



A SOVEREIGN, OPEN-SOURCE FOUNDATION MODEL FOR GERMAN AND ENGLISH

SOOFI S PRETRAINING REPORT v1.0

The Soofi-Team*

Core Team: Benedikt Droste¹⁰, David Fitzek^{3,9}, Ruben Härle⁵, Lukas Helff^{2,5}, Maximilian Idahl¹⁰, Alex Jude^{3,9}, Abbas Goher Khan³, Maurice Kraus⁵, Timm Ruland^{3,9}, Richard Rutmann^{3,9}, Sebastian Sztwiertnia⁵

Contributors: Markus Frey^{3,9}, Daniil Gurgurov², Jan Pfister⁶, Tom Röhr⁷, Sebastian von Rohrscheidt⁷

Advisors: Jörg Bienert¹, Nicolas Flores-Herr³, Simon Gottschalk⁸, Andreas Hotho⁶, Kristian Kersting^{2,5,11}, Joachim Köhler³, Alexander Löser⁷, Wolfgang Nejdl⁸, Simon Ostermann², Jan Plogsties⁴, Patrick Putzky¹²

Technical Leads: Mehdi Ali^{3,9}, Michael Fromm^{3,9}, Max Lübbering^{3,9}

Affiliations: ¹KI Bundesverband, ²DFKI, ³Fraunhofer IAIS, ⁴Fraunhofer IIS, ⁵Technische Universität Darmstadt, ⁶Universität Würzburg, ⁷Berliner Hochschule für Technik, ⁸L3S Research Center, ⁹Lamarr, ¹⁰ellamind, ¹¹hessian.AI, ¹²Merantix Momentum


Coordination & Funding: Consortium coordinated by the KI Bundesverband. Funded by the German Federal Ministry for Economic Affairs and Energy (BMWE).

* Authors are listed alphabetically. Detailed Contributions in Appendix A.

Abstract

We present Soofi S 30B-A3B, a sovereign, open-source Mixture-of-Experts (MoE) hybrid Mamba Transformer foundation model for German and English. Its hybrid design activates only 3B of 30B parameters per token and keeps the inference cache near-constant as context grows, giving it a decisive throughput advantage over dense models for long-context, high-concurrency deployment. Pretrained on roughly 27 trillion tokens with deliberately up-weighted German, Soofi S matches dense 14 to 27B models on aggregate English and German benchmarks while achieving the best code aggregates in both languages among 17 open base models, and outperforms every European sovereign baseline in our comparison, including ones far larger in active parameters. Among fully open models, Soofi S obtains the highest English and German evaluation scores, ahead of Olmo 3 32B and Apertus 70B. Soofi S was built end-to-end on the German Industrial AI Cloud, a sovereign HPC-scale AI infrastructure operated by Deutsche Telekom in Munich. Soofi S will be released under highly permissive, open-access terms: weights, selected intermediate checkpoints¹, full per-source data accounting, hyperparameters, and training and evaluation code. Where source licenses permit, data-construction artifacts are released under permissive licenses; commercially licensed sources are documented with aggregate statistics and exact mixture accounting.

1 Introduction



(a) Capability vs. measured aggregate decode TPS. (b) Aggregate decode TPS scaling with context.

Figure 1: Long-context serving efficiency. Soofi S combines frontier-level capability with the highest measured aggregate long-context decode TPS, and unlike full-attention dense baselines maintains high throughput as context grows. Panel (a) plots Capability Index versus measured aggregate decode TPS/GPU at 40K context and batch 32. The Capability Index averages five benchmark groups, i.e., Code, GSM8K, GPQA-Diamond, English aggregate, and German aggregate, after normalizing each group to the best plotted model. Aggregate decode TPS/GPU is measured with a TP=1, one-B200 vLLM latency-subtraction protocol. Panel (b) shows measured aggregate decode TPS/GPU as a function of input context length under the same batch-32 protocol. For both panels higher is better.

Open language models have improved at remarkable speed, yet three gaps remain conspicuous for anyone deciding what to actually deploy.

¹<https://huggingface.co/Soofi-Project>

The first is openness in substance rather than name: despite a proliferation of capable models, the majority of releases remain weight-only releases, documenting their training with little more than an aggregate token count and omitting the data, recipes, and decisions needed to reproduce or audit them.

The second is language: general-purpose multilingual models are either English-centric or spread their capacity thinly across dozens of languages, leaving German underrepresented relative to its economic and scientific weight. Dedicated European efforts to date [3, 26, 53, 5] have prioritized openness and language coverage over frontier capability.

The third gap is the one that most directly governs the deployment cost, and where our chosen architecture is aimed. At economic concurrency the price of generation is set not by how many parameters a model nominally contains, nor even by how many it activates per token, but by memory bandwidth: every decoded token must re-read the model weights and, for a Transformer, the attention cache of every sequence in the batch. As contexts grow into the tens or hundreds of thousands of tokens and many requests are served in parallel, this key-value (KV) cache comes to dominate, and full-attention dense models slow down accordingly. A model that keeps its per-sequence state small and near-constant in context length therefore enjoys a structural advantage that compounds in exactly the regime **long context**, **high concurrency** that matters most in production.

Soofi S 30B-A3B addresses all three at once: It is a Mixture-of-Experts (MoE) hybrid Mamba Transformer [61, 60] trained to excel in both German and English and will be released radically open: not weights alone, but the full set of artifacts required to audit every stage of training and, where source licenses permit, rebuild the data mixture, in the spirit of recent fully open efforts [5, 64]. Architecturally it adopts the openly published Nemotron 3 Nano reference design [61] (Section 2): Mamba-2 layers [16] carry most of the sequence mixing with a fixed-size recurrent state, only 6 of its 52 layers maintain a KV cache, and sparse MoE layers activate just 3.2 of 31.6 billion parameters per token—the capacity of a 30B network at roughly the inference cost of a 3B one.

Contributions. We summarize our contributions below:

- **German–English champion.** Soofi S achieves best-in-class performance on aggregate English and German base model benchmarks, including strong performance on code, mathematics and German regional-knowledge. It is the strongest fully open model in our evaluation on English and German benchmarks, and matches or outperforms every European sovereign baseline in our comparison on every German benchmark in our suite, while matching dense 14–27B international models on English and German aggregate performance at a fraction of their active-parameter cost (Section 4).
- **Full data transparency.** We release the complete pretraining corpus statistics (Table 6, Table 8, Table 10) and reproducible construction scripts² with *per-source and per-language token accounting*, distinguishing dataset-card estimates from tokenizer-exact consumed-token counts. We also provide the German : English : code mixing ratio, and the rationale behind it, in contrast to reports that disclose only an aggregate token count.
- **Reproducible recipe.** We publish the full learning-rate schedule (Warmup–Stable–Decay), optimizer, all hyperparameters, the per-phase token budgets, and the phase boundaries, so a

²<https://github.com/soofi-project/Soofi-Pretraining>

third party can rebuild the run.

- **Long-context serving efficiency.** The hybrid Mamba–MoE design keeps the per-sequence cache near-constant in context length, yielding measured aggregate decode TPS/GPU 8–9× that of dense 14–24B models at 40K context and batch 32, and aggregate decode TPS that stays essentially flat from 4K to 256K where full-attention models degrade (Figure 1, Section 4.3).
- **Documented design.** We report the data and design ablations behind each choice, exposing the *why* rather than only the final configuration.
- **Complete artifact release.** We will release selected Soofi S 30B-A3B Base checkpoints³, training⁴ and evaluation code^{5,6}, all under permissive licenses for transparent audit and extension.

The remainder of this report documents Soofi S 30B-A3B Base end to end. Section 2 presents the model and the training recipe: the Nemotron 3 Nano reference architecture we adopt and the rationale for adopting it, the Warmup–Stable–Decay optimization schedule, all hyperparameters and per-phase token budgets, the training dynamics of the run, and the compute infrastructure. Section 3 details the three-phase, 26.68T-token German–English data curriculum, with full per-source token accounting for every stage. Section 4 then evaluates the resulting base model against 16 open models of comparable or larger active size along two axes: capability across parallel English and German benchmarks, and serving efficiency. Section 5 situates the work relative to prior open and European efforts, and Section 6 concludes.

2 Model Architecture and Training

Soofi S 30B-A3B Base adopts the hybrid Mamba–Transformer Mixture-of-Experts (MoE) reference architecture of Nemotron 3 Nano [61, 60]: a 52-layer network interleaving 23 Mamba-2 sequence-mixing layers [16], 23 granular MoE layers with shared experts [77, 19, 15, 41], and 6 Grouped-Query Attention (GQA) layers [1] distributed sparsely through the network depth.

The model totals ~31.6B parameters, of which only ~3.2B are active per forward pass (~3.6B including embeddings), and only the 6 GQA layers maintain a KV cache. Because we reuse the reference design without modification, we refer the reader to the Nemotron reports [60, 61] for the design motivation and the ablations behind each architectural choice; Table 1 records the exact configuration for reproducibility.

Why a reference architecture. Reusing an established, openly specified architecture rather than designing a bespoke one was a deliberate decision, on three grounds. First, *deployability*: the Nemotron 3 Nano architecture is already integrated into the major open inference and serving stacks, including the vLLM stack used for our own serving measurements (Section 4.3), with mature, heavily optimized kernels for its Mamba-2, GQA, and MoE components, so Soofi S can be hosted efficiently

³<https://huggingface.co/Soofi-Project>

⁴<https://github.com/soofi-project/Soofi-Pretraining>

⁵<https://github.com/ellamind/base-eval>

⁶<https://github.com/ellamind/eval-hive>

Soofi S pre-training dynamics




Figure 2: **Training dynamics over the full ~ 27 T-token run.** The quantity is plotted against the number of consumed tokens (in trillions); the pretraining-to-annealing transition occurs at ~ 20 T tokens. The solid line is a rolling median over 1,000 steps and the faint trace is the raw per-step signal.

by existing software from the day of release, without bespoke integration work by downstream users. Second, *serving efficiency*: the predominantly Mamba-2 backbone makes the architecture exceptionally fast in exactly the regime we target, prefill and decode costs grow near-linearly in sequence length and the per-sequence cache stays near-constant, which is what underpins the long-context throughput results of Section 4.3 and makes the 1M-token context extension of Section 3.4 practical. Third, *scientific control*: sharing the backbone with Nemotron 3 Nano turns that model into an architecture-identical baseline, so the effect of our German–English data recipe can be measured in isolation (Section 4.2).

2.1 Optimization and Hyperparameters

We train Soofi S with the Megatron-Bridge framework⁷ using AdamW [49] under a Warmup–Stable–Decay (WSD) [35, 30] learning-rate schedule whose decay segment follows a `minus_sqrt` shape. With the exception of the long-context phase (Section 2.2), all stages share the same parallelism

⁷For this training run, we adopted Megatron-Bridge (<https://github.com/NVIDIA/Megatron-Bridge>) rather than integrating the Nemotron-3 reference architecture into our open-source stack [50], as Megatron-Bridge already provided mature, reliable support for the newly released architecture.

Table 1: Soofi S 30B-A3B architecture, following the Nemotron 3 Nano reference configuration [61]. The layer pattern interleaves Mamba-2 and MoE layers, with 6 GQA layers distributed through the network depth.

Hyperparameter	Value
Total parameters	31.6B
Active parameters per token	3.2B (3.6B incl. embeddings)
Num layers	52 (23 Mamba-2, 23 MoE, 6 GQA)
Model dimension	2688
Attention Q-heads	32
Attention KV-heads	2
Attention head dimension	128
Mamba-2 state dimension	128
Mamba-2 groups	8
Mamba-2 heads	64
Mamba-2 head dimension	64
Expert dimension	1856
Routed experts	128
Activated experts per token	6
Shared experts	2
MoE activation	squared ReLU
Router	learned MLP, sigmoid gating
Positional embeddings	none
Normalization	RMSNorm
Embedding / projection tying	untied

and batching configuration: tensor-model-parallel (TP) size 1, expert-parallel (EP) size 8, sequence parallelism disabled, micro-batch size 2, and global batch size 3072 at a sequence length of 8192 tokens, with the remaining GPUs used for data parallelism and a distributed (optimizer-state-sharded) optimizer. This corresponds to 25,165,824 tokens per optimizer step. All stages are trained in bf16 mixed precision. Since the granular MoE layers route each token to experts held on different expert-parallel ranks, all-to-all communication sits on the critical path, which motivates the node topology and interconnect reported in Section 2.4. Manual garbage collection is triggered every 101 iterations. The schedule is realised across one warmup-plus-stable phase and three successive annealing continuations, summarised in Table 2; the data mixture consumed by each stage is documented in Section 3:

- **Base pretraining (stable).** 794,728 iterations (19,999,984,975,872 tokens, $\sim 20T$) on the 20T EN/DE mixture (see Section 3.2). After a 254-iteration warmup to the peak learning rate of $1e-3$, training proceeds on the WSD stable plateau.
- **Main annealing (decay).** A continuation of 198,682 iterations (4,999,996,243,968 tokens, $\sim 5T$) on the 5T high-quality EN/DE mixture (see Section 3.3), applying the WSD `minus_sqrt` decay from $1e-3$ to $1e-5$. The configured total after this continuation is 993,410 iterations.
- **Constant annealing.** Since the slope at the end of the decay segment remained relatively steep, we append a further 62,590 iterations (1,575,128,924,160 tokens (see Section 3.3),

$\sim 1.58\text{T}$) at a *constant* learning rate of $1\text{e-}5$ on the same 5T mixture, allowing additional training at the tail of the annealing curve.

- **Final annealing (discarded).** A final WSD `minus_sqrt` decay of 11,920 iterations ($\sim 0.30\text{T}$, see Section 3.3) from $1\text{e-}5$ to 0, beginning at iteration reference 1,056,000 and targeting 1,067,920. We report this stage for completeness but *do not* use its checkpoints, as they showed no clear additional benchmark improvement over the constant annealing stage.

Table 2: Training stages and learning-rate schedule. “Iters” are the iterations added in each stage; all 8K-context stages use 25,165,824 tokens per iteration. The final annealing stage was run but its checkpoints were not used. The long-context stage uses a distinct configuration (Section 2.2).

Stage	Role	Iters	Tokens	LR (peak $\rightarrow$ min)
Base pretraining	Stable (warmup 254)	794,728	$\sim 20\text{T}$	$1\text{e-}3$ (plateau)
Main annealing	WSD decay (<code>minus_sqrt</code>)	198,682	$\sim 5\text{T}$	$1\text{e-}3 \rightarrow 1\text{e-}5$
Constant annealing	Constant LR	62,590	$\sim 1.58\text{T}$	$1\text{e-}5$
Final annealing [†]	WSD decay (<code>minus_sqrt</code>)	11,920	$\sim 0.30\text{T}$	$1\text{e-}5 \rightarrow 0$
Long context	Constant (warmup 100)	2,000	$\sim 0.10\text{T}$	$1\text{e-}5 \rightarrow 1\text{e-}7$

[†] Run for completeness; checkpoints not used (no clear benchmark gain).

2.2 Long-Context Extension

The long-context phase (Phase 3, Section 3.4) is trained with a distinct parallelism configuration to accommodate the 1M -token sequences. We use a context-parallel size of 16, a micro-batch size of 1, and a global batch size of 48 at a sequence length of 1,048,576 tokens, giving 50,331,648 tokens per optimizer step. The stage runs for 2,000 iterations, for an approximate total of 100.66B tokens. Optimization again uses AdamW. The learning rate follows a constant schedule with a 100-iteration warmup, holding at $1\text{e-}5$ (minimum $1\text{e-}7$). Two settings distinguish this phase from the earlier stages: HybridEP is enabled (it was not used previously), and we apply *no* intra-document masking, matching NVIDIA’s long-context recipe [61] rather than the alternative adopted by some other efforts.

2.3 Training Dynamics

Figure 2 summarizes the four central signals we logged over the full $\sim 27\text{T}$ -token run, each plotted against the number of consumed tokens. The pretraining-to-annealing transition at $\sim 20\text{T}$ tokens—the end of the stable plateau and the start of the main WSD decay (Table 2)—is the reference point for reading all four panels. For the loss, gradient-norm, and throughput traces, the solid line is a rolling median over 1,000 steps and the faint trace is the raw per-step signal; loss and gradient norm are clipped at 2.0 for readability.

Training logs. In addition to the static training-dynamics plot in Figure 2, we provide the corresponding Weights & Biases dashboard for the full pretraining run⁸. The dashboard contains the

⁸<https://api.wandb.ai/links/soofi-exchange/j11vi7rg>

raw and smoothed traces used to inspect optimization stability, throughput, checkpointing effects, and the transitions between the base-pretraining, annealing, and long-context phases.

The learning-rate schedule (Figure 2a) follows the Warmup–Stable–Decay (WSD) shape: a short linear warmup to the peak of 1×10^{-3} , a constant plateau held throughout base pretraining, and a `minus_sqrt` decay toward ~ 0 once annealing begins. The language-modeling loss (Figure 2b) declines steadily across the stable phase and then drops sharply at the onset of annealing due to the dataset switch, tracking the learning-rate decay as the model is concentrated on the high-quality mixture. The gradient norm (Figure 2c) remains stable across the entire run, with only a mild increase during the annealing phase and no divergence or sustained spikes. Training throughput (Figure 2d), measured in tokens per second, is steady for most of the run; the downward excursions correspond to checkpointing, evaluation, and restart iterations rather than to changes in the training dynamics themselves.

2.4 Compute Infrastructure

Soofi S was trained on the Industrial AI Cloud [17, 62] operated by Deutsche Telekom in Munich built together with NVIDIA and brought into operation in February 2026. Our run used up to 512 NVIDIA B200 GPUs, i.e. 64 DGX B200 nodes of 8 GPUs each. Within a node the 8 B200s are fully connected by fifth-generation NVLink/NVSwitch, while nodes are interconnected by an eight-rail NVIDIA Quantum-2 NDR InfiniBand fabric, one 400 Gb/s ConnectX-7 port per GPU, 3.2 Tb/s of scale-out bandwidth per node. Keeping expert-parallel groups within a node lets the MoE all-to-all (Section 2.1) run over intra-node NVLink rather than the slower inter-node fabric, which matters for MoE throughput at this scale. The run took place from 24 March 2026 to 13 May 2026 and consumed approximately 253,000 B200 GPU-hours across the stable, annealing, and long-context stages. Training on this infrastructure is part of the sovereignty of the model: it was executed on German soil under European operational and data-protection requirements rather than on extra-European hyperscale compute. Soofi S was one of the first flagship workloads on the Industrial AI Cloud, whose infrastructure was procured for the sovereign open-source foundation-model effort under which this model was developed [17]. The Munich facility is powered entirely by renewable energy, designed for high energy efficiency, cooled with water drawn from the nearby Eisbach canal, and integrated with a waste-heat-reuse concept that feeds the surrounding Tucherpark district.

3 Pretraining Data

Consistent with our commitment to full reproducibility, we document the pretraining corpus of Soofi S at the granularity of individual source datasets. For every constituent, we report its public identifier, its raw token count, the number of epochs it was repeated, the resulting effective token count, and its share of the phase. We deliberately also list sources that were enumerated but *excluded* from training (epoch count of zero), so that the mixture can be audited and rebuilt end to end. This stands in contrast to the common practice of disclosing only an aggregate token count for the training data, and follows the openness ethos of recent fully open efforts [5, 64].

Soofi S is trained on a three-phase curriculum, consistent with the Warmup–Stable–Decay (WSD) learning-rate schedule (Section 2.1). Phase 1 (see Section 3.2) maximizes diversity over a large, quality-tiered mixture of web, synthetic, code, mathematics, and multilingual data. Phase 2 (Section 3.3) is an annealing (decay) phase that concentrates the highest-quality web data together

Effective-token mixture across training phases

Band width = share of that phase’s tokens.




Figure 3: **Effective-token mixture across the three training phases.** A single flow diagram tracing seven data categories (English Web, Academic & Wiki, SFT, Reasoning, Code, Math, and German) from left to right across the phases. **Phase 1** (diverse pretraining), 23,051.13B effective tokens. **Phase 2** (high-quality annealing), 6,303.0B effective tokens, showing increased density of skill-oriented and German data relative to Phase 1. **Phase 3** (long-context extension), 188B effective tokens, where the SFT band branches into its General, Code, and Math SFT components.

with skill-focused code, mathematics, STEM, reasoning, and instruction data, while further up-weighting German to 15.3% of the constructed annealing pool. Phase 3 (Section 3.4) extends the usable context length up to 1M tokens via length-bucketed up-sampling. Table 3 summarizes the token budget of each phase. Across all phases, the corpus comprises approximately 27 trillion tokens, of which a deliberately elevated fraction is German, reflecting the design goal of a German–English model rather than a broadly multilingual one. The full per-source composition of every phase is documented in Appendix B; the figure in this section summarizes those tables as a mixture flow diagram.

3.1 Quality Tiers, Synthetic Data, and Epoching

Most of the web data is drawn from the openly released Nemotron-CC datasets [80, 61] (v1.0, v2.0, and v2.1), which provide documents pre-sorted into quality tiers (High, Medium-High, Medium) and include several synthetically rephrased variants (-SYNTHETIC), diverse question–answer reformulations (DIVERSE QA / DQA), and English translations of non-English documents (TRANSLATED-TO-ENGLISH). We exploit these tiers directly: the highest-quality and synthetic tiers are repeated for multiple epochs to increase their effective contribution, whereas the large Medium-Quality pools

Table 3: Three-phase pretraining curriculum and token budget. The *Pool* column is the effective-token mixture constructed for each phase (documented per source in Appendix B); for Phases 1–2 these are dataset-card counts and are approximate, whereas the Phase 3 pool is tokenized with our own tokenizer. The *Consumed* column is the number of tokens actually trained on, counted exactly from the optimizer schedule (iterations $\times$ tokens/iteration; Table 2). A phase may consume less than one epoch of its pool (Phase 1) or more than one (Phase 2); the headline $\sim 27\text{T}$ figure refers to consumed tokens.

Phase	Role	Pool (eff.)	Consumed
Phase 1	Diverse pretraining (stable)	$\sim 23.05\text{T}$	$\sim 20\text{T}$
Phase 2	High-quality annealing (decay)	$\sim 6.30\text{T}$	$\sim 6.58\text{T}$
Phase 3	Long-context extension (4K–1M)	$\sim 0.19\text{T}$	$\sim 0.10\text{T}$
Total		$\sim 29.5\text{T}$	$\sim 26.68\text{T}$

are listed but set to zero epochs and thus excluded from the final mixture. An epoch count greater than one, therefore, denotes deliberate up-sampling of a high-value source, and an epoch count of zero denotes a source we evaluated but chose not to train on. All effective-token figures and shares below are reported *after* applying these epoch multipliers.

Token accounting and its precision. For Phases 1 and 2, the raw and effective token counts below are taken from the token statistics reported on each source’s public (HuggingFace) dataset card. Since different datasets are tokenized with different tokenizers, these counts are not all expressed in our model’s tokens; they should be read as close approximations that document the *relative composition* of the mixture rather than an exact token ledger. By contrast, all per-iteration training-token figures (Section 2.1, Table 2) and the entire Phase 3 long-context pool (Section 3.4, Tables 10–11) were obtained by tokenizing the data with the Nemotron-3 tokenizer and are therefore exact. This distinction also explains why the constructed-pool totals reported below (e.g. $\sim 23.05\text{T}$ effective for Phase 1) differ from the exactly-counted tokens actually consumed during training (e.g. $\sim 20\text{T}$; Table 3): the former are dataset-card estimates of the *available* pool, the latter are exact counts of what the optimizer *saw*, and a phase may consume less than one epoch of its pool (Phase 1) or slightly more than one (Phase 2).

3.2 Phase 1: Diverse Pretraining

Phase 1 provides the bulk of the training signal. Its composition is given in full in Table 6, grouped by source family with per-family subtotals. The phase totals $\sim 16.35\text{T}$ raw tokens, which after epoching yield $\sim 23.05\text{T}$ effective tokens in dataset-card terms; the run consumes $\sim 20\text{T}$ of this pool (Section 2.1). The mixture is anchored on quality-filtered and synthetic web text from the three Nemotron-CC releases (CC-v2.1, CC-v2.0, and CC-v1.0 contribute $\sim 2.84\text{T}$, $\sim 5.60\text{T}$, and $\sim 3.15\text{T}$ effective tokens respectively), supplemented by a large code component ($\sim 3.38\text{T}$ across the Nemotron code datasets [61]), specialized STEM and scientific data (Nemotron-Pretraining-Specialized-v1 [61], $\sim 1.35\text{T}$), pretraining-stage SFT mixtures ($\sim 1.57\text{T}$ of Math, Code, and General SFT), and dedicated mathematics corpora ($\sim 1.06\text{T}$ from Nemotron-CC-Math v1/v2 [51] and a further 352B from UltraData-Math [91]). PDF-derived text from FinePDFs [44] and Dolma3_POOL [64] adds high-value long-form document data. Figure 3 (Left) summarises the Phase 1 mixture by source

family.

For the German–English objective, German is intentionally over-represented relative to the base Nemotron recipe: German sources contribute $\sim 1.65\text{T}$ effective tokens, or 7.2% of Phase 1, against the 5% multilingual share of the reference Nemotron 3 Nano mixture [61]. The German component combines naturally occurring web and document text (HPLT Monolingual Datasets 3.0 [63], German Commons [24], Genios [23], the German subset of FinePDFs [44] and FineWiki [67]) with machine-translated (MT) and synthetic German (MultiSynt/MT [37], MT-Reasoning [25, 58], MT of Nemotron-Multilingual-Reasoning [29], PleIAs/Synth [69]). Table 7 compares our full Phase 1 mixture against the published Nemotron 3 Nano mixture; relative to it we raise German and academic share, lean more heavily on high-quality web, and trim synthetic web and code-SFT.

3.3 Phase 2: High-Quality Annealing

The annealing phase coincides with the learning-rate decay and is composed exclusively of high-value data. The per-source breakdown is given in Table 8, organised by category with per-category subtotals; the constructed pool totals $\sim 6.30\text{T}$ effective tokens (dataset-card counts), of which the annealing schedule trains for $\sim 6.58\text{T}$. Relative to Phase 1, we drop the lower-tier web pools entirely, retain only the High-Quality and High-Quality-Synthetic web tiers, and substantially increase the density of skill-oriented data: code ($\sim 1.03\text{T}$), mathematics ($\sim 330\text{B}$), and a broad SFT mixture ($\sim 0.93\text{T}$) spanning math, code, agentic, competitive-programming, instruction-following, science, finance, software-engineering, safety, and multilingual subsets. A dedicated reasoning bucket—Nemotron-Pretraining-Specialized-v1 and related sources ($\sim 353\text{B}$)—strengthens chain-of-thought ability ahead of post-training [61]. English Web accounts for 36.4% of the annealing mixture; including Academic & Wiki, English web/document data accounts for 42.8%. Figure 3 (Mid) summarizes the Phase 2 mixture by category.

German is up-weighted again during the annealing phase. The German category alone contributes 965.37B effective tokens, drawn from a pre-release version of HPLT-4⁹, a German translation of ClimbMix [18] produced with the KletterMix pipeline [42], German FinePDFs-Edu [44] and FineWiki [67], and synthetic/translated German reasoning sources; this brings the multilingual share to 15.32%, more than triple the 5% of the reference mixture. Table 9 reports the category-level composition of the annealing phase against Nemotron 3 Nano.

As part of the SFT mixture, we additionally include QA-base ($\sim 0.05\%$ of the pool; the 1.43B English tokens are counted under the SFT category and the 1.87B German tokens under the German category in Table 8), paraphrased training splits of 25 standard NLP benchmarks in English and German, analogous to the paraphrase-augmented benchmark training data in Olmo 3’s mid-training mix (e.g. TinyMATH, Dolmino Flan) [64] and the benchmark-seeded synthetic data in Nemotron 3 [61].

3.4 Phase 3: Long-Context Extension

To extend the usable context window up to 1M tokens, we assemble a long-context data pool of approximately **188.5B tokens** drawn from ~ 21 million documents, partitioned into nine sequence-length buckets (4K, 8K, 16K, 32K, 64K, 128K, 256K, 512K, and 1M tokens). The schema targets

⁹<https://hplt-project.org/datasets/v4.0>

balanced exposure across context lengths by allocating comparable token mass to each bucket; in practice the realized mass per bucket varies with source availability, and several buckets at the extremes are sparsely populated or empty (Table 11). Seven domains contribute to the pool: general web (28.78B effective tokens), code (9.51B), mathematics (6.73B), German (6.68B), and three supervised-fine-tuning streams, general SFT (31.25B), code SFT (29.24B), and mathematics SFT (76.31B), which together account for roughly 73% of the pool. Table 10 lists the per-domain token budgets, document counts, and source priorities; Table 11 gives the per-bucket document counts. Unlike Phases 1–2, every Phase 3 token count was produced by tokenizing the data with the Nemotron-3 tokenizer, so the figures in this section and in Tables 10–11 are exact rather than dataset-card estimates.

The long-context stage trained for 2,000 optimizer steps at a 1M-token sequence length, i.e. approximately 100.66B tokens (Section 2.2)—about 53% of the pool, or roughly half an epoch. We stopped at this point because no further loss improvement was observed from additional long-context training, mirroring the rationale for the discarded final-annealing stage (Section 2.1).

When a bucket draws from several sources, we fill it according to a fixed priority order: for web, ClimbMix [18] is preferred over OlmoOCR [70], which is preferred over FinePDFs [44]; for code, SwallowCode [20] is preferred over Nemotron-Pretraining-Code-v1/v2 [61]; and German is composed of 40% HPLTv4 and 60% German translation of ClimbMix [18] with the pipeline of KletterMix [42].

We note the population irregularities in the interest of full disclosure. Our mathematics sources do not provide documents beyond the 64K bucket, so the longer mathematics buckets are effectively unpopulated (2 documents at 128K, none beyond). The web and German domains have no populated 256K bucket, and the code domain’s 4K and 8K buckets are effectively empty (whole-file repository code is scarce at these lengths after packing); these short sequence lengths are instead carried by the web, mathematics, and SFT streams, the latter being densest at 8K–64K. The SFT mathematics buckets nominally extend to 256K but are negligible above 64K (19 and 2 documents at 128K and 256K, respectively).

3.5 Data Provenance and Release

The overwhelming majority of our pretraining data is openly available. The web, code, mathematics, specialized, and SFT components build on NVIDIA’s publicly released Nemotron pre-training datasets [80, 51, 61] (Nemotron-CC v1.0/v2.0/v2.1, Nemotron-CC-Math, Nemotron-CC-Code, Nemotron-Pretraining-Code/Specialized/SFT, and the Nemotron post-training SFT collections), complemented by open corpora including the Dolma 3 pools [64], the Fine* family [68, 44] (FinePDFs, FinePDFs-edu, FineWiki), ClimbMix [18], SwallowCode [20], UltraData-Math [91], and AceReason [48]. German coverage is provided by open resources such as HPLT (v3 and v4) [11, 63], German-Commons [24], KletterMix [42], and the Fine* German subsets, augmented with machine-translated and synthetic German data (the MultiSynt/MT-* sources, Soofi-Think-SFT [28]), plus the commercially licensed Genios [23] corpus. Consistent with the openness goals of this work, we release the complete mixture specification; every source, its raw token count, epoch multiplier, and effective contribution for all three phases, so that the corpus can be independently reconstructed where source licenses permit; for commercially licensed Genios, we release aggregate statistics and exact mixture accounting rather than redistributing raw text.

Openness classification. The definition of open-source AI remains contested: the OSI’s Open Source AI Definition 1.0 [65] permits documented but unsharable training data, whereas stricter proposals for a European definition require every training token to be redistributable [46]. Soofi S satisfies OSAID 1.0: we will release weights, intermediate checkpoints, training and evaluation code, and exact per-source data accounting under permissive licenses. Under the stricter open-data standard, Soofi S falls short in exactly one documented component: the commercially licensed Genios [23] corpus (1.3% of Phase 1 effective tokens, reported in aggregate in Appendix D). Every other source is publicly obtainable, so $\sim 99\%$ of the mixture can be independently reconstructed. We state this boundary explicitly rather than claim a stronger openness status than the release supports.

3.5.1 Code Web Data

Our code-from-web component is Nemotron-CC-Code-v1 [61], a corpus of code and code-adjacent documents recovered directly from Common Crawl rather than from repository hosting platforms. Standard web-extraction pipelines tend to mangle source code embedded in HTML, collapsing the whitespace and indentation that is syntactically meaningful in languages such as Python, so this corpus is built with a code-aware extraction pipeline that preserves the formatting of code blocks, retains the surrounding natural-language context (tutorials, documentation, Q&A threads, blog posts), and applies dedicated cleaning, language identification, and deduplication stages. The result is code in its *instructional habitat*: implementations interleaved with the prose that explains them, which is a complementary signal to raw repository files. We use the ACTUAL subset at 3 epochs (427.9B raw, 1,283.7B effective tokens; Table 6), making it the single largest code source in Phase 1.

3.5.2 Curated Code Data

Repository-sourced code is drawn from Nemotron-Pretraining-Code-v1 [59] and -v2 [61], which curate permissively licensed source code from public repositories with quality filtering, per-language balancing, and aggressive deduplication. Both releases pair the curated ACTUAL code with synthetic derivatives generated from it, and we train on five such variants from v2: *code review* (critiques and improvement suggestions for real code), *question answering* (Q&A pairs grounded in repository code), *rewriting* (refactorings and alternative implementations of the same functionality), *student-teacher* (dialogues that explain code step-by-step), and *transpilation* (translations of programs between programming languages). These variants convert passive code exposure into bidirectional code-language supervision during pretraining itself. In Phase 1, the curated-code family contributes 2,091.14B effective tokens (the ACTUAL subsets at 3 epochs, synthetic subsets at 2), and the full family is retained at 1 epoch during annealing (Table 8). During annealing, we additionally include Swallow-Code-v2 [20] (stage 5 subset for 2 epochs) and the Dolma 3 Dolmino code pool [64] as high-quality curated complements.

3.5.3 German and English Web Data

General web text is the backbone of the corpus. The English side builds on the three Nemotron-CC releases [80, 61] (v1.0, v2.0, v2.1), which classify Common Crawl documents into quality tiers using ensembles of model-based quality classifiers, and augment the high-value tiers with synthetic

transformations: rephrased variants of high-quality pages (-SYNTHETIC), diverse question-answer reformulations (DIVERSE QA), and English translations of high-quality non-English crawl data (TRANSLATED-TO-ENGLISH). As described in Section 3.1, we up-sample the High-Quality and synthetic tiers and exclude the Medium tier entirely; the three releases together contribute 11,601.4B effective tokens to Phase 1. Long-form English document data comes from PDF-derived corpora (FinePDFs [44] and the Dolma 3 PDF pool [64]), which supply the book-, report-, and paper-style text that is underrepresented in HTML crawls.

The German web component is assembled from sources of complementary character. Naturally occurring German is provided by quality-filtered crawl data—the HPLT corpora [11, 63] restricted to the top decile of their educational-quality score HPLT-3-Top10% in Phase 1 (edu-scores based on JQL [2], HPLT-4-Top10% in Phase 2 (edu-scores based on Propella [36])—together with German-Commons [24] (openly licensed German text), Genios [23], a commercially licensed corpus of 916 German newspaper and trade-press archives comprising 193M articles (57.6B words, 2010–2025; see Appendix D), the German FinePDFs/FinePDFs-Edu subsets, German FineWiki, and the curated mixture used during annealing.

Since the supply of high-quality native German text is far smaller than for English, we extend it with translated and synthetic German: KletterMix [42] (machine translations of ClimbMix [18]), MultiSynt/MT [37] (machine translations of high-quality Nemotron-CC [80] English documents into German), MT-Reasoning [58], and Nemotron-Multilingual-Reasoning [29] (translated reasoning traces), and the German subset of PleIAs/Synth [69].

Phase 1 German up-sampling leans on the scarce high-quality crawl (the HPLT v3 top-decile pool runs 8.4 epochs), while annealing switches to the fresher HPLT v4 pool and a German translation of ClimbMix produced with the KletterMix pipeline [42]. In total, German receives 7.2% of Phase 1 and 15.3% of Phase 2 effective tokens, well above the $\sim 5\%$ *total* multilingual share of the reference recipe, concentrated in a single language.

3.5.4 Specialized Synthetic Data

Beyond web and code, we train on three families of specialized, largely synthetic data. First, nvidia/Nemotron-CC-Math [51] (v1 and v2) recovers mathematical content from Common Crawl with a rendering-based pipeline that preserves equations and converts them to a uniform LaTeX representation, sorted into quality bands; we train on all bands for 4 epochs in Phase 1 and keep the top v1 bands during annealing, complemented by UltraData-Math [91]. Second, Nemotron-Pretraining-Specialized-v1 [61] provides targeted synthetic corpora for STEM and reasoning, including synthetic math, Wikipedia-style rewrites of encyclopedic content, and competition-style problems, which we up-sample aggressively (5 epochs, 1,353.5B effective tokens, 5.9% of Phase 1) and retain at 1 epoch during annealing alongside the Dolma 3 THINKING pool [64]. Third, we include SFT-formatted data already at the pretraining stage: Nemotron-Pretraining-SFT-v1 (Math, Code, and General splits) in Phase 1, broadened during annealing by the Nemotron post-training collections¹⁰ (agentic, competitive programming, instruction following, math proofs, science, finance, software engineering, safety, and multilingual subsets; Table 8). Exposing the model to instruction- and reasoning-formatted text before post-training shortens the distribution shift at the SFT stage and measurably strengthens chain-of-thought behaviour of the base model.

¹⁰<https://huggingface.co/collections/nvidia/nemotron-post-training-v3>

3.6 Data Mixture and Ordering

We organise all sources into seven unified categories, English Web, Academic & Wiki, Code, Mathematics, SFT, Reasoning, and German, and steer the mixture at this category level (the per-source realisation is given in Appendix B). The appendix tables retain the finer Nemotron-style accounting categories; Figure 3 folds them into the seven unified categories used in the main text. In Phase 2, the main-text Reasoning category corresponds to the appendix row “Reasoning / STEM-SFT”, while the narrower longtable “Reasoning subtotal” reports only the non-SFT reasoning sources. The guiding principle of the ordering is a quality- and skill-based curriculum aligned with the WSD learning-rate schedule: breadth while the learning rate is high, concentration while it decays.

During the stable phase (Phase 1), the mixture is dominated by diverse web text (50.3% English Web plus 8.0% academic and wiki text), with code at 14.6%, reasoning at 10.4%, German at 7.2%, mathematics at 6.0%, and SFT at 3.5% (Figure 3). Up-sampling via epoch multipliers, rather than the inclusion of lower-quality pools, is the primary lever for hitting these targets: high-quality and synthetic tiers run 2–6 epochs, the medium tiers run zero. At the onset of learning-rate decay (Phase 2), the mixture shifts decisively toward skill density and German depth: English Web drops to 36.4%, Academic & Wiki accounts for 6.4%, SFT contributes 8.5%, Reasoning accounts for 11.8%, Code accounts for 16.4%, Mathematics accounts for 5.2%, and German rises to 15.3% (Figure 3).

This places the highest-value tokens in the regime where the decaying learning rate consolidates them most effectively. Sources seen for multiple epochs in Phase 1 are re-weighted downward in Phase 2 (typically to a single epoch, or fresh replacements are substituted, e.g. HPLT v3 $\rightarrow$ v4) so that the annealing phase adds new signal rather than repeating saturated data.

Phase 3 builds on top of the highest-quality documents from Phase 2 and spans seven document-level domains (web, code, math, German, Math SFT, General-SFT, and Code-SFT). It is organized by *length*, using the length-bucketed scheme of Section 3.4.

Relative to the reference Nemotron 3 Nano mixture, our category targets differ in one deliberate respect: German replaces the broad multilingual bucket and is raised from 5% (Phase 1 & Phase 2) to 7.2% (our Phase 1) and 15.32% (our Phase 2), funded by reductions in synthetic web and code-SFT share (Tables 7 and 9). Within each phase, category proportions are held stationary: data of all categories is shuffled and interleaved uniformly at the batch level, so the curriculum acts between phases, not within them.

4 Evaluations

We evaluate Soofi S 30B-A3B Base against 16 open base models with the same harness, prompts, and few-shot configuration; the complete per-task results for all models are released alongside this report. Figure 1 includes the full set of models in the throughput-capability comparison, while Tables 4 and 5 report detailed results for the main comparison subsets.¹¹ In the main text, we report detailed results for the strongest baselines in two comparison settings. The first compares against large open-source base models (Alia 40B [26], EuroLLM 22B [53, 52], Apertus 70B [5], and Olmo 3 32B [64]); the second compares against its architectural reference, Nemotron 3 Nano 30B-A3B [61], and large open-weight models of comparable or larger active size (Qwen3.5 35B-A3B [88, 71], Ministral 3 14B [47],

¹¹Baselines: Nemotron 3 Nano 30B-A3B; Qwen3.5 35B-A3B/9B; Gemma 3 12B/27B; Ministral 3 3B/8B/14B; Olmo 3 32B; Alia 40B; Apertus 8B/70B; EuroLLM 9B/22B; Teuken 7B; Salamandra 7B.

and Gemma 3 27B [83]). All models are evaluated with the same `lm-evaluation-harness` [22] pipeline, using the same task configurations (prompts, number of few-shot samples, etc), on English and German benchmark suites closely following the evaluation setup of Olmo 3 [64] covering code, mathematics, knowledge, reasoning, science, reading comprehension, and German language proficiency. This includes a held-out set of benchmarks that were not used during any ablation experiments. Soofi S is evaluated at the selected base checkpoint `iter_1056000`, the final checkpoint of the constant-annealing stage (Section 2.1, Appendix F); for other checkpointed models we report the highest available training step. The complete per-task results files are released alongside the model; headline numbers are collected in Tables 4 and 5 and Figures 4 and 9.

4.1 Soofi S vs Open-Source Models

The open-source comparison reports results against Alia 40B, EuroLLM 22B, Apertus 70B, and Olmo 3 32B. We separate these models from the larger open-weight baselines in Section 4.2 so that each table compares models with similar release categories. Within this set, Soofi S is the strongest model overall: it obtains the highest English aggregate (+2.8 over Olmo 3 32B), German aggregate (+6.3 over Apertus 70B), and the highest held-out English (+8.3) and German (+5.6) scores (Figure 4).




Figure 4: Evaluation overview for the open-source comparison. Soofi S is compared against large open-source models (Alia, EuroLLM, Apertus, and Olmo 3). Aggregates are the harness-level English and German suite means. Code EN averages HumanEval and MBPP, Code DE averages HumanEval-DE and MBPP-DE, and LBPP is reported separately.

Table 4 gives the per-task results. The largest margins over the next-best open-source baseline

occur on German and technical benchmarks: German aggregate (+6.3), held-out English (+8.3), held-out German (+5.6), HumanEval (+10.8), MBPP-DE (+13.4), GSM8K-Platinum-DE (+9.7), INCLUDE-DE (+10.1), and GLP-DE (+7.6). Olmo 3 32B is the strongest non-Soofi model overall and leads this subset on LBPP, SocialIQA, SQuAD, DROP, and NaturalQuestions. These results identify the main residual gaps for Soofi S in this comparison as contamination-aware code evaluation, open-domain factual recall, and extractive reading comprehension.

Code performance. Soofi S ranks first on four of the five code benchmarks in the open-source comparison (Figure 5). It exceeds the next-best open-source baseline by 10.8 points on HumanEval, 7.4 points on MBPP, 3.0 points on HumanEval-DE, and 13.4 points on MBPP-DE. LBPP is the only code benchmark in this subset where Soofi S is not first; Olmo 3 32B scores 32.1 compared with 31.0 for Soofi S.




Figure 5: Code generation results (pass@1) against large open-source models on English and German benchmarks. Soofi S leads this comparison on HumanEval, MBPP, HumanEval-DE, and MBPP-DE; Olmo 3 32B is strongest on LBPP.

Mathematics, knowledge, and reasoning. Soofi S also ranks first within the open-source subset on all mathematics benchmarks shown in Table 4 and Figure 6. The largest mathematics margins are on GSM8K-Platinum-DE (+9.7 over Olmo 3 32B), Minerva-500 (+24.2 over Olmo 3 32B), Minerva Math-EN (+27.0 over Olmo 3 32B), and Minerva MATH-DE (+7.5 over Olmo 3 32B). In

Table 4: Base model evaluation results (%) against large *open-source* models. Best result per row in **bold**, second best underlined. All models evaluated with identical harness, prompts, and few-shot settings; “-DE” denotes the German variant of a benchmark. Aggregates are harness-level suite means; Column shading: Soofi S (ours), open-source (Alia, EuroLLM, Apertus, Olmo 3).

Benchmark	Soofi S 30B-A3B	Alia 40B	EuroLLM 22B	Apertus 70B	Olmo 3 32B
Aggregates					
English aggregate	70.1	59.0	61.2	62.4	<u>67.3</u>
German aggregate	79.1	68.4	70.6	<u>72.8</u>	69.2
Held-out (EN)	41.4	28.0	30.8	27.6	<u>33.1</u>
Held-out (DE)	41.8	29.4	33.9	33.5	<u>36.2</u>
Code (pass@1)					
HumanEval	73.8	23.8	39.3	30.2	<u>63.0</u>
MBPP	70.2	35.5	47.6	41.5	<u>62.8</u>
LBPP	<u>31.0</u>	8.6	10.7	6.4	32.1
HumanEval-DE	65.5	29.9	34.7	29.9	<u>62.5</u>
MBPP-DE	84.2	45.6	59.4	50.9	<u>70.8</u>
Mathematics					
GSM8K	86.1	65.4	25.1	65.4	<u>80.7</u>
GSM8K-Platinum-DE	87.1	65.0	62.6	62.6	<u>77.4</u>
Minerva 500	79.4	14.3 [†]	28.7	33.7	<u>55.2</u>
Minerva Math-EN	81.0	14.4 [†]	26.4	34.5	<u>54.0</u>
Minerva MATH-DE	56.0	12.9 [†]	28.4	29.0	<u>48.5</u>
Knowledge					
MMLU stem	75.9	54.0	58.2	58.4	<u>70.9</u>
MMLU-Pro	51.4	33.0	34.3	37.5	<u>47.4</u>
MMLU-Pro-DE	49.4	29.4	33.4	34.0	<u>38.0</u>
INCLUDE-DE	61.2	43.9	<u>51.1</u>	50.4	48.2
NaturalQuestions (acc)	<u>79.0</u>	69.0	70.1	73.0	79.1
Commonsense and reading comprehension					
PIQA	85.7	83.7	82.4	<u>85.3</u>	83.5
PIQA-DE	91.5	85.7	80.5	<u>86.4</u>	67.4
SocialIQA	<u>60.5</u>	55.2	57.4	59.3	62.4
SocialIQA-DE	87.6	<u>86.6</u>	85.0	87.6	75.9
SQuAD (EM)	<u>87.5</u>	78.7	81.2	83.2	88.1
DROP (EM)	<u>66.5</u>	41.5	43.0	50.0	72.2
Reasoning and science					
BBH (CoT)	78.8	53.9	57.8	60.5	<u>77.4</u>
AGIEval	66.9	48.1	52.2	55.7	<u>63.2</u>
GPQA-Diamond	43.4	29.8	30.3	27.3	<u>33.3</u>
GPQA-Diamond-DE	41.9	25.7	30.2	26.3	<u>34.1</u>
ARC-Challenge	90.6	79.8	81.1	85.8	<u>89.9</u>
German language proficiency					
GLP-DE	88.8	65.4	78.2	<u>81.2</u>	73.0
ARC-Challenge-DE	92.3	81.5	83.6	<u>85.3</u>	83.6

knowledge, reasoning, and science (Figure 7), Soofi S ranks first on MMLU-STEM, MMLU-Pro, MMLU-Pro-DE, INCLUDE-DE, BBH, AGIEval, GPQA-Diamond, GPQA-Diamond-DE, and ARC-Challenge. The largest margins in this group include GPQA-Diamond (+10.1 over Olmo 3 32B), INCLUDE-DE (+10.1 over EuroLLM 22B), and GPQA-Diamond-DE (+7.8 over Olmo 3 32B).




Figure 6: Mathematics results against large open-source models on English and German benchmarks. Soofi S leads the open-source comparison on GSM8K, GSM8K-Platinum-DE, Math-EN, and Math-DE.

German capabilities. Figure 8 isolates the German benchmarks for the open-source comparison. Soofi S ranks first on every German row in this subset, with margins of +6.3 on the German aggregate, +7.6 on GLP-DE, +7.0 on ARC-Challenge-DE, +10.1 on INCLUDE-DE, +9.7 on GSM8K-Platinum-DE, +13.4 on MBPP-DE, and +5.6 on the held-out German suite relative to the strongest open-source baseline for each metric.




Figure 7: Knowledge (left) and reasoning/science (right) benchmarks against large open-source models. Soofi S leads the open-source comparison on MMLU-STEM, MMLU-Pro, INCLUDE-DE, BBH, AGIEval, GPQA-Diamond, GPQA-Diamond-DE, and ARC-Challenge.




Figure 8: German benchmark results against large open-source models. Soofi S ranks first on the German aggregate, GLP-DE, ARC-Challenge-DE, INCLUDE-DE, GSM8K-Platinum-DE, MBPP-DE, and the held-out German suite in this comparison.

4.2 Soofi S vs Open-Weight Models

The open-weight comparison contains two types of baselines: the architecture-identical Nemotron 3 Nano reference, and larger public-weight models from Qwen, Ministral, and Gemma. This split separates the effect of the German–English data recipe from comparisons against larger models with different architectures and training corpora. In this subset, Soofi S is not the top model on aggregate scores, but it is competitive with larger dense baselines and improves over Nemotron 3 Nano on all aggregate and held-out metrics.




Figure 9: Base model evaluation overview for Soofi S against Nemotron (same architecture) and large open-weight models (Qwen, Ministral, and Gemma). Aggregates are the harness-level English and German suite means. Code EN averages HumanEval and MBPP, Code DE averages HumanEval-DE and MBPP-DE, and LBPP is reported separately.

At the aggregate level, Qwen3.5 35B-A3B has the highest English, German, and held-out means in this subset (Table 5). Soofi S scores 70.1 on the English aggregate, compared with 70.3 for Gemma 3 27B and Ministral 3 14B. On the German aggregate, Soofi S scores 79.1, compared with 78.3 for Ministral 3 14B and 78.4 for Gemma 3 27B. Relative to Nemotron 3 Nano, Soofi S improves the English aggregate by +1.8, the German aggregate by +4.2, held-out English by +6.7, and held-out German by +2.0.

Code performance. Figure 10 reports the code-generation results for the open-weight comparison. Soofi S sets the best score in this subset on HumanEval [12] (73.8), MBPP [6] (70.2), and MBPP-DE (84.2), and is second only to its architectural reference on HumanEval-DE (65.5 vs. 68.8). The German code benchmarks (problem statements and docstrings in German) show the

Table 5: Base model evaluation results (%) against Nemotron 3 Nano and large *open-weight* models. Best result per row in **bold**, second best underlined. All models evaluated with identical harness, prompts, and few-shot settings; “-DE” denotes the German variant of a benchmark. Aggregates are harness-level suite means; Nemotron 3 Nano shares the same 30B-A3B Mixture-of-Experts architecture as Soofi S. Column shading: Soofi S (ours), Nemotron (same architecture), open-weight (Qwen, Ministral, Gemma).

Benchmark	Soofi S 30B-A3B	Nemotron 3 Nano 30B-A3B	Qwen3.5 35B-A3B	Ministral 3 14B	Gemma 3 27B
Aggregates					
English aggregate	70.1	68.3	74.6	<u>70.3</u>	<u>70.3</u>
German aggregate	<u>79.1</u>	74.9	81.6	78.3	78.4
Held-out (EN)	<u>41.4</u>	34.7	47.8	40.0	35.5
Held-out (DE)	41.8	39.8	46.0	<u>42.2</u>	38.3
Code (pass@1)					
HumanEval	73.8	<u>72.2</u>	67.1	60.9	60.2
MBPP	70.2	<u>67.5</u>	65.8	58.8	66.7
LBPP	31.0	38.1	<u>32.4</u>	27.0	22.4
HumanEval-DE	<u>65.5</u>	68.8	59.5	55.5	57.7
MBPP-DE	84.2	<u>79.9</u>	79.0	72.0	75.6
Mathematics					
GSM8K	<u>86.1</u>	86.5	82.6	83.7	79.9
GSM8K-Platinum-DE	87.1	<u>87.7</u>	91.2	81.3	80.7
Minerva 500	<u>79.4</u>	64.0	83.0	61.8	75.7
Minerva Math-EN	<u>81.0</u>	64.2	82.4	61.3	73.7
Minerva MATH-DE	56.0	58.1	76.5	55.1	<u>65.6</u>
Knowledge					
MMLU stem	75.9	74.8	81.9	<u>76.4</u>	74.7
MMLU-Pro	51.4	51.6	60.3	<u>53.3</u>	50.7
MMLU-Pro-DE	49.4	47.1	56.6	<u>50.2</u>	47.2
INCLUDE-DE	61.2	59.7	61.2	<u>60.4</u>	57.6
NaturalQuestions (acc)	79.0	80.3	<u>81.0</u>	77.1	83.5
Commonsense and reading comprehension					
PIQA	85.7	84.8	<u>84.9</u>	82.9	<u>84.9</u>
PIQA-DE	91.5	85.7	<u>87.6</u>	81.3	86.7
SocialIQA	60.5	57.6	<u>58.0</u>	53.9	56.2
SocialIQA-DE	87.6	82.7	<u>87.2</u>	83.0	86.7
SQuAD (EM)	87.5	<u>87.3</u>	80.4	82.7	85.1
DROP (EM)	66.5	64.2	62.6	62.9	<u>65.0</u>
Reasoning and science					
BBH (CoT)	78.8	77.8	84.2	<u>79.6</u>	77.8
AGIEval	66.9	65.1	71.7	67.0	<u>68.3</u>
GPQA-Diamond	<u>43.4</u>	33.8	50.0	40.9	35.9
GPQA-Diamond-DE	41.9	37.4	46.4	<u>42.5</u>	31.8
ARC-Challenge	90.6	89.4	93.1	<u>91.6</u>	<u>91.6</u>
German language proficiency					
GLP-DE	88.8	73.7	94.0	<u>89.5</u>	88.3
ARC-Challenge-DE	92.3	91.4	96.6	<u>93.5</u>	93.3

largest positive margins for Soofi S in this subset: on MBPP-DE, Soofi S leads the best dense open-weight model by 8.6 points. The contamination-aware LBPP benchmark [54] (31.0) is the one code benchmark where Soofi S trails, behind Nemotron (38.1) and Qwen3.5 35B-A3B (32.4), while staying ahead of Ministral and Gemma.




Figure 10: Code generation results (pass@1) against Nemotron 3 Nano and large open-weight models on English and German benchmarks. Soofi S achieves the best HumanEval, MBPP, and MBPP-DE scores in this comparison. The code aggregates used in Figure 9 average HumanEval with MBPP and HumanEval-DE with MBPP-DE; LBPP is reported separately.

Mathematics. On grade-school reasoning, Soofi S scores 86.1 on GSM8K [14] and 87.1 on GSM8K-Platinum-DE [86]—within 0.4 and 4.1 points of the best open-weight result on the respective tasks (Figure 11). On competition-style mathematics [32, 45], Soofi S reaches 79.4 on Minerva-500 and 81.0 on Minerva Math-EN—the second-best results in this comparison, behind only Qwen3.5 35B-A3B—lifting its Math-EN score to 82.8, essentially tied with Qwen3.5 35B-A3B under this suite definition. German competition mathematics is weaker: on Minerva MATH-DE (56.0) Soofi S trails Qwen3.5 35B-A3B (76.5), Gemma 3 27B (65.6), and its architectural reference (58.1), leaving Math-DE at 71.5.

Knowledge. On broad academic knowledge, Soofi S is competitive with dense models several times its active size: MMLU-STEM [31] 75.9 and MMLU-Pro [87] 51.4 sit between Gemma 3 27B




Figure 11: Mathematics results against Nemotron 3 Nano and large open-weight models on English and German benchmarks. Soofi S is second-best on GSM8K [14]; on the German side it is essentially level with its architectural reference on GSM8K-Platinum-DE, while competition-style Math-DE trails Qwen3.5 35B-A3B.

and Ministral 3 14B, and the German MMLU-Pro-DE (49.4) shows the same pattern (Figure 12, left). On INCLUDE-DE [76]—regional, Germany-specific knowledge spanning driving-licence, social-science, and STEM exams—Soofi S ties Qwen3.5 35B-A3B for the best score in the table (61.2), consistent with the up-weighted native German data described in Sections 3.2 and 3.3. Open-domain factual recall as measured by NaturalQuestions [43] (79.0) narrowly trails the largest dense models, consistent with storing world knowledge in 3B active parameters; we return to this in the Limitations paragraph.

Reasoning and science. Soofi S achieves strong general-reasoning performance, with scores of 78.8 on BBH [81], 66.9 on AGIEval [90], and 90.6 on ARC-Challenge [13]. These results place it consistently within a point of Ministral 3 14B (Figure 12, right). Its graduate-level science performance is particularly strong: GPQA-Diamond [75] increases to 43.4, a +9.6-point improvement over the reference Nemotron recipe, trailing Qwen3.5 35B-A3B among the open-weight models shown here. On the German GPQA-Diamond-DE benchmark, Soofi S reaches 41.9, a +4.5-point improvement over its reference and within 0.6 points of the best remaining dense model (Ministral 3 14B, 42.5).

We attribute these gains primarily to the reasoning-oriented annealing phase and the high density

of STEM-SFT data described in Section 3.5.4.




Figure 12: Knowledge (left) and reasoning/science (right) benchmarks against Nemotron 3 Nano and large open-weight models. Soofi S ties Qwen3.5 35B-A3B on INCLUDE-DE and improves GPQA-Diamond by +9.6 points over its architectural reference, while matching dense 14–27B models on MMLU-STEM, MMLU-Pro, BBH, and AGIEval.

German capabilities. Figure 13 reports the German benchmarks for the open-weight comparison. Soofi S ranks first on MBPP-DE (84.2) and ties Qwen3.5 35B-A3B on INCLUDE-DE (61.2). It is close to the larger dense models on the German aggregate (79.1), behind Qwen3.5 35B-A3B (81.6) and ahead of Gemma 3 27B (78.4), Ministral 3 14B (78.3), and Nemotron 3 Nano (74.9). Relative to Nemotron, Soofi S improves GLP-DE by +15.1, INCLUDE-DE by +1.5, and the held-out German suite by +2.0, while remaining slightly below Nemotron on GSM8K-Platinum-DE and HumanEval-DE.

Effect of the German–English recipe. Since Soofi S shares its architecture with Nemotron 3 Nano 30B-A3B, the pairwise comparison in Figure 14 isolates the effect of our data recipe from architecture. The German interventions deliver large, targeted gains—GLP-DE +15.1, German aggregate +4.2, INCLUDE-DE +1.5, German code +0.5—without the regression on English that monolingual specialisation usually incurs: the English aggregate *improves* by +1.8, code by +2.2, CommonsenseQA [82] by +5.1, and GPQA-Diamond by +9.6. The held-out benchmark suite scores improve by +6.7 (EN) and +2.0 (DE). The only notable regressions are German competition mathematics (Math-DE aggregate −1.4), open-domain factual recall (NaturalQuestions −1.3), and a −0.4 on GSM8K, consistent with German tokens displacing a portion of English web knowledge. Overall, the Nemotron comparison indicates that the German-focused data recipe substantially




Figure 13: German benchmark results against Nemotron 3 Nano and large open-weight models. Soofi S ranks first on MBPP-DE, ties Qwen3.5 35B-A3B on INCLUDE-DE, and improves the German aggregate over the architecture-identical Nemotron baseline by +4.2 points.

improves German capability while preserving or improving the English aggregate, code, reasoning, and held-out scores reported here.

Limitations. We report three caveats in the interest of full transparency. First, the Minerva mathematics [32, 45] scores required a corrected evaluation protocol. Our initial harness configuration used a `\n\n` stop sequence and a generation limit of 1,024 tokens, which truncated long chain-of-thought solutions before a final answer was produced and collapsed the Minerva scores of the strongest models (Soofi S measured 9.8 on Minerva-500 under this configuration, and the architecture-identical Nemotron 3 Nano 5.6). Removing the stop condition and raising the generation limit to 4,096 tokens yields the results reported in Tables 4 and 5 and the figures (79.4 and 81.0 for Soofi S on Minerva-500 and Minerva Math-EN); all models were re-evaluated under the same corrected configuration. Second, competition-style mathematics in German remains the clearest capability gap to the frontier: on Minerva MATH-DE (56.0) Soofi S trails Qwen3.5 35B-A3B (76.5) and Gemma 3 27B (65.6), even though its German grade-school mathematics (GSM8K-Platinum-DE, 87.1) is essentially level with the architectural reference. Third, open-domain factual recall remains capacity-limited: NaturalQuestions (79.0) trails the largest dense baselines (Gemma 3 27B 83.5), consistent with storing world knowledge in 3B active parameters; we expect retrieval-augmented deployment to close this gap in practice.




Figure 14: Per-benchmark score difference between Soofi S and the architecture-identical Nemotron 3 Nano 30B-A3B, isolating the effect of the German–English data recipe. German-focused benchmarks improve by up to 15 points while English capability is preserved or improved; the main cost is open-domain English factual recall.

4.3 Serving efficiency.

Active parameter count is only a proxy for inference cost. In deployment, the relevant quantity is sustained throughput at realistic context lengths and request concurrency, where decoding is often limited by memory bandwidth: each generated token must stream the active model weights and read the attention cache for every active sequence. This is the regime targeted by the hybrid Mamba–MoE architecture. In Soofi S, only 6 of 52 layers maintain a KV cache, with 2 KV heads each, while the 23 Mamba-2 layers carry a fixed-size recurrent state. Consequently, the incremental attention-cache footprint is only about 6 KB per token per sequence, which is 11–53× lower than for the dense models in our comparison. As context length grows, only this small attention component scales with sequence length; the Mamba recurrent state remains constant-size.

Figure 1 quantifies this effect using measured aggregate decode tokens per second (TPS) per GPU. We use a TP=1, single-B200 vLLM latency-subtraction protocol. For each context length, we measure fixed-batch latency as a function of output length O . Let $t(O)$ denote this latency. We estimate aggregate decode TPS per GPU by subtracting the $O_{\text{short}} = 1$ run from the $O_{\text{long}} = 1024$ run:

$$\text{TPS}_{\text{decode}}^{\text{agg/GPU}} = \frac{B(O_{\text{long}} - O_{\text{short}})}{N_{\text{GPU}}[t(O_{\text{long}}) - t(O_{\text{short}})]} \quad (1)$$

Here $O_{\text{short}} = 1$, $O_{\text{long}} = 1024$, $N_{\text{GPU}} = 1$, and $B = 32$ for all plotted measurements. This subtraction removes most prompt-prefill cost and isolates the per-GPU cache-bandwidth pressure that tensor parallelism can otherwise mask. At a 40K context and batch size 32, Soofi S sustains a measured aggregate decode rate of 4.82k TPS/GPU, which is 9.2× higher than Ministral 3 14B, while fitting the weights and all 32 sequence states on a single GPU. In contrast, Apertus 70B does

not fit on a single GPU under the current fixed-256K, TP=1 setup.

The same measurements also expose the prefill/TTFT side of the hybrid architecture: $t(O_{\text{short}}) = t(1)$ is the measured batch-32 latency to process the prompt and produce the first output token, i.e., the TTFT-like measurement in this protocol. At 40K context, Soofi S reaches $t(1) = 22.7$ s, compared with 71.9 s for Ministral 3 14B, 92.9 s for Gemma 3 27B, 101.7 s for OLMo 3.1 32B, and 213.4 s for the Qwen3 32B dense control. At 256K context, Soofi S remains the fastest complete sweep in our batch-32 protocol (372.7 s), ahead of Qwen3.5 35B-A3B (579.2 s), Gemma 3 27B (999.4 s), OLMo 3.1 32B (1,487.3 s), Ministral 3 14B (2,058.9 s), and Qwen3 32B dense (6,428.6 s). This is the prefill analogue of the decode-cache advantage: most sequence mixing in Soofi S is carried by Mamba layers rather than full attention, so long-context TTFT grows much more favorably than in full-attention dense baselines.

Panel (b) shows that the decode gap widens with context length: dense-model throughput decreases as KV-cache reads dominate decoding, whereas Soofi S remains nearly flat from 4K to 256K in the current measurement snapshot, with endpoint aggregate decode rate changing from 4.29k to 4.30k TPS/GPU and no point more than $\sim 34\%$ below the 4K value. Among the comparison models, only Qwen3.5, a Gated-DeltaNet hybrid [71], exhibits similar scaling behavior. In its published 9B configuration, 8 of 32 layers remain full-attention layers, with 4 KV heads of dimension 256, corresponding to a per-sequence cache of 32 KB per token, or $5.3\times$ that of Soofi S. For the 35B-A3B variant, we measure 2.60k aggregate decode TPS/GPU at 40K under the same protocol, which is $1.9\times$ lower than Soofi S. This qualitative separation is consistent with NVIDIA’s Nemotron-H measurements, where a related hybrid engine achieved $3.3\times$ higher throughput than Qwen3-30B-A3B on production inference engines [61].

5 Related Work

We situate Soofi S with respect to five strands of prior work: open language-model pretraining, data acquisition and filtering, training curricula, efficient sparse and hybrid architectures, and multilingual European foundation models. We focus on base-model pretraining; instruction tuning, preference optimization, and reasoning-specific reinforcement learning are complementary and outside the scope of this report.

Open language-model pretraining. The modern causal language-modeling paradigm was established by GPT-style pretraining: GPT-2 showed that a decoder-only Transformer trained with next-token prediction on large web text could perform a wide range of tasks in a zero-shot setting, and GPT-3 demonstrated the emergence of strong in-context few-shot behavior at a substantially larger scale [72, 10]. Subsequent scaling-law work made model size, data size, and compute budget explicit design variables, while the Chinchilla analysis shifted attention from parameter count alone toward data–compute balance and the importance of training sufficiently long on enough tokens [40, 33]. These results motivate the central design principle of recent pretraining recipes: capability is a joint function of architecture, token budget, data quality, curriculum, and serving cost, not merely of total parameters.

Open releases have followed a separate but related trajectory. OPT released a family of decoder-only models together with code and a detailed account of training infrastructure, GPT-NeoX-20B provided a permissively released large autoregressive model trained on the Pile, and Pythia

made training dynamics inspectable by releasing many intermediate checkpoints and a fixed data order [89, 9, 7, 21]. BLOOM extended this community-scale approach to multilingual pretraining, combining a large open-access model with the ROOTS multilingual corpus and a collaborative development process [8]. These efforts established the scientific value of releasing more than a final checkpoint: intermediate states, training code, data documentation, and evaluation recipes make it possible to study memorization, bias, scaling behavior, and data effects.

The broader open-weight ecosystem has since produced increasingly strong general-purpose baselines, including LLaMA and Llama 2, Falcon, Mistral 7B, Mixtral, Qwen, Gemma, and recent Mistral/Ministral releases [84, 85, 4, 38, 39, 88, 71, 83, 47, 55]. However, many such models are open-weight rather than fully reproducible: the weights and inference code are available, but the exact data mixture, filtering pipeline, training order, logs, intermediate checkpoints, and rejected data sources are often missing. OLMo, OLMoE, OLMo 3, and Apertus mark a stronger notion of openness by releasing training data or data recipes, code, model artifacts, and broader documentation for scientific audit and reuse [27, 57, 64, 5]. Soofi S follows this fully open line, but differs in scope and architecture: it is a German–English base model trained with exact per-source token accounting on a sparse hybrid Mamba–Transformer MoE backbone, rather than a dense general-purpose or broadly multilingual Transformer.

Pretraining data acquisition and filtering. Pretraining data pipelines have evolved from relatively simple web-crawl selection toward multi-stage corpus construction. Early influential corpora such as WebText, C4, and the Pile used web-scale collection, heuristic filtering, deduplication, and source balancing to turn noisy internet data into training data suitable for language modeling [72, 73, 21]. BLOOM’s ROOTS corpus added a multilingual community-curated data effort, combining web data with manually selected sources across many languages [8]. More recent corpora make the filtering pipeline itself a central contribution: FineWeb and Dolma document large-scale web cleaning, deduplication, and mixture construction; Nemotron-CC provides quality-tiered Common Crawl data with synthetic and translated variants; and JQL shows that language-model-based quality judgments can be distilled into multilingual data filters that transfer across languages more robustly than English-centric heuristics [68, 79, 80, 2].

Domain-specific acquisition has become equally important. PDF- and OCR-derived corpora such as FinePDFs and OlmoOCR add long-form reports, books, papers, and educational material that are underrepresented in HTML-only crawls [44, 70]. Code corpora such as Nemotron-Pretraining-Code-v1/v2 and SwallowCode complement web-extracted code with repository-derived or rewritten programming data [61, 59, 20]. Mathematical data pipelines such as Nemotron-CC-Math and UltraData-Math target equation-rich documents and problem-solving text, where generic web extraction often destroys structure [51, 91]. For German, large multilingual crawls and national resources such as HPLT, German Commons, and KletterMix are especially relevant because high-quality native German tokens are scarcer than English web tokens [11, 63, 24, 42]. Soofi S builds on these trends but reports the corpus at finer granularity: for every phase, we disclose source identifiers, raw tokens, epoch multipliers, effective token counts, and sources considered but excluded. This makes the data mixture auditable in a way that aggregate token counts cannot provide.

Training curricula and optimization recipes. Large-scale pretraining recipes increasingly separate the problem of collecting tokens from the problem of ordering and weighting them. Compute-optimal scaling results motivate training smaller or sparse models for more tokens when inference

efficiency matters, while modern open reports show that curriculum phase boundaries, learning-rate schedules, and annealing mixtures can have large effects on downstream behavior [33, 35, 30, 61]. A common pattern is to allocate early training to broad coverage and later training to higher-quality or skill-focused data, often under a Warmup–Stable–Decay schedule. OLMo-style releases emphasize logging, intermediate checkpoints, and reproducible recipes; Nemotron-style recipes emphasize quality-tiered web data, synthetic skill data, and high-quality annealing [27, 64, 61]. Soofi S follows the same high-level philosophy but adapts it to a German–English target: Phase 1 maximizes diversity, Phase 2 concentrates high-quality web, code, mathematics, reasoning, SFT-formatted, and German data, and Phase 3 extends context length with document-length buckets. The difference is the explicit bilingual reallocation of token budget and the disclosure of the exact realized mixtures.

Efficient sparse and hybrid architectures. Most open LLMs remain dense Transformers, but dense attention becomes costly at long context and high concurrency because decoding must repeatedly read both model weights and the per-sequence KV cache. Sparse Mixture-of-Experts models address the weight side of this problem by increasing total capacity while activating only a small expert subset per token. Foundational MoE work, Switch Transformers, DeepSeekMoE, and fine-grained MoE studies show that sparse routing can improve the capability–compute trade-off when routing and load balancing are stable [77, 19, 15, 41]. OLMoE demonstrates that this sparse route can also be made fully open at smaller active-parameter scales [57].

A complementary line of work reduces the sequence-state cost of attention. Mamba-2 and related state-space models replace much of the quadratic attention machinery with recurrent state updates, yielding linear-time sequence mixing and a near-constant state during decoding [16]. Nemotron-H and Nemotron 3 combine Mamba-style sequence mixing, sparse attention, and MoE layers to obtain strong long-context serving efficiency, while Qwen3.5 explores a different hybrid sequence-modeling path with Gated DeltaNet layers [60, 61, 71]. Soofi S adopts the Nemotron-style 30B-A3B hybrid Mamba–Transformer MoE design, but evaluates it under a distinct sovereign bilingual pretraining recipe. This makes the architectural comparison unusually clean: relative to Nemotron 3 Nano, gains and trade-offs can largely be attributed to data mixture, German up-weighting, annealing, and long-context continuation rather than to a different backbone.

Multilingual and European language models. Multilingual pretraining aims to reduce the English bias of large language models, but it requires careful allocation of model capacity and high-quality tokens across languages. BLOOM was an early large-scale open-access demonstration of multilingual decoder-only pretraining [8]. More recent European efforts add requirements around sovereignty, transparency, language coverage, and regulatory compatibility. Teuken-7B targets the official EU languages with a European multilingual tokenizer and a large non-English data share [3]. EuroLLM develops a family of European multilingual models across several scales, with multilingual data filtering, tokenizer design, and evaluation as central components [53, 52, 74]. Salamandra and the subsequent ALIA family focus on European and Iberian language modeling [26]. Apertus emphasizes fully open and compliant multilingual foundation models, and OpenEuroLLM extends this direction as a coordinated European initiative [5, 66].

Soofi S is complementary to these broad-coverage efforts. Rather than optimizing for many languages at once, it studies a narrower German–English setting in which the data mixture, annealing phase, and evaluation suite are designed around bilingual depth. Broad European models address coverage across languages, while Soofi S tests how much capability and efficiency can be gained

when one bilingual deployment setting is given a dedicated, fully documented pretraining recipe.

Positioning of Soofi S. The closest architectural reference for Soofi S is Nemotron 3 Nano, because both use a 30B-A3B hybrid Mamba–Transformer MoE architecture [61]. The closest openness references are OLMo, OLMoE, OLMo 3, and Apertus, because they emphasize releases that enable audit and reconstruction rather than merely inference [27, 57, 64, 5]. The closest European language references are EuroLLM, Teuken, Salamandra, Apertus, and OpenEuroLLM [53, 52, 74, 3, 26, 5, 66]. Soofi S combines these lines in a configuration not covered by prior work: a fully documented European pretraining run, a German–English data curriculum with exact per-source accounting, and a sparse hybrid architecture designed for long-context, high-concurrency serving. The resulting model fills a gap between broadly multilingual European sovereignty efforts and efficient international open-weight baselines: it asks whether a sovereign model can be simultaneously open, bilingual-depth-oriented, and competitive in capability per active parameter.

6 Conclusion

We presented Soofi S 30B-A3B, a sovereign, open-source MoE hybrid Mamba–Transformer foundation model for German and English. Built on the Nemotron 3 Nano architecture—52 layers combining Mamba-2, Grouped-Query Attention, and granular MoE layers that activate roughly 3B of $\sim$ 30B parameters per token—Soofi S was pretrained on approximately 27 trillion tokens under a three-phase Warmup–Stable–Decay curriculum: 20T tokens of diverse, quality-tiered pretraining, $\sim$ 7T tokens of high-quality annealing, and a length-bucketed long-context phase extending the usable context to 1M tokens. Throughout, German was deliberately up-weighted—to 7.2% of the stable phase and 15.32% of the annealing mixture, more than triple the multilingual share of the reference recipe—realizing the design goal of a German–English champion rather than a thinly spread multilingual model.

The result is a model that reaches the capability frontier at a fraction of the inference cost of dense alternatives. On a unified evaluation of 17 open base models (Section 4), Soofi S achieves the best English and German code aggregates among the measured models (HumanEval/MBPP averages, with LBPP reported separately), a Math-EN score essentially tied with the strongest model, second-best scores on GSM8K [14], the Minerva mathematics benchmarks, and INCLUDE-DE (tied-best within the reported open-weight comparison), and English and German aggregates that match dense 14–27B models—all while activating only 3B parameters per token. It outperforms every European sovereign baseline in our comparison—including those an order of magnitude larger in active parameters—matching or outperforming them on every German benchmark in the suite, often by 10–30 points (Figure 1, Figure 8, and Figure 13). To our knowledge, this makes Soofi S the first European sovereign model to sit on the same capability-per-active-parameter frontier as the strongest international open-weight releases, the strongest open German base model in its inference-cost class, and the strongest fully open base model in our evaluation on both English and German aggregates.

Equally central to this work is *how* the model is released. Trained end-to-end on the German Industrial AI Cloud by a consortium of German research institutions, Soofi S ships not only weights but the complete set of artifacts needed to audit and rebuild it: the full per-source token accounting of all three pretraining phases (including sources we evaluated and excluded), every hyperparameter and learning-rate stage—including a discarded final annealing stage, reported for completeness—

and the training and evaluation code under permissive licenses, with licensed data sources documented through aggregate statistics and exact mixture accounting rather than redistributed raw text. We hope this level of transparency moves the open ecosystem from open-weight toward genuinely open-source, and provides a reproducible template for other language communities seeking capable, efficient, sovereign foundation models.

Future work will extend Soofi S along three axes: open post-training (SFT and large-scale RL) toward instruct and reasoning variants (including modular reasoning in the spirit of FlexOlmo and Bar [78, 56]), broader and deeper German evaluation suites, and continued scaling of the high-quality German data pipeline that this release identified as the principal bottleneck for further gains.

Acknowledgments

This work was supported by the German Federal Ministry for Economic Affairs and Energy (BMWE) in the context of IPCEI-CIS and 8ra through “Soofi: Souveräne KI für Europa” (grant number 13IPC040A-J). Parts of it have benefited from the hessian.AI Service Center (funded by the Federal Ministry of Research, Technology and Space, BMFTR, grant no. 16IS22091) and the hessian.AI Innovation Lab (funded by the Hessian Ministry for Digital Strategy and Innovation, grant no. SDIW04/0013/003). We are also grateful to all the many people who have supported and enabled this project, including the Telekom Industrial AI Cloud and NVIDIA teams. In particular, we would like to thank Pramod Kumbhar and Oleg Sudakov, whose in-depth expertise and tremendous dedication have been invaluable to our work. We further thank Miroslav Shaltev and Oleh Astappiev from the L3S Research Center for operating the CPU cluster and its Slurm scheduling. We would also like to thank Christian Kotulek, Marek Soha and all their colleagues for their hard work in ensuring that the cluster runs round the clock at full capacity. Finally, we thank Lara Lawniczak, Nora Malke and Synje Jungbehr for their project coordination, organizing project meetings, and keeping the project running smoothly.

References

- [1] J. Ainslie, J. Lee-Thorp, M. De Jong, Y. Zemlyanskiy, F. Lebrón, and S. Sanghai. 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, 2023.
- [2] M. Ali, M. Brack, M. Lübbering, E. Wendt, A. G. Khan, R. Rutmann, A. Jude, M. Kraus, A. A. Weber, F. Stollenwerk, D. Kaczér, F. Mai, L. Flek, R. Sifa, N. Flores-Herr, J. Koehler, P. Schramowski, M. Fromm, and K. Kersting. Judging quality across languages: A multilingual approach to pretraining data filtering with language models. In C. Christodoulopoulos, T. Chakraborty, C. Rose, and V. Peng, editors, *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 8859–8898, Suzhou, China, Nov. 2025. Association for Computational Linguistics.
- [3] M. Ali, M. Fromm, K. Thellmann, J. Ebert, A. A. Weber, R. Rutmann, C. Jain, M. Lübbering, D. Steinigen, J. Leveling, K. Klug, J. S. Buschhoff, L. Jurkschat, H. Abdelwahab, B. J. Stein, K.-H. Sylla, P. Denisov, N. Brandizzi, Q. Saleem, A. Bhowmick, L. Helmer, C. John, P. O. Suarez, M. Ostendorff, A. Jude, L. Manjunath, S. Weinbach, C. Penke, O. Filatov, F. Barth, P. Mirza, L. Weber, I. Wendler, R. Sifa, F. Küch, A. Hertten, R. Jäkel, G. Rehm, S. Kesselheim,

- J. Köhler, and N. Flores-Herr. Teuken-7b-base & teuken-7b-instruct: Towards european llms. *arXiv preprint arXiv:2410.03730*, 2025.
- [4] E. Almazrouei, H. Alobeidli, A. Alshamsi, A. Cappelli, R. Cojocaru, M. Debbah, É. Goffinet, D. Hesslow, J. Launay, Q. Malartic, D. Mazzotta, B. Noune, B. Pannier, and G. Penedo. The falcon series of open language models. *arXiv preprint arXiv:2311.16867*, 2023.
- [5] P. Apertus, A. Hernández-Cano, A. Hägele, A. H. Huang, A. Romanou, A.-J. Solergibert, B. Pasztor, B. Messmer, D. Garbaya, E. F. Durech, I. Hakimi, J. G. Giraldo, M. Ismayilzada, N. Foroutan, S. Moalla, T. Chen, V. Sabolčec, Y. Xu, M. Aerni, B. AlKhamissi, I. A. Mariñas, M. H. Amani, M. Ansaripour, I. Badanin, H. Benoit, E. Boros, N. Browning, F. Bösch, M. Böther, N. Canova, C. Challier, C. Charmillot, J. Coles, J. Deriu, A. Devos, L. Drescher, D. Dzenhaliov, M. Ehrmann, D. Fan, S. Fan, S. Gao, M. Gila, M. Grandury, D. Hashemi, A. Hoyle, J. Jiang, M. Klein, A. Kucharavy, A. Kucherenko, F. Lübeck, R. Machacek, T. Manitaras, A. Marfurt, K. Matoba, S. Matrenok, H. Mendonça, F. R. Mohamed, S. Montariol, L. Mouchel, S. Najem-Meyer, J. Ni, G. Oliva, M. Pagliardini, E. Palme, A. Panferov, L. Paoletti, M. Passerini, I. Pavlov, A. Poiroux, K. Ponkshe, N. Ranchin, J. Rando, M. Sausser, J. Saydaliev, M. A. Sayfiddinov, M. Schneider, S. Schuppli, M. Scialanga, A. Semenov, K. Shridhar, R. Singhal, A. Sotnikova, A. Sternfeld, A. K. Tarun, P. Teiletche, J. Vamvas, X. Yao, H. Zhao, A. Ilic, A. Klimovic, A. Krause, C. Gulcehre, D. Rosenthal, E. Ash, F. Tramèr, J. VandeVondele, L. Veraldi, M. Rajman, T. Schulthess, T. Hoeffler, A. Bosselut, M. Jaggi, and I. Schlag. Apertus: Democratizing open and compliant llms for global language environments. *arXiv preprint arXiv:2509.14233*, 2025.
- [6] J. Austin, A. Odena, M. Nye, M. Bosma, H. Michalewski, D. Dohan, E. Jiang, C. Cai, M. Terry, Q. Le, et al. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*, 2021.
- [7] S. Biderman, H. Schoelkopf, Q. G. Anthony, H. Bradley, K. O’Brien, E. Hallahan, M. A. Khan, S. Purohit, U. S. Prashanth, E. Raff, A. Skowron, L. Sutawika, and O. Van Der Wal. Pythia: A suite for analyzing large language models across training and scaling. In A. Krause, E. Brunskill, K. Cho, B. Engelhardt, S. Sabato, and J. Scarlett, editors, *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 2397–2430. PMLR, 23–29 Jul 2023.
- [8] BigScience Workshop. BLOOM: A 176b-parameter open-access multilingual language model. *arXiv preprint arXiv:2211.05100*, 2022.
- [9] S. Black, S. Biderman, E. Hallahan, Q. Anthony, L. Gao, L. Golding, H. He, C. Leahy, K. McDonnell, J. Phang, M. Pieler, U. S. Prashanth, S. Purohit, L. Reynolds, J. Tow, B. Wang, and S. Weinbach. GPT-NeoX-20B: An open-source autoregressive language model. In A. Fan, S. Ilic, T. Wolf, and M. Gallé, editors, *Proceedings of BigScience Episode #5 – Workshop on Challenges & Perspectives in Creating Large Language Models*, pages 95–136, virtual+Dublin, May 2022. Association for Computational Linguistics.
- [10] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901, 2020.

- [11] L. Burchell, O. de Gibert, N. Arefyev, M. Aulamo, M. Bañón, P. Chen, M. Fedorova, L. Guillou, B. Haddow, J. Hajič, J. Helcl, E. Henriksson, M. Klimaszewski, V. Komulainen, A. Kutuzov, J. Kytöniemi, V. Laippala, P. Mæhlum, B. Malik, F. Mehryary, V. Mikhailov, N. Moghe, A. Myntti, D. O’Brien, S. Oepen, P. Pal, J. Piha, S. Pyysalo, G. Ramírez-Sánchez, D. Samuel, P. Stepachev, J. Tiedemann, D. Variš, T. Vojtěchová, and J. Zaragoza-Bernabeu. An expanded massive multilingual dataset for high-performance language technologies (HPLT). In W. Che, J. Nabende, E. Shutova, and M. T. Pilehvar, editors, *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 17452–17485, Vienna, Austria, July 2025. Association for Computational Linguistics.
- [12] M. Chen, J. Tworek, H. Jun, Q. Yuan, H. P. D. O. Pinto, J. Kaplan, H. Edwards, Y. Burda, N. Joseph, G. Brockman, et al. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.
- [13] P. Clark, I. Cowhey, O. Etzioni, T. Khot, A. Sabharwal, C. Schoenick, and O. Tafjord. Think you have solved question answering? try arc, the ai2 reasoning challenge, 2018.
- [14] K. Cobbe, V. Kosaraju, M. Bavarian, M. Chen, H. Jun, L. Kaiser, M. Plappert, J. Tworek, J. Hilton, R. Nakano, et al. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.
- [15] D. Dai, C. Deng, C. Zhao, R. Xu, H. Gao, D. Chen, J. Li, W. Zeng, X. Yu, Y. Wu, et al. Deepseekmoe: Towards ultimate expert specialization in mixture-of-experts language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1280–1297, 2024.
- [16] T. Dao and A. Gu. Transformers are ssms: Generalized models and efficient algorithms through structured state space duality. *arXiv preprint arXiv:2405.21060*, 2024.
- [17] Deutsche Telekom AG. Germany’s first AI factory for industry officially goes into operation in Munich. <https://www.telekom.com/en/media/media-information/archive/germany-s-first-ai-factory-for-industry-1101670>, Feb. 2026. Industrial AI Cloud, Munich; accessed 2026-07-08.
- [18] S. Diao, Y. Yang, Y. Fu, X. Dong, D. Su, M. Kliegl, Z. Chen, P. Belcak, Y. Suhara, H. Yin, et al. Nemotron-climb: Clustering-based iterative data mixture bootstrapping for language model pre-training. *Advances in Neural Information Processing Systems*, 38, 2026.
- [19] W. Fedus, B. Zoph, and N. Shazeer. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *Journal of Machine Learning Research*, 23(120):1–39, 2022.
- [20] K. Fujii, Y. Tajima, S. Mizuki, M. Kawamura, H. Shimada, T. Shiotani, K. Saito, M. Oi, T. Nakamura, T. Okamoto, S. Ishida, K. Hattori, Y. Ma, H. Takamura, R. Yokota, J. Sakuma, and N. Okazaki. Rewriting pre-training data boosts llm performance in math and code. *arXiv preprint arXiv:2505.02881*, 2026.
- [21] L. Gao, S. Biderman, S. Black, L. Golding, T. Hoppe, C. Foster, J. Phang, H. He, A. Thite, N. Nabeshima, S. Presser, and C. Leahy. The pile: An 800gb dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*, 2020.
- [22] L. Gao, J. Tow, B. Abbasi, S. Biderman, S. Black, A. DiPofi, C. Foster, L. Golding, J. Hsu, A. Le Noac’h, H. Li, K. McDonell, N. Muennighoff, C. Ociepa, J. Phang, L. Reynolds,

- H. Schoelkopf, A. Skowron, L. Sutawika, E. Tang, A. Thite, B. Wang, K. Wang, and A. Zou. The language model evaluation harness, 07 2024.
- [23] GBI-Genios Deutsche Wirtschaftsdatenbank GmbH. GENIOS – german business and press database. <https://www.genios.de/browse/Alle>, 2025. Commercially licensed corpus of German newspaper and trade-press archives; obtained under a license that does not permit redistribution.
 - [24] L. Gienapp, C. Schröder, S. Schweter, C. Akiki, F. Schlatt, A. Zimmermann, P. Genêt, and M. Potthast. The german commons - 154 billion tokens of openly licensed text for german language models. *arXiv preprint arXiv:2510.13996*, 2025.
 - [25] GlaiveAI. Reasoning-v1-20m, 2025. A synthetic reasoning dataset containing 22mil+ general reasoning questions and responses generated using deepseek-ai/DeepSeek-R1-Distill-Llama-70B.
 - [26] A. Gonzalez-Agirre, M. Pàmies, J. Llop, I. Baucells, S. D. Dalt, D. Tamayo, J. J. Saiz, F. Espuña, J. Prats, J. Aula-Blasco, M. Mina, I. Pikabea, A. Rubio, A. Shvets, A. Sallés, I. Lacunza, J. Palomar, J. Falcão, L. Tormo, L. Vasquez-Reina, M. Marimon, O. Pareras, V. Ruiz-Fernández, and M. Villegas. Salamandra technical report. *arXiv preprint arXiv:2502.08489*, 2025.
 - [27] D. Groeneveld, I. Beltagy, E. Walsh, A. Bhagia, R. Kinney, O. Tafjord, A. Jha, H. Ivison, I. Magnusson, Y. Wang, S. Arora, D. Atkinson, R. Authur, K. Chandu, A. Cohan, J. Dumas, Y. Elazar, Y. Gu, J. Hessel, T. Khot, W. Merrill, J. Morrison, N. Muennighoff, A. Naik, C. Nam, M. Peters, V. Pyatkin, A. Ravichander, D. Schwenk, S. Shah, W. Smith, E. Strubell, N. Subramani, M. Wortsman, P. Dasigi, N. Lambert, K. Richardson, L. Zettlemoyer, J. Dodge, K. Lo, L. Soldaini, N. Smith, and H. Hajishirzi. OLMo: Accelerating the science of language models. In L.-W. Ku, A. Martins, and V. Srikumar, editors, *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15789–15809, Bangkok, Thailand, Aug. 2024. Association for Computational Linguistics.
 - [28] D. Gurgurov and T. Röhr. ReasonXL: A multilingual cross-domain reasoning corpus. <https://huggingface.co/datasets/toroe/Soofi-Think-SFT-10B-multilingual>, 2026.
 - [29] D. Gurgurov, T. Röhr, and S. Ostermann. Nemotron-multilingual-reasoning: A multilingual science reasoning dataset. <https://huggingface.co/datasets/DGurgurov/Nemotron-Multilingual-Reasoning>, 2025. Derived from nvidia/Llama-Nemotron-Post-Training-Dataset with machine-translated reasoning traces.
 - [30] A. Hägele, E. Bakouch, A. Kosson, L. B. Allal, L. Von Werra, and M. Jaggi. Scaling laws and compute-optimal training beyond fixed training durations. *Advances in Neural Information Processing Systems*, 37:76232–76264, 2024.
 - [31] D. Hendrycks, C. Burns, S. Basart, A. Zou, M. Mazeika, D. Song, and J. Steinhardt. Measuring massive multitask language understanding, 2021.
 - [32] D. Hendrycks, C. Burns, S. Kadavath, A. Arora, S. Basart, E. Tang, D. Song, and J. Steinhardt. Measuring mathematical problem solving with the math dataset. In J. Vanschoren and S. Yeung, editors, *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, volume 1, 2021.

- [33] J. Hoffmann, S. Borgeaud, A. Mensch, E. Buchatskaya, T. Cai, E. Rutherford, D. de Las Casas, L. A. Hendricks, J. Welbl, A. Clark, T. Hennigan, E. Noland, K. Millican, G. van den Driessche, B. Damoc, A. Guy, S. Osindero, K. Simonyan, E. Elsen, O. Vinyals, J. Rae, and L. Sifre. An empirical analysis of compute-optimal large language model training. In *Advances in Neural Information Processing Systems*, volume 35, pages 30016–30030. Curran Associates, Inc., 2022.
- [34] C.-P. Hsieh, S. Sun, S. Krizan, S. Acharya, D. Rekesh, F. Jia, Y. Zhang, and B. Ginsburg. Ruler: What’s the real context size of your long-context language models?, 2024.
- [35] S. Hu, Y. Tu, X. Han, C. He, G. Cui, X. Long, Z. Zheng, Y. Fang, Y. Huang, W. Zhao, et al. Minicpm: Unveiling the potential of small language models with scalable training strategies. *arXiv preprint arXiv:2404.06395*, 2024.
- [36] M. Idahl, B. Droste, B. Plüster, and J. P. Harries. propella-1: Multi-property document annotation for llm data curation at scale, 2026.
- [37] M. Idahl, J. Tiedemann, S. Pyysalo, D. Salinas, T. Galica, S. Qian, T. N. Mateiu, Z. Li, A. Lokrantz, F. Vitiugin, A. F. T. Martins, J. Kanerva, F. Ginter, M. Lindemann, T. Isbister, B. Moell, J. Lindh, J. Hajič, J. Jitsev, A. Kutuzov, S. Oepen, and G. Ramírez-Sánchez. Multisynth/mt: Trillion-token multi-parallel pre-training data translated across 36 languages, 2026.
- [38] A. Q. Jiang, A. Sablayrolles, A. Mensch, C. Bamford, D. S. Chaplot, D. de las Casas, F. Bressand, G. Lengyel, G. Lample, L. Saulnier, L. R. Lavaud, M.-A. Lachaux, P. Stock, T. Le Scao, T. Lavril, T. Wang, T. Lacroix, and W. El Sayed. Mistral 7b. *arXiv preprint arXiv:2310.06825*, 2023.
- [39] A. Q. Jiang, A. Sablayrolles, A. Roux, A. Mensch, B. Savary, C. Bamford, D. S. Chaplot, D. de las Casas, E. Bou Hanna, F. Bressand, G. Lengyel, G. Bour, G. Lample, L. R. Lavaud, L. Saulnier, M.-A. Lachaux, P. Stock, S. Subramanian, S. Yang, S. Antoniak, T. Le Scao, T. Gervet, T. Lavril, T. Wang, T. Lacroix, and W. El Sayed. Mixtral of experts. *arXiv preprint arXiv:2401.04088*, 2024.
- [40] J. Kaplan, S. McCandlish, T. Henighan, T. B. Brown, B. Chess, R. Child, S. Gray, A. Radford, J. Wu, and D. Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.
- [41] J. Krajewski, J. Ludziejewski, K. Adamczewski, M. Pióro, M. Krutul, S. Antoniak, K. Ciebiera, K. Król, T. Odrzygóźdź, P. Sankowski, et al. Scaling laws for fine-grained mixture of experts. *arXiv preprint arXiv:2402.07871*, 2024.
- [42] M. Kraus, R. Härle, S. Sztwiertnia, A. G. Khan, M. Ali, M. Fromm, and K. Kersting. Klettermix: Climbing toward high-quality german pretraining data. *arXiv preprint arXiv:2606.03773*, 2026.
- [43] T. Kwiatkowski, J. Palomaki, O. Redfield, M. Collins, A. Parikh, C. Alberti, D. Epstein, I. Polosukhin, J. Devlin, K. Lee, K. Toutanova, L. Jones, M. Kelcey, M.-W. Chang, A. M. Dai, J. Uszkoreit, Q. Le, and S. Petrov. Natural questions: A benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:452–466, 2019.
- [44] H. Kydliček, G. Penedo, and L. von Werra. Finepdfs. <https://huggingface.co/datasets/HuggingFaceFW/finepdfs>, 2025.

- [45] A. Lewkowycz, A. Andreassen, D. Dohan, E. Dyer, H. Michalewski, V. Ramasesh, A. Slone, C. Anil, I. Schlag, T. Gutman-Solo, Y. Wu, B. Neyshabur, G. Gur-Ari, and V. Misra. Solving quantitative reasoning problems with language models. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 3843–3857. Curran Associates, Inc., 2022.
- [46] A. Liesenfeld, M. Dingemanse, D. Blankvoort, N. Kalra, and A. R. Golkhandan. European open source ai definitions. <https://osai-index.eu/osai-definitions>, 2026. Accessed: 2026-07-06.
- [47] A. H. Liu, K. Khandelwal, S. Subramanian, V. Jouault, A. Rastogi, A. Sadé, A. Jeffares, A. Jiang, A. Cahill, A. Gavaudan, A. Sablayrolles, A. Héliou, A. You, A. Ehrenberg, A. Lo, A. Eliseev, A. Calvi, A. Sooriyarachchi, B. Bout, B. Rozière, B. D. Monicault, C. Lanfranchi, C. Barreau, C. Courtot, D. Grattarola, D. Dabert, D. de las Casas, E. Chane-Sane, F. Ahmed, G. Berrada, G. Ecrepont, G. Guinet, G. Novikov, G. Kunsch, G. Lample, G. Martin, G. Gupta, J. Ludziejewski, J. Rute, J. Studnia, J. Amar, J. Delas, J. S. Roberts, K. Yadav, K. Chandu, K. Jain, L. Aitchison, L. Fainsin, L. Blier, L. Zhao, L. Martin, L. Saulnier, L. Gao, M. Buyl, M. Jennings, M. Pellat, M. Prins, M. Poirée, M. Guillaumin, M. Dinot, M. Futral, M. Darrin, M. Augustin, M. Chiquier, M. Schimpf, N. Grinsztajn, N. Gupta, N. Raghuraman, O. Bousquet, O. Duchenne, P. Wang, P. von Platen, P. Jacob, P. Wambergue, P. Kurylowicz, P. R. Muddireddy, P. Chagniot, P. Stock, P. Agrawal, Q. Torroba, R. Sauvestre, R. Soletskyi, R. Mennner, S. Vaze, S. Barry, S. Gandhi, S. Waghjale, S. Gandhi, S. Ghosh, S. Mishra, S. Aithal, S. Antoniak, T. L. Scao, T. Cachet, T. S. Sorg, T. Lavril, T. N. Saada, T. Chabal, T. Foubert, T. Robert, T. Wang, T. Lawson, T. Bewley, T. Edwards, U. Jamil, U. Tomasini, V. Nemychnikova, V. Phung, V. Maladière, V. Richard, W. Bouaziz, W.-D. Li, W. Marshall, X. Li, X. Yang, Y. E. Ouahidi, Y. Wang, Y. Tang, and Z. Ramzi. Ministral 3. *arXiv preprint arXiv:2601.08584*, 2026.
- [48] Z. Liu, Z. Yang, Y. Chen, C. Lee, M. Shoenybi, B. Catanzaro, and W. Ping. Acereason-nemotron 1.1: Advancing math and code reasoning through sft and rl synergy. *arXiv preprint arXiv:2506.13284*, 2025.
- [49] I. Loshchilov and F. Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- [50] M. Lübbering, T. Ruland, R. Rutmann, F. Stollenwerk, D. Fitzek, M. Fromm, A. Weber, R. Sifa, N. Flores-Herr, J. Köhler, et al. Modalities, a pytorch-native framework for large-scale llm training and research. *arXiv preprint arXiv:2602.08387*, 2026.
- [51] R. K. Mahabadi, S. Satheesh, S. Prabhume, M. Patwary, M. Shoenybi, and B. Catanzaro. Nemotron-cc-math: A 133 billion-token-scale high quality math pretraining dataset. *arXiv preprint arXiv:2508.15096*, 2025.
- [52] P. H. Martins, J. Alves, P. Fernandes, N. M. Guerreiro, R. Rei, A. Farajian, M. Klimaszewski, D. M. Alves, J. Pombal, N. Boizard, M. Faysse, P. Colombo, F. Yvon, B. Haddow, J. G. C. de Souza, A. Birch, and A. F. T. Martins. Eurollm-9b: Technical report. *arXiv preprint arXiv:2506.04079*, 2025.
- [53] P. H. Martins, P. Fernandes, J. Alves, N. M. Guerreiro, R. Rei, D. M. Alves, J. Pombal, A. Farajian, M. Faysse, M. Klimaszewski, P. Colombo, B. Haddow, J. G. C. de Souza, A. Birch,

- and A. F. T. Martins. Eurollm: Multilingual language models for europe. *arXiv preprint arXiv:2409.16235*, 2024.
- [54] A. Matton, T. Sherborne, D. Aumiller, E. Tommasone, M. Alizadeh, J. He, R. Ma, M. Voisin, E. Gilsonan-McMahon, and M. Gallé. On leakage of code generation evaluation datasets. In Y. Al-Onaizan, M. Bansal, and Y.-N. Chen, editors, *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 13215–13223, Miami, Florida, USA, Nov. 2024. Association for Computational Linguistics.
- [55] Mistral AI. Mistral small 3.1. <https://mistral.ai/news/mistral-small-3-1/>, 2025. Model release, March 2025.
- [56] J. Morrison, S. Adhikesaven, A. Bhagia, M. Zaharia, N. A. Smith, and S. Min. Train separately, merge together: Modular post-training with mixture-of-experts, 2026.
- [57] N. Muennighoff, L. Soldaini, D. Groeneveld, K. Lo, J. Morrison, S. Min, W. Shi, E. P. Walsh, O. Tafjord, N. Lambert, Y. Gu, S. Arora, A. Bhagia, D. Schwenk, D. Wadden, A. Wettig, B. Hui, T. Dettmers, D. Kiela, A. Farhadi, N. A. Smith, P. W. Koh, A. Singh, and H. Hajishirzi. OLMoE: Open mixture-of-experts language models. In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24–28, 2025*. OpenReview.net, 2025.
- [58] MultiSynt. Mt reasoning: Automatic translations into 2 languages of glaive ai reasoning dataset, 2025. An automatic translations into 2 languages of Glaive AI reasoning dataset.
- [59] NVIDIA, :, A. Basant, A. Khairnar, A. Paithankar, A. Khattar, A. Renduchintala, A. Malte, A. Bercovich, A. Hazare, A. Rico, A. Ficek, A. Kondratenko, A. Shaposhnikov, A. Bukharin, A. Taghibakhshi, A. Barton, A. S. Mahabaleshwarkar, A. Shen, A. Tao, A. Guan, A. Shors, A. Mandarwal, A. Mehta, A. Venkatesan, A. Sharabiani, A. Aithal, A. Poojary, A. Dattagupta, B. Buddharaju, B. Zhu, B. Simkin, B. Kartal, B. D. Rouhani, B. Chen, B. Ginsburg, B. Norick, B. Yu, B. Catanzaro, C. Wang, C. Truong, C. Mungekar, C. Patel, C. Alexiuk, C. Munley, C. Parisien, D. Su, D. Afrimi, D. Korzekwa, D. Rohrer, D. Gitman, D. Mosallanezhad, D. Narayanan, D. Reakesh, D. Yared, D. Pykhtar, D. Ahn, D. Riach, E. Long, E. Ning, E. Chung, E. Galinkin, E. Bakhturina, G. Prasad, G. Shen, H. Qian, H. Elisha, H. Sharma, H. Ross, H. Ngo, H. Sahota, H. Wang, H. C. Shin, H. Huang, I. Cunningham, I. Gitman, I. Moshkov, J. Jung, J. Kautz, J. P. Scowcroft, J. Casper, J. Zhang, J. Zeng, J. Zhang, J. Xue, J. Huang, J. Conway, J. Kamalu, J. Cohen, J. Jennings, J. V. Vialard, J. Yi, J. Parmar, K. Briski, K. Cheung, K. Luna, K. Wyss, K. Santhanam, K. Kong, K. Pawelec, K. Anik, K. Li, K. Ahmadian, L. McAfee, L. Sleiman, L. Derczynski, L. Vega, M. R. de Melo, M. N. Sreedhar, M. Chochowski, M. Cai, M. Kliegl, M. Stepniewska-Dziubinska, M. Novikov, M. Samadi, M. Price, M. Boudir, M. Boone, M. Evans, M. Bien, M. Zawalski, M. Martinez, M. Chrzanowski, M. Shoeybi, M. Patwary, N. Dhameja, N. Assaf, N. Habibi, N. Bhatia, N. Pope, N. Tajbakhsh, N. K. Juluru, O. Rybakov, O. Hrinchuk, O. Kuchaiev, O. Olabiyi, P. Ribalta, P. Subramanian, P. Chadha, P. Molchanov, P. Dykas, P. Jin, P. Bialecki, P. Januszewski, P. Thalasta, P. Gaikwad, P. Varshney, P. Gundecha, P. Tredak, R. K. Mahabadi, R. Patel, R. El-Yaniv, R. Rajan, R. Cheruvu, R. Shahbazyan, R. Borkar, R. Gala, R. Waleffe, R. Zhang, R. J. Hewett, R. Prenger, S. Jain, S. Krizan, S. Satheesh, S. Kaji, S. Yurick, S. Muralidharan, S. Narenthiran, S. Bak, S. Sameni, S. Han, S. Ramasamy, S. Ghosh, S. T. Sreenivas, S. Thomas, S. Diao, S. Gopal, S. Prabhumoye, S. Toshniwal, S. Ding, S. Singh, S. Jain, S. Majumdar, S. Singhal, S. Alborghetti, S. N. Akter, T. Kong, T. Moon, T. Hliwiak, T. Asida, T. Wang, T. Konuk, T. Vashishth, T. Poon, U. Karpas,

- V. Noroozi, V. Srinivasan, V. Korthikanti, V. Fugro, V. Kalluru, V. Kurin, V. Lavrukhin, W. U. Ahmad, W. Du, W. Byeon, X. Lu, X. Dong, Y. Karnati, Y. Choi, Y. Zhang, Y. Lin, Y. Fu, Y. Suhara, Z. Dong, Z. Li, Z. Zhu, and Z. Chen. Nvidia nemotron nano 2: An accurate and efficient hybrid mamba-transformer reasoning model, 2025.
- [60] NVIDIA, :, A. Blakeman, A. Basant, A. Khattar, A. Renduchintala, A. Bercovich, A. Ficek, A. Bjorlin, A. Taghibakhshi, A. S. Deshmukh, A. S. Mahabaleshwarkar, A. Tao, A. Shors, A. Aithal, A. Poojary, A. Dattagupta, B. Buddharaju, B. Chen, B. Ginsburg, B. Wang, B. Norick, B. Butterfield, B. Catanzaro, C. del Mundo, C. Dong, C. Harvey, C. Parisien, D. Su, D. Korzekwa, D. Yin, D. Gitman, D. Mosallanezhad, D. Narayanan, D. Fridman, D. Rekesh, D. Ma, D. Pykhtar, D. Ahn, D. Riach, D. Stosic, E. Long, E. Segal, E. Evans, E. Chung, E. Galinkin, E. Bakhturina, E. Dobrowolska, F. Jia, F. Liu, G. Prasad, G. Shen, G. Liu, G. Chen, H. Qian, H. Ngo, H. Liu, H. Li, I. Gitman, I. Karmanov, I. Moshkov, I. Golan, J. Kautz, J. P. Scowcroft, J. Casper, J. Seppanen, J. Lu, J. Sewall, J. Zeng, J. You, J. Zhang, J. Zhang, J. Huang, J. Xue, J. Huang, J. Conway, J. Kamalu, J. Barker, J. Cohen, J. Jennings, J. Parmar, K. Sapra, K. Briski, K. Chumachenko, K. Luna, K. Santhanam, K. Kong, K. Sivamani, K. Pawelec, K. Anik, K. Li, L. McAfee, L. Derczynski, L. Pavao, L. Vega, L. Voegtle, M. Bala, M. R. de Melo, M. N. Sreedhar, M. Chochowski, M. Kliegl, M. Stepniewska-Dziubinska, M. Le, M. Novikov, M. Samadi, M. Andersch, M. Evans, M. Martinez, M. Chrzanowski, M. Ranzinger, M. Blaz, M. Smelyanskiy, M. Fawzy, M. Shoeybi, M. Patwary, N. Lee, N. Tajbakhsh, N. Xu, O. Rybakov, O. Kuchaiev, O. Delalleau, O. Nitski, P. Chadha, P. Shamis, P. Micikevicius, P. Molchanov, P. Dykas, P. Fischer, P.-Y. Aquilanti, P. Bialecki, P. Varshney, P. Gundecka, P. Tredak, R. Karimi, R. Kanduri, R. El-Yaniv, R. Joshi, R. Waleffe, R. Zhang, S. Kavanaugh, S. Jain, S. Krizan, S. Lym, S. Satheesh, S. Muralidharan, S. Narenthiran, S. Anandaram, S. Bak, S. Kashirsky, S. Han, S. Acharya, S. Ghosh, S. T. Sreenivas, S. Clay, S. Thomas, S. Prabhumoye, S. Pachori, S. Toshniwal, S. Prayaga, S. Jain, S. Das, S. Kierat, S. Majumdar, S. Han, S. Singhal, S. Niverty, S. Alborghetti, S. Panguluri, S. Bhendigeri, S. N. Akter, S. Migacz, T. Shiri, T. Kong, T. Roman, T. Ronen, T. Saar, T. Konuk, T. Rintamaki, T. Poon, U. De, V. Noroozi, V. Singh, V. Korthikanti, V. Kurin, W. U. Ahmad, W. Du, W. Ping, W. Dai, W. Byeon, X. Ren, Y. Xu, Y. Choi, Y. Zhang, Y. Lin, Y. Suhara, Z. Yu, Z. Li, Z. Li, Z. Zhu, Z. Yang, and Z. Chen. Nemotron-h: A family of accurate and efficient hybrid mamba-transformer models. *arXiv preprint arXiv:2504.03624*, 2025.
- [61] NVIDIA, :, A. Blakeman, A. Grattafiori, A. Basant, A. Gupta, A. Khattar, A. Renduchintala, A. Vavre, A. Shukla, A. Bercovich, A. Ficek, A. Shaposhnikov, A. Kondratenko, A. Bukharin, A. Milesi, A. Taghibakhshi, A. Liu, A. Barton, A. S. Mahabaleshwarkar, A. Klein, A. Zuker, A. Geifman, A. Shen, A. Bhiwandiwalla, A. Tao, A. Guan, A. Mandarwal, A. Mehta, A. Aithal, A. Poojary, A. Ahamed, A. K. Thekkumkate, A. Dattagupta, B. Zhu, B. Sadeghi, B. Simkin, B. Lanir, B. Schifferer, B. Nushi, B. Kartal, B. D. Rouhani, B. Ginsburg, B. Norick, B. Soubasis, B. Kisacanin, B. Yu, B. Catanzaro, C. del Mundo, C. Hwang, C. Wang, C.-P. Hsieh, C. Zhang, C. Yu, C. Mungekar, C. Patel, C. Alexiuk, C. Parisien, C. Neale, D. Mosk-Aoyama, D. Su, D. Corneil, D. Afrimi, D. Rohrer, D. Serebrenik, D. Gitman, D. Levy, D. Stosic, D. Mosallanezhad, D. Narayanan, D. Nathawani, D. Rekesh, D. Yared, D. Kakwani, D. Ahn, D. Riach, D. Stosic, E. Minasyan, E. Lin, E. Long, E. P. Long, E. Lantz, E. Evans, E. Ning, E. Chung, E. Harper, E. Tramel, E. Galinkin, E. Pounds, E. Briones, E. Bakhturina, F. Ladhak, F. Wang, F. Jia, F. Soares, F. Chen, F. Galko, F. Siino, G. H. Agam, G. Ajjanagadde, G. Bhatt, G. Prasad, G. Armstrong, G. Shen, G. Batmaz, G. Nalbandyan, H. Qian, H. Sharma,

- H. Ross, H. Ngo, H. Sahota, H. Wang, H. Soni, H. Upadhyay, H. Mao, H. C. Nguyen, H. Q. Nguyen, I. Cunningham, I. Shahaf, I. Gitman, I. Loshchilov, I. Moshkov, I. Putterman, J. Kautz, J. P. Scowcroft, J. Casper, J. Mitra, J. Glick, J. Chen, J. Oliver, J. Zhang, J. Zeng, J. Lou, J. Zhang, J. Huang, J. Conway, J. Guman, J. Kamalu, J. Greco, J. Cohen, J. Jennings, J. Daw, J. V. Vialard, J. Yi, J. Parmar, K. Xu, K. Zhu, K. Briski, K. Cheung, K. Luna, K. Santhanam, K. Shih, K. Kong, K. Bhardwaj, K. C. Puvvada, K. Pawelec, K. Anik, L. McAfee, L. Sleiman, L. Derczynski, L. Ding, L. Liebenwein, L. Vega, M. Grover, M. V. Segbroeck, M. R. de Melo, M. N. Sreedhar, M. Kilaru, M. Ashkenazi, M. Romeijn, M. Cai, M. Kliegl, M. Moosaei, M. Novikov, M. Samadi, M. Corpuz, M. Wang, M. Price, M. Boone, M. Evans, M. Martinez, M. Chrzanowski, M. Shoeby, M. Patwary, N. Mulepati, N. Hereth, N. Assaf, N. Habibi, N. Zmora, N. Haber, N. Sessions, N. Bhatia, N. Jukar, N. Pope, N. Ludwig, N. Tajbakhsh, N. Juluru, O. Hrinchuk, O. Kuchaiev, O. Delalleau, O. Olabiyi, O. U. Argov, O. Xie, P. Chadha, P. Shamis, P. Molchanov, P. Morkisz, P. Dykas, P. Jin, P. Xu, P. Januszewski, P. P. Thombre, P. Varshney, P. Gundecha, Q. Miao, R. K. Mahabadi, R. El-Yaniv, R. Zilberstein, R. Shafipour, R. Harang, R. Izzo, R. Shahbazy, R. Garg, R. Borkar, R. Gala, R. Islam, R. Waleffe, R. Watve, R. Koren, R. Zhang, R. J. Hewett, R. Prenger, R. Timbrook, S. Mahdavi, S. Modi, S. Krizan, S. Kariyappa, S. Satheesh, S. Kaji, S. Pasumarthi, S. Narentharen, S. Narenthiran, S. Bak, S. Kashirsky, S. Poulos, S. Mor, S. Ramasamy, S. Acharya, S. Ghosh, S. T. Sreenivas, S. Thomas, S. Fan, S. Gopal, S. Prabhume, S. Pachori, S. Toshniwal, S. Ding, S. Singh, S. Sun, S. Ithape, S. Majumdar, S. Singhal, S. Alborghetti, S. Ge, S. D. Devare, S. K. Barua, S. Panguluri, S. Gupta, S. Priyadarshi, S. N. Akter, T. Bui, T.-D. Ene, T. Kong, T. Do, T. Blankevoort, T. Balough, T. Asida, T. B. Natan, T. Konuk, T. Vashishth, U. Karpas, U. De, V. Noorozi, V. Noroozi, V. Srinivasan, V. Elango, V. Korthikanti, V. Kurin, V. Lavrukhin, W. Jiang, W. U. Ahmad, W. Du, W. Ping, W. Zhou, W. Jennings, W. Zhang, W. Prazuch, X. Ren, Y. Karnati, Y. Choi, Y. Meyer, Y.-F. Wu, Y. Zhang, Y. Lin, Y. Geifman, Y. Fu, Y. Subara, Y. Suhara, Y. Gao, Z. Moshe, Z. Dong, Z. Liu, Z. Chen, and Z. Yan. Nemotron 3 nano: Open, efficient mixture-of-experts hybrid mamba-transformer model for agentic reasoning. *arXiv preprint arXiv:2512.20848*, 2025.
- [62] NVIDIA. Deutsche telekom and NVIDIA launch industrial AI cloud. <https://blogs.nvidia.com/blog/germany-industrial-ai-cloud-launch/>, Nov. 2025. Accessed 2026-07-08.
- [63] S. Oepen, N. Arefev, M. Aulamo, M. Bañón, M. Buljan, L. Burchell, L. Charpentier, P. Chen, M. Fedorova, O. de Gibert, B. Haddow, J. Hajič, J. Helcl, A. Kutuzov, V. Laippala, Z. Li, R. Luukkonen, B. Malik, V. Mikhailov, A. Myntti, D. O’Brien, L. Poláková, S. Pyysalo, G. R. Sánchez, J. Siewert, P. Stepachev, J. Tiedemann, T. Vahtola, D. Variš, F. Vitiugin, T. Vojtěchová, and J. Zaragoza. Hplt 3.0: Very large-scale multilingual resources for llms and mt. mono- and bi-lingual data, multilingual evaluation, and pre-trained models. *arXiv preprint arXiv:2511.01066*, 2026.
- [64] T. Olmo, :, A. Ettinger, A. Bertsch, B. Kuehl, D. Graham, D. Heineman, D. Groeneveld, F. Brahman, F. Timbers, H. Ivison, J. Morrison, J. Poznanski, K. Lo, L. Soldaini, M. Jordan, M. Chen, M. Noukhovitch, N. Lambert, P. Walsh, P. Dasigi, R. Berry, S. Malik, S. Shah, S. Geng, S. Arora, S. Gupta, T. Anderson, T. Xiao, T. Murray, T. Romero, V. Graf, A. Asai, A. Bhagia, A. Wettig, A. Liu, A. Rangapur, C. Anastasiades, C. Huang, D. Schwenk, H. Trivedi, I. Magnusson, J. Lochner, J. Liu, L. J. V. Miranda, M. Sap, M. Morgan, M. Schmitz, M. Guerquin, M. Wilson, R. Huff, R. L. Bras, R. Xin, R. Shao, S. Skjonsberg, S. Z. Shen, S. S. Li, T. Wilde, V. Pyatkin, W. Merrill, Y. Chang, Y. Gu, Z. Zeng, A. Sabharwal, L. Zettle-

- moyer, P. W. Koh, A. Farhadi, N. A. Smith, and H. Hajishirzi. Olmo 3. *arXiv preprint arXiv:2512.13961*, 2026.
- [65] Open Source Initiative. The open source ai definition, version 1.0. <https://opensource.org/ai/open-source-ai-definition>, 2024. Accessed: 2026-07-06.
- [66] OpenEuroLLM Consortium. OpenEuroLLM: A series of foundation models for transparent ai in europe, 2025. Project website, accessed 2026-06-25.
- [67] G. Penedo. Finewiki, 2025. Source: Wikimedia Enterprise Snapshot API (<https://api.enterprise.wikimedia.com/v2/snapshots>). Text licensed under CC BY-SA 4.0 with attribution to Wikipedia contributors.
- [68] G. Penedo, H. Kydlíček, A. Lozhkov, M. Mitchell, C. Raffel, L. Von Werra, T. Wolf, et al. The fineweb datasets: Decanting the web for the finest text data at scale. *Advances in Neural Information Processing Systems*, 37:30811–30849, 2024.
- [69] Pleias and AI Alliance. SYNTH: An open generalist synthetic dataset for training small reasoning models. <https://huggingface.co/datasets/PleIAs/SYNTH>, 2025. Dataset comprising 79,648,272 text samples (over 41 billion words) derived from the synthetic amplification of 58,698 Wikipedia and Wikibooks articles. Licensed under CC-BY-SA 4.0.
- [70] J. Poznanski, A. Rangapur, J. Borchardt, J. Dunkelberger, R. Huff, D. Lin, C. Wilhelm, K. Lo, and L. Soldaini. olmocr: Unlocking trillions of tokens in pdfs with vision language models. *arXiv preprint arXiv:2502.18443*, 2025.
- [71] Qwen Team. Qwen3.5: Towards native multimodal agents, February 2026.
- [72] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever. Language models are unsupervised multitask learners. Technical report, OpenAI, 2019.
- [73] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140):1–67, 2020.
- [74] M. M. Ramos, D. M. Alves, H. Gisserot-Boukhlef, J. Alves, P. H. Martins, P. Fernandes, J. Pombal, N. M. Guerreiro, R. Rei, N. Boizard, A. Farajian, M. Klimaszewski, J. G. C. de Souza, B. Haddow, F. Yvon, P. Colombo, A. Birch, and A. F. T. Martins. Eurollm-22b: Technical report. *arXiv preprint arXiv:2602.05879*, 2026.
- [75] D. Rein, B. L. Hou, A. C. Stickland, J. Petty, R. Y. Pang, J. Dirani, J. Michael, and S. R. Bowman. Gpqa: A graduate-level google-proof q&a benchmark. *arXiv preprint arXiv:2311.12022*, 2023.
- [76] A. Romanou, N. Foroutan, A. Sotnikova, S. H. Nelaturu, S. Singh, R. Maheshwary, M. Altomare, Z. Chen, M. Haggag, S. A. A. Amayuelas, A. H. Amirudin, D. Boiko, M. Chang, J. Chim, G. Cohen, A. K. Dalmia, A. Diress, S. Duwal, D. Dzenhalio, D. Florez, F. Farestam, J. M. Imperial, S. Islam, P. Isotalo, M. Jabbarishivari, B. F. Karlsson, E. Khalilov, C. Klammer, F. Koto, D. Krzemiński, G. de Melo, S. Montariol, Y. Nan, J. Niklaus, J. Novikova, J. S. Obando Ceron, D. Paul, E. Ploeger, J. Purbey, S. Rajwal, S. S. Ravi, S. Rydell, R. Santhosh, D. Sharma, M. Prifti Skenduli, A. Soltani Moakhar, B. moakhar, A. Tarun, A. T. Wasi,

- T. Weerasinghe, S. Yilmaz, M. Zhang, I. Schlag, M. Fadaee, S. Hooker, and A. Bosselut. Include: Evaluating multilingual language understanding with regional knowledge. In Y. Yue, A. Garg, N. Peng, F. Sha, and R. Yu, editors, *International Conference on Learning Representations*, volume 2025, pages 83291–83322, 2025.
- [77] N. Shazeer, A. Mirhoseini, K. Maziarz, A. Davis, Q. Le, G. Hinton, and J. Dean. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv preprint arXiv:1701.06538*, 2017.
- [78] W. Shi, A. Bhagia, K. Farhat, N. Muennighoff, J. Morrison, E. Walsh, D. Schwenk, S. Longpre, J. Poznanski, A. Ettinger, et al. Flexolmo: Open language models for flexible data use. *Advances in Neural Information Processing Systems*, 38:165943–165974, 2026.
- [79] L. Soldaini, R. Kinney, A. Bhagia, D. Schwenk, D. Atkinson, R. Authur, B. Bogin, K. Chandu, J. Dumas, Y. Elazar, V. Hofmann, A. Jha, S. Kumar, L. Lucy, X. Lyu, N. Lambert, I. Magnusson, J. Morrison, N. Muennighoff, A. Naik, C. Nam, M. Peters, A. Ravichander, K. Richardson, Z. Shen, E. Strubell, N. Subramani, O. Tafjord, E. Walsh, L. Zettlemoyer, N. Smith, H. Hajishirzi, I. Beltagy, D. Groeneveld, J. Dodge, and K. Lo. Dolma: An open corpus of three trillion tokens for language model pretraining research. In L.-W. Ku, A. Martins, and V. Srikumar, editors, *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15725–15788, Bangkok, Thailand, Aug. 2024. Association for Computational Linguistics.
- [80] D. Su, K. Kong, Y. Lin, J. Jennings, B. Norick, M. Kliegl, M. Patwary, M. Shoeybi, and B. Catanzaro. Nemotron-CC: Transforming Common Crawl into a refined long-horizon pre-training dataset. In W. Che, J. Nabende, E. Shutova, and M. T. Pilehvar, editors, *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2459–2475, Vienna, Austria, July 2025. Association for Computational Linguistics.
- [81] M. Suzgun, N. Scales, N. Schärli, S. Gehrmann, Y. Tay, H. W. Chung, A. Chowdhery, Q. Le, E. Chi, D. Zhou, and J. Wei. Challenging BIG-bench tasks and whether chain-of-thought can solve them. In A. Rogers, J. Boyd-Graber, and N. Okazaki, editors, *Findings of the Association for Computational Linguistics: ACL 2023*, pages 13003–13051, Toronto, Canada, July 2023. Association for Computational Linguistics.
- [82] A. Talmor, J. Herzig, N. Lourie, and J. Berant. CommonsenseQA: A question answering challenge targeting commonsense knowledge. In J. Burstein, C. Doran, and T. Solorio, editors, *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4149–4158, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics.
- [83] G. Team, A. Kamath, J. Ferret, S. Pathak, N. Vieillard, R. Merhej, S. Perrin, T. Matejovicova, A. Ramé, M. Rivière, L. Rouillard, T. Mesnard, G. Cideron, J. bastien Grill, S. Ramos, E. Yvinec, M. Casbon, E. Pot, I. Penchev, G. Liu, F. Visin, K. Kenealy, L. Beyer, X. Zhai, A. Tsitsulin, R. Busa-Fekete, A. Feng, N. Sachdeva, B. Coleman, Y. Gao, B. Mustafa, I. Barr, E. Parisotto, D. Tian, M. Eyal, C. Cherry, J.-T. Peter, D. Sinopalnikov, S. Bhupatiraju, R. Agarwal, M. Kazemi, D. Malkin, R. Kumar, D. Vilar, I. Brusilovsky, J. Luo, A. Steiner, A. Friesen, A. Sharma, A. Sharma, A. M. Gilady, A. Goedeckemeyer, A. Saade, A. Feng,

- A. Kolesnikov, A. Bendebury, A. Abdagic, A. Vadi, A. György, A. S. Pinto, A. Das, A. Bapna, A. Miech, A. Yang, A. Paterson, A. Shenoy, A. Chakrabarti, B. Piot, B. Wu, B. Shahriari, B. Petrini, C. Chen, C. L. Lan, C. A. Choquette-Choo, C. Carey, C. Brick, D. Deutsch, D. Eisenbud, D. Cattle, D. Cheng, D. Paparas, D. S. Sreepathihalli, D. Reid, D. Tran, D. Zelle, E. Noland, E. Huizenga, E. Kharitonov, F. Liu, G. Amirkhanyan, G. Cameron, H. Hashemi, H. Klimczak-Plucińska, H. Singh, H. Mehta, H. T. Lehri, H. Hazimeh, I. Ballantyne, I. Szpektor, I. Nardini, J. Pouget-Abadie, J. Chan, J. Stanton, J. Wieting, J. Lai, J. Orbay, J. Fernandez, J. Newlan, J. yeong Ji, J. Singh, K. Black, K. Yu, K. Hui, K. Vodrahalli, K. Greff, L. Qiu, M. Valentine, M. Coelho, M. Ritter, M. Hoffman, M. Watson, M. Chaturvedi, M. Moynihan, M. Ma, N. Babar, N. Noy, N. Byrd, N. Roy, N. Momchev, N. Chauhan, N. Sachdeva, O. Bunyan, P. Botarda, P. Caron, P. K. Rubenstein, P. Culliton, P. Schmid, P. G. Sessa, P. Xu, P. Stanczyk, P. Tafti, R. Shivanna, R. Wu, R. Pan, R. Rokni, R. Willoughby, R. Vallu, R. Mullins, S. Jerome, S. Smoot, S. Girgin, S. Iqbal, S. Reddy, S. Sheth, S. Pöder, S. Bhatnagar, S. R. Panyam, S. Eiger, S. Zhang, T. Liu, T. Yacovone, T. Liechty, U. Kalra, U. Evci, V. Misra, V. Roseberry, V. Feinberg, V. Kolesnikov, W. Han, W. Kwon, X. Chen, Y. Chow, Y. Zhu, Z. Wei, Z. Egyed, V. Cotruta, M. Giang, P. Kirk, A. Rao, K. Black, N. Babar, J. Lo, E. Moreira, L. G. Martins, O. Sanseviero, L. Gonzalez, Z. Gleicher, T. Warkentin, V. Mirrokni, E. Senter, E. Collins, J. Barral, Z. Ghahramani, R. Hadsell, Y. Matias, D. Sculley, S. Petrov, N. Fiedel, N. Shazeer, O. Vinyals, J. Dean, D. Hassabis, K. Kavukcuoglu, C. Farabet, E. Buchatskaya, J.-B. Alayrac, R. Anil, Dmitry, Lepikhin, S. Borgeaud, O. Bachem, A. Joulin, A. Andreev, C. Hardin, R. Dadashi, and L. Hussenot. Gemma 3 technical report. *arXiv preprint arXiv:2503.19786*, 2025.
- [84] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar, A. Rodriguez, A. Joulin, E. Grave, and G. Lample. LLaMA: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- [85] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale, D. Bikel, L. Blecher, C. Canton Ferrer, M. Chen, G. Cucurull, D. Esioibu, J. Fernandes, J. Fu, W. Fu, B. Fuller, C. Gao, V. Goswami, N. Goyal, A. Hartshorn, S. Hosseini, R. Hou, H. Inan, M. Kardaş, V. Kerkez, M. Khabsa, I. Kloumann, A. Korenev, P. S. Koura, M.-A. Lachaux, T. Lavril, J. Lee, D. Liskovich, Y. Lu, Y. Mao, X. Martinet, T. Mihaylov, P. Mishra, I. Molybog, Y. Nie, A. Poulton, J. Reizenstein, R. Rungta, K. Saladi, A. Schelten, R. Silva, E. M. Smith, R. Subramanian, X. E. Tan, B. Tang, R. Taylor, A. Williams, J. X. Kuan, P. Xu, Z. Yan, I. Zarov, Y. Zhang, A. Fan, M. Kambadur, S. Narang, A. Rodriguez, R. Stojnic, S. Edunov, and T. Scialom. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- [86] J. Vendrow, E. Vendrow, S. Beery, and A. Madry. Do large language model benchmarks test reliability? *arXiv preprint arXiv:2502.03461*, 2025.
- [87] Y. Wang, X. Ma, G. Zhang, Y. Ni, A. Chandra, S. Guo, W. Ren, A. Arulraj, X. He, Z. Jiang, et al. Mmlu-pro: A more robust and challenging multi-task language understanding benchmark. *Advances in Neural Information Processing Systems*, 37:95266–95290, 2024.
- [88] A. Yang, A. Li, B. Yang, B. Zhang, B. Hui, B. Zheng, B. Yu, C. Gao, C. Huang, C. Lv, C. Zheng, D. Liu, F. Zhou, F. Huang, F. Hu, H. Ge, H. Wei, H. Lin, J. Tang, J. Yang, J. Tu, J. Zhang, J. Yang, J. Yang, J. Zhou, J. Zhou, J. Lin, K. Dang, K. Bao, K. Yang, L. Yu, L. Deng, M. Li, M. Xue, M. Li, P. Zhang, P. Wang, Q. Zhu, R. Men, R. Gao, S. Liu, S. Luo, T. Li, T. Tang, W. Yin, X. Ren, X. Wang, X. Zhang, X. Ren, Y. Fan, Y. Su, Y. Zhang, Y. Zhang,

- Y. Wan, Y. Liu, Z. Wang, Z. Cui, Z. Zhang, Z. Zhou, and Z. Qiu. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025.
- [89] S. Zhang, S. Roller, N. Goyal, M. Artetxe, M. Chen, S. Chen, C. Dewan, M. Diab, X. Li, X. V. Lin, T. Mihaylov, M. Ott, S. Shleifer, K. Shuster, D. Simig, P. S. Koura, A. Sridhar, T. Wang, and L. Zettlemoyer. OPT: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*, 2022.
- [90] W. Zhong, R. Cui, Y. Guo, Y. Liang, S. Lu, Y. Wang, A. Saied, W. Chen, and N. Duan. AGIEval: A human-centric benchmark for evaluating foundation models. In K. Duh, H. Gomez, and S. Bethard, editors, *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 2299–2314, Mexico City, Mexico, June 2024. Association for Computational Linguistics.
- [91] C. Zhou, H. Lyu, X. Lin, H. Zhao, J. Guo, X. Zhang, S. Xue, Q. Ma, J. Zhou, Y. Wang, and Z. Liu. Ultradata-math, 2026.

A Author Contributions

A.1 Training

1. Pretraining stack development and evaluation: Max Lübbering, Richard Rutmann, Timm Ruland, David Fitzek, Mehdi Ali
2. Model architecture, training methodology and framework-correctness validation: Timm Ruland, David Fitzek, Max Lübbering, Richard Rutmann
3. Compute infrastructure, cluster benchmarking and interconnect tuning: David Fitzek, Timm Ruland, Richard Rutmann, Max Lübbering
4. Distributed-training scaling and memory/throughput optimization: David Fitzek, Timm Ruland, Max Lübbering, Richard Rutmann
5. Training execution, stability analysis, emergency debugging, framework bug fixes and experiment tracking: Timm Ruland, David Fitzek, Max Lübbering, Richard Rutmann

A.2 Data

1. Base model data acquisition: Michael Fromm, Alex Jude, Abbas Khan, Ruben Härle, Maurice Kraus, Jan Pfister, Daniil Gurgurov
2. Pretraining data mixture: Michael Fromm
3. Data curation infrastructure and experimentation: Michael Fromm, Alex Jude, Abbas Khan, Ruben Härle, Maurice Kraus, Richard Rutmann, Mehdi Ali, Max Lübbering, Maximilian Idahl
4. Data preprocessing / tokenization pipeline: Richard Rutmann, Max Lübbering, Alex Jude
5. Mid- and long-context data curation and experimentation: Michael Fromm, Alex Jude, Abbas Khan, Ruben Härle, Maurice Kraus, Sebastian Sztwiertnia, Tom Röhr, Sebastian von Rohrscheidt

A.3 Evaluation

1. Evaluation methodology and infrastructure: Maximilian Idahl, Benedikt Droste, Alex Jude, Abbas Khan

A.4 Other

1. Mentorship, advising, program management, and broader strategy: Nicolas Flores-Herr, Simon Gottschalk, Jörg Bienert, Kristian Kersting, Andreas Hotho, Alexander Löser, Wolfgang Nejdl, Simon Ostermann, Jan Plogsties, Patrick Putzky
2. Technical leadership and cross-workstream contributions: Mehdi Ali, Michael Fromm, Max Lübbering, Sebastian Sztwiertnia, Tom Röhr

B Detailed Pretraining Data Composition

This appendix gives the full per-source token accounting underlying the mixture flow diagram in Section 3 (Figure 3). Tables 6 and 7 cover Phase 1, Tables 8 and 9 cover Phase 2 (annealing), and Tables 10 and 11 cover the Phase 3 long-context extension. Rows with zero epochs are sources we enumerated but excluded from training.

Table 6: Phase 1 (diverse pretraining) data composition. “Raw” is the source token count in billions; “Ep.” is the number of epochs; “Eff.” is the effective token count after epoching; “Share” is the percentage of Phase 1 effective tokens. Rows with zero epochs are enumerated but excluded from training. Raw and effective counts are taken from the source dataset cards (HuggingFace) and may reflect different tokenizers; they are approximate. Exact tokenizer counts of consumed tokens appear in Table 2.

Source	Subset / Quality	Raw	Ep.	Eff.	Share
nvidia/Nemotron-CC-v2.1	High-Quality	26.0	3	78.0	0.3%
nvidia/Nemotron-CC-v2.1	Medium-High-Quality	16.9	1	16.9	0.1%
nvidia/Nemotron-CC-v2.1	Medium-Quality	53.5	0	0.0	0.0%
nvidia/Nemotron-CC-v2.1	High-Quality-Synthetic	93.5	2	187.0	0.8%
nvidia/Nemotron-CC-v2.1	Medium-High-Quality-Synthetic	2122.8	1	2122.8	9.2%
nvidia/Nemotron-CC-v2.1	HQ-Translated-To-English	39.6	2	79.2	0.3%
nvidia/Nemotron-CC-v2.1	MHQ-Translated-To-English	26.8	1	26.8	0.1%
nvidia/Nemotron-CC-v2.1	HQ-Translated-To-English-Synthetic	157.8	2	315.6	1.4%
nvidia/Nemotron-CC-v2.1	High-Quality-DQA	8.0	2	16.0	0.1%
<i>Nemotron-CC-v2.1 subtotal</i>				<i>2842.3</i>	<i>12.3%</i>
nvidia/Nemotron-CC-v2.0	High-Quality	613.7	3	1841.1	8.0%
nvidia/Nemotron-CC-v2.0	Medium-High-Quality	545.6	1	545.6	2.4%
nvidia/Nemotron-CC-v2.0	Medium-Quality	2200.8	0	0.0	0.0%
nvidia/Nemotron-CC-v2.0	High-Quality-Synthetic	1257.0	2	2514.0	10.9%
nvidia/Nemotron-CC-v2.0	Diverse QA	692.4	1	692.4	3.0%
nvidia/Nemotron-CC-v2.0	Translated-Diverse QA (DE)	2.0	2	4.0	0.0%
<i>Nemotron-CC-v2.0 subtotal</i>				<i>5597.1</i>	<i>24.3%</i>
nvidia/Nemotron-CC-v1.0	High-Quality	553.0	3	1659.0	7.2%
nvidia/Nemotron-CC-v1.0	Medium-High-Quality	504.0	1	504.0	2.2%
nvidia/Nemotron-CC-v1.0	Medium-Quality	2023.0	0	0.0	0.0%
nvidia/Nemotron-CC-v1.0	High-Synthetic-Diverse QA Pairs	499.5	2	999.0	4.3%
<i>Nemotron-CC-v1.0 subtotal</i>				<i>3162.0</i>	<i>13.7%</i>
nvidia/Nemotron-CC-Code-v1	Actual	427.9	3	1283.7	5.6%
nvidia/Nemotron-Pretraining-Code-v1	Synthetic	174.9	2	349.8	1.5%
nvidia/Nemotron-Pretraining-Code-v1	Actual	125.0	3	375.0	1.6%
nvidia/Nemotron-Pretraining-Code-v2	synthetic-code-review	71.38	2	142.76	0.6%
nvidia/Nemotron-Pretraining-Code-v2	synthetic-question-answering	212.52	2	425.04	1.8%
nvidia/Nemotron-Pretraining-Code-v2	synthetic-rewriting	76.84	2	153.68	0.7%
nvidia/Nemotron-Pretraining-Code-v2	synthetic-student-teacher	28.34	2	56.68	0.2%
nvidia/Nemotron-Pretraining-Code-v2	synthetic-transpilation	24.09	2	48.18	0.2%
nvidia/Nemotron-Pretraining-Code-v2	Actual	180.0	3	540.0	2.3%
<i>Code subtotal</i>				<i>3374.84</i>	<i>14.6%</i>
nvidia/Nemotron-Pretraining-Specialized-v1	—	270.7	5	1353.5	5.9%
nvidia/Nemotron-Pretraining-SFT-v1	Math SFT	190.6	5	953.0	4.1%
nvidia/Nemotron-Pretraining-SFT-v1	Code SFT	58.5	6	351.0	1.5%
nvidia/Nemotron-Pretraining-SFT-v1	General SFT	87.5	3	262.5	1.1%
<i>Specialized + SFT subtotal</i>				<i>2920.0</i>	<i>12.7%</i>
nvidia/Nemotron-CC-Math-v1	3plus	133.0	4	532.0	2.3%
nvidia/Nemotron-CC-Math-v1	4plus	52.0	4	208.0	0.9%

continued on next page

Table 6 – continued

Source	Subset / Quality	Raw	Ep.	Eff.	Share
nvidia/Nemotron-CC-Math-v1	v1	73.0	4	292.0	1.3%
nvidia/Nemotron-Math-v2	high	2.7	4	10.8	0.0%
nvidia/Nemotron-Math-v2	medium	2.0	4	8.0	0.0%
nvidia/Nemotron-Math-v2	low	1.2	4	4.8	0.0%
openbmb/UltraData-Math	en	88.0	4	352.0	1.5%
<i>Mathematics subtotal</i>				1407.6	6.1%
MultiSynt/MT-Reasoning	en	35.8	2	71.6	0.3%
nvidia/AceReason-1.1-SFT	en	30.0	2	60.0	0.3%
HuggingFaceFW/finewiki	en	10.0	5	50.0	0.2%
allenai/dolma3_pool	pdfs	600.0	1	600.0	2.6%
PleIAs/Synth	en	60.0	2	120.0	0.5%
HuggingFaceFW/finepdfs	en	1190.65	1	1190.65	5.2%
<i>English (other) subtotal</i>				2092.25	9.1%
HuggingFaceFW/finepdfs	de	177.56	2	355.12	1.5%
HPLT-3-Top10%	de	60.9	8.4	511.56	2.2%
coral-nlp/german-commons	de	154.56	1	154.56	0.7%
MultiSynt/MT-Nemotron-CC	de	117.0	2	234.0	1.0%
Genios	de	150.0	2	300.0	1.3%
HuggingFaceFW/finewiki	de	3.5	2	7.0	0.0%
MultiSynt/MT-Reasoning	de	42.0	2	84.0	0.4%
DGurgurov/Nemotron-Multilingual-Reasoning	de	2.0	2	4.0	0.0%
PleIAs/Synth	de	2.4	2	4.8	0.0%
<i>German subtotal</i>				1655.04	7.2%
Total		16,350.04		23,051.13	100.0%

Table 7: Phase 1 composition by category, with Nemotron 3 Nano’s reported shares for comparison. These categories cover the full Phase 1 mixture of 23,051.13B effective tokens (100%).

Category	Soofi (B)	Soofi	Nemotron 3 Nano
Web Medium	0.0	0.00%	6.8%
Academic	1790.65	7.77%	4.1%
German/Multiling	1655.04	7.18%	5.0%
Web High-Q-Synthetic	4110.8	17.83%	20.4%
Mathematics	1407.6	6.11%	6.4%
Crawl++	0.0	0.00%	2.9%
Wiki	50.0	0.22%	0.6%
Web Medium-High	1066.5	4.63%	5.7%
Web Med-H-Synthetic	2846.0	12.35%	11.7%
Code-SFT	351.0	1.52%	3.3%
Web High	3578.1	15.52%	6.5%
Code	3374.84	14.64%	15.3%
General SFT	454.1	1.97%	0.2%
STEM SFT	2366.5	10.27%	11.1%
Total	23,051.13	100.0%	100.0%

Table 8: Phase 2 (high-quality annealing) data composition. Columns as in Table 6. Rows with zero epochs are enumerated but excluded from training. Raw and effective counts are taken from the source dataset cards (HuggingFace) and may reflect different tokenizers; they are approximate. Exact tokenizer counts of consumed tokens appear in Table 2.

Source	Subset	Raw	Ep.	Eff.
Web				
nvidia/Nemotron-CC-v2.1	High-Quality	26.0	1	26.0
nvidia/Nemotron-CC-v2.1	High-Quality-Synthetic	93.5	1	93.5
nvidia/Nemotron-CC-v2.1	HQ-Translated-To-English	39.6	1	39.6
nvidia/Nemotron-CC-v2.1	HQ-Translated-To-English-Synthetic	157.8	0	0.0
nvidia/Nemotron-CC-v2.1	High-Quality-DQA	8.0	1	8.0
nvidia/Nemotron-CC-v2.0	High-Quality	613.7	0	0.0
nvidia/Nemotron-CC-v2.0	High-Quality-Synthetic	1257.0	0.5	628.5
nvidia/Nemotron-CC-v2.0	Diverse QA	692.4	1	692.4
nvidia/Nemotron-CC-v2.0	Translated-Diverse QA (DE)	2.0	0	0.0
nvidia/Nemotron-CC-v1.0	High-Quality	553.0	0	0.0
nvidia/Nemotron-CC-v1.0	High-Synthetic-Diverse QA Pairs	499.5	0	0.0
allenai/dolma3_dolmino_pool	Web	5.21	1	5.21
allenai/dolma3_dolmino_pool	pdfs	240.0	1	240.0
karpathy/climbmix-400b-shuffle	en	400.0	2	800.0
HuggingFaceFW/finepdfs-edu	en	142.0	1	142.0
Web subtotal			2675.21	
Code				
nvidia/Nemotron-Pretraining-Code-v1	Synthetic	174.9	1	174.9
nvidia/Nemotron-Pretraining-Code-v1	Actual	125.0	1	125.0
nvidia/Nemotron-Pretraining-Code-v2	synthetic-code-review	71.38	1	71.38
nvidia/Nemotron-Pretraining-Code-v2	synthetic-question-answering	212.52	1	212.52
nvidia/Nemotron-Pretraining-Code-v2	synthetic-rewriting	76.84	1	76.84
nvidia/Nemotron-Pretraining-Code-v2	synthetic-student-teacher	28.34	1	28.34
nvidia/Nemotron-Pretraining-Code-v2	synthetic-transpilation	24.09	1	24.09
nvidia/Nemotron-Pretraining-Code-v2	Actual	180.0	1	180.0
tokyotech-llm/swallow-code-v2	stage5	49.8	2	99.6
allenai/dolma3_dolmino_pool	Code	40.0	1	40.0
Code subtotal			1032.67	
Mathematics				
nvidia/Nemotron-CC-Math-v1	3plus	133.0	1	133.0
nvidia/Nemotron-CC-Math-v1	4plus	52.0	1	52.0
allenai/dolma3_dolmino_pool	Math	21.34	1	21.34
openbmb/UltraData-Math	en	88.0	1.4	123.2
Mathematics subtotal			329.54	
SFT				
nvidia/Nemotron-Pretraining-SFT-v1	Math SFT	190.6	2	381.2
nvidia/Nemotron-Pretraining-SFT-v1	Code SFT	58.5	2	117.0
nvidia/Nemotron-Pretraining-SFT-v1	General SFT	87.5	1	87.5
nvidia/Nemotron-Agentic-v1	agentic	1.0	2	2.0
nvidia/Nemotron-Competitive-Programming-v1	code-sft	50.0	2	100.0
nvidia/Nemotron-Instruction-Following-Chat-v1	General SFT	1.5	2	3.0
nvidia/Nemotron-Math-Proofs-v1	stem-sft	6.0	2	12.0
nvidia/Nemotron-Math-v2	stem-sft	30.0	2	60.0
nvidia/Nemotron-RLHF-GenRM-v1	—	0.5	2	1.0
nvidia/Nemotron-Science-v1	stem-sft	0.5	2	1.0
nvidia/Nemotron-SFT-Agentic-v2	agentic	1.5	2	3.0

continued on next page

Table 8 – continued

Source	Subset	Raw	Ep.	Eff.
nvidia/Nemotron-SFT-Competitive-Programming-v2	code-sft	20.0	2	40.0
nvidia/Nemotron-SFT-Instruction-Following-Chat-v2	General SFT	3.0	2	6.0
nvidia/Nemotron-SFT-Math-v3	stem-sft	1.0	2	2.0
nvidia/Nemotron-SFT-Multilingual-v1	multilingual-sft	3.5	2	7.0
nvidia/Nemotron-SFT-OpenCode-v1	code-sft	7.0	2	14.0
nvidia/Nemotron-SFT-Safety-v1	safety-sft	0.03	2	0.06
nvidia/Nemotron-SpecializedDomains-Finance-v1	stem-sft	4.0	2	8.0
nvidia/Nemotron-SWE-v1	code-sft	0.7	2	1.4
nvidia/Nemotron-SFT-SWE-v2	code-sft	2.5	2	5.0
nvidia/AceReason-1.1-SFT	en	30.0	1	30.0
allenai/dolma3_dolmino_pool	QA	25.8	1	25.8
allenai/dolma3_dolmino_pool	Instruction-Data	18.41	1	18.41
AIML-TUDA/QA-base	en	0.143	10	1.43
<i>SFT subtotal</i>				<i>926.80</i>
Reasoning				
nvidia/Nemotron-Pretraining-Specialized-v1	—	270.7	1	270.7
allenai/dolma3_dolmino_pool	Thinking	37.6	1	37.6
nvidia/Nemotron-Pretraining-Specialized-v1.1	en	9.3	1	9.3
MultiSynt/MT-Reasoning	en	35.8	1	35.8
<i>Reasoning subtotal</i>				<i>353.40</i>
Wiki				
HuggingFaceFW/finewiki	en	10.0	2	20.0
<i>Wiki subtotal</i>				<i>20.00</i>
German				
HuggingFaceFW/finepdfs-edu	de	20.0	2	40.0
AIML-TUDA/QA-base	de	0.187	10	1.87
HPLT-4-Top10%	de	291.0	1	291.0
coral-nlp/german-commons	de	154.56	0	0.0
German Translation of ClimbMix	de	571.0	1	571.0
Genios	de	150.0	0	0.0
HuggingFaceFW/finewiki	de	3.5	1	3.5
MultiSynt/MT-Reasoning	de	42.0	1	42.0
DGurgurov/Nemotron-Multilingual-Reasoning	de	2.0	1	2.0
toroe/Soofi-Think-SFT-10B-multilingual	de	7.0	2	14.0
<i>German subtotal</i>				<i>965.37</i>
Total			6,303.0	

C Proxy Data-Mixture Ablations

Before fixing the data recipe used for Soofi S, we ran controlled small-scale data-mixture ablations. These runs were designed as a proxy study: their role was to select a robust German–English data family under a much cheaper setup, not to predict the absolute performance of the final 30B-A3B hybrid-Mamba MoE model. To make the comparison realistic for the final setting, we held the English backbone fixed and changed only the German or multilingual part of the mixture. English therefore remained the dominant language in every candidate, contributing roughly 90–93% of the 100B-token proxy budget, while the remaining budget tested different allocations among German web, PDF, synthetic, commons, news, and lightly multilingual sources. This isolates the main

Table 9: Phase 2 (annealing) composition by category, with Nemotron 3 Nano’s reported shares for comparison. Percentages are shares of the 6,303 B annealing pool. For the seven-category main-text view in Figure 3, the Reasoning/STEM-SFT bucket is folded into Reasoning.

Category	Soofi (B)	Soofi	Nemotron 3 Nano
Web High	831.21	13.19%	6.50%
Web Medium-High	0.0	0.00%	4.00%
Web High-Q-Synthetic	1462.0	23.20%	20.40%
Web Med-H-Synthetic	0.0	0.00%	5.00%
Code	1032.67	16.38%	14.00%
Code-SFT	282.4	4.48%	6.70%
Reasoning / STEM-SFT	744.20	11.81%	22.30%
General-SFT	253.60	4.02%	0.40%
Academic / PDFs (High)	382.0	6.06%	2.00%
Wiki	20.0	0.32%	1.30%
Multilingual	965.37	15.32%	5.00%
Math	329.54	5.23%	12.50%
Total	6303.0	100.0%	100.0%

design question for Soofi S: how to spend the scarce non-English budget without confounding the result with changes in English data.

Ablated mixtures. The English backbone follows a fixed Nemotron-style composition of web, code, math, SFT, and PDF sources [61]. The ablation identifiers D01, D03–D13 correspond to the non-English mixtures summarized in Table 12.

The candidates cover several qualitatively different hypotheses. D01 combines all German source families with a strong HPLT-3 [63] web anchor and small amounts of German-Commons [24] and Genios [23]. D12 uses the same high-level shares but swaps the German HPLT quality classifier for the Propella-filtered variant [36]. D04 is the only explicitly multilingual candidate, adding Spanish, Italian, and French HPLT-3 top-decile data to the otherwise German tail. D08 tests an almost pure German FinePDFs [44] tail, D09 and D10 emphasize MultiSynt [37], and D11 tests a broader HPLT-3 German filter by moving from the top decile to the top two deciles.

Proxy training setup. All ablations used the same training configuration and differed only in the chosen tokenized dataset blend. The proxy model was based on the Qwen3-1.7B pretraining recipe [88]. We trained with bf16 mixed precision, global batch size 256, and micro-batch size 1. The runs were launched on four Leonardo nodes with four GPUs per node. Checkpoints were written during training and evaluated offline with the `lm-evaluation-harness` [22] pipeline. We used a training batch size of 2,097,152 tokens per optimizer step, and the final comparable checkpoint at step 47,683 corresponds to $\sim 100B$ trained tokens. Checkpoints were evaluated at approximately 10B-token intervals, giving a learning curve rather than a single endpoint for each candidate mixture.

Evaluation protocol. For mixture selection we used the same English–German harness family as in the main report. The primary scalar criterion was the average rank over four suite-level signals:

Table 10: Long-context phase: per-domain token budgets, document counts, and source priorities for the released data pool. The run consumed ~ 100.66 B of the 188.5B pool (Section 3.4).

Domain	Tokens (B)	Documents	Sources (priority order)
Web	28.78	3,051,984	ClimbMix > OlmoOCR > FinePDFs
Code	9.51	344,625	Swallow-Code-v2 > Nemotron-Pretraining-Code-v1 , v2
Mathematics	6.73	2,526,637	openbmb/UltraData-Math-L3 ; no data >64K
German	6.68	714,989	40% HPLT-4-Top10% + 60% German Translation of ClimbMix
General SFT	31.25	4,386,847	nvidia/Nemotron-Pretraining-SFT-v1
Code SFT	29.24	2,541,896	nvidia/Nemotron-Pretraining-SFT-v1
Math SFT	76.31	7,990,170	nvidia/Nemotron-Pretraining-SFT-v1
Total	188.49	21,557,148	

Table 11: Long-context phase: number of documents per sequence-length bucket and domain in the released pool. “–” denotes a bucket left unpopulated. Unlike the idealized symmetric schema, the realized counts do not halve cleanly across buckets.

Seq. len.	Web	Code	Math	German	Code SFT	Gen SFT	Math SFT
4K	1,535,801	–	1,781,457	367,330	–	–	–
8K	765,057	3	598,808	178,417	1,002,824	3,307,656	4,091,855
16K	386,188	173,557	123,606	87,248	969,107	851,104	2,991,068
32K	198,283	86,572	22,347	46,283	544,057	225,083	842,927
64K	99,484	43,424	417	23,253	25,908	3,004	64,299
128K	49,102	22,099	2	9,867	–	–	19
256K	–	10,840	–	–	–	–	2
512K	12,046	5,420	–	2,523	–	–	–
1M	6,023	2,710	–	68	–	–	–
Total	3,051,984	344,625	2,526,637	714,989	2,541,896	4,386,847	7,990,170

English bits-per-byte, German bits-per-byte, English normalized rank-choice accuracy, and German normalized rank-choice accuracy. Bits-per-byte is lower-better; normalized rank-choice accuracy is higher-better. We preferred ranks over raw-score averaging because the four metrics have different scales. All terminal averages in Table 13 are computed using only the step-47,683 checkpoint. We explicitly ignore the near-duplicate step-47,680 endpoint so that the final evaluation is not double-counted.

Results. The rank-based selection criterion gives a clear winner. D01 has the best terminal average rank ($R_{100B} = 1.75$) and the best average rank across the full training trace ($R_{all} = 3.55$). At the final checkpoint it is best on English bits-per-byte and German normalized accuracy, second on English normalized accuracy, and close to the strongest group on German bits-per-byte. This balance matters because several alternatives win a single metric but are less stable overall. D06, for example, achieves the lowest German bits-per-byte, but its average rank is only sixth at the final checkpoint and seventh over the full trace. D12 is the strongest runner-up and is slightly better on English normalized accuracy, but it gives back performance on both bits-per-byte suites and on German normalized accuracy relative to D01.

Table 12: Non-English mixture shares for the evaluated proxy ablations. Entries are percentages of the 100B-token proxy budget recorded in the ablation sheet. “0.1q” denotes the top 10% HPLT-3 filter, and “0.2q” denotes the top 20% filter. The “Other HPLT” column is non-German HPLT-3 data (Spanish, Italian, and French) and is nonzero only for the explicitly multilingual ablation D04.

Run	HPLT-DE filter	HPLT-DE %	FinePDFs-DE %	MultiSynt-DE %	German-Commons %	Genios %	Other HPLT %	Non-EN total %
D01	0.1q	3.04	1.77	0.93		0.72	0.80	7.26
D03	0.1q	3.04	1.77	2.33		0.00	0.00	7.14
D04	0.1q	3.04	1.71	1.99		0.00	0.00	9.71
D05	0.1q	0.30	1.77	1.40		2.16	1.20	6.83
D06	0.1q	1.52	3.54	1.86		0.00	0.00	6.92
D07	0.1q	0.30	5.31	1.40		0.00	0.00	7.01
D08	0.1q	0.00	7.08	0.00		0.00	0.00	7.08
D09	0.1q	0.61	1.77	4.65		0.00	0.00	7.03
D10	0.1q	1.52	2.66	4.65		4.32	0.00	13.15
D11	0.2q	6.09	0.00	0.93		0.00	0.00	7.02
D12	0.1q Propella	3.04	1.77	0.93		0.72	0.80	7.26
D13	0.1q	0.91	4.42	1.86		0.00	0.00	7.20

Table 13: Proxy data-mixture ablations at the terminal 100B-token checkpoint, computed using only step 47,683. “EN/DE bpb” are suite-level bits-per-byte scores (lower is better); “EN/DE acc.” are normalized rank-choice accuracies in percent (higher is better). R_{100B} is the average rank over these four terminal suite metrics. Best values are bolded and second-best values are underlined.

Run	EN bpb ↓	DE bpb ↓	EN acc. ↑	DE acc. ↑	R_{100B} ↓
D01	0.800	0.796	<u>46.8</u>	67.4	1.75
D12	0.803	0.799	46.8	<u>67.3</u>	<u>3.25</u>
D04	0.806	<u>0.795</u>	46.5	67.2	4.50
D09	0.805	0.796	46.6	66.4	5.00
D13	<u>0.803</u>	0.800	46.6	66.6	5.50
D06	0.807	0.791	46.3	66.7	6.00
D03	0.810	0.797	46.5	67.1	6.50
D11	0.808	0.807	46.4	67.1	7.50
D08	0.805	0.851	46.6	65.5	8.50
D10	0.815	0.797	46.3	66.9	8.50
D07	0.814	0.800	46.1	66.4	10.50
D05	0.809	0.812	46.2	66.1	10.50

The source-level patterns are also informative. A mixture consisting almost entirely of German FinePDFs (D08) performs competitively on English but collapses on German bits-per-byte, suggesting that document-style PDF text alone is too narrow for the German tail. Heavy MultiSynt mixtures improve some German likelihood tasks but do not produce the best bilingual aggregate (D09, D10). The explicitly multilingual mixture (D04) is a strong candidate and obtains the second-best German bits-per-byte score, but it does not match D01 on the combined English–German selection criterion. The comparison between D01 and D12 further suggests that the original HPLT-DE component is preferable to the Propella-filtered substitute under this proxy setup, even when the high-level source shares are otherwise unchanged.

Task-group aggregates support the same conclusion. D01 is best on the English likelihood suite, best on the English math and English QA bits-per-byte aggregates, and best on the German code bits-per-byte and German grammar-fluency aggregates. Its weaker German QA and German math likelihood scores are offset by stronger German rank-choice and fluency performance, which better

Source	Docs (M)	Words/doc	Years
Rheinische Post	5.48	220	2010–2025
Rhein-Zeitung	5.32	242	2010–2025
Neue Westfälische	5.05	208	2010–2025
Schwäbische Zeitung	3.15	272	2011–2025
Märkische Allgemeine	3.15	223	2010–2025
Thüringer Allgemeine	3.07	250	2010–2025
Passauer Neue Presse	2.97	303	2010–2025
Mitteldeutsche Zeitung	2.65	309	2010–2025
Ostthüringer Zeitung	2.64	259	2010–2025
Sächsische Zeitung	2.32	312	2010–2025
Other (906 sources)	157.3	—	2010–2025
Total (916 sources)	193.1	298 (avg.)	2010–2025

Table 14: Composition of the Genios corpus [23]: the ten largest publications by document count, out of 916 German newspaper and trade-press sources totalling 193.1M articles and ~ 57.6 B words.

matches the intended downstream profile of a German–English base model.

Selection for the full run. We therefore used the mixture represented by D01 as the basis for the full Soofi S data recipe. The proxy result was not copied mechanically into the final 26.68T-token curriculum: the final training plan still separates broad pretraining from high-quality annealing, up-weights German in both phases, and adds the long-context extension described in Section 2.2. Nevertheless, the ablations provide the empirical justification for selecting a balanced German data family—HPLT-DE, German FinePDFs, MultiSynt, German-Commons, and Genios—rather than optimizing a single benchmark, a single German source, or a broad multilingual tail in isolation.

D Further Dataset Information

Genios. The Genios corpus [23] is a commercially licensed collection of German-language newspaper and trade-press archives obtained from GBI-Genios. It comprises 916 distinct publications with 193.1M articles (~ 57.6 B words) spanning 2010–2025, with per-year volumes between 9.7M and 14.4M documents (Figure 15). The collection is dominated by regional daily newspapers (e.g. *Rheinische Post*, *Rhein-Zeitung*, *Neue Westfälische*), complemented by a long tail of national outlets and specialist trade and academic periodicals (Table 14). Articles average ~ 298 words. Because the data was purchased under a commercial license, it cannot be redistributed; we therefore report aggregate corpus statistics only.

E Further Base Model Evaluations

Long-context. We evaluate long-context behaviour with RULER [34] on the checkpoint after the long-context stage (Section 2.2), comparing against Nemotron 3 Nano 30B-A3B. Because the two models share an identical backbone, any difference in long-context accuracy isolates the effect of the long-context *data* and continuation recipe rather than the architecture. Figure 16 reports accuracy




Figure 15: Temporal distribution of the 193.1M Genios articles [23]. Coverage is roughly uniform across 2010–2025 (9.7–14.4M documents per year).




Figure 16: RULER accuracy averaged over the selected subtasks, by input context length (4K–1M). (a) all 13 subtasks; (b) all subtasks except CWE; (c) CWE (common-word extraction) only, where Soofi S degrades sharply past 32K. Shaded band = accuracy gap between the two models.

averaged over RULER subtasks as a function of input length from 4K to 1M tokens.

On the full 13-subtask suite Soofi S trails the reference by a mean of 6.8 points across lengths, reaching 50 versus 60 at 1M (Figure 16a). The gap is almost entirely attributable to a single subtask, common-word extraction (CWE), which requires aggregating and reproducing the most frequently occurring words across the entire input rather than retrieving a localized span. Excluding CWE, the two models track each other to within a mean of 4.4 points across all lengths and to within ~ 6 points at 1M (54 versus 60), i.e. effectively on par out to the full context window (Figure 16b). On CWE itself Soofi S matches the reference up to 32K but degrades sharply at longer inputs (Figure 16c), falling to $\sim 3\%$ at 256K–1M while Nemotron retains 60–64%.

We attribute this regression to the long-context data mixture rather than the backbone, for two reasons. First, the two models are architecturally identical, so the fixed-size Mamba-2 recurrent state and the 6-of-52 attention layers cannot by themselves explain a gap against the same architecture

trained on a different long-context blend. Second, the reference recipe deliberately targets this exact class of task: the Nemotron long-context phase devotes its blend to long-context document-QA data (scaled $3\times$ over the previous generation) together with a dedicated slice of synthetic *retrieval-focused* data at up to 256K tokens, added specifically to improve RULER-style subtasks, with the remainder being down-weighted high-quality pretraining data [61]. Our long-context pool (Section 3.4), by contrast, is dominated by length-bucketed SFT and general document text sampled uniformly, and contains no dedicated retrieval- or aggregation-oriented long-context data. The two runs consume a comparable long-context token budget ($\sim 100.66\text{B}$ here versus $\sim 121\text{B}$ for the reference [61]), so the difference reflects mixture composition rather than long-context training volume. We flag CWE beyond 32K as a known limitation and a concrete target for the next iteration of the long-context data pipeline, specifically, adding retrieval- and aggregation-style synthetic data in the 32K–1M range, and note that, this single subtask aside, Soofi S retains near-parity with its architecture-matched reference across the full 1M-token range.

F Checkpoint Merging Ablations

In addition to selecting a single late-stage checkpoint, we evaluated a series of post-hoc checkpoint merges over the annealing trajectory. These experiments were intended to test whether weight-space averaging could reduce checkpoint noise at the end of training and improve robustness without adding any new training tokens. In all cases, checkpoints came from the same training run or from direct continuations of it, so the models had identical architecture, tokenizer, tensor layout, and optimizer history. The merges therefore average model weights only; optimizer states were not included.

The relevant baseline for this comparison is `iter_1056000`, the checkpoint selected for the Section 4 base-model evaluation and used to initialize the long-context extension. This checkpoint is the model we would have released without any post-hoc merge. Later final-annealing checkpoints and their merges are therefore treated as ablations against this selected checkpoint, not as replacements by default.

Merge implementation. All merges were performed in the original distributed-checkpoint layout with a shared reference checkpoint. Floating-point tensors were accumulated in `float32`, while non-floating tensors were copied from the reference checkpoint. We considered three families of merge weights. First, uniform averaging assigns equal weight to each checkpoint in a window. Second, exponential averaging assigns larger weight to later checkpoints, with decay values $\alpha \in \{0.2, 0.8\}$; larger α keeps the merge closer to the end of the trajectory. Third, two-checkpoint manual blends mix the penultimate or earlier endpoint with the final checkpoint using either 0.2/0.8 or 0.1/0.9 weights. All manual weights were normalized and constrained to be non-negative.

Evaluation. We evaluated 22 merged checkpoints with the same benchmark harness used for late-stage checkpoint selection and compared them directly to the selected `iter_1056000` checkpoint. For the main comparison we used four aggregate suite metrics: English bits-per-byte, German bits-per-byte, English normalized rank-choice accuracy, and German normalized rank-choice accuracy. Bits-per-byte metrics are lower-better, while normalized accuracies are higher-better. As in the data-mixture ablations, we summarize the trade-off using an average rank over the four aggregate metrics rather than averaging raw scores with different scales. One two-checkpoint final merge had

incomplete bits-per-byte coverage in the filtered evaluation file and was therefore excluded from the rank table.

Outcome. The merge ablation did not reveal a uniformly better model than the selected `iter_1056000` checkpoint. The best merged variants were concentrated in the final annealing window: the uniform average over all 13 final-window checkpoints had the best aggregate rank, and the late-biased exponential average with $\alpha = 0.8$ was nearly tied. These merges slightly improved the German suite metrics relative to `iter_1056000`: German bits-per-byte improved from 0.3656 to 0.3637, and German normalized accuracy improved from 80.10 to 80.34. However, the selected checkpoint remained better on the English suite metrics, with English bits-per-byte 0.4390 and English normalized accuracy 77.39, compared with 0.4396 and 77.28 for the best uniform merge. This trade-off explains why we did not treat checkpoint merging as a decisive source of additional capability. The final-window merges provide useful evidence that the end of annealing lies in a stable basin and that small German-side gains are possible through weight averaging. At the same time, those gains come with small regressions on the English aggregate metrics and do not clearly dominate the checkpoint already used for the long-context continuation and the main Section 4 evaluation. We therefore report the merge results for transparency, but keep `iter_1056000` as the primary selected base checkpoint.

Table 15: Checkpoint-merge strategies evaluated after annealing. “Reference” is the checkpoint whose metadata and non-floating tensors were used for the merged export.

Family	Window / checkpoints	Variants	Reference
annealing_last_1t_*	12 checkpoints from iter_0952000 through iter_0993410 in the-big-run-4.	uniform, exp. $\alpha = 0.2$, exp. $\alpha = 0.8$	iter_0993410
annealing_last_2_*	Two-checkpoint tail blend of iter_0992000 and iter_0993410.	0.2/0.8, 0.1/0.9	iter_0993410
more_annealing_last_1t_*	16 checkpoints from iter_0996000 through iter_1056000 in the more/checkpoints continuation.	uniform, exp. $\alpha = 0.2$, exp. $\alpha = 0.8$	iter_1056000
more_annealing_last_2_*	Two-checkpoint tail blend of iter_1052000 and iter_1056000.	0.2/0.8, 0.1/0.9	iter_1056000
final_annealing_all_1t_*	13 checkpoints from iter_1056000_original through iter_1067920.	uniform, exp. $\alpha = 0.2$, exp. $\alpha = 0.8$	iter_1067920
final_annealing_all_2_*	Two-checkpoint tail blend of iter_1067000 and iter_1067920.	0.2/0.8, 0.1/0.9	iter_1067920
final_annealing_div4000_plus_last_1t_*	Sparse final-window merge over iter_1056000_original, iter_1060000, iter_1064000, and iter_1067920.	uniform, exp. $\alpha = 0.2$, exp. $\alpha = 0.5$, exp. $\alpha = 0.8$	iter_1067920
final_annealing_div4000_plus_last_2_*	Two-checkpoint endpoint blend of iter_1064000 and iter_1067920.	0.2/0.8, 0.1/0.9	iter_1067920
top3_uniform	Uniform average of three evaluation-selected final checkpoints: iter_1056000_original, iter_1060000, and iter_1064000.	uniform	iter_1064000

Table 16: Checkpoint merges compared against the selected base checkpoint used for Section 4 and long-context initialization. Accuracies are reported as percentages. R is the average rank over English bpb, German bpb, English normalized accuracy, and German normalized accuracy, computed over the selected checkpoint and all complete merge evaluations.

Candidate	EN bpb ↓	DE bpb ↓	EN acc. ↑	DE acc. ↑	R ↓
final_annealing_all_1t_uniform	0.4396	0.3637	77.28	80.34	3.00
final_annealing_all_1t_exp_avg_0.8	0.4394	0.3637	77.25	80.32	3.25
final_annealing_all_2_0.9	0.4396	0.3637	77.22	80.30	5.75
more_annealing_last_1t_exp_avg_0.8	0.4397	0.3641	77.37	80.27	6.75
top3_uniform	0.4396	0.3644	77.29	80.24	7.50
final_annealing_div4000_plus_last_2_0.8	0.4396	0.3638	77.19	80.26	7.75
final_annealing_all_1t_exp_avg_0.2	0.4395	0.3639	77.24	80.13	8.75
Selected checkpoint iter_1056000	0.4390	0.3656	77.39	80.10	9.25
final_annealing_div4000_plus_last_2_0.9	0.4397	0.3637	77.11	80.26	9.25