

# Convex Dataset Valuation for Post-Training

ICML 2026

Siqi Zeng, Christopher Jung, Rui Li, Zhe Kang, Ming Li, Nima Noorshams,  
Zhigang Wang, Fuchun Peng, Han Zhao, Xue Feng

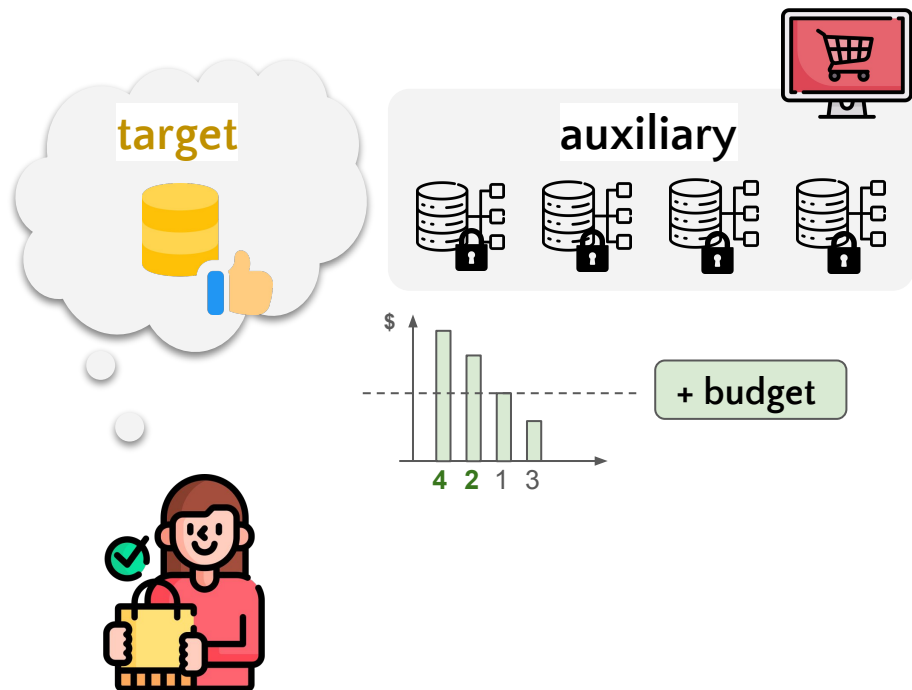


UNIVERSITY OF  
**ILLINOIS**  
URBANA-CHAMPAIGN




**ICML**  
International Conference  
On Machine Learning

# Dataset Marketplace & Dataset Valuation Scores



## Key assumptions

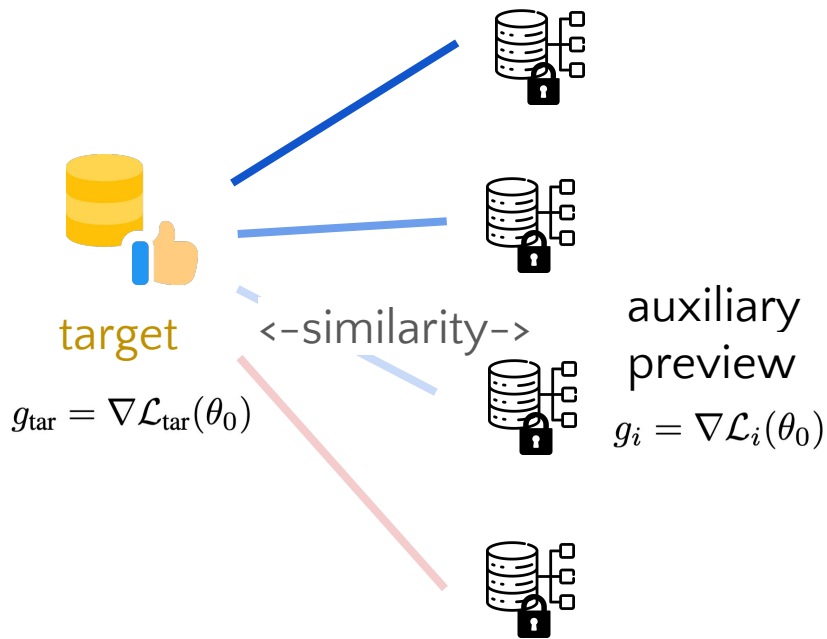
- Post-training (SFT/RL)
- Subset selection
  - One-time, upfront dataset license
- Limited data access 

