

On-device Content-based Recommendation with Single-shot Embedding Pruning: A Cooperative Game Perspective

Anonymous Author(s)

Abstract

Content-based Recommender Systems (CRSs) play a crucial role in shaping user experiences in e-commerce, online advertising, and personalized recommendations. However, due to the vast amount of categorical features, the embedding tables used in CRS models pose a significant storage bottleneck for real-world deployment, especially on resource-constrained devices. To address this problem, various embedding pruning methods have been proposed, but most existing ones require expensive retraining steps for each target parameter budget, leading to enormous computation costs. In reality, this computation cost is a major hurdle in real-world applications with diverse storage requirements, such as federated learning and streaming settings. In this paper, we propose **Shapley Value-guided Embedding Reduction (Shaver)** as our response. With Shaver, we view the problem from a cooperative game perspective, and quantify each embedding parameter’s contribution with Shapley values to facilitate contribution-based parameter pruning. To address the inherently high computation costs of Shapley values, we propose an efficient and unbiased method to estimate Shapley values of a CRS’s embedding parameters. Moreover, in the pruning stage, we put forward a field-aware codebook to mitigate the information loss in the traditional zero-out treatment. Through extensive experiments on three real-world datasets, Shaver has demonstrated competitive performance with lightweight recommendation models across various parameter budgets. The source code is available at <https://anonymous.4open.science/r/shaver-E808>.

1 Introduction

In the modern era, recommender systems (RSs) play a vital role in assisting users to identify relevant information. Meanwhile, on top of user-item interactions, the increasingly available side features of users and items (e.g., user gender and product category) have made content-based RSs (CRSs) [36] a step above collaborative filtering-based RSs, especially in cold-start settings. According to the Statista report [50], in the US, online advertising – a major CRS application – has doubled its market size from \$107.5 billion in 2018 to \$225 billion in 2023. More recently, with privacy legislation and cyber security awareness on the rise, on-device deployment of CRSs is quickly gaining popularity, where examples include Brave Browser’s federated news recommendation [37], as well as Kuaishou’s short video recommendation on mobile devices [15].

Most modern CRSs depend on sparse categorical features, such as movie genres and user regions in movie recommendation [36], where a dense embedding table that hosts all possible features’ embeddings (i.e., vector representations) is trained. Given the enormous amount of sparse categorical features, embedding tables have become the storage bottleneck for modern CRSs, which contradicts the surging demand for scalability and efficiency during on-device deployment [41, 60]. For illustration, in click-through rate (CTR)

prediction, a typical CRS task that predicts users’ clicking intention on an item, the RS deployed by Meta has been reported to consume 12T parameters and can demand up to 96TB of memory and multiple GPUs to train [39]. Consequently, lowering the parameter footprint has attracted immense attention in both industry applications [24, 41] and academic research [19, 43]. For instance, PEP [32] prunes the embedding table by searching a pruning mask through a learnable threshold. After the initial training of both the threshold and the model, the model is retrained from scratch to work under the pruned setting. Another example is AutoSrh [26], which adaptively groups features into blocks. After bi-level optimization, a sparse embedding parameterization can be selected according to the memory budget and then optimized via retraining.

Generally, these methods focus on pruning the embedding table to a fixed size, hence an expensive training cycle is required for every single target sparsity. On one side, in most existing solutions, optimal trade-offs between recommendation accuracy and parameter efficiency require tailoring models according to each device’s specific computational and storage capacities. However, on the flip side of the coin, such expensive, repetitive training procedures are not affordable in practical settings with varying on-device storage needs. We provide two scenarios below:

- **Scenario 1:** Federated recommendation ensures privacy by assigning each user an on-device recommender. Such paradigms commonly suffer from hardware heterogeneity, where the variation of computing resources across devices (e.g., TV boxes, smart watches, and mobile phones) [7, 60] prohibits a uniform model architecture from being used. This naturally calls for more efficient means of embedding pruning to provide multiple versions of a trained model.
- **Scenario 2:** When deployed on-device, a model needs to operate under different computational resource budgets due to software parallelization and varying battery status [4, 61], and CRSs are no exception. Considering the time-sensitive nature of recommendation services, a sophisticated embedding pruning mechanism that requires reinforcement learning [43] or even post-pruning retraining [32] will lead to a significant overhead, hindering the usability of the service pipeline.

Ideally, considering those real-life scenarios, a competent pruning technique is expected to quickly adapt a well-trained CRS model to any specified parameter budget, without undergoing iterative search or retraining that are both time-consuming.

Therefore, this paper focuses on a more practical, yet largely neglected setting of CRS compression for on-device deployment, which is to prune the embedding table of a recommender in one single shot¹. To achieve single-shot pruning, a common practice is based on attribution scores [2, 38]. In a nutshell, by quantifying

¹To clarify, different from transfer learning, our single-shot setting refers to performing model pruning in only one forward (and sometimes backward) pass [27, 42].

the contribution from each individual model weight (i.e., embedding parameter in our case) to a performance metric, parameters with the lowest scores can be nullified to meet an arbitrary size constraint. In content-based recommendation, existing solutions [42, 55] perform parameter attribution by measuring the loss value change after removing the parameter from the full embedding table, which can be approximated with first-order gradients [27]. If we treat all embedding parameters as cooperative “players” in a game, then this is essentially asking the question: *In a team formed by all players, what is each individual player’s contribution to the whole team?* However, this formulation hardly aligns with the pruning context, where a player (i.e., embedding parameter), if not pruned, will actually form a smaller team with only a subset of players. In a single-shot pruning setting, instead of retaining embedding parameters that are most important to the full parameter set, we should ultimately identify parameters that can be of high value to any parameter subsets. So, the right question to be asked is: *What is the expected contribution of each individual player, when it joins any team formed by a subset of players?* Due to the intertwined dependency among parameters, it is clear that this question does not share the same answer with the former one.

To answer this question, we aim to investigate embedding pruning for CRSs via the lens of cooperative games, where we employ Shapley value [47] to perform embedding parameter attribution. Shapley value is a well-defined method in the cooperative game theory for fairly allocating the total game revenue to players based on their individual contributions [12, 23, 34]. It quantifies each player’s contribution by considering all possible team combinations, resulting in the expected contribution of the player in any group. As such, compared with existing practices [27, 42], Shapley value is a better fit for generating unbiased attribution scores of embedding parameters to facilitate fast recommendation model pruning.

However, some major challenges are yet to be resolved before we can enjoy the benefit of Shapley values. Firstly, the computational costs of Shapley value are prohibitively expensive. In our context, when a CRS only activates a subset of embedding parameters, each parameter’s contribution in this subset can be measured by comparing the CRS’s recommendation performance with and without it. Then, a parameter’s Shapley value is its average contribution to all possible subsets of parameters. To calculate all the exact Shapley values, we need to enumerate over all 2^{nd} parameter subsets, where n and d are respectively the numbers of features and embedding dimensions. Considering that in CRSs nd is intractable (e.g., reaching 100 million in our biggest dataset), and reliably measuring the recommendation performance requires a sufficiently large dataset, the number of resulted forward passes has rendered the straightforward calculation of Shapley values an infeasible option. Secondly, most pruning methods employ a simple zero-out approach to remove less important parameters and create a sparse embedding matrix. While this reduces the model size, it can result in significant information loss as recommenders rely heavily on dense feature representations for accurate predictions. This further deteriorates in CRSs due to their extensive use of dot product [18, 46] and element-wise multiplication [53] between embeddings to model feature interactions. By setting parameters to zero, pruning can disrupt the complex relationships between features, leading to suboptimal performance.

Motivated by these challenges, we propose **Shapley Value-guided Embedding Reduction (Shaver)**, a novel embedding pruning method for compressing a CRS model to any target size in one shot. Shaver assigns each embedding parameter a Shapley value as its attribution score, such that a well-trained model can prune its embeddings to any parameter budget in a single shot. To address the high computational costs of obtaining exact Shapley values, through a series of theoretical analyses, we develop an unbiased approximation method to reliably estimate the Shapley values of all embedding parameters. Efficiency-wise, Shaver reduces the computational complexity from $O((nd)! \times |\mathcal{D}|)$ to the much lower $O(md \times |\mathcal{D}|)$, with \mathcal{D} and m respectively being the dataset and number of feature fields (e.g., user gender and movie genre). Notably, $m \ll n$ in our case – tens of feature fields versus millions of features (see Table 1). To reduce the information loss caused by the default zero-padding on pruned parameters, we propose an alternative that searches for an optimal placeholder value per feature field, termed the codebook, which is then used for padding nullified embedding parameters during inference. In addition, the codebook values can be solved in a closed form, making it better at retaining the recommendation performance with a negligible computational overhead.

In summary, our main contributions are:

- We revisit one important but largely neglected task of single-shot embedding pruning, which is more relevant in the real-world, on-device deployment of CRSs.
- We propose Shaver, a novel embedding compression method that reduces the time complexity of Shapley value estimation for CTR models and provides a more effective alternative to the widely used zero-out approach in model pruning with a field-aware codebook.
- We conduct extensive experiments on three datasets to compare Shaver with state-of-the-art embedding compression methods. The results show that even without an extra fine-tuning step, Shaver can achieve competitive performance.

2 Related Work

2.1 CTR Prediction

CTR prediction plays a vital role in the development of online advertisements. One of the first and most influential backbones for CTR prediction is Factorization Machine (FM) [46], where the interaction between features are explicitly model upto 2nd order. Recently, with the rapid development of deep neural networks (DNN), various deep models have also been proposed. NeuMF [21] combines the introduced Generalized Matrix Factorization (GMF) with a DNN network branch. DeepFM [18] is another prime example, which adds a DNN branch into the FM model. Unlike DeepFM, which only attempts to model up to 2nd-order interaction, DCNv2 [53] explicitly models the $l + 1$ -th order interaction by l crossnet layers. DCNv2 also incorporates a DNN to model more complex interactions. To further reduce the computation cost of DCNv2, Wang et al. [53] also propose integrating Mixture-of-Experts for modeling the cross-layers, creating DCN-Mix.

