



Interactive Social Media Federated Recommendation System via Chrome Extension

Prepared By

Name: Odai Atef Atiya Abdalbaky

ID No.: 252000312

Submitted in fulfillment of the requirements for the degree Master of
Science in Artificial Intelligence

At

Collage of Informatics

Midocean University

UAE

Supervisor Name: Dr.Ashraf Alsaed

Collage of Informatics, Midocean University, Comoros

Acknowledgements

I want to take this opportunity to express my gratitude to Dr. Ashraf Alsaed for his tremendous support and invaluable assistance with my thesis work. I appreciate all the help he has provided me with throughout this process. My thesis has benefited from Dr. Alsaed's help and support in every way imaginable.

Also, I would like to extend my gratitude to Midocean University, College of Informatics faculty and staff for the resources and academic climate they created which allowed me to conduct my research successfully. The support I received from them during my master's program has been greatly appreciated.

In addition to thanking Dr. Alsaed and Midocean University, I want to take this opportunity to thank all my supporters: family, friends, coworkers, and individuals who were willing to lend their time and energy to provide feedback on my research effort. Their support and encouragement greatly aided in the completion of my project.

Finally, I want to especially thank my family for the love, support, patience, and encouragement they have provided me throughout my academic career; they have always been my greatest source of inspiration and strength.

Abstract

Social media sites create abundant, ever-growing flows of information, making it increasingly hard for users to assess information of interest to them efficiently when it aligns with their interests, goals, and information needs. Current recommendation systems within these platforms mostly rely on passive behavioral signals likes, clicks, scrolling behavior, and dwell time to learn about user preferences. While effective for maximizing user engagement to the platform, these systems lack transparency, build a limit on user autonomy, and offer little to no information on the actual reasoning behind recommendations. Conventional recommendation algorithms also remain closely tied to the social media platform, which stops users from maintaining a consistent preference profile across social media platforms. All these difficulties culminate to lead to information overload, diminished trust, privacy violations, and feelings of distancing between recommended information and users' interests. This project showcases an alternative, user-directed recommendation framework that is delivered as a Chrome extension, enabling users to explicitly shape their content experience. Rather than depending on platform-native algorithms, the proposed framework allows users to manually annotate posts that they see from their social media feeds as either "**Recommended**" or "**Unwanted**." The user's annotations will form their own data set reflective of their own explicit preferences as opposed to inferred behavioral patterns. A simple ML model that has been locally repurposed via Fine Tuning with techniques such as low-rank adaptation (LoRA) has the ability to use these annotations to learn user-specific relevance patterns and anticipate likely user interactions with upcoming posts based on what users have shared previously. Additionally, since all fine-tuning has taken place on the user device, it helps protect the user's privacy by eliminating dependency on centralized collections of user data and dictating that sensitive information about user behavior should not be collected or shared. The cross-platform website uses a Chrome extension, which provides the ability to customize your experience on both platforms without needing to rely on Private API or Backend Service access. Instead, it relies only on publicly visible text the browser extracts. This model provides compatibility and transparency and is not ethically objectionable, unlike some paradigms based on surveillance-inspired tracking. The recommendation model is inherently dynamic, developing over

time as users contribute more annotations, and will become more accurate and personalized while allowing users to retain agency over the recommendations. The recommendation framework is able to surface relevant posts, deprioritize unwanted content, and provide a transparent alternative to conventional opaque recommendation models. In sum, this project puts forward a novel, privacy-centric, user-centered approach to content recommendation based on explicit feedback, transparency, and user agency. It combines interactive annotation, lightweight learning, and cross-platform functionality to arise at a viable and ethical approach to improve content consumption online when algorithmic filtering and information overload dominate the field.

