

Lightweight Embeddings for Graph Collaborative Filtering

Xurong Liang
xurong.liang@uq.edu.au
The University of Queensland
Brisbane, Australia

Tong Chen
tong.chen@uq.edu.au
The University of Queensland
Brisbane, Australia

Lizhen Cui
clz@sdu.edu.cn
Shandong University
Jinan, China

Yang Wang
yangwang@hfut.edu.cn
Hefei University of Technology
Hefei, China

Meng Wang
eric.mengwang@gmail.com
Hefei University of Technology
Hefei, China

Hongzhi Yin*
h.yin1@uq.edu.au
The University of Queensland
Brisbane, Australia

ABSTRACT

Graph neural networks (GNNs) are currently one of the most performant and versatile collaborative filtering methods. Meanwhile, like in traditional collaborative filtering, owing to the use of an embedding table to represent each user/item entity as a distinct vector, GNN-based recommenders have inherited its long-standing defect of parameter inefficiency. As a common practice for scalable embeddings, parameter sharing enables the use of fewer embedding vectors (which we term meta-embeddings), where each entity is represented by a unique combination of meta-embeddings instead. When assigning meta-embeddings, most existing methods are a heuristically designed, predefined mapping from each user/item entity's ID to the corresponding meta-embedding indexes (e.g., double hashing), thus simplifying the optimization problem into learning only the meta-embeddings. However, in the context of GNN-based collaborative filtering, such a fixed mapping omits the semantic correlations between entities that are evident in the user-item interaction graph, leading to suboptimal recommendation performance. To this end, we propose **Lightweight Embeddings for Graph Collaborative Filtering (LEGCF)**, a parameter-efficient embedding framework dedicated to GNN-based recommenders. LEGCF innovatively introduces an assignment matrix as an additional learnable component on top of meta-embeddings. To jointly optimize these two heavily entangled components, aside from learning the meta-embeddings by minimizing the recommendation loss, LEGCF further performs efficient assignment update by enforcing a novel semantic similarity constraint and finding its closed-form solution based on matrix pseudo-inverse. The meta-embeddings and assignment matrix are alternately updated, where the latter is sparsified on the fly to ensure negligible storage overhead. Extensive experiments on three benchmark datasets have verified LEGCF's smallest trade-off between size and performance, with consistent accuracy gain over state-of-the-art baselines. The codebase of LEGCF is available in <https://github.com/xurong-liang/LEGCF>.

ACM Reference Format:

Xurong Liang, Tong Chen, Lizhen Cui, Yang Wang, Meng Wang, and Hongzhi Yin*. 2024. Lightweight Embeddings for Graph Collaborative Filtering. In *Proceedings of The 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '24)*. ACM, New York, NY, USA, 11 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 INTRODUCTION

In e-commerce and social media applications, the success of recommender systems (RSs) in predicting preferred items for users [48, 53, 63, 65, 66, 69] has attracted immense revenue. Classic RSs deploy factorization-based methods [23, 40] or multilayer perceptrons (MLPs) [15] for predicting user-item affinity via collaborative filtering. Recent advances in graph neural networks (GNNs) [14, 50, 53] further facilitate graph-based collaborative filtering [14, 50], which propagates collaborative signals through connected nodes to model user-item interactions. To date, GNN-based recommenders are arguably one of the most popular variants of collaborative filtering methods in both academia [14, 50] and industry [12, 61].

While GNN-based recommenders bring in promising performance and versatility, they still struggle to bypass a long-standing challenge carried over from conventional RSs – that is, the scalability of embedding representations. In the very basic form of ID-based recommendation where the only entities¹ to represent are users and items, each entity corresponds to an individual embedding vector that can be looked up from the embedding table based on a unique feature value (e.g., ID). With the sheer volume of users and items in e-commerce sites, the embedding table inevitably introduces a huge number of parameters that render a recommender hard to scale. The embedding parameters are at the scope of tens of millions in common benchmark datasets [3, 29, 54] and even thousands of billions in industry-level applications [35]. Thus, to keep enjoying the performance benefits of GNN-based recommenders at scale, improving the parameter efficiency of embeddings is the key.

To alleviate the excessive parameterization in the embedding layer of RSs, one mainstream research direction focuses on parameter sharing, in which all entity embeddings are drawn from a parameter pool with only a fraction of the footprint of the full embedding table. As a typical way of parameter sharing, *compositional embeddings* are facilitated via a set of meta-embedding

Permission to make digital or hard copies of all or part 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 components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).
SIGIR '24, July 14–18, 2024, Washington, D.C.

© 2024 Association for Computing Machinery.
ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00
<https://doi.org/XXXXXXX.XXXXXXX>

*Hongzhi Yin is the corresponding author.

¹In this paper, we term both users and items as entities for convenience.

vectors, where some hash functions are handcrafted to map each entity to a unique set of meta-embeddings based on their feature value (*i.e.*, user/item IDs in common RSs) [6, 8, 52, 55, 62]. Then, an entity embedding is composed with those meta-embeddings, *e.g.*, by performing sum pooling [42]. As each entity is assigned a distinct set of meta-embeddings, their final compositional embeddings naturally avoid collisions and can be easily used for subsequent recommendation tasks.

Nevertheless, in these methods, it is noticed that the mapping from entities to their corresponding meta-embeddings is solely based on the hash functions given [26, 42], which in turn makes all entities' meta-embedding assignments a predefined and fixed decision. Although this narrows the optimization target of this lightweight embedding paradigm down to just the meta-embeddings, such an assignment mapping unfortunately retains minimal semantic association between entities since it is solely determined by the user/item IDs that are mutually independent. That is to say, even two dissimilar entities, *e.g.*, a user and her disliked item, might have two similar meta-embedding compositions where most meta-embeddings are the same, resulting in a misleadingly high affinity between their composed embeddings. Similarly, two highly correlated entities can hardly preserve such correlations in their embeddings if they diverge heavily in meta-embedding assignments. In the meantime, most compositional embedding schemes allocate a unanimous weight to all meta-embeddings [27, 29] with no mechanisms to differentiate their contributions in an entity's compositional embedding, indulging the homogeneity of entity embeddings. As a result, risks are incubated in terms of insufficient embedding expressiveness and suboptimal recommendation performance, and they are likely to deteriorate in GNN-based collaborative filtering where user/item links are explicitly accounted for message passing.

To this end, a more flexible compositional embedding scheme that not only accounts for the semantic connections among entities, but also provides a more nuanced, weighted composition of assigned meta-embeddings for each entity, is highly anticipated. Ideally, this involves learning the right meta-embedding assignments as a parallel goal to the optimization of meta-embeddings. However, this also means selecting a subset of meta-embeddings for every entity, translating into a typical combinatorial optimization problem that is known to be NP-hard. At the same time, for each entity embedding, any update on either its meta-embedding assignments or the meta-embedding themselves will incur changes in its expressiveness. Given the strong entanglement between the two, simultaneously learning them is prone to instability during optimization and heavily constrained recommendation efficacy.

In this paper, to achieve the desiderata of parameter-efficient embeddings with GNN-based recommenders, we propose our work **Lightweight Embeddings for Graph Collaborative Filtering (LEGCF)**. Specifically, to remedy the defects of predefined hash function methods, we propose a more adaptive assignment scheme by introducing a learnable, real-valued assignment matrix, of which each row encodes which (selection) and how (weighting) every meta-embedding is involved in the composition of one entity's embedding. To learn the assignment matrix while preserving the semantic relationships among entities, we make use of GNNs' ability to propagate messages between associated nodes. Specifically,

we construct an expanded message passing graph that additionally considers the weighted connections between entities and meta-embeddings, and frame a representation similarity constraint as an objective function that can be elegantly and quickly solved in a closed form via matrix pseudo-inverse. This not only transforms the combinatorial optimization task into inferring the relevance of different meta-embeddings to each entity, but also decouples it from the gradient-based update of meta-embedding weights. Furthermore, in LEGCF, dynamic pruning is enabled for the assignment matrix, such that it can be efficiently stored as a sparse matrix with a negligible memory footprint.

We summarize our main contributions as follows:

- We point out the deficiencies of existing embedding parameter sharing methods when they are applied in GNN-based recommenders. Instead of using predefined meta-embedding assignments, we highlight the necessity of a flexible meta-embedding assignment mechanism that can preserve entity-wise semantic correlations and differentiate the nuanced contributions of each entity's assigned meta-embeddings.
- To achieve the desiderata, we propose a novel lightweight embedding framework dedicated to GNN-based recommenders, namely LEGCF. On top of learning all meta-embeddings, LEGCF adaptively learns weighted meta-embedding assignment for each entity. A novel, closed-form solution is derived from the similarity among graph-propagated embeddings, so as to facilitate an alternate optimization scheme.
- We conduct extensive experiments on three benchmark datasets. The results have demonstrated LEGCF's advantageous accuracy amid its top-tier parameter efficiency among state-of-the-art lightweight baselines, and verified its smallest performance compromise compared with GNN-based recommenders using a full ($40\times$ larger) embedding table.

