

CohEx: A Generalized Framework for Cohort Explanation

Fanyu Meng¹, Xin Liu¹, Zhaodan Kong¹, Xin Chen²

¹University of California, Davis

²Georgia Institute of Technology

fymeng@ucdavis.edu, xinliu@ucdavis.edu, zdkong@ucdavis.edu, xinchen@gatech.edu

Abstract

eXplainable Artificial Intelligence (XAI) has garnered significant attention for enhancing transparency and trust in machine learning models. However, the scopes of most existing explanation techniques focus either on offering a holistic view of the target model (global explanation) or on individual instances (local explanation), while the middle ground, i.e., cohort-based explanation, is less explored. Cohort explanations offer insights into the target model’s behavior on a specific group or cohort of instances, enabling a deeper understanding of model decisions within a defined context. In this paper, we discuss the unique challenges and opportunities associated with measuring cohort explanations, define their desired properties, and create a generalized framework for generating cohort explanations based on supervised clustering.

Code — <https://github.com/fy-meng/cohex>

Extended version — <https://arxiv.org/abs/2410.13190>

1 Introduction & Motivation

eXplainable Artificial Intelligence aims to improve the existing machine learning system by providing transparent explanations for model predictions and decisions (Arrieta et al. 2020; Adadi and Berrada 2018). The *scope* of explanation refers to the degree to which the explanation provided is generalized. The scope, as defined by (Sokol and Flach 2020), is usually (1) *global explanation*, which aims to explain the target model comprehensively; or (2) *local explanation*, which focuses on explaining the behavior in one instance of model prediction. However, there is a less explored middle ground, namely (3) *cohort explanation*, or sometimes also called *regional explanation*, which generalizes the explanation on a subset of a dataset, a subspace in the model’s input space or decision space. Cohort explanation can be seen as a spectrum that generalizes both local and global explanation methods, where global and local explanations are at the two extremes of conciseness and faithfulness. This generalization allows users to uncover patterns, biases, and context-specific insights that may not be apparent when examining individual instances in isolation or the model as a whole.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Lack of Robustness of Feature Attribution Methods

Many existing explanation methods in XAI, notably feature attribution methods, face limitations regarding robustness. The output can be sensitive to small perturbations in input or model parameters, leading to unstable or unreliable insights (Alvarez-Melis and Jaakkola 2018; Slack et al. 2020). Some explainer would be produce biased output under simple transformations, failing to only capture relevant information for decision-making (Kindermans et al. 2019). A common practice to improve robustness in machine learning is via bootstrapping - averaging the output of multiple weak learners to reduce overfitting (Breiman 1996, 2001; Sagi and Rokach 2018). A similar approach could also be beneficial for XAI. By averaging the explanations on a group of similar samples, we could reduce the noises in explanations, providing more stable, interpretable insights that account for patterns across relevant subsets of data.

Faithfulness vs. Conciseness Global XAI methods provide one general explanation, but may not generalize well on subgroups; local methods assign a specific explanation for each data sample, but it is not concise and may not help to understand the entire model. Cohort explanation methods strive for a balance between the two, with the aim of creating meaningful splittings of the feature space and assigning explanations to each cohort that generalize well within each of the groups. Thus, the cohort explanation generates “*contextful*” information: the explanation of each cohort recognizes that it can be applied to a limited region. This helps users better understand the conditions that must be maintained for the output explanation to be valid and to offer a more concise and transparent explanation (Sokol and Flach 2020).

To illustrate this motivation, consider a simplified medical patient classification task in Fig. 1a. The two features are age and the fraction of family members with this disease. The likelihood of the condition occurring increases with respect to both features. The target model is a decision tree with a depth of two. Note that we treat the decision tree as a black-box, i.e. the explainer does not know that the model is a tree.

To generate cohort explanation to the tree, we can first apply a simplified version of the counterfactual-based local explanation (Wachter, Mittelstadt, and Russell 2017). We consider the importance of each feature as the minimum perturbation on that feature to change the prediction,

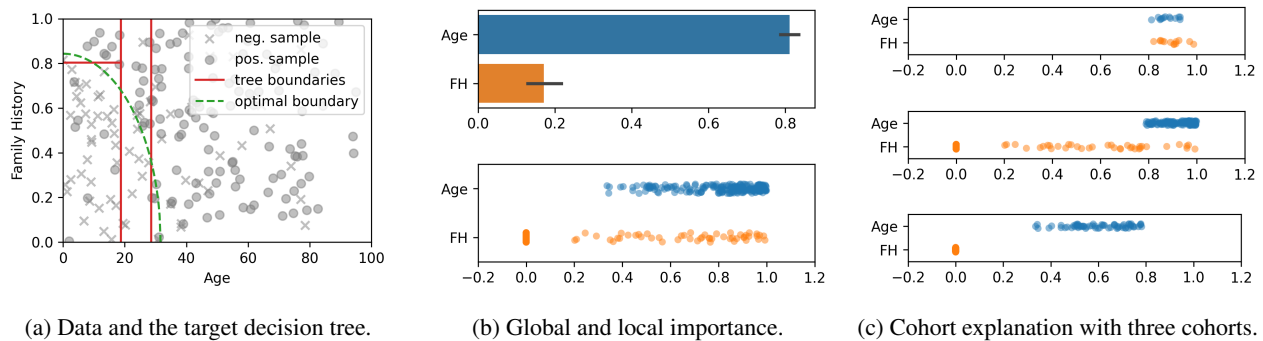


Figure 1: A motivating example: a synthesized patient binary classification problem.

