

Learning Dynamics of Meta-Learning in Small Model Pretraining

David Demitri Africa* Yuval Weiss
Paula Butterly Richard Diehl Martinez
University of Cambridge

Abstract

Large language models are powerful but costly. We ask whether meta-learning can make the pretraining of small language models not only faster but also more interpretable. We integrate first-order MAML with subset-masked LM pre-training, producing four LLama-style decoder-only models (11M–570M params), and evaluate on multilingual Universal NER. Compared with vanilla training, our hybrid setup (i) reaches the same loss up to 1.6x sooner, (ii) yields modest but consistent average gains on Universal NER at medium/large scales under equal compute (+2–3 percentage points), and (iii) and (iii) reveals phase-like learning dynamics: models first diversify their representations, then compress them in a pattern that aligns with improved episodic accuracy. These observations are correlational, not causal, and we do not claim generality beyond NER or across seeds. We also document a trade-off: perplexity on Paloma (a diverse language modeling benchmark spanning 18 domains; [Magnusson et al. \(2024\)](#)) is worse at most scales. Code, checkpoints and analysis logs are released.



davidafrica/pico-maml

DavidDemitriAfrica/pico-maml-train

1 Introduction

Small language models (SLMs) are attractive for privacy and energy reasons, but trail large models partly because they converge slowly and plateau early ([Godey et al., 2024](#); [Biderman et al., 2023](#); [Diehl Martinez et al., 2024](#)). As opposed to the common method of brute-force scaling, we explore a different axis: learning rules. First-order Model-Agnostic Meta-Learning (MAML) ([Finn et al., 2017](#)) promises a learn-to-learn initialization, yet has rarely been applied to decoder models, and

its effect on learning dynamics are poorly understood.

We address this by adding meta-learning in model pretraining,¹ interleaving ordinary next-token loss (keeps fluency) with 32-way subset-mask ([Bansal et al., 2020](#); [Li and Zhang, 2021](#)) episodes (forces rapid binding). Only a tiny MLP head is adapted in the inner loop, so we can track backbone weights without gradient noise. Our contributions are:

1. Four open SLMs (11M → 570M) trained with this **hybrid MAML rule**.
2. A public trainer that logs per-checkpoint **singular-value spectra**, head entropies and query accuracy.
3. A candid evaluation on Universal NER: **modest gains at medium/large scales** (+2–3 pp), alongside a **perplexity trade-off**.
4. Observational evidence of a **diversify-then-compress phase transition** in effective rank.

Reporting and scope. All pretraining and fine-tuning results are from a single shared seed per condition due to compute limits; we therefore report averages across datasets where applicable, avoid statistical claims, and treat learning-dynamics findings as exploratory. We limit generalization claims to NER and to our training regime.

2 Related Work

Meta-learning for NLP. (MAML; [Finn et al., 2017](#)) is an optimisation-based form of meta-learning that learns an initialisation from which a few gradient steps solve new tasks. It has been particularly successful in computer

*Corresponding
david.demitri.africa@gmail.com

author:

¹Using a lightweight modification of PICO-TRAIN ([Diehl Martinez, 2025](#)), a language model pretraining suite.

vision classification and reinforcement learning settings (Nichol et al., 2018). Within NLP, MAML has been adapted to a wide spectrum of supervised problems—including text classification, natural language inference, question answering, summarisation and named entity recognition—where a pre-trained encoder such as BERT is further fine-tuned on small datasets (Rajeswaran et al., 2019; Raghu et al., 2021; Hou et al., 2022). These studies operate (i) on encoder-only, masked-language models and (ii) at parameter counts close to the original 110M-parameter BERT. They leave open whether optimisation-based meta-learning helps decoder LMs and whether its benefits persist at larger parameter scales.

Meta-learning for pretraining. Initial NLP attempts applied MAML only at fine-tuning scale (Raghu et al., 2021; Hou et al., 2022). More recent work embeds bilevel objectives directly in pre-training (Miranda et al., 2023; Ke et al., 2021). While promising, these efforts evaluate only a single model size, focus on one downstream task, or release neither code nor weights, limiting reproducibility and obscuring scale trends. We embed meta-learning directly into the pretraining loop, evaluate on various unseen domains in an unseen task, and provide open weights (11M-570M) and layer-wise spectra, filling that gap.

Subset-Mask LMs (SMLMT). SMLMT constructs pseudo-tasks using a subset of vocabulary words (Bansal et al., 2020). Given an unlabeled text corpus, one selects a set of N words and builds an N -way classification task. For each chosen word, sentences containing it are collected and the word is masked out. The task is then to predict the masked word from the N candidates. Li and Zhang (2021) interleaves it with ProtoNet tasks; we interleave with vanilla LM updates and scale to 570M params.

Interpretable training dynamics. Various works discuss the training of language models in phase transitions (Olsson et al., 2022; Hoogland et al., 2024), describing broad changes in indicators as the model gains rapidly in capabilities over a short period of time. We study such phase transitions in the context of meta-learning in pretraining.

Effective-rank probes (entropy of singular values) highlight learning behavior in deep nets (Diehl Martinez et al., 2024). Lower-rank structure and rank compression are well documented in the literature (Huh et al., 2021; Galanti et al., 2022; Jaderberg et al., 2014), and we focus on the

timing and co-evolution of the measurements of effective-rank probes with episodic generalization under the hybrid objective (§5).

3 Method

We pretrain four decoder models at 11M, 65M, 181M and 570M parameters with a hybrid objective (Li and Zhang, 2021) that alternates conventional next-token prediction and first-order MAML episodes (Finn et al., 2017). The episodes are generated with Subset-Masked Language Modelling Tasks (SMLMT) (Bansal et al., 2020). This section details the backbone, the meta-learning episode, the optimisation schedule, and the downstream evaluation harness.

3.1 Baselines

The starting point is the open Pico decoder (Diehl Martinez, 2025), a LLAMA-style (Touvron et al., 2023) stack implemented in plain PyTorch. To maintain apples-to-apples comparability with the original models (and as such isolate the effect of introducing MAML to pretraining), we maintain the design choices and hyperparameter choices of the original Pico decoder models. A sequence of $L = 12$ decoder blocks receives 2048 input tokens. Each block performs RMSNorm (Zhang and Sennrich, 2019), grouped-query self-attention (Ainslie et al., 2023) with rotary position embeddings (Su et al., 2024), and a SwiGLU feed-forward network (Shazeer, 2020) that expands to $4d$ before projecting back to the model width d . Width is the only scale-dependent hyper-parameter: $d \in \{96, 384, 768, 1536\}$ for the tiny, small, medium and large variants. All models use 12 heads, 4 key-value heads and causal masking.

