MAWI Rec: Leveraging Severe Weather Data in Recommendation

Brendan Duncan
University of California, San Diego
La Jolla, CA, USA
baduncan@ucsd.edu

Surya Kallumadi
Lowe's Companies, Inc.
Mooresville, NC, USA
surya@ksu.edu

Taylor Berg-Kirkpatrick
University of California, San Diego
La Jolla, CA, USA
tbergkirkpatrick@ucsd.edu

Julian McAuley
University of California, San Diego
La Jolla, CA, USA
jmcauley@ucsd.edu

ABSTRACT

Inferring user intent in recommender systems can help performance but is difficult because intent is personal and not directly observable. Previous work has leveraged signals to stand as a proxy for intent (e.g. user interactions with resource pages), but such signals are not always available. In this paper, we instead recognize that certain events, which *are* observable, directly influence user intent. For example, after a flood, home improvement customers are more likely to undertake a renovation project to dry out their basement. We introduce MAWI Rec, a recommender system that leverages severe weather data to improve recommendation. Our weather-aware system achieves a significant improvement over a state-of-the-art baseline for online and in-store datasets of home improvement customers. This gain is most significant for weather-related product categories such as roof panels and flashings.

CCS CONCEPTS

• **Information systems** → **Recommender systems**; *Personalization*; • **Computing methodologies** → *Neural networks.*

KEYWORDS

Recommender systems, personalization, transformers

ACM Reference Format:
Brendan Duncan, Surya Kallumadi, Taylor Berg-Kirkpatrick, and Julian McAuley. 2024. MAWI Rec: Leveraging Severe Weather Data in Recommendation. In *18th ACM Conference on Recommender Systems (RecSys '24), October 14–18, 2024, Bari, Italy*. ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3640457.3688157

1 INTRODUCTION

Recommender systems that infer user intent can greatly improve performance (see Section 2). If we for example consider two users browsing hammers online, a user whose intent is to build a fence will have different shopping behavior from a user whose intent is to install a bookshelf.

RecSys '24, October 14–18, 2024, Bari, Italy

ACM ISBN 979-8-4007-0505-2/24/10
https://doi.org/10.1145/3640457.3688157

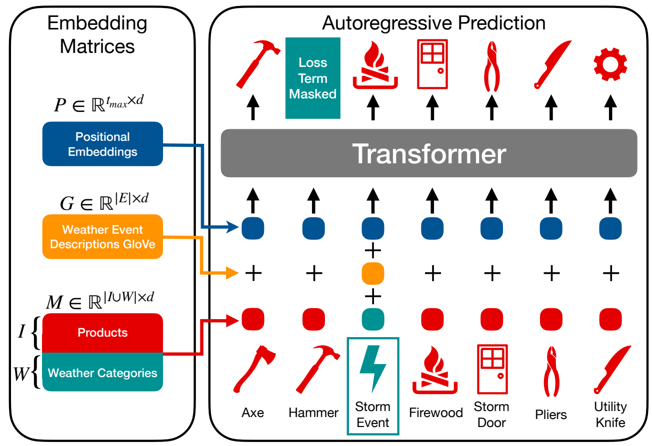

Figure 1: We take the sequence of a user's interactions (red) and local weather events (teal), and predict the next interaction (weather event loss terms are masked). To encode sequences, we sum together three embeddings: a unique embedding for each product (red) or weather event category (teal), a positional embedding (blue), and the sum of GloVe word embeddings for the weather event description (orange).

Traditional collaborative filtering recommender systems, which model user-item compatibility, are not powerful enough to deeply capture user intent, which changes through time. Transformer-based recommender systems like SASRec are sequential and allow recommendations to change across time, so they are better suited to capturing intent, and can be augmented to explicitly model it. The difficulty with modeling intent is that it is largely subjective and personal (see e.g. discussion in Tanjim et al. [25]). To solve this, some authors such as Tanjim et al. [25] treat intent as a latent variable that is not directly observed. Other work instead makes use of data that serves as a proxy for intent, such as user interactions with resource pages[13].

