Graph Regularized Encoder Training for Extreme Classification

Anshul Mittal anshulmittal@microsoft.com Microsoft Noida, India

> Siddarth Asokan Microsoft Research Bengaluru, India

> Pankaj Malhotra Microsoft Noida, India

Sumeet Agarwal IIT Delhi New Delhi, India Shikhar Mohan Microsoft Bengaluru, India

Suchith Chidananda Prabhu IIT Delhi New Delhi, India

Jian Jiao Microsoft Mountain View, United States

Soumen Chakrabarti IIT Bombay Mumbai, India

Manik Varma manik@microsoft.com Microsoft Research, IIT Delhi New Delhi, India Deepak Saini Microsoft Mountain View, United States

> Lakshya Kumar Microsoft Bengaluru, India

Amit Singh Microsoft Bengaluru, India

Purushottam Kar IIT Kanpur Kanpur, India

Abstract

Extreme Classification (XC) offers a scalable and efficient solution for retrieving highly relevant ads in Sponsored Search settings, significantly enhancing user engagement and ad performance. Most tasks in sponsored search involve highly skewed distributions over the data point (query) and label (ads) space with limited or no labeled training data. One approach to tackle this long-tail classification problem is to use additional data, often in the form of a graph such as similar queries, same session queries etc. that are associated with user queries/ads, called graph metadata. Graph-based approaches, particularly Graph Convolutional Networks (GCNs), have been successfully proposed to leverage this graph metadata and improve classification performance. However, for tail inputs/labels, GCNs induce graph connections that can be noisy, leading to downstream inaccuracies while also incurring significant computation and memory overheads. To address these limitations, we introduce a novel approach, RAMEN, that harnesses graph metadata as a regularizer while training a lightweight encoder rather than a compute- and memory- intensive GCN-based method. This avoids the inaccuracies incurred by noisy graph induction and sidesteps the computational costs of GCNs via an easy-to-train and deploy encoder. The proposed approach is a scalable and efficient solution that significantly outperforms GCN-based methods. Extensive A/B tests conducted on live multi-lingual Bing Ads search engine

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License. *WWW Companion '25, Sydney, NSW, Australia* © 2025 Copyright held by the owner/author(s). ACM ISBN 979-8-4007-1331-6/2025/04 https://doi.org/10.1145/3701716.3715230 traffic revealed that RAMEN increases revenue by **1.25-1.5**% and click-through rates by **0.5-0.6**% while improving quality of predictions across different markets. Additionally, evaluations on public benchmarks show that RAMEN achieves up to **5**% higher accuracy compared to state-of-the-art methods while being **50**% faster to infer, and having **70**% fewer parameters.

CCS Concepts

• Information systems \rightarrow Sponsored search advertising.

Keywords

Extreme classification, Large scale recommendation, Metadata, Sponsored search, ads, intelligent advertisement

ACM Reference Format:

Anshul Mittal, Shikhar Mohan, Deepak Saini, Siddarth Asokan, Suchith Chidananda Prabhu, Lakshya Kumar, Pankaj Malhotra, Jian Jiao, Amit Singh, Sumeet Agarwal, Soumen Chakrabarti, Purushottam Kar, and Manik Varma. 2025. Graph Regularized Encoder Training for Extreme Classification. In *Companion Proceedings of the ACM Web Conference 2025 (WWW Companion* '25), April 28-May 2, 2025, Sydney, NSW, Australia. ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3701716.3715230

1 Introduction

Sponsored search is a key revenue driver for search engines, displaying ads alongside organic results. The task in Sponsored Search is that of understanding the user query and recommending relevant ads. One of the popular approaches to predicting these recommendations is extreme classification. Extreme classification (XC) refers to a supervised machine learning paradigm wherein multi-label learning is performed on extremely large label spaces. The ability of XC to handle enormous label sets with millions of labels makes it WWW Companion '25, April 28-May 2, 2025, Sydney, NSW, Australia



Figure 1: A snapshot from LF-WikiSeeAlsoTitles-320K dataset for the article on "Cladistics." The related article "Common descent" is tagged but the ground truth is missing the link to "Crown group". Traversal on the hyperlink edges can help discover missing link but can also lead to irrelevant link such as "Vestigial organs".

an attractive choice for applications such as product recommendation [12, 30, 37, 42], search & advertisement [12, 24, 45], and query recommendation [6, 25]. The key appeal of XC comes because of two reasons: (a) The ability to recommend relevant ads to user queries even if queries are previously unseen (tail queries), and (b) The ability to accurately tag rare/tail ads relevant to a user query. An ad/query is considered to be part of the tail if very few training data points are associated with it. The tail problem is further aggravated due to the issue of missing data, since tail query/ads are also at a higher risk of going missing [24] in the ground truth. In solving the tail-data problem, XC approaches rely on metadata. Beyond textual ads descriptions [11, 13, 40], this auxiliary metadata can augment the limited supervision available for tail query/ads and is typically available in the form of multi-modal descriptions such as images [42], or graphs [41, 43, 47]. In this paper, we focus on graph metadata which is available in several applications, e.g. for online search and recommendation, users asking the multiple queries in the same search session (co-session queries) can be a metadata graph over queries.

Graph metadata in XC: Graph metadata has been used in XC to (a) handle tail ads [1, 30, 41], and (b) enhance the user query representations [10, 43, 47, 53]. To handle tail ads, these approaches use textual descriptions of an ad along with graph metadata to learn ad embeddings via graph convolutional networks (GCN). To enhance user query representation, these algorithms rely on a two stage retrieval pipeline wherein, for an unseen query, first the graph metadata nodes are retrieved, and subsequently, a GCN combines them with the query representation. The new combined representation is then used in a second stage to retrieve the relevant set of ads. The retrieval of graph nodes or graph traversal can also help discover missing ads associated with query. To showcase this, let us consider an example from a publicly available dataset from XC [3] called LF-WikiSeeAlsoTitles-320K. Here, task is similar to that of query-to-ads recommendation *i.e.* for a Wikipedia document title retrieve related Wikipedia page. For training, the ground truth comprises links to multiple Wikipedia page under the "See Also" section of the Wikipedia page. This ground truth is often incomplete and contains lots of missing links. Each Wikipedia page also has multiple hyperlinks to other Wikipedia page, which can be

Same Session Similar Query Query A Triplet Loss Triplet Loss Query Loss Triplet Loss ${\cal E}_{m heta}$ Triplet Loss Triplet Loss ${\cal E}_{ heta}$ A Same Session Similar Clicked Ad Ad

Figure 2: RAMEN uses graph metadata to regularize encoder during training. Training framework of RAMEN is robust to noise in graph. Here \mathcal{R}_x and \mathcal{R}_y are regularization loss terms and \mathcal{L} is the task based loss term.

used to construct the metadata graph. A snapshot of the dataset is provided in Figure 1, and shows how a rare Wikipedia page (which was previously unseen in the training ground truth) "Crown group" can be recovered for the Wikipedia article "Cladistics" by traversing the aforementioned graph. Similar to hyperlink graphs, in sponsored search, for queries/ads, we have graphs of (a) slight textual perturbations (similar queries/ads) but having similar search intent; (b) queries that were asked in the same session, or (c) ads that were clicked in the same search session. RAMEN utilize these metadata graphs to make faster and accurate predictions in comparison to GCNs like OAK [43].

