NULL-SPACE CALIBRATION OF AUXILIARY EVIDENCE FOR MULTI-BEHAVIOR RECOMMENDATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Multi-behavior recommendation learns user and item representations from target and auxiliary interactions to improve target prediction. Auxiliary representations may encode preferences not fully reflected in target interactions, but association strength alone does not determine the ranking value of this evidence. A key challenge is to retain target-associated information while learning to use potentially complementary evidence. To this end, we propose ***T**arget-referenced **N**ull-space **C**alibration* (***TNC***), a plug-in module that calibrates auxiliary representations through null-space constraints. Using cross-covariance between paired auxiliary and target representations, ***TNC*** constructs a target-associated subspace through truncated singular value decomposition. It retains both the target-associated and null-space components and learns separate calibration gains under target ranking supervision. For a fixed representation, reference, and target-associated gain, adjusting the null-space gain calibrates potentially complementary evidence without changing the prescribed target-associated projection. The module operates between behavior encoding and fusion and is jointly trained with compatible encoders, allowing target supervision to guide both representation learning and calibration. Experiments on three real-world datasets show that integrating ***TNC*** improves ranking performance across multiple multi-behavior recommendation backbones. To facilitate reproducibility, our implementation is available in an anonymous repository at this url.

1 INTRODUCTION

Multi-behavior recommendation learns user and item representations from target interactions and auxiliary feedback, such as views, favorites, and cart additions, to predict target behaviors such as purchases (Gao et al., 2019; Jin et al., 2020). Auxiliary behavior representations can encode different aspects of user interests and capture latent preferences not fully reflected in target interactions, providing additional evidence for recommendation (Ding et al., 2020; Meng et al., 2024). As illustrated in Figure 1, purchasing running shoes indicates an interest in running, while browsing trail-running shoes and adding a hiking backpack to the cart may suggest a more specific preference for trail running. When both road-running and trail-running shoes match the observed running interest, these auxiliary cues may help the model distinguish between the candidates. A key question is therefore how to use the preference evidence encoded in auxiliary representations to support target ranking.

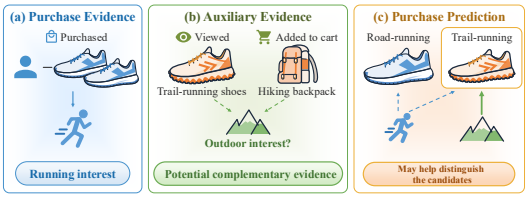

Figure 1: Illustration of potentially complementary auxiliary evidence for purchase prediction. Auxiliary interactions may reveal preferences not fully reflected in purchase history and help distinguish candidate items.

Prior work improves auxiliary representation learning through behavioral dependency modeling (Gao et al., 2019; Yan et al., 2024; Cheng et al., 2023), graph-based interaction modeling (Jin et al., 2020; Chen et al., 2021; Xia et al., 2021), and self-supervised learning (Wei et al., 2022; Gu et al., 2022; Xu et al., 2023). Projection-based methods such as HPMR and PKEF further distinguish shared and behavior-specific information within behavior representations and assign them different

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

roles in transfer, aggregation, and supervision (Meng et al., 2023b;a). However, association with target representations does not by itself determine ranking value. Strongly associated information can reinforce existing preference estimates, while information outside the dominant associated subspace may provide supplementary cues about latent preferences. As auxiliary representations can also encode variation unrelated to target prediction, the contributions of their components need to be learned from the target task (Zhang et al., 2023; Gong et al., 2025). Association statistics should therefore guide the organization of auxiliary information, but should not alone determine whether potentially complementary evidence is discarded. A single scaling factor applied to an entire auxiliary representation jointly rescales target-associated information and other cues, coupling their adjustments. We therefore study *how to calibrate auxiliary representations to retain target-associated information while separately learning the contribution of potentially complementary evidence to target ranking*.

To this end, we propose ***T**arget-referenced **N**ull-space **C**alibration* (***TNC***), a plug-in module that calibrates auxiliary representations through null-space constraints to better use potentially complementary evidence. Using cross-behavior associations, ***TNC*** constructs a target-associated subspace and retains both the target-associated and null-space components of each auxiliary representation. Inspired by null-space constraints for preserving specified linear responses (Fang et al., 2025; Sheng et al., 2026), we formulate calibration as uniform rescaling plus an additional adjustment confined to the null space of the target-associated projection. For a fixed input representation, reference, and target-associated gain, this constraint allows the contribution of potentially complementary evidence in the null-space component to be adjusted without changing the prescribed target-associated projection. Association statistics establish the calibration reference, while target ranking supervision learns separate gains to control how the two components contribute to prediction. The module operates between behavior encoding and fusion and is jointly trained with compatible encoders within a unified training framework, so target supervision guides both representation learning and calibration. Our main contributions are summarized as follows:

- We propose ***TNC***, a plug-in auxiliary representation calibration module that learns to use potentially complementary preference evidence for target prediction.
- We use cross-behavior associations to construct a target reference and constrain additional calibration adjustments to the null space of its projection.
- Experiments with six encoders on three real-world datasets show that the complete integration improves performance in most evaluated settings.

2 Related Work

Multi-behavior representation learning. Multi-behavior representation learning exploits behavioral dependencies (Xia et al., 2020), graph structure (Chen et al., 2021; Xia et al., 2021), and self-supervision. NMTR (Gao et al., 2019) combines shared embeddings, cascaded predictions, and multi-task learning to capture predictive dependencies. MBGCN (Jin et al., 2020) models behavior influence strengths through user–item propagation and behavior semantics through item–item propagation, while MB-CGCN (Cheng et al., 2023) transfers transformed representations across cascaded behavior-specific graph encoders. Cross-behavior contrastive objectives in S-MBRec (Gu et al., 2022) and MBSSL (Xu et al., 2023) provide additional supervision. ***TNC*** calibrates auxiliary representations before fusion and supports joint training with encoders that provide compatible behavior-specific representations.

Auxiliary information decomposition and target adaptation. Methods also distinguish components of behavioral information (Yan et al., 2025a;b; Li et al., 2025). CKML (Meng et al., 2024) learns shared and behavior-specific interests with cross-behavior dependencies among shared interests, while Disen-CGCN (Cheng et al., 2024) combines interest-factor decomposition, cross-behavior transformations, and factor-level aggregation. Projection-based HPMR (Meng et al., 2023b) transfers shared representations downstream and supervises and reprojects unique representations, while PKEF (Meng et al., 2023a) uses shared and unique expert components for aggregation and auxiliary supervision, respectively. Target adaptation includes interaction denoising and target-specific prompt tuning in DPT (Zhang et al., 2023), auxiliary–target user-representation alignment and cross-behavior prediction-risk regularization in RMBRec (Cai et al., 2026b), and frequency-domain filtering with global-context attention in SpectraMB (Cai et al., 2026a). ***TNC*** calibrates

auxiliary representations under a target-associated reference, retaining both the dominant associated and null-space components for scoring and learning separate gains under target ranking supervision.

Null-space constraints. Null-space and subspace constraints restrict updates to limit interference with specified linear responses (Zeng et al., 2019; Saha et al., 2021; Wang et al., 2021). AlphaEdit (Fang et al., 2025) constrains knowledge-editing weight updates using representations of preserved knowledge, while AlphaSteer (Sheng et al., 2026) uses benign-prompt activations to constrain refusal steering and limit its effects on benign inputs. ***TNC*** uses the leading left singular directions of auxiliary–target cross-covariance to construct a target-associated projection in the auxiliary representation space. Its null space constrains additional calibration adjustments. With the input representation, reference, and target-associated gain fixed, changing the null-space gain preserves the prescribed target-associated projection, enabling separate calibration of potentially complementary evidence. Linear concept erasure aims to remove linearly decodable concept information through projection (Ravfogel et al., 2020; 2022; Belrose et al., 2023), whereas ***TNC*** retains both components from the original auxiliary representation for prediction.

3 PRELIMINARIES

Multi-behavior recommendation uses auxiliary interactions, such as views and cart additions, to improve prediction of a target behavior, such as purchases (Gao et al., 2019; Jin et al., 2020; Cheng et al., 2023). Let $\mathcal{U}$, $\mathcal{I}$, and $\mathcal{B}$ denote the user, item, and behavior sets, respectively, with $K = |\mathcal{B}|$ behavior types and target behavior $t \in \mathcal{B}$. The auxiliary behaviors are $\mathcal{B} \setminus \{t\}$. Following the implicit-feedback setting (Hu et al., 2008), the training matrix $\mathbf{R}_b \in \{0,1\}^{|\mathcal{U}|\times|\mathcal{I}|}$ records observed interactions under behavior b as $\mathbf{R}_{b,ui} = 1$ and unobserved interactions as 0.

A common approach learns behavior-specific user and item representations for target prediction (Cheng et al., 2023; Cai et al., 2026b). We describe this process through a behavior encoder f_θ that maps the interaction matrices to user and item representations:

$$\{\mathbf{H}_{U,b}, \mathbf{H}_{I,b}\}_{b\in\mathcal{B}} = f_\theta(\{\mathbf{R}_b\}_{b\in\mathcal{B}}), \quad \mathbf{H}_{U,b} \in \mathbb{R}^{|\mathcal{U}|\times d}, \quad \mathbf{H}_{I,b} \in \mathbb{R}^{|\mathcal{I}|\times d}, \tag{1}$$

where θ denotes the encoder parameters and d is the representation dimension. For a user or item x, $\mathbf{h}_{x,b} \in \mathbb{R}^{1\times d}$ denotes its representation under behavior b.

Downstream mappings $\mathcal{F}_{\psi,U}$ and $\mathcal{F}_{\psi,I}$ construct user and item preference representations from these behavior-specific encodings (Cheng et al., 2023; Cai et al., 2026b; Xia et al., 2021; Cai et al., 2026a). We estimate target behavior preferences using inner-product scoring between user and item representations (He et al., 2020; Cheng et al., 2023):

$$\boldsymbol{\phi}_u = \mathcal{F}_{\psi,U}((\mathbf{h}_{u,b})_{b\in\mathcal{B}}), \quad \boldsymbol{\phi}_i = \mathcal{F}_{\psi,I}((\mathbf{h}_{i,b})_{b\in\mathcal{B}}), \quad \widehat{y}_{ui} = \boldsymbol{\phi}_u \boldsymbol{\phi}_i^\top, \tag{2}$$

where $\boldsymbol{\phi}_u$ and $\boldsymbol{\phi}_i$ are the representations used for scoring, and ψ denotes learnable parameters of the downstream mappings when present.

To learn target behavior preferences, the Bayesian Personalized Ranking (BPR) objective (Rendle et al., 2009) encourages observed target items to rank above sampled unobserved items:

$$\mathcal{L}_{\mathrm{BPR}} = -\frac{1}{|\mathcal{D}_t|} \sum_{(u,i^+,i^-)\in\mathcal{D}_t} \log \sigma(\widehat{y}_{ui^+} - \widehat{y}_{ui^-}), \tag{3}$$

where σ is the sigmoid function and $\mathcal{D}_t$ is a nonempty set of sampled triplets satisfying $\mathbf{R}_{t,ui^+} = 1$ and $\mathbf{R}_{t,ui^-} = 0$. During inference, candidate items are ranked by their target behavior scores.

4 METHOD

Auxiliary representations can encode both shared preferences and behavior-specific interests, providing potentially complementary evidence for target prediction (Ding et al., 2020; Meng et al., 2024; 2023b). We propose ***T**arget-referenced **N**ull-space **C**alibration* (***TNC***), a plug-in module that calibrates auxiliary representations to retain target-associated information while learning how potentially complementary evidence contributes to target ranking. The module constructs a reference

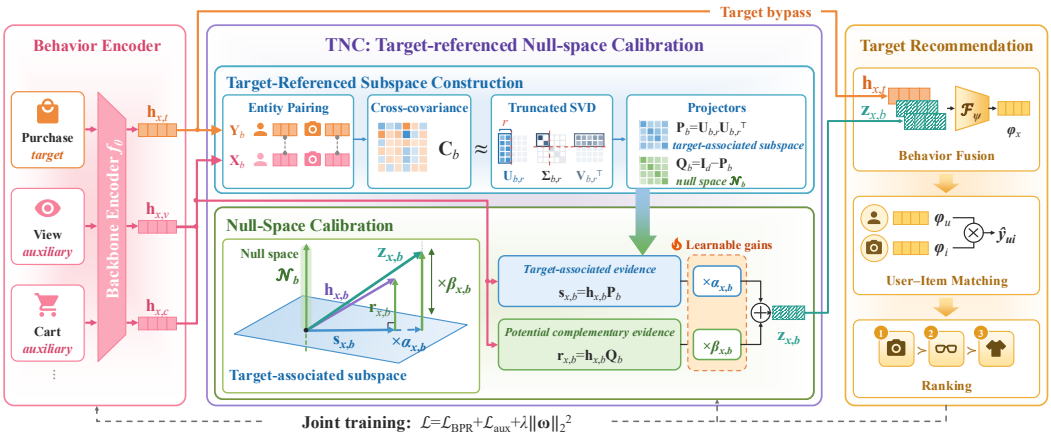

Figure 2: Overview of ***TNC***, a plug-in module for auxiliary representation calibration.

from cross-behavior associations and uses its null space to separate the calibration of information within and outside the target-associated subspace. Both components remain available for prediction, with their contributions learned through target ranking supervision. As illustrated in Figure 2, ***TNC*** transforms each auxiliary encoder output $\mathbf{h}_{x,b}$ into a calibrated representation $\mathbf{z}_{x,b}$ of the same dimension for downstream target prediction.

4.1 Target-Referenced Subspace Construction

We construct the calibration reference from paired auxiliary and target representations of the same users or items. For an auxiliary behavior $b \neq t$, let $\mathcal{U}_{b,t}$ and $\mathcal{I}_{b,t}$ denote the users and items with at least one training interaction under each of b and t. Their sizes are $n_{b,U} = |\mathcal{U}_{b,t}|$ and $n_{b,I} = |\mathcal{I}_{b,t}|$, and data preprocessing ensures that both sets are nonempty.