3.2 Task construction via SMLMT

SMLMT converts unlabelled text into few-shot classification tasks. From the corpus we sample a set of N content words, collect sentences that contain each word and replace that word with a single `<mask>`. The goal is to predict which of the N candidates was masked. Each episode supplies K support sentences and a disjoint query set. Table 1 shows an episode with $N = 4$ city names and $K = 2$ supports per class; the query asks the model to complete a new sentence about cherry blossoms. In practice we use $N = 32$ and $K = 4$ so the task entropy matches the five-bit next-token uncertainty

Set	Input (masked)	Label
Support ($K=2$ each)	I visited __ last summer. The sushi festival in __ was unforgettable. The Big Ben is in __. I caught the tube at __ yesterday. The Seine runs through __. She admired the art at the Louvre in __. The Forbidden City is in __. I sampled Peking duck in __.	Tokyo Tokyo London London Paris Paris Beijing Beijing
Query	I plan to travel to __ to see the cherry blossoms.	Tokyo

Table 1: Example SMLMT episode with $N=4$ classes and $K=2$ support sentences per class.

of English text.²

3.3 Optimiser, data, and monitoring

Training runs for 6000 outer updates on four A100 GPUs, with the original Pico-decoder models evaluated at the checkpoint after 6000 steps. Each GPU streams micro batches of 256 sequences from the 30 percent English subset of Dolma (Soldaini et al., 2024) that is already tokenised and chunked by Pico (Diehl Martinez, 2025). The outer optimiser is AdamW with peak learning rate 3×10^{-4} , 2500-step warm-up and cosine decay. Micro batches of 256 sequences are accumulated eight times giving an effective batch of 2048 (1024 for the 11M model). Every 100 steps we evaluate Paloma perplexity (Magnusson et al., 2024) and log the singular values of three attention and three feed-forward matrices to compute effective rank (Diehl Martinez et al., 2024). Query and support accuracies are also tracked.

3.4 Downstream protocol

Named entity recognition (NER), the downstream task for this study, is a fundamental NLP task that identifies and categorizes entities (e.g., persons, organizations, locations) within unstructured text (Chinchor and Robinson, 1997), and is used in healthcare (Kundeti et al., 2016; Polignano et al., 2021; Shafqat et al., 2022), law (Leitner et al., 2019; Au et al., 2022; Naik et al., 2023), business (Putthividhya and Hu, 2011; Alvarado et al., 2015; Zhao et al., 2021), and knowledge graph systems (Al-Moslmi et al., 2020). Specifically, we evaluate our models on Universal NER benchmark (Mayhew et al., 2024). UNER v1 comprises three categories of NER evaluation data, each built on top of Universal Dependencies (UD) (Nivre et al.,

²Shannon’s estimate of printed-English entropy is about 1.3 bits per character (Shannon, 1951); since English BPE tokens span on average about 4 characters (OpenAI, 2025), this implies roughly ≈ 5.2 bits/token. We therefore use 5 bits per token as a conservative rule of thumb.

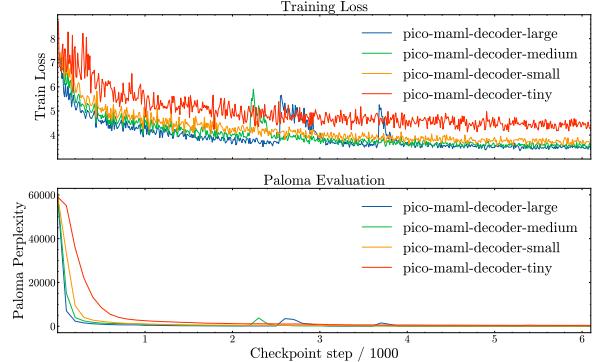


Figure 1: Training loss and Paloma perplexity across pretraining steps for all MAML models. Two-panel plot showing the evolution of (top) cross-entropy training loss and (bottom) Paloma perplexity, each as a function of global pretraining step.

Model	Train Loss @ 6k	Paloma Perplexity @ 6k
pico-decoder-tiny	5.31	786.85
pico-maml-decoder-tiny	4.44	422.42
pico-decoder-small	4.14	80.25
pico-maml-decoder-small	3.67	113.76
pico-decoder-medium	3.89	77.90
pico-maml-decoder-medium	3.49	78.63
pico-decoder-large	3.69	49.86
pico-maml-decoder-large	3.49	66.62

Table 2: For each model (rows) under vanilla vs. MAML pretraining (columns), shows cross-entropy loss and Paloma perplexity measured at exactly 6000 steps.

2016, 2020) tokenization and annotations: publicly available in-language treebanks, parallel UD (PUD) evaluation, and other eval-only sets (Appendix B).

After pretraining we load the checkpoint at step 6000 and attach a fresh linear classifier for UniversalNER. Two fine-tuning settings are used: head-only and full. In the head-only setting the Transformer is frozen so fine-tuning mirrors the inner loop, in the full setting all weights update. Fine-tuning uses AdamW at 3×10^{-5} for at most ten epochs with early stopping on development F_1 .

4 Model Pretraining

Training-perplexity tradeoff across scales. The prerequisite for modifying a pretraining method is ensuring the model still learns. All four Pico-MAML variants reach their respective vanilla loss 1.3–1.6× sooner (faster optimization), but Paloma perplexity is worse at most scales by 6000 steps (Table 2).

Contrary to expectation, MAML’s inductive bias may favor optimization over regularization. MAML accelerates convergence but degrades out-of-task fluency at most scales. However, this pat-

Model	Seen		Test-Only (PUD)		Test-Only (Other)	
	Head	Full	Head	Full	Head	Full
tiny (%)	-8.3	-3.0	+6.7	0.0	-37.5	+3.8
small (%)	+2.2	0.0	-17.2	-0.6	+46.7	+7.0
medium (%)	+1.9	+2.3	-4.6	+1.8	+14.8	+3.8
large (%)	+6.2	+4.8	+7.2	+3.5	+2.1	+8.1

Table 3: Relative percentage improvement of micro-F1 (higher = better) for head-only vs. full fine-tuning across seen, test-only (PUD), and low-resource language groups (other). Demonstrates MAML’s consistent 2–3 pp lift at medium/large scales under full tuning. **Green cells** indicate MAML improvements; **red cells** show degradations.

tern is consistent with known multi-task interference: the episodic discriminative objective improves adaptation signals but can conflict with next-token distributional modeling under fixed compute and a single set of hyperparameters (Kendall et al., 2017; Yu et al., 2020; Standley et al., 2020). Hence, it is unclear if the perplexity gap is an objective-mixing artifact or evidence that meta-learning inherently harms LM fluency.