and a small perturbation signals a high importance. Fig. 1b shows the global and distribution local importance on this model. These two types of explanations illustrates that in general, age has a higher importance, but they cannot used to explain whether they are patterns within the explanations. However, such pattern exists: if we separate the samples into three groups: (1) young patients with high family history, (2) middle-aged or young patients with low family history, and (3) older patients, the distribution of feature importance are distinct between the groups. Fig. 1c shows this cohort explanation, and we can provide a deeper analysis that the features are (1) both are significant; (2) only age is significant; (3) neither is significant respectively on the three groups. Naturally, this analysis correspond to the slope of the optimal decision boundary in Fig. 1a. Cohort explanation simplifies the mechanisms of the model into a few regions in which different features are the main factor behind decision making. Additionally, it provides insights that cannot be provided by global or local explanations.

Our contributions in this paper are as follows:

1. We propose desired properties of cohort explanation and devise their evaluation metrics;
2. We develop a generalized framework to convert existing data-driven local feature importance methods to cohort explanations;
3. We conduct experiments on several examples and demonstrate the superior performance of our algorithm over existing benchmarks based on the proposed metrics.

2 Related Works

Cohort Explanation Cohort explainability aims to strike a balance between local and global methods, which provide explanations based on a group of instances. It can be classified into inherent and post-hoc methods. Inherent cohort explanation, sometimes referred to as a local surrogate model, involves dividing the feature space and developing a local, interpretable model for each subspace. Examples of this approach include Hu et al. (2020), which constructs local-surrogate linear or generalized additive models by recursively partitioning the feature space.

Post-hoc cohort explanation methods partition the feature space and provide an explanation for each cluster. In general, such explainers run local explanation methods on a

dataset and then cluster the samples based on the similarity between the local explanations. Such methods include Cavus, Stando, and Biecek (2023), which averages local SHAP values within predefined groups. Other methods employ techniques to automatically identify the cohorts, such as VINE, which clusters individual conditional expectation (ICE) via unsupervised clustering methods (Britton 2019). Similarly, REPID (Herbinger, Bischl, and Casalicchio 2022) and Molnar et al. (2023), which use tree-based partitioning to generate the cohorts either based on ICE or conditional permutation feature importance. Additionally, GADGET extends previous methods by incorporating functional decomposition during partitioning, allowing for the consideration of more features (Herbinger, Bischl, and Casalicchio 2023). In contrast to the partitioning aspect, there are also works that focuses on how to aggregate local explanations, such as GALE (Van Der Linden, Haned, and Kanoulas 2019) which uses homogeneity to re-weight LIME importance before aggregating for multi-class classification problems.

A common challenge for XAI in general is a lack of evaluation metrics to compare different methods, since the explanation template varies between methods. The most common metrics used to qualify cohort explanation is generalizability/coherence, the similarity between local explanations within cohorts. In this paper, we propose additional desiderata and evaluation metrics for cohort explanation. We also introduce a supervised-clustering-based approach that achieves better performance based on the metrics.

Subgroup Discovery Subgroup discovery is a prominent area in data mining and machine learning that focuses on identifying interesting and interpretable subgroups within a dataset, which is valuable for uncovering hidden structures and actionable insights in various domains (Helal 2016). Such works include Sutton et al. (2020) which uses rule-based partitioning based on the loss of surrogate models to identify interesting regions. Similarly, Hedderich et al. (2022) also use a rule-based algorithm on natural language models to recognize regions with significant response to certain variables. These methods can be thought of as a type of *inherent* cohort explanation, though the identified clusters of interest may not cover the entire feature space.

Explainable Clustering Explainable clustering aims to enhance interpretability of clustering by providing understandable explanations for the derived clusters. For example, Moshkovitz et al. (2020) modifies the results from k -means into decision trees to increase the interpretability of the clusters. L’Yi et al. (2015) creates a visualization tool to compare and analyze clustering results. Guilbert, Vrain, and Dao (2024) formulate a constraint optimization program to clustering to allow for the injection of domain knowledge. Though both domains utilize explainability and clustering, their goals are different. Cohort explanation aims to apply clustering to simplify the result of local explanations; contrastively, explainable clustering’s goal is to introduce explainability methods to existing clustering algorithms.

3 Notation

We will use the following notation throughout the paper.

- M : target model, or the model to be explained;
- x : a specific sample;
- X or X_j : a group of samples. We will use X to denote the entire evaluation dataset and X_j to denote the j -th cohort. $\{X_1, \dots, X_k\}$ is referred as a *cohort definition*, which partitions the dataset X into k cohorts. The cohort definition must satisfy $\bigcup_{j=1}^k X_j = X$ and $X_j \cap X_l = \emptyset$ for $j \neq l$;
- ω : a local explanation method. It is data-driven and outputs local feature importance. Specifically, $\omega_M(X, x)$ represents the result of evaluating ω for the sample x , using the model M and the dataset X as context. The method should treat the target model M as a black-box, meaning that ω does not depend on the specific architecture of M ;
- e_j : the explanation for cohort j ;
- $a_x \in \{1, \dots, k\}$: the cohort assignment for sample x . In other words, $a_x = j \Leftrightarrow x \in X_j$.

