

Time-Aware Prospective Modeling of Users for Online Display Advertising

Djordje Gligorijevic
djordje@verizonmedia.com
Yahoo Research
Sunnyvale, CA

Jelena Gligorijevic
jelenas@verizonmedia.com
Yahoo Research
Sunnyvale, CA

Aaron Flores
aaron.flores@verizonmedia.com
Yahoo Research
Sunnyvale, CA

ABSTRACT

Prospective display advertising poses a great challenge for large advertising platforms as the strongest predictive signals of users are not eligible to be used in the conversion prediction systems. To that end efforts are made to collect as much information as possible about each user from various data sources and to design powerful models that can capture weaker signals ultimately obtaining good quality of conversion prediction probability estimates. In this study we propose a novel time-aware approach to model heterogeneous sequences of users' activities and capture implicit signals of users' conversion intents. On two real-world datasets we show that our approach outperforms other, previously proposed approaches, while providing interpretability of signal impact to conversion probability.

KEYWORDS

prospective advertising, deep learning, time-aware prediction

ACM Reference Format:

Djordje Gligorijevic, Jelena Gligorijevic, and Aaron Flores. 2019. Time-Aware Prospective Modeling of Users for Online Display Advertising. In *AdKDD '19, August 05, 2019*. Anchorage, Alaska, USA, 6 pages. <https://doi.org/10.1145/nnnnnnnn.nnnnnnnn>

1 INTRODUCTION

Online Display advertising has been one of the fastest growing industries in the world. In the U.S. alone, this industry amassed \$100 billion dollars in 2018¹. The concept of online display advertising (DA) is developed with the purpose of showing the most relevant ads to users anywhere online. The DA industry is composed of three major components: Supply Side Platforms (SSPs) who realize ad display opportunities on registered websites with user traffic and send ad requests to the following component, the online ad exchanges, who organize online auctions and forward the ad calls to several Demand Side Platforms (DSPs), a third component of the system, to bid on them. In order to have their ads shown to users, advertisers rely on DSPs to reach relevant users through ad display opportunities, bid on the auctions and display advertisers' ads. It is

¹<https://www.iab.com/wp-content/uploads/2019/05/Full-Year-2018-IAB-Internet-Advertising-Revenue-Report.pdf>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
AdKDD '19, August 05, 2019, Anchorage, Alaska
© 2019 Copyright held by the owner/author(s).
ACM ISBN 978-x-xxxx-xxxx-x/YY/MM.
<https://doi.org/10.1145/nnnnnnnn.nnnnnnnn>

the job of the DSPs to learn which users could be interested in the advertisers' products and could become their business in the near future. In order to achieve that, DSPs try to learn as much as possible about users, by collecting their online footprints through the data collected from advertisers websites, won auctions, third-party data providers and from owned-and-operated (O&O) properties.

Much of the DAs' business historically has been *retargeting*, a special case where ads are displayed to the users who have already shown interest in advertisers business. The goal of retargeting is to periodically remind users of the advertisers' products and hopefully generate conversions. However, as this particular form of DA is unlikely to bring *new* customers to the advertisers, they have shown increased interest in *prospective* targeting of users. The goal of *prospective* targeting is opposite of retargeting – users who have shown interest into advertisers business in the recent past should be excluded, and the goal becomes to generate new users as both visitors and converters for the advertiser. While the definition of retargeting users may significantly vary from one advertiser to another, in terms of the general advertising funnel (stages in which users are placed with respect to their probability of purchases of advertiser products [18]), prospective targeting should focus on users who are in the upper funnels (users further away from the conversion funnel). Conversely, in terms of advertising funnel, retargeting focuses on users in the very low funnel stages (users very close to conversion).

Prospective modeling of users poses a particularly difficult task for DSPs, as the direct signals of users interests (such are visits to advertisers website or recent conversions with the same advertiser) are no longer viable to be used. To maintain the high performance of user modeling, DSP is given a challenging task to generate powerful models which are able to detect relevant, yet weaker, signals users leave in their online trails and use them to the fullest extent. An example of such signals could be users' recent wedding related invoice which could signal potential interest in user purchasing furniture or flight ticket to the honeymoon, whereas any signals related to furniture or flight browsing on advertiser's website could not be consumed.