5 Downstream NER Evaluation

Models are fine-tuned on each dataset in Universal NER (Mayhew et al., 2024; Nivre et al., 2016, 2020) with publicly available train and dev sets³. Results (averaged across each finetuning dataset) are shown as micro-F1 scores in Table 3, organized by evaluation group: seen (language with full train/test/dev splits), test-only (using Parallel Universal Dependencies PUD), and test-only low-resource languages (e.g., Cebuano, Tagalog). We report delta F1 as percentage points (pp) unless explicitly marked as percent change (%).

The most striking takeaway from this stage is that, when averaged across all evaluation steps in a category, absolute F1 remains low (≤ 0.35 , i.e., $\leq 35\%$) due to poor zero-shot transfer, especially for logographic scripts. Overall, MAML improves mean uplift is approximately +2–3 pp when averaged over all in-language datasets at medium/large scales, confirming a modest “learning-to-learn” effect under full adaptation.⁴

³Namely, ddt, ewt, set, bosque, snk, set, talbanken, gsd, gsdsimp, all.

⁴While these results are much worse in comparison to the baseline in the original Universal NER paper (Mayhew et al., 2024), this is likely because XLM-R_{large} is a multilingual model (Conneau et al., 2020) and the pretraining dataset for Pico is entirely in English.

Model	Danish	English	Croatian	Portuguese	Swedish
large (%)	+8.1	+14.8	+10.7	+8.6	+18.0

Table 4: Percentage relative improvement of MAML over vanilla for head-only tuning in the large model.

Model	Danish	English	Croatian	Portuguese	Swedish
tiny (%)	+3.4	+0.2	-1.6	-0.7	+6.1
small (%)	-3.9	-4.7	-1.9	-2.6	+4.9
medium (%)	+0.8	+4.8	+3.9	+1.2	+3.7
large (%)	+3.6	+4.4	-0.5	+4.2	+2.8

Table 5: Percentage-wise relative improvement of MAML over vanilla under full tuning for each language.

In-language NER gains suggest capacity-dependent meta-learning. To better understand how meta-initialization influences cross-lingual transfer on seen languages, F1 scores are broken down by dataset within the in-language group. The results are separated by tuning regime to clarify the extent to which meta-learned representations help when only the classifier is updated (head-only) versus when the entire model is fine-tuned.

In the head-only setting (Table 7), absolute F1 scores remain low across most datasets. Tiny models fail to generalize altogether. MAML shows the strongest and most consistent gains at large scales (Table 4)—most prominently on en_ewt, hr_set, and sv_talbanken—suggesting that episodic pre-training creates more adaptable feature spaces, particularly for common entity types and scripts. On Chinese (zh_gsd, zh_gsdsimp), performance is uniformly poor, confirming the baseline result in (Mayhew et al., 2024) that transfer from phonographic to logographic scripts is difficult.

In the full setting (Table 5), both vanilla and MAML-pretrained models achieve higher F1 scores across the board. MAML confers consistent +0.01–0.03 gains at medium and large scales, especially for structurally complex languages like Croatian. These relative gains grow as model capacity increases, indicating that larger models benefit more from MAML pretraining. Even in Chinese, where scores are lowest, MAML nudges performance upward. These gains confirm that meta-pretraining does more than support shallow transfer: it reshapes the optimization landscape of the full model in a way that accelerates convergence and improves generalization.

Taken together, these tables validate that MAML pretraining injects a scalable and tunable learning-to-learn signal. However, these average metrics do not tell the full story. Some settings, entity

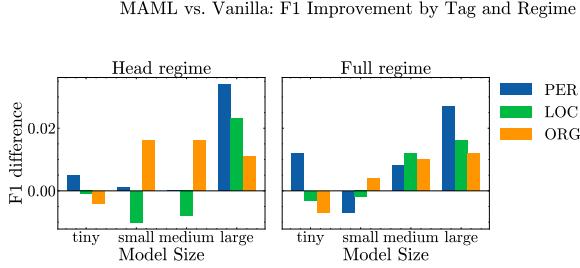


Figure 2: MAML-Vanilla micro-F1 difference by entity class and tuning regime, averaged across in-language datasets. Grouped bar charts reporting $\Delta F1 = F1_{MAML} - F1_{(\text{Vanilla})}$ for three named-entity classes—PERSON (PER), LOCATION (LOC) and ORGANIZATION (ORG)—for pico-MAML decoders of four sizes (tiny, small, medium, large), averaged over nine in-language NER datasets, over two fine-tuning regimes.

classes, and fine-tuning conditions benefit substantially more than others.

Class-specific prototype bias in entity recognition. We characterize the specific way MAML pretraining improves performance in NER by breaking down F1 score by entity class in Figure 2.

Meta-pretraining yields a clear capacity threshold in head-only adaptation. Under a frozen backbone, only the large model consistently converts its learned initialization into PER (+0.034) and LOC (+0.023) gains; medium and smaller variants lack the representational bandwidth to rewire person and place distinctions via a shallow classifier. By contrast, even medium and small models see gains in ORG (+0.016 F1) likely because organization names often include distinctive tokens (e.g., “Inc.”, “Corp.”, or “University”) that form rigid, token-level co-occurrence patterns. These simple patterns mirror the pseudo-classification episodes SMLMT generates, so a shallow classifier can latch onto them without requiring deep feature reconfiguration.

Full fine-tuning broadens and amplifies these effects. In the full setting, PER sees the largest MAML-induced lift (up to +0.027 in the large model). LOC improvements (+0.016 at large scale) climb more gradually: place names often span heterogeneous contexts and scripts (e.g. Zagreb vs. Beijing), so meta-pretraining must be supplemented by full gradient flow for location-specific embeddings. ORG continues to enjoy gains (+0.012 at large), reinforcing that organization recognition remains the simplest class to bootstrap from episodic tasks.

Model	Regime	Overall	Cebuano	Tagalog (TRG)	Tagalog (Ugnayan)
tiny	head	-100.0%	-100.0%	N/A	N/A
small	head	+151.1%	+209.6%	+315.7%	-15.7%
medium	head	+24.3%	+16.7%	-20.7%	+534.3%
large	head	+9.0%	+0.0%	+57.3%	-37.5%
tiny	full	-6.2%	-4.7%	-25.0%	+109.5%
small	full	+7.3%	-6.4%	+28.8%	+41.1%
medium	full	+0.0%	-1.0%	+1.4%	-2.1%
large	full	-8.0%	-14.5%	-1.6%	-0.8%

Table 6: Percentage change of MAML over vanilla zero-shot NER transfer (from English) F1 on low-resource languages (OTHER).

Significant zero-shot transfer gains in low-resource languages. Now, we discuss how inductive biases manifest in zero-shot cross-lingual transfer to low-resource languages—namely, Tagalog (t1) and Cebuano (ceb).