4 Desiderata

Desired properties of XAI have been well analyzed in works such as Adadi and Berrada (2018), Arrieta et al. (2020) and Sokol and Flach (2020). In this section, however, we focus on properties that are unique to cohort explanations. Table 1 lists these properties and compares them desired properties of cohort explanations with their counterparts on a global or local scale.

Generalizability Global explanation is prone to aggregation bias, where the average importance may not reflect the representative population in the dataset (Herbinger, Bischl, and Casalicchio 2023). Cohort explanation should aim to reduce such bias and to generate explanations that generalize well in each of the cohorts. If we apply a local method to all samples in a specific region, the explanation should be of low variance and close to the cohort explanation.

To evaluate generalizability, we use the average error w.r.t. local importance as the metric. Given a dataset X , a local importance method ω , cohorts $\{X_j\}$ and the respec-

tive explanations $\{\bar{e}_j\}$, the loss is defined as

$$\mathcal{L}_{\text{generalizability}} = \frac{1}{|X|} \sum_{j=1}^k \sum_{x \in X_j} \|\omega_M(X_j, x) - \bar{e}_j\|_2^2 \quad (1)$$

Conciseness One key goal of cohort explanation, especially cohort explanations, is to generate a concise explanation. It is evaluated using the number of cohorts and should be low without significantly sacrificing generalizability.

Disjoint Cohorts Cohort explanation should yield meaningful cohort definitions. Although cohorts can be defined using samples, the regions represented by these assignments should divide the feature space into disjoint subspaces. Otherwise, it would be difficult to describe what a cohort represents. This adds an additional challenge in defining the cohorts in addition to the generalizability criteria.

Cohort Locality The explanation of a cohort should emphasize the decision-making *local* or *specific* to this region. That is, if two models have the same behavior within a given cohort X (i.e., both models predict the labels for all samples in X), but they differ outside X , the cohort explanation for the two models in this fixed cohort should be the same.

However, this property does not hold in general for local explanation methods. Popular methods such as LIME (Ribeiro, Singh, and Guestrin 2016) and KernelSHAP (Lundberg and Lee 2017) require sampling, which may sample outside the cohort X . In this case, a good cohort explanation should try to preserve locality by stabilizing the explanation if a cohort is fixed.

To evaluate locality, given a fixed cohort X_j , we create alternative models \tilde{M}_j that behave the same as the original model M within X_j , but have the probability to be random otherwise. Specifically,

$$\tilde{M}_j(x) = \begin{cases} M(x), & x \in X_j, \\ M(x), & x \notin X_j, \text{ w/ probability } 1 - p, \\ \text{random}, & x \notin X_j, \text{ w/ probability } p. \end{cases} \quad (2)$$

where the last case denotes a random class for classification tasks or a random value sampled from the distribution of labels for regression tasks. We then consider the average difference in explanation on cohort X_j w.r.t. a probability p :

$$\mathcal{L}_{\text{locality}}(p) = \mathbb{E} \left[\sum_{j=1}^k \|e_j - \tilde{e}_j\|_2^2 \right]. \quad (3)$$

where e_j is evaluated on target M and \tilde{e}_j is on \tilde{M}_j , and the expectation is over the randomness from p . An ideal cohort explanation should have a low locality loss and thus focus on localized information.

Stability The explanation should be robust w.r.t. (1) the target model and (2) the cohort definitions. The cohort definitions should be stable and the importance should not drastically differ when the definitions are slightly changed.

To evaluate cohort stability, we tested the robustness of cohort definitions against the same input and then use the

	Global	Cohort	Local
Generalizability	w.r.t the whole model	w.r.t. the cohorts	-
Conciseness	most concise	adjustable	least concise
Disjoint Cohorts	-	applicable	-
Cohort Locality	-	applicable	-
Stability	w.r.t. target model	w.r.t target model and cohort definitions	w.r.t. target model

Table 1: Desired properties of explanations at different scope.

Adjusted Rand Index (ARI) (Santos and Embrechts 2009) to compare the similarity between two assignments. To evaluate, we conduct t run on the same dataset, and the metric is defined as

$$\mathcal{S}_{\text{cohort stability}} = \frac{1}{t} \sum_{i=2}^t \text{ARI}(\{a_x^1\}, \{a_x^i\}). \quad (4)$$

where a_x^i denotes the assignment computed on the i -th run.

To evaluate importance stability, we consider the change in explanation in a particular cohort if a new random sample is added to this cohort. Ideally, we want all samples to have localized, stable local importance scores. Thus, adding a sample to a cohort should not drastically affect its average importance. Therefore, we evaluate importance stability as

$$\mathcal{L}_{\text{importance stability}} = \mathbb{E}_x \left[\sum_{j=1}^k \|e_j - e_{X_j \cup \{x\}}\|_2^2 \right]. \quad (5)$$

where $e_{X_j \cup \{x\}}$ is evaluated on the artificial cohort by appending sample x to the cohort X_j . Note that Eqs. 4 and 5 measure two different aspects of cohort explanations. Cohort stability estimates the robustness on how the data are clustered, while importance stability concerns the robustness of explanations on predefined cohorts.