In this paper we take a different approach: we recognize that certain events directly influence user intent, and we build a model that leverages data about such events. For example, when a store announces a sale, this may influence users' intent to purchase certain items; in this case, we could leverage information about the sale begin date, end date, and type of sale (which categories of items are included, or which season in the case of seasonal sales) to improve recommendation. As another example, in the domain of home improvement, severe weather incidents are events that can cause a user to purchase certain products and undertake certain renovation projects, such as cleaning up the yard after a storm, repairing a broken screen door, purchasing and installing a storm door, or weatherproofing a house with flashings. This paper focuses on the second use case by leveraging data about severe weather to improve home improvement recommendation.

Our system is called MAWI Rec, for "Merchandise & Weather Interleaved". It incorporates weather events as tokens interleaved into the input of SASRec[19], a state-of-the-art transformer-based sequential recommender system. The main dataset for our experiments consists of historical online customer interaction data from Lowes.com. Lowe's is one of the largest and oldest home improvement store chains in the US. We also include experiments on a secondary dataset of in-store purchases from Lowe's stores. These datasets are augmented with NOAA data about severe weather events affecting customers. [1] Our results show significant improvement over a SASRec baseline that does not include weather data. The improvement over the baseline is most significant for weather-related product categories, such as storm doors and flashings.

2 RELATED WORK

As explained above, we view severe weather events as influencing the high-level intent of home improvement customers–namely, their intended renovation project. Therefore relevant previous work includes recommender systems that infer user intent. Some such methods have sought to leverage low-level intent signals, such as product categories. Examples of this type include CoCoRec [7], DSSRec [21], ICL [9], and ASLI [25]. Query-SeqRec [15] investigates the addition of search queries to better determine intent. As described further in Section 4, our method is most influenced by Duncan et al. [13], which includes user interactions with resource pages as an intent signal. Resource pages include recipes for grocery recommendation or DIY how-to guides for home improvement recommendation. We treat weather events similarly to resource pages in that work.

To the best of our knowledge, no previous work has incorporated severe weather event data to improve product recommendation. However, several previous works have utilized more general weather data. For example, Jaafar et al. [17] makes recommendations of appliances such that the device operating temperature and humidity ranges are compatible with local climate. Behl et al. [2] incorporates weather forecast information into a recommender system to help facilities managers predict electric grid load in buildings and choose appropriate control actions to manage this load. Similarly, Kamatchi and Parvathi [18] incorporates weather information to make climate predictions and recommend actions to increase crop production.

Li et al. [20] have discovered a statistically significant relationship between weather and mood by studying the effect of weather on user posts to Twitter. This effect of weather on user mood has been leveraged for music recommendation [12], clothing style choices [10], and grocery purchases [8]. These methods incorporate general weather information into context-aware, non-transformer based systems. As we discuss below in Section 3, including general weather information in this way does not significantly improve results for our home improvement dataset; we instead get better performance from leveraging detailed descriptions of severe weather events.

By far the largest area of research to have combined recommendation with weather data is weather-aware navigation systems and tour guides. These generally use current and forecasted weather data to improve the planning of sightseeing and navigation. Many such papers view sightseeing as a POI recommendation problem, and either make weather-aware POI or path recommendations or allow users to filter POI results or navigation path based on weather or travel time, where travel time is affected by weather events. Examples of weather aware tour guides and POI recommender systems include Bellotti et al. [3], Gavalas and Kenteris [14], Meehan et al. [22], Braunhofer et al. [5], Braunhofer et al. [6], and Hussain et al. [16], all of which predate transformer models. Other methods include Costa et al. [11] and Trivedi et al. [26], both of which are path-planning navigational systems to minimize travelling cost functions which include penalties for exposure to negative weather. Bahramian et al. [1] is an interactive artificial neural network-based navigational system, in which a user selects and rates tours. These models are not relevant to our domain of home improvement product recommendation.

