50% predictive probability, which separates the correctly from the wrongly classified instances. The other two thresholds partition the prior subspace into their half areas for the default option. However, the user can alter the vertical gray lines as indicated in Fig. 5(a.1–a.4), with a degree of freedom

set to $-/+20\%$ from the defaults. The vertical positioning of the instances is purely used to avoid—as much as possible—overlapping/cluttering issues via jittering. The data space will always be divided into four parts conveying extra information to the user. If no instances belong to a slice of the data space, the system works normally, but there will be no values for the statistical measures (see Section 4.3). Overall, the user’s goal is to move as many instances as possible from the left side (Worst and Bad subspaces) to the right side (Good and Best subspaces) while avoiding the opposite. Nevertheless, the primary purpose of this view is to provide better local and global explainabilities of the impact of features according to the user-defined slices. In the future, we plan to enable users to determine the number of slices customly (cf. Section 7.1).

In our running example, recognizing that the instances are uniformly distributed from Fig. 3(a.1–a.3), we use the default splitting with 25% predictive probability intervals. It is essential to focus on the aforementioned slices (especially (a.1) and (a.2)) without ruining the prediction for the correctly classified instances. Once the thresholds are set, we click the *Feature Exploration* button to start our search for features.

4.2 Feature Selection Techniques

Fig. 3(b) is a table heatmap view with five automatic feature selection techniques, their *Average* contribution, and an *# Action* button to exclude any number of features. As we originally train our ML algorithm with all features, the yellow color (one of the standard colors used for highlighting [77]) in the last column

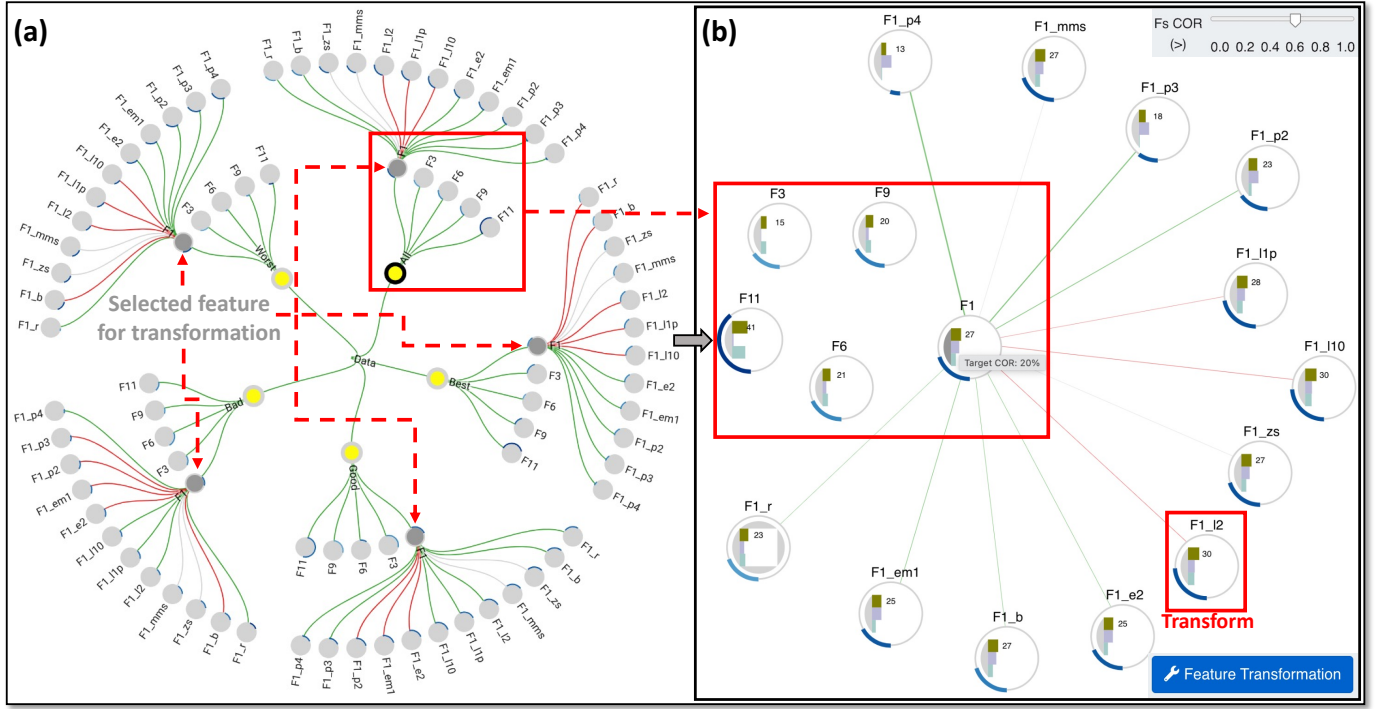


Fig. 4. The feature transformation phase in order to empower features and improve the prediction. In (a), we have selected F1, and we check the impact of different transformation strategies in all slices of the data space. The entire data space is shown in more detail with the graph visualization in (b). The F1's feature transformation that appears the best is the logarithmic transformation ($_{I2}$ or $_{I10}$).

symbolizes that all features are included in the current phase (if excluded, then B/W stripe patterns appear). The first technique is *Univariate FS* (Feature Selection) [78] which uses the ANOVA *F-value* test for selecting the k best features. The value k is always set to the maximum in order to retrieve scores for all features, since we want to avoid removing features automatically (instead, we visualize the scores for the user to decide). The *Impurity-based FI* (Feature Importance) method [79] is connected to the intrinsic nature of ensemble algorithms to export feature importance scores after their training. Hence, we extract the feature importance scores from the best model we found. *Permutation FI* [80] is another technique in which the decrease in a model's score is monitored while a single feature value is randomly shuffled [81]. The last two techniques are rather similar with one difference: the former method can be biased toward high cardinality features (many unique values) over low cardinality features such as binary features. In *Accuracy-based FI* [82], we also fit the ML model using one feature at a time and compute the accuracy to evaluate every feature's performance. Next, as XGBoost [29] is a nonlinear ML algorithm, we also train a linear classifier (a logistic regression [83] model with the default Scikit-learn's hyperparameters [84]) to compute the coefficients matrix and then use Recursive Feature Elimination (RFE) [40] to rank the features from the best to the worst in terms of contribution. This technique is referred to as *Ranking-based FS* [85] in our VA system. We would like to include further techniques in the future, however, the current selection is specifically assembled to contain one candidate for each of the high-level categories of feature selection methods introduced in Section 1. For every method, we normalize the output from 0 to 1 to set a common ground for the user to compare them, as indicated in the legend of Fig. 1(b). Hence, their average is calculated and displayed in the penultimate column.

Following the design guidelines from the conventions introduced by prior works [86], [87], we choose red and green colors for the table heatmap. This view also automatically extends for the newly-generated features from combinations of already existing features (cf. Fig. 1(b)). The original features used for the creation of new features are depicted in dark gray in the last column of the table heatmap view. The table is automatically sorted based on the average; however, Impurity-based FI is selected by the user for the Fig. 1(b)) scenario. Due to this selection, the table heatmap resorts the features from the highest to the lowest importance only according to the XGBoost model's inherent feature importance. More details can be found in Section 4.4.

At this phase, we want to identify any number of features that can be excluded from the analysis because they contribute only slightly to the final outcome. This will help the system to reduce the computational time needed to train the model—something significant in a real-world scenario. Indeed, if we take a closer look, the last five features underperform and are having a shallow impact on the final result (see Fig. 3(b)). Thus, we exclude those features one by one with the interactive cells from the *# Action* column. In Fig. 3(c), we can notice the black-and-white stripe pattern indicating that a feature is excluded. The remaining open question that we will examine in Section 4.3 and Section 4.4 is: shall we continue with the removal of features (e.g., F4 marked with the red ellipsoid shape) or stop at this point?

