CRCV REU – Week #8

Project: **Geo-HiPO**

Advisors:
Rohit Gupta, Parth Kulkarni, Dr. Shah

David Rivard

07/18/2025

Project Overview



Motivation:

- Relative spatial reasoning as new fine-tuning strategy for VLMs.
  - By introducing preference pairs.
- VLMs won't have to be retrained on massive geo-tagged dataset.
  - Should generalize to new scenes better.

Research Question:

- Can we improve geolocation accuracy by **teaching vision-language models to reason location like humans**?

This Week's Focus

- DPO training and evaluation
- Qwen vs Kimi model comparison
- Dataset debugging
- Building a recreatable pipeline

Model Setup & Experiments

- Kimi-VL working for inference + evaluation
- Baseline SFT for Qwen2.5-VL-3B and 7B
- Prompt structure/output differences (Qwen vs Kimi)
- Repo rebuild after weight update issues

DPO Training & Debugging

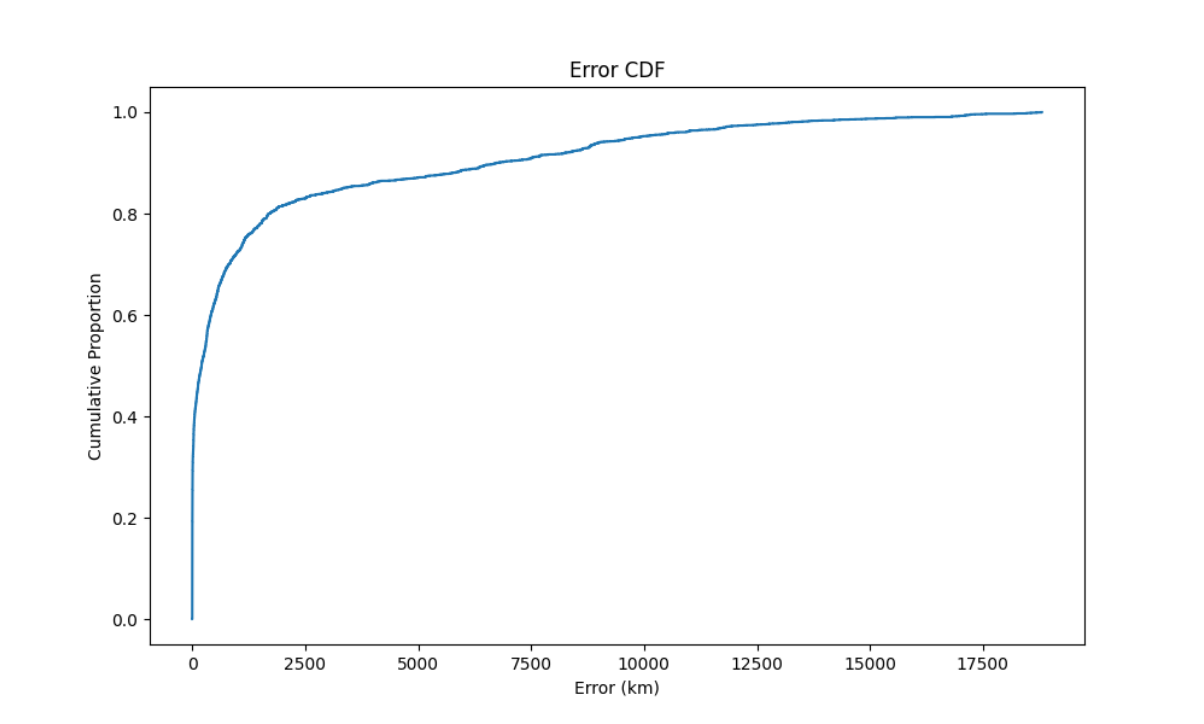
- LoRA and full finetuning for Qwen 3B/7B
- Fixed broken weight updates and config bugs
- Applied DPO loss properly to multimodal model
- Verified fine-tuning effects on output