3 DATASET

We use two types of data for our experiments: historical online interaction and sales data from the website of a major US home improvement company, Lowe's, and severe weather information from the National Oceanic and Atmospheric Administration (NOAA). The severe weather event datapoints have been filtered and associated with nearby store locations. Online users have also self-reported their local store, so we use local store location to match severe weather events to customer interaction data. In addition to location information, we also extract the event start date and the text description of the weather event from the NOAA severe weather dataset.

In addition to this severe weather data, we also considered augmenting our recommender system with several other types of weather data: daily temperature and precipitation data, forecasted temperature and precipitation data. We found little improvement by including non-severe weather information, which indicates that only severe events had a strong effect on intended renovation project. Weather alert data, such as heat advisories and flood warnings also failed to significantly improve results. This suggests that for the most part weather alerts on their own do not significantly affect users' renovation projects. It should be noted that not all alerts end up corresponding to severe weather events, since the volume of alert data is much larger.

We only consider users who have at least one product interaction that occurs within 30 days of a severe weather event in their local area. We consider the weather event to be in a user's local area if it occurred within 30 kilometers of their specified local store location. We chose these values because they worked well empirically for metrics on the validation set. Because DIY projects and

[1] Severe weather dataset and code released at https://github.com/brendanduncan/recsys2024.

Table 1: Data statistics after pre-processing

# users	275,887
# products	3,005
# average interactions/user	13.40
# average weather events/user	1.40

weather events vary significantly by season, our dataset includes one month of data from each season: in particular, we include data from March 2021, June 2021, September 2021, and January 2022. We chose one month from each season, under the assumption that weather patterns across different seasons will differ substantially from one another, while weather patterns from months within a single season will differ much less substantially. See Table 1 for dataset statistics.

For the online interaction data, we reduce data sparsity by recommending the lowest-level product category, rather than individual items. Therefore, we would recommend vent pipe flashing rather than SuperVent 6-in x 24-in Aluminum Vent and Pipe Flashing. This also accords with Duncan et al. [13], which also recommends lowest-level home improvement product category for DIY project-aware recommendation. For simplicity, we refer to lowest-level product categories as "products" throughout this paper.

Our method does not distinguish between different interactions types (click, add-to-cart, and purchase). We remove all duplicate interactions for each user by keeping the earliest such action, which ensures a given product appears at most once in each user's sequence. Thus the system is trained to predict novel products rather than re-recommend previously seen items.

In Section 6, we also consider the performance of our method on an additional dataset of in-person purchases at store locations. This dataset includes 3,574,916 purchase events, 2,295 items, 24,302,137 product purchases and covers the same four months as the online dataset. Although the in-store dataset has more volume than the online interaction data, it lacks sequential information. Thus it requires different processing, as described in Section 6.

4 METHOD

Given a temporally-ordered sequence of product interactions $S^u = (S^u_1, S^u_2, ...S^u_{|S^u|})$ for each user u, sequential recommendation models the probability of the next item interaction at each stage in the sequence:

$$P(S^u_t = i | S^u_1, S^u_2, ..., S^u_{t-1}). \quad (1)$$

We augment these sequences with local weather events affecting each user u, and use κ to indicate whether a given interaction S_t is a product interaction ($\in I$) or weather event ($\in W$):

$$S^u_t \in \begin{cases} I, & \text{if } \kappa(S^u_t) = 0 \\ W, & \text{if } \kappa(S^u_t) = 1. \end{cases} \quad (2)$$

This idea is similar to Duncan et al. [13], who include user interactions with resource pages in the interaction sequence. A crucial difference is that, while that work aims to model resource page interactions, we do not want to predict weather events. Therefore, $P(S^u_t = i | S^u_1, ..., S^u_{t-1})$ is only defined for $i \in I$.

We compute the probability in Equation 1 using the SASRec architecture [19], which is an auto-regressive transformer model [27], in which interaction embeddings replace word embeddings. To avoid training our system to model weather events, we mask out loss terms corresponding to weather events ($\kappa(S^u_t) = 1$), as shown in Figure 1. We include learned positional embeddings (matrix $P \in \mathbb{R}^{t_{max} \times d}$) to capture sequential information.