2 RELATED WORK

In this section, we present the background of real-valued embedding compression in recommender systems.

2.1 Parameter Sharing

A large proportion of work achieves embedding compression by allowing multiple entities to share embeddings from a meta-embedding pool and hashing is the most notable trick for this purpose. These methods often utilize multiple hash functions to gather several meta-embeddings and compose them together to form entity embeddings (*i.e.*, compositional embedding) to avoid hash collision.

Compositional Embedding Methods. Zhang *et al.* [62] and Shi *et al.* [43] adopt multiple hashing to generate composed embeddings for parts of the entities. Q-R trick [26, 42] deploys the quotient and remainder codebooks to generate complementary composed embeddings. Yan *et al.* [55] propose to break the binary code of ID value into multiple small hash values, each of which yields a meta-embedding for the entity. SCMA [8] relies on locality sensitive hashing (LSH) to allocate segments of a single memory chunk to entities. Embeddings generated by tensor train (TT) decomposition methods [49, 54, 58, 59] can be treated as an alternative form of compositional embedding, as every entity embedding is a multiplication product of a series of tensors. Quantization [17, 27, 28, 30, 32, 64]

methods find the best fit codes from multiple codebooks and compose them to form entity embeddings, which is also compositional.

Apart from compositional embedding methods, some work facilitates the goal of parameter sharing in other ways. DHE replaces the hash embedding table with a deep neural network (DNN) [19]. LCE [13] only stores the embeddings of users/items and the embeddings of their counterparts are inferred at training time. Alternatively, clustering can be a strategy for entity meta-embedding assignment [6]. CEL [4] takes a different approach to split one giant cluster into multiple small clusters based on entity similarity. Entities in each cluster have a unified embedding.

Despite all methods mentioned above generating a memory-efficient embedding layer, no work leverages the user-item interaction graph for meta-embedding allocation. Hence, there is no direct mapping between entity interaction and representation proximity in the latent space, potentially leading to inadequately learned entity embeddings and hindering the downstream task.

2.2 AutoML Embedding Optimization

Another hot spot of embedding optimization studies lies in the application of automatic machine learning (AutoML) [60, 69]. One main direction of research is the automated dimension search [5, 10, 18, 31, 39, 67], in which a set of hand-crafted entity embedding candidate dimension sizes is supplied beforehand. The algorithms can then devise a differentiable dimension search [5, 67], or a reinforcement learning agent [18, 31, 38, 39] to determine the embedding dimension size for each entity. Despite these methods being effective in finding the balance between memory budget and the recommendation performance, the low computational efficiency of NAS [36] and the huge search space for reinforcement learning often raise concern. Pruning is another popular direction [3, 7, 22, 25, 29, 33, 34, 37, 56, 57, 68]. Work that purely relies on pruning to reach the memory target [7, 33, 37, 56, 57] may face the challenge of null-value embedding vector for entities, in which all elements lied in the embedding vector are zero, severely impacting recommendation performance. Later work [3, 22, 25, 34] attempts to combat the negative impact of pruning technique by combining other embedding optimization strategies, such as AutoML dimension search or embedding hashing. However, these hybrid methods do not overcome the inherited drawback of low computational efficiency, making it intractable to large-scale deployment.

3 METHOD

We now present LEGCF, our proposed lightweight embedding method for GNN-based recommenders. In Figure 1, we provide visual aids regarding LEGCF's key components, whose details are introduced in this section.

3.1 Replacing Full Embedding Table with Compositional Embeddings

We denote the sets of users and items as \mathcal{U} and \mathcal{I} , respectively. As we focus on ID-based recommendation in this paper, the total number of entities is $N = |\mathcal{U}| + |\mathcal{I}|$. The user-item interactions are encoded in a binarized matrix $\mathbf{R} \in \{0, 1\}^{|\mathcal{U}| \times |\mathcal{I}|}$, indicating whether an interaction is observed between a user-item pair. To minimize the memory footprint of the embedding layer, the full embedding

table $\mathbf{E} \in \mathbb{R}^{N \times d}$ with embedding dimensionality d is replaced by a memory-efficient meta-embedding codebook $\mathbf{E}_{meta} \in \mathbb{R}^{c \times d}$, where $c \ll N$ is the number of meta-embeddings, *a.k.a.* bucket size. To compose an embedding for each entity, a common approach [26, 29, 42] is to hash an entity's index p into t non-overlapping meta-embedding indexes $\{q_1, q_2, \dots, q_t\}$, based on which the corresponding meta-embeddings are drawn from the codebook \mathbf{E}_{meta} . Each entity's embedding is then formed by merging those t meta-embeddings, usually via pooling or element-wise product.

In our work, we abandon the use of fixed hash functions and instead introduce a learnable assignment matrix $\mathbf{S} \in \mathbb{R}_{\geq 0}^{N \times c}$. Essentially, each row of the assignment matrix \mathbf{S} corresponds to a user/item entity, where the non-zero elements indicate which meta-embeddings are selected out of the c choices, as well as their weights in the composition of the entity embedding. The entries in \mathbf{S} are updated during training, providing high flexibility in embedding generation. With the meta-embedding codebook \mathbf{E}_{meta} and learnable assignment matrix \mathbf{S} , all entities' compositional embeddings $\hat{\mathbf{E}} \in \mathbb{R}^{N \times d}$ can be computed via the following matrix multiplication:

$$\hat{\mathbf{E}} = \mathbf{S}\mathbf{E}_{meta}. \quad (1)$$

Notably, considering the size of a dense assignment matrix, we constrain \mathbf{S} to be highly sparse, where corresponding details are deferred to Section 3.3.

In GNN-based recommenders, a user-item interaction graph is built to propagate collaborative signals among connected entities for embedding learning [53], and a common practice [14, 50] is to represent those connections via an adjacency matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$:

$$\mathbf{A} = \begin{bmatrix} \mathbf{0} & \mathbf{R} \\ \mathbf{R}^\top & \mathbf{0} \end{bmatrix}, \quad (2)$$

where \mathbf{R} is the user-item interaction matrix. For each user $u \in \mathcal{U}$ and item $i \in \mathcal{I}$, we first retrieve their embeddings \mathbf{e}_u and \mathbf{e}_i from $\hat{\mathbf{E}}$, and treat them as the inputs (*i.e.*, layer-0 embeddings) to GNN. In short, each GNN layer involves the propagation of embeddings over the graph \mathbf{A} , where the final graph-propagated embeddings $\mathbf{h}_u, \mathbf{h}_i \in \mathbb{R}^d$ are computed by applying a pooling operation to the embeddings of each u/i obtained from each GNN layer. In our paper, we use LightGCN [14] as the base recommender and follow its configuration to use mean pooling over all layers' propagated embeddings as the final representation of each entity. As our main innovation lies in the new compositional embedding paradigm for GNN-based recommenders, we omit other details of LightGCN for simplicity. For each (u, i) pair, their recommendation affinity prediction can be calculated as the dot product of their graph-propagated embeddings:

$$\hat{y}_{ui} = \mathbf{h}_u^\top \mathbf{h}_i, \quad (3)$$

which facilitates the Bayesian personalized ranking loss (BPR) [41] for optimizing each entity's associated meta-embeddings:

$$\mathcal{L}_{BPR} = \sum_{(u, i^+, i^-) \in \mathcal{B}} -\ln \sigma(\hat{y}_{ui^+} - \hat{y}_{ui^-}) + \lambda \|\Theta\|^2, \quad (4)$$

where \mathcal{B} is either the whole training set or a training batch, (u, i^+, i^-) is a triplet that contains sampled user u 's observed interacted item i^+ and unvisited item i^- , $\|\Theta\|^2$ is the L_2 regularization over trainable parameters and λ controls its weight in the loss.

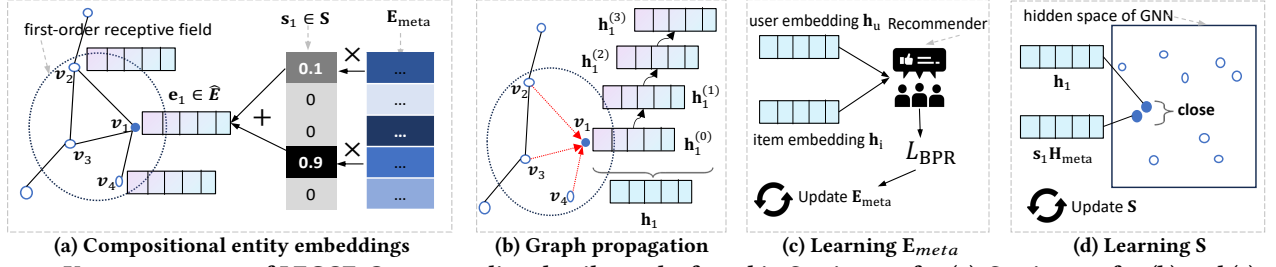


Figure 1: Key components of LEGCF. Corresponding details can be found in Section 3.1 for (a), Section 3.2 for (b) and (c), and Section 3.3 for (d). Note that we use e and s to denote one row of the codebook E_{meta} and assignment matrix S , respectively.