4.3 Feature Space Overview

After the initial removal of features, as described in Section 4.2, we take a look into the radial tree that presents statistical information about the impact of the currently included features. The core aim of this view is to examine the impact on various subspaces, since removing a feature might appear the right choice globally,

but can cause issues for instances in a particular subspace. This hierarchical visualization exploits the connections of these features (see Fig. 3(d.1–d.4)) with the four subspaces we defined in Section 4.1, which is the *inner layer*. The top part highlighted in a rectangular red box is the whole data space with all the slices (text in bold), and it is currently active (black stroke instead of gray) for the Feature Space Detail panel we will explore in Section 4.4. One of the purposes of this view is to serve as an *overview* for the graph visualization described in the next subsection (i.e., the more detailed *context*). The radial tree's gray nodes are the six remaining features that constitute the *middle layer*. Finally, the *outer layer* is used for the feature transformations, as depicted in Fig. 4(a). The user can click on any gray circle visible in Fig. 3(d.1–d.4) to activate another slice, thus updating the graph visualization (discussed in the next subsection).

The impact of the features is visually represented here in different ways. Each node is enclosed by a circular bar indicating the absolute value of *Pearson's correlation coefficient* [88] with the dependent variable, from 0% to |100%|. The light blue to dark blue color is the *MI* computed between each feature and the target. The colors of the edges represent the aggregated difference of a feature, if it was transformed at all. Green means that the correlation between features decreases with most transformations, while red depicts the increase, respectively. With those coherent design choices, the consistency between the radial tree and the graph visualization described in Section 4.4 is preserved, and the users' cognitive load is reduced (see Section 7.1 for further discussion of the core design choices).

The radial tree representation and the graph visualization (see below) use a layout panel visible at the bottom of Fig. 1(c). It supports several interactions such as *zooming*, *rotating*, *returning to the initial view*, and expanding/retracting particular slices with the *Toggle* options. Zooming enables users to compare effortlessly the statistical measures computed for the different subspaces depicted in the radial tree representation. Rotating allows users to read the text labels easily. Returning to the initial view can also be a helpful shortcut during the interactive exploration. The toggle options assist users in collapsing unimportant data subspaces in particular cases or vice versa. Furthermore, this functionality facilitates FeatureEnVi to scale for data sets with many more features (see Section 7.2 for the respective discussion).

Next, we focus on the overall inspection of features for all instances (see Fig. 3(d.1–d.4)). F4 (the ellipsoid shape) appears the worst in terms of target correlation (the small circular bar), and it has one of the lowest MI values (light blue color). The impact is also the weakest for the Worst and the Best subspaces compared to the other features. However, for the remaining slices, it has a higher impact than other features, as can be seen in Fig. 3(d.2 and d.3). The detailed data space view should thus be examined next.

4.4 Feature Space Detail

The graph visualization in Fig. 3(e) uses glyphs to encode further details for every feature, i.e., *node*. The per target correlation and MI are the same as in the radial tree, but an additional horizontal bar chart encodes the *per class correlation* of each feature with three colors (olive for fine class, purple for superior class, and turquoise for inferior class). Moreover, there are four states that indicate if the *VIF* per feature was greater than 10, between 10 and 5, between 5 and 2.5, and finally, less than 2.5. They are represented by up to four circular segments in gray (see Fig. 3(e))

TABLE 3
Summary of the Currently Supported Feature Transformations

Category	Suffix	Description
Rounding	r	Round evenly
Binning	b	Optimal binning
Scaling	mms	Min-max scaling
	zs	Z-score
Logarithmic transformations	l2	Binary logarithm
	l1p	Natural logarithm plus 1
	l10	Decimal logarithm
Exponential transformations	e2	Exponent with base 2
	em1	Exponent with a natural base minus 1
Power functions	p2	Power of 2
	p3	Power of 3
	p4	Power of 4

laid out clockwise within the node. We decided for these states as the prior research suggests they reflect concerns or problematic cases of colinearity [89], [90], [91]. *Edges* are displayed between feature nodes with the correlation above the current threshold (0.6 in our example), with edge widths encoding correlation values.

The graph visualization uses a force-based layout algorithm, and we allow the user to select different timestamps of the graph generated by a dynamic optimization approach [92]. Whenever a node in focus is not clearly visible, the user can choose another pre-stored layout to reduce clutter. We tweaked the parameters to put more weight to the avoidance of edge crossings ($\times 1.00$), then the deviation of angular resolution ($\times 0.75$), next the minimum angular resolution ($\times 0.40$), and finally the edge crossing angle ($\times 0.10$) in the graph visualization. We found empirically that these values produce good graph layouts that match the goal of our tool. The *weighted graph readability widget* is located at the bottom of the panel in Fig. 1(c), which also includes an edge filtering slider based on minimum *features' correlation* (Fs COR) value.

In the panel below the graph view (see Fig. 1(d)), the user is able to select between two modes: (1) *Feature Transformation* and (2) *Feature Generation*. The user is also informed about the currently navigated data space and provided with static legends describing this view. When the *Feature Transformation* mode is active, clicking on a node will expand it, exposing all supported transformations which are based on mathematical functions (Fig. 4(a)). The available transformations can be divided into six categories listed in Table 3. For Fig. 4(a and b), the red and green colors have the same meaning as explained in Section 4.3. The magnitude of increase/decrease is mapped to the line width in Fig. 4(b). For *Feature Generation*, FeatureEnVi supports four basic arithmetical operations (+, −, ×, /). The selected features are highlighted in the dark gray color (because it matches the default color, which is gray) of the VIF metric's region, as demonstrated in Fig. 1(d), and the combinations are generated for the two or three selected features automatically, as can be seen in Fig. 1(b). It is up to the user to select the best generation; however, he/she is being assisted by the automated selection techniques, as described in Section 4.2. In the future, we plan to support custom transformation and generation of features (see Section 6).

By comparing the lengths of the circular bars in Fig. 3(e), we see that the lowest overall target correlation is reported for F4 (on hover shown as 10%). Also, the MI exhibits low values for both F3 and F4. As we proceed, we observe that F3, F4, and F6 may cause problems regarding collinearity based on the VIF heuristic (2 out of 4 pieces of the symbol). Additionally, the per-class correlation is 12% for the fine label, and weak for the other classes. Finally,




Fig. 5. The process of features' exploration in a vehicle recognition scenario. (a.1) to (a.4) depict the change of the thresholds for the different data slices to intensify the responses for borderline instances. In (b), the user excluded unimportant features and then validates the remaining features through the radial tree visualization. (c.1) to (c.4) contain the main statistical measures (target correlation and mutual information) and allow the user to discover that F8 and F16 should be excluded, as shown in (d). Finally, (e) illustrates the transformation alternatives of the most contributing feature (i.e., F4); the user employs a power function since it increases the correlation with the target and reduces slightly the correlation between features.

the correlation between F4 and F6 (or F4 and F3) is above 0.6; thus, it can be considered moderate to strong. This leads us to excluding F4, with the results in Fig. 4(b) being promising since the remaining features (in the red box) show no signs of strong correlation with each other or any collinearity if we consider VIF.

From Fig. 1(d), we can already see the transformations we applied in the different features with the suffixes attached. At this stage, we choose the feature generation mode and pick three features (F1_l2, F6_b, and F9_l10) that correlate in a similar way with the individual classes. Moreover, F1_l2 is moderately correlated with F6_b. The selected nodes are highlighted with dark gray color, and the combinations of the different mathematical operations are prepared dynamically (see Fig. 1(b)). From this subfigure, we spot that $F1_l2 \times F6_b$ contributes much more than the source features if we sort the table according to the Impurity-based FI. Therefore, we add this new feature.

4.5 Process Tracker and Predictive Results

Each action of the user is registered in the punchcard visualization (see Fig. 1(e)). The basic recorded steps are: (1) *Include*, (2) *Exclude*, (3) *Transform*, and (4) *Generate* for each feature. The size of the circle encodes the order of the main actions, with larger radii for recent steps. The brown color is used only if the overall performance increases. The calculation is according to three validation metrics after we subtract their standard deviations. The grouped bar chart presents the performance based on accuracy, weighted precision, and weighted recall and their standard deviations due to cross-validation (error margins in black). Teal color encodes the current action's score, and brown the best result


reached so far. The choice of colors was made deliberately because they complement each other, and the former denotes the current action since it is brighter than the latter. If the list of features is long, the user can scroll this view. Finally, there is a button to *extract* the best combination of features (i.e., the "new" data set).

The brown circles in the punchcard in Fig. 1(e) enable us to acknowledge that the feature generation boosted the overall performance of the classifier. High scores were reached in terms of accuracy, precision, and recall. **All in all** with FeatureEnVi, we improve the total combined score by using 6 well-engineered features instead of the original 11. On the contrary, Rojo et al. [33] reported a slight decrease in performance when selecting 6 features for this task as a regression problem.