Moreover, a very important aspect of prospecting user modeling is explainability. Advertisers often require DSPs to provide insights into how predictions were made, what individual signals and what signal combinations seemed important during the modeling process. For the case of prospective modeling, these signals when interpreted can bring exceptional value to the advertiser, as they would be able to fine tailor future campaigns for different user groups that resonate better and reach more consumers which they potentially couldn't before.



Figure 1: Visualization of user activity sequence with different groups of activities ordered by the time they occurred and ending with the action of advertiser's interest.

To create a generic view on signals users leave, the most natural choice is to create a time-ordered sequence of activities user performed collected by the DSP. An example of one such sequence is provided in the Figure 1 where we observe multiple interactions of the user with different online properties such as mobile and desktop search, email receipts, reading news and interacting with ads. These trails of user's activities provide insight into the sequence of actions rather than sequence-oblivious features, moreover, actions always have assigned timestamp which carries a significant amount of additional information in terms of how close the subsequent events were or how much time passed between activity and event of interest (i.e. conversion). Modeling sequences of user events have been proposed in the past with great success [7, 18], however, to the best of our knowledge, never for prospective modeling of users. Moreover, utilizing activities data to the full extent such as temporal aspect has been largely ignored when modeling conversions in DA.

We summarize the contributions of this work below:

We motivate and propose the problem of prospective targeting in display advertising. To the best of our knowledge we are the first to discuss research for prospective modeling of users' interests.

We propose sequence learning approach to model time-ordered heterogeneous user activities coming from multiple data sources.

We propose a novel time-aware mechanism to capture temporal aspect of events and thus capturing better their relevance to the conversion.

2 BACKGROUND AND RELATED WORK

A brief overview of online advertising is given to stress the importance of predicting future conversions and to place it in the grander ecosystem. Additionally, relevant prior works on conversion prediction will be mentioned and their contributions discussed with respect to this study.

2.1 Online Advertising

Major DSP platforms for display advertising (e.g., Google DoubleClick, Verizon Media DSP) allow advertisers to sign up and run campaigns and lines. Every advertiser can create multiple campaigns and multiple lines within each campaign that target certain activity. Activities, for example, can be ad clicks or conversion activities (definition of which varies from one advertiser to another). The task for DSP platforms becomes to run advertiser's lines and serve users such that key performance indicators (KPIs) goals are reached. This is achieved by participating in online auctions for different ad opportunities. An important aspect of participating on

online auctions is deciding on the value of the ad opportunity as the maximum bid. For conversion predictions the maximum bid is often controlled by the probability of user converting not long after ad is displayed, or more precisely, the maximum bid is defined as a factored conversion probability:

$$\text{maximum_bid} = pCVR: \quad (1)$$

Thus, the estimate of conversion probability $pCVR$ is one of the key components in the DSP business that drives performance and directs the system towards displaying ads to relevant users. Similar relation can be used for click prediction lines.

2.2 Modeling users' conversion prediction

In large scale advertising setups, conversion probability estimation has been successfully tackled through logistic regression models [12]. However, manually designing and selecting features requires substantial investment of human time and effort, and utility of such generated features is largely dependent on the domain knowledge of human experts curating the features. Moreover, since typical applications are nonlinear, considering feature interactions (e.g. cross-features) quickly becomes prohibitively expensive due to a combinatorial explosion [11].

Recently, representation powerful deep learning models have also been proposed for CTR and CVR prediction, e.g., factorization machines [13] for CVR or deep residual networks [14] for CTR that tackle problems of learning non-linear interactions of features. Also, models that capture information from the sequence such as RNNs have been proposed recently [1, 5, 7, 17] and the reportedly perform significantly better than their non-sequential counterparts. Moreover [1] and [18] have used sequences of events from heterogeneous data sources, while [1] has additionally proposed adding temporal information of events as an additional source of information to better model sequence for conversion attribution task.