Batch Processing on Full Embedding Table Inference. Note that while we define the forward pass for computing all entities' embeddings as a unified matrix operation in Eq. 1, in cases where a large user and item number N challenges memory usage during training, our compositional embeddings also supports batch processing like in standard GNN-based recommenders. To do so, in the b -th batch, we can draw $n \ll N$ rows from S , forming a $n \times c$ matrix S_b for n entities. As such, for the b -th batch, Eq. 1 can be rewritten as $\hat{E}_b = S_b E_{meta}$ with $\hat{E}_b \in \mathbb{R}^{n \times c}$, which is compatible with all the subsequent computations.

3.2 Graph-Propagated Meta-Embeddings

The effectiveness of GNNs in refining entity embeddings via message passing has been widely acknowledged [9, 53]. In LEGCF, the compositional embedding table \hat{E} is generated from Eq. 1 with the meta-embedding codebook E_{meta} . As each meta-embedding is essentially linked to a set of nodes within the user-item interaction graph A , we represent such connections by incorporating the assignment matrix S into A :

$$A' = \begin{bmatrix} A & S \\ S^T & 0 \end{bmatrix}, \quad (5)$$

where $A' \in \mathbb{R}^{(N+c) \times (N+c)}$ is termed an *expanded interaction graph*. Intuitively, A' appends meta-embeddings as additional virtual nodes to the user-item interaction graph, thus taking advantages of the GNN for E_{meta} to propagate. The graph propagation of meta-embeddings resembles the same message passing process described in Section 3.1, where the input embeddings (*i.e.*, when layer $l = 0$) to GNN $H^{(0)} \in \mathbb{R}^{(N+c) \times d}$ is generated by stacking the compositional embeddings \hat{E} and the meta-embedding codebook E_{meta} :

$$H^{(0)} = \begin{bmatrix} \hat{E} \\ E_{meta} \end{bmatrix}. \quad (6)$$

The propagation operation at layer $l + 1$ is thus defined as:

$$H^{(l+1)} = (D^{-\frac{1}{2}} A' D^{-\frac{1}{2}}) H^{(l)}, \quad (7)$$

where A' is the adjacency matrix from the expanded interaction graph, $D' \in \mathbb{R}^{(N+c) \times (N+c)}$ is the diagonal degree matrix of A' . $D^{-\frac{1}{2}} A' D^{-\frac{1}{2}}$ creates a symmetrical, degree-normalized adjacency matrix. The final GNN embeddings can be gathered by taking the mean of embeddings from all layers:

$$H = \frac{1}{L+1} \sum_{l=0}^L H^{(l)}, \quad (8)$$

from which the graph-propagated representations of the full entity set $H_{full} \in \mathbb{R}^{N \times d}$ and meta-embeddings $H_{meta} \in \mathbb{R}^{c \times d}$ can then be respectively retrieved by splitting H into the following form:

$$H_{full} \leftarrow H[:, N:], H_{meta} \leftarrow H[N:, :]. \quad (9)$$

Specifically, the graph-propagated meta-embeddings H_{meta} are used for assignment matrix update as described in Section 3.3. In general, as H_{meta} now conceals collaborative information from semantically similar user/item neighbors, it facilitates learning a more sensible meta-embedding assignment for each user/item. To conduct the recommendation task, we retrieve the entities' graph-propagated embeddings from H_{full} , and then calculate the similarity score between each user-item pair as defined in Eq. 3. With a fixed assignment S in each iteration, the predicted scores of a training batch are passed to Eq. 4 to calculate the BPR loss \mathcal{L}_{BPR} , thus enabling back-propagation to update the values in the meta-embedding codebook E_{meta} .

3.3 Learning Assignment Weights

While the meta-embedding codebook E_{meta} is being learned, the weights in S should be updated over time as well in pursuit of a more reasonable mapping between entities and their meta-embedding assignments. As discussed earlier, a sensible meta-embedding assignment for each entity should be able to account for its own semantic relations to other entities. That is to say, two closely correlated entities are expected to have their embeddings composed by a similar set of meta-embeddings and vice versa – an intuitive analogy to the core idea behind collaborative filtering.

In this regard, a straightforward solution is to directly update the weights within S via gradient descent in the same back-propagation loop with E_{meta} . However, as the recommendation loss is directly affected by their product \hat{E} , when simultaneously updating the weights in both S and E_{meta} , the co-adaptation between their updates can introduce difficulties in finding a stable solution [16, 24]. Also, for each entity, instead of assigning weights to all c meta-embeddings and thus making S a dense matrix, a more desirable objective is to make it highly sparse, *i.e.*, for every row of S , only a few entries are non-zero. On the one hand, it makes less sense to associate each entity with all meta-embeddings in a compositional paradigm. On the other hand, if the $N \times c$ assignment matrix is dense, even a moderate choice of c (*e.g.*, $c = 100$) can make the storage cost of S comparable to the $N \times d$ full embedding table (*e.g.*, $d = 128$ as a popular choice), which defeats the purpose of compressing it at the first place. As a possible path to the discrete, weighted selection of meta-embeddings, reinforcement learning [46] unfortunately

requires a specifically crafted action sampler and reward function, and the optimization of meta-embedding combinations from a vast action space is hardly tractable.

To resolve the challenges, recall that \mathbf{S} defines the mapping from each entity to its corresponding meta-embeddings in the input layer of the GNN, where each $\mathbf{S}[p, q] \neq 0$ can be viewed as a similarity score between the p -th entity and the q -th meta-embedding. Optimally, if such similarity holds in the input layer, it should still be preserved between the entity embeddings and meta-embeddings after L layers of propagation. Therefore, we describe this property as the following semantic similarity constraint:

$$\mathbf{H}_{full} \approx \mathbf{S} \mathbf{H}_{meta}, \quad (10)$$

where \mathbf{H}_{full} and \mathbf{H}_{meta} are the graph-propagated embeddings as in Section 3.2. Assuming \mathbf{H}_{full} and \mathbf{H}_{meta} are in optimal form, the assignment matrix \mathbf{S} which reflects the relationship between \mathbf{H}_{full} and \mathbf{H}_{meta} can be solved as:

$$\mathbf{S} = \mathbf{H}_{full} \mathbf{H}_{meta}^{-1}, \quad (11)$$

with $\mathbf{H}_{meta}^{-1} \in \mathbb{R}^{d \times c}$ being the inverse of $\mathbf{H}_{meta} \in \mathbb{R}^{c \times d}$. However, as it is impractical to assume $c = d$, \mathbf{E}_{meta} is commonly a non-square matrix, neither does \mathbf{H}_{meta} . Consequently, the direct inverse of \mathbf{H}_{meta} may not exist. As a workaround, we approximate \mathbf{H}_{meta}^{-1} with the pseudo-inverse of \mathbf{H}_{meta} . We resort to the Moore-Penrose inverse [1] and solve the pseudo-inverse of \mathbf{H}_{meta} through singular value decomposition (SVD) [44] due to its computational simplicity and accuracy [1]. To be specific, we first decompose \mathbf{H}_{meta} into the following form:

$$\mathbf{H}_{meta} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^*, \quad (12)$$

where $\mathbf{U} \in \mathbb{R}^{c \times c}$ is a unitary matrix, $\mathbf{\Sigma} \in \mathbb{R}^{c \times c}$ is the square matrix that contains singular values along the diagonal, $\mathbf{V}^* \in \mathbb{R}^{d \times d}$ is the conjugate transpose of unitary matrix $\mathbf{V} \in \mathbb{R}^{d \times d}$. The Moore-Penrose pseudo-inverse of \mathbf{H}_{meta} , denoted as $\mathbf{H}_{meta}^\dagger \in \mathbb{R}^{d \times c}$, is computed as:

$$\mathbf{H}_{meta}^\dagger = \mathbf{V} \mathbf{\Sigma}^{-1} \mathbf{U}^*, \quad (13)$$

where $\mathbf{\Sigma}^{-1} \in \mathbb{R}^{c \times c}$ is the inverse matrix of $\mathbf{\Sigma}$, which can be easily obtained by replacing the singular values along the diagonal of $\mathbf{\Sigma}$ with their reciprocal, $\mathbf{U}^* \in \mathbb{R}^{c \times c}$ is the conjugate transpose of \mathbf{U} . Once the pseudo-inverse of \mathbf{H}_{meta} is calculated, we can plug it back into Eq. 11 to find the closed-form solution for \mathbf{S} :

$$\mathbf{S} = \mathbf{H}_{full} \mathbf{H}_{meta}^\dagger = \mathbf{H}_{full} \mathbf{V} \mathbf{\Sigma}^{-1} \mathbf{U}^*. \quad (14)$$

In this way, utilizing the constraint defined in Eq. 10, we can derive the updated assignment matrix \mathbf{S} via matrix pseudo-inverse. Such a gradient-free learning strategy prevents gradient-tracking on \mathbf{S} along with \mathbf{E}_{meta} , and is very efficient to compute. Moreover, it can effectively preserve the entities' semantic associations in their meta-embedding assignments, as two similar rows in \mathbf{H}_{full} will lead to similar results after the multiplication in Eq. 14.