4.6 Use Case: Understanding Feature Engineering

We now describe a different case of how FeatureEnVi could be used to explore the feature space and engineer 18 features for the multi-class *vehicle silhouettes* data set [93], comprising vehicle data assigned to one of three classes: *van*, *car* or *bus*. The class distribution is rather imbalanced, with 199 vans, 429 cars, and 218 buses. For this use case, we use the same ML algorithm, hyperparameter optimization method, and cross-validation strategy as in the previous application. This use case was performed by us, and it was the first time we explored this particular data set.

Slicing the data space and selecting subsets of features. Similar to the workflow described above, we start by choosing the appropriate thresholds for slicing the data space. As we want to concentrate more on the instances that are close to being predicted correctly, we move the left gray line from 25% to 35%



Generating features from similarly behaving features. The newly-transformed F4_p4 and F18_llp are both powerful

Tracking the process and evaluating the results. To verify each of our interactions, we continuously monitor the process through the punchcard, as shown in Fig. 6(c). From this visualization, we acknowledge that when F16 was excluded, we reached a better result. The feature generation process (described previously) led to the best predictive result we managed to accomplish. The new feature is appended at the end of the list in the punchcard. From the grouped bar chart in Fig. 6(c), the improvement is prominent for all validation metrics because the brown-colored bars are at the same level as the teal bars. **To summarize,** FeatureEnVi supported the exploration of valuable features and offered transparency to the process of feature engineering. The following case study is another proof of this concept.

In addition to the use cases discussed above, we present a case study with a focus on helping users to understand the feature engineering process and improve the classification results. The case study also provides useful insights that can be used to improve our visualization approach in the future.




Fig. 7. Engineering features for improved predictive performance. From the pre-training phase, we detect that most of the instances belong to the Best slice (a.4), then the Worst slice (a.1), followed by the remaining slices (a.3 and a.2). In view (b), we validate every feature by working in synergy with the table heatmap and the radial tree shown in (c). We pick the whole data space to be juxtaposed against the Worst slice since it contains many wrongly classified instances (see c.1). Interestingly, F4 appears unimportant for the Worst slice, but it is somewhat correlated with the target variable in the All space. Later, we explored this particular feature in the Bad slice, and we confirmed that it is valuable there. Consequently, we decided that it should be retained. F37 is an important feature for both spaces under investigation, hence, it is also kept in the pool of features. In (d), we observed that the logarithmic transformations are the most impactful in terms of increase in target correlation, but at the same time, F8_I2 and F8_I10 are the only features that reduce the correlation between features. Thus, we choose one of the best candidates. After the transformations are applied in different features, F3_I1p and F39_I2 are connected due to the new logarithmic transformations, as seen in (e). The between these features correlation has been increased compared to before. However, the links between F3_I1p with F10 and F39_I2 with F7 have disappeared because the transformations altered the instances' values. Finally, we experimented with feature generation by selecting strongly correlated features, and we discovered that $|F2 - F8_I2|$ and $F7 + F10$ combinations are better than their stock versions, as seen in (f).

This case study was carried out together with an invited ML expert (PhD in *computer science*, currently *associate professor* in *artificial intelligence*, with 20 years of experience in ML, but no prior experience in VA). The expert collaborated with us using FeatureEnVi to improve the results of a prior study about the relationships between *chemical structure* and *biodegradation of molecules*, described in the work by Mansouri et al. [94]. The *QSAR (Quantitative Structure Activity Relationships) biodegradation* data set represents a binary classification problem where molecules are assigned to either the *biodegradable* or *non-biodegradable* classes. The class distribution is relatively imbalanced, with 284 degradable and 553 non-degradable molecules for the training set containing 41 different features. The test set

contains 72 and 146 instances in each class. At the same time, the external validation set consists of 191 and 479 instances, respectively. For their solution, the authors trained three ML models (KNN [95], PLSDA [96], and SVM [39]) and then combined their results using two consensuses. We worked hand-in-hand with FeatureEnVi to surpass their predictive results by solely manipulating the features of this data set. Only for this case study, the cross-validation strategy was adapted to 5-folds instead of 8-folds for comparison reasons with Mansouri et al. [94].

Selection of impactful features for the Worst subspace. We began our investigation by examining the distribution of instances in the explorable subspaces. We noticed that most instances are correctly classified with more than 75% predicted probability (i.e.,

high confidence), as shown in Fig. 7(a.4). The invited ML expert found the 25% predicted probability intervals a considerably rational choice. Thus, it was adopted as the default threshold strategy for this analysis. The focus was on the Worst slice (cf. Fig. 7(a.1)) because it contained many instances that could be hard to classify. These instances could lead to low generalizability and hurt the overall predictions since such rare cases can occur in the test and external validation sets, and it is usually not trivial to classify them correctly. At this point, we worked with the table heatmap view (Fig. 7(b)) to validate all features and exclude a subset of features. The automatic feature selection techniques cannot be trusted when left alone due to the contradictions arising in various situations. Therefore, we utilized the statistical measures visible in the radial tree (see Fig. 7(c)). The radial tree had three collapsed data subspaces (a.2–a.4) except for All and Worst subspaces. We performed this action because there are too many features to be explored at once, and FeatureEnVi provides this capability to alter the layouts in order to scale for high-dimensional data sets. Basically, the core statistical measure to examine in the radial tree is the target correlation depicted as a circular bar of different sizes for each feature in every data subspace. Moreover, a supportive measurement is the MI per feature that should have a light blue color in the cases of a potentially removable feature. Going from the bottom to the top of the feature importance list sorted by the Average importance, we excluded all features with a target correlation value greater than 15% according to the All space. Two interesting examples of features are F4 and F37. The former was powerful only for the Accuracy-based FI and slightly better for the Univariate FS than the remaining feature selection techniques. The latter was moderately good for all techniques except for Permutation FI and Ranking-based FS. We concentrated on the conjunction of those automatic approaches with the statistical measures offered by FeatureEnVi. Specifically, F4 resembled an unimportant feature for the Worst subspace, as shown in Fig. 7(c.1). Albeit that, when closely explored in the whole data space, it was more impactful than other features (with a target correlation value of more than 15%). Afterwards, we collapsed the Worst slice and expanded the Bad slice to explore the impact of the feature. Indeed in that particular slice, this feature appeared to be rather important (not shown due to space limits). These recursive actions reassured us that F4 should be kept in the pool of features. On the other hand, F37 was very impactful in both All space and Worst subspace (cf. Fig. 7(c and c.1)), hence, it remained in the list of active features. Without the support of additional statistical measurements such as target correlation, we could have made possible mistakes because automatic feature selection techniques sometimes contradict each other, and they are not fully trustworthy in certain situations. After this phase, 24 out of the 41 original features remained in use.

Transformation of features with guided decisions. For the feature transformation phase, the first step was to analyze every feature with a top-down approach according to the sorting implied in the table heatmap (Fig. 7(b)). FeatureEnVi facilitates the exploration of multiple contradictory criteria again, even while choosing the individual transformations. In Fig. 7(d), we acknowledged a case where, according to the target correlation measure, all logarithmic transformations were sufficient. However, F8_12 and F8_110 were better candidates for transformation since the correlation between features decreased (green lines instead of red-colored) when choosing these options against the F8_11p transformation. All performed transformations are displayed in Fig. 7(e).

F26 was the only feature transformed using a power function, while F21, F8, F3, F20, F6, F32, F37, and F39 were transformed with logarithmic transformations. The rest of the features remained stable. During this exploration process, we applied the following rule: transformation of not so impactful features with extremely low target correlation value (less or equal to 20%) was skipped. As mentioned earlier in Section 4, most useful transformations appear in impactful features. We noticed that when one feature of a pair of strongly correlated features gets transformed, then there is a possibility that their correlation will decrease. The opposite effect is also plausible. Consequently, not all features that are strongly correlated with each other should be transformed, which is another gained insight. In Fig. 7(e), we transformed F3_11p and chose to keep the same F10 to remove the initial link between these two features. The same approach was followed for features F39_12 and F7, which led to an identical outcome. These transformations, though, created a new link of strongly correlated features with F3_11p being connected to F39_12 (see Fig. 7(e)). With the aforementioned rule and insight in mind, we managed to decrease the number of correlated features with more than 0.8 Fs COR to 8 connections (shown in gray lines) instead of the initial 10 links. Also, the VIF was reduced from greater than 5 to greater than 2.5 (i.e., from 3 to 2 symbols active) for F10 compared to before the transformation phase. Similarly, F26, F21, F20, F16, and F13 went from greater than 10 to greater than 5 in VIF (from 4 to 3 symbols active), while F4's VIF had a bigger difference gap from greater than 10 to just greater than 2.5 (i.e., from 4 to 2 symbols active). The increase in target correlations along with decreases in-between feature correlations and VIF were the main steps to reach a higher performance alongside the steps discussed next. At this point, 9 out of the 24 features were transformed.

