Poster Presentation Project

1 Goal and spirit

You will select a conference paper and prepare a **5-minute poster presentation** supported by:

(i) a short report (**max 5 pages**),

(ii) **unlimited source code**.

A poster (several printed A4 pages) *or* a single-slide "poster-style" board is acceptable, as long as it is readable and well-structured. Your poster must be understandable by a classmate who has *not* read the paper. The objective is **not** to retrain huge models or replicate the entire paper at scale, but to:

- extract and explain in depth the **key ideas** from the paper;
- **verify** those ideas via a small but convincing experiment and/or a clean theoretical sanity check;
- present the material with **clarity** and **creativity**, ideally relating it to concepts from the class;
- demonstrate sound scientific practice (ablations, controls, failure modes, reproducibility).

Submission. Send your **report**, **poster**, and **code** by email to `edouard.oyallon@cnrs.fr` with the tag *[MVA] Poster presentation* in the subject line.

Two examples of suitable projects (examples only). These papers illustrate the expected scope and are simply examples:

- Verify a neat ML-systems idea on *tiny* networks (e.g., a perceptron-like model) and a minimal distributed setup (e.g., `gloo`): *Ladder-Residual: Parallelism-Aware Architecture for Accelerating Large Model Inference with Communication Overlapping*.
- Study activation-memory savings and recomputation strategies by connecting to activation checkpointing as in: *Reformer: The Efficient Transformer*. Complement this with small numerical experiments to understand the method's memory footprint and for instance, reconstruction stability.

2 Paper selection rules

Eligible papers. Your starting point must be a paper corresponding to a conference submission (typically **8 pages** of main content) from venues such as MLSys, ICML, NeurIPS (or similarly scoped ML, ML-systems venues). You must pick a paper that has at least one of the following:

- a quantitative claim (speed, memory, accuracy, scaling trend, sample complexity, etc.);
- an algorithmic contribution (new operator, schedule, estimator, pipeline, systems mechanism);

- a theoretical statement with implications for the system part (bound, equivalence, invariance, stability, convergence claim, etc.).

Notes.

- A long appendix is fine; your project must still focus on a *small subset* of claims.
- Surveys/tutorials are generally *not* a good fit unless you pick a precise claim to verify.

If the original paper uses large-scale training. You must **scale down intelligently**. Examples:

- verify a systems mechanism on tiny models (the mechanism should still appear);
- test a memory or runtime claim with synthetic workloads;
- reproduce a trend in a reduced setting (smaller model, smaller dataset, fewer steps);
- illustrate a toy version of a theorem, or validate assumptions empirically.

3 Poster and report.

Expected content. Your report should answer:

- **What is the paper's claim?** State 1–2 claims precisely.
- **What exactly did you test?** Describe your verification protocol, baselines, and controls.
- **What results did you get?** Provide plots or tables and interpret them (not just raw numbers).
- **Why does it make sense?** Provide a short derivation / intuition / sanity check.
- **How does it relate to class?** Try at least one explicit connection.

Include at least one of the following (more is better):

- an ablation (remove a component and observe the effect);
- a baseline comparison (standard method vs. proposed);
- a stress test (vary a parameter; check robustness);
- a failure mode (when does it break or stop helping?).

This project is explicitly designed to reward **clever verification with small resources**.

Bonus. Clear negative results or well-explained discrepancies with the paper can be great if the analysis is rigorous. If you use external code (including public repositories), you must cite it and clearly state what you changed.

AI tools. If you use an AI assistant for writing or coding, you remain fully responsible for correctness and for understanding what you present.

4 FAQ-style clarifications

- **Do I need to reproduce the full paper's result?**
  No. You must reproduce or verify *one or two* precise claims that you have identified as the paper's key claims and explain what you learned.

- **Can I focus on theory only?**
  Possibly, if you include a real verification component (e.g., check assumptions numerically, derive a toy case, provide a counterexample, or connect to an implementable implication). A small simulation is still strongly recommended.
- **Can I focus on systems only?**
  Yes. Strong systems projects typically include a minimal prototype, microbenchmarks, and careful measurement methodology (timers, memory profiling, variance, etc.).
- **What if the paper needs special hardware (multi-node, TPU, etc.)?**
  You must redesign the experiment to be meaningful on limited resources (toy distributed setting, synthetic workloads, smaller models). The redesign is part of the creativity.
- **How many experiments are expected?**
  Quality over quantity. One strong, well-controlled experiment is better than five superficial ones.
- **What makes a poster "good" for this course?**
  A poster that teaches: clear visuals, minimal text, correct claims, and a memorable takeaway supported by evidence.