Limitations of GCN Methods: In online scenarios, new users and new advertisers are continuously using the Bing search engine. The **new user queries and ads account for 75-80% for the search volume. Having graph metadata associated with each of the items is a challenging tasks. GCNs circumvent this challenge by inducing graphs for both queries and ads. These induced graphs are noisy and can lead to wrong predictions. To illustrate this, recall the LF-WikiSeeAlsoTitles-320K hyperlink graph - Traversal over the graph can also lead to irrelevant links such as "Vestigial organs" and extracting meaningful information from such noisy graphs is a challenge. It has been noted that having a high-quality graph can offer enhanced model accuracy in recommendation settings [10, 41, 47, 53]. However, such quality is rarely available for a new query or ad. In addition to this, the use of GCN architectures makes both training and inference 50% and 100% more expensive (Table 5) as compared to RAMEN which is a dual encoder model. Noisy graphs, along with expensive inference architecture, make GCNs inaccurate for tail and costly to deploy. Our primary research question is: How can we use graph metadata for accurate predictions on tail

Graph Regularized Encoder Training for Extreme Classification

WWW Companion '25, April 28-May 2, 2025, Sydney, NSW, Australia



Figure 3: RAMEN does not require additional information to compute an accurate representation of the test point. In OAK (an instance of a GCN), inference is a computationally expensive two-stage pipeline, where the test point is first embedded into the graph and then, the linked nodes are used to compute the final representation (induced graph). Since RAMEN uses the graph for regularization, and does not traverse the graph at inference time, it can be 2× faster, and 5% more accurate, than OAK.

queries and ads while saving computational cost compared to GCNs?

1.1 Our Contributions

To address the above question, we propose gRaph regulArized encoder training for extreME classificatioN (RAMEN). RAMEN is a framework to effectively utilize graph metadata during training, with minimal overheads while saving computational cost in model size and inference time (Table 5) as compared to GCNs. In particular RAMEN's contributions are as follows:

- **Graph-based regularization**: Using graph metadata to regularize the training of a dual encoder (Figure 2). This improves the model's ability to capture semantic relationships between queries/ads, even for tail queries/ads which resulted in 1.25-1.5% revenue gains in over 160+ countries and up to 4.8% reduction in brand mismatch rates (Table 2).
- **Robustness to graph noise**: Developing a mechanism to dynamically adjust the reliance on graph metadata during training based on the confidence in the graph links (Table 8). This helps mitigate the impact of noisy graphs (Table 4) in training model parameters.
- Efficient inference: RAMEN provides accurate retrievals for tail queries and avoids the costly step of inducing graph as used in GCNs like OAK (Figure 3). This reduces computation cost at inference time by 50%. (Table 5).

RAMEN scales to datasets with up to 360M ads and can offer up to 5% higher prediction accuracies over state-of-the-art methods including those that use graph metadata to train GCN.

2 Related work

Extreme classification (XC) is a key paradigm in several areas such as ranking and recommendation. The literature on XC methods is vast [2, 4, 12, 15, 16, 18, 24–27, 29, 32, 37, 39–41, 45, 47, 49–51, 55–57, 59]. Early XC methods used fixed (bag-of-words) [2, 24, 29, 45, 50, 56] or pre-trained [25] features and focused on learning only a classifier architecture. Recent advances have demonstrated significant gains by using task-specific features obtained from a variety of deep encoders such as bag-of-embeddings [12, 13], CNNs [32], LSTMs [57], and transformers [4, 27, 55, 59]. Training is scaled to

millions of labels and training points by performing encoder pretraining followed by classifier training [12]. A data point is trained only on its relevant labels (that are usually few in number) and a select few irrelevant labels deemed most informative (known as *hard negatives*) and obtained using a process known as *negative mining* [8, 11, 13–15, 20, 22, 28, 31, 36, 38, 46, 52].

Label Metadata in XC: Most XC methods use textual representations as label metadata since they facilitate scalable training and inference and allow for leveraging good-quality pre-trained deep encoders such as RoBERTa [34], DistilBERT base [48], etc. Examples include encoder-only models such as DEXML [17], Twin-BERT [35] and ANCE [52], and encoder+classifier architectures such as DECAF [40], SiameseXML [11], X-Transformer [5], XR-Transformer [6], LightXML [27], and ELIAS [59], amongst many others [4, 33, 55, 57]. There is far fewer works in the literature on the use of other forms of label metadata. For instance, ECLARE [41] and GalaxC [47] use graph convolutional networks whereas MUFIN [42] explores multi-modal label metadata in the form of textual and visual descriptors for labels.

Graph Neural Networks in Related Areas: A sizeable body of work exists on using graph neural networks such as graph convolutional networks (GCN) for recommendation [7, 9, 19, 21, 23, 43, 53, 54, 58, 60, 61]. Certain methods, e.g., FastGCN [7], KGCL [54], and LightGCN [21] learn item embeddings as (functions of) free vectors. This makes them unsuitable for making predictions on a novel test point. Other GCN-based methods such as OAK [43], PINA [10], GraphSAGE [19] and GraphFormers [53] learn node representations as functions of node metadata e.g. textual descriptions. Consequently, these methods work in zero-shot settings, but they still incur the high storage and computational cost of GCNs. Moreover, diminishing returns are observed with increasing the number of layers of the GCN [9, 41] with at least one model, namely LightGCN [21] foregoing all non-linearities in its network, effectively opting for a single-layer GCN. It must be noted that GCNs can be highly accurate if one has access to an oracle for predicting relevant nodes (Table 4). However, such oracle is never available online and the slightest error in first stage retrieval leads to poor retrieval quality (Sec. 4).