## What makes good valuation scores?

- Task-dependent
- Signed
- Comparative

# Gradient-based baseline

Score:  $w_i = \langle g_{\text{tar}}, g_i \rangle$



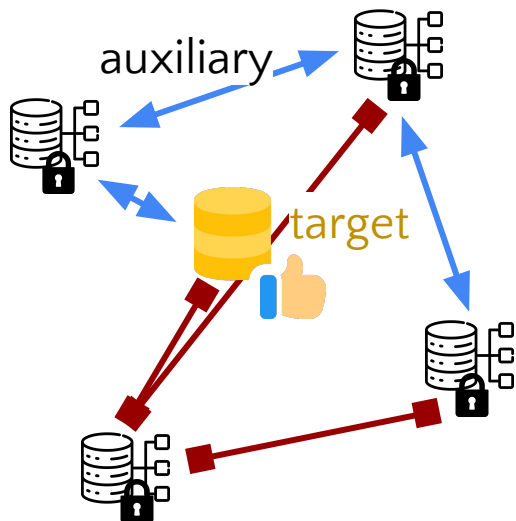
- Additive utility model
- **Assume auxiliary datasets are independent**
- Not true in general
  - Math vs. Coding



**How to include relationships between auxiliary datasets?**

# Dataset Valuation with Kernel Mean Matching

Non-additive utility model  $\Rightarrow \min_{w \in \mathcal{W}_k} \|Gw - \lambda g_{\text{tar}}\|_2^2$



Convex Quadratic Program

$$\min_{\text{budget}} \frac{1}{2} w^\top K w - \lambda \beta^\top w$$

$$K_{ij} = \langle g_i, g_j \rangle \quad \text{Penalize redundancy}$$

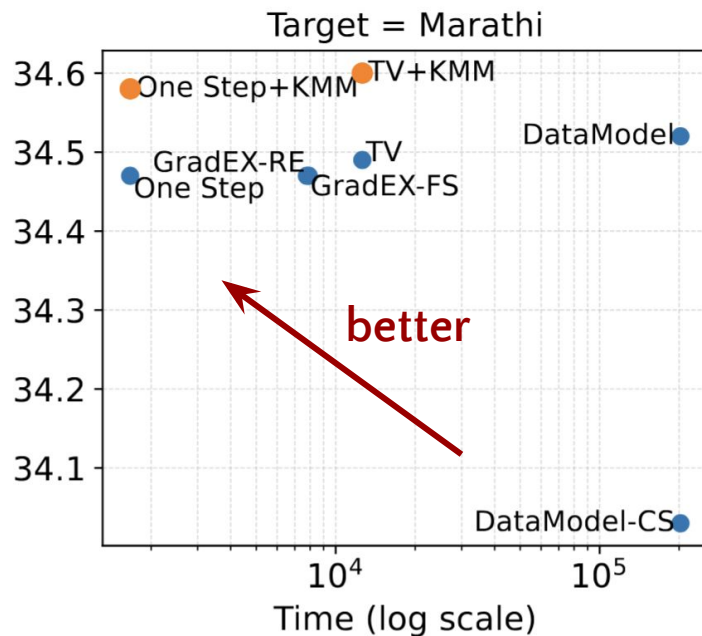
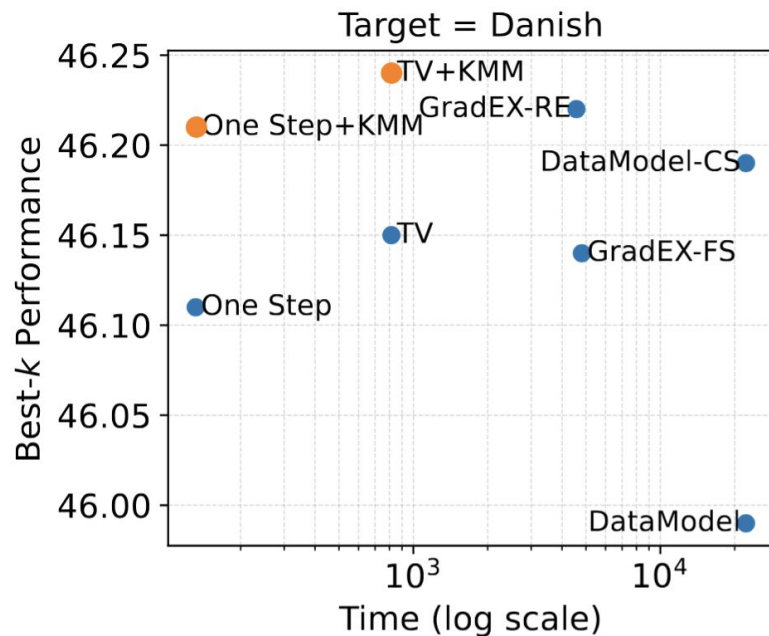
$$\beta_i = \langle g_i, g_{\text{tar}} \rangle \quad \text{Encourage alignment}$$