Tagalog and Cebuano are the two most widely spoken native languages in the Philippines, with tens of millions of first-language speakers each. Both are typologically Austronesian and low-resource, but differ significantly. Tagalog is a morphologically rich, predicate-initial language with a complex voice system that encodes syntactic roles (agent, patient, locative, etc.) through verbal affixes and aspect-marking (Kroeger, 1993; Schachter and Otanes, 1983; Ramos, 2021). Word order is flexible and often pragmatically driven, which weakens the utility of positional cues for tasks like named entity recognition. Cebuano is similarly Austronesian but morphologically simpler than Tagalog, with fewer voice alternations and less affixal variation (Tanangkingsing, 2011). It also does not consistently mark syntactic roles with overt case particles; entities must be inferred from context rather than surface markers (Sityar, 2000). Additionally, Cebuano exhibits a distinct orthographic tradition and more conservative vocabulary (e.g., less Spanish borrowing) (Bunye and Yap, 1971), which further distances it from the English-centric token distributions that dominate cross-lingual pretraining datasets. These characteristics make them ideal stress tests for testing the inductive bias of pretraining strategies like MAML.

In the head-only setting, MAML delivers its greatest impact on small and medium models. For example, the small head jumps from 0.088 to 0.221 overall—an absolute gain of 0.133 F1—and sees particularly large lifts in Cebuano (+0.153) and Tagalog-TRG (+0.262). The medium head also benefits substantially, improving from 0.259 to 0.322. Even the large head picks up a modest +0.030 F1. Only the tiny head collapses, reflecting

its inability to form reliable prototypes during meta-training. These patterns suggest that MAML’s episodic learning instills useful, language-agnostic representations in the classifier layers, enabling mid-size heads to generalize token-level cues to new languages without modifying the backbone.

Once we allow full fine-tuning, however, most of MAML’s advantages disappear at higher capacities. The small model retains a small $+0.026$ F1 edge, but the medium shows no net change and the large actually drops by 0.034 . This reversal implies that when every parameter is free to update, the strong gradient signals of full fine-tuning quickly override the meta-learned inductive biases, erasing or even inverting MAML’s earlier head-only gains. The tiny model again underperforms, consistent with its tendency to overfit during meta-training when unconstrained by a fixed backbone.

In the UNER benchmark, Tagalog and Cebuano serve as canonical low-resource, typologically distinct evaluation settings. Overall NER performance remains modest, but, as Table 6 shows, MAML provides meaningful zero-shot boosts in the head-only regime for small and medium models. These gains suggest that even without training exposure to these languages, the inductive biases from English episodic training transfer surprisingly well, at least for token-level prototypes.

6 Learning Dynamics

Despite clear convergence gains, the pretraining metrics alone leave several observations unexplained: the mid-training rebound and double-descent in Paloma perplexity, the abrupt jumps in support versus query accuracy, and the sudden collapse in representation rank. To understand this further, we now turn to a learning-dynamics analysis: tracking episodic support/query performance, classifier head statistics, and proportional effective rank throughout pretraining.

Effective meta-learning has a capacity threshold. To understand how MAML updates influence learning dynamics during pretraining, we track both support (training set in the inner loop) and query (held out final step in the inner loop) accuracy across training steps (Figure 3).

The small and medium models show clear signs of effective meta-learning. Support accuracy gradually increases and stabilizes around 6–7%, while query accuracy climbs steadily above 40%. This pattern indicates that the models are internalizing



Figure 3: Average support and query accuracy across pretraining steps for all models. Top: Average support-set accuracy (%) measured at the end of each inner-loop adaptation, as a function of the global pretraining step. Bottom: Corresponding average query-set accuracy (%) after adaptation.

a useful task prior, and show smooth convergence with relatively little instability.

The tiny model displays a distinct failure mode. While its support accuracy rises modestly, its query accuracy remains stagnant, hovering just above chance (10%). This suggests the model memorizes support examples but fails to learn task-generalizable features—a canonical symptom of underparameterization in meta-learning (Finn et al., 2017; Rajeswaran et al., 2019). In effect, it lacks the representational bandwidth to encode a shared inductive bias across tasks.

The large model shows a late-phase rise in query accuracy after 4,500 steps, coinciding with stabilization of head-weight variance. This suggests a phase-like reorganization where the model consolidates a useful episodic prior after a prolonged plateau. In the MAML setting, this may correspond to the model first learning how to adapt, before learning to generalize from adaptation.

Taken together, these patterns confirm that meta-learning is most stable within a mid-capacity regime. Models must be large enough to encode reusable structure, but not so large that their learning becomes erratic. These insights help contextualize downstream findings: the best generalization often arises from models that strike a balance between representational power and stable task-level adaptation.

Classifier head weight variance reveals adaptation behavior. To probe how episodic adaptation reshapes the backbone’s feature geometry, we track the mean and standard deviation of the

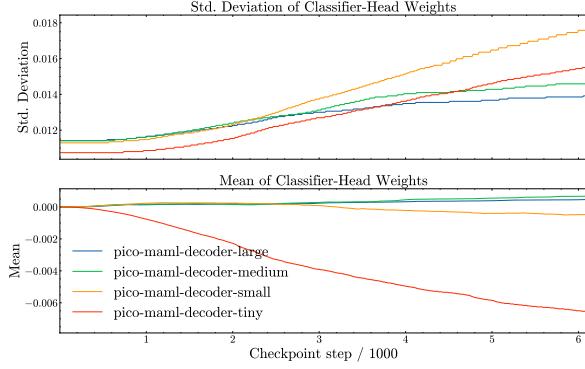


Figure 4: Evolution of classifier head weights during meta-training. Top: Standard deviation of the task-specific classifier head weights (in logits space). Bottom: Mean of the classifier head weights.

episodically adapted classifier head across training (Figure 4). Because the inner loop updates only this shallow head on frozen backbone features, its across-episode weight statistics act as a lightweight linear-probe proxy for class separability: under softmax on fixed features, class weight vectors tend to align with differences between class means, so greater dispersion (std) across head weights indicates larger between-class margins induced by the backbone, while transient spikes without sustained query gains suggest support overfit rather than stable generalization. We therefore relate inflections in mean/std to simultaneous changes in support/query accuracy to contextualize adaptation quality.

The top panel shows the standard deviation of head weights. All models exhibit growth in weight variance, indicating increasing expressivity in the task-specific head. The small model diverges most sharply, with its weight variance surpassing all others after 2k steps. This suggests an over-specialization effect: the model learns to adapt aggressively to each task, potentially at the cost of stability. In the lower panel, the mean of the head weights remains near zero for most models, but the tiny model is an outlier. It accumulates a strong bias in one direction over training, indicating that its head converges toward a fixed mapping that is minimally updated across episodes. This aligns with earlier diagnostics showing that its gradient norms collapse early in training.