We now develop the RAMEN method that offers a far more scalable alternative to GCNs and other popular graph-based architectures in XC settings, while significantly reducing the overheads WWW Companion '25, April 28-May 2, 2025, Sydney, NSW, Australia

of graph-based learning, and offering sustained and significant performance boosts in prediction accuracies.

3 RAMEN: <u>gRaph</u> regul<u>A</u>rized encoder training for extreME classificatioN

Notation: Let *L* be the number of labels in the recommendation task. Let $\mathbf{x}_i, \mathbf{z}_l$ be the textual descriptions of the queries or data point *i* and ads or label *l* respectively. From this point, we will be using query as data point and ads as labels interchangeably as per context need. For each data point $i \in [N]$, its ground-truth label vector is $\mathbf{y}_i \in \{-1, +1\}^L$, where $\mathbf{y}_{il} = +1$ if label *l* is relevant to the data point *i* and otherwise $\mathbf{y}_{il} = -1$. The training set is comprised of *N* labeled data points and *L* labels as $\mathcal{D} := \{\{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^N, \{\mathbf{z}_l\}_{l=1}^L\}$. Let $\mathcal{X} \stackrel{\text{def}}{=} \{\mathbf{x}_i\}_{i=1}^N$ denote the set of training data points and $\mathcal{Z} \stackrel{\text{def}}{=} \{\mathbf{z}_l\}_{l=1}^L$ denote the set of labels. The metadata graph over the anchor sets \mathcal{A} (**hyper-links, co-bidded queries) is denoted by \mathcal{G}_{XA} and \mathcal{G}_{ZA} for data point (query) and label (ads), respectively.

Metadata Graphs: In recommendation scenario, RAMEN obtains metadata graphs over queries and ads using user session or textual similarity as described in introduction section. These graph is essentially links between queries/ads to their relevant node in the graph. These are also called *Anchor Sets*. Refer to Figure 1, here hyper-linked Wikipedia pages are called anchor set. Let $\mathcal{A} = \{a_1, a_2, \ldots, a_M\}$ denote an anchor set of M elements *e.g.* pages connected via hyperlink for LF-WikiSeeAlsoTitles-320K dataset. We abuse notation to let a_m denote the textual representation of anchor item $m \in [M]$ as well. Each query and ad is associated with metadata graph over anchor sets:

- Query metadata graph: This is denoted as *G_{XA}* = (*V_{XA}*, *E_{XA}*) with *V_{XA}* ^{def} *X* ∪ *A* i.e., the union of training data points and anchor points. The matrix *E_{XA}* = {*e_{im}*} ∈ {0, 1}^{*N*×*M*} encodes whether data point **x**_{*i*} has an edge to anchor item **a**_{*m*} or not.
 Ads metadata graph: This is denoted as *G_{ZA}* = (*V_{ZA}*, *E_{ZA}*) with
 - $V_{ZA} \stackrel{\text{def}}{=} \mathbb{Z} \cup \mathcal{A}$ i.e. the union of labels and anchor points. The matrix $E_{ZA} = \{e_{lm}\} \in \{0, 1\}^{L \times M}$ encodes whether label \mathbf{z}_l has an edge to anchor item \mathbf{a}_m or not.

3.1 Regularized Training Framework

RAMEN incorporates graph based regularization while training model parameters. RAMEN training requires two components: (a) A base XC (\mathcal{M}) component which consists of an encoder block (\mathcal{E}_{θ}), and (b) The metadata graph (\mathcal{A}_c) component as described above. The encoder $\mathcal{E}_{\theta} : \mathcal{X} \to \mathcal{S}^{D-1}$ with trainable parameters θ is used to embed query and ads using their textual descriptions. Here, \mathcal{S}^{D-1} denotes the *D*-dimensional unit sphere, *i.e.*, the encoder provides unit-norm embeddings (unless stated otherwise). RAMEN uses a DistilBERT [48] encoder as \mathcal{E}_{θ} . In the following sections, we first explain training framework of RAMEN followed by incorporating graph regularization for robustness.

RAMEN loss function: Figure 2 shows overall training framework for RAMEN. The encoder is trained using triplet loss $(\mathcal{L}(\boldsymbol{\theta}))$ over query and ad representation, regularized using two components: a) The anchor set on the query side $(\mathcal{R}_x(\boldsymbol{\theta}))$, and b) The anchor set on ad side $(\mathcal{R}_z(\boldsymbol{\theta}))$, as explained in metadata graph section. The

 $\mathcal{L}(\boldsymbol{\theta})$ function is then given by:

$$\mathcal{L}(\boldsymbol{\theta}) = \sum_{i=1}^{N} \sum_{\substack{l:y_{il}=+1\\k:y_{ik}=-1}} [\mathcal{E}_{\boldsymbol{\theta}}(\mathbf{z}_k)^{\top} \mathcal{E}_{\boldsymbol{\theta}}(\mathbf{x}_i) - \mathcal{E}_{\boldsymbol{\theta}}(\mathbf{z}_l)^{\top} \mathcal{E}_{\boldsymbol{\theta}}(\mathbf{x}_i) + \gamma]_{+}$$

Note that this loss function encourages the encoder to embed a query close to its relevant ad and far from irrelevant ones. The optimized encoder (θ^*) through RAMEN framework is obtained by minimizing the following objective

$$\boldsymbol{\theta}^* = \min_{\boldsymbol{\theta}} \left\{ \mathcal{L}(\boldsymbol{\theta}) + \sum_{t=1}^T \left(\lambda_x \mathcal{R}_x^t(\boldsymbol{\theta}) + \lambda_z \mathcal{R}_z^t(\boldsymbol{\theta}) \right) \right\},\$$

where λ_x and λ_z are regularization constants. $\mathcal{R}_x^t(\boldsymbol{\theta})$ and $\mathcal{R}_z^t(\boldsymbol{\theta})$ are regularization loss applied over anchor set \mathcal{A}_t as explained below. **Metadata Graph Regularizers**: Given an encoder $\mathcal{E}_{\boldsymbol{\theta}}$, an anchor set \mathcal{A} , and graphs G_{XA} , G_{ZA} , we define the following two regularization functions over the encoder parameters:

$$\begin{aligned} \mathcal{R}_{x}(\boldsymbol{\theta}) &= \sum_{i=1}^{N} \sum_{\substack{p:e_{ip}=1\\n:e_{in}=0}} [\mathcal{E}_{\boldsymbol{\theta}}(\mathbf{x}_{i})^{\top} \mathcal{E}_{\boldsymbol{\theta}}(\mathbf{a}_{n}) - \mathcal{E}_{\boldsymbol{\theta}}(\mathbf{x}_{i})^{\top} \mathcal{E}_{\boldsymbol{\theta}}(\mathbf{a}_{p}) + \gamma]_{+} \\ \mathcal{R}_{z}(\boldsymbol{\theta}) &= \sum_{l=1}^{L} \sum_{\substack{p:e_{lp}=1\\n:e_{ln}=0}} [\mathcal{E}_{\boldsymbol{\theta}}(\mathbf{z}_{l})^{\top} \mathcal{E}_{\boldsymbol{\theta}}(\mathbf{a}_{n}) - \mathcal{E}_{\boldsymbol{\theta}}(\mathbf{z}_{l})^{\top} \mathcal{E}_{\boldsymbol{\theta}}(\mathbf{a}_{p}) + \gamma]_{+} \end{aligned}$$

Here, *p* is the positive anchor and *n* are in-batch negatives anchors (explained in next section). Note that these two regularizers encourage the encoder to keep data points and labels closely embedded to their related anchor points and far away from unrelated anchor points. If we have more than one anchor set, say $\mathcal{A}^1, \mathcal{A}^2$, we can define corresponding regularizers $\mathcal{R}_{x}^{t}(\boldsymbol{\theta}), \mathcal{R}_{z}^{t}(\boldsymbol{\theta}), t=1, 2$. Note that, while we are explaining using two anchor sets, RAMEN can be easily extended to multiple anchor sets without the loss of generality. Training RAMEN: RAMEN utilizes in-batch negative mining [8, 11, 13-15, 20](Figure 2). Specifically, a mini-batch is created using uniformly chosen set of queries and for each query, a relevant ad and a related anchor from query-anchor graph is chosen randomly. Similarly, for each of the chosen ad, a related anchor from ad-anchor graph is chosen randomly. Then, hard negative ads for a query are chosen only amongst those ads present in that mini-batch. Similarly, hard negative anchors for query/ad are chosen from only those anchors present in that mini-batch. Without the loss of generality training of RAMEN can be scaled to any number of anchor sets. RAMEN uses dual encoder models like DistilBERT [48] and uses graph regularized training framework as described above to learn robust query/ad representation.

Inference with RAMEN: Inference of RAMEN is same as that of any dual encoder model. During inference, query/ads representation is computed using the trained encoder and top-*k* relevant ads are retrieved based on cosine similarity. Encoder trained through proposed graph regularization sees the gains in accuracy due to additional metadata graph without computational cost of GCNs.

RAMEN introduces a novel approach to incorporate graph metadata into XC, addressing the limitations of GCN-based methods. Unlike GCNs, which rely on computationally expensive graph induction, RAMEN employs a more efficient dual encoder architecture (Figure 3). This architecture, coupled with a graph-based regularization, empowers the model to learn robust representations of queries and ads (Figure 2). By focusing on relevant anchor points

Anshul Mittal et al

for regularization during training, RAMEN eliminates the need to induce graph at inference time, eliminating the impact of noise at inference time. Additionally, the proposed method demonstrates enhanced ability to handle new entities as graph-based regularization improves the representation of tail queries/ads. These combination of factors results in improved accuracy and efficiency, making RAMEN favorable for real-world recommendation tasks.

4 Experiments

In this section, we compare RAMEN against deployed algorithms in Bing Ads, as well as established baselines on public benchmark datasets on the XML Repository [3]. In particular, this paper benchmarks RAMEN on the LF-WikiSeeAlsoTitles-320K dataset where the recommendation task is similar to Bing Ads. The dataset is curated from Wiki dumps (link). The scenario involves recommending related articles. Articles under the "See Also" section were used as ground-truth labels. Internal hyperlinks and category links were used to create two sets of metadata graphs - one using hyperlinked Wikipedia articles as anchors and the other, using Wikipedia categories as anchors. On the public dataset, all queries in the test set are unseen, making this a cold start problem, while on the Bing data, user queries could either belong to the head/tail, or be unseen (cold start). In online A/B tests, RAMEN shows significant improvement in revenue, click through rates (CTR) and impression yield (IY). In addition to that, RAMEN observes significant reduction in bad match rate (BMR) with respect to brands (Brand BMR) and locations (Location BMR). This is crucial for advertisers who want to ensure that their ads are shown only for relevant queries, based on brand terms and location context. Please refer to Tab. 9 of the appendix for dataset statistics. For details on evaluation metrics please refer to appendix C.

Implementation details: For public datasets we initialize the encoder with a pre-trained DistilBERT and fine-tune it. During training, we prune the metadata graph using the fine-tuned encoder to eliminated noise in the training graph. Table 10 in the appendix summarizes all hyper-parameters used for each dataset. We reiterate that, even though RAMEN uses a graph at training time, inference does not require any such information, making it highly suitable for long-tail queries and ads. We compare variants of RAMEN against baseline XC and dense retrieval approaches. In particular, we consider RAMEN (ANCE) and RAMEN (NGAME) where ANCE and NGAME where used as base XC model, on top of which we add RAMEN's regularization terms during training. All RAMEN variants and most baseline variants use the PyTorch [44] framework and were trained on 4 Nvidia V100 GPUs. DEXML [16] was trained on 16 Nvidia A100 GPUs. Refer to Appendix A for additional details.

Using this section we answer the following questions:

- What is the impact of graph regularization on accuracy and quality for real world applications?
- What is the computational cost of training RAMEN over GCN based approach (OAK)?
- What key design choices in RAMEN which lead to its success?

Case-study for Sponsored Search: Matching user queries with relevant advertiser. For example, for the query "cheap nike shoes",

a valid ad is "nike sneakers" but "adidas shoes" or "nike shorts" are not.

We study the effectiveness of RAMEN in this application by comparing it against the state-of-the-art encoder in production by conducting A/B test on live search-engine traffic. The click logs were mined to gather graph metadata for RAMEN, including two types of signals:

- Co-session queries/Ads: Queries that were asked in the same search session by multiple users as well as Ads that were clicked together in same session.
- (2) Similar queries/ Ads: Queries/Ads that are slight perturbations of each other but have the similar search intent.