It is worthwhile noting that there are currently no notable papers describing the use case of prospective user conversion modeling.

3 METHODOLOGY

In this section we discuss proposed model and its interpretability.

3.1 Proposed Approach

We propose a novel model - Deep Time Aware conversion (DTAIN) model (Fig. 2).

The DTAIN model takes sequence of events $f_{e_j} = 1 :: N$ and time difference of events' timestamps and the time point of prediction (usually timestamp of last event in a sequence) as inputs. It then forwards this information through 5 blocks specifically designed for this task to learn conversion rate prediction.

3.1.1 Blocks of the DTAIN model.

Events and Temporal information embedding. Embeddings of events and temporal information are performed in two separate parts of the network. First, I_n events in the user's trail and I_n timestep information associated to them are embedded into vectors h_{e_i} of $h_{e_i} \in \mathbb{R}^{d_w=300}$ dimensional common space (*Embedding block*).

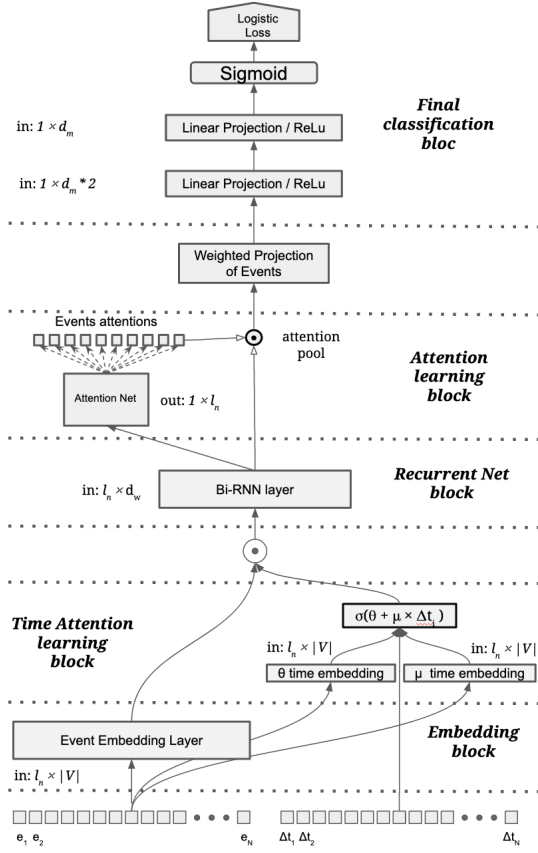


Figure 2: Graphical representation of the DTAIN model

Temporal attention learning. Each event e_i is also associated with two additional single-dimensional learnable parameters as μ_{e_i} and e_i of $\mu_{e_i}; e_i \in \mathbb{R}^{d_t=1}$. These parameters are designed to model the temporal increment Δ_t as time difference between current state i and the state of interest j (i.e. timestep when pCVR is served):

$$\Delta_t = e_j - e_i \quad (2)$$

$$e_j; \Delta_t^0 = S^1 e_i - \mu_{e_i} \Delta_t^0 \quad (3)$$

$$S^1 X^0 = \frac{1}{1 + e^{-X}} \quad (4)$$

$e_j; \Delta_t^0$ captures the influence of the current event to conversion with e_i measuring initial influence and μ_{e_i} measuring the change of the influence of the event with the time difference. Smaller $|\mu_{e_i}|$ refers to events whose influence does not change as we observe the event through different points in the users trails, while larger $|\mu_{e_i}|$ means that position and time of the event is very important for measuring its effect on conversion probability (given that the Δ_t is always positive and provided that e_i doesn't change, larger positive values of μ_{e_i} would mean that the temporal scores is closer to 0, and larger negative values that is closer to 1). Similar ways of modeling temporal increments can be seen in known results of Euler's forward method [4] for modeling change of state in dynamic linear systems. In our case, we opted for using time information as an event-level contribution to the final task, thus Sigmoid function

was used to transform $e_i - \mu_{e_i} \Delta_t$ into probability between 0 and 1. Rather than choosing Softmax layer which would force total influence of all events to be equal to one, we opted in to use Sigmoid to model influence of each event individually given their own specific influence factors. This approach allows us to model same events that happened multiple times withing same user trail differently, i.e. giving more attention to events that happened more recently.

