

Multilinguality as Sense Adaptation

Jan Christian Blaise Cruz^{1,5*} David Ifeoluwa Adelani^{2,3,4} Alham Fikri Aji^{1,5}

¹MBZUAI ²McGill University ³Mila - Quebec AI Institute

⁴Canada CIFAR AI Chair ⁵SEACrowd

{jan.cruz, alham.fikri}@mbzuai.ac.ae david.adelani@mila.quebec

Abstract

We approach multilinguality as *sense adaptation*: aligning latent meaning representations across languages rather than relying solely on shared parameters and scale. In this paper, we introduce **SENSe-based Symmetric Interlingual Alignment** (SENSIA), which adapts a Backpack language model from one language to another by explicitly aligning sense-level mixtures and contextual representations on parallel data, while jointly training a target-language language modeling loss to preserve fluency. Across benchmarks on four typologically diverse languages, SENSIA generally outperforms comparable multilingual alignment methods and achieves competitive accuracy against monolingual from-scratch baselines while using 2–4× less target-language data. Analyses of learned *sense geometry* indicate that local sense topology and global structure relative to English are largely preserved, and ablations show that the method is robust in terms of design and scale. Our results show that sense adaptation provides a simple, data-efficient, and linguistics-motivated path to multilingual fluency that is useful to low-resource languages and motivates further linguistics-motivated research.

1 Introduction

Multilingual language models now underpin cross-lingual NLP, from encoder-only systems such as XLM-R (Conneau et al., 2020) to generative models like mT5 (Xue et al., 2021), XGLM (Lin et al., 2021), and BLOOMZ (Muennighoff et al., 2023). However, despite strong empirical gains, we still lack a clear understanding of how cross-lingual competence arises: alignment often emerges only implicitly, can be uneven across layers, and is sometimes brittle across languages and tasks (Hämmerl et al., 2024; Wang et al., 2024; Liu and Niehues,

2025; Li et al., 2024). Many existing approaches are largely engineering-driven: they often derive performance from parameter sharing and increasing data or compute, without explicitly specifying what units of meaning should be aligned.

In this work, we adopt a linguistically motivated view: a model is multilingual not just because it processes several languages in a shared parameter space, but because it represents similar meanings with comparable internal ‘senses’ across languages. Building on this view, we propose **SENSe-based Symmetric Interlingual Alignment** (SENSIA), a method for aligning sense-level representations across languages. Concretely, SENSIA starts from an English Backpack (Hewitt et al., 2023) language model, which represents each word as a mixture of latent sense vectors. It adapts this model to a target language by aligning sense mixtures and contextual states with a symmetric InfoNCE objective (Oord et al., 2018) on parallel data, while jointly training a target-side LM loss to preserve fluency.

Across benchmarks on four typologically diverse languages, SENSIA-adapted models outperform comparable dense-representation multilingual alignment methods. They are also competitive with monolingual from-scratch baselines in the same parameter regime while using 2 to 4× less target-language data. Structural ablations show that the original English local and global sense representations remain largely intact, providing evidence that SENSIA operates at the level of meaning. Scaling tests show consistent improvements for all languages and tasks as model size increases. Additional empirical ablations indicate that our training recipe is robust, with each component contributing to final performance and fluency. Finally, we find that SENSIA degrades gracefully in low-data regimes, making it attractive for low-resource languages where curated monolingual data is scarce for large-scale pretraining.

Our contributions are as follows:

*Work was done while at Mila - Quebec AI and McGill University.




Figure 1: **SENSIA overview.** (A) A Backpack represents each input token as a context-dependent mixture of K latent sense vectors, weighted by the attention-based mixture weight $\pi_{t,k}$ (Section 2). (B) The *sense loss* pools sense vectors into a sentence-level embedding using a separate, magnitude-based pooling weight $\beta_{t,k}$ (Section 2.1) — distinct from Panel A’s $\pi_{t,k}$ — before aligning sense-level representations of meaning. (C) The *context loss* aligns context-level representations via the last non-pad tokens. (D) A target-language *language modeling loss* preserves fluency while senses and contexts are aligned across languages.

- We **model multilinguality as sense adaptation**, grounding cross-lingual transfer in representation-level alignment of meaning rather than in parameter sharing alone.
- We propose a **data-efficient** sense-based alignment method that adapts a pretrained model from one language to another using only parallel data, avoiding retraining from scratch.
- We demonstrate **stronger performance** than dense-representation alignment and competitive accuracy with monolingual from-scratch baselines, together with empirical evidence of meaning preservation, design robustness, and consistent performance at scale.

Ultimately, we argue that multilinguality can be modeled *as* sense adaptation, and that SENSIA is a simple, robust, data-efficient, and linguistically-motivated operationalization of this principle. We hope this reframing encourages more linguistically motivated work on *what* to align (senses vs. parameters) and *where* to align (layers/subspaces), yielding more reliable and inclusive cross-lingual transfer.