RAMEN was trained on a dataset containing 540M queries and 360M ads mined as described above but over a longer period to conduct an A/B test on the search engine. RAMEN was found to increase the Impression-Yield (relevant ad impressions per user query) by 0.70% and the CTR (clicks through rates) by 0.86% when compared against a control ensemble containing state-of-the-art embedding-based, generative, GCN, and XC algorithms Table 1. Gains in RAMEN are primarily from rare query deciles, as shown in Table 1, which is the intended decile for RAMEN. In addition to gains in revenue metrics, RAMEN leads to significant reduction in bad match rates in terms of brands and location by -4.8% and -7.2% respectively (Table 2). The A/B test was conducted on 160+ countries and Table 3 shows that RAMEN leads to gains in 1.25-1.44% increase in revenue, 0.3-0.66% increase in CTR and 0.28-0.9% increase in IY.

Results on benchmark datasets: Table 6 compares RAMEN variants with graph and XC methods on standard XC metrics like Propensity scored precision and nDCG. RAMEN is 5 points more accurate over the best baseline numbers. In particular RAMEN is 2-3 points more accurate than traditional graph-based methods. Additionally, the RAMEN variants are 3-4 points more accurate over OAK [43] & PINA [10], which use both XC and graph metadata.

Analysis of gains and computation cost: These experiments where conducted on public benchmark datasets for reproducibility and access to broader audience. Recall that the GCN (OAK) twostage retrieval pipeline can be noisy. Table 4 demonstrates that, if we replace the first stage with the oracle linker (graph induction with zero error), the performance of these graph-based methods starts to outperform RAMEN variants. However, *the oracle linker is never available for a novel test point*, and RAMEN variants achieved a similar performance in a fraction of the cost of training and prediction time as shown in Table 5. In addition to reduction in inference time, RAMEN requires $3.5 \times$ less model parameters compared to OAK, as can be seen in Table 5, while being 2 absolute points more accurate in Precision@1. Table 7 shows that RAMEN (ANCE) could make tail label predictions such as "Crown group" which was a missing label ground truth in the training data.

Ablations on design choices: To understand the impact of noisy edges in metadata, experiment "– *No Pruning*" disabled the trimming of noisy edges using cosine similarity filtering. A 4% loss in P@1 was observed which underscores the necessity of pruning unhelpful edges during training. RAMEN (ANCE) uses multiple metadata graphs for both query and ads. To ascertain the contributions of the query-anchor and ads-anchor metadata graphs, the ablations "– *No Doc. Graph*" and "– *No Lbl. Graph*" were conducted.

Table 1: Query decile wise-comparison of RAMEN versus an ensemble of deployed Bing Ads algorithms (a.k.a. Control)

Decile	$ \%\Delta$ Revenue	$\%\Delta \ CTR$	$\%\Delta$ IY
HEAD	0.42%	0.22%	0.27%
TORSO	0.41%	0.30%	0.19%
TAIL	0.52%	0.34%	0.24%
TOTAL	1.35%	0.86%	0.70%

Table 2: RAMEN leads to significant reduction in brand as well as location bad match rate in comparison to Bing Ads control over 160+ countries.

$\%\Delta$ Brand BMR	$\%\Delta$ Location BMR
-4.8	-7.2

Table 3: Results of RAMEN in comparison to Bing Ads control in over 160+ countries grouped in 4 Clusters by geography

Market Cluster	$\%\Delta$ Revenue	$\%\Delta \ CTR$	$\%\Delta$ IY
MC-1	1.25	0.31	0.28
MC-2	1.44	0.64	0.19
MC-3	1.46	0.66	0.92
MC-4	1.27	0.63	0.58

Table 4: Results using Oracle Linker for GCN Vs RAMEN (ANCE) on LF-WikiSeeAlsoTitles-320K.

Method	P@1	P@5	N@5	PSP@1	PSP@5
RAMEN (ANCE)	35.2	18.4	35.2	29.0	34.5
OAK	33.7	17.1	34.4	25.8	30.8
OAK + Oracle	38.9	19.4	40.4	29.7	34.8

These experiments reveal that information from these graphs plays a significant role, as disabling either leads to a 1.5-2% reduction in P@1. The information from these graphs can be incorporated in baseline methods like ANCE. To understand its impact, experiment "AugGT" trains ANCE with augmented ground truth. The ground truth was expanded by using label propagation wherein a label and a training point are linked by an edge if the label shares a neighbor in the metadata graph of the said training point. RAMEN (ANCE) outperformed the "AugGT" setup by 15%. This suggests that while leveraging graph information for ground truth enhancement is convenient, it may not be as effective due to noisy edges.

5 Conclusion

This paper presents RAMEN, a novel approach to incorporating graph metadata into extreme classification tasks. RAMEN addresses the limitations of GCN-based methods by leveraging a more efficient dual encoder architecture and a robust graph regularization technique. The experimental results demonstrate the effectiveness of RAMEN in enhancing accuracy and computational efficiency particularly for tail queries and ads.

By comparing RAMEN with state-of-the-art methods on public benchmarks and real-world datasets like Bing Ads, we have shown

Table 5: RAMEN (ANCE)'s computation relative to baselines on LF-WikiSeeAlsoTitles-320K.

Method	Train Time	Pred Time	Model size	P@1
RAMEN (ANCE)	1×	$1 \times$	$1 \times$	35.2
OAK	1.5×	$2 \times$	$3.5 \times$	33.7
ANCE	0.9×	$1 \times$	$1 \times$	30.8
DEXML	2.1×	1×	1×	29.9

Table 6: Results comparing RAMEN against XC and Graphbased baselines on the short-text *LF-WikiSeeAlsoTitles-320K* dataset. RAMEN variants is up to 15% more accurate as compared to both text-based and graph-based baselines. Here PSP, P and N refers to propensity score precision, precision and nDCG respectively.