Other formulations of time information were given in [1], however their approach only includes the cases of strict time decay effect where only events which happened close to prediction time may pass full information through the classifier. Similarly to [3] we learn event-specific initial and time influence factors which we use to control how much information passes from each event embedding into the first non-linear layer of the model.

The learned embeddings and contributions of each event are then summarized to obtain new event representation e_i :

$$\forall h_{e_i} \in \mathbb{R}^{d_f=1 \dots N} \forall \mu_{e_i, \Delta_t} \in \mathbb{R}^{d_f=1 \dots N} e_i = h_{e_i} - \mu_{e_i} \Delta_t^0 \quad (5)$$

Resulting again in $e_i \in \mathbb{R}^{d_w=300}$ dimensional space. This way of modeling allows for model interpretability, as for each event we can measure its initial and time influence factor and interpret their values as described above.

Recurrent Net block. The resulting embeddings of events are then fed into bi-directional RNN model (with GRU cells used for both forward and backward pass networks), our first non linear layer in the model:

$$e_1; e_2; \dots; e_N = biRNN^1 e_1; e_2; \dots; e_N; GRU^0 \quad (6)$$

Bi-directional RNN's ensure that the model learns complex relations between events, which is in particular important for user trails where events may be grouped by sessions which carry higher order information than the events themselves [9]. The resulting embeddings e_i are embedded into $e_i \in \mathbb{R}^{d_m=200}$ dimensional space.

Attention learning block. In order to learn rich representations of user's trail, it is imperative to focus on events that carry the most information. To learn representations that focus on important parts of the user trail we employ a dedicated attention mechanism on top of sequence modeling features [9]. Employed attention block yields event scores, that highlight events of greater importance for the task at hand. In our particular case, attention model is implemented as two-layered individual neural network $S_q^1 e_i; e_i^0$ with Softmax at it's final layer:

$$t_{e_i} = \frac{\exp^1 S_{e_i}^1 e_i; e_i^0}{\sum_{i=1}^{I_n} \exp^1 S_{e_i}^1 e_i; e_i^0} \quad (7)$$

Neural networks $S_{e_i}^1 e_i; e_i^0$ learns real valued scores for each i^{th} event in a given user trail. Attentions learning in the DTAIN model is coupled with the entire network (end-to-end).

Event attentions t_{e_i} are then used to re-weight their input representations e_i and to obtain compact representation of the entire sequence $s = \sum_i t_{e_i} e_i$. There are other ways of obtaining compact representations S , such as sum, average or max of individual event vectors. However, our experiments, as well as available literature [8, 16], demonstrate that such strategies are inferior to using attention.

Learning to predict from the resulting representation. The summarized user trail representation from previous block is finally fed to a sequence of fully connected layers with ReLU nonlinearities before finally passing through a sigmoid layer σ to obtain the probability of conversion (pCVR).

Finally, to optimize the parameters of DSM (denoted as W in remainder of the text), we have obtained logistic loss \mathcal{L} for the CTR prediction based on logits from the topmost layer:

$$\mathcal{L}(W) = -\frac{1}{N} \sum_{n=1}^N \left[y_n \log \sigma(\hat{z}_n) + (1 - y_n) \log (1 - \sigma(\hat{z}_n)) \right], \quad (8)$$

where \hat{z}_n are obtained logits after final sigmoid layers and y_n is conversion label for the n^{th} user trail.

Weights are initialized by a truncated normal initializer. To optimize \mathcal{L} , we use Adam [10] with a decaying gradient step.

4 DATA DESCRIPTION

RecSys 2015 challenge. We conducted conversion prediction experiments on publicly available dataset obtained from RecSys Challenge in 2015. This dataset contains a collection of sequences of click events with respective timesteps from Yoochoose website. Some of the click sessions ended with a purchase event (if so, label was set as positive, otherwise negative). This dataset reflects on reproducibility of retargeting results from this study only, as there is no publicly available prospecting dataset to the best of our knowledge.