2 Methodology

Overview. SENSIA uses the Backpack architecture (Hewitt et al., 2023) as a base for cross-

lingual adaptation. Backpack augments a GPT-2 transformer with a sense layer that represents each input token t as a mixture of K latent sense vectors $\mathbf{s}_{t,1}, \dots, \mathbf{s}_{t,K}$.¹ Given the token embedding $\mathbf{x}_t$, the sense module produces these K vectors, and an attention-based sense-weight layer assigns, independently at each position t , a distribution of importance weights $\pi_{t,k}$ over the K senses conditioned on context; there is no aggregation across positions in this mixture. The resulting context-dependent token representation is the soft mixture $\mathbf{h}_t = \sum_{k=1}^K \pi_{t,k} \mathbf{s}_{t,k}$, which replaces the standard token embedding and is consumed by the language-modeling head. Aligning these soft sense mixtures and their contextualized representations across languages allows SENSIA to transfer conceptual structure while permitting language-specific reorganization of the sense space.

2.1 Training Objectives

We perform contrastive learning on parallel sentence pairs using three complementary objectives: a **sense loss**, a **context loss**, and a **language modeling loss**. A visual sketch of how these are computed can be found in Figure 1.

¹In Backpack models, “senses” are not discrete, interpretable categories. Instead, they are continuous latent representations that capture possible variations in meaning. They provide a structured approximation of sense distinctions rather than explicit symbolic senses.

Symmetric InfoNCE. Given a batch of B normalized sentence-level vectors from the source language $\{\mathbf{u}_i^{\text{src}}\}_{i=1}^B$ and target language $\{\mathbf{u}_i^{\text{tgt}}\}_{i=1}^B$, and a temperature τ , we define bidirectional InfoNCE probabilities

$$p_{ij}(\tau) = \frac{\exp(\mathbf{u}_i^{\text{src}} \cdot \mathbf{u}_j^{\text{tgt}} / \tau)}{\sum_{k=1}^B \exp(\mathbf{u}_i^{\text{src}} \cdot \mathbf{u}_k^{\text{tgt}} / \tau)}, \quad (1)$$

$$p'_{ij}(\tau) = \frac{\exp(\mathbf{u}_i^{\text{tgt}} \cdot \mathbf{u}_j^{\text{src}} / \tau)}{\sum_{k=1}^B \exp(\mathbf{u}_i^{\text{tgt}} \cdot \mathbf{u}_k^{\text{src}} / \tau)}. \quad (2)$$

The associated symmetric InfoNCE loss is

$$\begin{aligned} & L_{\text{InfoNCE}}(\mathbf{u}^{\text{src}}, \mathbf{u}^{\text{tgt}}; \tau) \\ &= -\frac{1}{2B} \sum_{i=1}^B [\log p_{ii}(\tau) + \log p'_{ii}(\tau)] \end{aligned} \quad (3)$$

where $p_{ii}(\tau)$ and $p'_{ii}(\tau)$ denote the probabilities assigned to the true translation pair (the i -th source and target sentences).

Sense loss. This objective encourages semantically equivalent words across languages to produce similar mixtures of sense embeddings. Given a token embedding $\mathbf{x}_t$, the model computes K sense vectors

$$\mathbf{s}_{t,k} = \mathbf{E}_k^\top \mathbf{x}_t, \quad k = 1, \dots, K,$$

where $\mathbf{E}_k$ are learned projection matrices that map token embeddings into distinct sense subspaces. The magnitude of each sense vector, $a_{t,k} = \|\mathbf{s}_{t,k}\|_2$, serves as a measure of its contextual relevance and is transformed into a soft pooling weight $\beta_{t,k} = \text{softmax}_k(a_{t,k} / \tau_{\text{pool}})$, where τ_{pool} is a temperature parameter controlling the sharpness of the weighting. This pooling weight $\beta_{t,k}$ is a distinct quantity from the context-conditioned mixture weight $\pi_{t,k}$ used by the Backpack module itself (Section 2): $\beta_{t,k}$ is computed solely from sense-vector magnitude and is used only to pool sense vectors into a sentence-level sense embedding for the sense loss below, whereas $\pi_{t,k}$ is attention-based, context-conditioned, and drives the token representation consumed by the language-modeling head. After computing $\mathbf{h}_t$ through the backpack, we obtain sentence-level sense embeddings $\mathbf{z}_i^{\text{sns}}$ by mean-pooling $\mathbf{h}_t$ over all non-padding tokens in sentence i and ℓ_2 -normalizing.

For normalized source and target sense embeddings $\{\mathbf{z}_i^{\text{src,sns}}, \mathbf{z}_i^{\text{tgt,sns}}\}_{i=1}^B$, the sense alignment loss instantiates Eq. (3) with temperature τ_{sns} :

$$L_{\text{sns}} = L_{\text{InfoNCE}}(\mathbf{z}^{\text{src,sns}}, \mathbf{z}^{\text{tgt,sns}}; \tau_{\text{sns}}).$$

Context loss. This objective aligns contextualized transformer representations, encouraging words with similar meanings to appear in comparable linguistic environments. For each sentence i , we take $\mathbf{z}_i^{\text{ctx}}$ to be the final hidden state of the last non-padding token (again ℓ_2 -normalized), computed separately for source and target. Using the same symmetric InfoNCE formulation in Eq. (3) with temperature τ_{ctx} , we define

$$L_{\text{ctx}} = L_{\text{InfoNCE}}(\mathbf{z}^{\text{src,ctx}}, \mathbf{z}^{\text{tgt,ctx}}; \tau_{\text{ctx}}).$$