5 Problem Definition

We emphasize on generate *post-hoc* cohort explanation through existing *local*, *data driven*, *feature importance* explanations. We formulate the cohort explanation as an optimization problem, with the goal of minimizing the generalizability loss in Eq. 1 while penalizing for having too many cohorts. Consider a local data-driven explanation method $\omega_M(X, x)$, where M is the target model, X is a context dataset and x is a sample. The problem is defined as

$$\min_{k, c_1, \dots, c_k} \left(\frac{1}{|X|} \sum_{j=1}^k \sum_{x \in X_j} \|\omega_M(X_j, x) - e_j\|_2^2 \right) + \lambda \sqrt{\max \left(\frac{k - k^*}{|X|}, 0 \right)}. \quad (6)$$

where k is the number of cohorts, c_j is the centroid that defines the cohort X_j :

$$X_j = \{x \in X \mid \|x - c_j\|_2 \leq \|x - c_l\|_2, \forall l \neq j\}. \quad (7)$$




Figure 2: Two 1-D regression models with the same behavior on a specific cohort. The horizontal axis represents the feature, and the vertical axis represents the model output. The orange circle denotes the cohort.

e_j is explanation for the cohort X_j , defined as the average local importance across all samples in X_j :

$$e_j = \frac{1}{|X_j|} \omega_M(X_j, x). \quad (8)$$

λ and k^* are hyperparameters and k^* is the number of expected cohorts. k^* should be chosen as the desired granularity of the explanations, similar to how the number of clusters is chosen in clustering algorithms. Note that setting $k^* = 1$ with a large enough λ would generate global explanation, while setting it to the number of samples would be equivalent to local explanation. The first component in Eq. 6 represents the *generalizability* criteria, while the second component corresponds to *conciseness*. Choosing λ would shift the balance between these two core desired properties of cohort explanations. Note that both the explanation in Eq. 8 and the objective function in Eq. 6 use the cohort X_j instead of the whole dataset X to compute local importance. This aims to satisfy the *cohort locality* criteria by reducing information leakage from outside the cohorts. Additionally, since cohorts are defined using centroids as shown in Eq. 7, the partitions are guaranteed to be *disjoint*.

6 Challenges

We now discuss two major challenges unique to cohort explanations, which our algorithm aims to overcome.

Naive Average of Local Explanation Does Not Suffice

Given a particular cohort definition, the naive way to compute the cohort explanation is to average the local explanation of all members. However, this may fail the *cohort locality* criteria discussed in Sec. 4, since local explanations still contain information outside the cohort. For example, consider two regression models trained on the same tasks as shown in Fig. 2. On the cohort represented by the orange circle, the models have the same behavior. Naturally,

Algorithm 1: Generalized cohort explanation conversion framework (CohEx)

Input : Number of iterations n , Expected number of cohorts k , Target model M , Dataset X , Data-driven local explanation method ω , Supervised clustering algorithm g .

Output: Cohort assignment $a_x \in [1, k], \forall x \in X$, Cohort explanations $\bar{w}_j, j \in [1, k]$.

```
1 for  $i = 1, \dots, n$  do
2   randomly select  $k$  centroids  $c_1, \dots, c_k$  from  $X$ ;
3   for  $x \in X$  do
4      $a_x^i \leftarrow \arg \min_{1 \leq a \leq k} \|x - c_a\|_2$ ;          /* assign each sample to the closest centroids */
5   end
6   repeat
7     for  $j = 1, \dots, k$  do
8        $X_j \leftarrow \{x | a_x^i = j\}$ ;                          /* The samples in this cohort */
9       for  $x \in X_j$  do
10         $w_x^i \leftarrow \omega_M(X_j, x)$ ;    /* Recompute explanations using samples only in the cohort */
11      end
12       $\bar{e}_j^i \leftarrow (\sum_{x \in X_j} w_x^i) / |X_j|$ ;          /* Compute the average local explanations */
13    end
14     $k, \{a_x^i\} \leftarrow g(k, X, \{w_x^i\})$ ;          /* Recluster using the new explanations */
15  until clustering loss does not decrease for  $t$  iterations;
16 end
17 return the best  $\{a_x^i\}$  and  $\{\bar{e}_j^i\}$  with the lowest clustering loss
```

we expect the cohort explanations to this cohort to be identical between the models as well. However, this is usually not the case based on the naive method. Data-driven local explanations either require a context dataset (e.g. SHAP) or need sampling (e.g. LIME) that would exceed the boundary of the cohort. If we approach cohort explanation hierarchically - first compute all local importance, then apply clustering and/or transformations, which is true for most existing literature - the different behavior outside this cohort would leak into it, causing the cohort explanation to diverge. To address the issue without constraining the underlying local explainer, we propose an iterative framework for the conversion of local explanations that takes into account leakage and evaluates its locality based on Eq. 3.