User activity trails from Verizon Media. We also conducted experiments using user activity trails data from Verizon Media. This includes activities done in chronological order by a user, and the activities are derived from heterogeneous sources, e.g., Yahoo Search and Mail, reading news and other content on publisher's webpages associated with Yahoo, advertising data from Yahoo Gemini and Verizon Media DSP and data from all advertisers (e.g., ad clicks, conversions, and site visits). The representation of an activity comprises of activity ID, time stamp, its type (e.g., search, invoice, reservation, content view, order confirmation, parcel delivery), and a raw description of the activity (e.g., the exact search query for search activities) after stripping personally identifiable information. To ensure legality of information used, datasets created for each advertiser strictly follow legal guidelines determined by the contract, i.e. data collected from advertiser A will never be used for any optimization task for advertiser B.

Site visits are events that are most commonly labeled as retargeting events, i.e. user who is browsing to buy a furniture item on advertiser's website will in the next several months be regarded as a retargeting user for furniture conversions for that advertiser. As mentioned in the Introduction, advertisers who focus on prospective advertising are interested in generating new converters from non-retargeting users (users who did not visit advertiser's website), however, learning to target prospecting users from the existing data is very difficult. Namely, a common theme for a majority of *retail advertisers* is that a user will visit their webpage at least once before purchasing anything. Conversions in terms of whether a user visited advertisers' website for a single major retail advertiser are characterized in Table 1. Table 1 clearly shows that the vast

Conversion	Adv. site visit	Site visit prior to conv.	Percentage
TRUE	FALSE	FALSE	0.01%
TRUE	TRUE	FALSE	0.02%
TRUE	TRUE	TRUE	99.97%

Table 1: Percentages of conversions with respect whether advertisers' site visit (a retargeting) event occurred, and if it occurred before the conversion or not.

majority of conversions happen after users visited advertiser's website, thus becoming retargeting users before conversion. The goal of DSP prospective targeting is to target users before they become retargeting users, thus bringing new users to the advertiser and boosting their sales.

Any algorithm trained on original data collected will be biased towards modeling retargeting signals only as simple rule-based modes such as predicting conversion for all users that visit advertisers' website will yield very high recall (i.e. 99.7%). To prevent this from happening, we are performing a retargeting events black-listing, as highlighted by the advertiser. The process is shown in Fig. 3 and it reflects use cases where the algorithm learns to predict whether a user is going to convert the next day or not based on all signals, and not simply looking at site visits.




Figure 3: Visualisation of trail cutting process event before retargeting event happens.

Dataset used in this study is collected from a single anonymized major advertiser who defined two different conversion rules and it comprises (after eligible users and events are selected and negatives downsampling is performed to maintain roughly 10% of positives) of 788; 551 users in train and 196; 830 for test set collected over an undisclosed period longer than 90 days.

5 EXPERIMENTS

We first describe baseline algorithms that can capture information from the sequences of events as such models reportedly outperform standard structured models. Evaluation metrics are then defined. And finally results on both public and proprietary dataset are provided and discussed.

5.1 User Modeling Baselines

The following models are selected to either represent previously published studies or as models that are expected to fit well with the given setup.

- (1) Recurrent Neural Network (RNN): A recurrent neural network with embedding layer and GRU cells to ensure fast convergence.

- (2) 1-dimensional Convolutional Neural Networks (CNN): A 1-dimensional Convolutional neural network on top of learned event embeddings.
- (3) RNN with attention layer (RNN+Attn): An extension of RNN model with additional attention layer used to summarize the sequence [8].
- (4) RNN with self attention layer (RNN+SelfAttn): Another extension of RNN model with self-attention layer used learn higher order interactions between events before the RNN block [15].

5.1.1 Evaluation metrics. For assessing the quality of estimated CVR probabilities, we use the area under the ROC curve (AUC) classification performance measure, in addition to Accuracy, Precision and Recall obtained after choosing the appropriate classification threshold.