Language modeling loss. To improve fluency in the target language, the model minimizes a label-smoothed cross-entropy loss over target tokens:

$$L_{\text{lm}} = - \sum_t q_t \log p_\theta(y_t | x_{<t}^{\text{tgt}}),$$

where y_t is the target token at position t , $p_\theta(\cdot | x_{<t}^{\text{tgt}})$ is the model’s conditional distribution, and $q_t = (1 - \varepsilon) \delta_{y_t} + \frac{\varepsilon}{|\mathcal{V}|}$ is the smoothed target distribution with vocabulary $\mathcal{V}$ and smoothing parameter ε . Here δ_{y_t} denotes the standard Kronecker delta, i.e. the one-hot distribution equal to 1 at the true token y_t and 0 elsewhere.

Total objective. The overall training loss combines all components:

$$L_{\text{total}} = w_{\text{sns}} L_{\text{sns}} + w_{\text{ctx}} L_{\text{ctx}} + w_{\text{lm}} L_{\text{lm}}.$$

Training proceeds in three phases with distinct weight schedules: an *alignment* phase with high w_{sns} and w_{ctx} and negligible w_{lm} to establish cross-lingual correspondence; a *joint* phase in which w_{lm} is gradually increased while w_{sns} and w_{ctx} are reduced, allowing target-language fluency to develop while preserving semantic alignment; and a *polish* phase where w_{lm} dominates but small alignment weights are retained as regularizers. During the polish phase, the sense vectors and the sense-weight network are frozen to prevent representational drift, while the transformer backbone and language-modeling head continue to update. Loss weights are smoothly annealed across phases using the cosine schedule described in Appendix B.

2.2 Experimental Setup

We adapt a pretrained English Backpack model to a target language using parallel data from OPUS (Tiedemann, 2012). Adaptation is evaluated on four typologically diverse languages—Estonian, Turkish, Indonesian, and Swahili—selected for their




Figure 2: **Training dynamics on the FLORES-200 validation set.** (a) Mean target-side sense entropy decreases and stabilizes as senses specialize without collapsing. (b–c) $\text{src} \rightarrow \text{tgt}$ and $\text{tgt} \rightarrow \text{src}$ recall@1 rapidly rise and remain high, indicating a stable, aligned bilingual space during adaptation.

Benchmark	Lang	XGLM-7B5	BLOOMZ-7B1	Goldfish	GPT-2 + FT	MCL	SENSIA (Ours)
XCOPA	Estonian	0.5700	0.4800	0.5200	0.4760 _{.432, .520}	0.4760 _{.432, .520}	0.5120 _{.468, .556}
	Indonesian	0.5800	0.6700	0.5840	0.4840 _{.440, .526}	0.5120 _{.468, .554}	0.5160 _{.472, .560}
	Swahili	0.5700	0.6000	0.5500	0.5280 _{.484, .572}	0.5120 _{.468, .556}	0.5520 _{.508, .596}
	Turkish	0.5600	0.4800	0.5620	0.5100 _{.466, .554}	0.5140 _{.470, .558}	0.5540 _{.510, .598}
XStoryCloze	Indonesian	0.5963	0.8382	0.5321	0.5076 _{.486, .529}	0.5162 _{.494, .539}	0.5288 _{.506, .552}
	Swahili	0.5533	0.6843	0.5321	0.5157 _{.493, .538}	0.5407 _{.518, .563}	0.5420 _{.520, .564}
Belebele	Estonian	0.3411	0.2988	0.2933	0.2522 _{.223, .281}	0.2667 _{.238, .297}	0.2911 _{.261, .321}
	Indonesian	0.3522	0.4744	0.2800	0.2344 _{.208, .262}	0.2789 _{.250, .309}	0.2911 _{.262, .321}
	Swahili	0.3244	0.4022	0.2711	0.2489 _{.221, .278}	0.2678 _{.239, .297}	0.2844 _{.256, .314}
	Turkish	0.3522	0.3311	0.3089	0.2656 _{.237, .294}	0.2589 _{.230, .288}	0.3067 _{.278, .338}

Table 1: **Accuracy on reasoning and comprehension benchmarks.** SENSIA is competitive against from-scratch training (Goldfish), outperforms both finetuning (GPT-2 + FT) and dense-representation multilingual alignment (MCL), and approaches the performance of much larger multilingual models. Subscripts show 95% bootstrap confidence intervals ($N = 500$ for XCOPA, $N = 1,510$ for XStoryCloze, $N = 900$ for Belebele) for the three models trained under our protocol; XGLM-7B5, BLOOMZ-7B1, and Goldfish figures are sourced from prior work (Appendix E) and are not bootstrapped here.

linguistic variety and coverage in our benchmark datasets. These languages span distinct families (Finno-Ugric, Turkic, Austronesian, and Bantu) and vary widely in morphology and syntax, offering a strong test of whether sense-based alignment generalizes beyond Indo-European structure. Evaluation is then done on two angles: semantic alignment and task performance.