Table of Contents

List of Abbreviations.....	viii
List of Figures	ix
List of Tables.....	x
Chapter 1 Introduction	1
Chapter 2 Background and Related Work	4
Chapter 3 Proposed Approach (Methodology)	7
Chapter 4 :Experimental Results and Discussion	10
Chapter 5 Conclusion and Future Work	11
References	13

List of Abbreviations

- API – Application Programming Interface
- CF – Collaborative Filtering
- CARS – Context-Aware Recommender Systems
- Chrome Ext. – Chrome Extension
- DB – Database
- DL – Deep Learning
- JSON – JavaScript Object Notation
- LLM – Large Language Model
- LoRA – Low-Rank Adaptation
- ML – Machine Learning
- MSN – Mobile Social Network
- MySQL – My Structured Query Language
- NLP – Natural Language Processing
- Node.js – JavaScript Runtime Environment
- OCR – Optical Character Recognition
- OSN – Online Social Network
- PEFT – Parameter-Efficient Fine-Tuning
- RS – Recommender System
- UI – User Interface
- URL – Uniform Resource Locator
- XLM-R – Cross-lingual Language Model, RoBERTa
- mT5 – Multilingual Text-to-Text Transfer Transformer

List of Figures

- Figure 3.1 Overall architecture of the proposed recommendation system
- Figure 3.2 Workflow of the Chrome extension for post collection and annotation
- Figure 3.3 User interface of the Chrome extension showing “Recommended” and “Unwanted” labels
- Figure 3.4 Data flow from social media platforms to backend API and MySQL database
- Figure 3.5 Structure of the labeled dataset collected from X, Facebook, and LinkedIn
- Figure 3.6 Backend API processing pipeline for receiving and storing labeled posts

List of Tables

- Table 3.1 Supported social media platforms and collected post attributes
- Table 3.2 Structure of labeled data collected through the Chrome extension
- Table 3.3 Description of metadata fields stored in the MySQL database
- Table 3.4 Public datasets used for initial experimentation and model preparation

Chapter 1 Introduction

Since social media has become the main source of communication, entertainment, education, and news for most, users can no longer keep track of the sheer volume of posts that are produced every day. So to keep up with that, social media sites use recommendation algorithm-based recommendation engines to organize the content in user feeds and prioritize that feed based on the user's previous interactions. The recommendation algorithm-based recommendation engines are a significant portion of what users view and ultimately what they interact with.

1.1 Problem Definition

Although mainstream recommendation systems have increased user engagement, four significant shortcomings of these approaches exist. The first limitation is that the primary method for inferring user preferences is to use implicit behavioral signals (e.g., clicks, watch time, scrolling); however, these types of signals tend to be noisy and often easily misinterpreted. Secondly, because most mainstream recommendation algorithms are opaque (i.e., users do not receive any explanations about the content recommended), users are less likely to trust the algorithm and correct undesired patterns. Thirdly, most mainstream recommendation algorithms tend to optimize engagement for the platform (i.e., encouraging more time spent using the platform) instead of for goals of users, resulting in the emergence of irrelevant, repetitive, and/or sensationalized content in user feeds. Fourthly, most personalization models require extensive centralized tracking, raising significant privacy concerns, and restrict the learned preference profiles to the single platform where they were created, limiting personalizing experiences across platforms.

This research project attempts to solve this issue via user-directed recommendations via a Chrome extension. Users will explicitly label posts as “Recommended” or “Unwanted,” creating clear preference datasets. Using these explicit datasets, a lightweight model that runs locally on the user’s device will learn from user-generated feedback in order to make personalized predictions regarding future ranking and filtering decisions. By transferring control of personalized recommendations from platform-inferred data to user-annotated data and keeping the training local, this proposed system will provide a more transparent, private, and less overwhelming experience for users while providing

more consistet personalized experiences across platforms.

1.2 Thesis Objectives

Develop a Chrome extension capable of overlaying annotation buttons (“Recommended” and “Unwanted”) on social media posts within supported platforms.

Collect labeled feedback from users to create a personalized training dataset reflecting individual content preferences.

Train and optimize a recommendation model that learns from these annotations to predict which future posts should be highlighted or hidden.

Implement local or lightweight model training to preserve user privacy and minimize data transfer.

Evaluate the system’s effectiveness in improving content relevance and user satisfaction through both quantitative (accuracy, precision, recall) and qualitative (user experience) metrics.

