

# Adherence and Constancy in LIME-RS Explanations for Recommendation

Vito Walter Anelli<sup>1</sup>, Alejandro Bellogín<sup>2</sup>, Tommaso Di Noia<sup>1</sup>, Francesco Maria Donini<sup>3</sup>, Vincenzo Paparella<sup>1</sup> and Claudio Pomo<sup>1</sup>

<sup>1</sup>Politecnico di Bari, via Orabona, 4, 70125 Bari, Italy

<sup>2</sup>Universidad Autónoma de Madrid, Ciudad Universitaria de Cantoblanco, 28049 Madrid, Spain

<sup>3</sup>Università degli Studi della Tuscia, via Santa Maria in Gradi, 4, 01100 Viterbo, Italy

## Abstract

Explainable Recommendation has attracted a lot of attention due to a renewed interest in explainable artificial intelligence. In particular, post-hoc approaches have proved to be the most easily applicable ones to increasingly complex recommendation models, which are then treated as black boxes. The most recent literature has shown that for post-hoc explanations based on local surrogate models, there are problems related to the robustness of the approach itself. This consideration becomes even more relevant in human-related tasks like recommendation. The explanation also has the arduous task of enhancing increasingly relevant aspects of user experience such as transparency or trustworthiness. This paper aims to show how the characteristics of a classical post-hoc model based on surrogates is strongly model-dependent and does not prove to be accountable for the explanations generated.

## Keywords

explainable recommendation, post-hoc explanation, local surrogate model

## 1. Introduction

The explanation of a recommendation list plays an increasingly important role in the interaction of a user with a recommender system: the pervasiveness of economic interest and the inscrutability of most Artificial Intelligence systems make users ask for some form of accountability in the behavior of the systems they interact with. Given the explanation that a system can provide to a user we identify at least two characteristics that the explanation part should enforce [1, 2, 3]:

- *Adherence* to reality: the explanation should mention only features that really pertain to the recommended item. For instance, if the system recommends the movie “Titanic”, it should not explain this recommendation by saying “because it is a War Movie” since it is by no means an adherent description of that movie;
- *Constancy* in the behavior: when the explanation is generated based on some sample, and such a sample is drawn with a probability distribution, the entire process should not exhibit a random behavior to the user. For instance, if the explanation for recommending the movie “The Matrix” to the same user is first “because it is a Dystopian Science Fiction”, and then “because it

is an Acrobatic Duels Movie”, for the same user, this behavior would be perceived as nondeterministic, and thus reducing its trustworthiness.

Among several ways of generating explanations, we study here the application of LIME [4] to the recommendation process. LIME is an algorithm that can explain the predictions of any classifier or regressor in a faithful way, by approximating it locally with an interpretable model. LIME belongs to the category of post-hoc algorithms and it sees the prediction system as a black box by ignoring its underlying operations and algorithms. Since we can consider the recommendation task as a particular Machine Learning task, the LIME approach can also be applied to recommendation. LIME-RS [5] is an adaptation of the general algorithm to the recommendation task and can be considered in all respects as a black-box explainer. This means that it generates an explanation by drawing a huge number of (random) calls to the system, collecting the answers, building a model of behavior of the system, and then constructing the explanation for the particular recommended item. While the fact of adopting a black-box approach lets LIME-RS to be applicable for every recommender system, the way of building a model by drawing a huge random sample of system behaviors makes it lose both adherence and constancy, as our experiments show later on this paper. This suggests that the direct application of LIME-RS to recommender systems is not advisable, and that further research is needed to assess the usefulness of LIME-RS in explaining recommendations.

The paper is organized as follows: Section 2 reviews

*3rd Edition of Knowledge-aware and Conversational Recommender Systems (KaRS) & 5th Edition of Recommendation in Complex Environments (ComplexRec) Joint Workshop @ RecSys 2021, September 27–1 October 2021, Amsterdam, Netherlands*

✉ claudio.pomo@poliba.it (C. Pomo)

© 2021 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)

the state of the art on explanation in recommendation; Section 3 details LIME to make the paper self-contained. Section 4 shows the results of experiments with two mainstream recommendation models: Attribute Item-kNN and Vector Space Model. We discuss the outcomes of the experiments in Section 5, and conclude with Section 6.

## 2. Related Work

In recent years, the theme of Explanation in Artificial Intelligence has come to the foreground, capturing the attention not only of the Machine Learning and related communities – that deal more specifically with the algorithmic part – but also of fields closer to Social Sciences, such as Sociology or Cognitivism, which look with great interest to this area of research [1]. The growing interest in this area is also dictated by new regulations of both Europe [6] and US [7] with respect to sensitive issues in the field of personal data processing, and legal responsibility. This trend has also touched the research field of recommender systems [8, 9, 10, 11]. However, topics such as explanation are by no means new to this field. In fact, we can date back to 2014 the introduction of the term “explainable recommendation” [12], although the need to provide an explanation that accompanies the recommendation is a need that emerged as early as 1999 by Schafer et al. [13], when people began trying to explain a recommendation with other similar items familiar to the user who received that recommendation.

Catalyzation of interest around the topic of explanation of recommendations coincides also with the awareness achieved in considering metrics beyond accuracy as fundamental in evaluating a recommendation system [14, 15]. Indeed, all of the well-known metrics of novelty, diversity, and serendipity are intended to improve the user experience, and in this respect, a key role is played by explanation [3, 16]. “Why are you recommending that?”—this is the question that usually accompanies the user when a suggestion is provided. Tintarev and Masthoff [2] detailed in a scrupulous way the aspects involved in the process of explanation when we talk about recommendation. They identified 7 aspects: user’s trust, satisfaction, persuasiveness, efficiency, effectiveness, scrutability, and transparency.

This is the starting point to define Explainable Recommendation as a task that aims to provide suggestions to the users and make them aware of the recommendation process, explaining also why that specific object has been suggested. Gedikli et al. [3] evaluated different types of explanations and drew a set of guidelines to decide what the best explanation that should equip a recommendation system is. This is due to the fact that popular recommendation systems are based on Matrix Factorization (MF) [17]; for this type of model, trying

to provide an explanation opens the way to new challenges [1, 18, 19, 20].