Semantic Alignment. We evaluate on FLORES-200 during training, measuring bidirectional recall@1 for bilingual retrieval and the mean entropy of the target-side sense weights. Entropy characterizes the specialization of the sense inventory: values near $\log_e K$ indicate diffuse, non-specialized senses, while values near zero suggest sense collapse. These metrics capture how effectively the model forms a stable, aligned bilingual space. We additionally track perplexity on the FLORES-200 devtest set as an internal metric to gauge polish

phase effectiveness².

Task Performance. Downstream performance is evaluated on XCOPA, XStoryCloze, and Belebele to probe reasoning and understanding.

Baselines. We employ three main baselines: Goldfish (Chang et al., 2024): a suite of monolingual GPT-2 models trained *from-scratch* with 1 GB of data; English GPT-2 *finetuned* with the same target-language data and number of tokens as the SENSIA-adapted models, and; Multilingual Contrastive Learning (MCL) (Li et al., 2024)³, which adapts a model by aligning transformer representations with parallel data. This last method is the closest to ours and represents a baseline difference

²Perplexity is not compared against other models as different models use different tokenizers, and thus will not provide an apples-to-apples comparison.

³We implement only the MCL phase of the AFP framework proposed by Li et al. (2024). We do not include the second phase, Cross-Lingual Instruction Following, because our models are not instruction-following LLMs. Accordingly, we refer to this baseline as MCL rather than AFP.

in *alignment target*: MCL wants to align *dense* transformer representations, while SENSIA wants to align *meaning* representations through senses. We also compare to larger multilingual models: XGLM-7B5 (Lin et al., 2021) and BLOOMZ-7B1 (Muennighoff et al., 2023), which serve as upper bounds on multilingual transfer. Hyperparameter, data-processing, and benchmarking details appear in Appendix C, D, and E.

3 Results and Discussion

SENSIA is stable and robust. During adaptation, mean target-side sense entropy decreases rapidly and stabilizes around 1.4–1.5, indicating specialization without collapse; bidirectional recall@1 on FLORES-200 converges near 0.97 for most languages, evidencing a stable bilingual mapping (Figure 2). Swahili and Indonesian reach this plateau fastest, while Estonian converges more slowly. These intrinsic signals align with the steady downstream behavior in Table 1.

SENSIA is competitive with from-scratch training, consistently beats finetuning, and narrows the gap to multilingual LMs. Against GPT-2 + FT, SENSIA is uniformly stronger across XCOPA, XStoryCloze, and Belebele (Table 1), with gains of +2.1 to +5.7 pp (mean +3.56 pp). Relative to Goldfish, SENSIA is within ± 1 pp on 7/10 cells and ahead on 4 (XCOPA-Swahili +0.2 pp; XStoryCloze-Swahili +0.99 pp; Belebele-Indonesian +1.11 pp; Belebele-Swahili +1.33 pp), with a single notable deficit on XCOPA-Indonesian (−6.8 pp). While absolute scores trail larger multilingual LMs, SENSIA narrows the gap in some settings (e.g., XCOPA-Turkish is within 0.6 pp of XGLM-7B5), highlighting the efficiency of sense alignment relative to broad multilingual pretraining. This is despite a substantial scale disparity: the Backpack backbone used by SENSIA has 124M parameters (Table 10, Appendix A), roughly 60 $\times$ smaller than XGLM-7B5 (7.5B) and BLOOMZ-7B1 (7.1B). Gains over both MCL and GPT-2 + FT are consistent rather than incidental: SENSIA outperforms each baseline in all 10 evaluated benchmark–language cells in Table 1, which is significant under a sign test in both comparisons ($p = 0.001$).

Sense-based adaptation is stronger than dense-representation alignment. Across the 10 benchmark–language combinations where it is

Lang	Total Tok.	Data (MB)	vs. GF
Estonian	83M	246	3.9 $\times$
Indonesian	149M	509	2.0 $\times$
Swahili	135M	426	2.3 $\times$
Turkish	98M	283	3.5 $\times$

Table 2: **Training data efficiency** of SENSIA relative to Goldfish-1GB (GF). SENSIA achieves comparable performance with 2–4 $\times$ less target language data. Tokens are counted based on English GPT-2 tokenizer. MB is approximated based on UTF-8 bytes per token.

evaluated, MCL improves over GPT-2 + FT by roughly +1.2 pp on average, and is particularly helpful on XStoryCloze and Belebele. However, SENSIA consistently matches or outperforms MCL in all cells, with an average gain of about +2.3 pp. The largest improvements appear on XCOPA–Swahili and XCOPA–Turkish (up to +4–5 pp), and on Belebele–Turkish (+4.8 pp). In a few settings (e.g., XCOPA–Indonesian and XStoryCloze–Swahili), the gap nearly closes (≤ 0.4 pp), suggesting that LM-level contrastive alignment can recover part of the benefit, but sense-based alignment never underperforms it and typically yields a stronger model.