2.2 Lightweight Recommender Systems

A diverse set of research endeavors has proposed various methods for learning light-weight embeddings [29, 52, 62]. Compositional

embedding methods decompose the original embedding table into multiple smaller meta-embedding tables and share parameters between features [8, 9, 25, 30, 41, 49, 59]. Quantization approaches are also applied in light-weight RSs [16, 28, 56, 58], which focus on reducing the loss of model quality in lower precision representation. However, these approaches generally suffer from limited compression ratios [62]. Another line of research focuses on finding the best parameter allocation through Neural Architecture Search (NAS), which usually involves reinforcement learning [24, 44, 45], bi-level optimization [26, 63, 64], or evolutionary search [7, 35]. Consequently, NAS-based approaches generally have high computation costs. Pruning is another popular approach to reduce the amount of embedding parameters by finding a binary mask [32, 42, 54, 57]. Recently, many hybrid methods [31, 35] have also been proposed to combine the benefit of various archetypes to achieve the best trade-off between performance and parameter budget.

Our method falls within the pruning category. Unlike most of other pruning methods [32, 57], we eliminate the hassle of simultaneously optimizing both the pruning budget mask and model performance. Compared to [7, 45], which also can adapt to arbitrary parameter budgets but focus on collaborative filtering, our method tailored towards CRSs. Additionally, unlike SSEDs [42], which mandates uniform dimensions across all features within the same field, our method offers greater flexibility by allowing varied budget allocations without such constraints.

2.3 Shapley Value for Machine Learning

Shapley values have been widely adopted in the explainable AI community due to its high correlation with human intuition [34]. Because of its high computation cost, various researches have been proposed to tackle this challenge through random order sampling [5, 12], solving weighted least square problem [11, 34], being model specific [33, 34], and recently amortization approaches [10, 23]. Recently, the success of Shapley values in explainable AI has inspired its application in pruning deep models. To start with, Ancona et al. [2] pointed out several desirable properties of Shapley values in pruning neurons:

- Null player: If the game (loss) doesn't depend on a player (parameter or group of parameters), its attribution is zero.
- Symmetry: If the game depends on two players equally, then the two players should receive the same attribution.
- Efficiency: Attributions sum up to the difference between the loss evaluated when all players are enabled and the loss evaluated when all players are removed.

Later, NeuronShapley [14] considered Shapley value estimation as a multi-arm bandit problem, and applied truncated Monte Carlo to estimate Shapley values. Leveraging the multi-arm bandit settings, Guan et al. [17] employ an ϵ -greedy approach. Troupe [48] proposes using Shapley values to select models for an ensemble based on a specific data point. Last, although it is not applied for compressing embedding tables, Shapley values also have been applied in recommender system designs [3, 13].

3 Preliminaries

In this section, we first introduce the key concepts of content-based recommendation by defining a classic task, namely click-through

rate (CTR) prediction. Then, we provide the definition of Shapley value, the theoretical foundation of our solution.

3.1 Content-based Recommendation

In CTR prediction, a dataset \mathcal{D} is a set of samples (\mathbf{x}, y) , where \mathbf{x} is a multi-dimensional vector combining both user and item features, and y is a binary label indicating positive ($y = 1$, clicked) or negative ($y = 0$, not clicked) user interaction with the item. For \mathbf{x} , its features are commonly a collection of sparse binary features from multiple fields (i.e., user occupation and product category) [36, 65], such as:

$$\mathbf{x} = \underbrace{[1, 0]}_{\mathbf{x}_1 \text{ (field 1)}} \underbrace{[0, \dots, 1, \dots, 0]}_{\mathbf{x}_2 \text{ (field 2)}} \dots \underbrace{[0, \dots, 1, \dots, 0]}_{\mathbf{x}_m \text{ (field } m)}, \quad (1)$$

where an example of \mathbf{x}_1 and \mathbf{x}_2 can be *user gender = female* in field 1 and *movie genre = action* in field 2. \mathbf{x} is then a concatenation of all categorical feature encodings $\{\mathbf{x}_j\}_{j=1}^m$ from all m fields, where numerical features are commonly discretized via bucketing [65] in CTR prediction tasks.

In real-world applications, the input vector \mathbf{x} is extremely sparse, therefore each field j 's feature is mapped into a d -dimensional, dense latent vector $\mathbf{e}_j \in \mathbb{R}^d$, called embedding:

$$\mathbf{e}_j = \mathbf{V}_j^\top \mathbf{x}_j, \quad (2)$$

where $\mathbf{V}_j \in \mathbb{R}^{n_j \times d}$ and $\mathbf{x}_j \in \mathbb{R}^{n_j}$ are respectively the embedding table and sparse encoding for field j , with n_j the number of features in field j . In practice, all m field-specific embedding tables \mathbf{V}_i are concatenated as a single embedding table $\mathbf{E} \in \mathbb{R}^{n \times d}$ for storage, where $\mathbf{E} = [\mathbf{V}_1, \mathbf{V}_2, \dots, \mathbf{V}_m]$, and $n = \sum_{j=1}^m n_j$ is the total number of features. After obtaining embeddings for \mathbf{x} from all fields $\{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_m\}$, they are further processed by a recommendation backbone model:

$$\hat{y} = f(\{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_m\}), \quad (3)$$

with \hat{y} is a binary classification prediction whether a user will click on an item or not. Its primary target is to model the complex interactions between different features. For instance, the model can combine the user gender and the movie genre to create a 2nd-order interaction, enhancing model performance [18, 46]. These high order interactions are commonly modeled explicitly by dot product [18, 46] or element-wise multiplications [53], and implicitly by deep neural networks [18, 53]. Finally, $f(\cdot)$ is trained with log loss:

$$L = -y \ln(\hat{y}) - (1 - y) \ln(1 - \hat{y}), \quad (4)$$

which quantifies the prediction error between y and \hat{y} .

3.2 Shapley Value

Shapley value originates from the cooperative game theory, whose aim is to assign each player a score that quantifies the player's contribution in a cooperative game. Formally, a cooperative game consists of a set of players denoted by \mathcal{N} , where any arbitrary subset of players can form a "team". The coalitional values generated by player subsets are measured by function $v : 2^{\mathcal{N}} \rightarrow \mathbb{R}$, with $v(\emptyset) = 0$. For player $i \in \mathcal{N}$, its Shapley value $\phi_v(i)$ w.r.t. function $v(\cdot)$ is:

$$\phi_v(i) = \frac{1}{|\mathcal{N}|} \sum_{S \subseteq \mathcal{N} \setminus i} \binom{|\mathcal{N}| - 1}{|S|}^{-1} (v(S \cup i) - v(S)), \quad (5)$$

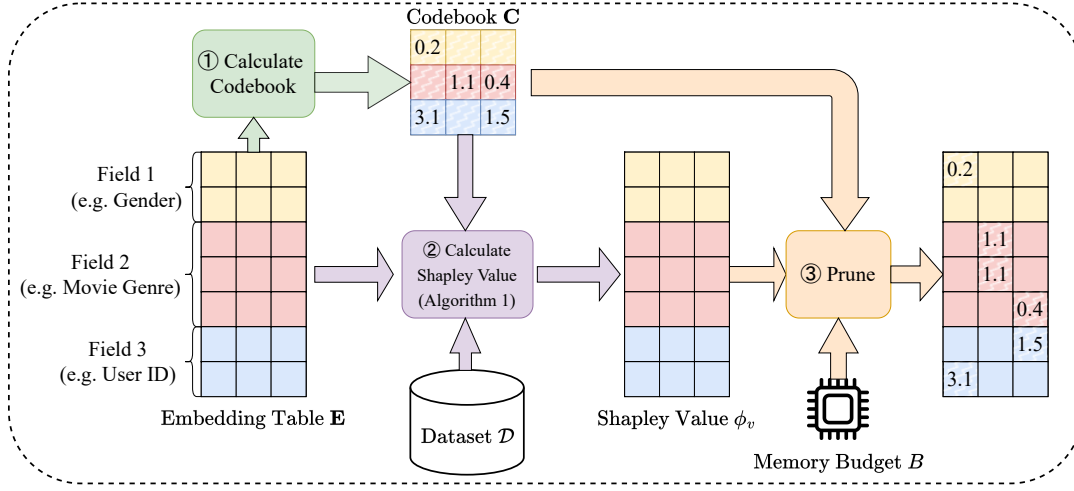


Figure 1: The overview of Shaver. We calculate placeholder values (codebook \mathbf{C}), and then compute Shapley values ϕ_v from the provided dataset \mathcal{D} . On any required memory budget B , we replace embedding parameters with the lowest attribution scores by placeholder values.

where $\binom{\cdot}{\cdot}$ is the binomial coefficient, and $v(\mathcal{S} \cup i) - v(\mathcal{S})$ quantifies the marginal contribution of i after joining team \mathcal{S} . In a nutshell, $\phi_v(i)$ is computed by enumerating over all possible teams $\mathcal{S} \subseteq \mathcal{N} \setminus i$. With that, a commonly used, permutation-based formulation of the Shapley value is:

$$\phi_v(i) = \frac{1}{|\mathcal{N}|!} \sum_{\forall \mathcal{R}} \left(v(\mathcal{P}_i^{\mathcal{R}} \cup i) - v(\mathcal{P}_i^{\mathcal{R}}) \right), \quad (6)$$

where the ordered set \mathcal{R} is a possible permutation of all players, $\mathcal{P}_i^{\mathcal{R}}$ is the set of players that precede i in \mathcal{R} . Thus, the Shapley value $\phi_v(i)$ can be calculated as the mean/expectation of i 's contributions in all $|\mathcal{N}|!$ possible permutations of \mathcal{N} .