Unifying Cohorts and Explanations Cohort explanations contain two distinct tasks: (1) creating definitions of the cohorts and (2) assigning an appropriate explanation on each of the cohorts. As discussed in Sec. 4, all the generalizability, conciseness, and representativeness properties put constraints on the definition of the cohorts.

Additionally, this is not a sequential task: one cannot first decide the cohort definitions and then compute their explanations, since the two tasks are interdependent. The cohort locality properties ensure that the explanation needs to be re-evaluated if the cohort definitions change, and the definitions depend on the explanations to achieve generalizability.

7 Approach

To solve these two problems, we propose a generalized cohort explanation conversion framework (CohEx) to adapt data-driven, local, feature importance methods into cohort explanation. We allow modification to local importance to account for locality, and design a feedback loop between recomputing importance and reclustering to refine cohort definitions and their importances.

Algorithm 1 shows the proposed method CohEx. The core idea is to provide only *data-driven* local explanation methods with information within the cohort. In each iteration, we first randomly create cohort assignments through randomly sampled centroids at lines 2-4. Then we recompute all importances at line 9, using only samples within each cohort. Then we run a supervised clustering algorithm to create new cohort assignments using the new importances at line 13. This is repeated until the clustering loss does not decrease for a certain number of iterations. The whole process is then repeated n times and the best result is reported.

To generate the cohort definitions, we choose supervised clustering. Compared to traditional clustering, it assumes that data are labeled and uses both features X and labels y to group data. Each data sample is clustered into the closest centroids in the *feature space*, while the goodness of an assignment is based on the *labels*, or in our case, feature importance (Eick et al., 2004). This ensures that the cohort definitions are always closed and well defined, while constraining local explanations within the cohort to be similar for the generalizability criteria. Specifically, in our evaluations, we use Single Representative Insertion/Deletion Steepest Decent Hill Climbing with Randomized Restart (SRIDHCR) as the supervised learning method, a greedy algorithm that attempts to solve the problem by randomly adding or removing centroids (Eick, Zeidat, and Zhao 2004). Note that standard, *unsupervised* clustering algorithms such as HDBSCAN does not apply here since we need a *supervised* mechanism to cluster both the feature values and the importance scores simultaneously.

For the convergence of the algorithm, since the cohort assignments depend on the base local explanation methods, the iterative loop between recomputing importance and reclustering is not guaranteed to converge. We resolved to use clustering loss as a stopping criterion. The algorithm terminates

if the loss does not decrease for a set number of iterations, and we repeat this process n times with different initialization. Thus, the algorithm is guaranteed to terminate, though the convergence speed may depend on the specifics of the chosen base method. The complexity analysis of CohEx and discussion related to non-data-driven methods can be found in Appendix A and B in the extended version.

8 Evaluation

To evaluate the proposed method, we experiment on three scenarios: (1) the patient classification problem in Sect. 1, a small scale example to demonstrate the importance of using supervised clustering and revising importance; (2) bike sharing, a classical ML regression problem to predict the hourly number of rented bikes (Fanaee-T 2013), and (3) MNIST digit classification (Deng 2012), a vision-based task with deep learning to mimic more realistic and complex scenarios. For baselines, we compare with the following methods:

- **VINE**, which applies supervised clustering on the local importances scores (Britton 2019);
- **REPID**, a tree-based partitioning method (Herbinger, Bischl, and Casalicchio 2022). For the sake of comparison, we modify the algorithm to use a local explainer that is consistent with other method, instead of the original partial dependency explainer¹;
- **GALE**, a homogeneity-based re-weighting mechanism to improve the quality of importance aggregation in classification problems (Van Der Linden, Haned, and Kanoulas 2019). GALE is not a cohort explainer by definition, but we feed the GALE importance into VINE and REPID to assess the quality of the reweighed importance. Note that GALE is not applicable to regression problems;
- **Hierarchical cohort explanation**: first compute the local importance, then run supervised clustering with SRIDHCR once. The difference between this method and the proposed CohEx is that it does not iteratively recompute the local importance scores.

8.1 Synthetic Patient Classification

We continue to use the motivating example in Sect. 1. For this evaluation, the expected number of cohorts k^* is set to 4, and the maximum depth of the tree in REPID is set to 2 so that the number of cohorts is consistent for all methods. We use LIME as the base method, specifically since it does not satisfy the cohort locality criteria and allows us to compare the effect of different algorithms on locality. Fig. 3 and Table 2 show the result of running the algorithms on this example. The cohort partitioning and explanations are shown in Fig. 3. Additionally, the evaluation against different k^* can be found in Appendix C.1 in the extended version.

8.2 Bike Sharing

This experiment focuses on a realistic machine learning regression problem. The bike sharing task asks to predict

¹Python’s SHAP package implements an auto-cohort feature. Though no reference is provided, we believe that it is equivalent to REPID.

the hourly number of rented bikes between years 2011 and 2012 in the Capital bikeshare system based on seasonal and weather information. For the target model, we train an XGBoost model (Chen and Guestrin 2016) on this dataset and achieve a R^2 score of 0.9453. We continue to use $k^* = 4$ or a depth of two for a similar number of cohorts and use SHAP as the local explainer.