SENSIA is more data efficient. SENSIA achieves competitive accuracy with substantially smaller target-language budgets than Goldfish: 2.0 $\times$ to 3.9 $\times$ less data across languages (Table 2), while maintaining parity or better on many rows of Table 1. Taken together, these results show that sense alignment preserves downstream quality while reducing the target-language data requirement by roughly 2–4 $\times$. We note that this comparison is stated in raw byte volume, which likely understates SENSIA’s effective supervision: each parallel sentence pair supplies an explicit bilingual training signal (both fluency and cross-lingual correspondence), whereas Goldfish’s monolingual bytes supply fluency alone. SENSIA’s total supervision cost also includes reuse of an already-pretrained English backbone, avoiding the full from-scratch pretraining compute that Goldfish requires. We therefore view total supervision cost, rather than raw target-language data volume alone, as the more accurate basis for comparison, particularly for languages where parallel data is available (e.g., via OPUS) but a full monolingual pretraining pipeline is not.

4 Sense Ablations

Is sense geometry truly preserved? We ask whether SENSIA truly adapts senses or merely aligns words. Direct word-to-word comparison is unreliable under the shared English BPE tokenizer, where a single target-language word can span many subwords. Instead, we test whether the English sense space is preserved after adaptation, both locally (within a word) and globally (across the vocabulary).

We therefore conduct two complementary analyses: *Sense Topology analysis*, which measures whether each word’s local sense cluster maintains its internal geometry, and *Procrustes Analysis*, which tests whether a single orthogonal transformation can align the global target sense manifold to English. Together, these analyses probe whether adaptation reorganizes meaning rather than memorizing lexical correspondences. Theoretical and practical implementation details, including sampling and statistics, are given in Appendix G.

Sense Topology Analysis. (Figure 3) For each aligned English–target word pair, we ask whether its senses are arranged in a similar way across languages. Rather than requiring a one-to-one correspondence between individual senses, we focus on the *pattern* of which senses are relatively close to or far from each other. Concretely, we compare the pairwise cosine-similarity structure of the K sense vectors in each language: we construct cosine Gram matrices for English and the target language, then compute the Spearman correlation ρ between their upper-triangular entries. Moderate ρ values indicate that the local sense geometry is preserved but not trivially copied: extremely high values would suggest that the model is effectively memorizing the English configuration, whereas very low values would indicate substantial semantic drift.

Procrustes Analysis. (Figure 4) To assess whether global sense organization is preserved, we view the English and target-language sense mixtures as two embedding spaces and ask whether one can be mapped onto the other by a rigid transformation such as a rotation. For each aligned word pair, we compute a mixture-weighted sense embedding and then learn an orthogonal matrix that best maps target-language embeddings onto English via an orthogonal Procrustes problem. We finally measure cosine similarity between the aligned embeddings. High post-alignment cosine indicates that, up to




Figure 3: **Sense Topology.** For each aligned word pair, we ask if its senses relate to each other similarly across languages, rather than requiring one-to-one matches between individual senses. We compute K sense vectors for English and the target language, pooling if a word is split into subwords. We then build cosine gram matrices over senses in each language and compare their upper triangles with a Spearman correlation ρ . A high ρ means that the pattern of inter-sense relationships is preserved even if the exact locations of individual senses differ, whereas exact one-to-one matches indicates simple memorization of the English space.

	SENSIA $\bar{\rho}$	Control $\bar{\rho}$	$\Delta\rho$
est	0.27	0.17	0.10
ind	0.30	0.21	0.09
swl	0.26	0.17	0.09
tur	0.25	0.16	0.09

Table 3: **Sense Topology results.** Moderate correlation indicates that the model preserves local sense geometry relative to the control setup, showing that SENSIA keeps English meaning intact while accounting for language-specific variations.

a rotation, the target manifold has essentially the same overall shape as the English sense space.

4.1 Results

Sense-aware adaptation preserves both local and global structure. As a control, we finetune an English Backpack to the target language using only the LM loss ($w_{\text{sns}} = w_{\text{ctx}} = 0$, $w_{\text{lm}} = 1.0$), which yields fluent target-language text but no explicit sense or context alignment.

Compared to this control, SENSIA consistently achieves higher local topology correlations ($\rho \approx 0.25$ – 0.30 , $\Delta\rho \approx 0.09$ – 0.10 ; Table 3) and higher global Procrustes cosine similarities (0.35 – 0.44 , $\Delta \cos \approx 0.08$ – 0.11 ; Table 4).

Analytic languages such as Indonesian show slightly tighter alignment than agglutinative ones like Turkish, consistent with richer morphology. Overall, local sense clusters remain coherent, and the global manifold is well aligned with English up to an orthogonal mapping, supporting the view that




Figure 4: **Procrustes Analysis.** Target-language sense embeddings may be rotated relative to English even if their structure is identical. We therefore learn an orthogonal transform Q^* that best aligns T to E , and measure cosine similarity between the aligned spaces to assess how similar their global geometry is.

	SENSIA $\overline{\text{cos}}$	Control $\overline{\text{cos}}$	Δcos
est	0.36	0.29	0.08
ind	0.44	0.33	0.11
swh	0.42	0.31	0.11
tur	0.35	0.26	0.09

Table 4: **Procrustes Analysis results.** Moderate cosine similarity indicates that, up to an orthogonal mapping, the target space retains similar global structure to the original English sense space, showing that SENSIA successfully preserves learned English meaning at a global level.