These dynamics reinforce the idea that episodic MAML indeed induces a scale-sensitive tradeoff: in higher-capacity models, episodic gradients drive generalizable structure into the shared initialization; in lower-capacity models, this same pressure

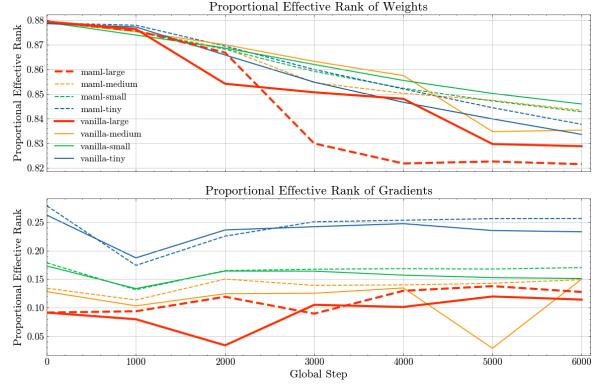


Figure 5: Proportional effective rank of MAML and vanilla models on available checkpoints until 6k steps. Top: weights; bottom: gradients.

can cause drift or collapse.

Evidence of representation collapse and reorganization. To understand how MAML alters internal representations, we track *proportional effective rank* (PER), a structure-sensitive metric during training applied to both weights and gradients in the attention layers (Figure 5).

Following Roy and Vetterli (2007) and Diehl Martinez et al. (2024), effective rank measures the entropy of the singular value spectrum of a matrix, while PER normalizes this by the total dimensionality:

$$\text{PER}(W) = \frac{\exp(-\sum_i p_i \log p_i)}{d}$$

where $p_i = \frac{\sigma_i}{\sum_j \sigma_j}$. PER captures the extent to which the model’s representations or updates span a full-dimensional space; a decline in PER indicates compression or structural specialization.

Key Finding: Phase Transition in Large Model

Across all MAML-pretrained models, PER declines over training, but the large model exhibits an **abrupt, synchronized drop** at step ~ 3000 in:

- **Proportional effective rank (PER)**
- Paloma perplexity (after initial rise)
- **Query accuracy** (sharp jump from plateau)

We interpret this behavior as a **representational phase transition**: the model initially fits the objective using diffuse, high-dimensional representations, which are later compressed into task-specialized, low-rank structures. The descent in

PER lags behind the initial perplexity gains, and only after this drop does the second descent in Paloma begin. There is no strong evidence of a comparable phase transition in the vanilla models. While the large and medium variants show mild inflection points in loss and perplexity around step 3000, these are gradual and lack the coordinated sharpness seen in the MAML-trained models.

This suggests that MAML’s bilevel updates and episodic task pressure may help reorganize the optimization landscape to favor discrete qualitative shifts in representation. As explored in Olsson et al. (2022); Wang et al. (2024); Hoogland et al. (2024), model training often proceeds in qualitatively distinct stages: from brute-force fitting, to intermediate rule memorization, to compressed algorithmic abstraction. The drop in PER may signal such a transition—from early diffuse representations to compressed heads tuned to solve the repeated structure of SMLMT episodes. This representational transition is also reflected in the model’s adaptation performance. Around the same step where PER and Paloma perplexity undergo a sharp drop (step ~ 3000), both support and query accuracies rise abruptly (see Figure 3). Prior to this point, query accuracy remains relatively flat, indicating that the model struggles to generalize from support to query examples. But after the phase transition, the model rapidly learns to extrapolate, with query accuracy climbing from near random to over 0.5.

This synchrony across metrics provides compelling evidence of a coordinated phase shift in the model’s learning trajectory. When looking into more granular checkpoints (Figure 6), there is clearer evidence that the model transitions from an early stage where it relies on diffuse representations to a later stage where it reorganizes both its representations and update paths into a lower-dimensional, more modular form capable of few-shot generalization. That said, this phase behavior appears scale-sensitive as it is absent in smaller scales. This suggests that the capacity to reorganize may be gated by scale, and that below a certain threshold, the inductive pressure of MAML induces collapse rather than modularization.

7 Conclusion

We interleaved first-order MAML episodes with decoder pretraining and analyzed dynamics across four SLM scales. Under equal compute, the hybrid objective accelerates optimization but trades

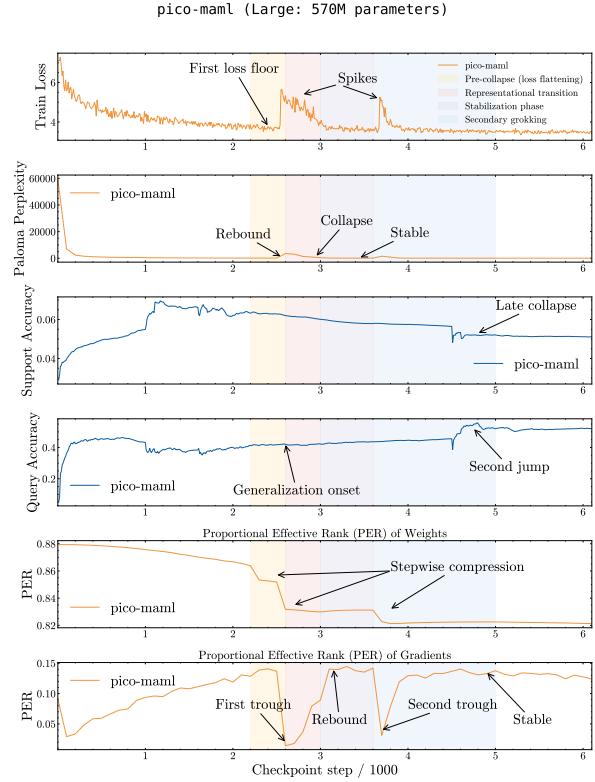


Figure 6: Dynamics of pico-maml-decoder-large over 6000 pretraining steps. **Pink shaded region** marks the phase transition (steps 2600-3200) where PER collapses, perplexity drops, and query accuracy jumps.

off perplexity at most scales; downstream, it brings modest average NER gains (+2–3 pp) at medium/large scales. Spectral logs expose a phase-like diversify–then–compress pattern that coincides with improving episodic query accuracy in the large model. Given our NER-only, single-seed scope, we present these as tools and observations rather than broad performance claims.

However, while our evaluation focuses exclusively on named entity recognition, the underlying mechanism—episodic adaptation via SMLMT—is task-agnostic. In principle, the same hybrid objective could be applied to other sequence labeling tasks (e.g., part-of-speech tagging, syntactic chunking) or even structured prediction problems that admit few-shot formulations. Whether the phase transitions and rank-compression patterns we observe generalize to non-linguistic domains (e.g., code generation, mathematical reasoning) remains an open question. Future work should explore whether MAML’s inductive bias is inherently suited to token-level structure learning or whether it confers broader benefits across modalities and task families.