There are two different approaches to address this type of issue.

- On the one hand, the model-intrinsic explanation strategy aims to create a user-friendly recommendation model or encapsulates an explaining mechanism. However, as Lipton [21] points out, this strategy will weigh in on the trade-off between the transparency and accuracy of the model. Indeed, if the goal becomes to justify recommendations, the purpose of the system is no longer to provide only personalized recommendations, resulting in a distortion of the recommendation process.
- On the other hand, we have a model-agnostic [22] approach, also known as post-hoc [23], which does not require to intervene on the internal mechanisms of the recommendation model and therefore does not affect its performance in terms of accuracy.

Most recommendation algorithms take an MF-approach, and thus the entire recommendation process is based on the interaction of latent factors that bring out the level of liking for an item with respect to a user. Many post-hoc explanation methods have been proposed for precisely these types of recommendation models. It seems evident that the most difficult challenge for this type of approach lies in making these latent factors explicit and understandable for the user [9]. Peake and Wang [23] generate an explanation by exploiting the association rules between features; Tao et al. [24] in their work, find benefit from regression trees to drive learning, and then explain the latent space; instead, Gao et al. [25] try a deep model based on attention mechanisms to make relevant features emerge. Along the same lines are Pan et al. [11], who present a feature mapping approach that maps the uninterpretable general features onto the interpretable aspect features. Among other approaches to consider, [12] proposes an explicit factor model that builds a mapping between the interpretable features and the latent space. On the same line we also find the work by Fusco et al. [26]. In their work, they provide an approach to identify, in a neural model, which features contribute most to the recommendation. However, these post-hoc explanation approaches turn out to be built for very specific models. Purely model-agnostic approaches include the recent work of Tsang et al. [27], who present GLIDER, an approach to estimate interactions between features rather than on the significance of features as in the original LIME [4] algorithm. This type of solution is constructed regardless of the recommendation model.

Our paper focuses on the operation of LIME, a model-agnostic method for a surrogate-based local explanation.

When a user-item pair is provided, this model returns as an outcome of the explanation a set of feature weights, for any recommender system. However, the recommendation task is very specific, so there is a version called LIME-RS [5] that applies the explanation model technique to the recommendation domain. In this way, any recommender is seen as a black box, so LIME-RS plays the role of a model-agnostic explainer whose result is a set of interpretable features and their relative importance.

The goal of LIME-RS is to exploit the predictive power of the recommendation (black box) model to generate an explanation about the suggestion of a particular item for a user. In this respect, it exploits a neighborhood drawn according to a generic distribution compared to the candidate item for the explanation. It seems obvious that the choice of the neighborhood plays a crucial role within the process of explanation generation by LIME-RS. We can compare this sample extraction action to a perturbation of the user-item pair we are using to generate the explanation. In the case of LIME-RS this perturbation must generate consistent samples with respect to the source dataset. We see that this choice represents a critical issue for all the post-hoc models which base their expressiveness on the locality of the instance to explain.

This trend is confirmed in several papers addressing this issue of surrogate-based explanation systems such as LIME and SHAP [28]. In two recent papers, Alvarez-Melis and Jaakkola [29] have shown how the explanations generated with LIME are not very robust: their contribution aims to bring out how small variations or perturbations in the input data cause significant variations in the explanation of that specific input [30]. In their paper, a new strategy is introduced to strengthen these methods by exploiting local Lipschitz continuity. By deeply investigating this drawback, they introduced self-explaining models in stages, progressively generalizing linear classifiers to complex yet architecturally explicit models.

Saito et al. [31] also explored this issue by turning their gaze to different types of sampling to make the result of an explanation generated through LIME more robust. In particular, in their work, they introduce the possibility of generating realistic samples produced with a Generative Adversarial Network. Finally, Slack et al. [32] adopt a similar solution in order to control the perturbation generating neighborhood data points by attempting to mitigate the generation of unreliable explanations while maintaining a stable black-box model of prediction.

### 3. Background Technology

From a formal point of view, we can define a LIME-generated explanation for a generic instance  $x \in \mathcal{X}$  produced by a model  $f$  as:

$$\xi(x) = \underset{e \in E}{\operatorname{argmin}} \mathcal{L}(f, e, \pi_x) + \Omega(e) \quad (1)$$

where  $\mathcal{L}$  represents the fidelity of the surrogate model to the original  $f$ , and  $e$  represents a particular instance of the class  $E$  of all possible explainable models. Among all the possible models, the one most frequently chosen is based on a linear prediction. In this case, an explanation refers to the weights of the most important interpretable features, which, when combined, minimize the divergence from the black-box model. The function  $\pi_x$  measures the distance between the instance to be explained  $x \in \mathcal{X}$ , and the samples  $x' \in \mathcal{X}$  extracted from the training set to train the model  $e$ . Finally,  $\Omega(e)$  represents the complexity of the explanation model.

Two pieces of evidence make the application of LIME possible: (i) the existence of a feature space  $\mathcal{Z}$  on which to train the surrogate model of  $f$ , (ii) and the presence of a surjective function that maps the space mentioned above ( $\mathcal{Z}$ ) to the original space of instances ( $\mathcal{X}$ ). Going into more detail, we consider the fidelity function  $\mathcal{L}$  as the mean square deviation between the prediction for a generic instance  $x' \in \mathcal{X}$  of the black-box model and that generated for the counterpart  $z' \in \mathcal{Z}$  by the surrogate model. Starting from these considerations we can express  $\mathcal{L}$  with the following formula:

$$\mathcal{L}(f, e, \pi_x) = \sum_{x' \in \mathcal{X}, z' \in \mathcal{Z}} \pi_x(x') \cdot (f(x') - e(z'))^2 \quad (2)$$

In the formula above  $\pi_x$  plays a fundamental role as it expresses the distance between the instance to be explained and the sampled instance used to build the surrogate model. From a generic perspective, we can express this function as a kernel function like  $\pi_x = \exp(-D(x, x')^2 / \sigma^2)$ , where  $D$  is any measure of distance.