3.4 Sparsifying, Batching, and Initializing \mathbf{S}

As a continuation of Section 3.3, we follow up with more design details about learning the assignment matrix \mathbf{S} .

Sparsification on Assignment Matrix. Since Eq. 14 yields a dense, real-valued matrix, to keep \mathbf{S} sparse, a sparsification process is required to make its storage footprint manageable. Fortunately,

our gradient-free assignment weight update strategy allows us to modify \mathbf{S} in place, where we conduct sparsification to retain the largest t non-zero weights in each row of \mathbf{S} so that the number of meta-embeddings used by each entity is kept at t at all time:

$$\mathbf{S}[p, q] = \begin{cases} \mathbf{S}[p, q] & \text{if } q \in \text{idx}_{\text{top-}t}(\mathbf{S}[p, :]) \\ 0 & \text{otherwise} \end{cases}, \quad (15)$$

where $\text{idx}_{\text{top-}t}(\cdot)$ returns the indexes of the top- t entries of a given vector (i.e., the p -th row of \mathbf{S} in our case). Essentially, only for meta-embeddings considered critical for an entity embedding, we retain their relatively larger weights in \mathbf{S} . In the worst case that every meta-embedding in \mathbf{E}_{meta} is selected in the sparsified \mathbf{S} , LEGCF exerts a space complexity of $O(tN + cd)$ for its lightweight embeddings. Considering $t \ll c$ ($t \leq 5$ in our settings), $O(tN + cd) \ll O(Nd)$ compared to the space complexity of a full embedding table, implying a significant improvement in parameter efficiency.

Batch Processing on Assignment Weight Update. In a similar vein to Eq. 1, Eq. 10 also supports batch computing, which can be enabled if lower memory consumption during training is desired. This is also facilitated by slicing the assignment matrix \mathbf{S} into batches, each of which has n rows ($n \ll N$). The operation in Eq. 14 to update \mathbf{S} is thus performed batch-wise:

$$\mathbf{S}_b = \mathbf{H}_b \mathbf{H}_{meta}^\dagger, \quad (16)$$

where b indexes the b -th batch. This substantially lowers the in-memory space complexity from $O(Nc + Nd + cd)$ to $O(nc + nd + cd)$, where the computed n dense rows of \mathbf{S} are immediately passed into Eq. 15 to obtain its sparse version.

Assignment Matrix Initialization. Given that similar entities are expected to own similar meta-embeddings, a good initialization of assignment matrix \mathbf{S} can potentially contribute to higher recommendation accuracy and a speedup in training. In \mathbf{S} , for each entity p , we term its highest-weighted meta-embedding (among the t selected) the *anchor meta-embedding* indexed by q_p^* , and it forms the base image of the entity's embedding \mathbf{e}_p since it takes the largest proportion in its composition. To let correlated entities have similar base embedding images during initialization, we propose to assign two entities p_1, p_2 the same anchor meta-embedding weight if they demonstrate high affinity with each other, i.e., $\mathbf{S}[p_1, q_{p_1}^*] = \mathbf{S}[p_2, q_{p_2}^*]$. Because the entities lying within each community in fact form tightly connected subgraphs in the original interaction graph, assigning them the same anchor meta-embedding helps pass the same proximity onto the latent space. In LEGCF, we take advantage of a well-established multilevel graph partitioning algorithm, namely METIS [20] for partitioning the user-item interaction graph. Although other advanced graph clustering methods can also be considered, METIS well serves the one-off initialization purpose due to its fast and accurate computation (no learning involved), balanced and non-overlapping partitions (all anchor meta-embeddings are fairly utilized), as well as deterministic results (ease of replication). We set the desired partition number to c , where entities in the same subgraph share one specific anchor meta-embedding in \mathbf{E}_{meta} . To reflect this, on initialization of the assignment matrix \mathbf{S} , we perform the following for every entity p :

- I. Given entity p 's subgraph index $c_p \in \{1, 2, \dots, c\}$, we set $\mathbf{S}[p, q_p^*] = w^*$, where $q_p^* = c_p$ and w^* is a universal hyperparameter.

- II. We uniformly sample $t - 1$ indexes from $\{1, 2, \dots, c\} \setminus c_p$ with replacement, denoted as a set Q . For every $q \in Q$, we set the corresponding assignment weight $\mathbf{S}[p, q] = \frac{(1-w^*)}{t-1}$.
- III. For all remaining entries at $q' \notin Q \cup c_p$, we set $\mathbf{S}[p, q'] = 0$.

In short, each entity receives one initial anchor meta-embedding based on the subgraph assigned, while the remaining $t - 1$ meta-embedding assignments are randomly initialized in favor of diversity. The impact of w^* will be examined later as a hyperparameter.

3.5 The Overall Algorithm

Algorithm 1: The algorithm of LEGCF.

```

1 Randomly initialize  $\mathbf{E}_{meta} \in \mathbb{R}^{c \times d}$ 
2 Initialize  $\mathbf{S} \in \mathbb{R}^{N \times c}$  following Section 3.4
3 Compute  $\mathbf{A}' \in \mathbb{R}^{(N+c) \times (N+c)}$  with Eq. 5
  /* Pretraining */
4 while not converged do
5   Optimize  $\mathcal{L}_{BPR}$  w.r.t. Eq. 5–9
  /* Assignment Update Enabled */
6 epoch_index  $\leftarrow$  0
7 while not converged do
8   Optimize  $\mathcal{L}_{BPR}$  w.r.t. Eq. 5–9
  /*  $m$  is a hyperparameter for update frequency */
9   if epoch_index mod  $m == 0$  then
10     Perform assignment update with Eq. 14
11     Regenerate  $\mathbf{A}'$  following Eq. 5
12   epoch_index  $\leftarrow$  epoch_index + 1

```

As depicted by the pseudocode in Algorithm 1, our proposed framework LEGCF consists of two stages. The first stage is the warm-up pretraining stage, wherein we randomly initialize the meta-embedding codebook \mathbf{E}_{meta} , generate the assignment matrix \mathbf{S} using METIS graph partitioning, and precompute the extended adjacency matrix \mathbf{A}' (lines 1-3). We freeze \mathbf{S} and only perform the downstream recommendation task to provide a stable environment for meta-embedding codebook learning (lines 4-5). Once the learning of \mathbf{E}_{meta} converges, we introduce the assignment matrix update to learn entities' assignment weights in \mathbf{S} and further refine the meta-embedding codebook \mathbf{E}_{meta} by alternating between the two procedures (lines 6-12). Note that as the meta-embedding assignments are computed with the graph-propagated representations, we only update \mathbf{S} every m epochs to ensure the graph-propagated representations are fairly stable by the time of updating \mathbf{S} . Furthermore, when the assignment matrix is updated, the expanded graph adjacency matrix \mathbf{A}' is also regenerated (line 11) to reflect updates in connections between entities and meta-embeddings.

4 EXPERIMENTS

In this section, we conduct experiments to validate the effectiveness of our work. We break down our experiments into the following research questions (RQs):

- **RQ1:** How is the recommendation accuracy of our work compared to other baselines under a tight memory budget?
- **RQ2:** Does METIS initialization and learnable assignment update improve recommendation performance?
- **RQ3:** How sensitive is our work to hyperparameters?
- **RQ4:** How does using different GNN recommenders affect our work's accuracy?

Table 1: Statistics of datasets used in our work.

Dataset	#User	#Item	#Interactions	Density
Gowalla	29,858	40,981	1,027,370	0.084%
Yelp2020	71,135	45,063	1,782,999	0.056%
Amazon-book	52,643	91,599	2,984,108	0.062%

4.1 Experimental Settings

4.1.1 Datasets. We use three publicly available benchmark datasets: **Gowalla**, **Yelp2020** and **Amazon-book**, which can be found in the official code repositories of [14] and [45]. The detailed statistics of datasets are summarized in Table 1.