We use empirical cross-covariance to estimate linear associations between the paired representations (Shao et al., 2023; Belrose et al., 2023). For the selected users, $\mathbf{X}_{b,U}, \mathbf{Y}_{b,U} \in \mathbb{R}^{n_{b,U} \times d}$ contain their auxiliary and target representations, respectively, with matching rows corresponding to the same user. We construct $\mathbf{X}_{b,I}, \mathbf{Y}_{b,I} \in \mathbb{R}^{n_{b,I} \times d}$ analogously for items. Subtracting the mean of each feature across the selected entities gives the centered matrices, denoted by tildes. We compute the user- and item-side cross-covariances and combine them with equal weights:

$$\mathbf{C}_{b,U} = \frac{\widetilde{\mathbf{X}}_{b,U}^\top \widetilde{\mathbf{Y}}_{b,U}}{n_{b,U}}, \quad \mathbf{C}_{b,I} = \frac{\widetilde{\mathbf{X}}_{b,I}^\top \widetilde{\mathbf{Y}}_{b,I}}{n_{b,I}}, \quad \mathbf{C}_b = \frac{1}{2}\mathbf{C}_{b,U} + \frac{1}{2}\mathbf{C}_{b,I}. \tag{4}$$

Each (p, q) entry of $\mathbf{C}_{b,U}$ or $\mathbf{C}_{b,I}$ measures the empirical covariance between auxiliary feature p and target feature q, computed over the paired entities, where $p, q \in \{1, \ldots, d\}$. Assuming compatible user and item feature coordinates within each behavior, their average $\mathbf{C}_b \in \mathbb{R}^{d \times d}$ provides a common association reference for users and items.

To identify dominant target-associated directions in the auxiliary representation space, we apply truncated singular value decomposition to $\mathbf{C}_b$ (Eckart & Young, 1936; Mirsky, 1960). Let $\mathbf{U}_{b,r} \in \mathbb{R}^{d \times r}$ contain the left singular vectors corresponding to the r largest singular values, where r is selected on the validation set and satisfies $1 \leq r < d$. We use the left singular vectors because the rows of $\mathbf{C}_b$ correspond to auxiliary features. Their span retains the largest cross-covariance energy in $\mathbf{C}_b$ among auxiliary subspaces of dimension r and defines the orthogonal projector:

$$\mathbf{P}_b = \mathbf{U}_{b,r}\mathbf{U}_{b,r}^\top. \tag{5}$$

This criterion identifies a reference for calibration rather than determining the predictive value of auxiliary information. Information outside the selected subspace is therefore retained for subsequent calibration under the target ranking objective.

4.2 Null-Space Calibration

Information outside the target-associated subspace may contain preference cues that supplement target prediction. We seek to calibrate its contribution without requiring the target-associated component to be rescaled alongside it. Inspired by null-space constraints for preserving specified linear

responses (Wang et al., 2021; Fang et al., 2025; Sheng et al., 2026), we restrict additional adjustments to the null space of the target-associated projection.

For a fixed projector $\mathbf{P}_b$, an additive adjustment $\boldsymbol{\delta} \in \mathbb{R}^{1\times d}$ preserves the target-associated projection if and only if $\boldsymbol{\delta}\mathbf{P}_b = \mathbf{0}$. We therefore define the calibration null space and its projector as:

$$\mathcal{N}_b = \left\{\mathbf{v} \in \mathbb{R}^{1\times d} : \mathbf{v}\mathbf{P}_b = \mathbf{0}\right\}, \qquad \mathbf{Q}_b = \mathbf{I}_d - \mathbf{P}_b, \tag{6}$$

where $\mathbf{I}_d$ is the $d \times d$ identity matrix and $\mathbf{Q}_b$ projects orthogonally onto the $(d-r)$-dimensional space $\mathcal{N}_b$. This null space is defined by $\mathbf{P}_b$ and need not coincide with the left null space of the full cross-covariance matrix $\mathbf{C}_b$.

We apply these projectors to the original, uncentered encoder output:

$$\mathbf{s}_{x,b} = \mathbf{h}_{x,b}\mathbf{P}_b, \qquad \mathbf{r}_{x,b} = \mathbf{h}_{x,b}\mathbf{Q}_b, \tag{7}$$

where $\mathbf{s}_{x,b}$ is the target-associated component and $\mathbf{r}_{x,b} \in \mathcal{N}_b$ is the null-space component. They are orthogonal and satisfy $\mathbf{h}_{x,b} = \mathbf{s}_{x,b} + \mathbf{r}_{x,b}$, so the decomposition preserves the complete auxiliary representation. The null-space component remains available as a source of potentially complementary evidence, whose contribution is learned from the target task rather than inferred from association strength alone.

For each user or item x under auxiliary behavior b, we introduce scalar gains $\alpha_{x,b}$ and $\beta_{x,b}$ to control the two contributions while preserving their subspace membership:

$$\mathbf{z}_{x,b} = \alpha_{x,b}\mathbf{s}_{x,b} + \beta_{x,b}\mathbf{r}_{x,b} = \alpha_{x,b}\mathbf{h}_{x,b} + (\beta_{x,b} - \alpha_{x,b})\mathbf{r}_{x,b}. \tag{8}$$

The second expression writes calibration as uniform rescaling by $\alpha_{x,b}$ plus an adjustment confined to $\mathcal{N}_b$. Consequently, $\mathbf{z}_{x,b}\mathbf{P}_b = \alpha_{x,b}\mathbf{s}_{x,b}$. For a fixed input representation, reference, and gain $\alpha_{x,b}$, changing $\beta_{x,b}$ therefore adjusts the contribution of potentially complementary evidence without altering the prescribed target-associated projection. The target-associated contribution remains learnable through $\alpha_{x,b}$, and equal gains recover uniform rescaling.

We parameterize the gains using learnable scalars $g_{x,b,s}$ and $g_{x,b,r}$:

$$\alpha_{x,b} = 1 + \tanh(g_{x,b,s}), \qquad \beta_{x,b} = 1 + \tanh(g_{x,b,r}). \tag{9}$$

Both gains lie in $(0, 2)$, allowing bounded attenuation or amplification of each component. We initialize the raw parameters to zero, yielding unit gains and $\mathbf{z}_{x,b} = \mathbf{h}_{x,b}$, so calibration starts from the identity mapping and preserves the original representation.

4.3 Target Prediction and Joint Optimization

The calibrated auxiliary representations are used for downstream target prediction. The target representation bypasses gain calibration, giving $\mathbf{z}_{x,t} = \mathbf{h}_{x,t}$. Using the prediction mappings introduced in Eq. (2), we compute:

$$\boldsymbol{\phi}_u = \mathcal{F}_{\psi,U}((\mathbf{z}_{u,b})_{b\in\mathcal{B}}), \quad \boldsymbol{\phi}_i = \mathcal{F}_{\psi,I}((\mathbf{z}_{i,b})_{b\in\mathcal{B}}), \quad \widehat{y}_{ui} = \boldsymbol{\phi}_u\boldsymbol{\phi}_i^\top. \tag{10}$$

The prediction configuration used in our experiments is specified in Appendix B.2.

The resulting scores enter the target ranking loss in Eq. (3), which is combined with the backbone's auxiliary objectives for joint optimization:

$$\mathcal{L} = \mathcal{L}_{\mathrm{BPR}} + \mathcal{L}_{\mathrm{aux}} + \lambda\|\boldsymbol{\omega}\|_2^2. \tag{11}$$

Here, $\mathcal{L}_{\mathrm{aux}}$ retains the original weights and inputs of the backbone's auxiliary objectives, excluding the target ranking loss and regularization, and is zero when no auxiliary objectives are present. Objectives defined on behavior encodings continue to use uncalibrated representations. The vector $\boldsymbol{\omega}$ contains all trainable model parameters, including the encoder and calibration parameters, and $\lambda \geq 0$ controls L_2 regularization.

Joint optimization allows target supervision to guide both auxiliary representation learning and the calibration of potentially complementary evidence. At every optimization step, subspace references are recomputed from current training representations and treated as fixed during gradient computation. Gradients propagate through target prediction and calibration to update the trainable parameters, but not through cross-covariance estimation or SVD. Representations under all behaviors, including the target, remain trainable throughout joint optimization.

4.4 Theoretical Analysis

The following proposition characterizes calibration as a minimum-change adjustment under the prescribed target-associated projection.

Proposition 1 (Minimal constrained calibration)**.** *Given $\mathbf{h}_{x,b} \in \mathbb{R}^{1\times d}$, an orthogonal projector $\mathbf{P}_b$, and scalar gains $\alpha_{x,b}$ and $\beta_{x,b}$, the calibrated representation in Eq. (8) uniquely solves*

$$\min_{\mathbf{v}\in\mathbb{R}^{1\times d}} \|\mathbf{v} - \beta_{x,b}\mathbf{h}_{x,b}\|_2^2 \quad \textit{subject to} \quad \mathbf{v}\mathbf{P}_b = \alpha_{x,b}\mathbf{h}_{x,b}\mathbf{P}_b, \tag{12}$$

where $\|\cdot\|_2$ denotes the Euclidean norm.

Here, $\beta_{x,b}\mathbf{h}_{x,b}$ is the uniformly scaled reference, while the constraint prescribes the target-associated contribution through $\alpha_{x,b}$. The unique closest representation satisfies this constraint while keeping the null-space component at $\beta_{x,b}\mathbf{r}_{x,b}$. This gives a geometric interpretation of separate calibration, allowing potentially complementary evidence to contribute at a learned strength without requiring the same gain for target-associated information.

The proposition characterizes the calibration rule and introduces no additional training objective for either encoders or gains. It holds for fixed inputs, references, and gains, which may evolve during joint training. Appendix A provides the theoretical results and reference sensitivity analysis, while Algorithm 1 summarizes joint training of ***TNC*** and the behavior encoder.

5 Experiments

5.1 Experimental Settings

Datasets. We evaluate on Taobao (Tianchi, n.d.a), Tmall (Tianchi, n.d.b), and Beibei (Gao et al., 2019), using purchases as the target behavior across all three benchmark datasets. Table 1 summarizes the datasets and auxiliary interaction types. Detailed descriptions are provided in Appendix B.1.

Table 1: Statistics of the three datasets.

Statistic	Taobao	Tmall	Beibei
Users	48,749	41,738	21,716
Items	39,493	11,953	7,977
View	1,548,162	1,813,498	2,412,586
Favorites	—	221,514	—
Cart	193,747	1,996	642,622
Purchase	259,771	287,158	304,576

Backbones and integration. We compare MBGCN (Jin et al., 2020), CRGCN (Yan et al., 2024), RMBRec (Cai et al., 2026b), UIPL (Yan et al., 2025b), SpectraMB (Cai et al., 2026a), and MBLFE (Yan et al., 2025a) with and without ***TNC***, denoted by ***+TNC*** and Base. Integrated models retain each backbone's encoder and auxiliary objectives, calibrate auxiliary representations, and use mean fusion and inner-product scoring. The encoder and calibration gains are jointly trained from scratch. Native aggregation or scoring operations are replaced where necessary. These comparisons therefore evaluate the complete integration. Integration details are provided in Appendix B.2.

Evaluation and parameter settings. We report Hit Ratio (HR) and Normalized Discounted Cumulative Gain (NDCG) at 10 and 20 (Cai et al., 2026b). We adopt full-ranking evaluation without sampling evaluation negatives (Krichene & Rendle, 2020). For each user, we exclude training purchases from the candidate set. Encoding and reference estimation use training interactions only. Backbone hyperparameters are selected on the validation set and shared between Base and ***+TNC***. We fix $d = 64$ and select $r \in \{4, 8, 16, 32, 48\}$ by validation performance. Further settings and ablation definitions appear in Appendices B.3 and B.4, respectively.

5.2 Overall Performance

To evaluate the applicability of auxiliary representation calibration, we compare six multi-behavior recommendation backbones with and without ***TNC*** on three datasets. Table 2 reports purchase recommendation performance, with relative changes calculated against the corresponding Base results. The results show improvements in most evaluated settings and lead to the following observations.

Table 2: Performance comparison of six backbones with and without ***TNC***. Bold indicates the better result within each backbone pair. Gain is the relative change (%) from Base, computed before rounding the scores; teal and red indicate increases and decreases.