We create a joint embedding matrix for both products and weather events $M \in \mathbb{R}^{|I \cup W| \times d}$, where d is the embedding vector length. We experiment with two W matrices: (1) a $1 \times d$ matrix, whose single embedding indicates that S^u_i corresponds to a severe weather event; and (2) a $k \times d$ matrix, with severe weather events distributed between k categories. We determine weather categories by performing Birch clustering[2] on tf-idf vectors generated from NOAA weather event descriptions. In addition to using the Birch algorithm for clustering, we also experimented with categorizing weather events into roughly snowy, hot, windy, rainy, and other, using simple keyword matching. We found that the Birch clustering performed slightly better, so we do not include our keyword matching experiments in this paper.

As seen in Figure 1, input positions that correspond to a weather event also receive an additional embedding from embedding matrix $G \in \mathbb{R}^{|E| \times d}$, where E is the set of individual weather events. Embeddings in G are computed by summing the GloVe word embeddings [24] for each word in the NOAA description of the weather event, after removing stopwords. These descriptions are on average 3.7 sentences long and contain an average of 45 words. See Table 4 for examples of weather event descriptions.

We sum all embeddings together to obtain the input embeddings. The sequence of input embeddings is then fed into the lowest layer of the SASRec transformer model.

5 EXPERIMENTS

We split the data into train, validation, and test sets as in the original SASRec paper. Each user sequence is split into three, with the last interaction $S^u_{|S^u|}$ placed into the test set, the next-to-last interaction $S^u_{|S^u|-1}$ into the validation set, and the remainder of interactions into the training set.

For our experiments, we use a two layer transformer and 100 dimensional embeddings. Our SASRec implementation uses PyTorch and a binary cross entropy loss.

Our experiments consider the following methods:

- **SASRec**: The SASRec architecture, which considers only product interactions without incorporating any weather information.
- **MAWI w/o cats**: A modified version of MAWI Rec which does not include multiple weather categories determined by Birch clustering. Instead, it includes only a single weather category embedding ($W \in 1 \times d$). It shows the partial improvement achieved by only adding GloVe embeddings.
- **MAWI w/o GloVe**: A modified version of MAWI Rec without including GloVe embeddings from matrix G. This method shows the partial improvement achieved by only adding categorical weather embeddings ($W \in 5 \times d$).

[2] Sklearn implementation [23] with threshold of 0.5, 5 clusters, and branching factor of 50.

Table 2: Results of our method, MAWI Rec, against a SASRec baseline for online interaction data.

Metric	NDCG		Hit Rate	
	$k = 5$	$k = 10$	$k = 5$	$k = 10$
SASRec	0.0942	0.1215	0.1676	0.2520
MAWI w/o GloVe	0.1229	0.1472	0.1811	0.2562
MAWI w/o cats	0.1253	0.1493	0.1838	0.2581
MAWI Rec	**0.1282**	**0.1523**	**0.1872**	**0.2621**

- **MAWI Rec**: The full weather-aware method as described in Section 4, which includes GloVe embeddings and categorical weather embeddings.

6 RESULTS

Table 2 compares the above four methods on the online dataset. We see that the three MAWI methods outperform the baseline, and that including both GloVe embeddings and weather categories works better than only including either one of these. We suspect that including weather categories helps prevent the system from overfitting to spurious signals from the GloVe features.

To further demonstrate that MAWI Rec's performance gain comes from knowledge about severe weather events, in Table 3 we break down the recommendations into higher-level product categories and show that weather-related categories have significantly higher than average percent improvements over the SASRec baseline. Interactions in flashings, storm doors, and exterior doors suggest users waterproofing and storm-proofing their house after severe weather, while screen door and roof panel purchases are likely to repair exterior damage after a storm.

