

Report on EARS'18: 1st International Workshop on Explainable Recommendation and Search

Yongfeng Zhang
Rutgers University
yongfeng.zhang@rutgers.edu

Yi Zhang
UC Santa Cruz
yiz@soe.ucsc.edu

Min Zhang
Tsinghua University
z-m@tsinghua.edu.cn

Abstract

This is a report on the first edition of the International Workshop on Explainable Recommendation and Search (EARS 2018), co-located with the 41st International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2018) held in Ann Arbor, Michigan, USA on July 8-12, 2018. The workshop was held on July 12, 2018, co-chaired by Yongfeng Zhang, Yi Zhang, and Min Zhang, with invited keynote speeches delivered by Prof. Paul Resnick from University of Michigan, and Dr. Qingsong Hua from Alibaba Inc. An invited panelist committee including Professors Matthew Lease, Paul Resnick, Mark Sanderson, Wlodek Zadrozny and Dr. Qingsong Hua made fruitful discussions about the past, present, and future on the research of explainable recommendation and search. The workshop accepted six contributed papers, and attracted over 80 registered participants. The scope of the workshop spans from technical approaches to explainable recommendation and search, to policy debates on the “principle of transparency” and the “right to explanation” of algorithmic decisions implied in recent EU General Data Protection Regulation (GDPR) and other similar regulations such as The California Consumer Privacy Act of 2018.

1 Introduction

EARS 2018, the 1st International Workshop on Explainable Recommendation and Search¹, was co-located with the 41st International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2018) held in Ann Arbor, Michigan, USA on July 12, 2018.

Explainable recommendation and search attempt to develop models or methods that not only generate high-quality recommendation or search results, but also intuitive explanations of the results, which can help users to better understand the results, or help system designers to better understand how the system or model works. It can eventually help to improve the system transparency, persuasiveness, trustworthiness, and effectiveness. Explainability is even more important in *personalized* search or recommendation scenarios, where users would like to know why a particular web page, product, news report, or friend suggestion exists in his or her own search or recommendation lists.

¹<https://ears2018.github.io/>

The motivation of the workshop is to promote the research and application of Explainable Recommendation and Search, under the background of Explainable AI in a broader sense. Early recommendation and search systems adopted intuitive yet easily explainable models to generate recommendation and search lists, such as user-based and item-based collaborative filtering for recommendation, which provide recommendations based on similar users or items, or TF-IDF based retrieval models for search, which provide document ranking lists according to word similarity between different documents.

However, state-of-the-art recommendation and search models extensively rely on complex machine learning and latent representation models such as matrix factorization or even deep neural networks, and they work with various types of information sources such as ratings, text, images, audio or video signals. The complexity nature of state-of-the-art models make search and recommendation systems as blank-boxes for both end users and system designers, and the lack of explainability weakens the transparency, persuasiveness, and trustworthiness of the system, making explainable recommendation and search important research issues to the research community.

In a broader sense, researchers in the broader artificial intelligence community have also realized the importance of Explainable AI, which aims to address a wide range of AI explainability problems in deep learning, computer vision, automatic driving systems, and natural language processing tasks. Very recently, a series of AI regulations have entered into force, such as the EU General Data Protection Regulation (GDPR) and The California Consumer Privacy Act of 2018, which emphasize the “principle of transparency” of intelligent algorithms, and imply the “right to explanation” of algorithmic decisions in AI systems. As an important branch of AI research, this further highlights the importance and urgency for our research community to discuss and address the explainability issues of various recommendation and search systems.

2 Keynotes

EARS 2018 workshop began with two invited keynote presentations, including a research keynote and an industry keynote.

Survey Equivalence: An Information-theoretic Measure of Classifier Accuracy When the Ground Truth is Subjective by Prof. Paul Resnick, the Michael D. Cohen Collegiate Professor of Information and Associate Dean for Research at the University of Michigan School of Information.

Many classification tasks have no objective ground truth. Examples include: which content or explanation is “better” according to some community, is this comment toxic, or what is the political leaning of this news article. The traditional modeling approach assumes each item has an objective true state that is perceived by humans with some random error. It fails to account for the fact that people have greater agreement on some items than others. Prof. Resnick described an alternative model where the true state is a distribution over labels that raters from a specified population would assign to an item. This leads to information gain as a theoretically justified and computationally tractable measure of a classifier’s quality, and an intuitive interpretation of information gain in terms of the sample size for a survey that would yield the same expected error rate.