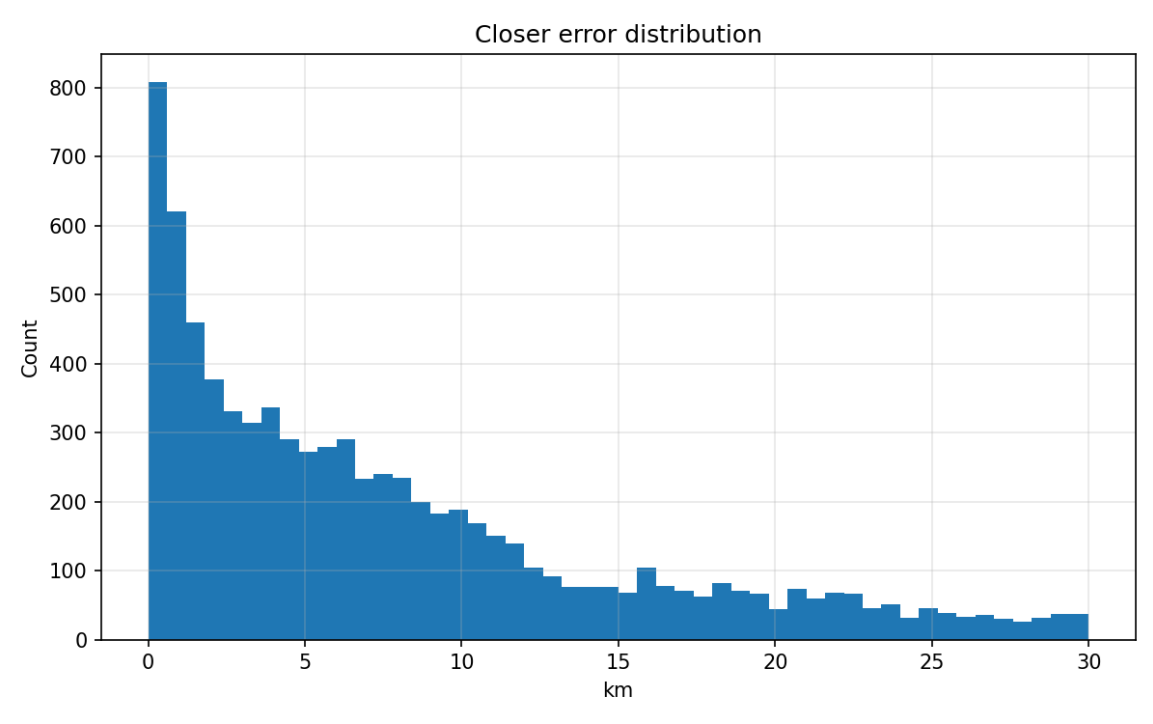
Dataset Preparation

- Built Geo-HiPO preference pairs
- Distance filtering and geocode cleanup
- Fixed 'chosen' hallucinations in pairs
- Created raw/filtered/final versions of dataset