SENSIA reorganizes meaning rather than destroying the original sense geometry.

5 Empirical Ablations

Where does SENSIA’s performance come from? We probe SENSIA’s design with four empirical ablations: (i) removing the sense-mixture mechanism, (ii) ablating loss terms and training phases, (iii) scaling the Backpack backbone, and (iv) varying the amount of parallel data, plus a stress test on Chinese with an English tokenizer. These experiments pinpoint which components drive fluency and reasoning, and how robust the method is across scale and data regimes. Practical and implementation details can be found in Appendix G.

Sense mixtures are necessary for fluent prediction. To test whether the model truly uses its sense mixtures at inference, we ablate the mixture weights $\pi_{t,k}$ while keeping all other parameters fixed. We compare three settings on FLORES devtest: the full learned mixture, a top-1 variant

Language	Full CE	Top-1 CE	Uniform CE
Swahili	2.98	8.77	10.10
Estonian	3.28	8.95	10.16
Indonesian	2.88	8.58	10.10
Turkish	2.82	8.83	10.13

Table 5: **Sense Mixture Ablation.** Cross-entropy on FLORES devtest (lower is better). Results indicate that the model’s performance relies on successful mixing and non-collapse of senses.

Variant	XCOPA	XStory	Belebele	PPL
No Context loss	0.498	0.523	0.278	13.83
No Sense loss	0.484	0.524	0.252	13.64
No LM loss	0.500	0.484	0.232	7.7e7
No Align phase	0.512	0.453	0.258	13.66
No Joint phase	0.498	0.456	0.278	14.57
No Polish phase	0.510	0.465	0.269	14.55
Full model	0.516	0.529	0.291	10.58

Table 6: **Loss and phase ablations** (Indonesian target). Results indicate that all phases and losses play a role in the final performance of the adapted model.

that renormalizes only the most probable sense per token, and a uniform mixture over all K senses (Table 5). Across four target languages, moving from the full mixture to top-1 increases token-level cross-entropy by roughly +6, and uniform mixtures degrade even further. This indicates that SENSIA’s fluency and predictive quality depend critically on softly combining multiple senses per token, and that this dependency is stable across typologically diverse languages.

Losses and training phases play complementary roles. We ablate each of the three objectives (Sense, Context, LM) and each phase in the Alignment–Joint–Polish schedule using Indonesian as the target language (Table 6). Removing the Context loss mostly hurts accuracy while leaving perplexity similar, showing that contextual alignment sharpens semantic decisions more than fluency. Dropping the Sense loss degrades comprehension benchmarks while leaving fluency relatively intact, indicating that sense-level anchoring primarily supports semantics. Eliminating the LM loss leads to extreme perplexity and lower accuracy, confirming that the generative objective is indispensable for stable training and coherent text. In the curriculum, skipping the Alignment phase causes moderate drops that can be partially recovered later, whereas removing the Joint phase yields the largest overall degradation, showing that jointly optimizing contrastive and generative signals is cru-

cial. The Polish phase mainly refines fluency and phrasing, providing smaller but consistent gains.

Task	Lang	Small	Medium	Large
XCOPA	<i>ind</i>	0.4740	0.4840	0.5140
	<i>swh</i>	0.5360	0.5440	0.5480
XStoryCloze	<i>ind</i>	0.5076	0.5156	0.5189
	<i>swh</i>	0.5107	0.5281	0.5301
Belebele	<i>ind</i>	0.2611	0.2689	0.2733
	<i>swh</i>	0.2567	0.2778	0.2911
FLORES	<i>ind</i>	10.58	9.90	8.18
	<i>swh</i>	11.76	11.68	11.52

Table 7: **Scaling test.** Results show predictable monotonic scaling, indicating that SENSIA benefits from scale.

Scaling improves reasoning and fluency in a predictable way. We pretrain Backpack models with GPT-2 small, medium (345M), and large (762M) backbones, following the same OpenWebText setup and Chinchilla-style token budgets (Hoffmann et al., 2022), then adapt each to Indonesian and Swahili under the same SENSIA recipe (details in Appendix G). Across both languages, moving from Small to Large increases macro-average accuracy over XCOPA, XStoryCloze, and Belebele by about +2.2 points and reduces FLORES perplexity (Table 7). These monotonic gains indicate that sense-based adaptation benefits from scale in a manner consistent with standard LM scaling laws.

SENSIA degrades gracefully with $20\times$ less bi-text. To approximate low-resource conditions, we downsample the Indonesian and Swahili parallel corpora used for adaptation to 100k sentence pairs, keeping the architecture, schedule, and loss weights fixed and retraining from the same English initialization. Across XCOPA, XStoryCloze, and Belebele, accuracy drops by roughly 2.5 points for Indonesian and 3.3 points for Swahili (Table 8), with macro accuracy decreasing by about 3 points overall despite a $20\times$ reduction in bitext. This indicates that sense-based adaptation remains competitive in realistically low-data regimes.