Relatedly, other natural extensions suggest themselves. Future work should also include multi-seed and hyperparameter sweeps (inner LR, episode frequency), multilingual pretraining to test cross-script transfer, varying which layers adapt in the inner loop, and evaluation on non-NER tasks (e.g., classification, QA, reasoning), as the architectural design space is rather large. In terms of exploratory work, a natural next step is to learn whether the same phase transition re-emerges when the corpus is multilingual, which would clarify why cross-script transfer remains the weak point of the present models. Varying which backbone layers adapt, how many steps they receive and how frequently episodes are interleaved may unlock better compute–capability trade-offs. Finally, the clear correlation between the effective-rank collapse and downstream utility hints that spectral diagnostics might serve as a self-supervised early-stopping signal.

Limitations

All training runs stop at exactly six thousand outer steps, a horizon that may be too short for the largest model, so the observed perplexity gap between MAML and vanilla training could shrink or even reverse if optimisation were allowed to continue. Our downstream evaluation focuses on a single task family, sequence labelling, so it remains unclear whether the same advantages would materialise on reasoning or generation-quality benchmarks. Because the corpus is predominantly English, improvements in low-resource or logographic languages remain modest; a more diverse corpus may alter both quantitative and qualitative conclusions. Hyper-parameters such as the hybrid episode probability, the inner-loop learning rate and the 32-way 4-shot episode size were transferred unchanged across scales; dedicated tuning might further modify the trade-off between convergence speed and final perplexity. Models were trained on academic budget, which limited training to 6000 outer steps. Some interesting training dynamics only appear after a very extended period of training, and future work should study this long-term behavior. Finally, each condition was run with a single random seed owing to GPU constraints, so although the phase transition appears robust, the exact magnitude of the gains should be interpreted with caution.

Acknowledgments

This work was supported by a grant from the Accelerate Programme for Scientific Discovery, made possible by a donation from Schmidt Futures. David Demitri Africa is supported by the Cambridge Trust and the Jardine Foundation. Richard Diehl Martinez is supported by the Gates Cambridge Trust (grant OPP1144 from the Bill & Melinda Gates Foundation). Suchir Salhan is supported by Cambridge University Press & Assessment.

References

Joshua Ainslie, James Lee-Thorp, Michiel de Jong, Yury Zemlyanskiy, Federico Lebron, and Sumit Sanghai. 2023. Gqa: Training generalized multi-query transformer models from multi-head checkpoints. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 4895–4901.

Tareq Al-Moslmi, Marc Gallofré Ocaña, Andreas L Opdahl, and Csaba Veres. 2020. Named entity ex-

traction for knowledge graphs: A literature overview. *IEEE Access*, 8:32862–32881.

Julio Cesar Salinas Alvarado, Karin Verspoor, and Timothy Baldwin. 2015. Domain adaption of named entity recognition to support credit risk assessment. In *Proceedings of the australasian language technology association workshop 2015*, pages 84–90.

Ting Wai Terence Au, Ingemar J Cox, and Vasileios Lampos. 2022. E-ner—an annotated named entity recognition corpus of legal text. *arXiv preprint arXiv:2212.09306*.

Trapit Bansal, Rishikesh Jha, Tsendsuren Munkhdalai, and Andrew McCallum. 2020. *Self-supervised meta-learning for few-shot natural language classification tasks*. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 522–534, Online. Association for Computational Linguistics.

Stella Biderman, Hailey Schoelkopf, Quentin Gregory Anthony, Herbie Bradley, Kyle O’Brien, Eric Halahan, Mohammad Afzal Khan, Shivanshu Purohit, Usvsn Sai Prashanth, Edward Raff, Aviya Skowron, Lintang Sutawika, and Oskar Van Der Wal. 2023. *Pythia: A suite for analyzing large language models across training and scaling*. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 2397–2430. PMLR.

Maria Victoria R. Bunye and Elsa Paula Yap. 1971. *Cebuano Grammar Notes and Cebuano for Beginners*. University of Hawaii Press, Honolulu.

Nancy Chinchor and Patricia Robinson. 1997. Muc-7 named entity task definition. In *Proceedings of the 7th Conference on Message Understanding*, volume 29, pages 1–21.

Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. *Unsupervised cross-lingual representation learning at scale*. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.

Richard Diehl Martinez. 2025. *Pico: A lightweight framework for studying language model learning dynamics*.

Richard Diehl Martinez, Pietro Lesci, and Paula Buttery. 2024. *Tending towards stability: Convergence challenges in small language models*. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 3275–3286, Miami, Florida, USA. Association for Computational Linguistics.

Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. *Model-agnostic meta-learning for fast adaptation of deep networks*. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1126–1135. PMLR.

Tomer Galanti, Zachary S Siegel, Aparna Gupte, and Tomaso Poggio. 2022. Sgd and weight decay secretly minimize the rank of your neural network. *arXiv preprint arXiv:2206.05794*.

Nathan Godey, Éric Villemonte de la Clergerie, and Benoît Sagot. 2024. *Why do small language models underperform? studying language model saturation via the softmax bottleneck*. In *First Conference on Language Modeling*.

Jesse Hoogland, George Wang, Matthew Farrugia-Roberts, Liam Carroll, Susan Wei, and Daniel Murfet. 2024. The developmental landscape of in-context learning. *arXiv preprint arXiv:2402.02364*.

Zejiang Hou, Julian Salazar, and George Polovets. 2022. *Meta-learning the difference: Preparing large language models for efficient adaptation*. *Transactions of the Association for Computational Linguistics*, 10:1249–1265.

Minyoung Huh, Hossein Mobahi, Richard Zhang, Brian Cheung, Pulkit Agrawal, and Phillip Isola. 2021. The low-rank simplicity bias in deep networks. *arXiv preprint arXiv:2103.10427*.

Max Jaderberg, Andrea Vedaldi, and Andrew Zisserman. 2014. Speeding up convolutional neural networks with low rank expansions. *arXiv preprint arXiv:1405.3866*.

Zhen Ke, Liang Shi, Songtao Sun, Erli Meng, Bin Wang, and Xipeng Qiu. 2021. *Pre-training with meta learning for Chinese word segmentation*. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5514–5523, Online. Association for Computational Linguistics.

Alex Kendall, Yarin Gal, and Roberto Cipolla. 2017. *Multi-task learning using uncertainty to weigh losses for scene geometry and semantics*. *Preprint*, arXiv:1705.07115.

Paul R. Kroeger. 1993. *Phrase Structure and Grammatical Relations in Tagalog*. Dissertations in Linguistics. Center for the Study of Language and Information (CSLI) Publications, Stanford, CA. 257 pp.