	PSP@1	PSP@5	P@1	P@5	N@5
RAMEN (ANCE)	28.9	34.5	35.2	18.4	36.5
RAMEN (NGAME)	28.6	34.4	35.46	18.5	36.8
OAK	25.8	30.8	33.7	17.1	34.3
GraphSage	21.5	23.5	27.3	12.9	28.3
GraphFormer	19.2	22.7	21.9	11.8	24.1
ANCE	25.1	28.7	30.7	15.3	31.4
NGAME	24.4	29.8	32.6	16.6	33.2
DEXA	24.4	28.6	31.7	15.8	32.2
DEXML	22.8	25.6	29.9	14.7	30.7
DECAF	16.7	21.1	25.1	12.8	25.9
Parabel	9.2	11.8	17.6	8.5	17.4
CascadeXML	12.6	17.6	23.4	12.1	23.4
XR-Transformer	10.1	12.7	19.4	8.9	18.5
AttentionXML	9.4	11.7	17.5	8.5	17.1
Bonsai	10.6	13.7	19.3	9.5	19.2
SiameseXML	26.8	30.3	31.9	16.2	32.5
ECLARE	22.1	26.2	29.3	15.1	30.2
ELIAS	13.4	17.6	23.3	11.8	23.6

Table 7: A subjective comparison of predictions made by RAMEN, the leading text-based method NGAME, and the leading graph-based method GraphFormers on LF-WikiSeeAlsoTitles-320K. Labels that are a part of the ground truth are formatted in black color. Labels not a part of the ground truth are formatted in light gray color. Relevant labels that are missing from the ground truth are marked in bold black. RAMEN (ANCE) could make predict highly relevant labels such as "Crown group", which were missing from the ground truth as well as omitted by other methods.

Method	Prediction
	Document: Clade
RAMEN (ANCE)	Cladistics, Phylogenetics, Crown group , Para- phyly, Polyphyly
ANCE	Cladistics, Linnaean taxonomy, Polyphyly, Paragroup, Molecular phylogenetics
OAK	Phylogenetic nomenclature, Molecular phylo- genetics, Haplotype, Cladistics, Paragroup

Table 8: Ablations were done using ANCE as base algorithm (RAMEN (ANCE)) on LF-WikiSeeAlsoTitles-320K to understand the impact of design choices on the quality of encoder training. RAMEN (ANCE)'s design choices are seen to be optimal and offer 2.5–13% improvement in the P@1 metric over alternate design choices.

RAMEN	P@1	P@3	P@5	N@3	N@5
RAMEN (ANCE)	35.2	24.1	18.3	35.3	36.5
– No Pruning	31.3	18.9	12.8	31.43	31.5
– No Doc. Graph	29.7	17.5	12.5	30.7	30.8
– No Lbl. Graph	34.1	22.7	14.6	32.0	34.1
AugGT	15.6	8.9	6.5	15.7	16.3

that it consistently outperforms existing approaches. The proposed graph regularization technique effectively mitigates the impact of noise in the graph data, leading to improved performance. In particular RAMEN can yield gains of up to 1-1.5% in Revenue as well as 0.86% and 0.70% in CTR and Impression Yield, respectively over Bing Ads control. RAMEN leads to enhanced user experience as well as satisfaction by significantly reducing bad match rates in brand and location sensitive queries by 4.8% and 7.2% respectively. Additionally, RAMEN's efficient architecture enables it to scale well to large-scale datasets. Performance gains in RAMEN over OAK (GCN based method) also comes with 70% reduced model parameters, 50% reduced inference time.

Future research directions include exploring the application of RAMEN to other domains, investigating further enhancements to the model architecture, and developing more efficient training and inference techniques.

References

- M. Artetxe, S. Bhosale, N. Goyal, T. Mihaylov, M. Ott, S. Shleifer, X.-V. Lin, J. Du, S. Iyer, R. Pasunuru, et al. Efficient large scale language modeling with mixtures of experts. arXiv preprint arXiv:2112.10684, 2021.
- [2] R. Babbar and B. Schölkopf. DiSMEC: Distributed Sparse Machines for Extreme Multi-label Classification. In WSDM, 2017.
- [3] K. Bhatia, K. Dahiya, H. Jain, A. Mittal, Y. Prabhu, and M. Varma. The Extreme Classification Repository: Multi-label Datasets & Code, 2016. URL http://manikvarma.org/downloads/XC/XMLRepository.html.
- [4] I. Chalkidis, M. Fergadiotis, P. Malakasiotis, N. Aletras, and I. Androutsopoulos. Extreme Multi-Label Legal Text Classification: A case study in EU Legislation. In ACL, 2019.
- [5] C. W. Chang, H. F. Yu, K. Zhong, Y. Yang, and I. S. Dhillon. A Modular Deep Learning Approach for Extreme Multi-label Text Classification. CoRR, 2019.
- [6] W.-C. Chang, Yu H.-F., K. Zhong, Y. Yang, and I.-S. Dhillon. Taming Pretrained Transformers for Extreme Multi-label Text Classification. In KDD, 2020.
- [7] J. Chen, T. Ma, and C. Xiao. FastGCN: Fast Learning with Graph Convolutional Networks via Importance Sampling. In *ICLR*, 2018.
- [8] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020.
- [9] W. Chiang, X. Liu, S. Si, Y. Li, S. Bengio, and C. Hsieh. Cluster-GCN: An Efficient Algorithm for Training Deep and Large Graph Convolutional Networks. In KDD, 2019.
- [10] E. Chien, J. Zhang, C. Hsieh, J. Jiang, W. Chang, O. Milenkovic, and H. Yu. PINA: Leveraging side information in eXtreme multi-label classification via predicted instance neighborhood aggregation. In *ICML*, 2023.
- [11] K. Dahiya, A. Agarwal, D. Saini, K. Gururaj, J. Jiao, A. Singh, S. Agarwal, P. Kar, and M. Varma. SiameseXML: Siamese Networks meet Extreme Classifiers with 100M Labels. In *ICML*, 2021.
- [12] K. Dahiya, D. Saini, A. Mittal, A. Shaw, K. Dave, A. Soni, H. Jain, S. Agarwal, and M. Varma. DeepXML: A Deep Extreme Multi-Label Learning Framework Applied to Short Text Documents. In WSDM, 2021.
- [13] K. Dahiya, N. Gupta, D. Saini, A. Soni, Y. Wang, K. Dave, J. Jiao, K. Gururaj, P. Dey, A. Singh, D. Hada, V. Jain, B. Paliwal, A. Mittal, S. Mehta, R. Ramjee, S. Agarwal,