1.3 Thesis Contribution

This thesis contributes to the domain of social media recommendation systems as it describes a user-centered and privacy-preserving recommendation framework through the implementation of a Chrome extension. Typical recommendation approaches are largely based on implicit behaviour feedback. However, the proposed system...will enable users to provide explicit either a "recommended" or an "unwanted" label to every social media post, and this will create a much clearer and more transparent preference model. Furthermore, the use of lightweight...Low-Rank Adaptation (LoRA) techniques will enable efficient model adaptation to a user device without having to collect large amounts of centralised data. The cross-platform browser-based architecture described in this thesis will also allow the system to operate on publicly-viewable material and will therefore make the approach more flexible and independent from proprietary platform APIs. This thesis provides a practical, transparent and ethical alternative to traditional opaque recommendation systems through the combination of explicit feedback, local learning and cross-platform usability.

1.4 Thesis Organization

The dissertation is organized into five sections. The first section is very introductory: it presents the foundation of the research; describes the principal aims and objectives of the research; provides a list of the primary contributions the study will make; and summarizes how the material in this study is arranged or structured. The second section addresses the theoretical background associated with the study of recommender systems, social media personalization, browser-based interventions, and lightweight adaptation techniques such as Low-Rank Adaptation (LoRA), in the context of existing literature. The third section details how data was collected using a Google Chrome extension by outlining the overall system architecture and design of the extension; describing the data collection process; presenting the backend RESTful API used for data storage; and providing details of how the data is stored in a database. The fourth section provides a description of how the experiments were established, what datasets were used, how they were evaluated, and what results were obtained. The final section provides a summary of the results obtained in each chapter to provide insight into future directions for research.

Chapter 2 Background and Related Work

A survey of Campana et al.(2023) [1] reported on the different types of recommender systems (RS) available for Online Social Networks (OSN) and Mobile Social Networks (MSN) that make use of social relationships and user mobility to improve the design of modern RSs. First, the paper introduced four main approaches to RSs: Collaborative Filtering (CF), Content-Based Filtering, Graph-Based Approaches, and Context-Aware Recommender Systems (CARS). As these approaches depend solely on existing conventional interaction data, each suffers the common shortcomings of this reliance: scalability issues, cold start problems, and sparser feedback due to fewer interactions. The authors argue that social networking environments, such as OSNs, have the potential to significantly boost the quality of recommendations as they offer rich social signals, including friendship links/follow Markup tags, trust channels, tagging behavior, geographical contexts, and the availability of user-generated content. Additionally, they categorised the published literature on recommendation systems into the following system types: Friend Recommendations Based on Link Prediction, Social Trust Enhanced Item Recommendations, Tag Based Recommendations, and Geolocation Based Recommendations. MSN environments create a set of unique challenges for RSs relating to the decentralised nature of communication between devices, the intermittent nature of the connections between devices, limited device resource availability, and privacy issues; thus, RSs are required to work primarily with local device data resources and to rely on opportunistic communication to connect devices, rather than following a centralised model. The authors therefore reviewed multiple distributed and decentralised approaches to RSs including OppCF, MobHinter, diffeRS, locPref, and p-PLIERS, illustrating how RSs may continue to be successful even with the absence of a centralised approach.

Piccardi et al.'s (2024) [2] current study describes a newly developed methodology for conducting field studies of social media feed ranking using web browsers as a platform to conduct experiments and/or interventions. The impetus for this study is based upon the growing body of research regarding how feed ranking systems affect user well-being and the way in which users access information through their feeds, as well as the increasing number of restrictions placed on the access that platforms have to feed ranking and data selection