Monte Carlo Approximation of Shapley Values. Based on Eq.(6), the default form of Shapley value calculation incurs a factorial time complexity w.r.t. the number of players. To scale it to a large \mathcal{N} , approximating the Shapley value via Monte Carlo approach is a widely adopted workaround [2, 14], which estimates $\phi_v(i)$ by sampling p permutations \mathcal{R}_k and take the average marginal contribution in these samples to approximate Shapley value:

$$\phi_v(i) \approx \frac{1}{p} \sum_{k=1}^p \left(v(\mathcal{P}_i^{\mathcal{R}_k} \cup i) - v(\mathcal{P}_i^{\mathcal{R}_k}) \right), \quad (7)$$

where $\phi_v(i)$ is now the Monte Carlo approximation [5] of player i 's Shapley value w.r.t. value function $v(\cdot)$.

4 Shaver: The Proposed Method

In this section, we unfold the design of Shaver, namely Shapley Value-guided Embedding Reduction for on-device CRSs, of which an overview is provided in Figure 1. We firstly formally define the problem, and subsequently provide corresponding theoretical analyses and pseudocode to efficiently estimate the Shapley value of each embedding parameter. Finally, we derive a codebook to replace the traditional zero-padding strategy, so as to compensate for the information loss brought by embedding parameter pruning and maintain maximum model performance. It is worth noting that, the full training process of $f(\cdot)$, Shapley value estimation, and

codebook computation are all performed on a cloud server, whose role is to find the optimal embedding parameter configuration to meet different parameter budgets for on-device model deployment.

4.1 Shapley Values of Embedding Parameters

Before defining the problem, we provide some clarifications for Shapley value in our specific task context. To facilitate selective embedding table pruning, we derive a Shapley value for every single parameter in the feature embedding table \mathbf{E} . Let $\mathbf{E}[i, c]$ denote the parameter at the i -th row and c -th column (i.e., the c -th embedding dimension of feature i), the Shapley value of player (i, c) is denoted as $\phi_v(i, c)$ with a value function $v(\cdot)$. Please note that the use case of Shapley value in our work differs from the mainstream, data-centric counterparts [23, 34] that focus on calculating values of individual data points or raw features, while we are specifically interested in the Shapley values of model parameters.

Value Function $v(\cdot)$. For measuring $\phi_v(i, c)$, we need to define a value function $v(\cdot)$ that can quantify the contribution from a group of players (i.e., parameters) $\mathcal{S} = \{\mathbf{E}[i_1, c_1], \mathbf{E}[i_2, c_2], \dots, \mathbf{E}[i_s, c_s]\}$. Given a CTR model $f(\cdot)$, we define $v(\cdot)$ as the log loss increase when parameters \mathcal{S} become absent from \mathbf{E} :

$$v(\mathcal{S}) = \frac{1}{|\mathcal{D}|} \sum_{(\mathbf{x}, y) \in \mathcal{D}} \left(L(f_{w/o-\mathcal{S}}(\mathbf{x}), y) - L(f(\mathbf{x}), y) \right), \quad (8)$$

where $f_{w/o-\mathcal{S}}(\cdot)$ denotes the CTR model without embedding parameters \mathcal{S} , L is the log loss, and \mathcal{D} denotes the dataset. With this definition, the Shapley value $\phi_v(i, c)$ of a single parameter $\mathbf{E}[i, c]$ can be quantified via an analogous process as in Section 3.2.

Definition 4.1 (Single-shot Pruning with Shapley Values). Given a trained CTR model $f(\cdot)$, our goal is to obtain the Shapley value of every single parameter $\phi_v(i, c)$ in the feature embedding table. Based on $\phi_v(i, c)$, we can efficiently prune a CTR model for any on-device budget by directly removing the lowest ranked entries from its embedding table to meet the specified parameter size.

4.2 Local and Global Values

If straightforward Monte Carlo approximation of $\phi_v(i, c)$ is in use, for all $|\mathcal{N}| = nd$ embedding parameters, $|\mathcal{N}||\mathcal{D}| = nd|\mathcal{D}|$ forward passes are needed per permutation, leading to a prohibitive computational cost of $O(pnd|\mathcal{D}|)$. One main cause is that Eq.(8) lays emphasis on embedding parameters' values to the overall data distribution. We term this *global values*, whose calculation involves a sufficiently large, if not the full dataset \mathcal{D} . With that said, the function of *local values* $v^{loc}(\cdot)$ is defined based on single data samples:

$$v^{loc}(\mathcal{S}|\mathbf{x}, y) = L(f_{w/o-\mathcal{S}}(\mathbf{x}), y) - L(f(\mathbf{x}), y). \quad (9)$$

In contrast to computing the global values, there are some efforts that draw an individual data sample at a time, pair it with one random permutation, and then derive each parameter's local value to approximate the global one [12]. This process can be described as follows:

$$\begin{aligned} \phi_v(i, c) &= \mathbb{E}_{\mathcal{R}} \left[v(\mathcal{P}_{i,c}^{\mathcal{R}} \cup (i, c)) - v(\mathcal{P}_{i,c}^{\mathcal{R}}) \right] \\ &= \mathbb{E}_{\mathcal{R}} \mathbb{E}_{(\mathbf{x}, y)} \left[v^{loc}(\mathcal{P}_{i,c}^{\mathcal{R}} \cup (i, c)|\mathbf{x}, y) - v^{loc}(\mathcal{P}_{i,c}^{\mathcal{R}}|\mathbf{x}, y) \right] \\ &= \mathbb{E}_{(\mathbf{x}, y), \mathcal{R}} \left[L(f_{w/o-\mathcal{S}_1}(\mathbf{x}), y) - L(f_{w/o-\mathcal{S}_2}(\mathbf{x}), y) \right], \\ &= \frac{1}{|\mathcal{D}|} \sum_{(\mathbf{x}, y) \in \mathcal{D}, \mathcal{R}(\mathbf{x}, y)} \left(L(f_{w/o-\mathcal{S}_1}(\mathbf{x}), y) - L(f_{w/o-\mathcal{S}_2}(\mathbf{x}), y) \right), \end{aligned} \quad (10)$$

with \mathcal{S}_1 and \mathcal{S}_2 defined as:

$$\begin{aligned} \mathcal{S}_1 &= \{E[i', c'] \mid (i', c') \in \mathcal{P}_{i,c}^{\mathcal{R}} \cup (i, c)\}, \\ \mathcal{S}_2 &= \{E[i', c'] \mid (i', c') \in \mathcal{P}_{i,c}^{\mathcal{R}}\}, \end{aligned} \quad (11)$$

where $\mathcal{R}(\mathbf{x}, y)$ denotes one random permutation paired with data sample (\mathbf{x}, y) . In a nutshell, Eq.(10) shows that, with a quality, decent-sized \mathcal{D} , the global Shapley value can be confidently approximated as the mean of local values [12]. For all parameters, this method runs in $O(nd|\mathcal{D}|)$. Despite being p times faster than vanilla Monte Carlo and a subsampled \mathcal{D} from the full training dataset can help further lower the complexity, the large feature size n (normally in million-level) remains the efficiency bottleneck for parameter attribution. The following section shows how we further reduce those $nd|\mathcal{D}|$ forward passes – which is essentially iterating over the dataset for nd times – to a much smaller number.

4.3 Efficient Shapley Value Computation for CRS Embeddings

Recall that in CRSs, features are categorized into m fields. On top of the above defined local value $v^{loc}(\cdot)$, we further define a local value function $u^{loc}(\cdot)$ specific to feature fields:

$$u^{loc}(\mathcal{S}'|\mathbf{x}, y) = L(f_{w/o-\mathcal{S}'}(\mathbf{x}), y) - L(f(\mathbf{x}), y), \quad (12)$$

Distinct from $v^{loc}(\cdot)$, in $u^{loc}(\cdot)$, a player (j, c) is a *group of parameters* that are located at the same column c in field j 's embedding table \mathbf{V}_j . To avoid cluttered notations, in this section, if not specified, both $v(\cdot)$ and $u(\cdot)$ are the *local value functions* calculated with a single instance (\mathbf{x}, y) . Then, we define our notion of corresponding player as follows.

Definition 4.2 (Corresponding Player). Given player (i, c) in $v(\cdot)$ and player (j, c') in $u(\cdot)$, and a data instance (\mathbf{x}, y) , we call (i, c)

Algorithm 1: Efficient Shapley Value Approximation for Embedding Parameters

Data: A trained CTR model $f(\cdot)$, dataset \mathcal{D} , local value function $u(\cdot)$, number of feature fields m , number of features n , embedding dimension d

Result: Shapley value of each player $\{\phi(i, c)\}_{i=1, c=1}^{n, d}$

```

1  $\{\phi(i, c)\}_{i=1, c=1}^{n, d} \leftarrow 0$ ;
2 for  $(\mathbf{x}, y) \in \mathcal{D}$  do
3    $\mathcal{R} \leftarrow$  a random order of players  $\{(j, c)\}_{j=1, c=1}^{m, d}$ ;
4    $\mathcal{S} \leftarrow \emptyset$ ;
5    $u_{prev} \leftarrow u(\mathcal{S}|\mathbf{x}, y)$ ;
6   for  $(j, c) \in \mathcal{R}$  do
7      $\mathcal{S} \leftarrow \mathcal{S} \cup (j, c)$ ;
8      $u_{next} \leftarrow u(\mathcal{S}|\mathbf{x}, y)$ ;
9     Find  $(i, c)$  as  $(j, c)$ 's corresponding player w.r.t.  $(\mathbf{x}, y)$ ;
10     $\phi(i, c) \leftarrow \phi(i, c) + u_{next} - u_{prev}$ ;
11     $u_{prev} \leftarrow u_{next}$ ;
12  end
13 end
14  $\{\phi(i, c)\}_{i=1, c=1}^{n, d} \leftarrow \{\frac{1}{|\mathcal{D}|} \phi(i, c)\}_{i=1, c=1}^{n, d}$ 
```

the *corresponding player* of (j, c') if $c = c'$ and the i -th feature is activated, i.e., $\mathbf{x}[i] = 1$.