P. Kar, and M. Varma. Ngame: Negative mining-aware mini-batching for extreme classification. In *WSDM*, March 2023.

- [14] F. Faghri, D.-J. Fleet, J.-R. Kiros, and S. Fidler. VSE++: Improving Visual-Semantic Embeddings with Hard Negatives. In *BMVC*, 2018.
- [15] C. Guo, A. Mousavi, X. Wu, D.-N. Holtmann-Rice, S. Kale, S. Reddi, and S. Kumar. Breaking the Glass Ceiling for Embedding-Based Classifiers for Large Output Spaces. In *NeurIPS*, 2019.
- [16] N. Gupta, D. Khatri, A. S Rawat, S. Bhojanapalli, P. Jain, and I. S Dhillon. Efficacy of dual-encoders for extreme multi-label classification. In *ICLR*, 2023.
- [17] N. Gupta, F. Devvrit, A. S. Rawat, S. Bhojanapalli, P. Jain, and I. S. Dhillon. Dualencoders for extreme multi-label classification. In *ICLR*, 2024.
- [18] V. Gupta, R. Wadbude, N. Natarajan, H. Karnick, P. Jain, and P. Rai. Distributional Semantics Meets Multi-Label Learning. In AAAI, 2019.
- [19] W. L. Hamilton, R. Ying, and J. Leskovec. Inductive Representation Learning on Large Graphs, 2018.
- [20] K. He, Haoqi Fan, Yuxin W., S. Xie, and R. Girshick. Momentum contrast for unsupervised visual representation learning. In CVPR, 2020.
- [21] X. He, K. Deng, X. Wang, Y. Li, Y. Zhang, and M. Wang. Lightgcn: Simplifying and powering graph convolution network for recommendation. In *SIGIR*, pp. 639–648, 2020.
- [22] S. Hofstätter, S.-C. Lin, J.-H. Yang, J. Lin, and A. Hanbury. Efficiently Teaching an Effective Dense Retriever with Balanced Topic Aware Sampling. In SIGIR, 2021.
- [23] W. Huang, T. Zhang, Y. Rong, and J. Huang. Adaptive Sampling Towards Fast Graph Representation Learning, 2018.
- [24] H. Jain, Y. Prabhu, and M. Varma. Extreme Multi-label Loss Functions for Recommendation, Tagging, Ranking and Other Missing Label Applications. In KDD, August 2016.
- [25] H. Jain, V. Balasubramanian, B. Chunduri, and M. Varma. Slice: Scalable Linear Extreme Classifiers trained on 100 Million Labels for Related Searches. In WSDM, 2019.
- [26] K. Jasinska, K. Dembczynski, R. Busa-Fekete, K. Pfannschmidt, T. Klerx, and E. Hullermeier. Extreme F-measure Maximization using Sparse Probability Estimates. In *ICML*, 2016.
- [27] T. Jiang, D. Wang, L. Sun, H. Yang, Z. Zhao, and F. Zhuang. LightXML: Transformer with Dynamic Negative Sampling for High-Performance Extreme Multilabel Text Classification. In AAAI, 2021.
- [28] V. Karpukhin, B. Oguz, S. Min, P. Lewis, L. Wu, S. Edunov, D. Chen, and W.-t. Yih. Dense passage retrieval for open-domain question answering. In *EMNLP*, 2020.
- [29] S. Khandagale, H. Xiao, and R. Babbar. Bonsai: diverse and shallow trees for extreme multi-label classification. ML, 2020.
- [30] S. Kharbanda, A. Banerjee, E. Schultheis, and R. Babbar. Cascadexml: Rethinking transformers for end-to-end multi-resolution training in extreme multi-label classification. In *NeurIPS*, 2022.
- [31] K. Lee, M.-W. Chang, and K. Toutanova. Latent retrieval for weakly supervised open domain question answering. In ACL, 2019.
- [32] J. Liu, W. Chang, Y. Wu, and Y. Yang. Deep Learning for Extreme Multi-label Text Classification. In SIGIR, 2017.
- [33] X. Liu, P. He, W. Chen, and J. Gao. Multi-Task Deep Neural Networks for Natural Language Understanding. In ACL, 2019.
- [34] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov. Roberta: A robustly optimized bert pretraining approach. arXiv preprint arXiv:1907.11692, 2019.
- [35] W. Lu, J. Jiao, and R. Zhang. TwinBERT: Distilling Knowledge to Twin-Structured Compressed BERT Models for Large-Scale Retrieval. In CIKM, 2020.
- [36] Y. Luan, J. Eisenstein, K. Toutanova, and M. Collins. Sparse, Dense, and Attentional Representations for Text Retrieval. In *TACL*, 2020.
- [37] T. K. R. Medini, Q. Huang, Y. Wang, V. Mohan, and A. Shrivastava. Extreme Classification in Log Memory using Count-Min Sketch: A Case Study of Amazon Search with 50M Products. In *NeurIPS*, 2019.
- [38] T. Mikolov, I. Sutskever, K. Chen, G. Corrado, and J. Dean. Distributed Representations of Words and Phrases and Their Compositionality. In NIPS, 2013.
- [39] P. Mineiro and N. Karampatziakis. Fast Label Embeddings via Randomized Linear Algebra. In ECML/PKDD, 2015.
- [40] A. Mittal, K. Dahiya, S. Agrawal, D. Saini, S. Agarwal, P. Kar, and M. Varma. DECAF: Deep Extreme Classification with Label Features. In WSDM, 2021.
- [41] A. Mittal, N. Sachdeva, S. Agrawal, S. Agrawal, P. Kar, and M. Varma. ECLARE: Extreme Classification with Label Graph Correlations. In WWW, 2021.
- [42] A. Mittal, K. Dahiya, S. Malani, J. Ramaswamy, S. Kuruvilla, J. Ajmera, K. Chang, S. Agrawal, P. Kar, and M. Varma. Multimodal extreme classification. In *CVPR*, June 2022.
- [43] S. Mohan, D. Saini, A. Mittal, S. R. Chowdhury, B. Paliwal, J. Jiao, M. Gupta, and M. Varma. Enriching Document Representations using Auxiliary Knowledge for Extreme Classification. In *ICML*, 2015.
- [44] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer. Automatic differentiation in PyTorch. In *NIPS-W*, 2017.
- [45] Y. Prabhu, A. Kag, S. Harsola, R. Agrawal, and M. Varma. Parabel: Partitioned label trees for extreme classification with application to dynamic search advertising.

WWW Companion '25, April 28-May 2, 2025, Sydney, NSW, Australia

Anshul Mittal et al.