Srinivasa Rao Kundeti, J Vijayananda, Srikanth Mujjiga, and M Kalyan. 2016. Clinical named entity recognition: Challenges and opportunities. In *2016 IEEE International Conference on Big Data (Big Data)*, pages 1937–1945. IEEE.

Elena Leitner, Georg Rehm, and Julian Moreno-Schneider. 2019. Fine-grained named entity recognition in legal documents. In *International conference on semantic systems*, pages 272–287. Springer.

Yue Li and Jiong Zhang. 2021. [Semi-supervised meta-learning for cross-domain few-shot intent classification](#). In *Proceedings of the 1st Workshop on Meta Learning and Its Applications to Natural Language Processing*, pages 67–75, Online. Association for Computational Linguistics.

Ian Magnusson, Akshita Bhagia, Valentin Hofmann, Luca Soldaini, Ananya Harsh Jha, Oyvind Tafjord, Dustin Schwenk, Evan Walsh, Yanai Elazar, Kyle Lo, and 1 others. 2024. Paloma: A benchmark for evaluating language model fit. *Advances in Neural Information Processing Systems*, 37:64338–64376.

Stephen Mayhew, Terra Blevins, Shuheng Liu, Marek Suppa, Hila Gonen, Joseph Marvin Imperial, Börje Karlsson, Peiqin Lin, Nikola Ljubešić, Lester James Miranda, Barbara Plank, Arij Riabi, and Yuval Pinter. 2024. [Universal NER: A gold-standard multilingual named entity recognition benchmark](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 4322–4337, Mexico City, Mexico. Association for Computational Linguistics.

Brando Miranda, Patrick Yu, Saumya Goyal, Yu-Xiong Wang, and Sanmi Koyejo. 2023. [Is pre-training truly better than meta-learning?](#) *Preprint*, arXiv:2306.13841.

Varsha Naik, Purvang Patel, and Rajeswari Kannan. 2023. Legal entity extraction: An experimental study of ner approach for legal documents. *International Journal of Advanced Computer Science and Applications*, 14(3).

Alex Nichol, Joshua Achiam, and John Schulman. 2018. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999*.

Joakim Nivre, Marie-Catherine de Marneffe, Filip Ginter, Yoav Goldberg, Jan Hajič, Christopher D. Manning, Ryan McDonald, Slav Petrov, Sampo Pyysalo, Natalia Silveira, Reut Tsarfaty, and Daniel Zeman. 2016. Universal dependencies v1: A multilingual treebank collection. In *Proceedings of the 10th Language Resources and Evaluation Conference (LREC 2016)*, pages 1659–1666, Portorož, Slovenia.

Joakim Nivre, Marie-Catherine de Marneffe, Filip Ginter, Jan Hajič, Christopher D. Manning, Sampo Pyysalo, Sebastian Schuster, Francis Tyers, and Daniel Zeman. 2020. [Universal Dependencies v2: An evergrowing multilingual treebank collection](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 4034–4043, Marseille, France. European Language Resources Association.

Catherine Olsson, Nelson Elhage, Neel Nanda, Nicholas Joseph, Nova DasSarma, Tom Henighan, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, and 1 others. 2022. In-context learning and induction heads. *arXiv preprint arXiv:2209.11895*.

OpenAI. 2025. What are tokens and how to count them? Accessed May 2025.

Marco Polignano, Marco de Gemmis, Giovanni Semeraro, and 1 others. 2021. Comparing transformer-based ner approaches for analysing textual medical diagnoses. In *CLEF (Working Notes)*, pages 818–833.

Duangmanee Putthividhya and Junling Hu. 2011. Bootstrapped named entity recognition for product attribute extraction. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 1557–1567.

Aniruddh Raghu, Jonathan Lorraine, Simon Kornblith, Matthew McDermott, and David K Duvenaud. 2021. Meta-learning to improve pre-training. *Advances in Neural Information Processing Systems*, 34:23231–23244.

Aravind Rajeswaran, Chelsea Finn, Sham M Kakade, and Sergey Levine. 2019. Meta-learning with implicit gradients. *Advances in neural information processing systems*, 32.

Teresita V Ramos. 2021. *Tagalog structures*. University of Hawaii Press.

Olivier Roy and Martin Vetterli. 2007. The effective rank: A measure of effective dimensionality. In *2007 15th European signal processing conference*, pages 606–610. IEEE.

Paul Schachter and Fe T. Otanes. 1983. *Tagalog Reference Grammar*. University of California Press, Berkeley, CA.

Sarah Shafqat, Hammad Majeed, Qaisar Javaid, and Hafiz Farooq Ahmad. 2022. Standard ner tagging scheme for big data healthcare analytics built on unified medical corpora. *Journal of Artificial Intelligence and Technology*, 2(4):152–157.

Claude E. Shannon. 1951. Prediction and entropy of printed english. *Bell System Technical Journal*, 30(1):50–64.

Noam Shazeer. 2020. Glu variants improve transformer. *arXiv preprint arXiv:2002.05202*.

Emily Sityar. 2000. *The Topic and Y Indefinite in Cebuano*, pages 145–165. Springer Netherlands, Dordrecht.

Luca Soldaini, Rodney Kinney, Akshita Bhagia, Dustin Schwenk, David Atkinson, Russell Author, Ben Boigin, Khyathi Chandu, Jennifer Dumas, Yanai Elazar, and 1 others. 2024. Dolma: an open corpus of three trillion tokens for language model pretraining research. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15725–15788.

Trevor Standley, Amir R. Zamir, Dawn Chen, Leonidas Guibas, Jitendra Malik, and Silvio Savarese. 2020. [Which tasks should be learned together in multi-task learning?](#) *Preprint*, arXiv:1905.07553.

Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. 2024. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063.

Michael Tanangkingsing. 2011. *A Functional Reference Grammar of Cebuano: A Discourse-Based Perspective*. Peter Lang, Berlin.

Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, and 1 others. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.

George Wang, Matthew Farrugia-Roberts, Jesse Hoogland, Liam Carroll, Susan Wei, and Daniel Murfet. 2024. Loss landscape geometry reveals stagewise development of transformers. In *High-dimensional Learning Dynamics 2024: The Emergence of Structure and Reasoning*.

Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn. 2020. **Gradient surgery for multi-task learning**. *Preprint*, arXiv:2001.06782.

Biao Zhang and Rico Sennrich. 2019. Root mean square layer normalization. *Advances in Neural Information Processing Systems*, 32.

Lingyun Zhao, Lin Li, Xinhao Zheng, and Jianwei Zhang. 2021. A bert based sentiment analysis and key entity detection approach for online financial texts. In *2021 IEEE 24th International conference on computer supported cooperative work in design (CSCWD)*, pages 1233–1238. IEEE.