With the above definition, we provide a key theorem for efficient Shapley value computation below:

THEOREM 4.3. Assume player (i, c) in $v(\cdot)$ is corresponding with player (j, c) in $u(\cdot)$ w.r.t. data instance (\mathbf{x}, y) . Then, denote their Shapley values based on local value functions respectively as $\phi_v^{x, y}(i, c)$ and $\phi_u^{x, y}(j, c)$, we have $\phi_v^{x, y}(i, c) = \phi_u^{x, y}(j, c)$.

Our proof is provided in Appendix A. In short, we observe that $\phi_v^{x, y}(i, c)$ is equivalent to $\phi_u^{x, y}(j, c)$, as long as (i, c) is the corresponding player of (j, c) . As computing $\phi_u^{x, y}(j, c)$ will suffice, only md rather than nd iterations over \mathcal{D} are now required to compute $\phi_u^{x, y}(j, c)$ for all players in $u(\cdot)$. Taking our largest evaluation dataset KDD as an example (see Table 1), with 11 (m) feature fields, over 6 million (n) features, and an embedding dimension of $d = 16$, the $O(md|\mathcal{D}|)$ complexity of our Shapley value computation over $O(nd|\mathcal{D}|)$ given that $m \ll n$.

The pseudo code is shown in Algorithm 1. For each data point, we first sample a random order \mathcal{R} for set of players in $u(\cdot)$ (line 3). We initialize an empty set to track the selected players (line 4) and record the value generated by local value function $u(\cdot)$ for the empty set (line 5). Then, for each player in the randomly ordered list, we add them to the selection set \mathcal{S} (line 7) and recalculate the value with this new team \mathcal{S} (line 8). With Theorem 4.3, we update the corresponding player (i, c) 's contribution accordingly (lines 9-10), and move on to the next player in the permutation. After iterating over all data points, the algorithm normalizes accumulated Shapley values with the total number of data samples in the dataset (line 14). To validate the low variance of our approximation method, an algorithm convergence rate analysis is provided in Appendix B.

4.4 Replacing Zero-padding with Field-aware Codebook

In various machine learning fields like explainable AI, the choice of placeholder value (i.e. the value representing the missing value) is crucial, and can greatly affect the final performance [6]. However, most if not all CRS pruning algorithms opt for a simple zero-padding strategy. This approach often leads to suboptimal recommendation performance, as interactions between feature embeddings are modeled using dot product and element-wise operations. Setting certain parameters to zero eliminates such interactions with all other features, compromising system effectiveness. Therefore, in this work, we introduce a novel approach to represent pruned features by a field-aware codebook $C \in \mathbb{R}^{m \times d}$. As it only comes with md elements, its storage overhead is negligible. For a parameter budget, let E_Q denote the sparsified embedding table that replaces the least important parameters identified in set Q , the following equation specifies how we impute the pruned embedding table into $E_{Q,C}$ with the codebook:

$$E_{Q,C}[i, c] = \begin{cases} E_Q[i, c], & \text{if } \{i, c\} \notin Q \\ C[j, c], & \text{else} \end{cases}, \quad (13)$$

with j the corresponding field of feature i . After calculating the attribution score, we will replace pruned parameters based on the above rule instead of only zeroing out.

To better preserve the model performance, we will choose codebook C^* that minimizes the expected Euclidean distance between the original embedding table and its imputed version:

$$C^* = \min_C \mathbb{E}_{\mathbf{x} \sim \mathcal{D}, |Q|=B} \left[\|\mathbf{E}^\top \mathbf{x} - \mathbf{E}_{Q,C}^\top \mathbf{x}\|_2^2 \right], \quad (14)$$

where $B \in (0, nd)$ represents any possible number of pruned parameters, and Q is uniformly sampled such that $|Q| = B$. With this formulation, the optimal codebook $C^* \in \mathbb{R}^{m \times d}$ can be solved via the weighted average of all feature in each field:

$$C^*[j, :] = \frac{\sum_{i \in \mathcal{F}_j} p_i \mathbf{V}_j[i, :]}{\sum_{i \in \mathcal{F}_j} p_i}, \quad (15)$$

where \mathcal{F}_j is the set of features in field j , and p_i is the occurrence frequency of feature i . The detailed derivation is provided in Appendix A. This closed-form solution allows for efficient computation of codebook C^* to compensate for the information loss from the pruning stage. When deployed on-device, both the pruned sparse embedding table E_Q and the codebook C^* are stored. When inference is needed, only the m embeddings $\{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_m\}$ of the feature \mathbf{x} will be imputed as the input to $f(\cdot)$ for on-device CTR prediction.

5 Experiments

In this section, we conduct experiments to study the effectiveness of Shaver. Specifically, we are interested in answering the following research questions (RQs):

- RQ1:** Compared with non-single-shot baselines, how does Shaver perform when compressing CRS models?
- RQ2:** Compared with other single-shot baselines, how does Shaver perform under different parameter budgets?
- RQ3:** In Shaver, the data \mathcal{D} used for estimating Shapley values can be subsampled to improve efficiency. What is the impact

Table 1: Statistics of the preprocessed datasets.

Name	#Instances	#Features	#Fields
Criteo	45,840,617	1,086,810	39
Avazu	40,428,967	4,428,511	22
KDD	149,639,105	6,019,761	11

of dataset size to the recommendation performance of the pruned models?

RQ4: Most on-device CRS compression methods require model retraining/finetuning post pruning. Does Shaver also benefit from such a practice?

RQ5: Does Shaver exhibit any patterns when pruning feature embeddings?

5.1 Experimental Settings

5.1.1 Datasets. We conduct our experiments on three public datasets. In all datasets, we randomly split them into 8:1:1 as the training, validation and test set respectively. We specify the dataset preprocess procedure in Appendix C.2. Table 1 provides the core statistics of pre-processed datasets.

5.1.2 Metrics. We evaluate all models using two commonly used evaluation metric in CTR prediction community [66]: LogLoss and AUC (Area under the ROC curve). In the CTR problem, a difference of **0.001** in AUC is generally considered significant [35, 52]. The lower LogLoss suggests better performance, while the higher AUC implies more accurate recommendations. The sparsity rate t indicates the portion of embedding parameters that have been removed compared to the original embedding table, defined as $t = |Q|/nd$.

5.1.3 Implementation Details. All methods are tested with two well-known CTR backbones: DeepFM [18] and DCN-Mix [53]. We implement two variants of Shaver, namely Shaver-Zero that uses the traditional zero-padding on pruned parameters, and Shaver-Codebook described in Section 4.4. Descriptions of all baselines are provided in Appendix C.1. We adhere strictly to the experimental settings described in [52] for reproducibility, adopting the same data splits and hyperparameter search ranges. To ensure a fair comparison, we applied the same hyperparameters used in the training of the original model during our fine-tuning step. We used both the training and validation datasets to calculate the attribution scores. We select final checkpoints according to validation AUC.

5.2 Comparison with Non-single-shot Pruning Methods (RQ1)

Table 2 shows the performance of Shaver and other non-single-shot pruning methods with the DCN-Mix backbone. The baselines compared are QR [49], TTRec [59], PEP [32], OptEmb [35], and CERP [31], whose detailed descriptions are provided in Appendix C.1. As DCN-Mix is generally a more performant backbone compared with DeepFM, results with the DeepFM backbone is deferred to Appendix D due to page limits. For this comparison, we test three sparsity rate settings with $t = \{50\%, 80\%, 95\%\}$ as all baselines under this category needs to be retrained for every setting.

Across three datasets, Shaver achieves competitive results with other non-single-shot pruning methods that require costly training steps for each parameter budget. Shaver demonstrates a strong

Table 2: Comparative results with non-single-shot baselines under three sparsity rates. All methods adopt DCN-Mix as the CRS backbone, where we use \uparrow and \downarrow to respectively mark higher-is-better and lower-is-better metrics. For OptEmb, we did not report results on some settings where it fails to converge. “#Params” indicates the parameter size of the embedding table (in millions). In each dataset, the best result is marked in bold and the second best one is underlined.

Dataset	Method	Single-shot?	$t = 50\%$			$t = 80\%$			$t = 95\%$		
			AUC \uparrow	LogLoss \downarrow	#Params	AUC \uparrow	LogLoss \downarrow	#Params	AUC \uparrow	LogLoss \downarrow	#Params
Criteo	QR [49]	\times	0.8095	0.4425	8.69M	0.8091	0.4427	3.48M	0.8064	0.4452	870K
	TTRec [59]	\times	0.8112	<u>0.4408</u>	8.26M	0.8103	0.4416	3.48M	0.8106	0.4414	870K
	PEP [32]	\times	<u>0.8113</u>	0.4407	8.69M	0.8108	0.4413	3.46M	0.8114	0.4406	831K
	OptEmb [35]	\times	-	-	-	<u>0.8112</u>	<u>0.4408</u>	1.45M	0.8109	<u>0.4408</u>	856K
	CERP [31]	\times	0.8110	0.4410	8.66M	0.8106	0.4416	3.44M	0.8107	0.4414	850K
	Shaver-Zero	\checkmark	0.8114	0.4407	8.69M	0.8114	0.4407	3.48M	0.8110	0.4412	869K
	Shaver-Codebook	\checkmark	0.8114	0.4407	8.69M	0.8114	0.4407	3.48M	<u>0.8112</u>	0.4410	870K
Avazu	QR [49]	\times	0.7743	0.3857	35.43M	0.7698	0.3922	14.17M	0.7652	0.3895	3.54M
	TTRec [59]	\times	0.7711	0.3873	35.43M	0.7710	0.3900	14.17M	0.7678	0.3903	3.54M
	PEP [32]	\times	0.7754	0.3852	33.55M	0.7669	0.3906	12.57M	0.7564	0.3938	3.52M
	OptEmb [35]	\times	-	-	-	0.7663	0.3877	13.24M	<u>0.7683</u>	0.3895	3.43M
	CERP [31]	\times	0.7723	0.3852	35.34M	0.7673	0.3873	13.05M	<u>0.7631</u>	0.3933	3.54M
	Shaver-Zero	\checkmark	0.7759	0.3837	35.43M	0.7745	0.3840	14.17M	0.7690	0.3865	3.54M
	Shaver-Codebook	\checkmark	0.7762	<u>0.3838</u>	35.43M	0.7749	<u>0.3842</u>	14.17M	<u>0.7683</u>	<u>0.3873</u>	3.54M
KDD	QR [49]	\times	0.7724	0.1581	48.16M	0.7709	0.1584	19.26M	0.7702	0.1588	4.82M
	TTRec [59]	\times	0.7862	0.1576	48.16M	<u>0.7843</u>	<u>0.1564</u>	19.26M	0.7815	0.1559	4.82M
	PEP [32]	\times	0.7754	0.1576	45.93M	<u>0.7731</u>	<u>0.1581</u>	18.88M	0.7737	0.1579	4.50M
	OptEmb [35]	\times	0.7877	0.1558	48.99M	0.7818	0.1593	15.10M	0.7826	0.1591	4.80M
	CERP [31]	\times	0.7841	0.1566	47.50M	0.7781	0.1573	19.18M	<u>0.7721</u>	0.1589	4.68M
	Shaver-Zero	\checkmark	0.7858	0.1563	48.16M	0.7826	0.1569	19.26M	0.7699	0.1601	4.82M
	Shaver-Codebook	\checkmark	<u>0.7870</u>	<u>0.1559</u>	48.16M	0.7870	0.1557	19.26M	0.7836	<u>0.1563</u>	4.82M