Generation of features from correlated pairs of features. At this stage, we concentrated on generating new and more effective features from the existing ones. FeatureEnVi provides assistance in this procedure by highlighting only the strong correlation of features (with the current value for the Fs COR) in gray lines (cf. Fig. 7(e)). We tested all old links between pairs of features that are visible in (e) because they are computed quicker than the examination of three features concurrently. The new connection was omitted because it was generated artificially from our interactions. In detail, F37_12 with F2, F8_12 with F12, F8_12 with F2, and finally, F7 with F10 were explored a pair at a time through the updatable table heatmap view. The first two pairs of features, when combined, gave only less important features compared to the original features used to create them. Thus, we did not include any new features from these pairs. On the other hand, F8_12 and F2 when combined with the $|F2 - F8_{12}|$ mathematical operation, generated a more impactful feature (see Fig. 7(f)). The same is true for the F7 and F10 pair of features when using the $F7 + F10$ operation. To sum up, two new features were generated, resulting in 26 well-engineered features instead of the 41 original features.

Evaluation with the test and external validation sets. Throughout the aforementioned phases, we utilized feature engineering to improve the most powerful XGBoost model found through hyperparameter tuning. To verify whether our findings were reliable, we applied the resulting ML model to the same test and external validation sets as Mansouri et al. [94], see Table 4. For the test data set, the reported accuracy was approximately 87%. In our case, we reached 89% for accuracy after the feature engineering process was completed. Using our approach, we managed to achieve the same accuracy as before, 89%, compared

TABLE 4
Results for the Test and External Validation Sets
for the Case Study

Validation Metric	FeatureEnVi		Mansouri et al. [94]	
	Test	External Validation	Test	External Validation
Accuracy	89% (+2%)	89%	87%	83%
Precision	88% (+2%)	87%	92%	91%
Recall	87% (+3%)	85%	82%	76%

The value in parentheses (+X%, where X is a number) highlights the increase in predictive performance due to the use of FeatureEnVi.

to 83% reported by Mansouri et al. [94] for the additional external data set. For precision and recall, we always use *macro-average*, which is identical to Mansouri et al. [94]. On the one hand, the precision was 4% lower in both test and external validation sets for our analysis. On the other hand, the recall was 5% higher for the test set and 9% higher for the external validation data set, which improves the generalizability of the ML model for similar data. **In summary**, our system not only supports the exploration of many different perspectives of the complex relationships between features, but also is capable of optimizing features (and consequently the underlying ML model) to efficiently use fewer well-engineered features and boost the predictive performance.

6 EVALUATION

Following the guidelines from prior works [68], [97], [98], [99], we conducted online semi-structured interviews with three experts to collect qualitative feedback about our system's effectiveness. The first ML expert (**E1**) is a senior lecturer in mathematics with a PhD in this field. He has approximately 3.5 years of experience with ML, and he currently works with reinforcement learning. The second ML expert (**E2**) is a senior researcher in software engineering and applied ML working in a governmental research institute as an adjunct professor. He has worked with ML for the past 8 years. The third expert (**E3**) is an associate professor in computer science and a visualization researcher, with 15 years of experience in ML and VA. The latter two experts have PhDs in computer science; no expert reported any colorblindness issues. Each interview lasted about 1 hour, and the interview procedure was: (1) presentation of the key goals of FeatureEnVi, (2) demonstration of the functionality of each view and experts' interaction with the system using the *iris flower* data set [75], and (3) explanation of the process of reaching the results for the *red wine quality* use case in Section 4. The first part serves as an introduction to the participants to understand the research problems that FeatureEnVi strives to address. The second part focuses on the overall usability as well as the interactions and the connections between views with a simple and well-known data set. The third part is about assessing the system's effectiveness and getting familiar with the workflow. The participants were free to comment on anything, and their main points can be found below.

Workflow. All experts commented that the workflow of FeatureEnVi is straightforward, because it is mainly linear despite involving optional iterative steps. **E2** stated that feature engineering is usually very time consuming, especially without the support of a system like ours. **E3** also agreed with us that the features have an important influence on the predictive model-based quality and affect the generalization ability of the final ML model. Furthermore,

he noticed that it is difficult to judge how each feature should be engineered when there is contradicting statistical evidence. Without FeatureEnVi, it would have been risky to make a deterministic decision. The connection between the features present in the radial tree visualization and the instances' reallocation in the data space at the top of FeatureEnVi helps to identify impactful features (as highlighted by **E1**). Hence, it is up to users to understand which features matter more for the subspaces (locally) and the entire data space (globally), which provides transparency and enhances the trustworthiness of the feature engineering process, as outlined by **E1** and **E2**. Although FeatureEnVi works better with a limited number of features (e.g., 41 for the case study in Section 5), **E2** suggested that a prephase with an AutoML system [100] or a DR algorithm [101] could be a solution, if used to set specific boundaries by investigating the relations between features. The ultimate goal is to exclude non-contributing features prior to extensive analyses within FeatureEnVi. Finally, **E3** discussed a recent example application in design materials [102] that he worked with, which could benefit from our tool. In particular, he said: "automatic ML can combine materials (i.e., features) without them having any physical meaning or being possible to exist; all that just for the sake of improving the predictive accuracy".

Visualization and interaction. **E1** and **E2** were surprised by the promising results we managed to achieve with the assistance of our VA system in the *red wine quality* use case of Section 4. Initially, **E1** was slightly overwhelmed by the number of statistical measures mapped in the system's glyphs. However, after the interview session was over, he acknowledged that FeatureEnVi assisted him with hover interactions, and he even requested more metrics. He stated: "the trick here is to concentrate on specific measurements at a time for multiple features". **E2** was particularly enthusiastic about how simple and informative the visualization of the data space is. **E1** also added that the radial tree visual metaphor encourages him to engage further with the tool since it follows a structural design found in nature. Formulating and testing hypotheses about features is possible through this view due to the several implemented interaction mechanisms, as noted by **E1**. Moreover, **E1** and **E3** appreciated the tool's ability to return back to a previous step, if necessary. The former was positively surprised by the ease of extracting the engineered features, and he specifically stated: "this is awesome". The latter suggested an alternative option for the data space visualization which could have been to aggregate the instances belonging to the different subspaces and use a histogram instead of the beeswarm plot. Another choice is a Sankey diagram which will make the view more complex than it currently is, but it can capture all the steps in between. Both visual representations are worthy options; however, this view has a two-fold purpose: we allow users to adjust the thresholds for the slices directly and, at the same time, we enable them to observe the instances' movement between different subspaces (while reducing cluttering with the use of the available vertical space). The same expert provided us with a further extension that can improve the efficacy of the beeswarm plot, which is described in the following paragraph. He also mentioned that a confusion matrix for visualizing the validation results might be a more detailed approach. Despite that, the bar chart fits the available space, and it scales better in multi-class classification problems. We plan to perform a user study to test diverse visualizations. We also intend to extend the support for user customization, e.g., to hide some of the statistical measurements if preferred by the user.

Limitations. **E1** and **E2** were concerned about the *scalability and efficiency* of the system. Nevertheless, the use of automatic feature selection techniques as one of the first workflow steps is designed exactly for such purposes. Another possible improvement is to utilize parallel processing on powerful cloud servers. Progressive VA and data science workflows [103], [104] could also be effective. Moreover, *alternative feature selection techniques* for computing feature importance could be incorporated in our tool (e.g., SHAP [105]). A customized beeswarm plot could facilitate selecting groups of instances and then *explaining why some instances migrated*. DR methods could also be helpful here, as noted by **E3**. Also, he proposed to include *additional filtering options* for all metrics. The punchcard visualization could allow users to *return at any step and follow another path* (**E2**). Finally, **E1** mentioned that it could be useful if our system supported *custom transformation and generation of features* for users to experiment with. We intend to implement the above functionalities.