In WWW 2018

- [46] Y. Qu, Y. Ding, J. Liu, K. Liu, R. Ren, W. X. Zhao, D. Dong, H. Wu, and H. Wang. Rocketqa: An optimized training approach to dense passage retrieval for opendomain question answering, 2021.
- [47] D. Saini, A.K. Jain, K. Dave, J. Jiao, A. Singh, R. Zhang, and M. Varma. GalaXC: Graph Neural Networks with Labelwise Attention for Extreme Classification. In WWW, 2021.
- [48] V. Sanh, L. Debut, J. Chaumond, and T. Wolf. DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. ArXiv, 2019.
- [49] W. Siblini, P. Kuntz, and F. Meyer. CRAFTML, an Efficient Clustering-based Random Forest for Extreme Multi-label Learning. In ICML, 2018.
- [50] T. Wei, W. W. Tu, and Y. F. Li. Learning for Tail Label Data: A Label-Specific Feature Approach. In IJCAI, 2019.
- [51] M. Wydmuch, K. Jasinska, M. Kuznetsov, R. Busa-Fekete, and K. Dembczynski. A no-regret generalization of hierarchical softmax to extreme multi-label classification. In NIPS, 2018.
- [52] L. Xiong, C. Xiong, Y. Li, K.-F. Tang, J. Liu, P. Bennett, J. Ahmed, and A. Overwijk. Approximate nearest neighbor negative contrastive learning for dense text retrieval. In ICLR, 2021.
- [53] J. Yang, Z. Liu, S. Xiao, C. Li, D. Lian, S. Agrawal, A. Singh, G. Sun, and X. Xie. Graphformers: Gnn-nested transformers for representation learning on textual

- graph. NeurIPS, 34:28798–28810, 2021. [54] Y. Yang, C. Huang, L. Xia, and C. Li. Knowledge graph contrastive learning for recommendation. In SIGIR Conference, pp. 1434-1443, 2022. URL https: //github.com/yuh-yang/KGCL-SIGIR22.
- [55] H. Ye, Z. Chen, D.-H. Wang, and B. .D. Davison. Pretrained Generalized Autoregressive Model with Adaptive Probabilistic Label Clusters for Extreme Multi-label Text Classification. In ICML, 2020.
- E.H. I. Yen, X. Huang, W. Dai, P. Ravikumar, I. Dhillon, and E. Xing. PPDSparse: [56] A Parallel Primal-Dual Sparse Method for Extreme Classification. In KDD, 2017.
- [57] R. You, S. Dai, Z. Zhang, H. Mamitsuka, and S. Zhu. AttentionXML: Extreme Multi-Label Text Classification with Multi-Label Attention Based Recurrent Neural Networks. In NeurIPS, 2019.
- [58] H. Zeng, H. Zhou, A. Srivastava, R. Kannan, and V. Prasanna. GraphSAINT: Graph Sampling Based Inductive Learning Method. In ICLR, 2020.
- [59] J. Zhang, W.-c. Chang, H.-f. Yu, and I. Dhillon. Fast multi-resolution transformer fine-tuning for extreme multi-label text classification. In NeurIPS, 2021.
- [60] J. Zhu, Y. Cui, Y. Liu, H. Sun, X. Li, M. Pelger, T. Yang, L. Zhang, R. Zhang, and H. Zhao. Textgnn: Improving text encoder via graph neural network in sponsored search. In theWebConf, pp. 2848-2857, 2021.
- [61] D. Zou, Z. Hu, Y. Wang, S. Jiang, Y. Sun, and Q. Gu. Layer-Dependent Importance Sampling for Training Deep and Large Graph Convolutional Networks, 2019.

Graph Regularized Encoder Training for Extreme Classification

Graph Regularized Encoder Training for Extreme Classification (Appendix)

A Implementation details

Links obtained on the metadata graph from raw data suffer from missing links in much the same way there are missing labels in the ground truth. To deal with this, RAMEN performs a random walk with restart on each anchor node. The random walk was performed for 400 hops with a restart probability of 0.8, thus ensuring that the walk did not wander too far from the starting node. This random walk could also introduce noisy edges, leading to poor model performance. To deal with such edges, in-batch pruning was performed and edges to only those anchors were retained which had a cosine similarity of > 0 based on the embeddings given the encoder. To get the encoder, RAMEN initialize the encoder with a pre-trained DistilBERT and fine-tuned it for 10 epochs(warmup phase) using unpruned metadata graphs. Then the metadata graphs were pruned using the fine-tuned encoder. Encoder fine-tuning was then was continued for 5 epochs using the pruned graphs after which the graphs were re-pruned. These alternations of 5 epochs of encoder fine-tuning followed by re-pruning were repeated till convergence. The learning rate for each bandit was set to 0.01. Table 10 in supplementary material summarizes all hyper-parameters for each dataset. It is notable that even though RAMEN uses a graph at training time, inference does not require any such information, making it highly suitable for long-tail queries.

B Data stats

Table 9: Dataset statistics summary for benchmark datasets used by RAMEN. Entries marked with ‡ were not disclosed because the dataset is proprietary.

# Train Pts N	# Labels	# Test Pts N'	Avg. docs. per label	Avg. labels per doc.	Graph Types	# Graph Nodes <i>G</i>	Avg. node neighbors per doc.	Avg. node neighbors per label
LF-WikiSeeAlsoTitles-320K								
693,082	312,330	177,515	4.67	2.11	Hyperlink Category	2,458,399 656,086	38.87 4.74	7.71 4.82

C Evaluation metrics

• Impression Yield or IY = $\frac{\text{Relevant Ads impressed}}{\text{Total number of queries}} \times 100$

• Click through rate or $CTR = \frac{Number of Clicks}{Number of Impressions} \times 100$

For more information about Precision@K (P@K) and nDCG@K (N@K) and their propensity scored variants, please refer to [3].

Table 10: Hyper-parameter values for RAMEN on all datasets to enable reproducibility. RAMEN code will be released publicly. Most hyperparameters were set to their default values across all datasets. LR is learning rate. Multiple clusters were chosen to form a batch hence B > C. Clusters were refreshed after 5 epochs. Cluster size C was doubled after every 25 epochs. Margin $\gamma = 0.3$ was used for contrastive loss. For training M2 number of positive samples and negative samples were kept at 2 and 12 respectively. A cell containing the symbol ↑ indicates that that cell contains the same hyperparameter value present in the cell directly above it.

Dataset	Batch	Encoder	Encoder LR	BERT seq.
	Size S	epochs	LR ₁	len L _{max}
LF-WikiSeeAlsoTitles-330K	1024	300	0.0002	32