The full impact of this distance is captured when the fidelity function also considers the transformation of the surrogate sample in the original space. As mentioned earlier, we consider a surjective function  $p$  that maps the original space into the feature space  $p : \mathcal{X} \rightarrow \mathcal{Z}$ . We can also consider the function that allows us to move in the opposite direction  $p^{-1} : \mathcal{Z} \rightarrow \mathcal{X}$ . At this point, Equation (2) becomes:

$$\mathcal{L}(f, e, \pi_x, p) = \sum_{z' \in \mathcal{Z}} \pi_x(p^{-1}(z')) \cdot (f(p^{-1}(z')) - e(z'))^2 \quad (3)$$

From this last equation, we can grasp the criticality of the surjective mapping function. Indeed, the neighborhood in  $\mathcal{Z}$ -space cannot be guaranteed with the transformation in  $\mathcal{X}$ -space. Thus, some samples selected to train the surrogate model could not satisfy the neighborhood criterion for which they were chosen.

We must therefore stress on the centrality of the sampling function: how do we extract the neighborhood of

our instance to be explained? If we look at the application of LIME to the recommendation domain, we can compare this sampling action to a local perturbation around our instance  $x$ ; however, this perturbation aims to generate  $n$  samples  $x'$ , which might contain inconsistencies: as an example, suppose we want to explain *James's* feeling about the movie *The Matrix*. The original triple of the instance to be explained associates to the user-item pair the genre of the movie (representing the explainable space) and in this case it is of the type  $\langle James, TheMatrix, Sci-Fi \rangle$ . A perturbation around this instance could generate inconsistencies of the type  $\langle James, TheMatrix, Western \rangle$ . For this reason, in LIME-RS the perturbation considers only real and not synthetic data. This choice is dictated by the avoidance of the out-of-sample (OOS) process phenomenon. Closely related to this problem predicted by OOS is that the interpretation examples selected in LIME-RS represent the ability to capture the locality through disturbance mechanisms effectively. One of the disadvantages of LIME-like methods is that they sometimes fail to estimate an appropriate local replacement model but instead generate a model that focuses on explaining the examples and is also affected by more general trends in the data.

This issue is central to our work, and it involves two aspects: (i) the first one concerns the sampling function of the samples precisely. In the LIME-RS implementation, this function is driven by the popularity distribution of the items within the dataset. (ii) The second critical issue concerns the model's ability to wittily discriminate the user's taste from the neighborhood extracted to build the surrogate model. A model that squashes too much on bias or is inaccurate cannot bring out the peculiarities of user taste that are critical in building the explainable model which are, in turn, useful in generating the explanation for the instance of interest.

These observations dictate the two research questions that motivated our work:

- RQ1** Can we consider the surrogate-based model on which LIME-RS is built to generate always the same explanations, or does the extraction of a different neighborhood severely impact the system's constancy?
- RQ2** Are LIME-RS explanations adherent to item content, despite the fact that the sampling function is uncritical and based only on popularity?

## 4. Experiments

This section is devoted to illustrating how the experimental campaign was conducted. The datasets used for this phase of experimentation are *Movielens 1M* [33], *Movielens Small* [33], and *Yahoo! Movies*<sup>1</sup>. Their characteristics are shown in Table 1.

*lens Small* [33], and *Yahoo! Movies*<sup>1</sup>. Their characteristics are shown in Table 1.

**Table 1**  
Characteristics of the datasets involved in the experiments.

	Users	Items	Transactions	Sparsity
Movielens 1M	6040	3675	797758	0,9640
Movielens Small	610	8990	80419	0,9853
Yahoo! Movies	7636	8429	160604	0,9975

As for the choice of the models to be used in this work is concerned, we selected two well-known recommendation models that are able to exploit the information content of the items to produce a recommendation: Attribute Item kNN (Att-Item-kNN) and Vector Space Model (VSM). The two chosen models represent the simplest solution to address the recommendation problem by exploiting the content associated with the items in the catalog.

Att-Item-kNN exploits the characteristics of neighborhood-based models but expresses the representation of the items in terms of their content and, based on this representation, it computes a similarity between users. Starting from this similarity and exploiting the collaborative contribution in terms of interactions between users and items, Att-Item-kNN tries to estimate the level of liking of the items in the catalog. VSM represents both users and items in a new space to link users and items to the considered information content. Once obtained this new representation, with an appropriate function of similarity, VSM estimates which are the most appealing items for a specific user. The implementation of both models are available in the ELLIOT [34] evaluation framework. This benchmarking framework was used to select the best configuration for the two recommendation models by exploiting the corresponding configuration file<sup>2</sup>.

Our experiments start by selecting the best configurations based on nDCG [35, 36] for the two models on the considered datasets. Then, we generate the top-10 list of recommendations for each user, and we take into account the first item  $i_1$  on these lists for each user  $u$ . Finally, each recommendation pair  $(u, i_1)$  is explained with LIME-RS. The explanation consists of a weighted vector  $(g, w)_i$  where  $g$  is the genre of the movies in the dataset – i.e., the features – and  $w$  is the weight associated to  $g$  by LIME-RS within the explanation. Then, this vector is sorted by descending weights. In this way, the genres of the movies which play a key role within the recommendation, as explained by LIME-RS, are highlighted at the first positions of the vector. These operations are then repeated  $n = 10$  times and changing the seed each time,

<sup>1</sup>R4 - Yahoo! Movies User Ratings and Descriptive Content Information, v.1.0 <http://webscope.sandbox.yahoo.com/>.

<sup>2</sup>[https://tmy.sh/basic\\_limers](https://tmy.sh/basic_limers)

as 10 is likely to be a good choice to detect a general pattern in the behavior of LIME-RS. At this point, for each pair  $(u, i_1)$ , we have a group of 10 explanations ordered by descending values of  $w$ , which we exploit to answer our two research questions.