		Taobao				Tmall				Beibei			
Backbone	Setting	HR@10	NDCG@10	HR@20	NDCG@20	HR@10	NDCG@10	HR@20	NDCG@20	HR@10	NDCG@10	HR@20	NDCG@20
MBGCN	Base	0.0222	0.0111	0.0345	0.0143	0.0454	0.0240	0.0633	0.0283	0.0548	0.0271	0.0962	0.0375
	+TNC	**0.0246**	**0.0126**	**0.0373**	**0.0158**	**0.0505**	**0.0269**	**0.0708**	**0.0318**	**0.0705**	**0.0356**	**0.1182**	**0.0476**
	Gain (%)	+10.83	+13.02	+8.02	+10.77	+11.27	+12.24	+11.97	+12.48	+28.55	+31.28	+22.89	+26.94
CRGCN	Base	0.0793	0.0443	0.1148	0.0531	0.0747	0.0381	0.1134	0.0475	0.0530	0.0257	0.0952	0.0362
	+TNC	**0.1010**	**0.0603**	**0.1370**	**0.0693**	**0.0833**	**0.0437**	**0.1231**	**0.0534**	**0.0575**	**0.0285**	**0.1005**	**0.0393**
	Gain (%)	+27.45	+36.26	+19.35	+30.47	+11.61	+14.78	+8.57	+12.41	+8.33	+11.04	+5.61	+8.48
RMBRec	Base	0.1453	0.0955	0.1774	0.1036	0.1165	0.0648	0.1602	0.0755	0.0848	0.0438	0.1359	0.0566
	+TNC	**0.1621**	**0.1070**	**0.1958**	**0.1155**	**0.1279**	**0.0719**	**0.1715**	**0.0825**	**0.0912**	**0.0472**	**0.1414**	**0.0599**
	Gain (%)	+11.58	+12.00	+10.39	+11.47	+9.78	+10.99	+7.12	+9.38	+7.49	+7.90	+4.07	+5.77
UIPL	Base	0.1294	0.0769	0.1672	0.0864	**0.1378**	**0.0745**	0.1885	0.0870	0.0586	0.0290	0.1010	0.0397
	+TNC	**0.1337**	**0.0792**	**0.1730**	**0.0891**	0.1372	0.0743	**0.1897**	**0.0871**	**0.0605**	**0.0297**	**0.1031**	**0.0404**
	Gain (%)	+3.28	+3.03	+3.45	+3.11	−0.45	−0.34	+0.65	+0.15	+3.30	+2.28	+2.10	+1.76
SpectraMB	Base	0.1404	0.0849	0.1757	0.0938	**0.1170**	**0.0637**	0.1640	0.0752	0.0522	0.0259	**0.0897**	**0.0353**
	+TNC	**0.1414**	**0.0852**	**0.1779**	**0.0944**	0.1167	0.0637	**0.1642**	**0.0753**	**0.0533**	**0.0262**	0.0890	0.0351
	Gain (%)	+0.67	+0.27	+1.24	+0.60	−0.23	−0.03	+0.10	+0.05	+2.21	+1.04	−0.77	−0.55
MBLFE	Base	0.1629	0.0912	0.2245	0.1067	0.1263	0.0661	0.1827	0.0800	0.0616	0.0307	0.1032	0.0411
	+TNC	**0.1634**	**0.0919**	**0.2255**	**0.1076**	**0.1339**	**0.0705**	**0.1931**	**0.0850**	**0.0668**	**0.0332**	**0.1117**	**0.0445**
	Gain (%)	+0.33	+0.84	+0.47	+0.86	+6.00	+6.66	+5.70	+6.22	+8.45	+8.30	+8.21	+8.25

- **Integrating *TNC* benefits diverse representation learning approaches.** MBGCN, CRGCN, RMBRec, and MBLFE improve on all four metrics across the three datasets, covering graph propagation, cascaded modeling, target alignment, and latent preference modeling. The largest relative improvement reaches 36.26% in NDCG@10 for CRGCN on Taobao. Moreover, models equipped with ***TNC*** achieve the highest scores on all four metrics on Taobao and Beibei, as well as HR@20 and NDCG@20 on Tmall. These results support the effectiveness of the complete integration across different encoder designs. Where prediction operations also change, the gains cannot be attributed solely to calibration.

- **Target consistency leaves room for better use of auxiliary evidence.** RMBRec encourages consistency with the target behavior through contrastive alignment and cross-behavior risk regularization. Integrating ***TNC*** improves its NDCG@10 by 12.00%, 10.99%, and 7.90% on Taobao, Tmall, and Beibei, respectively, while retaining mean fusion and inner-product scoring. ***TNC*** separately calibrates target-associated information and potentially complementary cues outside the target-associated subspace. These gains suggest that consistency-based representation learning can benefit from further calibration that allows potentially useful auxiliary evidence beyond the selected subspace to support target ranking.

- **Integration gains vary across encoders.** On Tmall, integrating ***TNC*** improves RMBRec's NDCG@10 by 10.99%, while SpectraMB remains nearly unchanged (-0.03%). SpectraMB already adjusts auxiliary representations through spectral filtering before combining them with global-context attention. This existing adjustment may partly overlap with calibration, leaving less room for additional gains. Replacing its native attention-based fusion with mean aggregation may also affect the outcome. These results suggest that the benefit of integration depends partly on how the backbone processes auxiliary evidence and uses it for prediction.

5.3 Ablation Study

To understand how ***TNC*** uses auxiliary evidence for purchase prediction, we examine its reference construction, component retention, gain learning, calibration scope, and auxiliary sources with RMBRec. *w/o Calibration* directly uses encoded auxiliary representations, while *w/o Target Reference* constructs an auxiliary PCA reference without removing purchase interactions or supervision. Fixing $\alpha = 0$ or $\beta = 0$ removes the target-associated or null-space component, respectively. Fixing either gain to one retains both components but disables learning of the corresponding gain. *User-only* and *Item-only* apply calibration to auxiliary representations on one side, leaving the other side uncalibrated. Finally, *w/o View* and *w/o Cart* remove the corresponding interactions, objectives, and fusion branches. Table 3 summarizes the Taobao results for the evaluated variants; complete results and detailed analysis across all three datasets are provided in Appendix C.1, with comparison scope specified in Appendix B.4.

Table 3: Ablation results on Taobao with RMBRec. Parenthesized percentages indicate relative changes from ***+TNC***, calculated from the displayed scores; decreases are red.

Metric	***+TNC***	w/o Calibration	w/o Target Reference	Fix $\alpha = 0$	Fix $\beta = 0$	Fix $\alpha = 1$	Fix $\beta = 1$	User-only	Item-only	w/o View	w/o Cart
HR@10	**0.1621**	0.1454 (-10.30%)	0.1456 (-10.18%)	0.1194 (-26.34%)	0.0486 (-70.02%)	0.1487 (-8.27%)	0.1255 (-22.58%)	0.1521 (-6.17%)	0.1573 (-2.96%)	0.1240 (-23.50%)	0.1289 (-20.48%)
NDCG@10	**0.1070**	0.0955 (-10.75%)	0.0956 (-10.65%)	0.0787 (-26.45%)	0.0279 (-73.93%)	0.0990 (-7.48%)	0.0821 (-23.27%)	0.1005 (-6.07%)	0.1039 (-2.90%)	0.0872 (-18.50%)	0.0760 (-28.97%)
HR@20	**0.1958**	0.1772 (-9.50%)	0.1772 (-9.50%)	0.1489 (-23.95%)	0.0681 (-65.22%)	0.1789 (-8.63%)	0.1566 (-20.02%)	0.1855 (-5.26%)	0.1891 (-3.42%)	0.1368 (-30.13%)	0.1677 (-14.35%)
NDCG@20	**0.1155**	0.1036 (-10.30%)	0.1036 (-10.30%)	0.0862 (-25.37%)	0.0328 (-71.60%)	0.1067 (-7.62%)	0.0899 (-22.16%)	0.1089 (-5.71%)	0.1119 (-3.12%)	0.0904 (-21.73%)	0.0858 (-25.71%)

The target reference yields higher reported scores than auxiliary PCA. NDCG@10 is 0.1070 with the target reference and 0.0956 with auxiliary PCA. The PCA variant retains purchase supervision, so it tests a different reference construction rather than the removal of target learning.

Null-space variants show larger performance gaps. Both component-removal variants score below w/o Calibration, with a larger gap for the variant without the null-space component. Fixing $\alpha = 1$ and $\beta = 1$ reduces NDCG@10 by 7.48% and 23.27%, respectively, relative to the complete model. Both component-removal and fixed-gain comparisons show larger gaps for the null-space variants. Fix $\beta = 1$ also underperforms w/o Calibration. Unit initialization applies only before training; subsequent updates to α and the encoder during joint optimization can change relative component weighting and test performance.

Calibrating both users and items improves recommendation. The complete model outperforms both single-sided variants on all four metrics, improving NDCG@10 by 6.47% over User-only and 2.98% over Item-only. Although Item-only performs better than User-only on Taobao, calibrating both sides yields further improvements. These results support adjusting auxiliary representations on both sides of user–item matching for target prediction.

Auxiliary sources contribute differently across ranking metrics. Removing carts causes a larger NDCG@10 decrease than removing views (28.97% versus 18.50%), whereas removing views causes a larger HR@20 decrease (30.13% versus 14.35%). Thus, both sources support purchase prediction on Taobao, but their relative contributions depend on the evaluation metric. The appendix further shows that these effects vary across datasets, including improvements in individual metrics after removing some auxiliary behaviors on Tmall.

5.4 Group-wise Performance Analysis

To identify which users and items benefit from calibration, we compare RMBRec with and without ***TNC*** on Taobao and Tmall. We group users and items by the auxiliary coverage of their training purchases. User coverage is the fraction of a user's training purchase items that also have at least one auxiliary interaction from that user. Item coverage is the fraction of an item's training purchasers who also have at least one auxiliary interaction with that item. Any auxiliary behavior counts toward coverage, and each user–item pair is counted only once. We construct Low-, Medium-, and High-coverage groups separately for users and items within each dataset. Items without training purchases form the *No Training Purchases* group, while *Overall* includes all test purchase pairs. Beibei is excluded because all defined coverage scores equal one in the evaluated split. Detailed definitions and grouping procedures are provided in Appendix B.5.

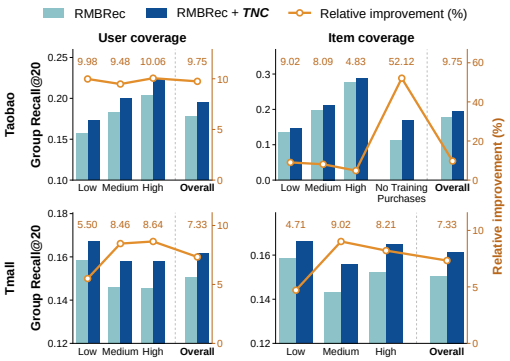

Figure 3: Group Recall@20 of RMBRec with and without ***TNC*** across user and item groups on Taobao and Tmall.

Calibration benefits users and items across coverage levels. As shown in Figure 3, ***TNC*** improves Group Recall@20 in every group, with overall gains of 9.75% on Taobao and 7.33% on Tmall. The Low-, Medium-, and High-coverage groups improve by 4.71%–10.06% across the two datasets. These gains show that calibration benefits entities with both limited and extensive overlap between auxiliary interactions and training purchases. They support learning how auxiliary evidence contributes to prediction across different coverage conditions.

Items without training purchases show a pronounced gain. On Taobao, Recall for these items increases from 0.11110 to 0.16900, a relative improvement of approximately 52.1%. Although they account for only 7.60% of test purchase pairs, they contribute approximately 25.42% of the overall absolute Recall increase. These results show that the integrated model benefits items lacking positive training purchases, without isolating the contribution of their own gains from changes to user representations and joint encoding. Appendix B.5 explains how such items participate in training. Meanwhile, items with training purchases improve by 7.63% when evaluated together. Thus, the benefits extend both to supplementing observed purchase histories and to predicting purchases for items without such histories.

5.5 Representation Visualization

To examine how ***TNC*** changes auxiliary representations, we randomly sample 200 users and 200 items from Tmall using RMBRec as the backbone. We extract behavior-specific representations immediately before and after calibration from the same trained model. For users and items separately, we jointly fit a two-dimensional PCA projection across all behaviors and both stages, keeping the sampled entities and axis limits fixed.

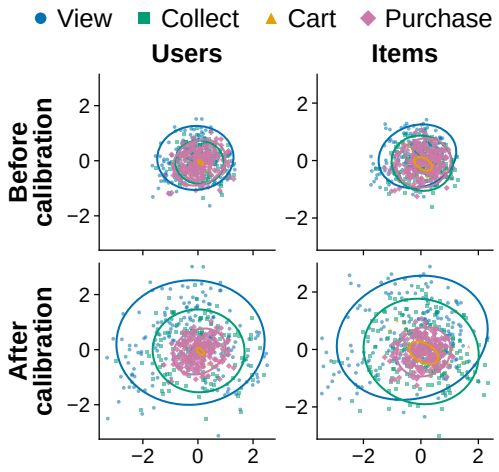

Figure 4: Behavior representations on Tmall with RMBRec + ***TNC***. Columns show users and items; rows show before and after calibration. Colors indicate behaviors (Collect denotes favorites), and ellipses summarize their projected distributions.

As shown in Figure 4, calibration expands the projected distributions of view and collect (favorites) representations on both the user and item sides, while cart representations remain relatively compact. Purchase representations remain unchanged because they bypass calibration. The resulting distributions retain differences across behaviors, which is consistent with our goal of preserving auxiliary cues that may not be fully captured by purchase representations. The projected distributions alone do not establish the ranking value of individual directions. By retaining both the target-associated and null-space components, ***TNC*** allows information outside the retained target-associated subspace to participate in prediction. The ranking improvements and fixed-gain ablations further support learning the contributions of these retained components under target supervision. Additional visualizations on Taobao and Beibei are provided in Appendix C.4.

6 Conclusion

In this work, we proposed ***T**arget-referenced **N**ull-space **C**alibration* (***TNC***), a plug-in module that calibrates auxiliary representations to use potentially complementary preference evidence. Using a target-associated reference and null-space constraints, the module retains both representation components and learns their contributions under target ranking supervision. Experiments with six backbones on three real-world datasets show that the complete integration improves performance in most evaluated settings. Notably, in the RMBRec-based analysis on Taobao, ***TNC*** improves Group Recall@20 by approximately 52.1% for items without training purchases. This finding suggests that calibrating auxiliary representations can help exploit preference evidence even when an item has no positive target interactions during training.

AI USE STATEMENT

We used generative AI tools to assist with literature search and organization, translation and language editing, reference formatting, and technical feedback on the method and mathematical arguments. The authors made the final research decisions and reviewed the AI-assisted content, including verifying references and checking mathematical claims and derivations. We take full responsibility for the final content of this work, including all text, claims, and artifacts developed with AI assistance.

REPRODUCIBILITY STATEMENT

The main text describes the model formulation and training procedure, with experimental settings summarized in Section 5.1. Dataset descriptions, backbone integration details, and evaluation and hyperparameter settings are provided in Appendices B.1, B.2, and B.3, respectively. The assumptions and proofs underlying our theoretical results are presented in Appendix A. To facilitate reproduction of our experiments, our implementation is available at `https://anonymous.4open.science/r/TNC_ICLR2027`.