Inference & Evaluation

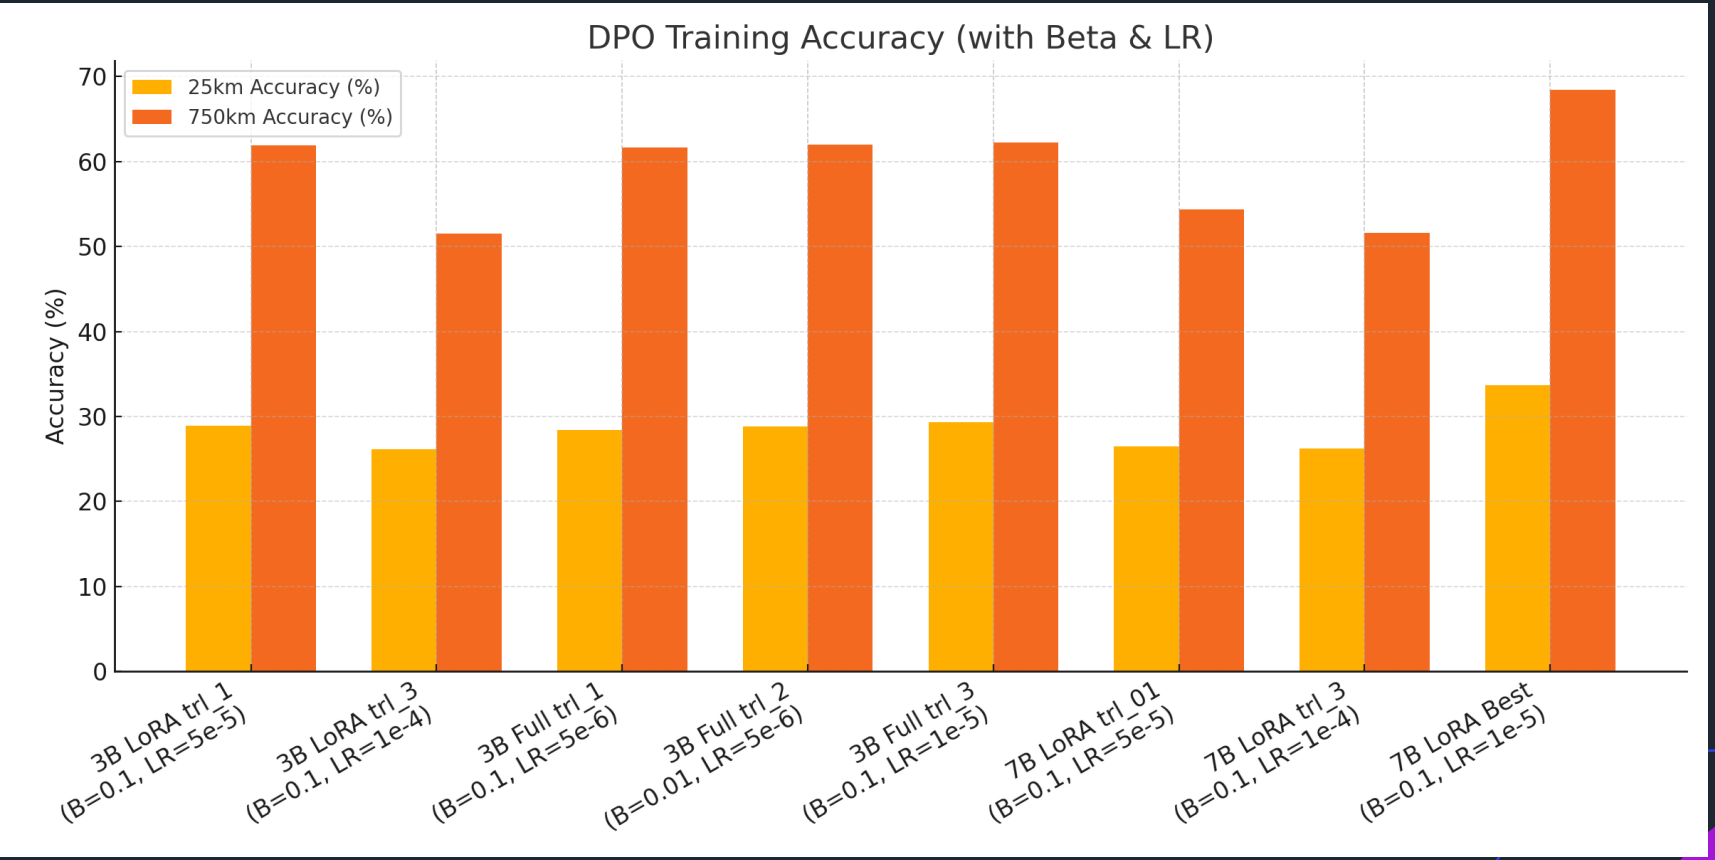
- Multi-checkpoint inference completed
- 25km and 750km accuracy measured
- Analyzed hallucinations and prediction quality





DPO Accuracy Summary

- Compared Qwen 3B/7B across LoRA + full finetuning
- Performance strongly influenced by Beta and LR