4.1.2 Base Recommender and Baselines. As described in Section 3.1, the GNN backbone of LEGCF is implemented based on LightGCN [14]. For evaluation fairness, we also adopt LightGCN as the base recommender in the following baselines:

- **Variable-Size Methods** Methods in this category allows the generation of embedding layers to satisfy various memory targets. This is accomplished by applying pruning (**PEP** [33]) or changing the hash embedding structure (**QR** [42]). Note that our work is also capable of generating embedding layers of various memory usage. This is done by modifying the number of meta-embedding buckets c and the number of composition embeddings assigned to each entity t .
- **One-Size Methods** Methods in this category do not formalize the final memory target as an optimization objective. Thus, it is impractical to ensure the final embedding structure meets the memory budget. Baselines that belong to this category include AutoML-based dimension search algorithms **AutoEmb** [67], **ESAPN** [31], **OptEmbed** [34], **CISS** [39]; **DHE** [19] which replaces embedding hashing with DNN network; and **NimbleTT** [58, 59] which approximates the conventional embedding table using tensor-train decomposition. We omit the result of **CEL** [4] in this section due to its poor performance in our empirical testing.

Apart from the above baselines, to compare the performance compromise of various lightweight embedding methods, we also conduct experiments on the non-compressed versions of LightGCN with relatively large dimension sizes (*i.e.*, 64 and 128), which we term the Unified Dimensionality (**UD**) setting.

4.1.3 Evaluation Metrics. we select **NDCG@N** and **Recall@N** with N set to $\{10, 20\}$ as the evaluation metrics [51].

4.1.4 Implementation Details. We follow the train/test/validation protocol specified in [50] for dataset splitting. To generate interaction samples for training, for each user-positive item interaction, we randomly draw 5 negative items to form the training set. We set the dimension size of embeddings d to 128. We initialize the meta-embeddings and recommender weights using Xavier Initialization [11] and then deploy ADAM optimizer [21] for value learning. The learning rate is chosen from $\{1e-2, 1e-3, 1e-4\}$. The L_2 penalty factor λ is selected from $\{0, 5e-4, 1e-3, 5e-3, 1e-2\}$. We use a 3-layer GNN for Gowalla and a 4-layer GNN for Yelp2020 and Amazon-book. The number of compositional embeddings per entity t is set to 2 to reduce the memory usage of the assignment matrix. On initialization of \mathbf{S} , the weight of anchor embedding w^* for all entities is searched from the range $\{0.5, 0.6, 0.7, 0.8, 0.9\}$ for each

dataset. When the pretrain stage is done, we update the assignment matrix \mathbf{S} once every epoch. We choose the default number of meta-embedding buckets in the codebook c to be 500 and reveal the overall performance evaluated under this setting in RQ1. For variable-size baselines, we first calculate the total number of parameters used by our work under the 500-bucket setting, then we use that as the memory budget to select the correct sparsity target/number of hashing buckets for PEP and QR respectively, so that the parameter sizes of their embedding layers do not exceed the nominated memory budget. For baselines incapable of controlling the final parameter size, we define appropriate search spaces to avoid extreme memory footprint on the embedding layer and meanwhile, provide flexibility for these methods to select the most optimal setting. For each experiment setting, we select 3 random seeds to obtain results from multiple runs. We record the average results in this section.

4.2 Overall Performance (RQ1)

The overall performance of LEGCF with meta-embedding bucket size $c = 500$ and baseline methods are shown in Table 2. Regarding recommendation performance, on three datasets, LEGCF achieves the most outstanding result across all evaluation metrics compared to all other embedding optimization baselines. This is verified as the performance result of LEGCF exceeds the second best baseline (*i.e.*, CIESS for Gowalla and Yelp2020 and AutoEmb for Amazon-book) by a large margin. The performance boost of LEGCF is especially significant in Amazon-book dataset, revealing LEGCF’s suitability in big-scale data application. Comparing the parameter sizes achieved by various embedding optimization methods, LEGCF belongs to the tier that yields the lowest parameter sizes for all three datasets. Yet among its variable-size competitors PEP and QR, LEGCF attains the best performance. As for one-size methods, NimbleTT and CIESS attain similar parameter sizes to variable-size methods but their performance is not comparable to ours. Embedding layers generated by OptEmbed, ESAPN and AutoEmb techniques have much higher parameter sizes than ours. DHE yields the worst performance on all datasets, indicating small dimension size hash codes do not improve embedding uniqueness. Our experiment implies that most one-size methods are not designed to work under a tight memory budget.

On the other hand, comparing the result of LEGCF to unified dimensionality (UD) settings, it is observed that our work sacrifices a tidy drop in performance but in return, the memory usage of the embedding table is significantly reduced. Using the result evaluated on the Gowalla dataset as an example, despite applying LEGCF causes performance degradation of 6.5%, the parameter size of the embedding table generated by LEGCF is 22× smaller than the full embedding table in UD setting with dimension size = 64 and 44× smaller than that with dimension size = 128. LEGCF demonstrates an optimal tradeoff between memory usage and recommendation performance, such that the resultant embedding layer can easily fit into resource-constrained devices, yet the recommendation performance is barely compromised.

4.3 Model Component Analysis (RQ2)

The innovative components in LEGCF include METIS assignment matrix initialization and the assignment update process. To verify

their effectiveness in improving performance, we conduct ablation studies on these two main components. The performance comparison is depicted in Table 3. We further carry out case studies to showcase the learning of entity embeddings and assignment weights in different training stages.

Community-based Assignment Initialization. We modify the initialization strategy of assignment matrix \mathbf{S} to randomly assign anchor embeddings to entities. Our experiment reveals that the settings without METIS initialization perform much worse than those with it. This means having a meaningful assignment matrix on initialization is the key step to train an accurate meta-embedding codebook for high-quality entity embeddings.

Assignment Update. We eliminate the assignment update process by evaluating the recommender immediately after the pretrain stage converges. The drop in performance without assignment update shows the necessity of updating \mathbf{S} as it provides the opportunity for entities to alter the compositional weights to generate a more meaningful and customized embedding.

Case Study on Learned Entity Embeddings. Due to the page limit, in this part, we only discuss the case study on Gowalla dataset. To study the relationship between entity embeddings and their anchor embeddings, we sample 5 meta-embeddings and for each meta-embedding, we randomly select 100 entities whose initial anchor embedding is the chosen meta-embedding on initialization. Then we devise T-SNE [47] dimensionality reduction technique to visualize all selected meta-embeddings and entity embeddings in stages of framework initialization, completion of pretrain and final model after assignment weight update. The scatter plots can be found in Figure 2. From the three subplots, we can observe that on initialization, most entity embeddings are scattered around their anchor embeddings closely. At this stage, the entity embeddings do not contain any semantic correlations, thus the poor model accuracy. When the model pretrain is completed, meta-embeddings are learned according to the METIS assignment allocation. At this stage, entity embeddings spread more evenly around their anchor embeddings due to the freeze of assignment matrix. The recommendation performance has been largely improved owing to the learning of meta-embeddings. Finally, with the assignment weight update conducted, more flexibility is given in generating entity embeddings. The entity representations are more diversified and the sampled meta-embeddings are not necessarily the centroid of the entity group labeled the same color. This is because the assignment update process provides the ability to reselect their anchor embeddings, further strengthening the quality of entity embeddings.

Case Study on Learned Assignments. To study the change of assignment weight distribution before and after the weight update process, we sampled 20 meta-embeddings and plotted the average assignment weight distribution graphs for the initial and final stages of the model training in Figure 3. A huge weight shift can be spotted after multiple pseudo-inverse-based update rounds were conducted on the METIS initialized assignment matrix. Initially, METIS assigns roughly the same number of entities to each meta-embedding. This can be verified as the average weights for all sampled meta-embedding are similar. In the final stage, only a few outliers can be found in the average weight charts. In contrast, all other meta-embeddings have extremely small average weight values, meaning during the weight update process, critical meta-embeddings can be

Table 2: Performance comparison between LEGCF and baselines. “#Param” indicates the total parameter size of the embedding layer of each method. “UD - dim 128” and “UD - dim 64” are the full embedding table setting with unified dimensions $d = 128$ and $d = 64$, respectively. In each column, we use bold font to mark the best result achieved by a lightweight embedding method.