REFERENCES

Nora Belrose, David Schneider-Joseph, Shauli Ravfogel, Ryan Cotterell, Edward Raff, and Stella Biderman. LEACE: Perfect linear concept erasure in closed form. In *NeurIPS*, volume 36, pp. 66044–66063. Curran Associates, Inc., 2023. doi: 10.52202/075280-2884.

Miaomiao Cai, Yunshan Ma, Fangqi Zhu, Junfeng Fang, Zhijie Zhang, Zhiyong Cheng, Xiang Wang, and See-Kiong Ng. Dynamic Spectral Denoising with Global-Context Attention for Multi-Behavior Recommendation. In *KDD, Vol. 2*, pp. 198–209. ACM, 2026a. doi: 10.1145/3770855.3818191.

Miaomiao Cai, Zhijie Zhang, Junfeng Fang, Zhiyong Cheng, Xiang Wang, and Meng Wang. RMBRec: Robust Multi-Behavior Recommendation towards Target Behaviors. In *WWW*, pp. 6731–6742. Association for Computing Machinery, 2026b. doi: 10.1145/3774904.3792617.

Chong Chen, Weizhi Ma, Min Zhang, Zhaowei Wang, Xiuqiang He, Chenyang Wang, Yiqun Liu, and Shaoping Ma. Graph Heterogeneous Multi-Relational Recommendation. In *AAAI*, volume 35, pp. 3958–3966, 2021. doi: 10.1609/aaai.v35i5.16515.

Zhiyong Cheng, Sai Han, Fan Liu, Lei Zhu, Zan Gao, and Yuxin Peng. Multi-Behavior Recommendation with Cascading Graph Convolution Networks. In *WWW*, pp. 1181–1189. ACM, 2023. doi: 10.1145/3543507.3583439.

Zhiyong Cheng, Jianhua Dong, Fan Liu, Lei Zhu, Xun Yang, and Meng Wang. Disentangled Cascaded Graph Convolution Networks for Multi-Behavior Recommendation. *ACM Trans. Recomm. Syst.*, 2(4):31:1–31:27, 2024. doi: 10.1145/3673244.

Chandler Davis and W. M. Kahan. The rotation of eigenvectors by a perturbation. III. *SIAM J. Numer. Anal.*, 7(1):1–46, March 1970. doi: 10.1137/0707001.

Jingtao Ding, Guanghui Yu, Yong Li, Xiangnan He, and Depeng Jin. Improving Implicit Recommender Systems with Auxiliary Data. *ACM Trans. Inf. Syst.*, 38(1):11:1–11:27, 2020. doi: 10.1145/3372338.

Carl Eckart and Gale Young. The approximation of one matrix by another of lower rank. *Psychometrika*, 1(3):211–218, September 1936. doi: 10.1007/BF02288367.

Ky Fan. On a theorem of Weyl concerning eigenvalues of linear transformations I. *Proc. Natl. Acad. Sci. U.S.A.*, 35(11):652–655, November 1949. doi: 10.1073/pnas.35.11.652.

Junfeng Fang, Houcheng Jiang, Kun Wang, Yunshan Ma, Jie Shi, Xiang Wang, Xiangnan He, and Tat-Seng Chua. AlphaEdit: Null-space constrained knowledge editing for language models. In *ICLR*, 2025.

Chen Gao, Xiangnan He, Dahua Gan, Xiangning Chen, Fuli Feng, Yong Li, Tat-Seng Chua, and Depeng Jin. Neural Multi-task Recommendation from Multi-behavior Data. In *ICDE*, pp. 1554–1557. IEEE, 2019. doi: 10.1109/icde.2019.00140.

Shuwei Gong, Yuting Liu, Yizhou Dang, Guibing Guo, Jianzhe Zhao, and Xingwei Wang. Multiple Purchase Chains with Negative Transfer Elimination for Multi-Behavior Recommendation. In *AAAI*, volume 39, pp. 11717–11725, 2025. doi: 10.1609/aaai.v39i11.33275.

Shuyun Gu, Xiao Wang, Chuan Shi, and Ding Xiao. Self-supervised Graph Neural Networks for Multi-behavior Recommendation. In *IJCAI*, pp. 2052–2058. International Joint Conferences on Artificial Intelligence Organization, 2022. doi: 10.24963/ijcai.2022/285.

Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yongdong Zhang, and Meng Wang. LightGCN: Simplifying and powering graph convolution network for recommendation. In *SIGIR*, pp. 639–648, New York, NY, USA, 2020. Association for Computing Machinery. doi: 10.1145/3397271.3401063.

Yifan Hu, Yehuda Koren, and Chris Volinsky. Collaborative filtering for implicit feedback datasets. In *ICDM*, pp. 263–272. IEEE, 2008. doi: 10.1109/ICDM.2008.22.

Kalervo Järvelin and Jaana Kekäläinen. Cumulated gain-based evaluation of IR techniques. *ACM Trans. Inf. Syst.*, 20(4):422–446, October 2002. doi: 10.1145/582415.582418.

Bowen Jin, Chen Gao, Xiangnan He, Depeng Jin, and Yong Li. Multi-behavior Recommendation with Graph Convolutional Networks. In *SIGIR*, pp. 659–668. Association for Computing Machinery, 2020. doi: 10.1145/3397271.3401072.

Walid Krichene and Steffen Rendle. On sampled metrics for item recommendation. In *KDD*, pp. 1748–1757. Association for Computing Machinery, 2020. doi: 10.1145/3394486.3403226.

Qingfeng Li, Wei Liu, Zaiqiao Meng, and Jian Yin. Decision-Aware Preference Modeling for Multi-Behavior Recommendation. In *IJCAI*, pp. 3063–3071. International Joint Conferences on Artificial Intelligence Organization, 2025. doi: 10.24963/ijcai.2025/341.

Chang Meng, Chenhao Zhai, Yu Yang, Hengyu Zhang, and Xiu Li. Parallel Knowledge Enhancement based Framework for Multi-behavior Recommendation. In *CIKM*, pp. 1797–1806. ACM, 2023a. doi: 10.1145/3583780.3615004.

Chang Meng, Hengyu Zhang, Wei Guo, Huifeng Guo, Haotian Liu, Yingxue Zhang, Hongkun Zheng, Ruiming Tang, Xiu Li, and Rui Zhang. Hierarchical Projection Enhanced Multi-behavior Recommendation. In *KDD*, pp. 4649–4660. Association for Computing Machinery, 2023b. doi: 10.1145/3580305.3599838.

Chang Meng, Ziqi Zhao, Wei Guo, Yingxue Zhang, Haolun Wu, Chen Gao, Dong Li, Xiu Li, and Ruiming Tang. Coarse-to-Fine Knowledge-Enhanced Multi-Interest Learning Framework for Multi-Behavior Recommendation. *ACM Trans. Inf. Syst.*, 42(1):30:1–30:27, 2024. doi: 10.1145/3606369.

L. Mirsky. Symmetric gauge functions and unitarily invariant norms. *Q. J. Math.*, 11(1):50–59, January 1960. doi: 10.1093/qmath/11.1.50.

Shauli Ravfogel, Yanai Elazar, Hila Gonen, Michael Twiton, and Yoav Goldberg. Null it out: Guarding protected attributes by iterative nullspace projection. In *ACL*, pp. 7237–7256, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.647.

Shauli Ravfogel, Michael Twiton, Yoav Goldberg, and Ryan Cotterell. Linear Adversarial Concept Erasure. In *ICML*, volume 162 of *PMLR*, pp. 18400–18421. PMLR, 2022.

Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. BPR: Bayesian personalized ranking from implicit feedback. In *UAI*, pp. 452–461. AUAI Press, 2009.

Gobinda Saha, Isha Garg, and Kaushik Roy. Gradient projection memory for continual learning. In *ICLR*, 2021.

Shun Shao, Yftah Ziser, and Shay B. Cohen. Gold doesn't always glitter: Spectral removal of linear and nonlinear guarded attribute information. In *EACL*, pp. 1611–1622, Dubrovnik, Croatia, May 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.eacl-main.118.

Leheng Sheng, Changshuo Shen, Weixiang Zhao, Junfeng Fang, Xiaohao Liu, Zhenkai Liang, Xiang Wang, An Zhang, and Tat-Seng Chua. AlphaSteer: Learning refusal steering with principled null-space constraint. In *ICLR*, 2026.

Tianchi. User behavior data from Taobao for recommendation. Alibaba Tianchi, Dataset 649, n.d.a. Accessed: 2026-09-24.

Tianchi. Rec-Tmall: Tmall recommendation dataset. Alibaba Tianchi, Dataset 140281, n.d.b. Accessed: 2026-09-24.

Shipeng Wang, Xiaorong Li, Jian Sun, and Zongben Xu. Training networks in null space of feature covariance for continual learning. In *CVPR*, pp. 184–193, June 2021. doi: 10.1109/CVPR46437.2021.00025.

Wei Wei, Chao Huang, Lianghao Xia, Yong Xu, Jiashu Zhao, and Dawei Yin. Contrastive Meta Learning with Behavior Multiplicity for Recommendation. In *WSDM*, pp. 1120–1128. Association for Computing Machinery, 2022. doi: 10.1145/3488560.3498527.

Hermann Weyl. Das asymptotische Verteilungsgesetz der Eigenwerte linearer partieller Differentialgleichungen (mit einer Anwendung auf die Theorie der Hohlraumstrahlung). *Math. Ann.*, 71 (4):441–479, December 1912. doi: 10.1007/BF01456804.

Lianghao Xia, Chao Huang, Yong Xu, Peng Dai, Bo Zhang, and Liefeng Bo. Multiplex Behavioral Relation Learning for Recommendation via Memory Augmented Transformer Network. In *SIGIR*, pp. 2397–2406. ACM, 2020. doi: 10.1145/3397271.3401445.

Lianghao Xia, Chao Huang, Yong Xu, Peng Dai, Xiyue Zhang, Hongsheng Yang, Jian Pei, and Liefeng Bo. Knowledge-Enhanced Hierarchical Graph Transformer Network for Multi-Behavior Recommendation. In *AAAI*, volume 35, pp. 4486–4493, 2021. doi: 10.1609/aaai.v35i5.16576.

Jingcao Xu, Chaokun Wang, Cheng Wu, Yang Song, Kai Zheng, Xiaowei Wang, Changping Wang, Guorui Zhou, and Kun Gai. Multi-behavior Self-supervised Learning for Recommendation. In *SIGIR*, pp. 496–505. Association for Computing Machinery, 2023. doi: 10.1145/3539618.3591734.

Mingshi Yan, Zhiyong Cheng, Chen Gao, Jing Sun, Fan Liu, Fuming Sun, and Haojie Li. Cascading Residual Graph Convolutional Network for Multi-Behavior Recommendation. *ACM Trans. Inf. Syst.*, 42(1):10:1–10:26, 2024. doi: 10.1145/3587693.

Mingshi Yan, Zhiyong Cheng, Yahong Han, and Meng Wang. Latent Factor Modeling With Expert Network for Multi-Behavior Recommendation. *IEEE Trans. Knowl. Data Eng.*, 37(10):6020–6031, 2025a. doi: 10.1109/TKDE.2025.3591503.

Mingshi Yan, Zhiyong Cheng, Fan Liu, Yingda Lyu, and Yahong Han. User Invariant Preference Learning for Multi-Behavior Recommendation. *ACM Trans. Inf. Syst.*, 43(4):110:1–110:24, 2025b. doi: 10.1145/3728465.

Yi Yu, Tengyao Wang, and Richard J. Samworth. A useful variant of the Davis–Kahan theorem for statisticians. *Biometrika*, 102(2):315–323, June 2015. doi: 10.1093/biomet/asv008.

Guanxiong Zeng, Yang Chen, Bo Cui, and Shan Yu. Continual learning of context-dependent processing in neural networks. *Nat. Mach. Intell.*, 1(8):364–372, August 2019. doi: 10.1038/s42256-019-0080-x.

Chi Zhang, Rui Chen, Xiangyu Zhao, Qilong Han, and Li Li. Denoising and Prompt-Tuning for Multi-Behavior Recommendation. In *WWW*, pp. 1355–1363. Association for Computing Machinery, 2023. doi: 10.1145/3543507.3583513.

A Theoretical Details

We analyze the empirical target reference, projection-preserving calibration, and sensitivity to reference perturbations. All representation vectors are row vectors in $\mathbb{R}^{1\times d}$, with $1 \le r < d$. We use $\|\cdot\|_2$ for the Euclidean vector norm or spectral matrix norm, and $\|\cdot\|_F$ for the Frobenius norm. The results are conditional on the given representations, references, and gains; they do not require independence among learned representations.

A.1 Target Reference and Retained Association

For $a \in \{U, I\}$ indexing the user or item side, the matrices $\mathbf{X}_{b,a}, \mathbf{Y}_{b,a} \in \mathbb{R}^{n_{b,a}\times d}$ contain paired auxiliary and target representations of the same entities in matching row order. Pairing requires each entity to have training interactions under both behaviors, but does not require identical user–item interactions across behaviors. After separately centering the columns of each matrix, we compute

$$\mathbf{C}_{b,a} = \frac{\widetilde{\mathbf{X}}_{b,a}^{\top}\widetilde{\mathbf{Y}}_{b,a}}{n_{b,a}}, \qquad \mathbf{C}_b = \frac{1}{2}\left(\mathbf{C}_{b,U} + \mathbf{C}_{b,I}\right). \tag{13}$$

We assume $n_{b,U} > 0$, $n_{b,I} > 0$, and compatible user and item feature coordinates within each behavior. Centering applies only to reference estimation, while calibration acts on the original, uncentered encoder outputs. The reference-estimation protocol is specified in Appendix B.3.