Best: 7B LoRA

25km: **33.70**%,

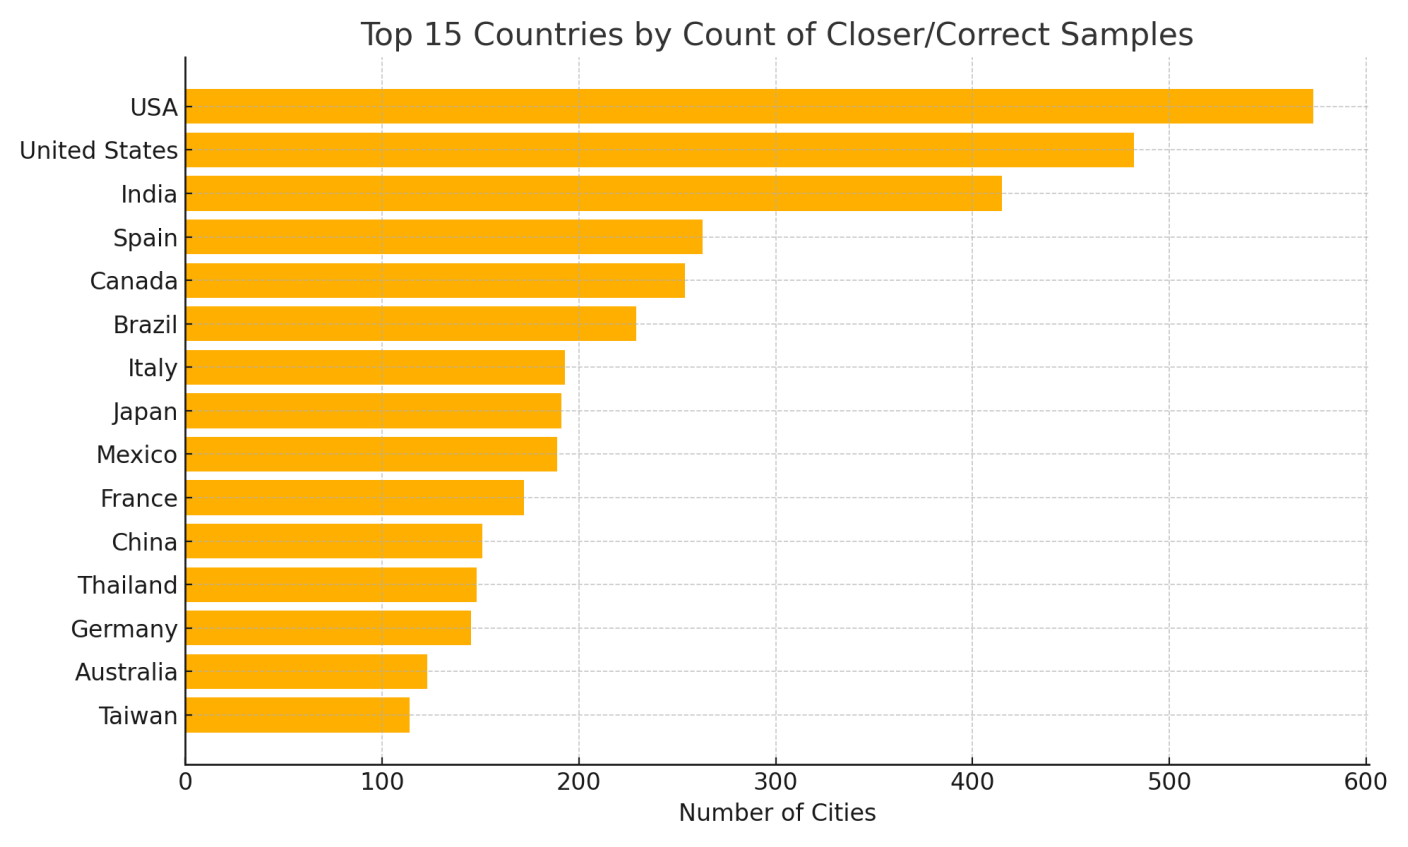
750km: **68.47**%

Beta: **0.1**

Learning Rate: **1e-5**

Top Predicted Countries

- Summarized predictions by country
- USA, India, Spain, Canada are top contributors



Plan – Week 9

- Drop Insight-V if not working
- Get hierarchical loss implemented
- Expand benchmarks and evals
- Finalize poster and finish paper draft

Thank You!

-

Any Questions?



David Rivard

07/18/2025