Method	Gowalla					Yelp2020					Amazon-book				
	#Param	N@10	R@10	N@20	R@20	#Param	N@10	R@10	N@20	R@20	#Param	N@10	R@10	N@20	R@20
UD - dim 128	9.07m	0.0901	0.1101	0.1059	0.1576	14.87m	0.0284	0.0426	0.0382	0.0721	18.46m	0.0172	0.0215	0.0230	0.0367
UD - dim 64	4.53m	0.0884	0.1088	0.1041	0.1557	7.44m	0.0274	0.0409	0.0366	0.0687	9.23m	0.0165	0.0204	0.0222	0.0353
ESAPN	0.78m	0.0268	0.0273	0.0305	0.0405	1.12m	0.0066	0.0083	0.0087	0.0141	1.44m	0.0045	0.0037	0.0051	0.0057
AutoEmb	0.92m	0.0273	0.0278	0.0319	0.0429	1.51m	0.0071	0.0089	0.0093	0.0148	1.44m	0.0048	0.0039	0.0059	0.0070
OptEmbed	1.42m	0.0375	0.0370	0.0405	0.0498	5.96m	0.0085	0.0090	0.0109	0.0159	2.11m	0.0038	0.0032	0.0050	0.0062
DHE	0.85m	0.0049	0.0051	0.0068	0.0105	1.39m	0.0020	0.0025	0.0028	0.0046	1.73m	0.0012	0.0010	0.0015	0.0019
NimbleTT	0.20m	0.0259	0.0296	0.0300	0.0431	0.28m	0.0055	0.0073	0.0071	0.0122	0.31m	0.0028	0.0031	0.0037	0.0055
CISS	0.29m	0.0643	0.0742	0.0745	0.1080	0.47m	0.0175	0.0253	0.0234	0.0437	0.57m	0.0021	0.0030	0.0029	0.0054
PEP	0.21m	0.0638	0.0568	0.0698	0.0807	0.30m	0.0170	0.0191	0.0218	0.0326	0.29m	0.0030	0.0024	0.0037	0.0043
QR	0.21m	0.0354	0.0416	0.0412	0.0606	0.30m	0.0067	0.0092	0.0088	0.0157	0.29m	0.0040	0.0046	0.0053	0.0082
LEGCF	0.21m	0.0846	0.0979	0.0988	0.1444	0.30m	0.0214	0.0310	0.0291	0.0548	0.35m	0.0134	0.0156	0.0172	0.0259

Table 3: Performance comparison between the default setting and settings with particular component modified. The default setting is the one with assignment update and METIS assignment initialization enabled. We use bold font to indicate the best result.

Assignment Update	Init Method	Gowalla		Yelp2020		Amazon-book	
		N@20	R@20	N@20	R@20	N@20	R@20
on	METIS	0.0988	0.1444	0.0291	0.0548	0.0172	0.0259
off	METIS	0.0901	0.1254	0.0272	0.0497	0.0147	0.0227
on	random	0.0788	0.1113	0.0183	0.0327	0.0109	0.0163

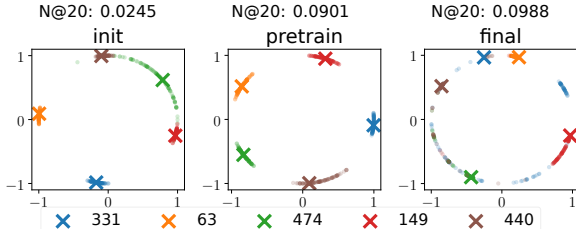


Figure 2: Visualization of sampled entity embeddings and their anchor meta-embeddings. The N@20 value annotated at the top of each plot is the NDCG@20 metric evaluated under the nominated setting. The sampled entities are represented as dots and each cross is a meta-embedding. The same color is used to scatter sampled entities assigned the identical anchor embedding on initialization. Note that we normalize the embeddings in $[-1, 1]$ scale for visualization consistency.

learned gradually, and more entities are assigned to these critical ones as the anchor embedding.

4.4 Hyperparameter Analysis (RQ3)

To study the sensitiveness of our work on hyperparameters, we conduct hyperparameter analysis on the number of meta-embeddings assigned to each entity t , the assignment matrix update frequency m , the initial weight of the anchor meta-embedding w^* as well as the bucket size of the meta-embedding codebook c .

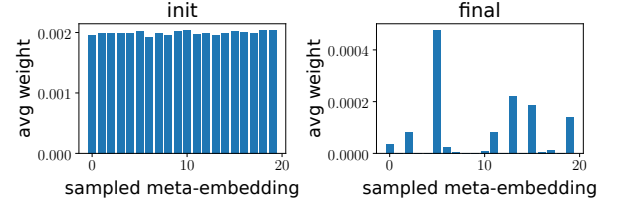


Figure 3: Average assignment weights of sampled meta-embeddings. The weights in the graphs have a small magnitude because we are taking the average weights over all users/items in the dataset.

Number of Meta-embeddings t per Entity. The recommendation performance regarding the number of compositional embeddings assigned to each entity t is plotted in Figure 4a. The set of t 's for testing is $\{1, 2, 3, 4, 5\}$. It is spotted that when each entity is assigned only 1 meta-embedding, the recommendation performance is severely impacted on all three datasets. This is because when each entity embedding is only generated using one meta-embedding, the meta-embedding mapping scheme is indifferent from single hash mapping, which causes serious hash collision. This statement can be verified as the change of t from 1 to 2 on both datasets leads to a great leap in performance immediately.

Assignment Matrix Update Frequency q . We alter the assignment matrix update frequency q to update every epoch and every 2, 3, 5 and 10 epochs respectively. The recommendation performance regarding different update frequencies is shown in Figure 4b. It is witnessed that updating the assignment matrix once every epoch in general provides a (nearly) optimal performance result for all three datasets. For Yelp2020, LEGCF is insensitive to the assignment weight update frequency q as the change of update frequency only causes a maximum performance difference of 2.4%. For Amazon-book dataset, changing the value of q may cause a performance difference of 11.3%. Our experiments indicate that the best assignment update frequency varies case by case. Some datasets require frequent updates to learn the weights of composed meta-embeddings in time, while some other datasets desire multiple

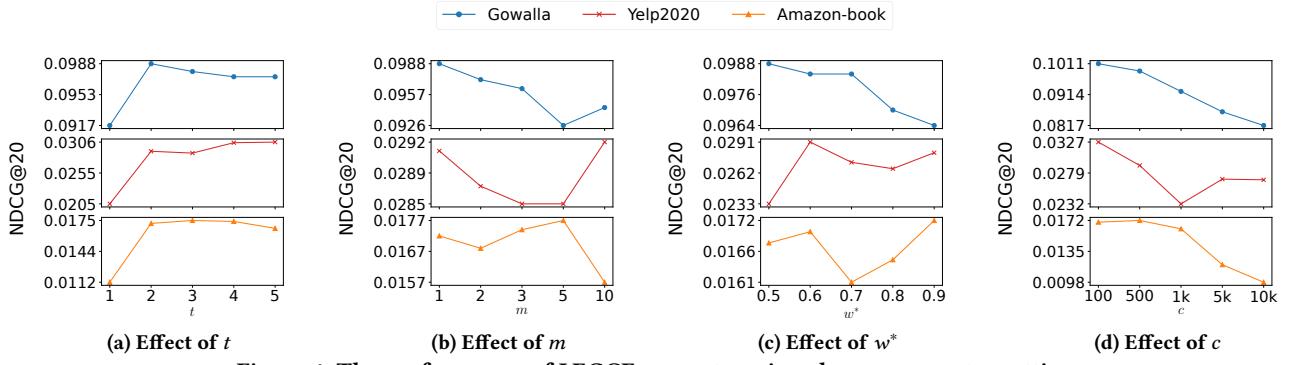


Figure 4: The performance of LEGCF on w.r.t. various hyperparameter settings.

Table 4: Performance comparison between various GNN base recommenders. The best performance under each setting is indicated with bold font.

	Gowalla		Yelp2020		Amazon-book	
Setting	N@20	R@20	N@20	R@20	N@20	R@20
LightGCN - LEGCF	0.0988	0.1444	0.0291	0.0548	0.0172	0.0259
LightGCN - UD	0.0584	0.0863	0.0169	0.0321	0.0074	0.0110
NGCF - LEGCF	0.0797	0.1092	0.0328	0.0602	0.0130	0.0204
NGCF - UD	0.0547	0.0810	0.0115	0.0230	0.0062	0.0096
LR-GCCF - LEGCF	0.0716	0.1052	0.0157	0.0272	0.0126	0.0192
LR-GCCF - UD	0.0001	0.0002	0.0002	0.0005	0.0006	0.0010

meta-embedding learning epochs between two assignment update processes to learn expressive meta-embeddings.

Initial Weight of the Anchor Meta-embedding w^* . The chart showing the relationship between model performance and the initial weight of anchor meta-embedding is depicted in Figure 4c. The set of w^* values for testing is $\{0.5; 0.6; 0.7; 0.8; 0.9\}$. It is observed that for Gowalla dataset, the best result is computed using the initial anchor weight of 0.5. As the value of w^* increases, a 2.4% drop in performance is caused. The performance of Yelp2020 and Amazon-book w.r.t. the value of w^* follows the same trend, that $w^* = 0.6$ yield satisfactory results, but further increasing w^* will first cause the decrease in performance, and then the performance is gradually improved again when w^* rises to 0.8 and 0.9.