systems. The authors have developed a methodology called a browser-based client-side ‘reranking pipeline’ that allows users to intercept the content of feeds as it is delivered to them, select which content they would like to see and apply the desired interventions to the selected content, thereby re-rendering the modified feed for their respective participants. The authors have also identified that there are three different types of intervention methods—namely, (1) downranking ‘bad’ content, (2) upranking ‘good’ and/or properly balanced content and (3) rewriting or editing content. Furthermore, the authors have discussed the technical challenges associated with implementation (e.g., inventory constraints, latency, consistency in handling user interface changes) as well as the ethical issues (consent, privacy and minimization of harm). The authors describe various ways to measure outcomes using a combination of survey data, passive behavioral metrics, and real-time ecological momentary assessments, as well as practical information about enrolling participants in the study, assigning participants to conditions, designing the interface, optimizing performance, and storing participant data securely. The study's conclusion is that extension-based reranking can facilitate ecologically-valid experimental research conducted in social media platforms, even if the social media company does not support such activity; therefore, the experimental designs used will still produce rigorous results for measuring the effect of interventions.

Hu et al.'s (2021) [\[3\]](#) study introduces LoRA (Low-Rank Adaptation) as a low-cost way of adapting a large number of language models (for example: GPT-3) for downstream tasks, without having to fine-tune/old model parameters for those tasks. The authors state that traditional fine-tuning for extremely large models (e.g., GPT-3) is costly in storage and computation, and that traditional fine-tuning methods (full tuning, adapter, prefix tuning) can add inference overhead (e.g., latency) and/or reduce available sequence length. To address these issues, LoRA uses the insight that many model fine-tuning updates are low-dimensional; therefore, if only low-dimensional updates are required to perform the task, then LoRA can be used without the need to fully update a model for all tasks, unlike adapter and prefix tuning. LoRA's method involves freezing the pre-trained weights and representing ΔW , the update, as a product of two small matrices (BA) whose rank is extremely low, generally 1 or 2. The only matrices trained will therefore be BA, which dramatically reduces the number of parameters that need to be trained as well as the amount of memory consumed. The authors of this article find that during their experimentation with RoBERTa, DeBERTa, GPT-2, and GPT-3, LoRA is able to reach performance levels comparable to or

higher than that of full fine-tuning, using as much as 10,000x fewer trainable parameters. Furthermore, LoRA allows for the merging of the LoRA updates into the base weights without introducing any notable inference latency. The analysis also proposes that LoRA learns low-rank and meaningful adaptations that are primarily concentrated in the attention projection matrices (e.g., W_q and W_v) and that they are stable across multiple trials. The conclusions drawn indicate that adaptation can be effectively captured in a low-rank space. Overall, the paper establishes LoRA as a valid way of scalably and practically adapting large pre-trained models to new tasks in cases where compute and memory limitations are a primary concern. According to Kong et al.'s (2025) [\[4\]](#) article, utilizing a single instance of a LoRA Module across an entire community of LLM-Based Sequential Recommenders creates a high risk for negative transfer due to model and user differences throughout the entire lifetime of the recommender system. The authors go on to describe how the common LLM-SR process creates the optimal training pathways for the creation of end-to-end training data via PEF-trained LoRA Modules for prompt-texts by converting user histories into prompts and utilizing prompt-context and expected translation prompts as potential output hypotheses within a retraining LLM process; however, they also show that different clusters of behavior have a tendency to present issues in the manner that their respective gradients misalign and lead to conflicts that arise during the creation of an optimal single LoRA Module'. They approach the challenge of large-scale sequence recommendation by developing a single multi-task training system with LoRA. Model updates use an MoE structure with a single gating network, driven by what a user has selected from prior recommendations and captured through a recurrent model like SARPS. Users identify "sequences," or sets of recommendations based on the user's selected sequence, and each sequence identifies its expert "weight." A recommendation generated by the MoE can then be combined with the user's prior selections to update formulations of LoRA for the specific user and sequence. iLoRA provides users with greater accuracy of their recommendations than does LoRA using fewer than 1% additional trainable parameters. A scalable method of providing instance-wise expert activation will also be beneficial to the future of personalization and elimination of negative transfer within sequential recommendation.

Chapter 3 Proposed Approach (Methodology)

The implementation and methodology for creating an interactive social media recommendation system using a web-based application framework for users to label their own preferred social media items. This differs from most current recommendation systems that utilize only implicit user signals, such as click-through rates and how long users view items. The proposed interactive recommendation system will use explicit user feedback through classifications of recommended and unwanted traffic as the basis for providing more transparent, user-centered, and privacy-aware recommendations.