Let $\mathbf{U}_{b,r}$ contain the leading r left singular vectors of $\mathbf{C}_b$, and define $\mathbf{P}_b = \mathbf{U}_{b,r}\mathbf{U}_{b,r}^{\top}$ and $\mathbf{Q}_b = \mathbf{I}_d - \mathbf{P}_b$. For a fixed orthogonal projector $\mathbf{\Pi}$, centering and projection commute: $\widetilde{\mathbf{X}_{b,a}\mathbf{\Pi}} = \widetilde{\mathbf{X}}_{b,a}\mathbf{\Pi}$. Thus, applying $\mathbf{\Pi}$ to the auxiliary matrices used for reference estimation gives cross-covariance

$$\frac{(\widetilde{\mathbf{X}}_{b,a}\mathbf{\Pi})^{\top}\widetilde{\mathbf{Y}}_{b,a}}{n_{b,a}} = \mathbf{\Pi}\mathbf{C}_{b,a},$$

and transforms the combined matrix into $\mathbf{\Pi}\mathbf{C}_b$. This motivates the following retained-association criterion.

Proposition 2 (Optimal retained cross-covariance energy)**.** *Among all rank-r orthogonal projectors, $\mathbf{P}_b$ maximizes retained cross-covariance energy, satisfying*

$$\|\mathbf{P}_b\mathbf{C}_b\|_F^2 = \max_{\substack{\mathbf{\Pi}^{\top}=\mathbf{\Pi},\ \mathbf{\Pi}^2=\mathbf{\Pi} \\ \operatorname{rank}(\mathbf{\Pi})=r}} \|\mathbf{\Pi}\mathbf{C}_b\|_F^2 = \sum_{j=1}^{r} \sigma_j(\mathbf{C}_b)^2, \tag{14}$$

where $\sigma_j(\mathbf{C}_b)$ is the jth largest singular value.

Proof. For any feasible projector,

$$\|\mathbf{\Pi}\mathbf{C}_b\|_F^2 = \operatorname{tr}\big(\mathbf{C}_b^{\top}\mathbf{\Pi}^{\top}\mathbf{\Pi}\mathbf{C}_b\big) = \operatorname{tr}\big(\mathbf{\Pi}\mathbf{C}_b\mathbf{C}_b^{\top}\big).$$

The Ky Fan variational principle bounds this trace by the sum of the largest r eigenvalues of $\mathbf{C}_b\mathbf{C}_b^{\top}$ (Fan, 1949). These eigenvalues are the squared leading singular values of $\mathbf{C}_b$, and $\mathbf{P}_b$ attains the bound. □

This criterion concerns association in the combined empirical matrix, rather than the predictive value of either representation component. It need not maximize the user- and item-side energies separately, since their associations can reinforce or cancel in the average. By the Eckart–Young–Mirsky theorem, $\mathbf{C}_{b,r} = \mathbf{P}_b\mathbf{C}_b$ is a best rank-at-most-r approximation in the Frobenius and spectral norms (Eckart & Young, 1936; Mirsky, 1960). The remaining cross-covariance satisfies

$$\|\mathbf{Q}_b\mathbf{C}_b\|_F^2 = \sum_{j=r+1}^{d} \sigma_j(\mathbf{C}_b)^2, \qquad \|\mathbf{Q}_b\mathbf{C}_b\|_2 = \sigma_{r+1}(\mathbf{C}_b). \tag{15}$$

Information outside the selected subspace can therefore retain target association. ***TNC*** keeps this information available as potentially complementary evidence and learns its contribution through target supervision.

Optimality requires neither full rank nor a spectral gap. The maximizing projector is unique when $\sigma_r(\mathbf{C}_b) > \sigma_{r+1}(\mathbf{C}_b)$. A boundary tie allows multiple optimal projectors. If $r > \operatorname{rank}(\mathbf{C}_b)$, the reference includes orthonormal directions associated with zero singular values that are not uniquely determined by the association criterion. The calibration identities below remain valid for any selected rank-r orthogonal projector, although different choices can produce different calibrated outputs when the gains differ.

A.2 Null-Space Calibration and Its Learning Signal

For a fixed entity and auxiliary behavior, abbreviate the input, projectors, and gains as $\mathbf{h}$, $\mathbf{P}$, $\mathbf{Q} = \mathbf{I}_d - \mathbf{P}$, α, and β. The calibration null space is $\mathcal{N}_b = \{\mathbf{v} \in \mathbb{R}^{1\times d} : \mathbf{vP} = \mathbf{0}\}$, with dimension $d - r$ and orthogonal projector $\mathbf{Q}$. It is defined by the retained projection and need not equal the left null space of the full cross-covariance matrix. The components $\mathbf{s} = \mathbf{hP}$ and $\mathbf{r} = \mathbf{hQ}$ are orthogonal and satisfy $\mathbf{h} = \mathbf{s} + \mathbf{r}$.

Projection-preserving calibration. The calibrated representation satisfies

$$\begin{aligned} \mathbf{z} &= \alpha\mathbf{s} + \beta\mathbf{r} = \alpha\mathbf{h} + (\beta - \alpha)\mathbf{hQ}, \\ \mathbf{zP} &= \alpha\mathbf{hP}, \qquad \mathbf{zQ} = \beta\mathbf{hQ}. \end{aligned} \tag{16}$$

For fixed $\mathbf{h}$, $\mathbf{P}$, and α, changing β adjusts the null-space contribution without changing the prescribed target-associated projection.

These identities follow directly from $\mathbf{P}^2 = \mathbf{P}$, $\mathbf{Q}^2 = \mathbf{Q}$, and $\mathbf{PQ} = \mathbf{QP} = \mathbf{0}$. They also yield the minimum-change characterization stated in the main text.

Proof of Proposition 1. Every vector satisfying $\mathbf{vP} = \alpha\mathbf{hP}$ can be written as $\mathbf{v} = \alpha\mathbf{hP} + \mathbf{w}$, with $\mathbf{wP} = \mathbf{0}$. Orthogonality gives

$$\|\mathbf{v} - \beta\mathbf{h}\|_2^2 = (\alpha - \beta)^2\|\mathbf{hP}\|_2^2 + \|\mathbf{w} - \beta\mathbf{hQ}\|_2^2. \tag{17}$$

The first term is fixed, and the second is uniquely minimized at the feasible vector $\mathbf{w} = \beta\mathbf{hQ}$. Hence $\mathbf{z} = \alpha\mathbf{hP} + \beta\mathbf{hQ}$ uniquely solves Eq. (12). □

Calibration is therefore the smallest change to $\beta\mathbf{h}$ that satisfies the prescribed projection. Equivalently, relative to $\alpha\mathbf{h}$, it adds an adjustment confined to $\mathcal{N}_b$. These interpretations explain the calibration geometry without introducing an additional training objective.

Bounded adjustment. With $\alpha = 1 + \tanh(g_s)$ and $\beta = 1 + \tanh(g_r)$, finite raw parameters give $0 < \alpha, \beta < 2$. Orthogonality yields

$$\begin{aligned} \|\mathbf{z}\|_2^2 &= \alpha^2\|\mathbf{s}\|_2^2 + \beta^2\|\mathbf{r}\|_2^2, \\ \|\mathbf{z} - \mathbf{h}\|_2^2 &= (\alpha - 1)^2\|\mathbf{s}\|_2^2 + (\beta - 1)^2\|\mathbf{r}\|_2^2. \end{aligned} \tag{18}$$

Consequently,

$$\|\mathbf{z}\|_2 \le 2\|\mathbf{h}\|_2, \qquad \|\mathbf{z} - \mathbf{h}\|_2 \le \|\mathbf{h}\|_2. \tag{19}$$

These bounds are relative to the current encoder output, rather than uniform bounds on learned representations throughout training. Initializing both raw gains at zero gives $\mathbf{z} = \mathbf{h}$ for every reference.

Learning evidence contributions. Let $\ell(\mathbf{z})$ denote the target ranking loss with its other inputs fixed. Holding $\mathbf{h}$ and $\mathbf{P}$ fixed, the chain rule gives

$$\begin{aligned} \frac{\partial \ell}{\partial g_s} &= (1 - \tanh^2 g_s)\langle \nabla_{\mathbf{z}}\ell, \mathbf{s}\rangle, \\ \frac{\partial \ell}{\partial g_r} &= (1 - \tanh^2 g_r)\langle \nabla_{\mathbf{z}}\ell, \mathbf{r}\rangle. \end{aligned} \tag{20}$$

Here the gradient is a row vector and $\langle \mathbf{v}, \mathbf{w}\rangle = \mathbf{vw}^\top$. The gradient $\nabla_{\mathbf{z}}\ell$ already includes backpropagation through fusion and scoring. These expressions describe the target-ranking contribution to the

gain gradients. The stated L_2 penalty additionally contributes $2\lambda g_s$ and $2\lambda g_r$ to the corresponding raw-parameter gradients. Each gain receives supervision through the loss gradient along its corresponding component. The two scalar gains calibrate the two components as wholes, rather than assigning a separate weight to every direction within either component. Orthogonal representation responses do not imply independent loss derivatives or uncoupled optimization.

Scope during joint training. The identities hold separately for every user and item, whose gains are learned under a shared behavior-specific reference. Pairing restricts reference estimation, not which entities can be calibrated. During joint training, gradients pass through calibration to the encoder outputs and gains, while the current reference is held fixed during backpropagation. The encoder, gains, and subsequently estimated reference can all change. Projection preservation is thus a conditional property of calibration, rather than a guarantee of unchanged target scores or fixed representations throughout training. The update protocol and treatment of entities without training purchases are described in Appendices B.3 and B.5, respectively.

A.3 Conditional Sensitivity to the Reference

For fixed gains, define $\mathcal{T}_{\alpha,\beta}(\mathbf{h}, \mathbf{P}) = \mathbf{h}[\beta \mathbf{I}_d + (\alpha - \beta)\mathbf{P}]$. For two inputs and orthogonal projectors, adding and subtracting $\mathcal{T}_{\alpha,\beta}(\mathbf{h}', \mathbf{P})$ gives

$$\begin{aligned} &\|\mathcal{T}_{\alpha,\beta}(\mathbf{h}, \mathbf{P}) - \mathcal{T}_{\alpha,\beta}(\mathbf{h}', \mathbf{P}')\|_2 \\ &\quad \le \max\{|\alpha|, |\beta|\}\|\mathbf{h} - \mathbf{h}'\|_2 + |\alpha - \beta|\, \|\mathbf{h}'\|_2 \|\mathbf{P} - \mathbf{P}'\|_2. \end{aligned} \tag{21}$$

The bound follows because the matrix in the calibration map acts as multiplication by α and β on the two orthogonal subspaces. With the chosen parameterization, calibration is at most 2-Lipschitz in its input for fixed gains and reference. Reference sensitivity is scaled by $|\alpha - \beta|$ and vanishes when the gains are equal.

Proposition 3 (Conditional reference stability)**.** *Let $\mathbf{C}, \mathbf{E} \in \mathbb{R}^{d\times d}$, $\mathbf{C}' = \mathbf{C} + \mathbf{E}$, and let $\mathbf{P}, \mathbf{P}'$ project onto the leading r-dimensional left singular subspaces of $\mathbf{C}, \mathbf{C}'$, respectively. Define*

$$\gamma = \sigma_r(\mathbf{C})^2 - \sigma_{r+1}(\mathbf{C})^2 > 0, \qquad \varepsilon = \|\mathbf{C}'\mathbf{C}'^\top - \mathbf{C}\mathbf{C}^\top\|_2. \tag{22}$$

If $\varepsilon < \gamma/2$, the perturbed leading subspace is unique and

$$\|\mathbf{P}' - \mathbf{P}\|_2 \le \frac{2\varepsilon}{\gamma}, \qquad \varepsilon \le (2\|\mathbf{C}\|_2 + \|\mathbf{E}\|_2)\|\mathbf{E}\|_2. \tag{23}$$

Proof. Set $\mathbf{A} = \mathbf{C}\mathbf{C}^\top$ and $\mathbf{A}' = \mathbf{C}'\mathbf{C}'^\top$, and order their eigenvalues decreasingly. Weyl's inequality (Weyl, 1912) gives

$$\lambda_r(\mathbf{A}') - \lambda_{r+1}(\mathbf{A}') \ge \gamma - 2\varepsilon > 0,$$

establishing uniqueness of the perturbed leading subspace. Moreover,

$$\lambda_r(\mathbf{A}) - \lambda_{r+1}(\mathbf{A}') \ge \gamma - \varepsilon > 0.$$

Thus, the leading eigenvalues of $\mathbf{A}$ are separated from the remaining eigenvalues of $\mathbf{A}'$ by at least $\gamma - \varepsilon$. The operator-norm Davis–Kahan bound (Davis & Kahan, 1970; Yu et al., 2015), together with the equality between the projector distance and the largest principal-angle sine for equal-dimensional subspaces, gives

$$\|\mathbf{P}' - \mathbf{P}\|_2 \le \frac{\varepsilon}{\gamma - \varepsilon} \le \frac{2\varepsilon}{\gamma}.$$

Finally,

$$\mathbf{A}' - \mathbf{A} = \mathbf{C}\mathbf{E}^\top + \mathbf{E}\mathbf{C}^\top + \mathbf{E}\mathbf{E}^\top.$$

The spectral-norm triangle inequality and submultiplicativity yield $\varepsilon \le 2\|\mathbf{C}\|_2\|\mathbf{E}\|_2 + \|\mathbf{E}\|_2^2$, which proves the second bound. □

Effect on calibrated outputs. For fixed $\mathbf{h}$, α, and β, the same assumptions give

$$\begin{aligned} &\|\mathcal{T}_{\alpha,\beta}(\mathbf{h},\mathbf{P}') - \mathcal{T}_{\alpha,\beta}(\mathbf{h},\mathbf{P})\|_2 \\ &= |\alpha-\beta|\,\|\mathbf{h}(\mathbf{P}'-\mathbf{P})\|_2 \\ &\le |\alpha-\beta|\,\|\mathbf{h}\|_2\|\mathbf{P}'-\mathbf{P}\|_2 \\ &\le \frac{2|\alpha-\beta|\,\|\mathbf{h}\|_2\,\varepsilon}{\gamma}. \end{aligned} \tag{24}$$

This bound connects perturbations of the empirical association matrix to changes in the calibrated output through the gain contrast, input norm, and reference spectral gap. When $\alpha = \beta$, calibration is independent of the reference. At a boundary tie, the reference perturbation bound does not apply, although the projection identities remain valid for any selected orthogonal projector. These results are conditional on the stated assumptions and do not guarantee ranking improvement or convergence of joint training.