Meta-embedding Bucket Size c . We draw the meta-codebook buckets c from the set $\{100; 500; 1,000; 5,000; 10,000\}$ for analysis. The graphs in this section are shown in Figure 4d. It is witnessed that for all three datasets, LEGCF attains the best performance under the small bucket size c settings. For Gowalla and Yelp2020, the best bucket size is 100 and for Amazon-book, the best bucket size is 500. As the bucket size increases to more than 1,000, its performance becomes less competitive. This phenomenon is especially obvious on the largest dataset Amazon-book, as the increase in bucket size leads to a rapid drop in performance. Our results indicate that LEGCF is designed to work under tight memory budget.

4.5 Framework Generalizability (RQ4)

We examine the generalizability of our framework in adopting various GNN-based recommenders. To do this, apart from LightGCN, we select two other popular GNN base recommenders, namely, the NGCF [50] and the LR-GCCF [2] to conduct performance testing. We apply appropriate dimension aggregation techniques to feed

their propagated embeddings to our assignment update process. We keep all other settings identical to LightGCN for consistency. In addition, we also implement the same GNN recommender settings on unified dimension (UD) embedding tables of similar parameter sizes for comparison. We report the performance in Table 4.

From the table, it is noted that LEGCF works well on all three GNN recommenders. LightGCN is the strongest base recommender for Gowalla and Amazon-book datasets. For Yelp2020 dataset, NGCF performs slightly better than LightGCN. Contrasting LEGCF’s strong ability in devising GNN to reinforce embedding quality under the memory-restrained scenario, UD settings are severely affected by the choice of the GNN recommenders on all three datasets. This is verified when comparing the performance between LEGCF and unified dimension settings under a tight memory budget. It is discovered that with the help of our method, the three base recommenders boost the performance substantially. Meanwhile, for unified dimension settings, the recommendation performance is heavily affected by the number of usable parameters. One extreme case is when the LR-GCCF is utilized to train unified dimension embedding tables for the three datasets, the performance is completely obsolete.

5 CONCLUSION

In this paper, we propose a novel graph-based compositional embedding framework LEGCF to resolve the challenges of fixed hash mapping and unanimous weight allocation to all composed meta-embeddings by introducing a compact meta-embedding codebook and a learnable assignment matrix. In LEGCF, we compute GNN propagated compositional embeddings for entities. We propose an efficient assignment update strategy to reflect the meta-embedding usage for composition in time. Both innovative designs significantly improve the quality of composed entity embeddings and boost the recommendation performance. Our comprehensive experiment indicates LEGCF’s superiority in performance over other embedding optimization work under memory-restrained scenarios.

ACKNOWLEDGMENTS

This work is partially supported by the Australian Research Council under the streams of Future Fellowship (Grant No. FT210100624), Discovery Early Career Researcher Award (Grant No. DE230101033), and Discovery Project (Grants No. DP240101108, and No. DP240101814).