The framework created to implement this interactive recommendation system has three major components: a Chrome extension, a backend Node.JS application to handle all business logic, and a MySQL database for storing user data. The Chrome extension works on the social media platforms of Facebook, LinkedIn, and X. The purpose of this extension is to identify visible posts in the user's social media feed, extract relevant content from these posts, and provide a straightforward interface for the user to annotate whether each post is recommended or unwanted. The interactive recommendation system will capture explicit preference signals that will later be used to create personalized recommendations and filtered content.

Each time a user labels a post, an automated extension pulls relevant text snippets and/or metadata related to the post from the interface. This data includes the content of posts, where they came from (platform), a unique local identifier for each post and the label assigned (as well as when a user has interacted with the post). After a user has labeled the post, the extension sends those user-labeled posts to a back end API built on Node.js. The back end API validates and processes the records, then stores them to a MySQL database. The MySQL database is the main source for finding examples users have labeled, and provides the basis for future analysis, machine learning model development, and evaluation.

Based upon this workflow, the system has three key components. In the data acquisition stage, user-generated content (e.g., posts) is acquired from X, Facebook and LinkedIn using the Chrome extension. In the user annotating stage, users annotate the acquired post based on the post's relevancy. In the back end processing and storage stage, user-labeled posts are sent through the API and stored in the database. Using this process, user-specific data matching actual user

preferences can be created, which can be done more accurately (due to direct interaction) than by traditional engagement data (i.e. impressions/click-through rates).

One of the significant advantages of this methodology is the identification of explicit user interaction as the means of produce data, rather than data being inferred through behaviour. This gives more transparency since the users are made familiar with how the data they provide is generated and used, in addition to enhancing the quality of the data since the data will be derived from intentional and interpretable feedback rather than from uncertain behavioural traces.

Having cross-platform functionality of the Chrome extension means that the identical process of annotating content would be applied to multiple social media platforms; therefore, an enhanced, versatile, and comprehensive profile of user preference could be developed by directly accessing the same authenticated database that stores the data produced through the Chrome extension.

The database of the data that is stored will be used to create or modify a recommended model that predicts whether or not a particular user will regard any newly encountered post as being relevant to or not relevant to them. Thus, the current implementation adequately provides the necessary systems to allow the development of an interactive and federated recommendation systems infrastructure. Collectively, the Chrome extension, database, and back end API comprise the basic structure through which large amounts of explicit preference user generated data can be collected for the development of user specific recommendation systems research.

This approach to collecting data and storing it also include a component that uses artificial intelligence, as defined by the technique of Low-Rank Adaptation (LoRA). By collecting labeled posts through a Chrome extension, we have access to a supervised dataset that can be used to train or adapt a recommendation model that incorporates the individual preferences of users. LoRA provides an efficient means of fine-tuning when compared to retraining a large language or text model, as it does not require a large number of parameters to be updated, but rather only requires that some portion of the model's parameters to be updated while the original model weight is fixed. This approach greatly reduces the computational resource requirements (cost/ memory), and the time taken to develop the recommendation model, and as such is more feasible when developing personalized recommendations. Furthermore, in terms of scalability and being privacy-aware, the use of LoRA is appropriate in the context of this research, because of its ability to perform user-specific adaptations with limited resources.

By continually collecting additional labeled examples, the recommendation model could be incrementally improved in its ability to make the distinction between content that is recommended, and content that is not wanted (i.e. spam), thereby increasing the overall accuracy and relevance of the recommendations being provided.