A.4 Model Analysis

Time complexity. Let N denote the total number of users and items, A the number of auxiliary behaviors, and d the representation dimension. Let M denote the total number of paired entities used for reference estimation, summed over auxiliary behaviors and the user and item sides. Given the encoder outputs, constructing the cross-covariance matrices costs $O(Md^2)$. Using dense SVD, reference decomposition and construction of the rank-r projectors together cost $O(Ad^3)$, where $1 \le r < d$. Applying dense projectors to all auxiliary representations costs $O(ANd^2)$ per calibration pass. Residual computation and gain scaling require $O(ANd)$ and do not change this leading term. With S optimization steps per epoch, references are recomputed and calibration is applied at every step, giving

$$\mathcal{T}_{\text{epoch}} = \mathcal{T}_{\text{BaseModel}} + S\mathcal{T}_{\text{ref-enc}} + O\left(S\left[ANd^2 + Md^2 + Ad^3\right]\right), \tag{25}$$

where $\mathcal{T}_{\text{BaseModel}}$ denotes the computation required by the backbone and its original training objectives. The term $\mathcal{T}_{\text{ref-enc}}$ accounts for the additional encoder forward pass per step under the reference-estimation settings in Appendix B.3; it is zero when compatible forward settings allow encoder outputs to be reused. Reference estimation is detached from gradient computation.

Parameter count and inference. The module adds $2AN$ trainable gain parameters, while reference matrices are computed from encoder outputs rather than learned as additional parameters. At evaluation, references are recomputed at the selected checkpoint and calibrated representations can be precomputed. Once the calibrated representations are cached, scoring uses the same representation dimension and prediction operation as RMBRec.

B Experimental Settings

This section specifies the datasets, backbone integration, training and evaluation procedures, and protocols for the supplementary analyses. Results are presented separately in Appendix C.

B.1 Dataset Configuration

We evaluate on Taobao (Tianchi, n.d.a), Tmall (Tianchi, n.d.b), and Beibei (Gao et al., 2019), using purchases as the target behavior. Taobao and Beibei provide views and cart additions as auxiliary feedback, while Tmall additionally includes favorites. We denote Tmall clicks as View throughout the paper. Table 1 summarizes the processed datasets used in our experiments.

B.2 Backbone Integration

Backbone models. **MBGCN** (Jin et al., 2020) models behavior influence through user–item propagation and behavior semantics through item–item propagation. **CRGCN** (Yan et al., 2024) progressively refines representations across ordered behaviors using residual graph convolutions. **RMBRec** (Cai et al., 2026b) combines auxiliary–target user-representation alignment with regularization of prediction-risk variance across behaviors. **UIPL** (Yan et al., 2025b) combines a variational

autoencoder with invariant risk minimization over environments formed from behavior combinations. **SpectraMB** (Cai et al., 2026a) applies feature-frequency filtering and aggregates behavior representations through global-context attention. **MBLFE** (Yan et al., 2025a) extracts latent preference factors from graph-enhanced embeddings using a gated expert network and self-supervised constraints.

Calibration interface and prediction. Base denotes each backbone with its original prediction operations. The integrated model applies ***TNC*** to encoded auxiliary user and item representations before prediction; target representations bypass calibration. For the evaluated models, the downstream mappings in Eq. (10) use mean aggregation, following RMBRec (Cai et al., 2026b), and inner-product scoring,

$$\boldsymbol{\phi}_x = \frac{1}{K} \sum_{b \in \mathcal{B}} \mathbf{z}_{x,b}, \qquad \widehat{y}_{ui} = \boldsymbol{\phi}_u \boldsymbol{\phi}_i^\top. \tag{26}$$

This interface introduces no additional prediction parameters and replaces native aggregation or scoring where necessary. Reference construction assumes compatible user and item feature coordinates within each behavior, while aggregation assumes compatible coordinates across behaviors.

Training objectives and comparison scope. Encoder and calibration parameters are jointly trained from scratch using Eq. (11). Retained auxiliary objectives use their original weights and inputs, excluding target BPR and regularization terms already included in the joint loss. Objectives defined on behavior encodings use uncalibrated representations. For fixed encoder outputs, unit gains make calibration an identity mapping, but do not necessarily recover a backbone's original predictor when its native aggregation or scoring has been replaced. Accordingly, comparisons with Base evaluate the complete integration. Calibration variants retain the integrated prediction interface; behavior-removal variants additionally change its set of inputs.

B.3 Training and Evaluation Protocols

Reference estimation during training. At every optimization step, we recompute each auxiliary behavior's projector $\mathbf{P}_b$ from the current encoder representations of all paired training users $\mathcal{U}_{b,t}$ and items $\mathcal{I}_{b,t}$. Dropout and graph augmentation are disabled for this estimation. We write $\mathbf{H}^{\mathrm{ref}}$ for these representations and $\mathbf{H}$ for the representations used in gradient-based training. Both are produced by the same encoder parameters. This notation distinguishes forward settings and gradient paths without introducing another trainable encoder.

Reference projectors are held fixed during backpropagation. Gradients pass through calibration and prediction to $\mathbf{H}$ and the gain parameters, but not through reference estimation, cross-covariance computation, or SVD. Target representations bypass gain calibration while remaining trainable through the ranking objective. Disabling stochastic operations for reference estimation does not change the backbone's other training computations. When forward settings coincide, encoder outputs may be reused as long as only the reference-construction path is detached and the training path remains differentiable. Algorithm 1 summarizes the optimization procedure.

Reference estimation for final evaluation. At the selected checkpoint, we recompute projectors using the same reference-estimation procedure rather than reusing cached training projectors. Encoding and reference estimation use training interactions only; validation and test interactions are excluded.

Ranking evaluation. We report HR@10, NDCG@10, HR@20, and NDCG@20, following RMBRec (Cai et al., 2026b). NDCG accounts for the ranked positions of relevant items (Järvelin & Kekäläinen, 2002). Evaluation ranks the full candidate set without sampling evaluation negatives (Krichene & Rendle, 2020). For each user, training purchases are excluded from the candidates. The pair-weighted metric used in grouped analyses is specified separately in Appendix B.5.

Parameter selection. Backbone hyperparameters are selected on the validation set and shared between Base and the integrated model. We set the embedding dimension to $d = 64$ and select $r \in \{4, 8, 16, 32, 48\}$ by validation performance. Test results are used for reporting, not for selecting

Algorithm 1: Training Procedure of ***TNC***

Input: Training interactions $\{\mathbf{R}_b\}_{b\in\mathcal{B}}$; target behavior t; encoder f_θ; dimension $1 \le r < d$; epochs E; regularization weight λ.
Output: Optimized parameters $\boldsymbol{\omega} = (\theta, \mathbf{g})$.

Initialize θ and set $g_{x,b,s} = g_{x,b,r} = 0$ for each entity x and $b \neq t$
Set $K \leftarrow |\mathcal{B}|$ and identify all paired training sets $\mathcal{U}_{b,t}, \mathcal{I}_{b,t}$ for $b \neq t$
for $e = 1, \ldots, E$ **do**
  foreach *nonempty mini-batch $\mathcal{D}_t$ of sampled target triplets* **do**
    `// Recompute references from all paired training entities`
    Compute $\{\mathbf{H}^{\mathrm{ref}}_{U,b}, \mathbf{H}^{\mathrm{ref}}_{I,b}\}_{b\in\mathcal{B}}$ with the current f_θ, disabling dropout and graph augmentation; detach these outputs for reference construction
    foreach $b \in \mathcal{B} \setminus \{t\}$ **do**
      Form paired matrices $\mathbf{X}_{b,a}, \mathbf{Y}_{b,a}$ from $\mathbf{H}^{\mathrm{ref}}$ using all paired entities in matching order, and center their columns, for $a \in \{U, I\}$
      $\mathbf{C}_{b,a} \leftarrow \widetilde{\mathbf{X}}^\top_{b,a}\widetilde{\mathbf{Y}}_{b,a}/n_{b,a}, \quad a \in \{U, I\}$
      $\mathbf{C}_b \leftarrow (\mathbf{C}_{b,U} + \mathbf{C}_{b,I})/2$
      $\mathbf{U}_{b,r} \leftarrow$ leading r left singular vectors of $\mathbf{C}_b$
      $\mathbf{P}_b \leftarrow \mathbf{U}_{b,r}\mathbf{U}^\top_{b,r}, \quad \mathbf{Q}_b \leftarrow \mathbf{I}_d - \mathbf{P}_b$
    end
    `// Differentiate through training representations, not references`
    Compute differentiable representations $\{\mathbf{H}_{U,b}, \mathbf{H}_{I,b}\}_{b\in\mathcal{B}}$ with the backbone's training settings
    foreach $b \in \mathcal{B} \setminus \{t\}$ **do**
      foreach *user or item x* **do**
        $\alpha_{x,b} \leftarrow 1 + \tanh(g_{x,b,s}), \quad \beta_{x,b} \leftarrow 1 + \tanh(g_{x,b,r})$
        $\mathbf{s}_{x,b} \leftarrow \mathbf{h}_{x,b}\mathbf{P}_b, \quad \mathbf{r}_{x,b} \leftarrow \mathbf{h}_{x,b}\mathbf{Q}_b$
        $\mathbf{z}_{x,b} \leftarrow \alpha_{x,b}\mathbf{s}_{x,b} + \beta_{x,b}\mathbf{r}_{x,b}$
      end
    end
    $\mathbf{z}_{x,t} \leftarrow \mathbf{h}_{x,t}$ for every user or item x
    $\boldsymbol{\phi}_x \leftarrow K^{-1}\sum_{b\in\mathcal{B}} \mathbf{z}_{x,b}$ for every user or item x
    Compute $\widehat{y}_{ui^+} = \boldsymbol{\phi}_u\boldsymbol{\phi}^\top_{i^+}$ and $\widehat{y}_{ui^-} = \boldsymbol{\phi}_u\boldsymbol{\phi}^\top_{i^-}$ for $(u, i^+, i^-) \in \mathcal{D}_t$
    Compute $\mathcal{L}_{\mathrm{aux}}$ using its retained original inputs and weights
    $\mathcal{L}_{\mathrm{BPR}} \leftarrow -\frac{1}{|\mathcal{D}_t|}\sum_{(u,i^+,i^-)\in\mathcal{D}_t} \log\sigma(\widehat{y}_{ui^+} - \widehat{y}_{ui^-})$
    $\mathcal{L} \leftarrow \mathcal{L}_{\mathrm{BPR}} + \mathcal{L}_{\mathrm{aux}} + \lambda\|\boldsymbol{\omega}\|^2_2$
    Update $\boldsymbol{\omega}$ using $\nabla_{\boldsymbol{\omega}}\mathcal{L}$, keeping $\{\mathbf{P}_b\}_{b\neq t}$ fixed during backpropagation
  end
end
return $\boldsymbol{\omega}$

r. The selected rank can differ across dataset–backbone configurations; all RMBRec calibration ablations within a dataset use the same rank.

B.4 Ablation Protocol

We conduct ablations with RMBRec to examine reference construction, component retention, gain learning, calibration scope, and auxiliary sources. Results are reported in Appendix C.1.

Calibration and reference construction. The complete model retains the target-associated and null-space components and learns $\alpha_{x,b}$ and $\beta_{x,b}$ as defined in Eq. (8). The w/o Calibration variant bypasses the module and directly uses encoded auxiliary representations. The w/o Target Reference variant replaces the auxiliary–target cross-covariance reference with an auxiliary PCA reference constructed from the same paired entities. We center auxiliary representations separately on each side and form

$$\mathbf{C}^{\mathrm{PCA}}_b = \frac{1}{2}\left(\frac{\widetilde{\mathbf{X}}^\top_{b,U}\widetilde{\mathbf{X}}_{b,U}}{n_{b,U}} + \frac{\widetilde{\mathbf{X}}^\top_{b,I}\widetilde{\mathbf{X}}_{b,I}}{n_{b,I}}\right).$$

Its leading r eigenvectors define the auxiliary PCA projector. This variant retains the subspace dimension and gain parameterization, together with purchase interactions and supervision. For RMBRec, Base and w/o Calibration denote the same uncalibrated model, including mean fusion and inner-product scoring. They are retained as separately reported results; their numerical differences do not represent different prediction interfaces.

Component retention and gain learning. Fix $\alpha = 0$ removes the target-associated component, while Fix $\beta = 0$ removes the null-space component. Fix $\alpha = 1$ and Fix $\beta = 1$ retain both components but disable learning of the corresponding gain. In each case, the remaining gain is learned jointly with the encoder. The fixed-at-one variants therefore test gain learning without removing either representation component.

Calibration scope. User-only calibrates auxiliary user representations, while Item-only calibrates auxiliary item representations. The opposite side bypasses calibration. The complete model calibrates both sides, and target-behavior representations bypass the module in all these configurations.

Auxiliary sources. Each behavior-removal variant excludes one auxiliary behavior's interactions, associated objectives, and fusion branch. We remove views and carts separately on all three datasets, and additionally remove favorites on Tmall. Purchase interactions are retained in every variant. Mean aggregation uses the remaining behaviors, with K denoting their number.

Comparison scope and result reporting. The variant definitions above specify the intended interventions. The tables compare the reported results; interpreting differences for w/o Calibration, w/o Target Reference, Fix $\alpha = 0$, and Fix $\beta = 0$ as controlled intervention effects requires verification of matched data splits, training settings, and model selection. All tabulated relative changes in the ablation results use the complete model on the same dataset as the denominator and are calculated from the reported four-decimal scores.

B.5 Group-wise Evaluation Protocol

We construct user and item groups separately from training interactions and keep assignments fixed across compared models. Auxiliary-coverage groups characterize the overlap between auxiliary feedback and purchases, while purchase-degree groups characterize the amount of target feedback.