Overall assessment. All experts were pleased with the effectiveness of FeatureEnVi in testing hypotheses about individual features or their combinations. Overall, **E2** was impressed with FeatureEnVi and mentioned: “I and colleagues would love to try it out with our own data”. Also, **E1** said: “it is so much fun using your tool”. Afterwards, he continued: “FeatureEnVi has too much value, even using it in a subset of your data, since you can still generate new knowledge”. **In conclusion**, the experts deemed FeatureEnVi a useful approach for supporting feature engineering due to the combination of human and computer intelligence.

7 DISCUSSION

In this section, we discuss several aspects of the design choices as well as limitations of our approach and the current VA system.

7.1 Design Choices

Here, we extend our discussion on the core design principles of our approach that were introduced in Section 4.

Ground truth versus per class predicted probability. In Section 4.1, we explain that the data space view presents the predicted probabilities for the ground truth class. A different approach could have been to visualize the predicted probability for every class. Although this idea appears valuable and straightforward for binary classification problems, it will not scale well with multiclass problems addressed by our VA system. It will be tough to present confusion between classes when more than a couple of class labels are available, which is typical in multiclass problems. Moreover, the limited amount of space in that view is another reason we abandoned this idea, despite being a valid alternative. While the local explainability offered by this data space view is valuable, the focus of our tool is on the feature engineering process.

Custom subspace slicing. Perception and cognition problems could emerge as the number of slices increases. We believe that four slices are already a good start to explore the vast majority of the data space because users will often focus on particular areas of interest. The two interactive thresholds are a key component here because they allow users to choose the size of each subspace depending on the problem. Additionally, the experts in Section 6 did not miss the custom slicing functionality. However, FeatureEnVi could be improved further by enabling users to select a custom number of slices, if that is preferred in certain scenarios.

Radial tree versus other visual representations. The radial tree is an essential supplementary view of the graph visualization

because it serves as a contextual anchor. The visual encodings remain the same, so the perception for users does not change overall. This design decision drastically reduces the cognitive load since users do not have to learn additional encodings from scratch. A potential alternative visual metaphor could have been a tabular form, e.g., an icicle plot. The benefit of this metaphor is the compact space that it captures in comparison to a radial tree. However, the numerous interactions available (e.g., the ability to collapse unimportant slices/central nodes) enable this view to scale well even with more than 40 features. A detailed list of possible layout modifications and interactions can be found in Section 4.3. Despite these functionalities, there is still a bottleneck in the number of features that FeatureEnVi allows users to explore due to visualizations and computational burdens. The first limitation in Section 7.2 provides further information on this matter.

7.2 Additional Limitations and Future Work

Next we focus on additional limitations identified for our approach, which also hint at possible future improvements.

Scalability for large number of features. The limits of FeatureEnVi have already been put to the test with the case study in Section 5. The used data set contains 41 features, which unfortunately are difficult to explore in multiple subspaces simultaneously. Despite that, the invited ML expert was able to monitor the global impact of features along with the local region (Worst subspace) that included a lot of hard-to-classify instances. For many more features, it would be cognitively cumbersome for a human to explore all these features in different subspaces concurrently. The strategy that could be followed is to initially limit the space under examination using an extra preprocessing phase in the pipeline before utilizing FeatureEnVi for a detailed exploration of features. The overall scalability of the system could be improved in various ways, as previously described in Section 6.

Other types of data. Regardless of the wide variety of application domains explored through the use cases and the case study, FeatureEnVi is not tested with other data types rather than structured tabular data consisting of numerical values [106]. One of our future intentions is to support further data types, but as our prototype tool is a proof of concept, the system’s workflow and theoretical ideas are generalizable in this regard. Nonetheless, the features of each data set under investigation should be meaningful because we concentrate on human expertise and knowledge to resolve the contradictory situations between features and to prevent the generation of non-existent features for the sake of performance; in any other case, automatic feature engineering approaches might be preferable. For instance, if we consider image data transformed to numerical values, it is challenging to derive their specific meaning. Also, image data sets are typically combined with neural networks, which could not benefit as much from manual or semi-automatic feature engineering [107], [108].

Targeted users. ML experts and practitioners are the main target groups that would benefit the most by using FeatureEnVi. We assume that they know at least the basics regarding the instances and features of their data, but they require further guidance to the feature engineering process. As seen in Section 6, the ML experts and the visualization researcher were able to understand and use FeatureEnVi within 1-hour interview sessions. An opportunity here is to develop another, more simplified version of our tool, specifically designed for novice users working with ML.

Completion time per activity. The frontend of FeatureEnVi is implemented in JavaScript using Vue.js [109], D3.js [110],

TABLE 5
Analysis Completion Time for Each Activity of Feature Engineering for the Use Cases and the Case Study

Data Set	Feature Engineering		
	Selection	Transformation	Generation
Red Wine Quality	1:12	3:29	2:25
Vehicle Silhouettes	2:58	6:14	5:24
QSAR Biodegradation	5:30	33:24	13:35

The wall-clock completion time is reported in the *minute:second* format.

and Plotly.js [111], and the backend is written in Python with Flask [112] and Scikit-learn [84]. The use cases and experiments were performed on a MacBook Pro 2019 with a 2.6 GHz (6-Core) Intel Core i7 CPU, an AMD Radeon Pro 5300M 4 GB GPU, 16 GB of DDR4 RAM at 2667 Mhz, and running macOS Big Sur. By taking into consideration the specifications of the machine, we captured the total wall-clock time spent to complete each of the use cases and the case study (see Table 5, rows). The time measured is combined for the performance of the user actions and computational analyses as reported in every use case or case study (cf. Sections 4 and 5). Table 5 columns map the time for each activity of feature engineering (i.e., selection, transformation, generation). Concretely, as the number of features analyzed increases, we observe an increase in the time required to compare the features and execute the user-defined commands. Interestingly, the feature transformation took longer in all cases, next was the generation, and finally the selection of subsets of features. These numbers can be rather subjective when different paths are tried out. In general, the efficiency of FeatureEnVi could be improved in numerous ways as already indicated before.

8 CONCLUSIONS

In this paper, we presented FeatureEnVi, a VA system with the aim to engineer features using stepwise feature selection and semi-automatic feature extraction approaches. Multiple coordinated views support users in selecting features, transforming them, and generating new features as part of an intensely iterative process. Exploring the impact of the features with several statistical measures and automatic feature selection techniques enables users to improve the predictive performance, reduce the need for computational resources, and decrease the time spent for training. Finally, our VA system is beneficial for feature engineering because it makes the complete process more transparent. The effectiveness of FeatureEnVi was examined using real-world data sets that demonstrated improvements in performance with fewer available but optimally-tuned features. Our system's workflow and visualizations obtained encouraging feedback from experts, who also helped us to identify the current limitations of FeatureEnVi. These limitations are considered as our future research directions.

ACKNOWLEDGMENTS

This work was partially supported through the ELLIIT environment for strategic research in Sweden. The authors also thank the anonymous reviewers for their constructive feedback.