Fig. 4 shows the result of running our proposed algorithm on the XGBoost model, including both the explanations and the characteristics to each cohort. Comparison with other methods can be found in the Appendix. C.2 in the extended version. Our algorithm can detect clusters with distinct characteristics and importance signals, while maintaining a coherent explanation within each cohort.

8.3 Digit Classification

We evaluate based on MNIST (Deng 2012). Note that the MNIST inputs are cropped, centered, and relatively simple, so measuring the distance between importance heatmaps is meaningful. We built the target model as a neural network with two convolutional layers and two fully-connected layers and achieved an accuracy of 93.93% on the test dataset. More details on the model and a visualization of the cohort explanation can be found in Appendix. C.3 in the extended version.

8.4 Evaluation Metrics and Analysis

Generalizability & Conciseness CohEx achieves the best performance based on the objective function in Eq. 6. Note that all methods yield $k = 4$, so the conciseness loss for all is zero. VINE yields a low generalizability loss since its clustering criterion focuses only on feature importance, while all other methods need to consider the distance in the feature space. However, CohEx is able to exceed VINE since they use different local importance scores w_x . Through the recomputing of importance, CohEx edits the importance by putting more emphasis on local structures, allowing the importance scores in each cohort to have lower variance. Introducing GALE re-weighting improves the performance of VINE and REPID since it also recomputes importance, but it does not perform well on low-dimensional features as shown in Fig. 3 since the importance of the feature with the lowest homogeneity will always be set to zero.

Cohort Locality CohEx yields a significantly lower locality loss compared to the other methods according to Eq. 3, which is as expected since CohEx introduces an iterative refinement process specifically for locality. To further justify the approach, in the patient classification problem, CohEx identifies a clustering that mostly match with the three clusters in our ideal explanation proposed in Sect. 1. In particular, CohEx is the only algorithm that identifies a region (the blue cohort in Fig. 3f) in which family history is the more important feature. Note that for almost all samples, the static LIME importance of age is higher than family history similar to the result in Fig. 1b. Therefore, since other methods are based on a static set of local importance scores, they cannot find any regions with a meaningful family history importance. We further justify that such a region should




Figure 3: Running the six cohort explanation algorithms on the motivational example in Fig. 1. Each color represents a different cohort. The labels in the figures denotes cohort explanation ($e_{\text{age}}, e_{\text{family}}$) of the cohort of the corresponding color.

	Clustering Penalty (\downarrow)	Closed Cohort	Cohort Locality (\downarrow)	Cohort Stability (\uparrow)	Importance Stability (\downarrow)
Synthetic Patient Classification (LIME)					
VINE	0.0011 ± 0.0001	×	0.0273 ± 0.0076	$42.33\% \pm 0.78\%$	0.0131 ± 0.0039
VINE+GALE	0.0006 ± 0.0000	×	0.1011 ± 0.0046	$42.96\% \pm 5.38\%$	0.0481 ± 0.0070
REPID	0.0012 ± 0.0000	✓	0.0301 ± 0.0021	$47.37\% \pm 0.56\%$	0.0169 ± 0.0015
REPID+GALE	0.0005 ± 0.0000	✓	0.0978 ± 0.0039	$53.52\% \pm 9.92\%$	0.0415 ± 0.0048
Hierarchical	0.0018 ± 0.0003	✓	0.0157 ± 0.0050	$36.74\% \pm 9.57\%$	0.0167 ± 0.0019
CohEx	0.0006 ± 0.0003	✓	0.0052 ± 0.0088	$40.76\% \pm 10.34\%$	0.0027 ± 0.0033
Bike Sharing (SHAP)					
VINE	0.6641 ± 0.0198	×	3.5983 ± 0.3986	$79.89\% \pm 17.87\%$	3.0773 ± 0.5905
REPID	1.1633 ± 0.0000	✓	2.5280 ± 0.0091	$100.00\% \pm 0.00\%$	2.1123 ± 0.0046
Hierarchical	0.4796 ± 0.1173	✓	3.1891 ± 0.0187	$39.17\% \pm 10.11\%$	2.5951 ± 0.0061
CohEx	0.3196 ± 0.1018	✓	0.5034 ± 0.2214	$43.68\% \pm 14.23\%$	0.1022 ± 0.0886
MNIST (DeepLIFT)					
VINE	1.5450 ± 0.1698	×	0.0528 ± 0.0253	$80.98\% \pm 22.07\%$	4.1130 ± 0.932
VINE+GALE	1.4729 ± 0.0918	×	0.1417 ± 0.0253	$81.03\% \pm 19.28\%$	2.1209 ± 0.634
REPID	1.3462 ± 0.0000	✓	0.0531 ± 0.0198	$100.00\% \pm 0.00\%$	1.4889 ± 0.0060
REPID+GALE	1.3290 ± 0.0000	✓	0.0999 ± 0.0203	$100.00\% \pm 0.00\%$	0.4139 ± 0.0046
Hierarchical	1.9371 ± 0.3825	✓	0.0127 ± 0.0047	$31.50\% \pm 32.08\%$	0.0956 ± 0.0475
CohEx	1.3112 ± 0.3037	✓	0.0010 ± 0.0019	$33.22\% \pm 29.13\%$	0.0210 ± 0.0419

Table 2: Evaluating the cohort explanation algorithms in the motivational example shown in Fig. 1 and the bike sharing dataset. The (\downarrow) symbol after the criteria represents that the lower the better, and (\uparrow) represents the higher the better.

exist, since if we only consider said region, the horizontal decision boundary is the dominant decision boundary. This shows that CohEx’s explanation is more *local* to its cohorts. We also note that applying GALE would harm locality, since the re-weighting process smoothes the importance over the whole dataset, specifically contradicts the goal of locality.