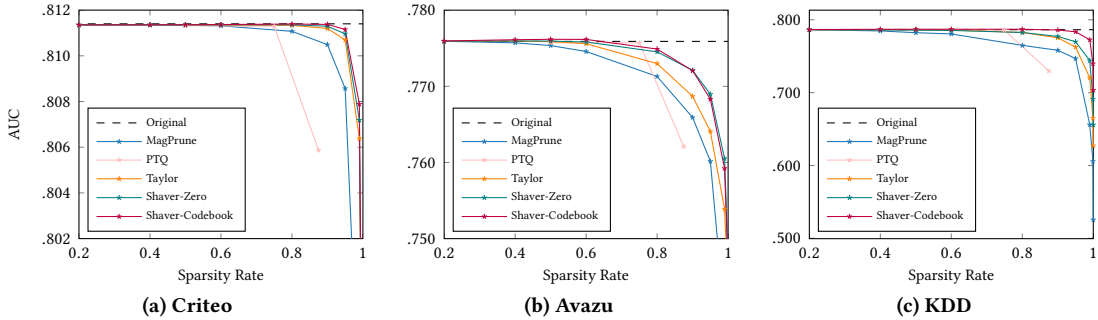


Figure 2: Comparative results with single-shot baselines, where DCN-Mix is used as the backbone.

performance compared to other methods, especially with an 80% sparsity rate. While in the low sparsity rate, the performance gap between methods is small as it is generally easier to reduce parameters. At 95% sparsity, Shaver delivers impressive results exceeding other methods, despite a lesser extent. Higher sparsity rates pose challenges for single-shot pruning due to removing a larger portion of original model parameters, thus, leading to a more significant model degradation.

5.3 Comparison with Single-shot Pruning Methods (RQ2)

To make further comparisons, we evaluate our approach against other single-shot compression methods. The baselines compared are MagPrune [52], PTQ [40], and Taylor [38] (see Appendix C.1 for

descriptions). Same with RQ1, DCN-Mix backbone is adopted by default, where results with DeepFM backbone are provided in Appendix D. Notably, different from RQ1, as all baselines tested here are all single-shot methods, we have been able to test with a wider variety of sparsity rates $t = \{0.2, 0.4, 0.5, 0.6, 0.8, 0.9, 0.95, 0.99, 0.999\}$.

Figure 2 shows the results. Shaver-Codebook is the most competitive method for KDD and Criteo, retaining the performance even with only 10% and 5% parameters of the original models, respectively. The second best is Shaver-Zero, whose performance is similar to Shaver-Codebook in the Avazu dataset. These results indicate promising potential in applying the Shapley value in compressing models for CTR problems. Taylor is the third best as they also utilize information from the training and validation datasets, especially in the Criteo dataset, where the difference is only noticeable under high-sparsity settings. PTQ provides strong performance initially

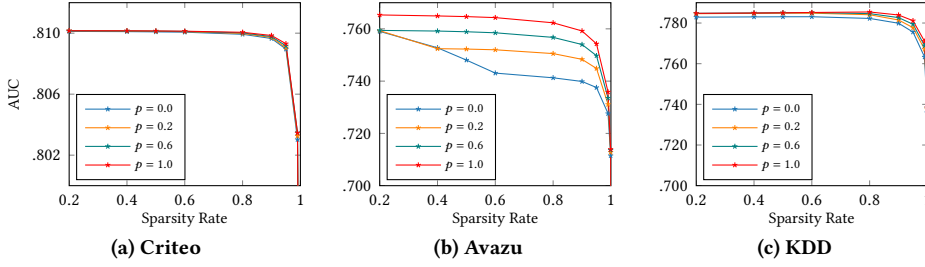


Figure 3: Trade-off between performance and data size used to estimate Shapley value. p is the portion of the full dataset used to calculate the Shapley value.

but could not hold up at a higher sparsity rate, as it allocates the same memory budget for all features, limiting their performance.

5.4 Effect of Post-Pruning Finetuning (RQ3)

To answer RQ3, Table 3 shows the effect of fine-tuning pruned models. This experiment also mimics scenarios where a deployed recommendation model has sufficient computing resources to perform further on-device updates. While there is some variation in results across different datasets, in most cases, fine-tuning models further enhance their performance or maintain it at its current level, or with minimal loss. However, occasionally, fine-tuning can lead to decreased model performance, particularly for some settings in KDD. The reason for the performance drop is that Shaver uses the validation set to estimate Shapley values, which can cause data leakage and potentially degrade the final checkpoint performance post-fine-tuning. Moreover, fine-tuning seems to have a larger effect on high sparsity rates with Shaver-Zero, where performance degradation is more pronounced.

5.5 Performance and Data Size Trade-off (RQ4)

Figure 3 depicts the trade-off between the performance and the proportion p of the full dataset \mathcal{D} employed. Note that \mathcal{D} combines both training and validation samples in our setting. Under all configurations, we guarantee the algorithms have converged and seen all data. Empirically, the required runtime is linearly scaled

Table 3: Effect of post-pruning finetuning, where \dagger indicates variants that are finetuned with the pruned embeddings.

Dataset	Method	$t = 50\%$	$t = 80\%$	$t = 95\%$
Criteo	Shaver-Zero	0.8114	0.8114	0.8110
	Shaver-Zero \dagger	0.8116	0.8114	0.8112
	Shaver-Codebook	0.8114	0.8114	0.8112
	Shaver-Codebook \dagger	0.8115	0.8114	0.8112
Avazu	Shaver-Zero	0.7759	0.7745	0.7690
	Shaver-Zero \dagger	0.7777	0.7735	0.7703
	Shaver-Codebook	0.7762	0.7749	0.7683
	Shaver-Codebook \dagger	0.7781	0.7717	0.7670
KDD	Shaver-Zero	0.7858	0.7826	0.7699
	Shaver-Zero \dagger	0.7858	0.7826	0.7775
	Shaver-Codebook	0.7870	0.7870	0.7836
	Shaver-Codebook \dagger	0.7754	0.7806	0.7836




Figure 4: Comparison of how different Shaver variants prune embeddings for features in different frequency buckets.

with the dataset size employed, as Shaver generally converges after iterating through the whole dataset once.

In the Criteo and KDD datasets, the amount of data employed has a negligible effect on model performance. Conversely, in the Avazu dataset, the volume of data has a markedly greater impact. One possible hypothesis for this discrepancy is that despite containing more features than the Criteo dataset, the Avazu dataset has fewer instances. Consequently, the validation dataset is insufficient to represent the actual data distribution. Nonetheless, as our methods only require forward passes, we can effortlessly parallelize the computation compared to other training-required methods. Moreover, the overhead remains minimal, even with the entire Criteo dataset – our worst case as our computation cost increases with the number of fields – Shaver takes approximately 2.5 hours to compute the Shapley value. In comparison, for any given parameter budget, training each method generally takes over 1 hour, without including the time needed for hyperparameter tuning.

5.6 Qualitative Analysis (RQ5)

We aim to unveil the key difference between the traditional zero-out approach and our codebook-based approach in compressed models. Therefore, employing DCN-Mix with a 50% sparsity rate, we split each dataset’s feature into three different bins based on frequency so that each bin had the same number of features. Figure 4 shows the average embedding size assigned per feature in each group.

The codebook approach assigns less parameter budget to high frequency. One plausible explanation is that the codebook already contained parameters representing the high-frequency features. Thus, the model could have a larger budget to allocate to the lower frequency features, improving the model’s performance.

6 Conclusions

In this paper, we propose **Shapley Value-guided Embedding Reduction** (Shaver), which employs Shapley value to fairly distribute each embedding parameter contribution and effectively prune the embedding table in CTR models. Shaver first provide an efficient method to calculate the Shapley value for CTR models, and then implement a novel field-aware codebook quantization to represent the removed parameters. Extensive experiments show that Shaver achieves competitive performance without expensive retraining and fine-tuning steps for each parameter budget.