**RQ1.** Empirically, since in a real scenario of recommendation a too verbose explanation is not useful, we consider only the first five features in the sorted vector representing the explanation of each recommendation. In order to verify the constancy of the behavior of LIME-RS, given a  $(u, i_1)$  pair, we exploit the  $n$  previously generated explanations for this pair. Then for  $k = 1, 2, \dots, 5$ , we define  $G_k$  as the multiset of genres that appear in  $k$ -th position – for instance, if “Sci-Fi” occurs in the first position of 7 explanations, then “Sci-Fi” occurs 7 times in the multiset  $G_1$ , and similarly for other genres and multisets. Then, we compute the frequency of genres in each position as follows: given a position  $k$ , a genre  $g$ , and the number  $n$  of generated explanations for a given pair  $(u, i_1)$ , the frequency  $f_{gk}$  of  $g$  in  $k$ -th position is computed as:

$$f_{gk} = \frac{|\{g \mid g \in G_k\}|}{n} \quad (4)$$

where  $|\cdot|$  denotes the cardinality of a multiset. Then, all this information is collected for each user in five lists – one for each of the  $k$  positions – of pairs  $\langle g, f_{gk} \rangle$  sorted by frequency. One can observe that the computed frequency is an estimation of the probability that a given genre is put in that position within the explanation generated by LIME-RS sorted by values. Hence, the pair  $\langle g, \max(f_{gk}) \rangle$  describes the genre with the highest frequency in the  $k$ -th position of the explanation for a pair  $(u, i_1)$ . Finally, it makes sense to compute the mean  $\mu_k$  of the highest probability values in each position  $k$  of the explanations for each pair  $(u, i_1)$ . Formally, by setting a position  $k$ , the mean  $\mu_k$  is computed as:

$$\mu_k = \frac{\sum_{j=1}^{|U|} \max(f_{gk})_j}{|U|} \quad (5)$$

where  $U$  is the set of users for whom it was possible to generate a recommendation for. Observing the value of  $\mu_k$ , we can state to what extent LIME-RS is constant in providing the explanations until the  $k$ -th feature: the higher the value of  $\mu_k$ , the higher the constancy of LIME-RS concerning the  $k$ -th feature.

By looking at Table 2, we can see that for Att-Item-kNN the LIME-RS explanation model is reliable as long as it considers at most three features in the weighted vector presented as an explanation of the recommendation. Extending the explanation to four features, we have a constancy that falls below 65%, while arriving at an explanation with five features is more likely to run into explanations that exhibit an unacceptably random behavior. On the other hand, we can see that for VSM

**Table 2**

Constancy of LIME-RS. A value equal to 0 means that the genre(s) provided by LIME-RS in the first  $k$  position(s) is always different (worst case: completely inconstant behavior); A value equal to 1 means that the genre(s) provided by LIME-RS in the first  $k$  position(s) is always the same (total constancy).

	$\mu_1$	$\mu_2$	$\mu_3$	$\mu_4$	$\mu_5$
Att-Item-kNN					
Movielens 1M	0,9130	0,7822	0,6927	0,6288	0,5727
Movielens Small	0,8830	0,7426	0,6639	0,60459	0,5616
Yahoo! Movies	<b>0,9230</b>	<b>0,8016</b>	<b>0,7232</b>	<b>0,6528</b>	<b>0,5830</b>
VSM					
Movielens 1M	0,8929	0,7953	0,7729	0,7726	0,7801
Movielens Small	0,9464	0,8636	0,8343	0,8138	0,8049
Yahoo! Movies	<b>0,9732</b>	<b>0,9209</b>	<b>0,8887</b>	<b>0,8884</b>	<b>0,9056</b>

the values are much more stable. In this case, we have a constancy that, regardless of the length of the weighted vector of the explanation, stabilizes on average around 80%. An aspect emerges that will be discussed in detail later: LIME-RS is conditioned by the ability of the black-box model to discriminate the user’s tastes locally.

**RQ2.** With the aim of providing an answer about the adherence to reality of LIME-RS, we make a comparison between the genres claimed to explain a recommended item and its actual genres. Indeed, the explanations about an item should fit the list of genres the item is characterized by. This means that, in an ideal case, all highly weighted features within the explanation should match the genres of the item. From the results in Table 2, we notice that using Att-Item-kNN the constancy of LIME-RS reaches a low value after the third feature. Hence, it is a futile effort to go deeper in the study of the explanation. To this aim, we intersected each explanation limited to the set  $E_k$  of its first  $k$  genres with the set of genres  $F_{i_1}$  characterizing the first recommended item, for  $k = 1, 2, 3$ . Upon completion of this operation for all the  $n$  explanations generated for each  $(u, i_1)$  pair, we computed the number of times we obtained an empty intersection of these sets, normalized by the total number of explanations  $n \times |U|$ , in order to understand to what extent an explanation is (not) adherent to the item. Formally, for a given value of  $k$ , the value  $adherence_k$  is computed as:

$$adherence_k = \frac{\sum_{j=1}^{n \times |U|} \left[ (E_k \cap F_{i_1})_j = \emptyset \right]}{n \times |U|} \quad (6)$$

where  $U$  is the set of users of the dataset for whom it was possible to generate a recommendation,  $n$  is the number of generated explanations for each pair  $(u, i_1)$ , and by  $\sum[\dots]$  we mean that we sum 1 if the condition inside  $[\dots]$  is true, and 0 otherwise. One can note that  $adherence_k \in [0, 1]$ , where a value of 1 indicates the worst case in which for none of the  $n$  explanations under consideration at least one genre of the item is in the first  $k$

features of the explanation. In contrast, the lower the value of  $adherence_k$ , the higher the adherence of LIME-RS.

**Table 3**

Adherence of LIME-RS. For value equals to 1 no genre provided by LIME-RS in the first  $k$  real genres of the movie (worst case); For value equals to 0 at least one genre provided by LIME-RS in the first  $k$  genres is always among the real genres of the movie.