In addition, for the proprietary data, we study the bias [2] of the predicted probabilities defined as ratio between sum of sample ($S \geq S$) conversion probabilities $P^1 S^0 \geq 0::1\%$, and sum of conversion labels $1^1 S^0 \geq 0::1\%$ as $Bias = \frac{\sum_{s \geq S} P^1 S^0}{\sum_{s \geq S} 1^1 S^0}$. Unbiasedness (Bias=1) is a desirable property, as higher than 1 bias implies overly-optimistic estimates and waste of resources (bidding where there is a lower chance of conversion), and lower than 1 bias implies to overly-conservative estimates and missed opportunity (not bidding where there is a higher probability of conversion).

Finally, for the cases of multi-task learning where class disbalance becomes prominent we report area under Precision-Recall curve, as a more representative metric [6].

5.2 Experimental results

The proposed algorithm and baseline methods are evaluated on two described datasets and the results are given below.

5.2.1 Results on public dataset. Results of the experiments on public data source are given in Table 2. The ROC AUC and PRC AUC

	ROC AUC	PRC AUC	Accuracy	Precision	Recall
CNN	0.7534	0.2870	0.6779	0.2087	0.7041
GRU	0.7504	0.2725	0.6958	0.2142	0.6746
GRU+SelfAttn	0.7029	0.2391	0.6734	0.1907	0.6184
GRU+Attn	0.7639	0.2973	0.6997	0.2195	0.6904
DTAIN	0.7666	0.3019	0.6943	0.2186	0.7047

Table 2: Performance metrics on the Youchoose dataset for all algorithms.

results show that the proposed GRU+TimeAttn model outperforms all of the baselines. The PRC AUC was reported as the ratio of positives was approximately 9% in the dataset. Competitive results in the rest of the metrics show that the temporal information can truly help the predictive task even in datasets such as this one, especially Given that all examples in the public dataset occur within one hour time window. It may be surprising that adding temporal information helps, however, as discussed in the Section 3.1, the temporal information has two aspects to it and thus can model initial impact of the events to the conversion thus providing additional information to the classifier.

5.2.2 Results on proprietary dataset - prospecting users conversion prediction.

Results on binary classification. In this section we conduct experiment on prediction task whether user converted for any of the conversion rules set by the advertiser. Similar, yet more prominent results are obtained on the proprietary dataset where temporal aspect plays a major role in prediction (Table 3). We can see that the

	ROC AUC	Accuracy	Precision	Recall	Bias
CNN	0.8806	0.8110	0.2457	0.7871	1.0161
GRU	0.9018	0.8520	0.3004	0.7972	1.1983
GRU+Attn	0.8968	0.8438	0.2882	0.7982	0.8047
GRU+SelfAttn	0.8804	0.8364	0.2743	0.7756	0.9273
DTAIN	0.9263	0.8602	0.3219	0.8537	0.9871

Table 3: Performance metrics on the proprietary user trails dataset for all algorithms.

DTAIN outperforms other baselines by a large margin on all metrics. The time aspect of the events is much more prominent in the proprietary dataset. Moreover, as the time window is significantly larger, the events may repeat multiple times, and time mechanism will be able to select the most important events out of the redundant ones and thus filter out the noise in the data.

Results on multi-task classification. Finally we show results for the multi-task classification setup where we predict whether user will not convert for the advertiser, or for which of the two conversion rules will the user convert. As the positives and negatives rate