Auxiliary-coverage groups. We adapt the Behavioral Alignment Ratio (Cai et al., 2026b) by measuring the auxiliary coverage of training purchases. A training purchase is covered when the same user–item pair also appears under at least one auxiliary behavior. Let $\mathcal{I}^{\mathrm{tr}}_{u,t}$ be user u's training purchase items and $\mathcal{I}^{\mathrm{tr}}_{u,\mathrm{aux}}$ the union of items with auxiliary interactions from that user. Similarly, let $\mathcal{U}^{\mathrm{tr}}_{i,t}$ be item i's training purchasers and $\mathcal{U}^{\mathrm{tr}}_{i,\mathrm{aux}}$ the union of its auxiliary interacting users. For nonempty purchase sets, define

$$c_u = \frac{|\mathcal{I}^{\mathrm{tr}}_{u,t} \cap \mathcal{I}^{\mathrm{tr}}_{u,\mathrm{aux}}|}{|\mathcal{I}^{\mathrm{tr}}_{u,t}|}, \qquad c_i = \frac{|\mathcal{U}^{\mathrm{tr}}_{i,t} \cap \mathcal{U}^{\mathrm{tr}}_{i,\mathrm{aux}}|}{|\mathcal{U}^{\mathrm{tr}}_{i,t}|}. \tag{27}$$

Each purchase pair is counted once regardless of how many auxiliary behaviors cover it. These scores describe observed interaction overlap, not similarity between learned representations or the fraction of auxiliary interactions that lead to purchases.

Within each dataset, users and items with defined scores are grouped separately into Low-, Medium-, and High-coverage groups using approximate tertiles, keeping tied scores together. Items without training purchases have undefined scores and form the *No Training Purchases* group; an undefined score is distinct from a defined score of zero. Beibei is excluded from this analysis because all defined scores are one in the evaluated split.

Purchase-degree groups. User purchase activity and item purchase popularity are measured by

$$d_u = |\mathcal{I}^{\mathrm{tr}}_{u,t}|, \qquad d_i = |\mathcal{U}^{\mathrm{tr}}_{i,t}|. \tag{28}$$

Within each dataset, users and items are grouped separately into Low-, Medium-, and High-degree groups. Taobao additionally includes a Zero group for items without training purchases. Zero

purchase degree does not imply an absence of auxiliary interactions. The long-tail analysis compares RMBRec with and without ***TNC*** using the same grouped metric as the coverage analysis.

Ranking and aggregation. Let $\mathcal{T}$ denote the test purchase pairs and $\mathcal{T}_G \subseteq \mathcal{T}$ the nonempty subset assigned to group G by the relevant user or item grouping. For each user u, let $\mathcal{R}_u^{20}$ be the top-20 list from the full candidate set after excluding training purchases. Group Recall@20 is

$$R_G@20 = \frac{\sum_{(u,i)\in\mathcal{T}_G} \mathbb{I}\big[i \in \mathcal{R}_u^{20}\big]}{|\mathcal{T}_G|}. \tag{29}$$

Overall uses $\mathcal{T}_G = \mathcal{T}$ and includes all test purchase pairs. Grouping changes the aggregation of test hits, not the ranking candidates. This metric weights purchase pairs equally and may differ from a metric that first aggregates within users and then averages across users.

Contribution to the overall improvement. For nonempty, mutually exclusive groups $\mathcal{G}$ that partition $\mathcal{T}$, let

$$w_G = \frac{|\mathcal{T}_G|}{|\mathcal{T}|}, \qquad \Delta R_G = R_G^{\mathrm{TNC}} - R_G^{\mathrm{Base}},$$

where the cutoff 20 is omitted for brevity. With the same test pairs and model comparison throughout,

$$\Delta R_{\mathrm{Overall}} = \sum_{G\in\mathcal{G}} w_G \Delta R_G. \tag{30}$$

When $\Delta R_{\mathrm{Overall}} > 0$, group G's contribution fraction is

$$\eta_G = \frac{w_G \Delta R_G}{\Delta R_{\mathrm{Overall}}}. \tag{31}$$

Its contribution percentage is $100\eta_G\%$. We apply this decomposition to one complete grouping at a time; user and item groups cannot be combined because they overlap. Contributions are signed and may fall outside $[0, 1]$ when some groups deteriorate. They differ from the within-group relative improvement $100\Delta R_G / R_G^{\mathrm{Base}}\%$, defined when $R_G^{\mathrm{Base}} > 0$.

Items without training purchases. The pairing requirement applies to reference estimation, not to which entities can be calibrated. An item without training purchases can use a shared behavior-specific reference without entering the paired item set that estimates it. Its entity-specific gains can receive target-ranking gradients if it is sampled as a negative in the BPR objective. Thus, no positive training purchases does not imply complete absence from training, and access to a shared reference alone does not supply supervision to the item's gains. This analysis concerns items in the evaluated entity set and does not assess inductive prediction for newly arriving entities.

B.6 Representation Visualization Protocol

Representation visualization. For each dataset, we randomly sample 200 users and 200 items and extract behavior-specific representations immediately before and after calibration from the same trained RMBRec + ***TNC*** model. Within each dataset, we fit a two-dimensional PCA projection separately for users and items, jointly using all behaviors and both calibration stages for the relevant entity type. Sampled entities, PCA projections, and axis limits are fixed across each before–after comparison. Columns show users and items, rows show the two calibration stages, and ellipses summarize the projected distributions. Tmall is shown in Section 5.5; Taobao and Beibei are reported in Appendix C.4.

C Additional Experimental Analyses

We extend the main-text experiments with complete ablation results and analyses of auxiliary coverage, purchase sparsity, and representation geometry. Experimental protocols are provided in Appendix B.

Table 4: Calibration ablations with RMBRec. Percentages indicate relative changes from ***+TNC***, calculated from the reported four-decimal scores, with decreases in red. The shaded row is the complete model, which calibrates both users and items.

	Taobao				Tmall				Beibei			
Variant	HR@10	NDCG@10	HR@20	NDCG@20	HR@10	NDCG@10	HR@20	NDCG@20	HR@10	NDCG@10	HR@20	NDCG@20
Base	0.1453	0.0955	0.1774	0.1036	0.1165	0.0648	0.1602	0.0755	0.0848	0.0438	0.1359	0.0566
+TNC	**0.1621**	**0.1070**	**0.1958**	**0.1155**	**0.1279**	**0.0719**	**0.1715**	**0.0825**	**0.0912**	**0.0472**	**0.1414**	**0.0599**
w/o Calibration	0.1454 (-10.30%)	0.0955 (-10.75%)	0.1772 (-9.50%)	0.1036 (-10.30%)	0.1167 (-8.76%)	0.0649 (-9.74%)	0.1597 (-6.88%)	0.0755 (-8.48%)	0.0848 (-7.02%)	0.0438 (-7.20%)	0.1359 (-3.89%)	0.0566 (-5.51%)
w/o Target Reference	0.1456 (-10.18%)	0.0956 (-10.65%)	0.1772 (-9.50%)	0.1036 (-10.30%)	0.1170 (-8.52%)	0.0649 (-9.74%)	0.1602 (-6.59%)	0.0755 (-8.48%)	0.0826 (-9.43%)	0.0421 (-10.81%)	0.1329 (-6.01%)	0.0548 (-8.51%)
Fix $\beta = 0$	0.0486 (-70.02%)	0.0279 (-73.93%)	0.0681 (-65.22%)	0.0328 (-71.60%)	0.0571 (-55.36%)	0.0291 (-59.53%)	0.0872 (-49.15%)	0.0364 (-55.88%)	0.0500 (-45.18%)	0.0249 (-47.25%)	0.0860 (-39.18%)	0.0339 (-43.41%)
Fix $\alpha = 0$	0.1194 (-26.34%)	0.0787 (-26.45%)	0.1489 (-23.95%)	0.0862 (-25.37%)	0.0994 (-22.28%)	0.0559 (-22.25%)	0.1390 (-18.95%)	0.0656 (-20.48%)	0.0704 (-22.81%)	0.0357 (-24.36%)	0.1154 (-18.39%)	0.0469 (-21.70%)
Fix $\alpha = 1$	0.1487 (-8.27%)	0.0990 (-7.48%)	0.1789 (-8.63%)	0.1067 (-7.62%)	0.1150 (-10.09%)	0.0649 (-9.74%)	0.1580 (-7.87%)	0.0754 (-8.61%)	0.0854 (-6.36%)	0.0448 (-5.08%)	0.1341 (-5.16%)	0.0570 (-4.84%)
Fix $\beta = 1$	0.1255 (-22.58%)	0.0821 (-23.27%)	0.1566 (-20.02%)	0.0899 (-22.16%)	0.0967 (-24.39%)	0.0531 (-26.15%)	0.1362 (-20.58%)	0.0628 (-23.88%)	0.0643 (-29.50%)	0.0325 (-31.14%)	0.1067 (-24.54%)	0.0431 (-28.05%)
User-only	0.1521 (-6.17%)	0.1005 (-6.07%)	0.1855 (-5.26%)	0.1089 (-5.71%)	0.1249 (-2.35%)	0.0694 (-3.48%)	0.1704 (-0.64%)	0.0806 (-2.30%)	0.0890 (-2.41%)	0.0459 (-2.75%)	0.1402 (-0.85%)	0.0588 (-1.84%)
Item-only	0.1573 (-2.96%)	0.1039 (-2.90%)	0.1891 (-3.42%)	0.1119 (-3.12%)	0.1203 (-5.94%)	0.0669 (-6.95%)	0.1632 (-4.84%)	0.0774 (-6.18%)	0.0882 (-3.29%)	0.0454 (-3.81%)	0.1364 (-3.54%)	0.0575 (-4.01%)

C.1 Ablation Analysis

To examine how the calibration design and auxiliary sources affect purchase prediction, we evaluate the RMBRec variants defined in Appendix B.4 on all three datasets. Tables 4 and 5 report calibration-level and behavior-level results, respectively.

The target reference yields higher reported scores than auxiliary PCA. The complete model has higher reported scores than w/o Target Reference on all four metrics across the three datasets. The PCA variant's NDCG@10 is lower by 10.65%, 9.74%, and 10.81% on Taobao, Tmall, and Beibei, respectively. The variant is defined to retain the paired entities, subspace dimension, gain parameterization, and purchase supervision; the comparison scope is specified in Appendix B.4. Both references recover uniform rescaling when $\alpha = \beta$, but yield reference-dependent transformations when the gains differ. Similar ranking scores alone do not establish that learned gains remain at their initial values.

Both component-removal variants score below the uncalibrated model. Both component-removal variants perform worse than w/o Calibration on all four metrics across the three datasets. The variant without the null-space component consistently has the larger gap. On Taobao, NDCG@10 is 0.0955 without calibration, 0.0787 for Fix $\alpha = 0$, and 0.0279 for Fix $\beta = 0$. These results are consistent with retaining auxiliary evidence both within and outside the selected target-associated subspace for purchase prediction. Because the components can differ in dimensionality and energy, the unequal score gaps do not measure their relative information content.

Learning both gains improves ranking, with larger losses when the null-space gain is fixed. The complete model outperforms both fixed-at-one variants on every metric across all three datasets. Fixing $\beta = 1$ reduces NDCG@10 by 23.27%, 26.15%, and 31.14% on Taobao, Tmall, and Beibei, respectively, compared with 7.48%, 9.74%, and 5.08% when fixing $\alpha = 1$. Both components remain available in these variants; the results favor jointly learning their gains over fixing either gain at one in the evaluated configurations. The larger losses under Fix $\beta = 1$ highlight the importance of calibrating the null-space contribution under purchase supervision.

Calibrating both users and items consistently outperforms calibrating either side alone. The complete model achieves higher scores than User-only and Item-only on all four metrics across the three datasets. Its NDCG@10 exceeds that of the stronger single-sided variant by 2.98%, 3.60%, and 2.83% on Taobao, Tmall, and Beibei, respectively. Item-only is stronger on Taobao, whereas User-only is stronger on Tmall and Beibei. Thus, the more effective side depends on the dataset, while calibrating both sides consistently improves user–item ranking.

Table 5: Behavior-level ablations with RMBRec. Percentages indicate relative changes from ***+TNC***, calculated from the reported four-decimal values. Increases are teal and decreases are red; bold marks the best value in each column. Favorites are evaluated only on Tmall.

	Taobao				Tmall				Beibei			
Variant	HR@10	NDCG@10	HR@20	NDCG@20	HR@10	NDCG@10	HR@20	NDCG@20	HR@10	NDCG@10	HR@20	NDCG@20
Base	0.1453	0.0955	0.1774	0.1036	0.1165	0.0648	0.1602	0.0755	0.0848	0.0438	0.1359	0.0566
+TNC	**0.1621**	**0.1070**	**0.1958**	**0.1155**	0.1279	0.0719	0.1715	0.0825	**0.0912**	**0.0472**	**0.1414**	**0.0599**
w/o View	0.1240 (-23.50%)	0.0872 (-18.50%)	0.1368 (-30.13%)	0.0904 (-21.73%)	0.0584 (-54.34%)	0.0358 (-50.21%)	0.0757 (-55.86%)	0.0400 (-51.52%)	0.0768 (-15.79%)	0.0395 (-16.31%)	0.1221 (-13.65%)	0.0509 (-15.03%)
w/o Cart	0.1289 (-20.48%)	0.0760 (-28.97%)	0.1677 (-14.35%)	0.0858 (-25.71%)	0.1267 (-0.94%)	**0.0723** (+0.56%)	0.1715 (0.00%)	**0.0833** (+0.97%)	0.0604 (-33.77%)	0.0308 (-34.75%)	0.1008 (-28.71%)	0.0409 (-31.72%)
w/o Favorites	—	—	—	—	**0.1298** (+1.49%)	0.0671 (-6.68%)	**0.1855** (+8.16%)	0.0808 (-2.06%)	—	—	—	—

The contribution of an auxiliary behavior depends on the dataset and ranking metric. Removing views or carts reduces all four metrics on Taobao and Beibei. Cart removal causes the larger NDCG@10 decrease, reaching 28.97% on Taobao and 34.75% on Beibei, compared with 18.50% and 16.31% for view removal. On Taobao, the ordering reverses for HR@20, with a 30.13% decrease after removing views and a 14.35% decrease after removing carts. These differences show that auxiliary sources affect retrieval and ranking to different degrees.