	$adherence_1$	$adherence_2$	$adherence_3$
Att-Item-kNN			
Movielens 1M	0,2774	0,1105	0,0488
Movielens Small	<b>0,2364</b>	<b>0,0651</b>	<b>0,0180</b>
Yahoo! Movies	0,3597	0,1202	0,0476
VSM			
Movielens 1M	0,5357	0,2539	0,1088
Movielens Small	0,4384	0,1674	0,0403
Yahoo! Movies	<b>0,1013</b>	<b>0,01348</b>	<b>0,0021</b>

Observing the results from Table 3, Att-Item-KNN performs well in terms of adherence since, in approximately 75% of cases, even considering only the main feature of the explanation, it falls into the set of the item genres, as for Movielens dataset family. This performance is a 10% lower for Yahoo! Movies. In contrast with this result, VSM shows poor performances on both dataset of the Movielens family, by failing half the time about Movielens 1M as regards adherence. A surprising result is achieved for Yahoo! Movies dataset because, enlarging the study to the first three features among the explanation, the error is almost completely absent. The reasons we found to explain this difference in the performances concern the characteristics and the quality of the dataset, as we highlight later on.

## 5. Discussion

This work investigates how well a *post-hoc* approach based on local surrogates – such as the LIME-RS algorithm – explains a recommendation. Instead of studying the impact of explanations on users (that is a well-studied topic in the literature and is beyond our scope), we focus on objective evidences that could emerge. In this respect, we have designed specific experiments, which introduced two different metrics, to evaluate adherence and constancy for this kind of algorithms. For instance, Table 2 shows a different behavior for Att-Item-kNN and VSM. On the one hand, Att-Item-kNN seems to guarantee a good constancy in explanations up to the third feature. This suggests that an explanation that exploits the first three features of the list produced by LIME-RS could be barely considered as reliable (i.e., reaching a constancy of 0.69 on Movielens 1M). On the other hand, VSM exhibits a much more “stable” behavior, demonstrating in all cases

(except for the first feature with Movielens 1M) better performance than Att-Item-kNN in terms of constancy, with peaks up to 97%. A straightforward consequence of these observations could be analyzed in terms of confidence or probability. If the constancy steadily decreases, it means that the probability that LIME-RS suggests the same explanatory feature decreases. In practical terms, we could say that LIME-RS is less confident about its explanation. In fact, this is the behavior of Att-Item-kNN. Conversely, VSM shows high values of constancy, resulting in a more “deterministic” behavior. With VSM, LIME-RS is more confident of its explanations. This could increase user’s trustworthiness, since LIME-RS behavior is more reliable.

However, these results could also be interpreted together with the ones from Table 3. They show how often at least one feature – out of  $k$  features provided by LIME-RS – adheres to the features that describe the item being explained. In other words, they measure the probability that LIME-RS succeeds in reconstructing at least one feature of a specific item. Combining the results of Table 2 and those of Table 3, Att-Item-kNN, as already mentioned, shows good performance regarding adherence and identifies 3 times out of 4 the first fundamental feature of the explanation among those present in the set of features originally associated with the item. As expected, if the number  $k$  of LIME-RS-reconstructed features increases, the number of times such a set has a nonempty intersection (with the features belonging to the item) – i.e., adherence – increases. It could be noted that Att-Item-kNN on Yahoo! Movies shows the worst behavior in terms of adherence. VSM shows a different behavior. Despite the excellent performance regarding constancy, it could be observed that on both Movielens datasets, the performance in terms of adherence is poor, and worse for Movielens 1M than for Movielens Small. Surprisingly, on Yahoo! Movies, VSM performs much better, and the errors are almost negligible.

**The difference between the two models** could be due to many reasons. In the following we analyze possible relations between such behaviors and two of them: *popularity bias in the dataset* and *characteristics of side information*. On the one hand, if the dataset is affected by popularity bias, it would be a well-studied cause of confusion for LIME-RS. On the other hand, the characteristics of the side information associated with the datasets could dramatically influence the performance of the two recommendation models. To assess these hypotheses, we have evaluated (see Table 4) the recommendation lists produced by Att-Item-kNN and VSM considering nDCG, Hit Rate (HR), Mean Average Precision (MAP), and Mean Reciprocal Rank (MRR). Table 4 shows that the chosen datasets are strongly affected by popularity bias. Indeed, MostPop is the best performing approach, and the two “personalized” models fail to produce accurate results. This triggers the second aspect that concerns the quality

**Table 4**

Results of the experiments on the models involved in the experiments. Models are optimized according to the value of nDCG.

model	nDCG	Recall	HR	Precision	MAP	MRR
Movielens 1m						
Random	0,0051	0,0028	0,0869	0,0098	0,0094	0,0264
MostPop	0,0845	0,0379	0,4548	0,104	0,115	0,2205
Att-Item-kNN	0,0229	0,0165	0,2425	0,0383	0,0387	0,0888
VSM	0,0173	0,0109	0,2106	0,0292	0,0306	0,0741
Movielens Small						
Random	0,0030	0,0013	0,0492	0,0049	0,0068	0,0205
MostPop	0,0715	0,0389	0,3902	0,0748	0,0912	0,1961
Att-Item-kNN	0,0124	0,0068	0,1459	0,0197	0,0191	0,0484
VSM	0,0085	0,0056	0,1000	0,0111	0,0123	0,0350
Yahoo! Movies						
Random	0,0005	0,0008	0,0051	0,0005	0,0005	0,0015
MostPop	0,2188	0,2589	0,596	0,1067	0,1501	0,3447
Att-Item-kNN	0,0215	0,0262	0,1198	0,0132	0,0155	0,0435
VSM	0,0131	0,0171	0,0754	0,0081	0,0092	0,0261

of the content. The results suggest that the side information is not good enough to boost the recommendation systems in producing meaningful recommendations. In fact, the three datasets seem to have an informative content that is not adequate to generate appealing recommendations. We observe that, from an informative point of view, the Yahoo! Movies dataset is slightly more complete: 22 genres against the 18 genres available on Movielens. Although the VSM model does not show excellent performance, in combination with LIME-RS, it provides explanations that are very reliable in terms of constancy (see Table 2) and adherence (see Table 3) to the actual content of the items being explained.