Disjoint Cohorts Only VINE fails this property since it clusters solely based on importance and it’s very likely that some distant samples happen to have similar explanations and thus are partitioned together. All other method guarantees that the cohorts can be separated into disjoint areas by hyperplanes.

Cohort Stability In all scenarios, REPID achieves higher ARI scores based on Eq. 4, which shows that our supervised clustering-based approaches introduce variance. We consider this variance to be inherent to the methods, as both the recomputing process and SRIDHCR are trial-and-error by nature. The question of how to improve the stability remains a challenge for future work.

Importance Stability It measures the change in cohort explanations, if a random sample is added to the cohort. For fair comparisons, we allow the baseline methods to recompute the importances similar to CohEx after adding the sample, which greatly improve their stabilities. CohEx still exceeds the baselines. This demonstrates that even though the partitions of CohEx is not as robust based on the previous *cohort stability* metric, on a fixed cohort, CohEx’s explanation is much less gullible to changes or noises in the dataset.

9 Conclusion

In this paper, we identify a less explored piece in the XAI literature, i.e., cohort explanation. We illustrate the desired properties and challenges unique to the cohort explanation. We proposed a unified framework CohEx to generate such explanation by transferring existing local explanation methods while preserving generalizability and locality. It showed an increase in most metrics compared to baseline methods.

The current limitations to the algorithm include (1) the

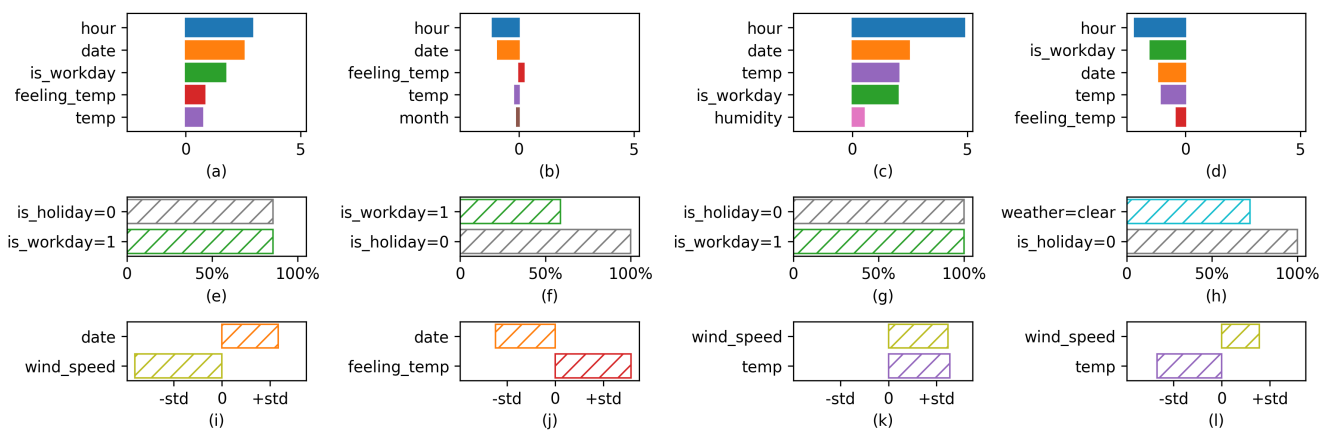


Figure 4: Running CohEx on the bike sharing dataset. Each column represents a cohort. (a-d) Cohort importance of the top five features in each cohort; (e-h) the top two most homogeneous categorical features in each cohort, measured in homogeneity; (i-l) the top two most abnormal continuous features, measured in the centroid difference w.r.t. dataset mean in standard deviations.

stability of cohort definitions; (2) the high time complexity in more complex scenarios, such as vision tasks; (3) the limits of the framework to local data-driven explainers. For future directions, alternatives can be considered, such as using non-data-driven, non-local, or non-post-hoc explanation methods to generate cohort explanation, increasing generalizability to more complex tasks. Additionally, the meaningfulness of the cohort definitions can be potentially increased through enforcing stricter constraints on the definitions, such as only partitioning the samples using tags, rules, or hyper-planes.