Tmall exhibits a different pattern. Removing views reduces NDCG@10 by 50.21%, whereas removing carts increases NDCG@10 and NDCG@20 by 0.56% and 0.97%, respectively. The latter change also reduces HR@10 and leaves the reported HR@20 unchanged. Removing favorites increases HR@10 and HR@20 by 1.49% and 8.16%, but decreases NDCG@10 and NDCG@20 by 6.68% and 2.06%. Together, these results motivate learning the use of auxiliary sources under the target objective and show that including an additional behavior does not guarantee better performance on every metric.

C.2 Auxiliary-Coverage Analysis

To determine whether calibration benefits entities with different degrees of auxiliary–purchase overlap, we evaluate the user and item groups defined in Appendix B.5. Figure 5 presents all six backbones on Taobao and Tmall; the quantitative findings below focus on RMBRec.

Calibration improves RMBRec across user and item coverage levels. ***TNC*** improves Group Recall@20 in every reported RMBRec group, with overall gains of 9.75% on Taobao and 7.33% on Tmall. The Low-, Medium-, and High-coverage groups improve by 4.71%–10.06% across the two datasets. The benefits therefore extend to entities with both limited and extensive overlap between auxiliary interactions and training purchases. Improvements in the High-coverage groups further show that substantial interaction overlap leaves room for better use of auxiliary representations.

Items without training purchases contribute disproportionately to the improvement on Taobao. For RMBRec, Group Recall@20 on these items increases from 0.11110 to 0.16900, corresponding to a relative gain of approximately 52.1%. Although they account for only 7.60% of test purchase pairs, they contribute approximately 25.42% of the overall absolute Recall increase. Items with training purchases also improve by 7.63% when evaluated together. Thus, calibration benefits both items with observed purchase histories and items without positive purchase observations during training.

C.3 Long-Tail Analysis

To assess the effect of target-feedback sparsity, we compare RMBRec with and without ***TNC*** across training purchase-degree groups. Figure 6 reports Group Recall@20 using the protocol in Appendix B.5.

Calibration benefits users across all purchase-activity groups. Every user group improves on all three datasets. Low-activity users obtain the largest relative gains on Taobao and Tmall, reaching 10.40% and 10.03%, respectively. On Beibei, the three groups improve by 1.15%–2.34%, with no

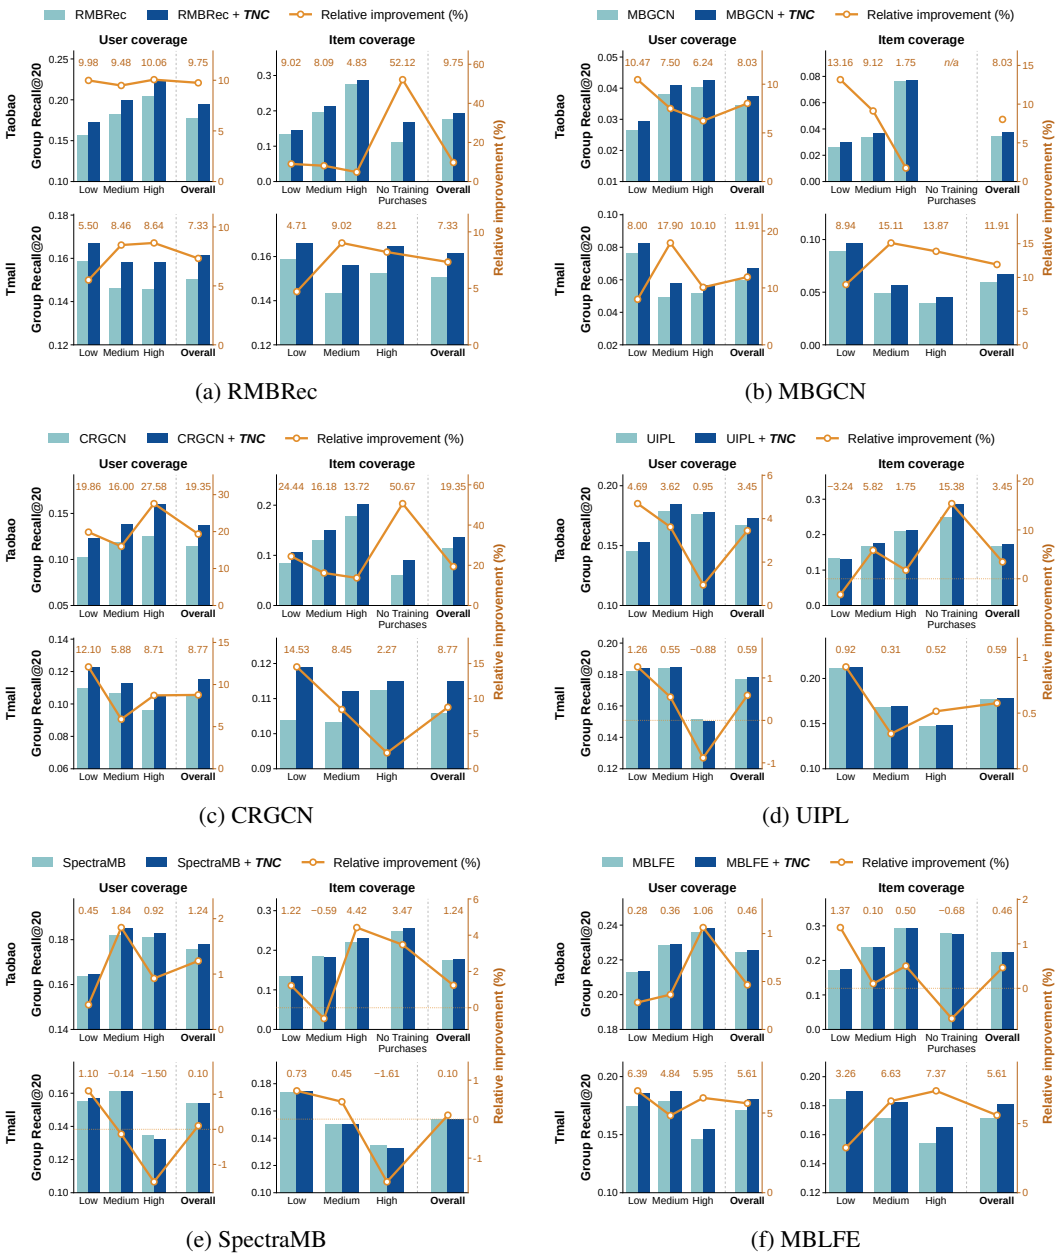

(a) RMBRec (b) MBGCN

(c) CRGCN (d) UIPL

(e) SpectraMB (f) MBLFE

Figure 5: Auxiliary-coverage analysis across six backbones. Group Recall@20 is reported with and without ***TNC***. Coverage is the fraction of training purchases with auxiliary interactions on the same user–item pair. Within each backbone panel, rows correspond to Taobao and Tmall, and columns correspond to user and item groups.

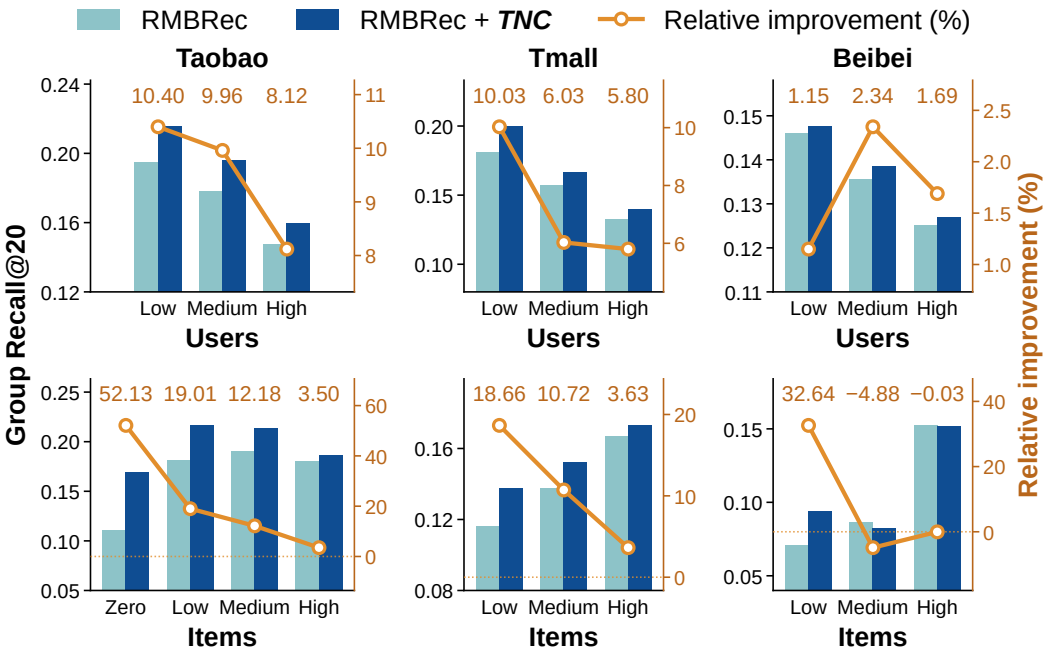

Figure 6: Performance across training purchase-degree groups. The top row groups users by purchase activity, and the bottom row groups items by purchase popularity. Bars report Group Recall@20 on the left axes, and orange lines show relative changes over RMBRec on the right axes. Zero denotes items without training purchase interactions.

monotonic relationship between activity and improvement. These findings show that auxiliary calibration benefits both users with sparse purchase histories and those with more extensive feedback.

Low-popularity items obtain the largest relative gains among items with training purchases. The Low groups improve by 19.01%, 18.66%, and 32.64% on Taobao, Tmall, and Beibei, respectively. On Taobao and Tmall, the relative gains decrease as purchase degree increases. Taobao's Zero group also improves by approximately 52.1%, extending the benefit to items without positive training purchases. On Beibei, the Medium group decreases by 4.88%, while the High group changes by -0.03%. The concentration of relative gains among low-degree items is consistent with using auxiliary evidence when purchase feedback is limited. These results characterize the complete calibration module's benefits; the component ablations examine the role of null-space calibration.

C.4 Representation Geometry Analysis

To examine how calibration changes behavior representations, Figure 7 extends the main-text Tmall visualization to Taobao and Beibei. Each comparison uses the same sampled 200 users or 200 items, PCA projection, and axis limits before and after calibration, following Appendix B.6.

Calibration preserves distributional differences between auxiliary and purchase representations. On Taobao, view and cart distributions expand on both the user and item sides, with a larger expansion for views. Substantial overlap remains, so the change concerns distributional extent rather than the formation of separate behavior clusters. On Beibei, auxiliary user representations spread over a larger region, and their positional differences from purchases become more visible. For items, view representations extend further beyond the purchase distribution along the second principal component, while cart representations continue to overlap with both behaviors. These observations show that calibration changes different behaviors and entity types differently, without uniformly concentrating auxiliary representations around the purchase distribution.

Purchase representations remain unchanged within each comparison because they bypass calibration, although the target encoder remains trainable during optimization. The observed geometry

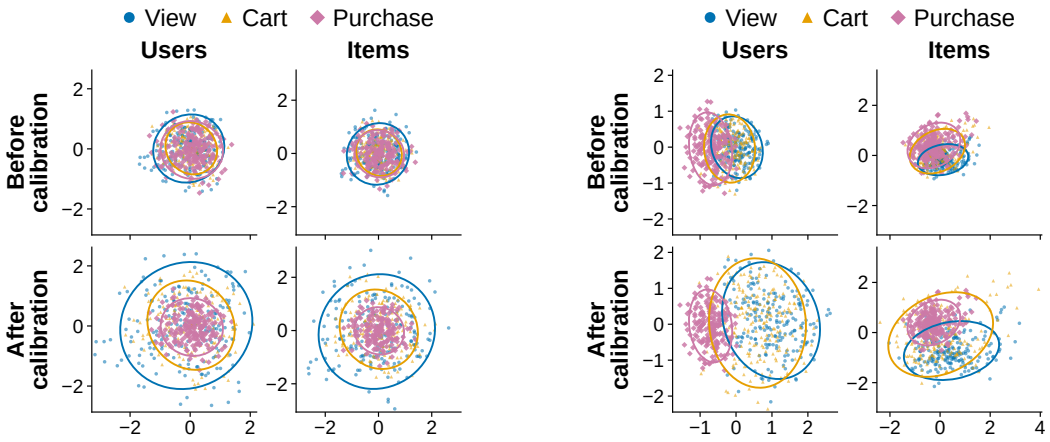

Figure 7: Additional representation visualizations with RMBRec + ***TNC*** on Taobao (left) and Beibei (right). Within each dataset, columns correspond to users and items, and rows show representations before and after calibration. Each column uses the same sampled entities, PCA projection, and axis limits across stages. Colors and markers distinguish behaviors; ellipses summarize their projected distributions.

is consistent with retaining behavior-specific structure while adjusting auxiliary representations for target prediction. Ranking and component-ablation results provide the accompanying evidence for recommendation effectiveness; the two-dimensional projections alone do not identify the predictive value of individual directions.

D Future Work

We view this study as an initial step toward target-referenced calibration as a general interface between heterogeneous representation learning and recommendation. Future work will investigate adaptive reference spaces that reflect user context and the evidence available across behaviors, modalities, and domains. Another promising direction is to model interactions among information sources during calibration, accounting for redundancy and allowing cues that become useful only in combination to support target prediction. Motivated by the gains for items without training purchases, we will further explore how such calibration can support preference learning when direct target feedback is scarce or unavailable.