FSLNet: Episodic Few-Shot Adaptation for Industrial Recommender Systems

Maolei Huang , Shuhan Song, Ming Dun, Zhenlong Yuan, Huawei Cao
Alibaba Group,
University of Chinese Academy of Sciences (UCAS) ,
State Key Lab of Processors, Institute of Computing Technology (CAS)

ICDM'2026

Background

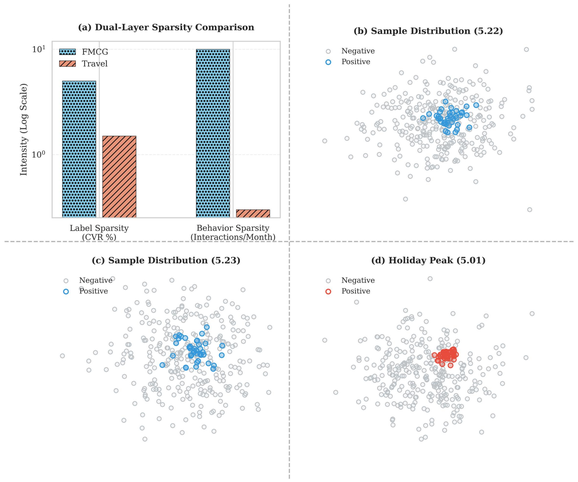

In large-scale industrial recommendation systems, accurately predicting Click-Through Rate (CTR) and Conversion Rate (CVR) is fundamental to optimizing advertising revenue and user experience. While deep learning has succeeded in high-frequency e-commerce, the travel industry presents a unique "high-ticket, low-frequency" challenge characterized by high transaction costs and prolonged decision cycles spanning weeks or months. Unlike standard retail where both exposure-to-click (CTR) and click-to-purchase (CVR) feedback loops are dense and rapid, travel scenarios suffer from an extreme imbalance: while exposure-to-click interactions are relatively observable, the actual click-to-purchase ratio drops to a negligible 1.5%. This massive gap between initial user clicks and final conversions results in extreme label sparsity, rendering traditional deep learning models that rely on dense, consecutive interaction sequences largely ineffective.

Challenge

LIMITATIONS

1. **Dual-Layer Sparsity (Label & Behavior) :** Extremely low click-to-purchase conversion rates coupled with highly limited annual user interaction footprints cause traditional deep learning models to overfit on negative samples and fail.
2. **Non-IID Distribution Shifts (Temporal Noise) :** Although aggregating long-window data mitigates sparsity, it introduces severe noise because travel intent shifts non-stationarily due to seasonality and holidays, rendering static models suboptimal.
3. **Metric-Optimization Gap (BCE vs. AUC) :** Traditional point-wise classification losses fail to directly optimize global ranking metrics like AUC, creating a gap between offline model training objectives and actual online ranking performance.
4. **Industrial Deployment Constraints (Sub-40ms Latency):** Any complex framework designed to solve sparsity must still satisfy stringent industrial production requirements, specifically maintaining a real-time inference latency of under 40 milliseconds under high concurrency.

Architecture Design

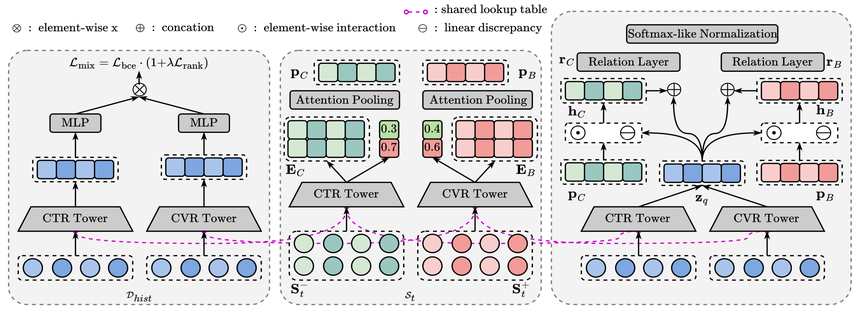

Problem Definition & Feature Embedding and Domain-Adaptive Pre-training

To handle non-IID distribution shifts in travel CVR estimation, the task is formulated as Context-Aware Few-Shot Learning. By utilizing a dynamic support set of the T most recent interactions,

$$\mathcal{S}_t = \{(\mathbf{x_j}, y_j)\}_{j=t-T}^{t-1}$$

partitioned into conversion $\mathcal{S}_t^+$ and exploratory $\mathcal{S}_t^-$ subsets, the model decouples predictions from stale statistics by mapping a query $\mathbf{x_q}$ to its conversion probability via

$$\widehat{y_q} = f(\mathbf{x_q}, \mathcal{S}_t; \Theta).$$

Dynamic Prototype Construction with Temporal Awareness

Recognizing that casual clicks in the support set introduce significant noise, a Salience-Aware Attention Pooling mechanism computes dynamic contribution weights α_i for each sample embedding $\mathbf{z_i} \in \mathcal{S}_t$ using the query-independent attention score:

$$e_i = \mathbf{u}^{\mathbf{T}} \tanh(\mathbf{W_a z_i} + \mathbf{b_a}), \quad \alpha_i = \frac{\exp(e_i)}{\sum_{j \in \mathcal{S}_t} \exp(e_j)}$$

Using these weights to filter out incidental clicks, the model constructs time-sensitive class-specific prototypes for booking intent

$$\mathbf{p_B} = \sum_{\mathbf{x_i} \in \mathcal{S}_t^+} \alpha_i \mathbf{z_i}$$

and exploratory clicks

$$\mathbf{p_C} = \sum_{\mathbf{x_j} \in \mathcal{S}_t^-} \alpha_j \mathbf{z_j}$$

, allowing the model to adapt instantaneously to sudden distribution shifts during holidays or seasonal promotions.

Metric-based Interaction

To overcome the limitations of rigid, static distance metrics (such as Cosine similarity) that assume an isotropic latent space, the Relation Layer computes complex non-linear interactions by constructing an augmented interaction space

$$\mathbf{h_{q,k}} = \text{Concat}\left(\mathbf{z_q}, \mathbf{p_k}, (\mathbf{z_q} \ominus \mathbf{p_k}), (\mathbf{z_q} \odot \mathbf{p_k})\right), k \in \{B, C\}.$$

A learnable MLP-based relation head then evaluates the affinity score through

$$r_k = \text{MLP}_{rel}(\mathbf{h_{q,k}}; \Psi),$$

which yields the final temporally-sensitive conversion probability via the softmax-like normalization:

$$\widehat{y_{rel}} = \frac{\exp(r_B)}{\exp(r_B) + \exp(r_C)}$$

Objective Alignment and Optimization

$$\mathcal{L}_{rank} = \frac{1}{|\mathcal{D}^+||\mathcal{D}^-|} \sum_{i \in \mathcal{D}^+} \sum_{j \in \mathcal{D}^-} \left[\text{clip}(\tau - (v_i - v_j), 0, \infty)\right]^2$$

Offline Evaluation

TABLE III
COMPREHENSIVE PERFORMANCE COMPARISON ON THE INDUSTRIAL PRODUCTION DATASET. **BOLD** INDICATES THE BEST PERFORMANCE; UNDERLINE DENOTES THE SECOND-BEST RESULTS.

Methods	$GAUC_{avg}$	$GAUC_{show}$	$GAUC_{click}$	AUC	LogLoss	Recall@10	MRR	RPM	Δ $GAUC_{avg}$
GBDT	65.41%	66.27%	70.01%	73.88%	0.1331	0.441	0.306	10.01	-
XGboost	66.31%	67.52%	70.98%	74.66%	0.1280	0.448	0.310	10.10	-
WDL	67.04%	68.35%	71.32%	75.43%	0.1245	0.452	0.312	10.23	-
DIN	67.89%	70.06%	72.06%	76.16%	0.1198	0.461	0.320	10.58	-
MMoE	67.56%	69.11%	71.12%	75.88%	0.1210	0.458	0.318	10.41	-
ESMM	68.05%	69.18%	71.23%	76.04%	0.1182	0.465	0.325	10.65	-
PLE	67.99%	69.15%	71.89%	75.91%	0.1190	0.463	0.322	10.62	-
DROS	68.19%	70.04%	72.31%	76.34%	0.1175	0.470	0.328	10.74	-
IDEA	69.08%	70.89%	72.91%	77.18%	0.1148	0.485	0.340	11.10	-
COR	69.79%	70.12%	72.89%	78.13%	0.1142	0.489	0.342	11.25	-
FSLNet (Ours)	**70.63%**	**71.33%**	**74.16%**	**79.32%**	**0.1120**	**0.501**	**0.358**	**11.58**	**+0.84%**

TABLE I
EFFICIENCY COMPARISON OF DIFFERENT METHODS (M: MINUTE; S: SECOND; MS: MILLISECOND).

Method	Training Time	Inference Latency
WDL	78m 58s	28.82ms
DIN	80m 42s	29.17ms
ESMM	78m 03s	28.74ms
PLE	85m 16s	29.16ms
DROS	81m 22s	26.31ms
FSLNet (Ours)	**83m 17s**	**28.62ms**

TABLE II
STATISTICS OF INDUSTRY PRODUCTION DATASET (M-MILLION).

Feature	Training	Validation	Testing
# of samples	11.5M	1.5M	1.8M
# of click samples	1.5M	0.5M	0.8M
# of pay samples	0.2M	0.1M	0.1M
# of negative samples	10M	1M	1M
# of users	2M	1M	1M
# of items	0.2M	0.2M	0.2M

TABLE IV
PERFORMANCE DEGRADATION (AUC) UNDER DISTRIBUTION SHIFTS. Δ QUANTIFIES THE SENSITIVITY TO NON-IID SHIFTS; SMALLER Δ INDICATES SUPERIOR ADAPTATION.

Variant	Stable	Holiday	Drop (Δ)
COR	78.13%	76.54%	1.59%
Full FSLNet	**79.32%**	**78.56%**	**0.76%**
w/o Dynamic Prototypes	78.55%	76.12%	2.43%
w/o Metric-based Relation	78.21%	76.68%	1.53%
w/o Mix Loss	78.78%	77.21%	1.57%

TABLE V
SENSITIVITY ANALYSIS OF t AND τ (AUC SCORE).

Margin (τ)	Window Size (t)				
	1	2	3	4	5
0.1	0.754	0.768	0.775	0.771	0.762
0.3	0.761	0.775	0.788	0.782	0.770
0.5	0.765	0.782	**0.793**	0.783	0.781
0.7	0.758	0.771	0.784	0.779	0.768
1.0	0.745	0.762	0.772	0.778	0.755

Tested on a highly sparse 180-day travel dataset (11.5M samples, 1:50 conversion ratio), FSLNet achieves SOTA ranking accuracy ($GAUC_{avg}$:70.63%, AUC:79.32%) and calibration (LogLoss:0.1120). It exhibits outstanding temporal robustness, limiting AUC drop to just 0.76% during holiday peaks. Ablation and sensitivity analyses further validate that our dynamic episodic modules and Mix Loss successfully bridge the metric-optimization gap while maintaining peak prediction stability.

System Deployment Design & Online A/B Test

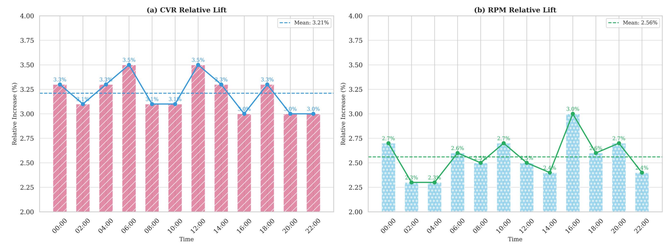

During a two-week online A/B test on a major travel platform (10% traffic), FSLNet achieved a +4.7% overall CVR increase, maintaining a 3.21% average CVR lift and 2.56% RPM gain across 24-hour cycles. By leveraging real-time intent anchors $\{\mathbf{p_B}, \mathbf{p_C}\}$to prevent peak-hour performance decay, these results confirm FSLNet's robustness and strong commercial viability for large-scale production deployment.