Shakespeare of Alibaba: Practice of Intelligent Recommendation Reason Generation in Alibaba by Dr. Qingsong Hua, who is responsible for the overall international search algorithm team in Alibaba.

Explainable recommendation and search is a very promising topic in both academia and industry in the recent years. There are many human generated recommendation reasons for products in Alibaba Taobao to improve user experience and to increase user stickiness. However, relying on human-generated content will result in low coverage, low quality stability, and high financial expenditure. With the rapid development of deep learning technology in NLP, especially in the nature language generation field, we tried natural language generation approach in recommendation reason generation and achieved good results. We created recommendation reasons for auction and auction lists, which covered millions of product categories in Taobao e-commerce, and the generated explanations were used for large-scale real-world transactions in “2017 Double 11 Shopping Festival” without any manual checking. Industry-level real system experiments show that it was very difficult to distinguish whether the explanations are machine-generated or manually-written, and the content generation can be controlled in multiple dimensions such as text style, text length, topics, etc. We will introduce our solution and technical details about generating free-text explanations in this keynote.

3 Contributed Papers

In this section, we provide brief overviews of the six contributed papers accepted for publication at the workshop. The paper presentation program was broadly classified into two sessions: three papers for Explainable Search, and the other three papers for Explainable Recommendation. Of course there is no clear technical boundary between the two sessions, the papers are assigned to different sessions according to their research task.

3.1 Explainable Search

Explainable Information Retrieval using X-rays of Documents by Noriaki Kawamae (The University of Tokyo) [1]. To help grasp a set of documents covering various views at a glance, this paper presents a hierarchical topic model, the Multiple Perspective Dirichlet Process (MPDP). The model constructs a hierarchy of topics that becomes more specialized toward the leaves, by the help of both supervised side-information, and a regression model. MPDP unveils parent-children hierarchies between these topics and constructs topic trees that can be taken as the bones of given documents. Each tree distinguishes multiple perspectives at the leaf level topics, unlike previous models which fail to offer this function due to the limitations of their structures. For example, applications using MPDP can help to find documents that best match a query at the document level, by presenting words that are located in the leaf level topics and indicates differences in terms of perspective. Experiments on various data sets show that MPDP can provide insights as explanations on given documents, such as X-rays of the texts, and realize a variety of opinion mining tasks.

Explaining Credibility in News Articles using Cross-Referencing by Dimitrios Bountouridis, Monica Marrero, Nava Tintarev, and Claudia Hauff (Delft University of Technology) [2]. The proliferation of online news sources has placed the issue of credibility at

the center of public and scholarly attention alike. Without an authoritative entity that can vouch and sufficiently explain the quality of a piece of information appearing in news articles, readers become skeptical. At the same time, computational solutions are typically founded on different and possibly narrow interpretations of the complex concept of credibility. As such, while significant progress has been made, computational efforts have yet to propose a widely accepted solution. This paper proposes an interactive interface alternative to the existing algorithmic solutions: an additional information layer that is applied to an article’s original textual contents. By contrasting heterogeneous articles of the same story, i.e., articles from different news outlets, the proposed approach reveals those pieces of information that are cross-referenced and thus – the authors argue – more likely to be credible. A demo of the tool is available at <http://fairnews.ewi.tudelft.nl/InCredible>, and the code is open-sourced at <https://github.com/dbountouridis/InCredible>.

Posthoc Interpretability of Learning to Rank Models using Secondary Training Data by Jaspreet Singh and Avishek Anand (L3S Research Centre) [3]. Predictive models are omnipresent in automated and assisted decision making scenarios. But for the most part they are used as black boxes, which output a prediction without understanding partially or even completely how different features influence the model prediction, which limits the algorithmic transparency. Rankings are ordering over items encoding implicit comparisons typically learned using a family of features based on learning-to-rank models. In this paper the authors focus on how best we can understand the decisions made by a ranker in a post-hoc model agnostic manner. The authors operate on the notion of interpretability based on explainability of rankings over an interpretable feature space. Furthermore they train a tree based model (inherently interpretable) using labels from the ranker, called secondary training data to provide explanations. Consequently, they attempt to study how well does a subset of features, potentially interpretable, explain the full model under different training sizes and algorithms. The authors did experiments on the learning to rank datasets with 30k queries and report results that serve show in certain settings we can learn a faithful interpretable ranker.