In Table 4, we give example recommendations for several weather-related product categories. For all these examples, the weather-aware system was able to predict the correct interaction category as its top recommendation, while the SASRec baseline failed to include it in its top ten suggestions. In the first example, hail and thunderstorms influenced the user's purchase of a closer for their outside door; in the second, heavy rainfall resulted in the purchase of flashing for waterproofing a house; and in the third, flooding resulted in the purchase a portable fan, possibly for drying out a flooded house.

While the above experiments confirm the benefit of our method for online users, we suspect severe weather causes many users to make in-store purchases rather than online purchases, due to the emergency nature of some repairs. To test this, we compare results from MAWI Rec to those of the SASRec baseline on the dataset of in-store purchases (see dataset statistics in Section 3). The in-store purchase dataset was more limited than the online purchase dataset: we only had access to the set of products bought at different store locations at different times, and user information was lost, preventing us from creating time-based product sequences for each user. In order to create sequences to be used with our method, we decided to randomly order the group of products for each store purchase and predict the last product of the sequence. Table 5 presents the results, and shows that MAWI Rec again outperforms the baseline.

Table 3: Percent improvements of MAWI Rec over the SASRec baseline for the test set. We break recommendations into higher-level product category, showing that our method does dramatically better for weather-related categories. All values are statistically significant ($p < 0.05$) according to the bootstrap test[4], unless indicated with a dagger (†).

Metric % improvement	NDCG % imp.		Hit Rate % imp.	
	$k = 5$	$k = 10$	$k = 5$	$k = 10$
Flashings	73%	47%	38%	15%
Roof Panels & Accessories	57%	49%	18%	18%
Screen & Storm Door HW	56%	41%	14%	0%†
Exterior Doors	38%	32%	11%	10%
Overall	36%	25%	12%	4%

Table 4: An example of three recommendations, along with weather event. The weather-aware method chose the correct product interaction as its top recommendation for each, while the SASRec baseline failed to choose the correct item in its top ten.

Weather event:
Scattered **thunderstorms** moved from interior south-central Florida to the Treasure Coast. One storm became severe as it moved into Saint Lucie County where it produced **ping pong to two inch hail** that resulted in damage to numerous vehicles.
Product interaction in test set:
Screen Door & Storm Door Closers

Weather event:
Heavy rainfall affected the region from the late morning of the 12th to the evening of the 14th as a cold front stalled over northern Arkansas moved back to the north as a warm front.
Product interaction in test set:
Drip Edge Flashing

Weather event:
... Marion County saw the most **impactful flooding** with numerous evacuations and water rescues taking place across the county. Southeast Kansas also saw **ssserious flooding** with the Neosho River seeing significant rises...
Product interaction in test set:
Portable Fans

This recommender system built from in-store purchases could be used by sales associates to recommend products to customers and to help plan product displays and inventory after a weather event.

7 CONCLUSION

In this work we developed a system, MAWI Rec, to make weather-aware product recommendations for users who recently experienced a local severe weather event. Our system leverages severe weather data in the form of detailed text descriptions of each event

Table 5: Results of MAWI Rec vs. a SASRec baseline on in-store purchase data.

Metric	NDCG		Hit Rate	
	$k = 5$	$k = 10$	$k = 5$	$k = 10$
SASRec	0.2646	0.3132	0.3922	0.5424
MAWI Rec	**0.2785**	**0.3284**	**0.4109**	**0.5649**

to significantly improve recommendation for home improvement customers. We view local severe weather incidents as events that directly influence user intent by causing them to purchase certain items or begin a new renovation project, such as purchasing and installing flashings to waterproof a house. Our method also generalizes to other cases in which events directly influence user intent, such as stores announcing sales events. MAWI Rec is thus a novel approach to the more general problem of inferring high-level user intent, which is particularly difficult due to the unobservable and personal nature of intent. We put forward this paper to assist other researchers model the influence of events such as severe weather on user intent.

REFERENCES

[1] Zahra Bahramian, Rahim Ali Abbaspour, and Christophe Claramunt. 2017. A cold start context-aware recommender system for tour planning using artificial neural network and case based reasoning. *Mobile Information Systems* 2017 (2017).