**From the designer perspective**, there is also a pragmatic way to look at the experimental results. Suppose a developer needs an off-the-shelf way of generating explanations for recommendations, and chooses LIME-RS to do that. Our results suggest that if the explainer employs a Movielens dataset with Att-Item-kNN model, then it is better to run the explainer several times. Indeed, the first feature obtained for the explanation could change around 1 time every 5 trials (first column of Table 2), and once such a feature is obtained, it is better to check whether this feature is really among the ones describing the item, since 1 time out of 4 the feature can be wrong (first column of Table 3). Moreover, if the explainer employs the Yahoo! Movies dataset with VSM model, then probably there is no need to run the explainer twice, since its behavior is constant 97% of the times, while the feature is wrong only 10% of the times. However, the low performance of such a model is to be taken into account.

## 6. Conclusion

In this paper we shed a first light on the effectiveness of LIME-RS as a black-box explanation model in a rec-

ommendation scenario. We propose two different measures to understand how reliable an explanation based on LIME-RS is: (i) *constancy* was used to assess the impact of the random sampling phase of LIME-RS on the provided explanation – ideally the explanation should remain constant in spite of the sample used to obtain it; (ii) *adherence* was proposed to understand the reconstructive power of LIME-RS with respect to the features that belong to the item involved in the explanation – ideally, LIME-RS should provide an explanation that always adheres to the actual features of the recommended item.

To test both constancy and adherence, we trained and optimized two content-based recommendation models: Attribute Item-kNN (Att-Item-kNN), and a classical Vector Space Model. For each model, and for all datasets exploited in the study, we generated recommendation lists for all users. We exploited the first item of these top-10 lists to produce the explanations that were then the subject of our investigation. It turned out that for models built with a large collaborative input such as Att-Item-kNN, LIME-RS produces fairly constant explanations up to a length of three features. Moreover, these explanations turn out to be adherent with respect to the item between 65% and 75% of the cases in which only the first feature of the weighted vector of explanations is considered. VSM shows a different behavior where explanations are much more constant, but suffer a lot in terms of adherence, except for the Yahoo! Movies dataset for which the explanation model showed outstanding performance despite the poor ability of VSM to provide sound recommendations to users.

In our experiments, some evidence started to emerge highlighting that the adopted explanation model is conditioned not only by the accuracy of the black-box model it tries to explain but also by the quality of the side information used to train the model. The latter result deserves to be adequately investigated to search for a link at a higher level of detail. We plan to apply our experiments also to other recommendation models, to see whether the problems with adherence and constancy that we found for the two tested models show up also in other situations. We will also investigate what impact structured knowledge has on this performance by exploiting models capable of leveraging this type of content. In addition, it would also be the case to try different reference domains with richer datasets of side information to understand what impact content quality has on this type of explainer.

## Acknowledgments

The authors acknowledge partial support of PID2019-108965GB-I00, PON ARS01\_00876 BIO-D, Casa delle Tecnologie Emergenti della Città di Matera, PON ARS01\_00821 FLET4.0, PIA Servizi Locali 2.0, H2020 Passapartout - Grant n. 101016956, PIA ERP4.0, and IPZS-PRJ4\_IA\_NORMATIVO.