**Theorem:** KMM valuation scores are still reliable with only finite auxiliary samples.

# KMM achieves the best accuracy-efficiency tradeoff

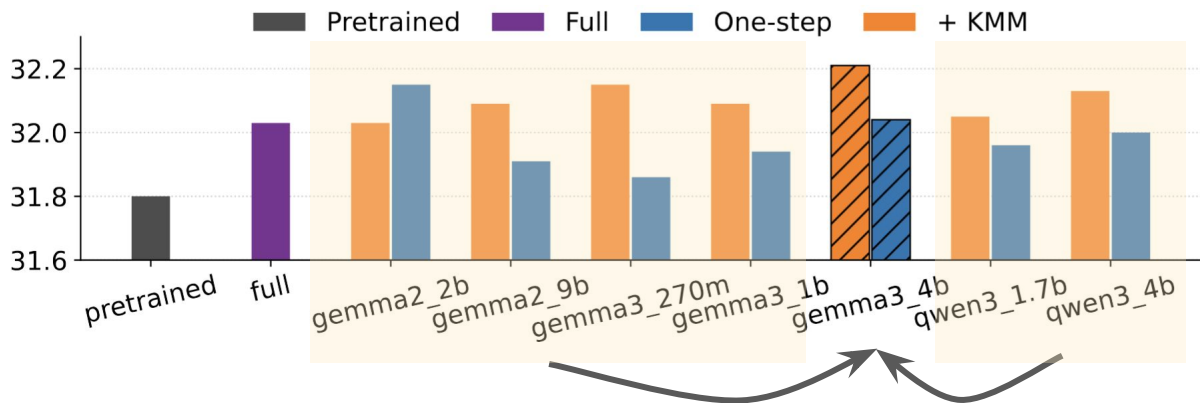
Setup: 25 auxiliary multilingual datasets



# KMM scores transfer across scales and architecture

- KMM scores on *smaller model*  $\Rightarrow$  *larger model*
- KMM scores on *older model*  $\Rightarrow$  *newer model*
- KMM scores on *architecture family A*  $\Rightarrow$  *family B*

## Multilingual



## Instruction Following

Model	TV	TV+KMM
Llama-3.2-1B	26.26	<b>27.70</b>
Llama-3.1-8B	52.64	<b>53.00</b>
Gemma-3-12B	47.00	<b>49.04</b>
Llama-3.1-70B	70.10	<b>71.20</b>

**1B scales to 70B!**

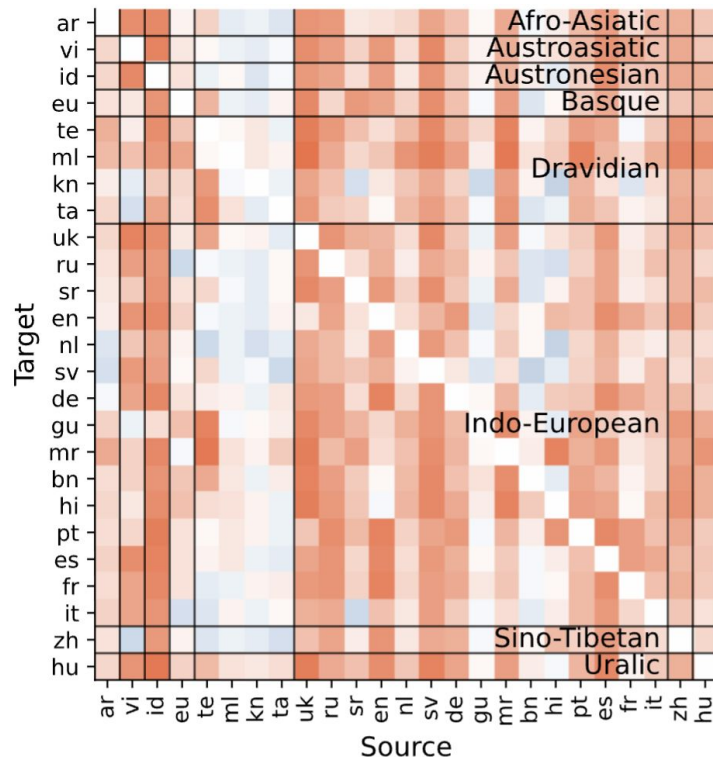
# KMM align with real utility and are interpretable

KMM-based utility better matches true utility

Method	Spearman $\rho$	$p$ -value	Top-10 Overlap
One-Step	-0.1746	0.198	0/10
One-Step+KMM	-0.1265	0.353	2/10
TV	+0.1810	0.182	2/10
TV+KMM	<b>+0.3960</b>	<b>0.003</b>	<b>4/10</b>

	$\theta_0$	$\theta_{tar}$	zh-pos	zh-pos-neg
Train CE	-	2.8542	2.7132	1.7787
Eval CE	3.5199	3.3881	3.3819 (-0.0062)	3.3871 (+0.0052)
M_MMLU	50.44	50.51	50.64 (+0.13)	50.49 (-0.15)

Positive/Negative scores  
 $\Leftrightarrow$  Task Performance Gain/Loss



# Takeaway

- Use **dataset valuation scores for post-training subset selection**
  - Select costly auxiliary datasets under a budget.
- Propose **KMM-based valuation, scores datasets by target alignment while accounting for redundancy.**
- KMM produces **scalable, faithful, interpretable rankings** that outperform baselines.
- Finite-sample reliability **guarantee under limited data access.**

Contact: [siqi6@illinois.edu](mailto:siqi6@illinois.edu)



Paper



Code