References

- [1] Yi Wang Aden. 2012. KDD Cup 2012, Track 2. <https://kaggle.com/competitions/kddcup2012-track2>
- [2] Marco Ancona, Cengiz Öztireli, and Markus Gross. 2020. Shapley value as principled metric for structured network pruning. *arXiv preprint arXiv:2006.01795* (2020).
- [3] Omer Ben-Porat and Moshe Tennenholtz. 2018. A game-theoretic approach to recommendation systems with strategic content providers. *NeurIPS* (2018).
- [4] Han Cai, Chuang Gan, Tianzhe Wang, Zhekai Zhang, and Song Han. 2020. Once-for-all: Train one network and specialize it for efficient deployment. *ICLR* (2020).
- [5] Javier Castro, Daniel Gómez, and Juan Tejada. 2009. Polynomial calculation of the Shapley value based on sampling. *Computers & operations research* (2009).
- [6] Hugh Chen, Ian C Covert, Scott M Lundberg, and Su-In Lee. 2023. Algorithms to estimate Shapley value feature attributions. *Nature Machine Intelligence* 5, 6 (2023), 590–601.
- [7] Tong Chen, Hongzhi Yin, Yujia Zheng, Zi Huang, Yang Wang, and Meng Wang. 2021. Learning Elastic Embeddings for Customizing On-Device Recommenders. In *SIGKDD*. 138–147.
- [8] 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*.
- [9] Benjamin Coleman, Wang-Cheng Kang, Matthew Fahrback, Ruoxi Wang, Lichan Hong, Ed H. Chi, and Derek Zhiyuan Cheng. 2023. Unified Embedding: Battle-tested feature representations for web-scale ML systems. In *NeurIPS*.
- [10] Ian Covert, Chanwoo Kim, Su-In Lee, James Zou, and Tatsunori Hashimoto. 2024. Stochastic Amortization: A Unified Approach to Accelerate Feature and Data Attribution. *arXiv preprint arXiv:2401.15866* (2024).
- [11] Ian Covert and Su-In Lee. 2021. Improving kernelshap: Practical shapley value estimation via linear regression. *Artificial Intelligence and Statistics* (2021).
- [12] Ian Covert, Scott M Lundberg, and Su-In Lee. 2020. Understanding global feature contributions with additive importance measures. *NeurIPS* (2020).
- [13] Yunfei Fang, Caihong Mu, and Yi Liu. 2023. AutoShape: Automatic Design of Click-Through Rate Prediction Models Using Shapley Value. In *Pacific Rim International Conference on Artificial Intelligence*.
- [14] Amirata Ghorbani and James Zou. 2020. Neuron Shapley: discovering the responsible neurons. In *NeurIPS*.
- [15] Xudong Gong et al. 2022. Real-time short video recommendation on mobile devices. In *CIKM*. 3103–3112.
- [16] Hui Guan, Andrey Malevich, Jiyan Yang, Jongsoo Park, and Hector Yuen. 2019. Post-training 4-bit quantization on embedding tables. *arXiv preprint arXiv:1911.02079* (2019).
- [17] Jiyang Guan, Zhuozhuo Tu, Ran He, and Dacheng Tao. 2022. Few-shot backdoor defense using shapley estimation. In *CVPR*.
- [18] Huifeng Guo, Ruiming Tang, Yunming Ye, Zhenguo Li, and Xiuqiang He. 2017. DeepFM: A Factorization-Machine based Neural Network for CTR Prediction. In *IJCAI, Carles Sierra* (Ed.). 1725–1731.
- [19] Jialiang Han, Yun Ma, Qiaozhu Mei, and Xuanzhe Liu. 2021. Deeprec: On-device deep learning for privacy-preserving sequential recommendation in mobile commerce. In *WWW*. 900–911.
- [20] Song Han, Huizi Mao, and William J. Dally. 2016. Deep Compression: Compressing Deep Neural Network with Pruning, Trained Quantization and Huffman Coding. In *ICLR, Yoshua Bengio and Yann LeCun* (Eds.).
- [21] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural collaborative filtering. In *WWW*. 173–182.
- [22] Olivier Chapelle Jean-Baptiste Tien, joycenv. 2014. Display Advertising Challenge. <https://kaggle.com/competitions/criteo-display-ad-challenge>
- [23] Neil Jethani, Mukund Sudarshan, Ian Connick Covert, Su-In Lee, and Rajesh Ranganath. 2021. FastSHAP: Real-time shapley value estimation. In *ICLR*.
- [24] Manas R. Joglekar, Cong Li, Jay K. Adams, Pranav Khaitan, and Quoc V. Le. 2019. Neural Input Search for Large Scale Recommendation Models. *SIGKDD* (2019). <https://api.semanticscholar.org/CorpusID:195874115>
- [25] Wang-Cheng Kang et al. 2021. Learning to Embed Categorical Features without Embedding Tables for Recommendation. In *SIGKDD*. 840–850.
- [26] Shuming Kong, Weiye Cheng, Yanyan Shen, and Linpeng Huang. 2022. Autosrh: An embedding dimensionality search framework for tabular data prediction. *IEEE Transactions on Knowledge and Data Engineering* 35, 7 (2022), 6673–6686.
- [27] Namhoon Lee, Thalaiyasingam Ajanthan, and Philip Torr. 2018. SNIP: Single-shot network pruning based on connection sensitivity. In *ICLR*.
- [28] Shiwei Li et al. 2023. Adaptive Low-Precision Training for Embeddings in Click-through Rate Prediction. In *AAAI*.
- [29] Shiwei Li et al. 2024. Embedding Compression in Recommender Systems: A Survey. *ACM Comput. Surv.* (jan 2024).
- [30] Xurong Liang, Tong Chen, Lizhen Cui, Yang Wang, Meng Wang, and Hongzhi Yin. 2024. Lightweight Embeddings for Graph Collaborative Filtering. In *SIGIR*. 1296–1306.
- [31] Xurong Liang, Tong Chen, Quoc Viet Hung Nguyen, Jianxin Li, and Hongzhi Yin. 2023. Learning Compact Compositional Embeddings via Regularized Pruning for Recommendation. *arXiv:2309.03518 [cs.IR]*
- [32] Siyi Liu, Chen Gao, Yihong Chen, Depeng Jin, and Yong Li. 2021. Learnable Embedding Sizes for Recommender Systems. In *ICLR*. <https://openreview.net/forum?id=vQzcqQWIS0q>
- [33] Scott M Lundberg, Gabriel Erion, Hugh Chen, Alex DeGrave, Jordan M Prutkin, Bala Nair, Ronit Katz, Jonathan Himmelfarb, Nisha Bansal, and Su-In Lee. 2020. From local explanations to global understanding with explainable AI for trees. *Nature machine intelligence* (2020).
- [34] Scott M Lundberg and Su-In Lee. 2017. A unified approach to interpreting model predictions. *NeurIPS* (2017).
- [35] Fuyuan Lyu et al. 2022. OptEmbed: Learning Optimal Embedding Table for Click-through Rate Prediction. In *CIKM*. 1399–1409.
- [36] Matteo Marcuzzo, Alessandro Zangari, Andrea Albarelli, and Andrea Gasparetto. 2022. Recommendation systems: An insight into current development and future research challenges. *IEEE Access* 10 (2022), 86578–86623.
- [37] Lorenzo Minto and Moritz Haller. 2021. Using Federated Learning to Improve Brave’s On-Device Recommendations While Protecting Your Privacy. (2021). <https://brave.com/blog/federated-learning/>
- [38] Pavlo Molchanov, Stephen Tyree, Tero Karras, Timo Aila, and Jan Kautz. 2017. Pruning Convolutional Neural Networks for Resource Efficient Inference. In *ICLR*.
- [39] 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.
- [40] Markus Nagel, Marios Fournarakis, Rana Ali Amjad, Yelysei Bondarenko, Mart Van Baalen, and Tijmen Blankevoort. 2021. A white paper on neural network quantization. *arXiv preprint arXiv:2106.08295* (2021).
- [41] Niketan Pansare et al. 2022. Learning compressed embeddings for on-device inference. *Proceedings of Machine Learning and Systems* 4 (2022), 382–397.
- [42] Liang Qu, Yonghong Ye, Ningzhi Tang, Lixin Zhang, Yuhui Shi, and Hongzhi Yin. 2022. Single-shot Embedding Dimension Search in Recommender System. In *SIGIR*. 513–522.
- [43] Yunke Qu, Tong Chen, Quoc Viet Hung Nguyen, and Hongzhi Yin. 2024. Budgeted Embedding Table For Recommender Systems. In *WSDM (WSDM ’24)*. 557–566.
- [44] Yunke Qu, Tong Chen, Xiangyu Zhao, Lizhen Cui, Kai Zheng, and Hongzhi Yin. 2023. Continuous Input Embedding Size Search For Recommender Systems. In *SIGIR*. 708–717.
- [45] Yunke Qu, Liang Qu, Tong Chen, Xiangyu Zhao, Quoc Viet Hung Nguyen, and Hongzhi Yin. 2024. Scalable Dynamic Embedding Size Search for Streaming Recommendation. *arXiv preprint arXiv:2407.15411* (2024).
- [46] Steffen Rendle. 2010. Factorization Machines. In *ICDM*. 995–1000. <https://doi.org/10.1109/ICDM.2010.127>
- [47] Alvin E Roth. 1988. *The Shapley value: essays in honor of Lloyd S. Shapley*. Cambridge University Press.
- [48] Benedek Rozemberczki and Rik Sarkar. 2021. The Shapley Value of Classifiers in Ensemble Games. In *CIKM*. 1558–1567.
- [49] Hao-Jun Michael Shi, Dheevatsa Mudigere, Maxim Naumov, and Jiyan Yang. 2020. Compositional Embeddings Using Complementary Partitions for Memory-Efficient Recommendation Systems. In *KDD*. 165–175.
- [50] Statista. 2024. Online advertising revenue in the United States from 2000 to 2023. (2024). <https://www.statista.com/statistics/183816/us-online-advertising-revenue-since-2000/>
- [51] Will Cukierski Steve Wang. 2014. The Avazu Dataset. <https://kaggle.com/competitions/avazu-ctr-prediction>
- [52] Hung Vinh Tran, Tong Chen, Quoc Viet Hung Nguyen, Zi Huang, Lizhen Cui, and Hongzhi Yin. 2024. A Thorough Performance Benchmarking on Lightweight Embedding-based Recommender Systems. *arXiv preprint arXiv:2406.17335* (2024).
- [53] Ruoxi Wang et al. 2021. Dcn v2: Improved deep & cross network and practical lessons for web-scale learning to rank systems. In *WWW*. 1785–1797.
- [54] Shuyao Wang, Yongduo Sui, Jiancan Wu, Zhi Zheng, and Hui Xiong. 2024. Dynamic Sparse Learning: A Novel Paradigm for Efficient Recommendation. In *WSDM*. 740–749.
- [55] Yejing Wang, Zhaocheng Du, Xiangyu Zhao, Bo Chen, Huifeng Guo, Ruiming Tang, and Zhenhua Dong. 2023. Single-shot feature selection for multi-task recommendations. In *SIGIR*. 341–351.
- [56] Zhiqiang Xu, Dong Li, Weijie Zhao, Xing Shen, Tianbo Huang, Xiaoyun Li, and Ping Li. 2021. Agile and accurate CTR prediction model training for massive-scale online advertising systems. In *Proceedings of the International Conference on Management of Data (SIGMOD)*.
- [57] 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 *CIKM*.
- [58] Jie Amy Yang, Jianyu Huang, Jongsoo Park, Ping Tak Peter Tang, and Andrew Tulloch. 2020. Mixed-precision embedding using a cache. *arXiv preprint arXiv:2010.11305* (2020).
- [59] Chunxing Yin, Bilge Acun, Xing Liu, and Carole-Jean Wu. 2021. TT-Rec: Tensor Train Compression for Deep Learning Recommendation Models. *arXiv abs/2101.11714* (2021). <https://api.semanticscholar.org/CorpusID:231719841>