## References

- [1] T. Miller, Explanation in artificial intelligence: Insights from the social sciences, *Artif. Intell.* 267 (2019) 1–38. URL: <https://doi.org/10.1016/j.artint.2018.07.007>. doi:10.1016/j.artint.2018.07.007.
- [2] N. Tintarev, J. Masthoff, Explaining recommendations: Design and evaluation, in: *Recommender Systems Handbook*, Springer, 2015, pp. 353–382.
- [3] F. Gedikli, D. Jannach, M. Ge, How should I explain? A comparison of different explanation types for recommender systems, *Int. J. Hum. Comput. Stud.* 72 (2014) 367–382. URL: <https://doi.org/10.1016/j.ijhcs.2013.12.007>. doi:10.1016/j.ijhcs.2013.12.007.
- [4] M. T. Ribeiro, S. Singh, C. Guestrin, “why should I trust you?”: Explaining the predictions of any classifier, in: B. Krishnapuram, M. Shah, A. J. Smola, C. C. Aggarwal, D. Shen, R. Rastogi (Eds.), *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, San Francisco, CA, USA, August 13-17, 2016, ACM, 2016, pp. 1135–1144. URL: <https://doi.org/10.1145/2939672.2939778>. doi:10.1145/2939672.2939778.
- [5] C. Nóbrega, L. B. Marinho, Towards explaining recommendations through local surrogate models, in: C. Hung, G. A. Papadopoulos (Eds.), *Proceedings of the 34th ACM/SIGAPP Symposium on Applied Computing*, SAC 2019, Limassol, Cyprus, April 8-12, 2019, ACM, 2019, pp. 1671–1678. URL: <https://doi.org/10.1145/3297280.3297443>. doi:10.1145/3297280.3297443.
- [6] S. Wachter, B. Mittelstadt, C. Russell, Counterfactual explanations without opening the black box: Automated decisions and the gdpr, *Harv. J. L. & Tech.* 31 (2017) 841.
- [7] J. Chakraborty, K. Peng, T. Menzies, Making fair ML software using trustworthy explanation, in: *35th IEEE/ACM International Conference on Automated Software Engineering*, ASE 2020, Melbourne, Australia, September 21-25, 2020, IEEE, 2020, pp. 1229–1233. URL: <https://doi.org/10.1145/3324884.3418932>. doi:10.1145/3324884.3418932.
- [8] Y. Zhang, X. Chen, Explainable recommendation: A survey and new perspectives, *Found. Trends Inf. Retr.* 14 (2020) 1–101. URL: <https://doi.org/10.1561/15000000066>. doi:10.1561/15000000066.
- [9] V. W. Anelli, T. D. Noia, E. D. Sciascio, A. Ragone, J. Trotta, How to make latent factors interpretable by feeding factorization machines with knowledge graphs, in: C. Ghidini, O. Hartig, M. Maleshkova, V. Svátek, I. F. Cruz, A. Hogan, J. Song, M. Lefrançois, F. Gandon (Eds.), *The Semantic Web - ISWC 2019 - 18th International Semantic Web Conference*, Auckland, New Zealand, October 26-30, 2019, *Proceedings, Part I*, volume 11778 of *Lecture Notes in Computer Science*, Springer, 2019, pp. 38–56. URL: [https://doi.org/10.1007/978-3-030-30793-6\\_3](https://doi.org/10.1007/978-3-030-30793-6_3). doi:10.1007/978-3-030-30793-6\_3.
- [10] G. P. Polletti, H. N. Munhoz, F. G. Cozman, Explanations within conversational recommendation systems: improving coverage through knowledge graph embedding, in: *2020 AAAI Workshop on Interactive and Conversational Recommendation System*. AAAI Press, New York City, New York, USA, 2020.
- [11] D. Pan, X. Li, X. Li, D. Zhu, Explainable recommendation via interpretable feature mapping and evaluation of explainability, in: C. Bessiere (Ed.), *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*, IJCAI 2020, ijcai.org, 2020, pp. 2690–2696. URL: <https://doi.org/10.24963/ijcai.2020/373>. doi:10.24963/ijcai.2020/373.
- [12] Y. Zhang, G. Lai, M. Zhang, Y. Zhang, Y. Liu, S. Ma, Explicit factor models for explainable recommendation based on phrase-level sentiment analysis, in: S. Geva, A. Trotman, P. Bruza, C. L. A. Clarke, K. Järvelin (Eds.), *The 37th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '14, Gold Coast, QLD, Australia - July 06 - 11, 2014, ACM, 2014, pp. 83–92. URL: <https://doi.org/10.1145/2600428.2609579>. doi:10.1145/2600428.2609579.
- [13] J. B. Schafer, J. A. Konstan, J. Riedl, Recommender systems in e-commerce, in: S. I. Feldman, M. P. Wellman (Eds.), *Proceedings of the First ACM Conference on Electronic Commerce (EC-99)*, Denver, CO, USA, November 3-5, 1999, ACM, 1999, pp. 158–166. URL: <https://doi.org/10.1145/336992.337035>. doi:10.1145/336992.337035.
- [14] S. M. McNee, J. Riedl, J. A. Konstan, Being accurate is not enough: how accuracy metrics have hurt recommender systems, in: G. M. Olson, R. Jeffries (Eds.), *Extended Abstracts Proceedings of the 2006 Conference on Human Factors in Computing Systems*, CHI 2006, Montréal, Québec, Canada, April 22-27, 2006, ACM, 2006, pp. 1097–1101. URL: <https://doi.org/10.1145/1125451.1125659>. doi:10.1145/1125451.1125659.
- [15] S. Vargas, Novelty and diversity enhancement and evaluation in recommender systems and information retrieval, in: S. Geva, A. Trotman, P. Bruza, C. L. A. Clarke, K. Järvelin (Eds.), *The 37th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '14, Gold Coast, QLD, Australia - July 06 - 11, 2014, ACM, 2014, p. 1281. URL: <https://doi.org/10.1145/2600428.2610382>. doi:10.1145/2600428.2610382.
- [16] N. Tintarev, J. Masthoff, A survey of explanations in recommender systems, in: *ICDE Workshops*, IEEE Computer Society, 2007, pp. 801–810.
- [17] Y. Koren, R. M. Bell, C. Volinsky, Matrix factorization techniques for recommender systems, *Computer* 42 (2009) 30–37. URL: <https://doi.org/10.1109/MC.2009.263>. doi:10.1109/MC.2009.263.
- [18] K. Tsukuda, M. Goto, Dualdiv: diversifying items and explanation styles in explainable hybrid recommendation, in: T. Bogers, A. Said, P. Brusilovsky, D. Tikk (Eds.), *Proceedings of the 13th ACM Conference on Recommender Systems*, RecSys 2019, Copenhagen, Denmark, September 16-20, 2019, ACM, 2019, pp. 398–402. URL: <https://doi.org/10.1145/3298689.3347063>. doi:10.1145/3298689.3347063.
- [19] X. Chen, H. Chen, H. Xu, Y. Zhang, Y. Cao, Z. Qin, H. Zha, Personalized fashion recommendation with visual explanations based on multimodal attention network: Towards visually explainable recommendation, in: B. Piwowarski, M. Chevalier, É. Gaussier, Y. Maarek, J. Nie, F. Scholer (Eds.), *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR 2019, Paris, France, July 21-