A hostile tokenizer limits sense adaptation on non-Latin scripts. We adapt the same English Backpack to Chinese using English–Chinese bi-text from OPUS while keeping the original English GPT-2 BPE tokenizer unchanged, so Chinese is represented as long, poorly aligned subword sequences. Under this setup, SENSIA trails the

Task	Lang	2M Pairs	100k Pairs
XCOPA	<i>ind</i>	0.5160	0.4880
	<i>swh</i>	0.5520	0.4940
XStoryCloze	<i>ind</i>	0.5288	0.5122
	<i>swh</i>	0.5420	0.5107
Belebele	<i>ind</i>	0.2911	0.2600
	<i>swh</i>	0.2844	0.2733

Table 8: **Low-resource adaptation test.** Results show that the method is reasonably resistant to low-data regimes, and thus is useful for low-resource settings.

Task	Goldfish	SENSIA (Zho)
XCOPA	0.5420	0.4800
XStoryCloze	0.4944	0.5069
Belebele	0.3144	0.2611

Table 9: **Hostile tokenizer test.** Results show the method’s relative weakness when adapting to a language that uses non-Latin scripts.

monolingual Goldfish baseline by 6.2 points on XCOPA and 5.3 points on Belebele, while slightly surpassing Goldfish on XStoryCloze (+1.3 points; Table 9). Averaged across tasks, the adapted model remains 3–4 points behind Goldfish, suggesting that when tokenization is badly misaligned, sense adaptation cannot fully compensate. We therefore view this experiment as a stress test highlighting tokenizer limitations rather than evidence that SENSIA currently supports non-Latin scripts.

6 Related Work

LM-level cross-lingual alignment in causal LMs.

A growing line of work targets the *internal states* of decoder-only LMs to improve cross-lingual transfer. Liu and Niehues (2025) show that aligning *middle layers* during task fine-tuning consistently boosts transfer, suggesting that layer choice is critical for multilingual sharing. Complementarily, Li et al. (2024) align *sentence-level* hidden states in generative LMs via multilingual contrastive learning, while Wang et al. (2024) probe training checkpoints of BLOOM and relate the emergence of cross-lingual alignment to neuron overlap and downstream performance. **SENSIA differs in both target and granularity:** instead of aligning generic hidden states, we align a *sense-decomposed generator*—optimizing symmetric InfoNCE on *latent sense mixtures* and a *context* vector, while pre-

serving target-side fluency—building directly on Backpack’s sense factorization (Hewitt et al., 2023). This design is *linguistically motivated*: the unit of transfer is meaning (senses), not parameters or sub-word overlap.

Activation-space and post-hoc interventions. Beyond gradient tuning, several works steer or map activation spaces to enhance multilinguality in LLMs. Sundar et al. (2025) show that neuron-level *model interventions* (“finding experts”) can measurably increase cross-lingual alignment in multilingual LLMs; Wang et al. (2025) propose *INCLINE*, learning linear alignment matrices from parallel data and applying them at inference; and Bandarkar et al. (2025) demonstrate *layer swapping* as a post-hoc merge that composes language and task experts for zero-shot cross-lingual transfer. Surveys synthesize broader evidence and cautions about what and where to align (Hämmerl et al., 2024). **In contrast**, we restructure the model’s *semantic basis* by aligning *latent senses*, then show that local topology and a near-orthogonal global mapping are preserved after adaptation—empirical evidence for a *linguistically motivated* view of multilinguality as *sense adaptation* rather than surface-form or parameter sharing.

Sense inventories and lexical semantics. A separate line of work performs cross-lingual sense alignment over discrete, human-defined sense inventories: Navigli and Ponzetto (2012) construct BabelNet by integrating WordNet and Wikipedia into a wide-coverage multilingual semantic network, while Scarlini et al. (2020) and Loureiro and Jorge (2019) (SensEmBERT and LMMS, respectively) produce sense embeddings tied to such inventories for word sense disambiguation. SENSIA instead operates on Backpack’s continuous latent senses, learned from data without reference to an explicit sense inventory (see footnote in Section 2). This makes SENSIA applicable in low-resource settings where resources like WordNet or BabelNet coverage are unavailable or incomplete, but it also means our use of “sense” is a structured approximation rather than a claim of equivalence to the lexicographic sense distinctions studied in this literature.

7 Conclusion

We approach multilinguality as sense adaptation: aligning latent meaning representations across lan-

guages rather than merely sharing parameters. We operationalize this with SENSIA, which yields competitive performance against from-scratch, alignment, and finetuning methods. We additionally provide systemic ablations that show the method’s soundness and robustness, as well as tests of scale and tokenizer limitations. Looking forward, we aim to scale to larger backbones and explore ways to adapt more smoothly to non-Latin text. More broadly, we view sense adaptation as a simple, effective, and linguistically motivated path to multilingual fluency inside generative LMs.

Limitations

While our reframing of multilinguality as *sense adaptation* yields competitive accuracy and strong data efficiency, several limitations remain.