- [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] Jiahui Yu, Linjie Yang, Ning Xu, Jianchao Yang, and Thomas Huang. 2019. Slimmable neural networks. *ICLR* (2019).
- [62] Hailin Zhang et al. 2023. Experimental Analysis of Large-scale Learnable Vector Storage Compression. *arXiv:2311.15578* [cs.LG]
- [63] Xiangyu Zhao et al. 2021. AutoDim: Field-aware Embedding Dimension Search in Recommender Systems. In *WWW*. 3015–3022.
- [64] Xiangyu Zhao et al. 2021. AutoEmb: Automated Embedding Dimensionality Search in Streaming Recommendations. In *ICDM*. 896–905.
- [65] Jieming Zhu et al. 2022. BARS: Towards Open Benchmarking for Recommender Systems. In *SIGIR*. 2912–2923.
- [66] Jieming Zhu, Jinyang Liu, Shuai Yang, Qi Zhang, and Xiuqiang He. 2021. Open benchmarking for click-through rate prediction. In *CIKM*. 2759–2769.

A Proofs

Figure 5: A graphic explanation on the notion of corresponding player.



THEOREM 4.3. Assume player (i, c) in $v(\cdot)$ is corresponding with player (j, c) in $u(\cdot)$ w.r.t. data instance (\mathbf{x}, \mathbf{y}) . Then, denote their Shapley values based on local value functions respectively as $\phi_v^{\mathbf{x}, \mathbf{y}}(i, c)$ and $\phi_u^{\mathbf{x}, \mathbf{y}}(j, c)$, we have $\phi_v^{\mathbf{x}, \mathbf{y}}(i, c) = \phi_u^{\mathbf{x}, \mathbf{y}}(j, c)$.

PROOF. Figure 5 shows two examples of corresponding players between game v and u . In the proof below, we first simplify the notations by using i and j to represent the player (i, c) and (j, c) . Let's define function $g : 2^{\mathcal{N}_v} \rightarrow 2^{\mathcal{N}_u}$ as follow:

$$g(\mathcal{S}_v)_j = \begin{cases} 1 & \text{if corresponding } i \in \mathcal{S}_v \\ 0 & \text{else} \end{cases} \quad (16)$$

In other words, $g(\mathcal{S}_v) = \mathcal{S}_u$ if and only if $\forall j \in \mathcal{S}_u$, corresponding $i \in \mathcal{S}_v$ and $\forall j \in \mathcal{N}_u \setminus \mathcal{S}_u$, corresponding $i \notin \mathcal{S}_v$. Intuitively, this g function provide a method to map from a set \mathcal{S}_v in v to the corresponding

set \mathcal{S}_u in u . Let's define marginal contribution Δ as:

$$\Delta(i|\mathcal{S}_v) = v(\mathcal{S} \cup i) - v(\mathcal{S}) \quad (17)$$

It is easy to show that if $g(\{i\}) = \{j\}$ and $g(\mathcal{S}_v) = \mathcal{S}_u$ then:

$$g(\mathcal{S}_v \cup \{i\}) = \mathcal{S}_u \cup \{j\}.$$

Consequently, if $g(\mathcal{S}_v) = \mathcal{S}_u$ and $g(\{i\}) = \{j\}$, then:

$$\Delta(i|\mathcal{S}_v) = \Delta(j|\mathcal{S}_u)$$

It is trivial to show that $\phi_v^{\mathbf{x}, \mathbf{y}}(i) = 0$ if $g(\{i\}) = \emptyset$ (feature i is not activated). We will only consider case that $\exists j \in \mathcal{N}_u, j = g(\{i\})$. By definition, we have Shapley value of a player i in game v as:

$$\phi_v^{\mathbf{x}, \mathbf{y}}(i) = \frac{1}{|\mathcal{N}_v|} \sum_{\mathcal{S}_v \subseteq \mathcal{N}_v \setminus i} \binom{|\mathcal{N}_v| - 1}{|\mathcal{S}_v|}^{-1} \Delta(i|\mathcal{S}_v), \quad (18)$$

It is easy to show that:

$$|\mathcal{N}_v| = nd; |\mathcal{N}_u| = md. \quad (19)$$

With known \mathcal{S}_u and $|\mathcal{S}_v|$, the number of \mathcal{S}_v s.t. $g(\mathcal{S}_v) = \mathcal{S}_u$ is

$$\binom{|\mathcal{N}_v| - |\mathcal{N}_u|}{|\mathcal{S}_v| - |\mathcal{S}_u|} = \binom{nd - md}{|\mathcal{S}_v| - |\mathcal{S}_u|},$$

as we only need to find $|\mathcal{S}_v| - |\mathcal{S}_u|$ from a pool of $|\mathcal{N}_v| - |\mathcal{N}_u|$ elements. Thus, we can write Eq. 18 as:

$$\begin{aligned} \phi_v^{\mathbf{x}, \mathbf{y}}(i) &= \frac{1}{nd} \sum_{\mathcal{S}_u \subseteq \mathcal{N}_u \setminus j} \Delta(j|\mathcal{S}_u) \sum_{k=|\mathcal{S}_u|}^{nd-md+|\mathcal{S}_u|} \binom{nd-md}{k-|\mathcal{S}_u|} \binom{nd-1}{k}^{-1} \\ &= \frac{1}{nd} \sum_{\mathcal{S}_u \subseteq \mathcal{N}_u \setminus j} \Delta(j|\mathcal{S}_u) \sum_{k=0}^{nd-md} \binom{nd-md}{k} \binom{nd-1}{k+|\mathcal{S}_u|}^{-1} \end{aligned} \quad (20)$$

With some rearrangements, we have the following equation:

$$\begin{aligned} &\binom{nd-md}{k} \binom{md-1}{|\mathcal{S}_u|} \binom{nd-1}{k+|\mathcal{S}_u|}^{-1} \\ &= \frac{nd}{md} \binom{k+|\mathcal{S}_u|}{|\mathcal{S}_u|} \binom{nd-1-k-|\mathcal{S}_u|}{md-|\mathcal{S}_u|-1} \binom{nd}{md}^{-1} \end{aligned} \quad (21)$$

And by Chu-Vandermonde Identity:

$$\sum_{k=0}^{nd-md} \binom{k+|\mathcal{S}_u|}{|\mathcal{S}_u|} \binom{nd-1-k-|\mathcal{S}_u|}{md-|\mathcal{S}_u|-1} = \binom{nd}{md} \quad (22)$$

With Eq. (20), (21), and (22), we have:

$$\begin{aligned} \phi_v^{\mathbf{x}, \mathbf{y}}(i) &= \frac{1}{nd} \sum_{\mathcal{S}_u \subseteq \mathcal{N}_u \setminus j} \Delta(j|\mathcal{S}_u) \frac{nd}{md} \binom{md-1}{|\mathcal{S}_u|}^{-1} \\ &= \phi_u^{\mathbf{x}, \mathbf{y}}(j) \end{aligned} \quad (23)$$

□

In the part below, we provide proof for the codebook optimization result:

$$\begin{aligned} C^* &= \min_C \mathbb{E}_{\mathbf{x} \sim \mathcal{D}, |Q|=B} \left[\|\mathbf{E}^\top \mathbf{x} - \mathbf{E}_{Q,C}^\top \mathbf{x}\|_2^2 \right], \\ C^*[j, :] &= \frac{\sum_{i \in \mathcal{F}_j} p_i \mathbf{V}_j[i, :]}{\sum_{i \in \mathcal{F}_j} p_i}, \end{aligned} \quad (24)$$

with B is any given memory budget, \mathcal{F}_j is the set of features that belong to field j , p_i is the appear frequency of feature i . To begin with, we have the following equation:

$$\begin{aligned} &\mathbb{E}_{\mathbf{x} \sim \mathcal{D}, |Q|=B} \left[\|\mathbf{E}^\top \mathbf{x} - \mathbf{E}_{Q,C}^\top \mathbf{x}\|_2^2 \right] \\ &= \sum_{j=1}^m \sum_{i \in \mathcal{F}_j} \sum_{c=1}^d P((i, c) \in Q | |Q| = B) p_i (\mathbf{E}[i, c] - C[j, c])^2 \end{aligned} \quad (25)$$

where $P((i, c) \in Q | |Q| = B)$ is the probability that parameter (i, c) is in Q . Since Q is sampled from a uniform distribution with a budget B , the probability of any embedding parameter appearing in Q is uniform across all parameters and can be treated as a constant. By removing the above constant, we have the following equation:

$$C^* = \min_C \sum_{j=1}^m \sum_{i \in \mathcal{F}_j} p_i \|\mathbf{V}_j[i, :] - C[j, :]\|_2^2 \quad (26)$$

Apply Cauchy–Bunyakovsky–Schwarz inequality to the above equation, we receive the required Eq. 24.