REFERENCES

- [1] Adi Ben-Israel and Thomas NE Greville. 2003. *Generalized inverses: theory and applications*. Vol. 15. Springer Science & Business Media.
- [2] Lei Chen, Le Wu, Richang Hong, Kun Zhang, and Meng Wang. 2020. Revisiting graph based collaborative filtering: A linear residual graph convolutional network approach. In *AAAI*.
- [3] Tong Chen, Hongzhi Yin, Yujia Zheng, Zi Huang, Yang Wang, and Meng Wang. 2021. Learning elastic embeddings for customizing on-device recommenders. In *SIGKDD*.
- [4] Yizhou Chen, Guangda Huzhang, Anxiang Zeng, Qingtao Yu, Hui Sun, Heng-Yi Li, Jingyi Li, Yabo Ni, Han Yu, and Zhiming Zhou. 2023. Clustered Embedding Learning for Recommender Systems. In *WWW*.
- [5] Weiyou Cheng, Yanyan Shen, and Linpeng Huang. 2020. Differentiable neural input search for recommender systems. *arXiv preprint arXiv:2006.04466* (2020).
- [6] Abhinandan S Das, Mayur Datar, Ashutosh Garg, and Shyam Rajaram. 2007. Google news personalization: scalable online collaborative filtering. In *WWW*.
- [7] Wei Deng, Junwei Pan, Tian Zhou, Deguang Kong, Aaron Flores, and Guang Lin. 2021. Deeplight: Deep lightweight feature interactions for accelerating ctr predictions in ad serving. In *WSDM*.
- [8] Aditya Desai, Yanzhou Pan, Kuangyuan Sun, Li Chou, and Anshumali Shrivastava. 2021. Semantically constrained memory allocation (scma) for embedding in efficient recommendation systems. *arXiv preprint arXiv:2103.06124* (2021).
- [9] Chen Gao, Yu Zheng, Nian Li, Yinfeng Li, Yingrong Qin, Jinghua Piao, Yuhuan Quan, Jianxin Chang, Depeng Jin, Xiangnan He, et al. 2023. A survey of graph neural networks for recommender systems: Challenges, methods, and directions. *ReSys* 1, 1 (2023).
- [10] Antonio A Ginart, Maxim Naumov, Dheevatsa Mudigere, Jiyan Yang, and James Zou. 2021. Mixed dimension embeddings with application to memory-efficient recommendation systems. In *ISIT*.
- [11] Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *AISTATS*.
- [12] Saket Gururkar, Nikil Pancha, Andrew Zhai, Eric Kim, Samson Hu, Srinivasan Parthasarathy, Charles Rosenberg, and Jure Leskovec. 2022. MultiBiSage: A Web-Scale Recommendation System Using Multiple Bipartite Graphs at Pinterest. *Vldb* 16, 4 (2022).
- [13] Mengyue Hang, Tobias Schnabel, Longqi Yang, and Jennifer Neville. 2022. Lightweight compositional embeddings for incremental streaming recommendation. *arXiv preprint arXiv:2202.02427* (2022).
- [14] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, YongDong Zhang, and Meng Wang. 2020. LightGCN: Simplifying and Powering Graph Convolution Network for Recommendation. In *SIGIR*.
- [15] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural Collaborative Filtering. In *WWW*.
- [16] Geoffrey E Hinton et al. 2012. Improving neural networks by preventing co-adaptation of feature detectors. (2012).
- [17] Gangwei Jiang, Hao Wang, Jin Chen, Haoyu Wang, Defu Lian, and Enhong Chen. 2021. xLightFM: Extremely memory-efficient factorization machine. In *SIGIR*.
- [18] Manas R. Joglekar, Cong Li, Mei Chen, Taibai Xu, Xiaoming Wang, Jay K. Adams, Pranav Khaitan, Jiahui Liu, and Quoc V. Le. 2020. Neural Input Search for Large Scale Recommendation Models. In *SIGKDD*.
- [19] Wang-Cheng Kang, Derek Zhiyuan Cheng, Tiansheng Yao, Xinyang Yi, Ting Chen, Lichan Hong, and Ed H. Chi. 2021. Learning to Embed Categorical Features without Embedding Tables for Recommendation. In *SIGKDD*.
- [20] George Karypis and Vipin Kumar. 1997. METIS: A software package for partitioning unstructured graphs, partitioning meshes, and computing fill-reducing orderings of sparse matrices. (1997).
- [21] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [22] Shuming Kong, Weiyou Cheng, Yanyan Shen, and Linpeng Huang. 2022. AutoSrh: An Embedding Dimensionality Search Framework for Tabular Data Prediction. *TKDE* (2022).
- [23] Yehuda Koren, Robert Bell, and Chris Volinsky. 2009. Matrix Factorization Techniques for Recommender Systems. *Computer* 42, 8 (2009).
- [24] Gustav Larsson, Michael Maire, and Gregory Shakhnarovich. 2016. FractalNet: Ultra-Deep Neural Networks without Residuals. In *ICLR*.
- [25] Pan Li, Maofei Que, and Alexander Tuzhilin. 2023. Dual Contrastive Learning for Efficient Static Feature Representation in Sequential Recommendations. *TKDE* (2023).
- [26] Yang Li, Tong Chen, Peng-Fei Zhang, and Hongzhi Yin. 2021. Lightweight self-attentive sequential recommendation. In *CIKM*.
- [27] Defu Lian, Haoyu Wang, Zheng Liu, Jianxun Lian, Enhong Chen, and Xing Xie. 2020. LightRec: A Memory and Search-Efficient Recommender System. In *WWW*.
- [28] Defu Lian, Xing Xie, Enhong Chen, and Hui Xiong. 2020. Product quantized collaborative filtering. *TKDE* 33, 9 (2020).
- [29] Xurong Liang, Tong Chen, Quoc Viet Hung Nguyen, Jianxin Li, and Hongzhi Yin. 2023. Learning Compact Compositional Embeddings via Regularized Pruning for Recommendation. In *ICDM*.
- [30] Chong Liu, Defu Lian, Min Nie, and Xia Hu. 2020. Online optimized product quantization. In *ICDM*.
- [31] Haochen Liu, Xiangyu Zhao, Chong Wang, Xiaobing Liu, and Jiliang Tang. 2020. Automated Embedding Size Search in Deep Recommender Systems. In *SIGIR*.
- [32] Qi Liu, Jin Zhang, Defu Lian, Yong Ge, Jianhui Ma, and Enhong Chen. 2021. Online additive quantization. In *KDD*.
- [33] Siyi Liu, Chen Gao, Yihong Chen, Depeng Jin, and Yong Li. 2021. Learnable Embedding Sizes for Recommender Systems. *ICLR* (2021).
- [34] Fuyuan Lyu, Xing Tang, Hong Zhu, Huifeng Guo, Yingxue Zhang, Ruiming Tang, and Xue Liu. 2022. OptEmbed: Learning Optimal Embedding Table for Click-through Rate Prediction. In *CIKM*.
- [35] Dheevatsa Mudigere et al. 2022. Software-hardware co-design for fast and scalable training of deep learning recommendation models. In *Proceedings of the 49th Annual International Symposium on Computer Architecture*. 993–1011.
- [36] Hieu Pham, Melody Guan, Barret Zoph, Quoc Le, and Jeff Dean. 2018. Efficient neural architecture search via parameters sharing. In *ICML*.
- [37] Liang Qu, Yonghong Ye, Ningzhi Tang, Lixin Zhang, Yuhui Shi, and Hongzhi Yin. 2022. Single-shot embedding dimension search in recommender system. In *SIGIR*.
- [38] Yunke Qu, Tong Chen, Quoc Viet Hung Nguyen, and Hongzhi Yin. 2024. Budgeted embedding table for recommender systems. In *WSDM*.
- [39] Yunke Qu, Tong Chen, Xiangyu Zhao, Lizhen Cui, Kai Zheng, and Hongzhi Yin. 2023. Continuous Input Embedding Size Search For Recommender Systems. In *SIGIR*.
- [40] Steffen Rendle. 2010. Factorization Machines. In *ICDM*.
- [41] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian Personalized Ranking from Implicit Feedback. (2009). 10.
- [42] Hao-Jun Michael Shi, Dheevatsa Mudigere, Maxim Naumov, and Jiyan Yang. 2020. Compositional Embeddings Using Complementary Partitions for Memory-Efficient Recommendation Systems. *SIGKDD* (2020).
- [43] Shaoyun Shi, Weizhi Ma, Min Zhang, Yongfeng Zhang, Xinxing Yu, Houzhi Shan, Yiqun Liu, and Shaoping Ma. 2020. Beyond user embedding matrix: Learning to hash for modeling large-scale users in recommendation. In *SIGIR*.
- [44] Gilbert W Stewart. 1993. On the early history of the singular value decomposition. *SIAM review* 35, 4 (1993).
- [45] Jianing Sun, Zhao Yue Cheng, Saba Zuberi, Felipe Pérez, and Maksims Volkovs. 2021. Hgcf: Hyperbolic graph convolution networks for collaborative filtering. In *WWW*.
- [46] Richard S Sutton and Andrew G Barto. 2018. *Reinforcement learning: An introduction*.
- [47] Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of machine learning research* (2008).
- [48] Hao Wang, Yanmei Fu, Qinyong Wang, Hongzhi Yin, Changying Du, and Hui Xiong. 2017. A location-sentiment-aware recommender system for both hometown and out-of-town users. In *KDD*.
- [49] Qinyong Wang, Hongzhi Yin, Tong Chen, Zi Huang, Hao Wang, Yanchang Zhao, and Nguyen Quoc Viet Hung. 2020. Next Point-of-Interest Recommendation on Resource-Constrained Mobile Devices. In *WWW*.
- [50] Xiang Wang, Xiangnan He, Meng Wang, Fuli Feng, and Tat-Seng Chua. 2019. Neural graph collaborative filtering. In *SIGIR*.
- [51] Yining Wang, Liwei Wang, Yuanzhi Li, Di He, Wei Chen, and Tie-Yan Liu. 2013. A theoretical analysis of NDCG ranking measures. In *COLT*.
- [52] Kilian Weinberger, Anirban Dasgupta, John Langford, Alex Smola, and Josh Attenberg. 2009. Feature hashing for large scale multitask learning. In *ICML*.
- [53] Shiwen Wu, Fei Sun, Wentao Zhang, Xu Xie, and Bin Cui. 2023. Graph Neural Networks in Recommender Systems: A Survey. *Comput. Surveys* (2023).
- [54] Xin Xia, Hongzhi Yin, Junliang Yu, Qinyong Wang, Guandong Xu, and Quoc Viet Hung Nguyen. 2022. On-Device Next-Item Recommendation with Self-Supervised Knowledge Distillation. In *SIGIR*.
- [55] Bencheng Yan, Pengjie Wang, Jinqian Liu, Wei Lin, Kuang-Chih Lee, Jian Xu, and Bo Zheng. 2021. Binary code based hash embedding for web-scale applications. In *CIKM*.
- [56] Bencheng Yan, Pengjie Wang, Kai Zhang, Wei Lin, Kuang-Chih Lee, Jian Xu, and Bo Zheng. 2021. Learning effective and efficient embedding via an adaptively-masked twins-based layer. In *ICKM*.
- [57] Yao Yao, Bin Liu, Haoxun He, Dakui Sheng, Ke Wang, Li Xiao, and Huanhuan Cao. 2022. i-Razor: A Neural Input Razor for Feature Selection and Dimension Search in Large-Scale Recommender Systems. *arXiv preprint arXiv:2204.00281* (2022).
- [58] Chunxing Yin, Bilge Acun, Carole-Jean Wu, and Xing Liu. 2021. Tt-rec: Tensor train compression for deep learning recommendation models. *Proceedings of Machine Learning and Systems* 3 (2021).
- [59] Chunxing Yin, Da Zheng, Israt Nisa, Christos Faloutsos, George Karypis, and Richard Vuduc. 2022. Nimble GNN Embedding with Tensor-Train Decomposition. In *SIGKDD*.
- [60] Hongzhi Yin, Liang Qu, Tong Chen, Wei Yuan, Ruiqi Zheng, Jing Long, Xin Xia, Yuhui Shi, and Chengqi Zhang. 2024. On-Device Recommender Systems: A Comprehensive Survey. *arXiv preprint arXiv:2401.11441* (2024).

- [61] Rex Ying, Ruining He, Kaifeng Chen, Pong Eksombatchai, William L Hamilton, and Jure Leskovec. 2018. Graph convolutional neural networks for web-scale recommender systems. In *SIGKDD*.
- [62] Caojin Zhang, Yicun Liu, Yuanpu Xie, Sofia Ira Ktena, Alykhan Tejani, Akshay Gupta, Pranay Kumar Myana, Deepak Dilipkumar, Suvadip Paul, Ikuhiro Ihara, Prasang Upadhyaya, Ferenc Huszar, and Wenzhe Shi. 2020. Model Size Reduction Using Frequency Based Double Hashing for Recommender Systems. In *RecSys*.
- [63] Junwei Zhang, Min Gao, Junliang Yu, Lei Guo, Jundong Li, and Hongzhi Yin. 2021. Double-scale self-supervised hypergraph learning for group recommendation. In *CIKM*.
- [64] Jin Zhang, Qi Liu, Defu Lian, Zheng Liu, Le Wu, and Enhong Chen. 2022. Anisotropic additive quantization for fast inner product search. In *AAAI*.
- [65] Shuai Zhang, Lina Yao, Aixin Sun, and Yi Tay. 2020. Deep Learning based Recommender System: A Survey and New Perspectives. *Comput. Surveys* 52 (2020).
- [66] Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. 2023. A survey of large language models. *arXiv preprint arXiv:2303.18223* (2023).
- [67] Xiangyu Zhao, Haochen Liu, Wenqi Fan, Hui Liu, Jiliang Tang, Chong Wang, Ming Chen, Xudong Zheng, Xiaobing Liu, and Xiwang Yang. 2021. Autoemb: Automated embedding dimensionality search in streaming recommendations. (2021).
- [68] Ruiqi Zheng, Liang Qu, Tong Chen, Kai Zheng, Yuhui Shi, and Hongzhi Yin. 2024. Personalized Elastic Embedding Learning for On-Device Recommendation. *TKDE* (2024).
- [69] Ruiqi Zheng, Liang Qu, Bin Cui, Yuhui Shi, and Hongzhi Yin. 2023. AutoML for Deep Recommender Systems: A Survey. *TOIS* (2023).