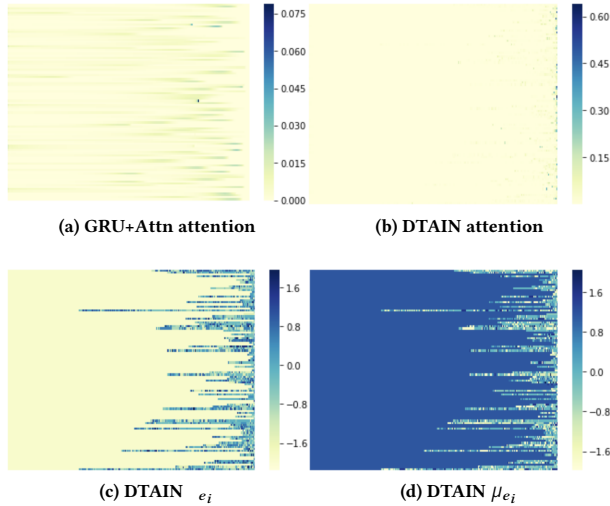
	PRC AUC	Accuracy	Precision	Recall	Bias
Task 0					
CNN	0.9880	0.8139	0.9810	0.8153	1.0069
GRU	0.9896	0.8544	0.9821	0.8588	1.0030
GRU+Attn	0.9907	0.8511	0.9837	0.8537	0.9933
GRU+SelfAttn	0.9877	0.8456	0.9795	0.8515	0.9941
DTAIN	0.9926	0.8613	0.9876	0.8614	0.9982
Task 1					
CNN	0.2523	0.9602	0.3161	0.2506	0.8836
GRU	0.2711	0.9629	0.3635	0.2720	0.9715
GRU+Attn	0.3013	0.9630	0.3788	0.3139	1.1163
GRU+SelfAttn	0.2452	0.9606	0.3277	0.2648	1.0645
DTAIN	0.2880	0.9652	0.4000	0.2539	1.0680
Task 2					
CNN	0.2495	0.9584	0.3287	0.2419	0.7588
GRU	0.2567	0.9597	0.3485	0.2464	0.8849
GRU+Attn	0.2374	0.9584	0.3355	0.2582	1.0453
GRU+SelfAttn	0.2081	0.9587	0.3081	0.1951	0.9887
DTAIN	0.2776	0.9633	0.4083	0.2348	0.9460

Table 4: Performance metrics on the proprietary user trails dataset for different tasks.

becomes very disbalanced when binary task is split to multi-task classification tasks, we report PRC AUC metric [6]. The DTAIN shows the best performance on majority of metrics across the three tasks always having the top Accuracy and Precision metrics. Overall evaluation shows that the DTAIN model is the best among the chosen baselines once again. The DTAIN model was prominently

the best approach for the Task 0 (prediction if the user is not going to convert) which is very important for the bidding system to know if it should bid for a user or not.

Figure 4: Heat maps of events attentions scores for 100 randomly sampled converters



5.2.3 Attention analysis and interpretation. To tap into the explainability of the models we randomly selected a hundred converters and analyzed attentions of their events. We compare DTAIN model primarily against the GRU+Attn model, which has shown properties of explainability in the past [8]. From Fig. 5a it can be seen that GRU+Attn model assigns attentions across the users trails, highlighting not only events that happened close to conversion which is a desirable property for prospective advertising. The DTAIN model has a slightly different mechanism of attention as time plays a major role in allowing information from different signals to be passed through the network. As discussed in Section 3 key parameters e_i and μ_{e_i} have interesting interpretability properties. To show this we plot scores of both the key parameters in Fig. 5c and 5d respectively, and of attentions from the attention block in Fig. 5b. Interestingly, we can see that there are plenty of high positive values of e_i and high negative values of μ_{e_i} further away from the end of sequences, in addition to the expected ones closer to the end of it. This means that the DTAIN is capturing long term as well as short term patterns and controls which events signals fully pass through the rest of the network. Moreover, interesting pattern shows when we observe the attention scores, which look very different from the ones from the GRU+Attn model as they are focused towards the end of the sequence. It is important to notice that less relevant events have been already filtered by the time-aware mechanism before being passed into the GRU layer and overall signals of sequences are then summarized in last few vectors in the GRU layer output, which was not possible in the other model.

These interesting findings allow us to use the attention scores for explainability to the advertisers by providing them insights into both long- and short-term patterns and important events that they can further use to improve their creatives and advertising strategies.