REFERENCES

- [1] U. Khurana, H. Samulowitz, and D. Turaga, "Feature engineering for predictive modeling using reinforcement learning," in *Proc. of AAAI*. AAAI, 2018.
- [2] P. Domingos, "A few useful things to know about machine learning," *Commun. ACM*, vol. 55, no. 10, pp. 78–87, Oct. 2012.
- [3] S. Kandel, A. Paepcke, J. M. Hellerstein, and J. Heer, "Enterprise data analysis and visualization: An interview study," *IEEE TVCG*, vol. 18, no. 12, pp. 2917–2926, Dec. 2012.
- [4] S. R. Hong, J. Hullman, and E. Bertini, "Human factors in model interpretability: Industry practices, challenges, and needs," *PACMHCI*, vol. 4, no. CSCW1, May 2020.
- [5] European Parliament and Council of the European Union, "Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation)," <https://eur-lex.europa.eu/eli/reg/2016/679/oj>, Apr. 2016, accessed Dec. 8, 2021.
- [6] J. Zhou and F. Chen, "2D transparency space—Bring domain users and machine learning experts together," in *Human and Machine Learning*. Springer, 2018, pp. 3–19.
- [7] J. Heaton, "An empirical analysis of feature engineering for predictive modeling," in *Proc. of SoutheastCon*. IEEE, 2016.
- [8] M. Brooks, S. Amershi, B. Lee, S. M. Drucker, A. Kapoor, and P. Simard, "FeatureInsight: Visual support for error-driven feature ideation in text classification," in *Proc. of IEEE VAST*. IEEE, 2015, pp. 105–112.
- [9] J. Cheng and M. S. Bernstein, "Flock: Hybrid crowd-machine learning classifiers," in *Proc. of ACM CSCW*. ACM, 2015, pp. 600–611.
- [10] H. Liu and H. Motoda, "Feature transformation and subset selection," *IEEE ISTA*, vol. 13, no. 2, pp. 26–28, Mar.–Apr. 1998.
- [11] K. Patel, J. Fogarty, J. A. Landay, and B. Harrison, "Investigating statistical machine learning as a tool for software development," in *Proc. of ACM CHI*. ACM, 2008, pp. 667–676.
- [12] S. Markovitch and D. Rosenstein, "Feature generation using general constructor functions," *Mach. Learn.*, vol. 49, no. 1, pp. 59–98, Oct. 2002.
- [13] B. Schuller, S. Reiter, and G. Rigoll, "Evolutionary feature generation in speech emotion recognition," in *Proc. of IEEE ICME*. IEEE, 2006, pp. 5–8.
- [14] Y. Kankanige and J. Bailey, "Improved feature transformations for classification using density estimation," in *PRICAI 2014: Trends in Artificial Intelligence*. Springer, 2014, pp. 117–129.
- [15] T. Mühlbacher and H. Piringer, "A partition-based framework for building and validating regression models," *IEEE TVCG*, vol. 19, no. 12, pp. 1962–1971, Dec. 2013.
- [16] I. Guyon and A. Elisseeff, "An introduction to variable and feature selection," *JMLR*, vol. 3, pp. 1157–1182, Mar. 2003.
- [17] G. Chandrashekar and F. Sahin, "A survey on feature selection methods," *J. Electr. Comput. Eng.*, vol. 40, no. 1, pp. 16–28, Jan. 2014.
- [18] R. Kohavi and G. H. John, "Wrappers for feature subset selection," *Artif. Intell.*, vol. 97, no. 1–2, pp. 273–324, Dec. 1997.
- [19] A. L. Blum and P. Langley, "Selection of relevant features and examples in machine learning," *Artificial Intelligence*, vol. 97, no. 1–2, pp. 245–271, Dec. 1997.
- [20] K.-B. Duan, J. C. Rajapakse, H. Wang, and F. Azuaje, "Multiple SVM-RFE for gene selection in cancer classification with expression data," *IEEE TNB*, vol. 4, no. 3, pp. 228–234, Sep. 2005.
- [21] J. Ding, J. Shi, and F.-X. Wu, "SVM-RFE based feature selection for tandem mass spectrum quality assessment," *Int. J. Data. Min. Bioinform.*, vol. 5, no. 1, pp. 73–88, Feb. 2011.
- [22] Y. Yang and J. O. Pedersen, "A comparative study on feature selection in text categorization," in *Proc. of ICML*. Morgan Kaufmann Publishers Inc., 1997, pp. 412–420.
- [23] Y. Aphinyanaphongs, L. D. Fu, Z. Li, E. R. Peskin, E. Efsthadiadis, C. F. Aliferis, and A. Statnikov, "A comprehensive empirical comparison of modern supervised classification and feature selection methods for text categorization," *JASIST*, vol. 65, no. 10, pp. 1964–1987, Oct. 2014.
- [24] G. Forman, "An extensive empirical study of feature selection metrics for text classification," *JMLR*, vol. 3, pp. 1289–1305, Mar. 2003.
- [25] T. May, J. Davey, and T. Ruppert, "SmartStripes — Looking under the hood of feature subset selection methods," in *Proc. of EuroVA*. The Eurographics Association, 2011.
- [26] T. May, A. Bannach, J. Davey, T. Ruppert, and J. Kohlhammer, "Guiding feature subset selection with an interactive visualization," in *Proc. of IEEE VAST*. IEEE, 2011, pp. 111–120.
- [27] F. L. Dennig, T. Polk, Z. Lin, T. Schreck, H. Pfister, and M. Behrisch, "FDive: Learning relevance models using pattern-based similarity measures," in *Proc. of IEEE VAST*. IEEE, 2019, pp. 69–80.