B Algorithm Analysis

This part provides prove for the converge rate of Algorithm. 1, based on [12]. The central limit theorem states that if dataset \mathcal{D} becomes large, the sample mean $\hat{\phi}_v(i)$ converges in distribution to a Gaussian with mean:

$$\mathbb{E}[\Delta(j, c | \mathcal{S}_u)] = \mathbb{E}[\Delta(i, c | \mathcal{S}_v)] = \phi_v(i, c). \quad (27)$$

The first equation is due to Theorem 4.3, while the second equation is due to Eq. (10). Eq. (27) shows that Algorithm 1 is unbiased. Also because of central limit theorem, the estimation variance is:

$$\frac{\text{Var}(\Delta(j, c | \mathcal{S}_u))}{|\mathcal{D}|}.$$

Although we don't have access to the numerator $\text{Var}(\Delta(j, c | \mathcal{S}_u))$, we can conclude that the variance behaves as $O(1/|\mathcal{D}|)$. Thus, the algorithm would converge with large enough $|\mathcal{D}|$.

C Experiment Settings

In this section, we provide more details on our experiment settings.

C.1 Baselines

We compare Shaver with the following embedding pruning methods for on-device content-based recommendation:

- QR [49] divides the original embedding table into two smaller embedding tables (commonly called meta-embedding tables). The meta-embedding vectors are aggregated by multiplication to create the final embedding vector.

- TTRec [59] tackles the challenge of excessive embedding table parameters with a sequence of matrix products, employing tensor-train decomposition.
- PEP [32] adopts a Soft-Thresholding Reparametrization (STR) trick to iteratively remove redundant parameters in the first training step.
- OptEmbed [35] involves three consecutive steps: Train supernet and remove excessive features, evolutionary search for optimal embedding size for each field, and finally, retrain the network from the found embedding mask.
- CERP [31] incorporate STR into two equal-size embedding tables E_1 and E_2 . To compensate for the two sparse meta-embedding tables, the authors integrate a regularization loss into the training loss and combine two meta-embedding vectors with a sum operation, resulting in a dense vector.
- MagPrune [52] prune the trained embedding table by their magnitude. In line with our method, MagPrune also compresses the model by pruning the trained model.
- Taylor [38] proposes to prune the model with the first-order Taylor approximation of the loss function. Similar to our method, this method also employs feedback from the dataset to calculate the attribution scores.
- Post-Training Quantization (PTQ) [20, 40] first linearly scales the weight into $[-2^{b-1}, 2^{b-1} - 1]$ range, where b is the bit width used. Then we round up the result, converting parameters from float to int. When performing inference, these compressed parameters would be scaled back to float32 with the stored bias and scale.

For Taylor and our methods, we pruned the original checkpoints as specified. For PTQ, we compressed the models from float32 to (4, 8, 16)-bit integers, the sparsity rate is calculated as $1 - b/32$. During training, Shaver's codebook C is freed.

C.2 Datasets

The dataset preprocess steps are described below:

- Criteo [22] contains ad click data over a week. Following the winning solution from the original Criteo Challenge, we discretize each value x in numeric feature fields to $\lceil \log_2(x) \rceil$ if $x > 2$. We replace infrequent features, which appear less than $\text{min_count} = 10$ times, with out-of-vocabulary (OOV) tokens for each field.
- Avazu [51] includes 10 days of click-through data. As a common practice, we first remove the 'id' field, which has a unique value for every record. Similar to Criteo dataset, we remove infrequent features and replace them with OOV tokens ($\text{min_count} = 2$).
- KDD [1] consists of records gathered from search session logs. Similar to previous two datasets, we utilize OOV tokens to represent infrequent features ($\text{min_count} = 10$).

For Criteo and Avazu dataset, we directly adopt the split and preprocessed data from [52].

D Further Experimental Results

Table 4 and Figure 6 show the comparison results with other non-single-shot and single-shot methods of Shaver for DeepFM backbone. Table 5 provides uncompressed model performance.

Table 4: Comparative results with non-single-shot baselines under three sparsity rates. All methods adopt DeepFM as the CRS backbone. In each dataset, the best result is marked in bold and the second best one is underlined.

Dataset	Method	Single-shot?	$t = 50\%$			$t = 80\%$			$t = 95\%$		
			AUC \uparrow	LogLoss \downarrow	#Params	AUC \uparrow	LogLoss \downarrow	#Params	AUC \uparrow	LogLoss \downarrow	#Params
Criteo	QR [49]	\times	0.8081	0.4435	8.69M	0.8078	0.4438	3.48M	0.8033	0.4482	870K
	TTRec [59]	\times	0.8075	0.4442	9.36M	0.8070	0.4446	3.15M	0.8087	0.4431	870K
	PEP [32]	\times	0.8105	0.4414	8.67M	0.8098	0.4419	3.47M	0.8084	0.4432	806K
	OptEmb [35]	\times	-	-	-	0.8088	0.4430	3.17M	0.8043	0.4471	863K
	CERP [31]	\times	0.8099	0.4419	8.66M	0.8095	0.4423	3.21M	0.8062	0.4454	864K
	Shaver-Zero	\checkmark	0.8101	0.4417	8.69M	0.8100	<u>0.4418</u>	3.48M	0.8085	0.4432	869K
	Shaver-Zero †	\times	0.8103	<u>0.4415</u>	8.69M	0.8105	0.4414	3.48M	0.8104	<u>0.4414</u>	869K
	Shaver-Codebook	\checkmark	0.8101	0.4417	8.69M	<u>0.8101</u>	<u>0.4418</u>	3.48M	0.8093	<u>0.4425</u>	870K
	Shaver-Codebook †	\times	0.8105	0.4414	8.69M	0.8105	0.4414	3.48M	0.8106	0.4413	870K
Avazu	QR [49]	\times	0.7697	0.3864	35.43M	0.7695	0.3871	14.17M	0.7644	0.3905	3.54M
	TTRec [59]	\times	0.7647	0.3915	35.43M	0.7651	0.3902	14.17M	0.7608	0.3924	3.54M
	PEP [32]	\times	0.7647	0.3886	29.67M	0.7633	0.3896	14.07M	0.7590	0.3922	3.50M
	OptEmb [35]	\times	-	-	-	-	-	-	0.7585	0.3936	3.41M
	CERP [31]	\times	0.7649	0.3885	35.30M	0.7638	0.3892	14.14M	0.7607	0.3920	3.54M
	Shaver-Zero	\checkmark	0.7655	0.3934	35.43M	0.7645	0.3942	14.17M	0.7575	0.3984	3.54M
	Shaver-Zero †	\times	0.7715	0.3855	35.43M	0.7715	0.3857	14.17M	0.7655	0.3902	3.54M
	Shaver-Codebook	\checkmark	0.7647	0.3939	35.43M	0.7623	0.3952	14.17M	0.7543	0.4002	3.54M
	Shaver-Codebook †	\times	0.7644	0.3903	35.43M	0.7602	0.3940	14.17M	0.7552	0.3974	3.54M
KDD	QR [49]	\times	0.7744	0.1579	48.16M	0.7753	0.1578	19.26M	0.7779	0.1573	4.82M
	TTRec [59]	\times	0.7762	0.1588	48.16M	0.7727	0.1583	19.26M	<u>0.7785</u>	<u>0.1572</u>	4.82M
	PEP [32]	\times	0.7801	0.1568	47.80M	0.7745	0.1578	19.22M	0.7730	0.1579	4.71M
	OptEmb [35]	\times	0.7742	0.1584	47.52M	0.7740	0.1584	18.29M	0.7663	0.1597	4.79M
	CERP [31]	\times	0.7818	0.1567	48.10M	0.7765	0.1578	18.57M	0.7727	0.1581	4.77M
	Shaver-Zero	\checkmark	0.7846	0.1563	48.16M	0.7836	0.1565	19.26M	0.7766	0.1578	4.82M
	Shaver-Zero †	\times	<u>0.7846</u>	<u>0.1563</u>	48.16M	<u>0.7836</u>	<u>0.1565</u>	19.26M	0.7766	0.1578	4.82M
	Shaver-Codebook	\checkmark	0.7850	0.1562	48.16M	0.7854	0.1560	19.26M	0.7812	0.1567	4.82M
	Shaver-Codebook †	\times	0.7850	0.1562	48.16M	0.7854	0.1560	19.26M	0.7812	0.1567	4.82M




Figure 6: Comparative results with single-shot baselines for DeepFM backbone.

Table 5: Original (Uncompressed) Model Performance

Dataset	Backbone	AUC	Loss	Params
Criteo	DeepFM	0.8102	0.4416	17.39M
	DCN	0.8114	0.4407	17.39M
KDD	DeepFM	0.7847	0.1563	96.31M
	DCN	0.7864	0.1562	96.31M
Avazu	DeepFM	0.7658	0.3932	70.86M
	DCN	0.7759	0.3839	70.86M

For DeepFM backbone, our methods achieves the best results in all settings. With the DeepFM backbone, where the original model

has poor performance, our method can remove excessive parameters and improve model performance with fine-tuning. Thus, model fine-tuning has a more positive effect on the DeepFM backbone. Except for the Avazu dataset, codebook approaches outperform zero baseline slightly on Criteo, and significantly on KDD. The reason for this exception is unclear. One plausible explanation is that zero padding can act as a regularization and better prevent overfitting compared to codebook approaches.

Regarding single-shot comparison, in general, the performance trend is similar to DCN, with few exceptions. First, PTQ achieves worse relative results in the Criteo dataset. Second, as mentioned above, Shaver-Zero outperforms Shaver-Codebook and other methods in Avazu settings.