Chapter 4 :Experimental Results and Discussion

Chapter 5 Conclusion and Future Work

This thesis developed a user-centric and interactive recommendation framework for social media platforms via a Chrome extension, which addressed many of the limitations of traditional recommendation systems; specifically, their reliance on implicit behavioral cues, limited transparency, and dependence on centralized data collection. By providing users with an option to label posts as either ‘recommended’ or ‘unwanted’, this system allows for more transparency and personalization in content recommendations than do currently available recommendations. The architecture is made up of a Chrome extension, Node.js backend API, and MySQL database which together support collecting, transmitting, storing, and managing user preference information across many different platforms. A significant contribution of this research is the provision of a practical infrastructure for explicit feedback-based recommendation systems. This proposal allows users to play an active role in shaping their experiences with the platform, not just relying on engagement-based algorithms driven by the platform. This will enhance the transparency of the recommendations made to users when using the platform and create a more usable and interpretable preference dataset for future use in building recommendation models. The research also demonstrated the value of using multilingual AI models and lightweight adaptation techniques (i.e., LoRA) to help create efficient, scalable, and privacy-aware personalized recommendations, including those based on Arabic-language data.

While this current implementation has successfully shown the ability to collect and store cross-platform labeled data, it serves merely as the first step toward a fully adaptive recommendation framework; currently, the system's focus is on creating the infrastructure for explicitly collecting data and for storing it in the backend. Future work will still expand this implementation into real-time recommendation generation, retraining, content re-ranking, and automatically predicting user preference through user input annotations. In particular, developing a recommendation model based on these collected annotations so that new posts can automatically be classified as either relevant or irrelevant will continue to evolve.

There are many other potential improvements to make as well. The next phase of system development could involve expanding the number of social media platforms supported within the system, improving the ability to process multilingual texts, especially Arabic-language content, and expanding the recommendation pipeline to provide for handling multimedia formats such as video and images. Finally, further research on integrating federated or on-device learning methods to lessen the risk of privacy violations by reducing the need for centralized model training will be an

important area of future work. Additionally, increasing transparency and accountability of the recommendation mechanism by providing users with explanations of why specific posts were suggested or filtered should be a part of the future development of this system in order to enhance users' confidence and trust in it and to increase the interpretation of the recommendations provided to them by the system.

In conclusion , This Research Has demonstrated that Browser Extension-Driven Explicit Feedback Recommendations Are a Viable Alternative to Traditional Opaque Recommendations Systems. With the Combination of Cross-Platform Data Collection, User-Based Annotation, and a Lightweight Personalized Model, The Proposed Social Media Recommender System Has the Potential to Provide a More Ethical, Transparent, and User-Focused Approach to Recommendations for Online Content. Results Suggest That When Users Empowered to Influence Recommendations Directly, The Relevance and Trustworthiness of Online Content Will be Increased.

References

- [1] M. G. Campana and F. Delmastro, “Recommender systems for online and mobile social networks: A survey,” *arXiv*, 2023. [Online]. Available: <https://arxiv.org/abs/2307.01207>.
- [2] T. Piccardi *et al.*, “Reranking social media feeds: A practical guide for field experiments,” *arXiv*, 2024. [Online]. Available: <https://arxiv.org/html/2406>.
- [3] E. J. Hu *et al.*, “LoRA: Low-rank adaptation of large language models,” *arXiv*, 2021. [Online]. Available: <https://arxiv.org/abs/2106.09685>.
- [4] X. Kong *et al.*, “Customizing language models with instance-wise LoRA for sequential recommendation,” *arXiv*, 2025. [Online]. Available: <https://arxiv.org/abs/2408.10159>.
- [5] A. Go, R. Bhayani, and L. Huang, “Twitter sentiment classification using distant supervision,” Stanford Univ., CS224N Project Rep., 2009.
- [6] CrowdFlower, “Twitter US Airline Sentiment,” *Kaggle*, 2015. [Online]. Available: <https://www.kaggle.com/datasets/crowdflower/twitter-airline-sentiment>.
- [7] K. Regmi, “Job Title and Job Description Dataset,” *Kaggle*. [Online]. Available: <https://www.kaggle.com/datasets/kshitizregmi/jobs-and-job-description>.
- [8] asaniczka, “1.3M LinkedIn Jobs & Skills (2024),” *Kaggle*, 2024. [Online]. Available: <https://www.kaggle.com/datasets/asaniczka/1-3m-linkedin-jobs-and-skills-2024>.