- [28] P. Duboue, *The Art of Feature Engineering: Essentials for Machine Learning*. Cambridge University Press, 2020.
- [29] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in *Proc. of ACM KDD*. ACM, 2016, pp. 785–794.
- [30] S. Liu, J. Xiao, J. Liu, X. Wang, J. Wu, and J. Zhu, "Visual diagnosis of tree boosting methods," *IEEE TVCG*, vol. 24, no. 1, pp. 163–173, Jan. 2018.
- [31] J. Fogarty and S. E. Hudson, "Toolkit support for developing and deploying sensor-based statistical models of human situations," in *Proc. of ACM CHI*. ACM, 2007, pp. 135–144.
- [32] J. Zhao, M. Karimzadeh, A. Masjedi, T. Wang, X. Zhang, M. M. Crawford, and D. S. Ebert, "FeatureExplorer: Interactive feature selection and exploration of regression models for hyperspectral images," in *Proc. of IEEE VIS*. IEEE, 2019, pp. 161–165.
- [33] D. Rojo, N. N. Htun, and K. Verbert, "GaCoVi: A correlation visualization to support interpretability-aware feature selection for regression models," in *Proc. of EuroVis — Short Papers*. The Eurographics Association, 2020.
- [34] D. Guo, "Coordinating computational and visual approaches for interactive feature selection and multivariate clustering," *Inf. Vis.*, vol. 2, no. 4, pp. 232–246, Dec. 2003.
- [35] Y. Perez-Riverol, M. Kuhn, J. A. Vizcaíno, M.-P. Hitz, and E. Audain, "Accurate and fast feature selection workflow for high-dimensional omics data," *PLoS ONE*, vol. 12, no. 12, p. e0189875, Dec. 2017.
- [36] M. A. Hall, "Correlation-based feature selection for discrete and numeric class machine learning," in *Proc. of ICML*. Morgan Kaufmann Publishers Inc., 2000, pp. 359–366.
- [37] D. Collaris and J. J. van Wijk, "ExplainExplore: Visual exploration of machine learning explanations," in *Proc. of IEEE PacificVis*. IEEE, 2020, pp. 26–35.
- [38] H. Piringer, W. Berger, and J. Krasser, "HyperMoVal: Interactive visual validation of regression models for real-time simulation," *CGF*, vol. 29, no. 3, pp. 983–992, Jun. 2010.
- [39] M. A. Hearst, S. T. Dumais, E. Osuna, J. Platt, and B. Scholkopf, "Support vector machines," *IEEE ISTA*, vol. 13, no. 4, pp. 18–28, Jul.–Aug. 1998.
- [40] I. Guyon, J. Weston, S. Barnhill, and V. Vapnik, "Gene selection for cancer classification using support vector machines," *Mach. Learn.*, vol. 46, no. 1, pp. 389–422, Jan. 2002.
- [41] J. Yang and V. Honavar, "Feature subset selection using a genetic algorithm," *IEEE ISTA*, vol. 13, no. 2, pp. 44–49, Mar.–Apr. 1998.
- [42] M. Friendly, "Corrgrams: Exploratory displays for correlation matrices," *Am. Stat.*, vol. 56, no. 4, pp. 316–324, Nov. 2002.
- [43] A. MacEachren, D. Xiping, F. Hardisty, D. Guo, and G. Lengerich, "Exploring high-D spaces with multiform matrices and small multiples," in *Proc. of IEEE INFOVIS*. IEEE, 2003, pp. 31–38.
- [44] A. Sanchez, C. Soguero-Ruiz, I. Mora-Jiménez, F. Rivas-Flores, D. Lehmann, and M. Rubio-Sánchez, "Scaled radial axes for interactive visual feature selection: A case study for analyzing chronic conditions," *Expert Syst. Appl.*, vol. 100, pp. 182–196, Jun. 2018.
- [45] E. Artur and R. Minghim, "A novel visual approach for enhanced attribute analysis and selection," *C&G*, vol. 84, pp. 160–172, Nov. 2019.
- [46] Y. Wang, J. Li, F. Nie, H. Theisel, M. Gong, and D. J. Lehmann, "Linear discriminative star coordinates for exploring class and cluster separation of high dimensional data," *CGF*, vol. 36, no. 3, pp. 401–410, Jun. 2017.
- [47] J. Yang, A. Patro, S. Huang, N. Mehta, M. O. Ward, and E. A. Rundensteiner, "Value and relation display for interactive exploration of high dimensional datasets," in *Proc. of IEEE INFOVIS*. IEEE, 2004, pp. 73–80.
- [48] N. Elmqvist, P. Dragicevic, and J.-D. Fekete, "Rolling the dice: Multidimensional visual exploration using scatterplot matrix navigation," *IEEE TVCG*, vol. 14, no. 6, pp. 1539–1548, Nov.–Dec. 2008.
- [49] J. Yang, M. O. Ward, E. A. Rundensteiner, and S. Huang, "Visual hierarchical dimension reduction for exploration of high dimensional datasets," in *Proc. of VISSYM*. The Eurographics Association, 2003, pp. 19–28.
- [50] J. Krause, A. Perer, and E. Bertini, "INFUSE: Interactive feature selection for predictive modeling of high dimensional data," *IEEE TVCG*, vol. 20, no. 12, pp. 1614–1623, Dec. 2014.
- [51] H. Piringer, W. Berger, and H. Hauser, "Quantifying and comparing features in high-dimensional datasets," in *Proc. of IV*. IEEE, 2008, pp. 240–245.
- [52] J. Seo and B. Shneiderman, "A rank-by-feature framework for interactive exploration of multidimensional data," *Inf. Vis.*, vol. 4, no. 2, pp. 96–113, Jun. 2005.
- [53] C. Turkay, P. Filzmoser, and H. Hauser, "Brushing dimensions — A dual visual analysis model for high-dimensional data," *IEEE TVCG*, vol. 17, no. 12, pp. 2591–2599, Dec. 2011.
- [54] S. Johansson and J. Johansson, "Interactive dimensionality reduction through user-defined combinations of quality metrics," *IEEE TVCG*, vol. 15, no. 6, pp. 993–1000, Nov.–Dec. 2009.
- [55] Y. Lu, R. Krüger, D. Thom, F. Wang, S. Koch, T. Ertl, and R. Maciejewski, "Integrating predictive analytics and social media," in *Proc. of IEEE VAST*. IEEE, 2014, pp. 193–202.
- [56] S. Gratzl, A. Lex, N. Gehlenborg, H. Pfister, and M. Streit, "LineUp: Visual analysis of multi-attribute rankings," *IEEE TVCG*, vol. 19, no. 12, pp. 2277–2286, Dec. 2013.
- [57] J. Yang, W. Peng, M. O. Ward, and E. A. Rundensteiner, "Interactive hierarchical dimension ordering, spacing and filtering for exploration of high dimensional datasets," in *Proc. of IEEE INFOVIS*. IEEE, 2003, pp. 105–112.
- [58] H. Lin, S. Gao, D. Gotz, F. Du, J. He, and N. Cao, "RCLens: Interactive rare category exploration and identification," *IEEE TVCG*, vol. 24, no. 7, pp. 2223–2237, Jul. 2018.
- [59] P. E. Rauber, R. R. O. d. Silva, S. Feringa, M. E. Celebi, A. X. Falcão, and A. C. Telea, "Interactive image feature selection aided by dimensionality reduction," in *Proc. of EuroVA*. The Eurographics Association, 2015.
- [60] P. Klemm, K. Lawonn, S. Glaßer, U. Niemann, K. Hegenscheid, H. Völzke, and B. Preim, "3D regression heat map analysis of population study data," *IEEE TVCG*, vol. 22, no. 1, pp. 81–90, Jan. 2016.
- [61] D. Dingen, M. van't Veer, P. Houthuizen, E. H. J. Mestrom, E. H. H. M. Korsten, A. R. A. Bouwman, and J. van Wijk, "RegressionExplorer: Interactive exploration of logistic regression models with subgroup analysis," *IEEE TVCG*, vol. 25, no. 1, pp. 246–255, Jan. 2019.
- [62] S. Barlowe, T. Zhang, Y. Liu, J. Yang, and D. Jacobs, "Multivariate visual explanation for high dimensional datasets," in *Proc. of IEEE VAST*. IEEE, 2008, pp. 147–154.
- [63] F. Hohman, K. Wongsuphasawat, M. B. Kery, and K. Patel, "Understanding and visualizing data iteration in machine learning," in *Proc. of ACM CHI*. ACM, 2020, pp. 1–13.
- [64] J. Bernard, M. Steiger, S. Widmer, H. Lücke-Tieke, T. May, and J. Kohlhammer, "Visual-interactive exploration of interesting multivariate relations in mixed research data sets," *CGF*, vol. 33, no. 3, pp. 291–300, Jun. 2014.
- [65] A. Chatzimparmpas, R. M. Martins, I. Jusufi, and A. Kerren, "A survey of surveys on the use of visualization for interpreting machine learning models," *Inf. Vis.*, vol. 19, no. 3, pp. 207–233, Jul. 2020.
- [66] A. Chatzimparmpas, R. M. Martins, I. Jusufi, K. Kucher, F. Rossi, and A. Kerren, "The state of the art in enhancing trust in machine learning models with the use of visualizations," *CGF*, vol. 39, no. 3, pp. 713–756, Jun. 2020.
- [67] A. Chatzimparmpas, R. M. Martins, and A. Kerren, "t-viSNE: Interactive assessment and interpretation of t-SNE projections," *IEEE TVCG*, vol. 26, no. 8, pp. 2696–2714, Aug. 2020.
- [68] A. Chatzimparmpas, R. M. Martins, K. Kucher, and A. Kerren, "Stack-GenVis: Alignment of data, algorithms, and models for stacking ensemble learning using performance metrics," *IEEE TVCG*, vol. 27, no. 2, pp. 1547–1557, Feb. 2021.
- [69] —, "VisEvol: Visual analytics to support hyperparameter search through evolutionary optimization," *CGF*, vol. 40, no. 3, pp. 201–214, Jun. 2021.
- [70] E. I. Altman, "Financial ratios, discriminant analysis and the prediction of corporate bankruptcy," *J. Finance*, vol. 23, no. 4, pp. 589–609, Sep. 1968.
- [71] Y. Li, Z.-F. Wu, J.-M. Liu, and Y.-Y. Tang, "Efficient feature selection for high-dimensional data using two-level filter," in *Proc. of ICMLC*, vol. 3. IEEE, 2004, pp. 1711–1716.
- [72] I. Dohoo, C. Ducrot, C. Fourichon, A. Donald, and D. Hurnik, "An overview of techniques for dealing with large numbers of independent variables in epidemiologic studies," *Prev. Vet. Med.*, vol. 29, no. 3, pp. 221–239, Jan. 1997.
- [73] F. Nogueira, "Bayesian Optimization," <https://git.io/vov5M>, 2014, accessed Dec. 8, 2021.
- [74] P. Cortez, A. Cerdeira, F. Almeida, T. Matos, and J. Reis, "Modeling wine preferences by data mining from physicochemical properties," *Decis. Support Syst.*, vol. 47, no. 4, pp. 547–553, Nov. 2009.
- [75] D. Dua and C. Graff, "UCI machine learning repository," <http://archive.ics.uci.edu/ml>, 2017, accessed Dec. 8, 2021.
- [76] A. Laughter and S. Omari, "A study of modeling techniques for prediction of wine quality," in *Intelligent Computing*. Springer, 2020, pp. 373–399.