A Pseudocode

Below is the pseudocode for the MAML and vanilla pretraining setup.

Algorithm 1 Distributed SMLMT Loop

```

Initialize model  $f_\theta$ , head  $h_\phi$ , outer optimizer, and inner
SGD on  $h_\phi$ 
step  $\leftarrow 0$ 
for each sub-batch  $B$  from dataloader do
     $X \leftarrow \text{AllGather}(B \text{ inputs})$   $\triangleright$  across devices
     $r \leftarrow \text{Broadcast}(\text{Uniform}(0, 1))$ 
    if  $r < \rho$  then
         $(S, Q, ys, y_Q) \leftarrow \text{MaskTokens}(X)$ 
         $\phi_{\text{snap}} \leftarrow \phi$   $\triangleright$  save head params
        for  $t = 1$  to  $T_{\text{inner}}$  do
             $\ell_S \leftarrow \text{CE}(h_\phi(f_\theta(S)), ys)$ 
             $\phi \leftarrow \phi - \alpha, \nabla_\phi \ell_S$ 
        end for
         $\ell \leftarrow \text{CE}(h_\phi(f_\theta(Q)), y_Q)$ 
         $\phi \leftarrow \phi_{\text{snap}}$   $\triangleright$  restore head
    else
         $X_{\text{in}} \leftarrow X$  without last token;  $Y \leftarrow X$  without first
        token
         $\ell \leftarrow \text{CE}(f_\theta(X_{\text{in}}), Y)$ 
    end if
    Backward( $\ell, /, \text{accum\_steps}$ )
    if  $(\text{step} + 1) \bmod \text{accum\_steps} = 0$  then
        OptimizerStep(); SchedulerStep(); ZeroGrad()
        AggregateMetrics( $\ell$ ); Barrier()
    end if
    step  $\leftarrow \text{step} + 1$ 
end for

```

Algorithm 2 Distributed AR Loop

```

Initialize configs, Fabric/strategy, tokenizer, model  $f_\theta$ , optimizer
Prepare dataloader and distribute it
step  $\leftarrow 0$ ; ZeroGrad()
for each sub-batch  $B$  from dataloader do
     $X \leftarrow \text{AllGather}(B \text{ inputs})$   $\triangleright$  across devices
     $X_{\text{in}} \leftarrow X$  without last token;  $Y \leftarrow X$  without first
    token
     $\ell \leftarrow \text{CE}(f_\theta(X_{\text{in}}), Y)$ 
    Backward( $\ell, /, \text{accum\_steps}$ )
    if  $(\text{step} + 1) \bmod \text{accum\_steps} = 0$  then
        OptimizerStep(); SchedulerStep(); ZeroGrad()
        Barrier()  $\triangleright$  optional
    end if
    step  $\leftarrow \text{step} + 1$ 
end for

```

A.1 Multi-GPU processing

Pico already uses Lightning-Fabric data parallelism but meta-learning introduces various demands that make multi-GPU processing complicated. A Bernoulli draw is done on one GPU and broadcast so all ranks choose the same objective. Support and query tensors are constructed on rank 0 then scattered, because per-rank random masks would destroy gradient equivalence. Every GPU

performs the same ten head updates before any gradient is communicated. A stray early all_reduce would mix gradients from different inner steps, so we place an explicit barrier between inner and outer phases.

B Universal NER Datasets

To comprehensively evaluate the pretraining method, each permutation of fine-tuning setup (<{head-only, full}, fine-tuning dataset (<{da_ddt, ..., zh_gsdssimp, all}>) (where all consists of all available training sets), model size (<{tiny, small, medium, large})), and pretraining setup (<{vanilla, MAML}>) is evaluated, for a total of 160 evaluation runs.

- **Publicly Available In-language treebanks** (9 langs): full train/dev/test splits, identical to the official UD partitions.

- da_ddt, en_ewt, hr_set, pt_bosque, sk_snk, sr_set, sv_talbanken, zh_gsd, zh_gsdssimp

- **Parallel UD (PUD) evaluation** (6 langs): single test.txt files, all sentence-aligned across German, English, Portuguese, Russian, Swedish and Chinese.

- de_pud, en_pud, pt_pud, ru_pud, sv_pud, zh_pud

- **Other eval-only sets** (3 langs): small test splits for low-resource languages.

- ceb_gja (Cebuano), tl_trg (Tagalog TRG), tl_ugnayan (Tagalog Ugnayan)

C Supplementary Figures

C.1 Supplementary Tables

Table 7: Micro-F1 scores (rows: selected datasets, columns: vanilla vs. MAML) under head-only tuning for large models. Highlights which languages benefit most from MAML without full adaptation.

Model	da_ddt	en_ewt	hr_set	pt_bosque	sk_snk	sr_set	sv_talbanken	zh_gsd	zh_gsdssimp
vanilla_tiny	0.004	0.031	0.011	0.000	0.004	0.009	0.000	0.005	0.009
maml_tiny	0.000	0.057	0.000	0.014	0.014	0.002	0.000	0.000	0.005
vanilla_small	0.000	0.196	0.123	0.099	0.047	0.056	0.020	0.000	0.003
maml_small	0.004	0.156	0.162	0.104	0.063	0.044	0.000	0.003	0.005
vanilla_medium	0.141	0.252	0.311	0.240	0.153	0.325	0.065	0.010	0.020
maml_medium	0.087	0.288	0.329	0.243	0.136	0.362	0.108	0.005	0.010
vanilla_large	0.247	0.366	0.401	0.337	0.178	0.422	0.261	0.034	0.039
maml_large	0.267	0.420	0.444	0.366	0.191	0.455	0.308	0.023	0.040

Table 8: Percentage relative improvement of MAML over vanilla for head-only tuning in the large model.

Model	da_ddt	en_ewt	hr_set	pt_bosque	sk_snk	sr_set	sv_talbanken	zh_gsd	zh_gdsimp
Large (%)	+8.1	+14.8	+10.7	+8.6	+7.3	+7.8	+18.0	-32.4	+2.6

Table 9: Percentage-wise relative improvement of MAML over vanilla under full tuning for each language.

Model	da_ddt	en_ewt	hr_set	pt_bosque	sk_snk	sr_set	sv_talbanken	zh_gsd	zh_gdsimp
tiny (%)	+3.4	+0.2	-1.6	-0.7	-2.4	+1.5	+6.1	-9.2	-2.7
small (%)	-3.9	-4.7	-1.9	-2.6	+3.4	+0.9	+4.9	+1.6	+4.9
medium (%)	+0.8	+4.8	+3.9	+1.2	+0.3	-0.3	+3.7	+5.0	+8.2
large (%)	+3.6	+4.4	-0.5	+4.2	+5.7	+1.3	+2.8	+3.4	+5.0