[2] Madhur Behl, Francesco Smarra, and Rahul Mangharam. 2016. DR-Advisor: A data-driven demand response recommender system. *Applied Energy* 170 (2016), 30–46. https://doi.org/10.1016/j.apenergy.2016.02.090

[3] Victoria Bellotti, Bo Begole, Ed H. Chi, Nicolas Ducheneaut, Ji Fang, Ellen Isaacs, Tracy King, Mark W. Newman, Kurt Partridge, Bob Price, Paul Rasmussen, Michael Roberts, Diane J. Schiano, and Alan Walendowski. 2008. Activity-Based Serendipitous Recommendations with the Magitti Mobile Leisure Guide. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Florence, Italy) *(CHI '08)*. Association for Computing Machinery, New York, NY, USA, 1157–1166. https://doi.org/10.1145/1357054.1357237

[4] Taylor Berg-Kirkpatrick, David Burkett, and Dan Klein. 2012. An empirical investigation of statistical significance in NLP. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*. 995–1005.

[5] Matthias Braunhofer, Mehdi Elahi, Mouzhi Ge, Francesco Ricci, and Thomas Schievenin. 2013. STS: design of weather-aware mobile recommender systems in tourism. In *Proceedings of the First International Workshop on Intelligent User Interfaces: Artificial Intelligence meets Human Computer Interaction (AI* HCI 2013); A workshop of the XIII International Conference of the Italian Association for Artificial Intelligence (AI* IA 2013), Turin, Italy, December 4, 2013*, Vol. 1125. CEUR-WS.

[6] Matthias Braunhofer, Mehdi Elahi, Francesco Ricci, and Thomas Schievenin. 2013. Context-Aware Points of Interest Suggestion with Dynamic Weather Data Management. In *Information and Communication Technologies in Tourism 2014*, Zheng Xiang and Iis Tussyadiah (Eds.). Springer International Publishing, Cham, 87–100.

[7] Renqin Cai, Jibang Wu, Aidan San, Chong Wang, and Hongning Wang. 2021. Category-Aware Collaborative Sequential Recommendation. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Virtual Event, Canada) *(SIGIR '21)*. Association for Computing Machinery, New York, NY, USA, 388–397. https://doi.org/10.1145/3404835.3462832

[8] Shampa Chakraverty and Aseem Mithal. 2018. IoT Based Weather and Location Aware Recommender System. In *2018 8th International Conference on Cloud Computing, Data Science & Engineering (Confluence)*. 636–643. https://doi.org/10.1109/CONFLUENCE.2018.8442454

[9] Yongjun Chen, Zhiwei Liu, Jia Li, Julian McAuley, and Caiming Xiong. 2022. Intent Contrastive Learning for Sequential Recommendation. In *Proceedings of the ACM Web Conference 2022* (Virtual Event, Lyon, France) *(WWW '22)*. Association for Computing Machinery, New York, NY, USA, 2172–2182. https://doi.org/10.1145/3485447.3512090

[10] Kyung-Yong Chung. 2011. Sensibility Ergonomics Design Recommendation System Using Weather WebBot. In *2011 International Conference on Information Science and Applications*. 1–6. https://doi.org/10.1109/ICISA.2011.5772371

[11] Constantinos Costa, Xiaoyu Ge, and Panos Chrysanthis. 2019. CAPRIO: Context-Aware Path Recommendation Exploiting Indoor and Outdoor Information. In *2019 20th IEEE International Conference on Mobile Data Management (MDM)*. 431–436. https://doi.org/10.1109/MDM.2019.000-7

[12] Iman Dolatkia and Fatemeh Azimzadeh. 2016. Music recommendation system based on the continuous combination of contextual information. In *2016 Second International Conference on Web Research (ICWR)*. 108–114. https://doi.org/10.1109/ICWR.2016.7498454

[13] Brendan Duncan, Surya Kallumadi, Berg-Kirkpatrick Taylor, and Julian McAuley. 2023. Jointly modeling products and resource pages for task-oriented recommendation. (2023). https://doi.org/10.1145/3543873.3584642

[14] Damianos Gavalas and Michael Kenteris. 2011. A web-based pervasive recommendation system for mobile tourist guides. *Personal and Ubiquitous Computing* 15 (2011), 759–770.