6 CONCLUSIONS AND FUTURE WORK

In this study we proposed a sequence based approach for modeling conversion prediction based on users' activity trails that leverages both the sequence and temporal information of heterogeneous events collected from many data sources. We proposed a new way to model temporal information for conversion prediction that preserves ability of interpretation, and finally we showed that the DTAIN mode outperforms baselines that represent state-of-the-art on both public and proprietary datasets. However, as the data is collected from many data sources, and different events may repeat often or periodically there it is still significant noise that the algorithms need to address and developing novel techniques to address these concerns will be the next steps in developing new solutions.

REFERENCES

- [1] S. K. Arava, C. Dong, Z. Yan, A. Pani, et al. Deep neural net with attention for multi-channel multi-touch attribution. *arXiv preprint arXiv:1809.02230*, 2018.
- [2] R. A. Baeza-Yates and B. Ribeiro-Neto. *Modern Information Retrieval*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1999.
- [3] T. Bai, S. Zhang, B. L. Egleston, and S. Vucetic. Interpretable representation learning for healthcare via capturing disease progression through time. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 43–51. ACM, 2018.
- [4] X. H. Cao, C. Han, and Z. Obradovic. Learning a dynamic-based representation for multivariate biomarker time series classifications. In *2018 IEEE International Conference on Healthcare Informatics (ICHI)*, pages 163–173. IEEE, 2018.
- [5] Y. Cui, R. Tobossi, and O. Vigouroux. Modelling customer online behaviours with neural networks: applications to conversion prediction and advertising retargeting. *arXiv preprint arXiv:1804.07669*, 2018.
- [6] J. Davis and M. Goadrich. The relationship between precision-recall and roc curves. In *Proceedings of the 23rd international conference on Machine learning*, pages 233–240. ACM, 2006.
- [7] D. Gligorijevic, J. Stojanovic, A. Raghuvveer, M. Grbovic, and Z. Obradovic. Modeling mobile user actions for purchase recommendations using deep memory networks. In *41st Int'l ACM SIGIR Conf. on Research and Development in Information Retrieval*, 2018.
- [8] D. Gligorijevic, J. Stojanovic, W. Satz, I. Stojkovic, K. Schreyer, D. Del Portal, and Z. Obradovic. Deep attention model for triage of emergency department patients. In *2018 SIAM International Conference on Data Mining (SDM 2018)*, 2018.
- [9] J. Gligorijevic, D. Gligorijevic, I. Stojkovic, X. Bai, A. Goyal, and Z. Obradovic. Deeply supervised model for click-through rate prediction in sponsored search. *Data Mining and Knowledge Discovery*, Apr 2019.
- [10] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [11] H. B. McMahan, G. Holt, D. Sculley, M. Young, D. Ebner, J. Grady, L. Nie, T. Phillips, E. Davydov, D. Golovin, et al. Ad click prediction: a view from the trenches. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1222–1230. ACM, 2013.
- [12] B. N., K. R., and M. S. A large scale prediction engine for app install clicks and conversions. In *International Conference on Information and Knowledge Management (CIKM)*, 2017.
- [13] J. Pan, Y. Mao, A. L. Ruiz, Y. Sun, and A. Flores. Predicting different types of conversions with multi-task learning in online advertising. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1834–1842. ACM, 2019.
- [14] Y. Shan, T. R. Hoens, J. Jiao, H. Wang, D. Yu, and J. Mao. Deep crossing: Web-scale modeling without manually crafted combinatorial features. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 255–262. ACM, 2016.
- [15] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.
- [16] S. Zhai, K.-h. Chang, R. Zhang, and Z. M. Zhang. Deepintent: Learning attentions for online advertising with recurrent neural networks. In *22nd ACM SIGKDD Intl. Conf. on Knowledge Discovery and Data Mining*, pages 1295–1304. ACM, 2016.
- [17] Y. Zhang, H. Dai, C. Xu, J. Feng, T. Wang, J. Bian, B. Wang, and T.-Y. Liu. Sequential click prediction for sponsored search with recurrent neural networks. In *Twenty-Eighth AAAI Conference on Artificial Intelligence*, 2014.
- [18] Y. Zhou, S. Mishra, J. Gligorijevic, T. Bhatia, and N. Bhamidipati. Understanding consumer journey using attention based recurrent neural networks. *KDD*, 2019.