References

- Adadi, A.; and Berrada, M. 2018. Peeking inside the black-box: a survey on explainable artificial intelligence (XAI). *IEEE access*, 6: 52138–52160.
- Alvarez-Melis, D.; and Jaakkola, T. S. 2018. On the robustness of interpretability methods. *arXiv preprint arXiv:1806.08049*.
- Arrieta, A. B.; Díaz-Rodríguez, N.; Del Ser, J.; Bennetot, A.; Tabik, S.; Barbado, A.; García, S.; Gil-López, S.; Molina, D.; Benjamins, R.; et al. 2020. Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information fusion*, 58: 82–115.
- Breiman, L. 1996. Bagging predictors. *Machine learning*, 24: 123–140.
- Breiman, L. 2001. Random forests. *Machine learning*, 45: 5–32.
- Britton, M. 2019. Vine: Visualizing statistical interactions in black box models. *arXiv preprint arXiv:1904.00561*.
- Cavus, M.; Stando, A.; and Biecek, P. 2023. Glocal Explanations of Expected Goal Models in Soccer. *arXiv preprint arXiv:2308.15559*.
- Chen, T.; and Guestrin, C. 2016. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, 785–794.
- Deng, L. 2012. The mnist database of handwritten digit images for machine learning research [best of the web]. *IEEE signal processing magazine*, 29(6): 141–142.
- Eick, C. F.; Zeidat, N.; and Zhao, Z. 2004. Supervised clustering-algorithms and benefits. In *16th IEEE international conference on tools with artificial intelligence*, 774–776. IEEE.
- Fanaee-T, H. 2013. Bike Sharing Dataset. <https://doi.org/10.24432/C5W894>. Accessed: 2024-11-10.
- Guilbert, M.; Vrain, C.; and Dao, T.-B.-H. 2024. Towards Explainable Clustering: A Constrained Declarative based Approach. *arXiv preprint arXiv:2403.18101*.
- Hedderich, M. A.; Fischer, J.; Klakow, D.; and Vreeken, J. 2022. Label-Descriptive Patterns and Their Application to Characterizing Classification Errors. In Chaudhuri, K.; Jegelka, S.; Song, L.; Szepesvari, C.; Niu, G.; and Sabato, S., eds., *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, 8691–8707. PMLR.
- Helal, S. 2016. Subgroup discovery algorithms: a survey and empirical evaluation. *Journal of computer science and technology*, 31: 561–576.
- Herbinger, J.; Bischl, B.; and Casalicchio, G. 2022. REPID: Regional Effect Plots with implicit Interaction Detection. In *International Conference on Artificial Intelligence and Statistics*, 10209–10233. PMLR.
- Herbinger, J.; Bischl, B.; and Casalicchio, G. 2023. Decomposing Global Feature Effects Based on Feature Interactions. *arXiv preprint arXiv:2306.00541*.
- Hu, L.; Chen, J.; Nair, V. N.; and Sudjianto, A. 2020. Surrogate locally-interpretable models with supervised machine learning algorithms. *arXiv preprint arXiv:2007.14528*.
- Kindermans, P.-J.; Hooker, S.; Adebayo, J.; Alber, M.; Schütt, K. T.; Dähne, S.; Erhan, D.; and Kim, B. 2019. The (un) reliability of saliency methods. *Explainable AI: Interpreting, explaining and visualizing deep learning*, 267–280.

- Lundberg, S. M.; and Lee, S.-I. 2017. A unified approach to interpreting model predictions. *Advances in neural information processing systems*, 30.
- L'Yi, S.; Ko, B.; Shin, D.; Cho, Y.-J.; Lee, J.; Kim, B.; and Seo, J. 2015. XCluSim: a visual analytics tool for interactively comparing multiple clustering results of bioinformatics data. *BMC bioinformatics*, 16: 1–15.
- Molnar, C.; König, G.; Bischl, B.; and Casalicchio, G. 2023. Model-agnostic feature importance and effects with dependent features: a conditional subgroup approach. *Data Mining and Knowledge Discovery*, 1–39.
- Moshkovitz, M.; Dasgupta, S.; Rashtchian, C.; and Frost, N. 2020. Explainable k-means and k-medians clustering. In *International conference on machine learning*, 7055–7065. PMLR.
- Ribeiro, M. T.; Singh, S.; and Guestrin, C. 2016. “Why should i trust you?” Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, 1135–1144.
- Sagi, O.; and Rokach, L. 2018. Ensemble learning: A survey. *Wiley interdisciplinary reviews: data mining and knowledge discovery*, 8(4): e1249.
- Santos, J. M.; and Embrechts, M. 2009. On the use of the adjusted rand index as a metric for evaluating supervised classification. In *International conference on artificial neural networks*, 175–184. Springer.
- Slack, D.; Hilgard, S.; Jia, E.; Singh, S.; and Lakkaraju, H. 2020. Fooling lime and shap: Adversarial attacks on post hoc explanation methods. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 180–186.
- Sokol, K.; and Flach, P. 2020. Explainability fact sheets: a framework for systematic assessment of explainable approaches. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 56–67.
- Sutton, C.; Boley, M.; Ghiringhelli, L. M.; Rupp, M.; Vreeken, J.; and Scheffler, M. 2020. Identifying domains of applicability of machine learning models for materials science. *Nature communications*, 11(1): 4428.
- Van Der Linden, I.; Haned, H.; and Kanoulas, E. 2019. Global aggregations of local explanations for black box models. *arXiv preprint arXiv:1907.03039*.
- Wachter, S.; Mittelstadt, B.; and Russell, C. 2017. Counterfactual explanations without opening the black box: Automated decisions and the GDPR. *Harv. JL & Tech.*, 31: 841.