[15] Zhankui He, Handong Zhao, Zhaowen Wang, Zhe Lin, Ajinkya Kale, and Julian Mcauley. 2022. Query-Aware Sequential Recommendation. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management* (Atlanta, GA, USA) *(CIKM '22)*. Association for Computing Machinery, New York, NY, USA, 4019–4023. https://doi.org/10.1145/3511808.3557677

[16] Muhammed Mas-ud Hussain, Besim Avci, Goce Trajcevski, and Peter Scheuermann. 2016. Incorporating Weather Updates for Public Transportation Users of Recommendation Systems. In *2016 17th IEEE International Conference on Mobile Data Management (MDM)*, Vol. 1. 333–336. https://doi.org/10.1109/MDM.2016.57

[17] Basim Amer Jaafar, Methaq Talib Gaata, and Mahdi Nsaif Jasim. 2020. Home appliances recommendation system based on weather information using combined modified k-means and elbow algorithms. *Indonesian Journal of Electrical Engineering and Computer Science* 19, 3 (2020), 1635–1642.

[18] S. Bangaru Kamatchi and R. Parvathi. 2019. Improvement of Crop Production Using Recommender System by Weather Forecasts. *Procedia Computer Science* 165 (2019), 724–732. https://doi.org/10.1016/j.procs.2020.01.023 2nd International Conference on Recent Trends in Advanced Computing ICRTAC -DISRUP - TIV INNOVATION , 2019 November 11-12, 2019.

[19] Wang-Cheng Kang and Julian McAuley. 2018. Self-Attentive Sequential Recommendation. In *2018 IEEE International Conference on Data Mining (ICDM)*. 197–206. https://doi.org/10.1109/ICDM.2018.00035

[20] Jiwei Li, Xun Wang, and Eduard Hovy. 2014. What a Nasty Day: Exploring Mood-Weather Relationship from Twitter. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management* (Shanghai, China) *(CIKM '14)*. Association for Computing Machinery, New York, NY, USA, 1309–1318. https://doi.org/10.1145/2661829.2662090

[21] Jianxin Ma, Chang Zhou, Hongxia Yang, Peng Cui, Xin Wang, and Wenwu Zhu. 2020. Disentangled Self-Supervision in Sequential Recommenders. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (Virtual Event, CA, USA) *(KDD '20)*. Association for Computing Machinery, New York, NY, USA, 483–491. https://doi.org/10.1145/3394486.3403091

[22] Kevin Meehan, Tom Lunney, Kevin Curran, and Aiden McCaughey. 2013. Context-aware intelligent recommendation system for tourism. In *2013 IEEE International Conference on Pervasive Computing and Communications Workshops (PERCOM Workshops)*. 328–331. https://doi.org/10.1109/PerComW.2013.6529508

[23] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12 (2011), 2825–2830.

[24] Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. GloVe: Global Vectors for Word Representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, Doha, Qatar, 1532–1543. https://doi.org/10.3115/v1/D14-1162

[25] Md Mehrab Tanjim, Congzhe Su, Ethan Benjamin, Diane Hu, Liangjie Hong, and Julian McAuley. 2020. Attentive Sequential Models of Latent Intent for Next Item Recommendation. In *Proceedings of The Web Conference 2020* (Taipei, Taiwan) *(WWW '20)*. Association for Computing Machinery, New York, NY, USA, 2528–2534. https://doi.org/10.1145/3366423.3380002

[26] Rajani Trivedi, Bibudhendu Pati, and Chhabi Rani Panigrahi. 2022. wPOI: Weather-Aware POI Recommendation Engine. *Computación y Sistemas* 26, 2 (2022).

[27] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Ł ukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems*, I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.), Vol. 30. Curran Associates, Inc. https://proceedings.neurips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf