Online Resource 1

Supplementary material for

Adaptive Lagrangian Control of Long-Tail Exposure in Graph Recommendation

Puskar Adhikari* and Chantri Polprasert*
Faculty of Advanced Science and Technology, Asian Institute of Technology, Khlong Luang, Pathum Thani 12120, Thailand
*Corresponding authors: puskar.adhikari@alumni.ait.asia, chantri@ait.asia

Knowledge and Information Systems

Section, table, and equation numbers without the prefix "S" refer to the main paper.

S1 Hyperparameters and Implementation Details

We separate hyperparameters into two groups: those held fixed across all (dataset, backbone) configurations (Table S1) and those selected once per configuration on a validation subset (Table S2). Tuned values are then frozen for all subsequent seeds and operating points of that configuration; we do not re-tune them per target or per seed.

Table S1: Hyperparameters held fixed across configurations for the in-training methods (BASE, POP, PBiLoss, ALF-S, and ALF-P).

Hyperparameter	**Value**
Embedding dimension d	64
Hidden dimension	64
Dropout	0.0
Optimizer	Adam
L_2 embedding regularization γ	1×10^{-4}
Negatives per positive N_{neg}	100
Negative-sampling popularity exponent	0.75
Ranking temperature t_{r}	0.10
Gradient clip norm	5.0
Candidate-pool cap $S_{\max}$	5000
Head fraction h (diagnostic reporting)	10%
Tail fraction ℓ	20%
Minimum interactions for tail eligibility $c_{\min}$	5
Random seeds	$\{42, 43, 44, 45, 46\}$

Two configuration arguments are accepted by the released implementation but do not affect any reported result. First, the soft tail-mass proxy accepts a tail-margin argument; this value is not subsequently used in the loss computation (§3.2, Eq. 5). Second, the batch sampler accepts an inverse-popularity distribution and a nominal head/tail mixing ratio of 0.8/0.2, but the reported runs draw negative candidates exclusively from the popularity-weighted distribution. These arguments are therefore inactive code paths in the reported protocol; the effective sampling mechanism is given in §3.5.

Table S2: Per-(dataset, backbone) hyperparameters. LGCN = LightGCN; GS = GraphSAGE. Head-related dual parameters are retained for diagnostic logging; the corresponding surrogate contributes no primal gradient in the reported experiments (see §3.2). Maximum epochs and patience are those of the ALF runs; BASE and POP on MovieLens-1M and Yelp2018 used a 500-epoch cap (§4.6).

	ML-100K		**ML-1M**		**Yelp2018**	
Hyperparameter	**LGCN**	**GS**	**LGCN**	**GS**	**LGCN**	**GS**
Architecture and training						
GNN layers L	2	2	2	2	2	2
Batch size B	2048	2048	2048	2048	8192	8192
Learning rate	1×10^{-3}	2×10^{-3}	1×10^{-3}	1×10^{-3}	1×10^{-3}	1×10^{-3}
Weight decay	10^{-4}	10^{-4}	10^{-4}	10^{-4}	10^{-4}	10^{-4}
Max epochs	1000	1000	1000	1000	1000	1000
Early-stop patience (eval rounds)	20	20	20	20	20	20
Reported ALF operating points						
Tail target τ_{tail}	0.10	0.10	0.10	0.10	0.15	0.15
Normalized entropy target $\bar{\tau}_{\text{ent}}$	0.85	0.85	0.91	0.91	0.97	0.97
Adaptive dual updates						
Damping β	0.60	0.60	0.60	0.60	0.60	0.60
Warmup T_{warm} (epochs)	50	50	50	50	50	50
Dual rate η_{tail} (ALF-S)	0.30	0.30	0.30	0.30	0.30	0.30
Dual rate η_{tail} (ALF-P)	1.0	–	1.0	–	1.0	–
Dual rate η_{ent}	0.03	0.03	0.03	0.03	0.03	0.03
Dual clamp $\lambda_{\max,\text{tail}}$	5.0	5.0	5.0	5.0	5.0	5.0
Dual clamp $\lambda_{\max,\text{ent}}$	0.50	0.50	0.50	0.50	0.50	0.50
Proxy scale and shape						
Tail proxy scale α_{tail}	1.0	1.0	1.0	1.0	1.0	1.0
Entropy proxy scale α_{ent}	1.0	1.0	1.0	1.0	1.0	1.0
Proxy temperature t_π	0.10	0.10	0.10	0.10	0.03	0.03
Candidate-pool cap $S_{\max}$	5000	5000	5000	5000	5000	5000

Protocol choices. Two protocol choices matter for interpretation. First, ALF uses a maximum of 1000 epochs and an early-stopping patience of 20. A shorter patience can terminate ALF before its adaptive dual has stabilized at demanding operating points (Sect. S4). Second, Yelp2018 uses a proxy temperature of 0.03, compared with 0.10 on the two MovieLens catalogues. The selection criterion is the same on all three datasets: the unconstrained model's proxy must lie below the targets being requested, so that the rectified violation $(\tau_{\text{tail}} - \widetilde{T})_+$ is non-zero (§3.2). On Yelp2018, $t_\pi = 0.10$ failed this criterion at the headline target itself: the softmax over the much larger catalogue was diffuse enough that the unconstrained proxy already exceeded the target, so the constraint supplied no gradient although the discrete target was unmet. Lowering the temperature to 0.03 restores headroom. On the MovieLens catalogues, $t_\pi = 0.10$ meets the criterion for $\tau_{\text{tail}} \geq 0.10$ but not at the lowest target: at $\tau_{\text{tail}} = 0.05$ on MovieLens-1M the tail term never becomes active, and its training trajectory is identical, seed for seed, to that of the entropy-only ablation, while the dual variable still grows to 0.13 (Table S13). A correlation check alone would not detect this, since an offset large enough to zero the rectifier leaves Pearson r unchanged.

Why several hyperparameters vary by configuration. Most dual and proxy parameters are shared across datasets and backbones. The remaining differences are limited to the learning rate for GraphSAGE on MovieLens-100K and the larger batch size and lower proxy temperature used on Yelp2018. The reported tail targets are the headline operating points in Table 2 together with the GraphSAGE ALF evaluations; they are not configuration-level defaults. In particular, the configuration default for Yelp2018 is $\tau_{\text{tail}} = 0.10$, whereas the reported LightGCN and GraphSAGE ALF results use 0.15, chosen so that the requested lift over the unconstrained backbone is comparable to the MovieLens headline targets (§5.2); the $\tau_{\text{tail}} = 0.10$ result is reported in Table S11. ALF-P uses a larger dual step than ALF-S because its surrogate has a different scale (§3.4); it is the same on every dataset. All other choices are frozen across seeds.

Candidate-pool size. The proxy candidate pool consists of $S_{\max} = 5000$ item indices drawn uniformly with replacement for each active fairness update, on every dataset (§3.2).

Hardware. All LightGCN and GraphSAGE runs were executed on a single NVIDIA GeForce RTX 4090. We do not report a mean wall-clock table because shared-cluster GPU contention makes elapsed-time comparisons noisy; the computational overhead and its asymptotic dependence on the candidate pool are described in §6.5. Full scripts and logs are provided in the released repository.

S2 GraphSAGE Results

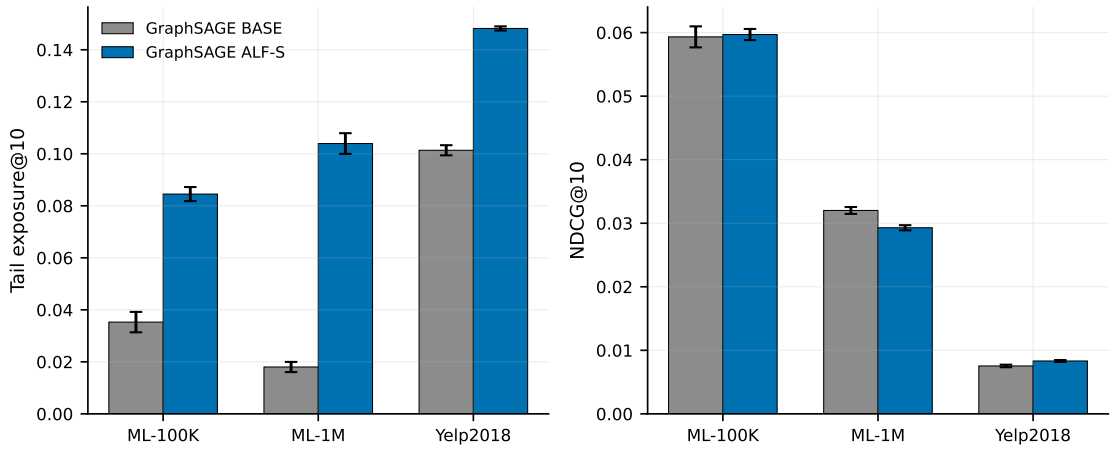


Figure S1: Transfer to a second encoder. TailExp@10 (left) and NDCG@10 (right) for the unconstrained GraphSAGE backbone and GraphSAGE with ALF-S at the headline target, on each dataset. Bars are means over five seeds and error bars are standard errors

This section evaluates BASE versus ALF-S with the GraphSAGE backbone under the same data splits and evaluation protocol as the LightGCN experiments, using five seeds ($\{42, 43, 44, 45, 46\}$) per dataset and condition (Fig. S1). The comparison is restricted to these two methods because the full POP and xQuAD sweeps were conducted with LightGCN. ALF-S uses the reported targets $\tau_{\text{tail}} = 0.10$ for the MovieLens datasets and 0.15 for Yelp2018. The proxy temperatures and entropy targets are transferred from the LightGCN calibration rather than re-tuned for GraphSAGE. A single-seed sanity check on each dataset confirmed that neither proxy was saturated or degenerate

under the GraphSAGE embeddings before the five-seed evaluation was run. Values are means ± standard deviations over five seeds, computed directly from the per-seed logs, and are reported in Tables S3–S5.

Graph construction. The GraphSAGE encoder operates on a heterogeneous graph with three node types (users, items, and item categories) and four relation types: user→item (*rates*), item→user (*rated_by*), item→category (*belongs_to*), and category→item (*contains*). User–item edges are built from training interactions only; validation and test interactions never enter the graph. Item–category edges come from item metadata and do not depend on the split. On MovieLens-100K and MovieLens-1M the categories are the movie genres distributed with each dataset, and a movie is linked to every genre it is tagged with. On Yelp2018 the categories are the entities connected to items by the business-category relation in the knowledge graph released with KGAT (`kg_final.txt`); all other knowledge-graph relations are discarded. Users, items, and categories receive learnable ID embeddings of dimension 64; no textual or pretrained features are used. Each layer applies one `SAGEConv` per relation, sums the relation outputs arriving at each node type, and applies a ReLU. Recommendation scores use the cosine similarity of the resulting user and item embeddings, exactly as for LightGCN, so category nodes influence scores only through message passing. The implementation contains an optional additive category-profile score term, which is disabled in all reported runs.

Tail-restricted recall, exposure entropy, and average rank were not logged for this GraphSAGE sweep and are therefore omitted.

Table S3: GraphSAGE: BASE vs. ALF-S ($\tau_{\text{tail}} = 0.10$) on MovieLens-100K. Mean ± sd over five seeds.

Metric	**BASE**	**ALF-S**	**Change**
Recall@10	0.1164 ± 0.0050	0.1177 ± 0.0036	+1.1%
NDCG@10	0.0593 ± 0.0037	0.0597 ± 0.0020	+0.6%
Coverage@10	0.5784 ± 0.0529	0.5981 ± 0.0163	+3.4%
TailExp@10	0.0353 ± 0.0088	0.0845 ± 0.0060	+139.6%
Exposure Gini ↓	0.7271 ± 0.0369	0.7146 ± 0.0155	−1.7%

Table S4: GraphSAGE: BASE vs. ALF-S ($\tau_{\text{tail}} = 0.10$) on MovieLens-1M. Mean ± sd over five seeds.

Metric	**BASE**	**ALF-S**	**Change**
Recall@10	0.0658 ± 0.0013	0.0604 ± 0.0021	−8.3%
NDCG@10	0.0320 ± 0.0012	0.0293 ± 0.0009	−8.5%
Coverage@10	0.6825 ± 0.0348	0.7070 ± 0.0338	+3.6%
TailExp@10	0.0180 ± 0.0044	0.1039 ± 0.0089	+477.0%
Exposure Gini ↓	0.7359 ± 0.0159	0.7259 ± 0.0236	−1.4%

Table S5: GraphSAGE: BASE vs. ALF-S ($\tau_{\text{tail}} = 0.15$) on Yelp2018. Mean ± sd over five seeds.

Metric	**BASE**	**ALF-S**	**Change**
Recall@10	0.0160 ± 0.0009	0.0177 ± 0.0007	+10.4%
NDCG@10	0.0075 ± 0.0004	0.0083 ± 0.0003	+10.4%
Coverage@10	0.8128 ± 0.0316	0.7696 ± 0.0555	−5.3%
TailExp@10	0.1014 ± 0.0044	0.1482 ± 0.0018	+46.2%
Exposure Gini ↓	0.5896 ± 0.0152	0.6241 ± 0.0326	+5.8%

The constraint transfers to the second backbone. ALF-S increases GraphSAGE tail exposure on all three datasets, by +140%, +477%, and +46% on MovieLens-100K, MovieLens-1M, and Yelp2018, and attains 0.085, 0.104, and 0.148 against targets of 0.10, 0.10, and 0.15. On the MovieLens datasets the trade-off matches LightGCN: NDCG is unchanged on MovieLens-100K (+0.6%) and falls by 8.5% on MovieLens-1M, while coverage rises by 3.4% and 3.6% and exposure Gini falls by 1.7% and 1.4%. Yelp2018 differs because of the unconstrained baseline's checkpoint: under the common validation rule, GraphSAGE BASE is reported at early epochs (160–480 of 1000) whose validation tail exposure exceeds that of the final state (0.101 against 0.087) at lower validation NDCG. The reported BASE therefore already has elevated exposure and lower accuracy, so ALF-S shows higher NDCG (+10.4%), lower coverage (−5.3%), and higher Gini (+5.8%) than BASE there. We read the Yelp2018 comparison as evidence that the constraint transfers, not as an accuracy gain from it.

The proportional tail-exposure increase is largest on MovieLens-1M and smallest on Yelp2018, where the BASE checkpoint already has elevated exposure. We do not attribute this ordering to a particular mechanism; as for the LightGCN results (§7.1), three datasets are too few to establish a relationship with sparsity, and this GraphSAGE experiment was not designed to test one. The result therefore supports transfer of the constraint mechanism to a second encoder without establishing a general relationship between catalogue sparsity and ALF's relative effect.

S3 xQuAD Results

Table S6 reports the xQuAD-Smooth grid used as the primary post-hoc baseline in the main text (Fig. 2, Table 4).

Table S6: xQuAD-Smooth on LightGCN: NDCG@10 / TailExp@10 / Coverage@10 by reranking weight λ_{xQ}. Means over five seeds; the $\lambda_{\text{xQ}} = 0$ row is the BASE ranking.

λ_{xQ}	**ML-100K**	**ML-1M**	**Yelp2018**
0.0	0.0647 / 0.0398 / 0.685	0.0333 / 0.0273 / 0.815	0.01136 / 0.0632 / 0.838
0.1	0.0626 / 0.0912 / 0.703	0.0324 / 0.0594 / 0.838	0.01091 / 0.1242 / 0.849
0.3	0.0490 / 0.3115 / 0.697	0.0268 / 0.2353 / 0.839	0.00860 / 0.3441 / 0.826
0.5	0.0360 / 0.5684 / 0.647	0.0178 / 0.5048 / 0.804	0.00630 / 0.5763 / 0.771
0.7	0.0276 / 0.7561 / 0.567	0.0118 / 0.7183 / 0.740	0.00474 / 0.7528 / 0.690
0.9	0.0216 / 0.9000 / 0.439	0.0071 / 0.8940 / 0.611	0.00345 / 0.8999 / 0.552

Results for the Binary variant are reported here for completeness, including the $\lambda = 0$ sanity check against the independently measured BASE NDCG (§4.3), in Table S7.

Table S7: xQuAD-Binary on LightGCN. The Binary variant reaches identical reported values for all $\lambda \geq 0.5$ on every dataset. The $\lambda = 0$ row applies no reordering and reproduces the independently measured BASE NDCG@10 (Table 2) up to floating-point and evaluation-batching differences.

Dataset	Configuration	NDCG@10	TailExp@10
ML-100K	xQuAD-Binary ($\lambda = 0.0$)	0.0647	0.0398
	xQuAD-Binary ($\lambda = 0.1$)	0.0629	0.0644
	xQuAD-Binary ($\lambda = 0.3$)	0.0545	0.1102
	xQuAD-Binary ($\lambda = 0.5$)	0.0514	0.1127
	xQuAD-Binary ($\lambda = 0.7$)	0.0514	0.1127
	xQuAD-Binary ($\lambda = 0.9$)	0.0514	0.1127
ML-1M	xQuAD-Binary ($\lambda = 0.0$)	0.0333	0.0273
	xQuAD-Binary ($\lambda = 0.1$)	0.0327	0.0446
	xQuAD-Binary ($\lambda = 0.3$)	0.0295	0.0991
	xQuAD-Binary ($\lambda = 0.5$)	0.0266	0.1082
	xQuAD-Binary ($\lambda = 0.7$)	0.0266	0.1082
	xQuAD-Binary ($\lambda = 0.9$)	0.0266	0.1082
Yelp2018	xQuAD-Binary ($\lambda = 0.0$)	0.0114	0.0632
	xQuAD-Binary ($\lambda = 0.1$)	0.0111	0.0887
	xQuAD-Binary ($\lambda = 0.3$)	0.0099	0.1196
	xQuAD-Binary ($\lambda = 0.5$)	0.0094	0.1203
	xQuAD-Binary ($\lambda = 0.7$)	0.0094	0.1203
	xQuAD-Binary ($\lambda = 0.9$)	0.0094	0.1203

The $\lambda = 0$ row matches the corresponding BASE NDCG@10 values in Table 2 on all three datasets (0.0647, 0.0333, and 0.0114), consistent with xQuAD applying no reordering at $\lambda = 0$. For the Binary variant, the reported NDCG and TailExp values become identical for all $\lambda \geq 0.5$ on each dataset. Thus, increasing the reranking weight beyond this range does not produce an additional change in the reported operating point. Compared with xQuAD-Smooth, the Binary variant consequently provides fewer distinct exposure–accuracy operating points over the tested grid, which is why Smooth is used as the primary post-hoc baseline in the main text.

S4 Reproducibility

Code release. All code, configuration files, and analysis scripts are released at `https://github.com/puskr99/Adaptive-Lagrangian-Fairness`. The repository contains the training harness, xQuAD reranking implementation, Pareto-analysis pipeline, and figure-generation scripts used to produce the reported results.

Reproducibility checklist. For each (dataset, backbone, method, seed) configuration, the released artifacts include:

- the training script, whose configuration class and experiment definitions fix all hyperparameters (Table S2),
- per-run JSON logs containing the reported metrics,
- aggregated CSV files used for Pareto and hypervolume analyses, and
- the random seed associated with each run.

Each dataset has one training script, `lgcn.py`, that implements both backbones; `gs.py` runs the GraphSAGE experiment. On MovieLens-1M and Yelp2018, each ALF-S run is evaluated on the test set at both the final training state and the validation-rule checkpoint, and both results are released (Table S14). The analysis scripts construct Pareto fronts, compute hypervolume and bootstrap confidence intervals, and generate the reported tables and figures from the released logs.

Data availability. MovieLens-100K and MovieLens-1M are publicly available from the GroupLens project (Harper and Konstan, 2015; `https://grouplens.org/datasets/movielens/`). The Yelp2018 data and test partition are the version released with KGAT (Wang et al., 2019), available from the original authors' repository (`https://github.com/xiangwang1223/knowledge_graph_attention_network`). Full references are given in the main paper.

Computational budget. All reported runs were trained on a single NVIDIA GeForce RTX 4090. The fairness candidate pool is 5000 item indices sampled with replacement per active fairness update on every dataset.

The primary LightGCN experiments comprise BASE (one configuration), POP (10 fixed-λ configurations), PBiLoss (6 fixed-w configurations), ALF-S ($7+9+10=26$ tail-target configurations), and ALF-P ($7+9+5=21$ tail-target configurations). With five seeds and three datasets, this gives $15+150+90+130+105=490$ LightGCN training runs. The GraphSAGE experiment adds BASE versus ALF-S on three datasets with five seeds, for 30 runs. The scoring diagnostic (Table 1) retrains BASE under inner-product and cosine scoring, for another 30 runs. SimGCL adds 15 runs, and the three-condition constraint ablation adds 45 runs. xQuAD is evaluated post-hoc on the existing BASE LightGCN checkpoints and requires no additional training: two variants, six λ values, five seeds, and three datasets yield 180 reranking evaluations. The single-seed PopNeg check on MovieLens-100K adds 5 runs.

Grid coverage. The POP grid uses the same ten λ values, $\{0.01, 0.05, 0.10, 0.15, 0.20, 0.25, 0.30, 0.40, 0.50, 0.75\}$, on every dataset. The grid was fixed before the dataset-specific results were examined and was chosen to span the exposure range reached by the ALF target sweep rather than being optimized separately for each dataset. This criterion depends only on the exposure axis and does not reference accuracy (§4.3). The PBiLoss grid, $\{0.05, 0.1, 0.2, 0.5, 1.0, 2.0\}$, is likewise shared across datasets.

The resulting attained exposure ranges are ML-100K: POP $[0.042, 0.666]$ versus ALF-S $[0.039, 0.675]$; ML-1M: POP $[0.027, 0.966]$ versus ALF-S $[0.027, 0.876]$; and Yelp2018: POP $[0.071, 0.994]$ versus ALF-S $[0.067, 0.948]$. PBiLoss and ALF-P attain $[0.054, 0.519]$ and $[0.050, 0.549]$ on ML-100K, $[0.044, 0.506]$ and $[0.049, 0.589]$ on ML-1M, and $[0.071, 0.201]$ and $[0.095, 0.214]$ on Yelp2018. The ranges within each pair substantially overlap, motivating the common-range hypervolume comparison in §5.3 rather than a simple range-containment comparison.

Training-budget sensitivity. The final protocol uses a maximum of 1000 epochs and an early-stopping patience of 20 evaluation rounds for ALF on every dataset and backbone (Table S2). Because ALF's dual variable can undergo an overshoot-and-recovery phase, we additionally examined whether a shorter patience could terminate a run before this correction had settled. Table S8 shows one fully logged trajectory for MovieLens-1M at $\tau_{\text{tail}} = 0.65$, seed 42. For this diagnostic experiment, the longer-patience runs were capped at 500 epochs to isolate the effect of allowing the trajectory to continue.

Table S8: Training-budget sensitivity for one illustrative trajectory (MovieLens-1M, $\tau_{\text{tail}} = 0.65$, seed 42). The final protocol uses patience 20 and a 1000-epoch maximum; the 500-epoch cap here is a diagnostic setting. Values are validation metrics at the stopping point.

Patience	**Stop epoch**	**val NDCG**	**val Tail**	**Tail/τ**
10	320	0.328	0.738	1.14×
20	500	0.345	0.658	1.01×

With patience 10, training stops at epoch 320 while validation tail exposure remains approximately 14% above the target. With patience 20, the trajectory reaches the 500-epoch diagnostic cap and the corresponding tail exposure is approximately 1% above target. A patience-30 diagnostic run on the same trajectory also reaches the 500-epoch cap rather than stopping early, indicating that the difference is caused by allowing the run to reach the cap rather than by the distinction between patience 20 and 30.

This experiment illustrates why the final protocol uses a uniform patience of 20 rather than the shorter value. It does not establish that every ALF configuration has fully converged under the final 1000-epoch budget; that question is treated as a limitation in §7.2. BASE and POP were also checked under the same extended budget and did not show material changes in their reported outcomes.

S5 Full Target Sweeps and Weight Grids

Tables S9–S11 list the ALF-S target sweeps and POP grids plotted in Fig. 2; Table S12 lists the ALF-P sweeps with their terminal dual variables and the PBiLoss grids; Table S13 lists the ALF-S terminal dual variables summarized in Sect. 6.4; and Table S14 compares the two ALF-S checkpoints on MovieLens-1M and Yelp2018 (Sect. 4.5).

Table S9: MovieLens-100K: full ALF-S target sweep and POP λ grid. All ALF-S points are non-dominated within their own method; the POP $\lambda = 0.01, 0.05, 0.20$ points are dominated within POP. Means over five seeds.

ALF-S		**POP**	
τ_{tail}	NDCG@10 / TailExp@10	λ	NDCG@10 / TailExp@10
0.05	0.0640 / 0.0390	0.01	0.0627 / 0.0420
0.10	0.0587 / 0.0817	0.05	0.0616 / 0.0468
0.15	0.0563 / 0.1477	0.10	0.0628 / 0.0539
0.20	0.0538 / 0.2119	0.15	0.0592 / 0.0671
0.35	0.0430 / 0.3760	0.20	0.0588 / 0.0809
0.50	0.0309 / 0.5357	0.25	0.0590 / 0.1033
0.65	0.0228 / 0.6748	0.30	0.0572 / 0.1335
		0.40	0.0527 / 0.2266
		0.50	0.0431 / 0.3555
		0.75	0.0247 / 0.6657

Table S10: MovieLens-1M: full ALF-S target sweep and POP λ grid. All reported operating points are non-dominated within their own method. Means over five seeds.

	ALF-S		**POP**
τ_{tail}	NDCG@10 / TailExp@10	λ	NDCG@10 / TailExp@10
0.05	0.0339 / 0.0267	0.01	0.0334 / 0.0272
0.10	0.0308 / 0.0982	0.05	0.0331 / 0.0352
0.15	0.0285 / 0.1479	0.10	0.0319 / 0.0602
0.20	0.0265 / 0.1993	0.15	0.0224 / 0.2974
0.35	0.0210 / 0.3494	0.20	0.0143 / 0.5820
0.50	0.0166 / 0.4999	0.25	0.0100 / 0.7307
0.65	0.0122 / 0.6490	0.30	0.0078 / 0.8114
0.75	0.0087 / 0.7856	0.40	0.0049 / 0.8925
0.85	0.0058 / 0.8758	0.50	0.0037 / 0.9301
		0.75	0.0024 / 0.9656

Table S11: Yelp2018: full ALF-S target sweep and POP λ grid. All reported operating points are non-dominated within their own method. Means over five seeds.

	ALF-S		**POP**
τ_{tail}	NDCG@10 / TailExp@10	λ	NDCG@10 / TailExp@10
0.05	0.0112 / 0.0674	0.01	0.0113 / 0.0708
0.10	0.0109 / 0.1009	0.05	0.0090 / 0.2214
0.15	0.0100 / 0.1509	0.10	0.0049 / 0.5760
0.20	0.0093 / 0.1974	0.15	0.0034 / 0.7310
0.35	0.0073 / 0.3491	0.20	0.0027 / 0.8145
0.50	0.0056 / 0.5003	0.25	0.0023 / 0.8660
0.65	0.0042 / 0.6461	0.30	0.0021 / 0.9017
0.75	0.0033 / 0.7480	0.40	0.0017 / 0.9465
0.85	0.0025 / 0.8488	0.50	0.0015 / 0.9710
1.00	0.0017 / 0.9480	0.75	0.0011 / 0.9937

Table S12: ALF-P target sweeps with the selected terminal dual variable λ_{tail}, and PBiLoss weight grids, on all three datasets. Means over five seeds.

	ALF-P			**PBiLoss**	
	τ_{tail}	NDCG@10 / TailExp@10	λ_{tail}	w	NDCG@10 / TailExp@10
ML-100K	0.05	0.0634 / 0.0504	0.048	0.05	0.0622 / 0.0541
	0.10	0.0624 / 0.1013	0.204	0.1	0.0619 / 0.0683
	0.15	0.0615 / 0.1575	0.336	0.2	0.0597 / 0.1026
	0.20	0.0596 / 0.2074	0.466	0.5	0.0561 / 0.2173
	0.35	0.0536 / 0.3564	0.935	1.0	0.0495 / 0.3708
	0.50	0.0459 / 0.4982	1.734	2.0	0.0410 / 0.5186
	0.65	0.0374 / 0.5491	2.596		
ML-1M	0.05	0.0319 / 0.0491	0.069	0.05	0.0326 / 0.0436
	0.10	0.0309 / 0.1000	0.180	0.1	0.0328 / 0.0629
	0.15	0.0297 / 0.1533	0.282	0.2	0.0309 / 0.1117
	0.20	0.0283 / 0.2036	0.386	0.5	0.0282 / 0.2473
	0.35	0.0249 / 0.3499	0.795	1.0	0.0240 / 0.3948
	0.50	0.0211 / 0.4607	1.342	2.0	0.0201 / 0.5058
	0.65	0.0189 / 0.5338	2.073		
	0.75	0.0175 / 0.5680	2.653		
	0.85	0.0167 / 0.5891	3.302		
Yelp2018	0.10	0.01095 / 0.0947	0.206	0.05	0.01125 / 0.0709
	0.15	0.01051 / 0.1299	0.472	0.1	0.01115 / 0.0787
	0.20	0.01007 / 0.1562	0.756	0.2	0.01102 / 0.0923
	0.35	0.00906 / 0.1964	1.719	0.5	0.01041 / 0.1310
	0.50	0.00826 / 0.2143	2.789	1.0	0.00970 / 0.1685
				2.0	0.00871 / 0.2011

Table S13: Selected terminal tail dual variable λ_{tail} (mean ± sd over five seeds) by target and dataset (§4.5).

Dataset	**Target τ_{tail}**	**λ_{tail} (mean ± sd)**
ML-100K	0.05	0.049 ± 0.009
	0.10	0.226 ± 0.005
	0.15	0.338 ± 0.005
	0.20	0.398 ± 0.006
	0.35	0.524 ± 0.017
	0.50	0.652 ± 0.011
	0.65	0.788 ± 0.009
ML-1M	0.05	0.131 ± 0.003
	0.10	0.139 ± 0.008
	0.15	0.124 ± 0.006
	0.20	0.132 ± 0.006
	0.35	0.154 ± 0.004
	0.50	0.173 ± 0.001
	0.65	0.207 ± 0.002
	0.75	0.256 ± 0.004
	0.85	0.374 ± 0.004
Yelp2018	0.05	0.000 ± 0.000
	0.10	0.034 ± 0.000
	0.15	0.042 ± 0.001
	0.20	0.050 ± 0.001
	0.35	0.068 ± 0.001
	0.50	0.091 ± 0.001
	0.65	0.127 ± 0.003
	0.75	0.164 ± 0.001
	0.85	0.241 ± 0.001
	1.00	0.442 ± 0.001

Table S14: ALF-S on MovieLens-1M and Yelp2018 evaluated at the final training state (reported in the main paper) and at the validation-rule checkpoint used for all other methods. Each run is trained once and tested at both checkpoints. Hypervolume differences ($\times 10^3$) use the same windows, reference points, and 1000 bootstrap resamples as Table 3 (relative differences for ALF-S vs. POP: +10.5% and +9.6% on MovieLens-1M, +13.2% and +13.1% on Yelp2018); target error is $|\tau_{\text{tail}} -$ mean TailExp@10| over the full sweep (median) and over $0.15 \le \tau_{\text{tail}} \le 0.65$ (range).

		Hypervolume difference		**Target error**	
Dataset	**Checkpoint**	ALF-S vs. POP	ALF-P vs. ALF-S	median	$\tau \in [0.15, 0.65]$
ML-1M	final	+1.27 [+0.92, +1.58]	+1.86 [+1.61, +2.07]	0.002	0.000–0.002
	rule	+1.16 [+0.86, +1.46]	+1.86 [+1.67, +2.06]	0.050	0.048–0.073
Yelp2018	final	+0.485 [+0.446, +0.531]	+0.025 [+0.013, +0.035]	0.002	0.000–0.004
	rule	+0.480 [+0.431, +0.531]	+0.026 [+0.012, +0.036]	0.038	0.020–0.094

At the rule's checkpoint, ALF-S overshoots its targets by 0.05–0.07 on MovieLens-1M and 0.02–0.09 on Yelp2018 for mid-range targets, as expected from the validation histories (Sect. 4.5), but the frontier comparisons are unchanged: ALF-S still exceeds POP with intervals excluding zero, and

ALF-P still exceeds ALF-S by the same margin. The final-state exception therefore affects target tracking, not the frontier results.

S6 Training and Dual Dynamics

Figure S2 shows validation trajectories and the selected checkpoints (Sect. 4.5), Fig. S3 the pooled agreement between the soft proxy and the discrete metric (Sect. 6.3), and Fig. S4 the dual-variable trajectories (Sect. 6.4).

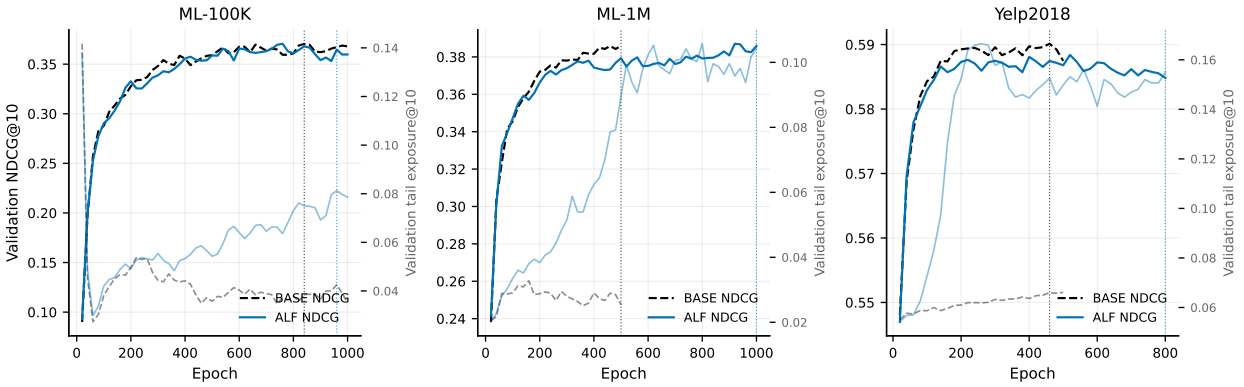


Figure S2: Validation trajectories for BASE (dashed) and ALF-S at the headline target (solid) on each dataset, seed 42. Dark curves show validation NDCG@10 on the left axis; faded curves show validation TailExp@10 on the right axis. Vertical dotted lines indicate the reported checkpoints

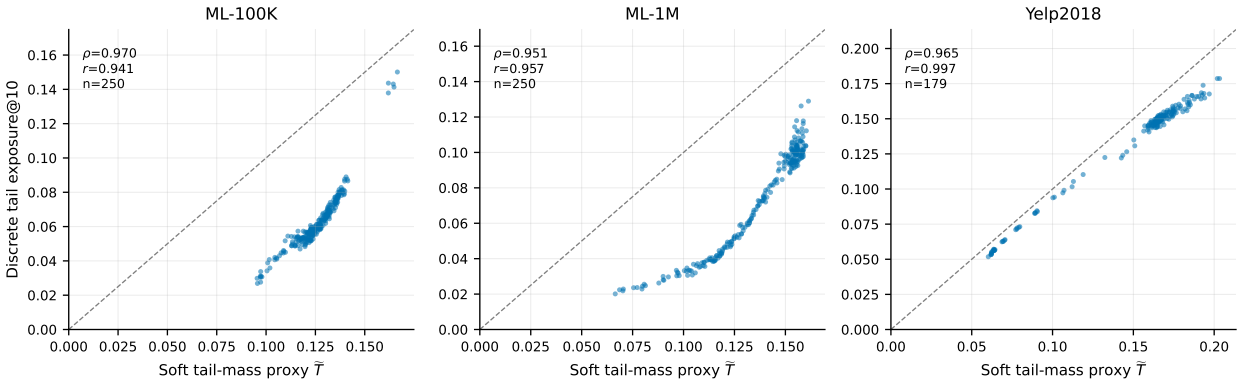


Figure S3: Agreement between the differentiable soft tail-mass proxy and the discrete TailExp@10 metric, pooled over five seeds and all logged epochs of the headline run on each dataset. The dashed line denotes the identity. Spearman ρ, Pearson r, and the number of observations are shown in each panel. The logged proxy is the masked, full-catalogue validation quantity. The proxy is higher in level than the discrete metric while retaining substantial rank agreement. Table 8 reports the within-seed statistics

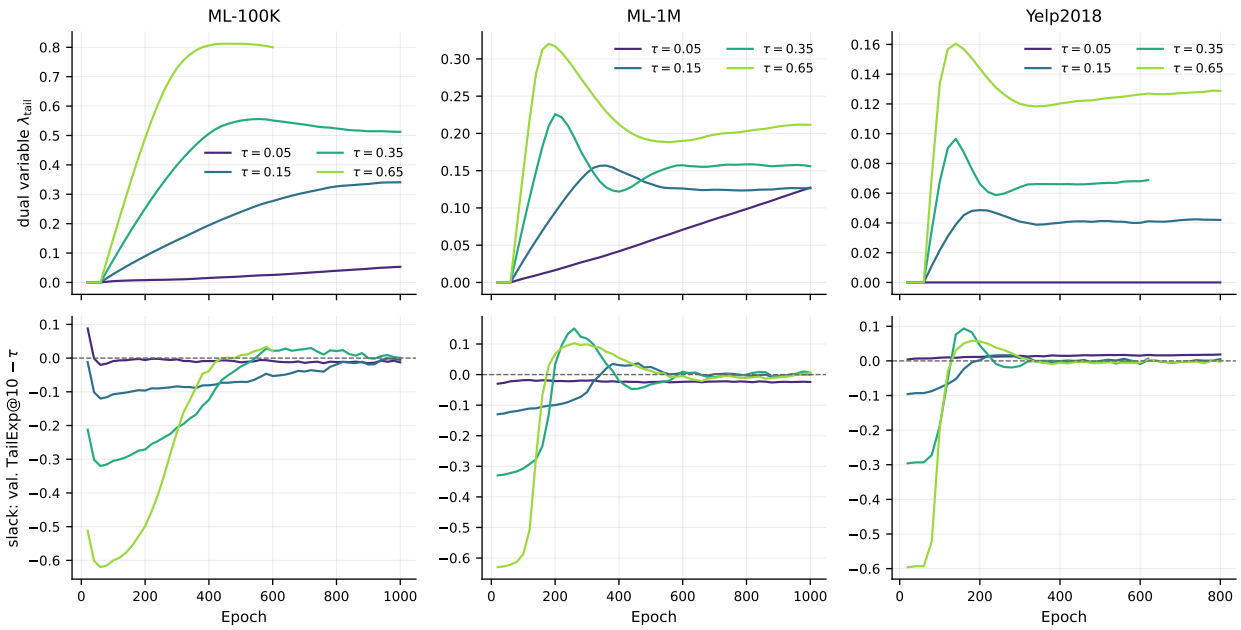


Figure S4: ALF-S dual ascent dynamics for seed 42 at four targets per dataset. Top row: the tail dual variable λ_{tail} against training epoch. Bottom row: the constraint slack, attained validation tail exposure minus the requested target, with the zero line dashed. Darker curves correspond to lower targets. On MovieLens-1M at $\tau_{\text{tail}} = 0.05$ the multiplier grows linearly without effect: the discrete metric stays below the target while the proxy stays above it, so the primal term is inactive (§6.4)