- 25, 2019, ACM, 2019, pp. 765–774. URL: <https://doi.org/10.1145/3331184.3331254>. doi:10.1145/3331184.3331254.
- [20] G. Cornacchia, F. M. Donini, F. Narducci, C. Pomo, A. Ragone, Explanation in multi-stakeholder recommendation for enterprise decision support systems, in: A. Polyvyanyy, S. Rinderle-Ma (Eds.), *Advanced Information Systems Engineering Workshops - CAiSE 2021 International Workshops*, Melbourne, VIC, Australia, June 28 - July 2, 2021, Proceedings, volume 423 of *Lecture Notes in Business Information Processing*, Springer, 2021, pp. 39–47. URL: [https://doi.org/10.1007/978-3-030-79022-6\\_4](https://doi.org/10.1007/978-3-030-79022-6_4). doi:10.1007/978-3-030-79022-6\_4.
- [21] Z. C. Lipton, The mythos of model interpretability, *Commun. ACM* 61 (2018) 36–43. URL: <https://doi.org/10.1145/3233231>. doi:10.1145/3233231.
- [22] X. Wang, Y. Chen, J. Yang, L. Wu, Z. Wu, X. Xie, A reinforcement learning framework for explainable recommendation, in: *IEEE International Conference on Data Mining, ICDM 2018*, Singapore, November 17-20, 2018, IEEE Computer Society, 2018, pp. 587–596. URL: <https://doi.org/10.1109/ICDM.2018.00074>. doi:10.1109/ICDM.2018.00074.
- [23] G. Peake, J. Wang, Explanation mining: Post hoc interpretability of latent factor models for recommendation systems, in: Y. Guo, F. Farooq (Eds.), *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2018*, London, UK, August 19-23, 2018, ACM, 2018, pp. 2060–2069. URL: <https://doi.org/10.1145/3219819.3220072>. doi:10.1145/3219819.3220072.
- [24] Y. Tao, Y. Jia, N. Wang, H. Wang, The fact: Taming latent factor models for explainability with factorization trees, in: B. Piwowarski, M. Chevalier, É. Gaussier, Y. Maarek, J. Nie, F. Scholer (Eds.), *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2019*, Paris, France, July 21-25, 2019, ACM, 2019, pp. 295–304. URL: <https://doi.org/10.1145/3331184.3331244>. doi:10.1145/3331184.3331244.
- [25] J. Gao, X. Wang, Y. Wang, X. Xie, Explainable recommendation through attentive multi-view learning, in: *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019*, Honolulu, Hawaii, USA, January 27 - February 1, 2019, AAAI Press, 2019, pp. 3622–3629. URL: <https://doi.org/10.1609/aaai.v33i01.33013622>. doi:10.1609/aaai.v33i01.33013622.
- [26] F. Fusco, M. Vlachos, V. Vasileiadis, K. Wardatzky, J. Schneider, Reconet: An interpretable neural architecture for recommender systems, in: S. Kraus (Ed.), *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019*, Macao, China, August 10-16, 2019, *ijcai.org*, 2019, pp. 2343–2349. URL: <https://doi.org/10.24963/ijcai.2019/325>. doi:10.24963/ijcai.2019/325.
- [27] M. Tsang, D. Cheng, H. Liu, X. Feng, E. Zhou, Y. Liu, Feature interaction interpretability: A case for explaining ad-recommendation systems via neural interaction detection, in: *8th International Conference on Learning Representations, ICLR 2020*, Addis Ababa, Ethiopia, April 26-30, 2020, *OpenReview.net*, 2020. URL: <https://openreview.net/forum?id=BkgnhTEtDS>.
- [28] E. Strumbelj, I. Kononenko, Explaining prediction models and individual predictions with feature contributions, *Knowl. Inf. Syst.* 41 (2014) 647–665. URL: <https://doi.org/10.1007/s10115-013-0679-x>. doi:10.1007/s10115-013-0679-x.
- [29] D. Alvarez-Melis, T. S. Jaakkola, Towards robust interpretability with self-explaining neural networks, in: S. Bengio, H. M. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, R. Garnett (Eds.), *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018*, December 3-8, 2018, Montréal, Canada, 2018, pp. 7786–7795. URL: <https://proceedings.neurips.cc/paper/2018/hash/3e9f0fc9b2f89e043bc6233994dfcf76-Abstract.html>.
- [30] D. Alvarez-Melis, T. S. Jaakkola, On the robustness of interpretability methods, *CoRR abs/1806.08049* (2018). URL: <http://arxiv.org/abs/1806.08049>. arXiv:1806.08049.
- [31] S. Saito, E. Chua, N. Capel, R. Hu, Improving LIME robustness with smarter locality sampling, *CoRR abs/2006.12302* (2020). URL: <https://arxiv.org/abs/2006.12302>. arXiv:2006.12302.
- [32] D. Slack, S. Hilgard, E. Jia, S. Singh, H. Lakkaraju, Fooling LIME and SHAP: adversarial attacks on post hoc explanation methods, in: A. N. Markham, J. Powles, T. Walsh, A. L. Washington (Eds.), *AIES '20: AAAI/ACM Conference on AI, Ethics, and Society*, New York, NY, USA, February 7-8, 2020, ACM, 2020, pp. 180–186. URL: <https://doi.org/10.1145/3375627.3375830>. doi:10.1145/3375627.3375830.
- [33] F. M. Harper, J. A. Konstan, The movielens datasets: History and context, *ACM Trans. Interact. Intell. Syst.* 5 (2016) 19:1–19:19. URL: <https://doi.org/10.1145/2827872>. doi:10.1145/2827872.
- [34] V. W. Anelli, A. Bellogin, A. Ferrara, D. Malitesta, F. A. Merra, C. Pomo, F. M. Donini, T. D. Noia, Elliot: A comprehensive and rigorous framework for reproducible recommender systems evaluation, in: F. Diaz, C. Shah, T. Suel, P. Castells, R. Jones, T. Sakai (Eds.), *SIGIR '21: The 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, Virtual Event, Canada, July 11-15, 2021, ACM, 2021, pp. 2405–2414. URL: <https://doi.org/10.1145/3404835.3463245>. doi:10.1145/3404835.3463245.
- [35] V. W. Anelli, T. D. Noia, E. D. Sciascio, C. Pomo, A. Ragone, On the discriminative power of hyperparameters in cross-validation and how to choose them, in: T. Bogers, A. Said, P. Brusilovsky, D. Tikk (Eds.), *Proceedings of the 13th ACM Conference on Recommender Systems, RecSys 2019*, Copenhagen, Denmark, September 16-20, 2019, ACM, 2019, pp. 447–451. URL: <https://doi.org/10.1145/3298689.3347010>. doi:10.1145/3298689.3347010.
- [36] W. Krichene, S. Rendle, On sampled metrics for item recommendation, in: R. Gupta, Y. Liu, J. Tang, B. A. Prakash (Eds.), *KDD '20: The 26th ACM SIGKDD Con-*

ference on Knowledge Discovery and Data Mining, Virtual Event, CA, USA, August 23-27, 2020, ACM, 2020, pp. 1748–1757. URL: <https://doi.org/10.1145/3394486.3403226>. doi:10.1145/3394486.3403226.