- [77] S. Silva, B. Sousa Santos, and J. Madeira, "Using color in visualization: A survey," *C&G*, vol. 35, no. 2, pp. 320–333, 2011.
- [78] S. Jain and A. Saha, "Rank-based univariate feature selection methods on machine learning classifiers for code smell detection," *Evol. Intell.*, pp. 1–30, 2021.
- [79] G. Louppe, L. Wehenkel, A. Suter, and P. Geurts, "Understanding variable importances in forests of randomized trees," in *Proc. of NIPS*, vol. 1. Curran Associates Inc., 2013, pp. 431–439.
- [80] P. Radivojac, Z. Obradovic, A. K. Dunker, and S. Vucetic, "Feature selection filters based on the permutation test," in *Machine Learning: ECML 2004*. Springer, 2004, pp. 334–346.
- [81] L. Breiman, "Random forests," *Mach. Learn.*, vol. 45, pp. 5–32, Oct. 2001.
- [82] A. Janecek, W. Gansterer, M. Demel, and G. Ecker, "On the relationship between feature selection and classification accuracy," in *Proc. of FSDM at ECML/PKDD 2008*, vol. 4. PMLR, 2008, pp. 90–105.
- [83] J. S. Cramer, "The origins of logistic regression," Tinbergen Institute, Discussion Paper 2002-119/4, Nov. 2002.
- [84] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-Learn: Machine learning in Python," *JMLR*, vol. 12, pp. 2825–2830, Nov. 2011.
- [85] H. Jeon and S. Oh, "Hybrid-recursive feature elimination for efficient feature selection," *Appl. Sci.*, vol. 10, no. 9, 2020.
- [86] A. Lex, M. Streit, E. Kruijff, and D. Schmalstieg, "Caleydo: Design and evaluation of a visual analysis framework for gene expression data in its biological context," in *Proc. of IEEE PacificVis*. IEEE, 2010, pp. 57–64.
- [87] M. B. Eisen, P. T. Spellman, P. O. Brown, and D. Botstein, "Cluster analysis and display of genome-wide expression patterns," *PNAS*, vol. 95, no. 25, pp. 14 863–14 868, 1998.
- [88] K. Pearson, "Note on regression and inheritance in the case of two parents," *Proc. R. Soc. Lond.*, vol. 58, no. 347–352, pp. 240–242, 1895.
- [89] E. Vittinghoff, D. V. Glidden, S. C. Shiboski, and C. E. McCulloch, *Regression Methods in Biostatistics*. Springer, 2012.
- [90] S. Menard, *Applied Logistic Regression Analysis*. SAGE, 2001.
- [91] R. Johnston, K. Jones, and D. Manley, "Confounding and collinearity in regression analysis: A cautionary tale and an alternative procedure, illustrated by studies of British voting behaviour," *Qual. Quant.*, vol. 52, no. 4, pp. 1957–1976, Jul. 2018.
- [92] R. Gove, "It pays to be lazy: Reusing force approximations to compute better graph layouts faster," in *Proc. of FMT*. St. Pölten University of Applied Sciences, 2018, pp. 43–51.
- [93] J. P. Siebert, "Vehicle recognition using rule based methods," Turing Institute, Research Memorandum TIRM-87-018, Mar. 1987.
- [94] K. Mansouri, T. Ringsted, D. Ballabio, R. Todeschini, and V. Consonni, "Quantitative structure–activity relationship models for ready biodegradability of chemicals," *J. Chem. Inf. Model.*, vol. 53, no. 4, pp. 867–878, Apr. 2013.
- [95] E. Fix and J. L. Hodges, "Discriminatory analysis. Nonparametric discrimination: Consistency properties," *Int. Stat. Rev.*, vol. 57, no. 3, pp. 238–247, Dec. 1989.
- [96] R. G. Brereton and G. R. Lloyd, "Partial least squares discriminant analysis: Taking the magic away," *J. Chemom.*, vol. 28, no. 4, pp. 213–225, Apr. 2014.
- [97] Y. Ma, T. Xie, J. Li, and R. Maciejewski, "Explaining vulnerabilities to adversarial machine learning through visual analytics," *IEEE TVCG*, vol. 26, no. 1, pp. 1075–1085, Jan. 2020.
- [98] Y. Ming, P. Xu, F. Cheng, H. Qu, and L. Ren, "ProtoSteer: Steering deep sequence model with prototypes," *IEEE TVCG*, vol. 26, no. 1, pp. 238–248, Jan. 2020.
- [99] K. Xu, M. Xia, X. Mu, Y. Wang, and N. Cao, "EnsembleLens: Ensemble-based visual exploration of anomaly detection algorithms with multidimensional data," *IEEE TVCG*, vol. 25, no. 1, pp. 109–119, Jan. 2019.
- [100] T. Swearingen, W. Drevo, B. Cyphers, A. Cuesta-Infante, A. Ross, and K. Veeramachaneni, "ATM: A distributed, collaborative, scalable system for automated machine learning," in *Proc. of IEEE Big Data*. IEEE, 2017, pp. 151–162.
- [101] M. Espadoto, R. M. Martins, A. Kerren, N. S. T. Hirata, and A. C. Telea, "Toward a quantitative survey of dimension reduction techniques," *IEEE TVCG*, vol. 27, no. 3, pp. 2153–2173, Mar. 2021.
- [102] S. M. Moosavi, K. M. Jablonka, and B. Smit, "The role of machine learning in the understanding and design of materials," *JACS*, vol. 142, no. 48, pp. 20 273–20 287, Nov. 2020.
- [103] C. D. Stolper, A. Perer, and D. Gotz, "Progressive visual analytics: User-driven visual exploration of in-progress analytics," *IEEE TVCG*, vol. 20, no. 12, pp. 1653–1662, Dec. 2014.
- [104] C. Turkay, N. Pezzotti, C. Binnig, H. Strobel, B. Hammer, D. A. Keim, J. Fekete, T. Palpanas, Y. Wang, and F. Rusu, "Progressive data science: Potential and challenges," *ArXiv e-prints*, vol. 1812.08032, 2018.
- [105] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Proc. of NIPS*, 2017, pp. 4768–4777.
- [106] R. Schwartz-Ziv and A. Armon, "Tabular data: Deep learning is not all you need," *ArXiv e-prints*, vol. 2106.03253, 2021.
- [107] D. Babaev, M. Savchenko, A. Tuzhilin, and D. Umerenkov, "E.T.-RNN: Applying deep learning to credit loan applications," in *Proc. of ACM KDD*. ACM, 2019, pp. 2183–2190.
- [108] E.-S. M. El-Kenawy, A. Ibrahim, S. Mirjalili, M. M. Eid, and S. E. Hussein, "Novel feature selection and voting classifier algorithms for COVID-19 classification in CT images," *IEEE Access*, vol. 8, pp. 179 317–179 335, Sep. 2020.
- [109] "Vue.js — The progressive JavaScript framework," <https://vuejs.org/>, 2014, accessed Dec. 8, 2021.
- [110] "D3 — Data-driven documents," <https://d3js.org/>, 2011, accessed Dec. 8, 2021.
- [111] "Plotly — JavaScript open source graphing library," <https://plot.ly>, 2010, accessed Dec. 8, 2021.
- [112] "Flask — A micro web framework written in Python," <https://flask.palletsprojects.com/>, 2010, accessed Dec. 8, 2021.



Angelos Chatzimparmpas is PhD student within the ISOVIS research group and the Linnaeus University Centre for Data Intensive Sciences and Applications at the Department of Computer Science and Media Technology, Linnaeus University, Sweden. His main research interests include visual exploration of the inner parts and the quality of machine learning (ML) models with a specific focus on making complex ML models better understandable and explainable, as well as to provide reliable trust in the ML

models and their results.



Rafael M. Martins is Senior Lecturer at the Department of Computer Science and Media Technology at Linnaeus University, Sweden. His PhD research involved mainly the visual exploration of the quality of dimensionality reduction (DR) techniques, a topic he continues to investigate, in addition to other related research areas such as the interpretation of DR layouts and the application of DR techniques in different domains including software engineering and digital humanities.



Kostiantyn Kucher is Associate Senior Lecturer at the Department of Computer Science and Media Technology at Linnaeus University and part-time Visiting Assistant Lecturer at the Department of Science and Technology at Linköping University, Sweden. His main research interests include visual text and network analytics, explainable AI, and domain applications of information visualization and visual analytics, especially in digital humanities and information science.



Andreas Kerren is Professor of Information Visualization at Linköping University (LiU) and Linnaeus University (LNU), Sweden. He holds the Chair of Information Visualization at LiU and is head of the ISOVIS research group at LNU. In addition, he is an ELLIIT professor supported by the Excellence Center at Linköping–Lund in Information Technology and key researcher of the Linnaeus University Centre for Data Intensive Sciences and Applications. His main research interests include several areas of information visualization and visual analytics, especially visual network analytics, text visualization, and the use of visual analytics for explainable AI.