3.2 Explainable Recommendation

Layer-wise Relevance Propagation for Explainable Recommendations by Homanga Bharadhwaj (Indian Institute of Technology Kanpur) [4]. In this paper, the author tackle the problem of explanations in a deep learning based model for recommendations by leveraging the technique of layer-wise relevance propagation. The author used a Deep Convolutional Neural Network to extract relevant features from the input images before identifying similarity between the images in feature space. Relationships between the images are identified by the model and layer-wise relevance propagation is used to infer pixel-level details of the images that may have significantly informed the model’s choice. The author evaluated the proposed method on an Amazon products dataset and demonstrate the efficacy of the proposed approach.

User-Oriented Explanation for Intelligent Candidate Recommendation in Liepin.com by Roy Shan, Xueyan Xu, and Zhijie Li (Liepin.com) [5]. Liepin.com is one of the largest e-recruitment platforms in China. Matching millions of jobs with more than 40 million talents effectively is critical to the customer success and user experience in this service. In this paper,

the authors first provided an overview of the architecture of their intelligent candidate recommender engine, and then presented a user-oriented approach to generating various types of explanations for recommendation. The experimental evaluation of the approaches show that the explanations significantly improved customer satisfaction and business metrics in real-world systems. Except for the industry keynote that shows how explainable recommendation system can help to generate natural language recommendation reasons in Alibaba, this talk is another contribution from industry that shows how explanations play important roles in commercial systems.

Correlated Topic Modeling via Householder Flow by Luyang Liu, Heyan Huang, and Yang Gao (Beijing Institute of technology) [6]. Topic models can be one of the prevalent unsupervised learning approaches in natural language processing. Recent works on neural variational inference offer a powerful framework combining neural networks and interpretable probability models. However, one fundamental assumption is that topics in the latent space are independent to each other, which is actually not the case in the reality. In this paper, the authors propose the Correlated Householder Topic Model (CHTM) to capture the correlations among topics, and model them via Householder flow. The experiments show that, by incorporating topic correlation, CHTM outperforms baseline methods on topic clustering tasks, and shows the possibility to provide better word-cloud based explanations for search and recommendation tasks.

4 Panel and Breakout Discussions

The workshop invited five distinguished researchers and practitioners as the panelist committee to discuss about the past, present, and future of explainable recommendation and search in both the industry and academic world. The panelist committee included Dr. Qingsong Hua (Alibaba), and Professors Matthew Lease (UT Austin), Paul Resnick (UMich), Mark Sanderson (RMIT), and Wlodek Zadrozny (UNC Charlotte), moderated by Yongfeng Zhang, as well as participation from audiences. The questions raised to the panelists stemmed from a wide range of topics from research problems to policy debates.

One of the first discussion points for the panel was what do we exactly mean by “explainability”. Recent research on explainable recommendation and search has spanned across a wide range of topics, including but not limited to designing explainable machine learning models, generating user-oriented textual/visual/statistical explanations for the search or recommendation results, or understanding the nature of user behaviors and user interaction with the results. The application scenario of explainable search and recommendation also spans a wide range of systems, such as explainable web search, product search, e-commerce recommendation, social recommendation, or other tasks that falls into broader information retrieval categories, such as explainable document clustering, classification, and summarization, as well as explainable multimedia retrieval systems. Broadly speaking, explainable recommendation and search try to answer questions that are not only about the *how* (such as how to design a model that improves performance), but also the *why*, such as why a model works in the way that it is supposed to, and why users interact with an IR model/system in the observed manner. The panelists agree that to achieve this goal, it needs extensive efforts from both the computer science and information science communities.

The panel also interacted with audiences to discuss what are important problems to solve

on explainable IR, information systems, and decision making systems, and reached a wide scope of answers. For example, a lot of neural approaches to search and recommendation rely on complex neural networks to learn matching scores, while it is important to understand how the network structure encodes query-document or user-item similarities and how intuitive explanations can be generated out of these structures. Another example is to explore new types of explanations in information retrieval systems – classical web search engines use snippet or document summarization as the explanation of search results, and representative e-commerce recommendation systems use statistical measures to construct explanations, such as “80% consumers that bought this product also bought the recommended product”. To provide more informed search or recommendation results, we can explore new types of explanations such as visual explanations, natural language explanations, or case-based explanations, and explanations can also be personalized. Explanations are also important to a broader scope of IR systems, such as legal retrieval, medical search and academic search systems, where new types of explanation are needed, and it also requires the system to have certain reasoning ability to provide informed explanations. Panelists also mentioned that the evaluation of explainable search and recommendation systems is important, including how to evaluate the explainability/transparency of the models, and how to evaluate the quality of the output explanations. Various approaches can be explored for evaluation, including benchmark-based evaluation measures, and crowd-sourcing approaches to evaluating the explanations.

The last question that raised extensive discussion is regarding policies on AI regulation. In May 2018, provisions of the EU General Data Protection Regulation (GDPR) became directly applicable in all EU member states, and in July 2018, GDPR also became valid in the EEA countries. In June 2018, the US state of California passed a similar bill called The California Consumer Privacy Act of 2018, which is expected to take effect in 2020. GDPR regulations emphasize the “principle of transparency” of intelligent algorithms, and imply the “right to explanation” of algorithmic decisions in AI systems, which could bring significant impact on the research and application of AI systems in the near future, including intelligent search and recommendation systems. The panelists agreed on the importance of the research on explainability issues of AI systems, and also pointed out that there still lacks a clear legal and technical definition of the right to explanation in the related regulations – this is a good opportunity for our research community to explore possible technical roadmaps towards explainable AI, and also highlights the responsibility of the academic world to participate in the policy making towards future explainable AI systems.

5 Concluding Remarks

Overall, the 1st International Workshop on ExplainAble Recommendation and Search (EARS 2018) was a success, in terms of both the number of participants and the research interests emerged during the keynotes, presentations, and discussions. We also formed tentative plans to organize the second workshop next year, as well as plans to organize a special issue in journals on this topic, which are expected to attract a series of research efforts on explainable recommendation, search, and beyond.

6 Acknowledgements

The workshop organizers would like to express since thanks to all the program committee members: Behnoush Abdollahi (University of Louisville), Qingyao Ai (University of Massachusetts Amherst), Rose Catherine (Carnegie Mellon University), Xu Chen (Tsinghua University), Michael Ekstrand (Boise State University), Ruining He (Pinterest), Xiangnan He (National University of Singapore), Bart Knijnenburg (Clemson University), Aonghus Lawlor (University College Dublin), Piji Li (The Chinese University of Hong Kong), Julian Mcauley (University of California San Diego), Sole Pera (Boise State University), Zhaochun Ren (JD Data Science Lab), Sungyong Seo (University of Southern California), Xiang Wang (National University of Singapore), Yao Wu (Twitter Inc.), and Hamed Zamani (University of Massachusetts Amherst).

Finally, we want to thank all the authors, speakers, volunteers, and other participants for making EARS 2018 a success.

References

- [1] Noriaki Kawamae. Explainable Information Retrieval using X-rays of Documents. In *Proceedings of the 1st International Workshop on ExplainAble Recommendation and Search (EARS 2018)*. Ann Arbor, MI, USA.
 - [2] Dimitrios Bountouridis, Monica Marrero, Nava Tintarev, and Claudia Hauff. Explaining Credibility in News Articles using Cross-Referencing. In *Proceedings of the 1st International Workshop on ExplainAble Recommendation and Search (EARS 2018)*. Ann Arbor, MI, USA.
 - [3] Jaspreet Singh and Avishek Anand. Posthoc Interpretability of Learning to Rank Models using Secondary Training Data. In *Proceedings of the 1st International Workshop on ExplainAble Recommendation and Search (EARS 2018)*. Ann Arbor, MI, USA.
 - [4] Homanga Bharadhwaj. Layer-wise Relevance Propagation for Explainable Recommendations. In *Proceedings of the 1st International Workshop on ExplainAble Recommendation and Search (EARS 2018)*. Ann Arbor, MI, USA.
 - [5] Roy Shan, Xueyan Xu, and Zhijie Li. User-Oriented Explanation for Intelligent Candidate Recommendation in Liepin.com. In *Proceedings of the 1st International Workshop on ExplainAble Recommendation and Search (EARS 2018)*. Ann Arbor, MI, USA.
 - [6] Luyang Liu, Heyan Huang, and Yang Gao. Correlated Topic Modeling via Householder Flow. In *Proceedings of the 1st International Workshop on ExplainAble Recommendation and Search (EARS 2018)*. Ann Arbor, MI, USA.
-