Tokenization and script coverage. Our pipeline assumes a shared GPT-2 BPE tokenizer trained on English and, in practice, targets Latin-script languages. All four evaluation languages (Estonian, Indonesian, Swahili, Turkish) use Latin script, and both data accounting and modeling rely on the English GPT-2 tokenizer (e.g., shared vocabulary of 50,264 tokens; data-efficiency counts based on the English tokenizer).⁴ This assumption restricts adaptation to languages that are not severely over-segmented by the English BPE; applying the method to non-Latin scripts (e.g., Chinese, Arabic, Devanagari) is understudied and out of scope in its current form. Future work will investigate adapting or retraining the tokenizer (e.g., Unicode/byte-level tokenization or lightweight target-specific merges) while preserving the learned sense geometry.

Nature of supervision (OPUS) and scaling limits. Our adaptation is driven by sentence-parallel bitext filtered from OPUS. These pairs are typically short and literal with limited discourse continuity and world knowledge, which strengthens local semantic alignment and fluency but provides little grounding for discourse-level reasoning. In the scaling study (Table 7), XSTORYCLOZE shows smaller Small→Large gains than XCOPA on average (+1.5 vs. +2.6 points), and BELEBELE is intermediate (+2.3 points overall, with Swahili much stronger than Indonesian), consistent with the view that narrative plausibility and reading comprehension require information

⁴See Table 7 for scaling results. Appendix A notes the shared vocabulary size.

beyond sentence-level semantics. We hypothesize that OPUS’s length and domain characteristics cap improvements on these tasks. Future work will explore longer, document-level parallel corpora (e.g., paragraph/episode-aligned subtitles, newswire, TED talks) and grounded parallel sources to inject discourse and world knowledge.

Parallel-data dependence. SENSIA presumes adequate, high-quality alignments for the target language. This constraint may limit effectiveness in low-resource settings where bitext is scarce or noisy, or where domain mismatch is severe. Extending the approach to weakly aligned or mined corpora—with alignment uncertainty modeled in the loss—is a natural next step.

Model size regime and instruction-level transfer. Our experiments intentionally do not test scaling to billion-parameter models; the goal is to validate the representation-level framing, not to chase absolute SOTA. It remains an open question whether sense adaptation scales to 7B+ backbones and whether it can transfer task performance when supervision comes from a *parallel instruction* corpus rather than sentence-parallel text. We leave both questions to future work.

On Chasing SOTA Results. We do *not* claim state-of-the-art results or parity with large multilingual LLMs—and that is **not the objective of this paper**. Our aim is to introduce a linguistically motivated way of thinking about multilinguality, demonstrate that the reframing works (SENSIA), show it is competitive and data-efficient, and establish robustness through analyses and ablations (sense topology, Procrustes, mixture and loss/phase ablations). The next step is to evaluate the approach at LLM scale and test whether sense adaptation can transfer instruction-following and downstream tasks when trained with parallel instruction data.

Ethical Considerations

Data. Our adaptation relies on parallel corpora from OPUS, whose sub-corpora vary in license, domain, and quality; we apply length/similarity filtering but do not add PII scrubbing or content moderation. Because we align target languages to an English-trained model, English-centric biases and OPUS domain skews may carry over, especially for dialects or topics under-represented in bitext.

Risk. Supervision is sentence-parallel and short-form, which chiefly improves *semantic alignment and fluency* and less so discourse grounding; correspondingly, our fixed-protocol scaling gains are larger on XCOPA than on XSTORYCLOZE or BELEBELE (+2.6 pp vs. +1.5 and +2.3 pp Small→Large), consistent with those tasks requiring narrative plausibility and comprehension beyond sentence-level semantics (Table 7). We will investigate longer, document-level parallel corpora and grounded sources. Our scope is GPT-2-scale models; we do not claim SOTA or parity with large multilingual LLMs, and caution that even small adapted generators can be misused (spam/misinformation). We therefore emphasize research-only release practices and recommend additional safeguards for any deployment. We track tokens and follow compute-aware training to limit environmental impact, and we plan to study extension to 7B+ backbones and *parallel-instruction* supervision to assess transfer and risk at LLM scale.

Use of AI. We acknowledge the use of Stanford’s Agentic AI Reviewer⁵ and OpenAI’s ChatGPT 5.1-Pro to assist with reviewing and proofreading the final manuscript.

Acknowledgments

This research was enabled in part by equipment and support provided by Mila - Quebec AI Institute and the Digital Research Alliance of Canada (alliancecan.ca). This research was supported in part by the Natural Sciences and Engineering Research Council (NSERC) of Canada. We are grateful for the support from IVADO and the Canada First Research Excellence Fund.

References

- Lucas Bandarkar, Benjamin Muller, Pritish Yuvraj, Rui Hou, Nayan Singhal, Hongjiang Lv, and Bing Liu. 2025. [Layer swapping for zero-shot cross-lingual transfer in large language models](#). In *Proceedings of the International Conference on Learning Representations (ICLR)*. Published as a conference paper at ICLR 2025.
- Tyler A Chang, Catherine Arnett, Zhuowen Tu, and Benjamin K Bergen. 2024. Goldfish: Monolingual language models for 350 languages. *arXiv preprint arXiv:2408.10441*.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco

⁵<https://paperreview.ai/>

- Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Marta R. Costa-Jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, and others. 2022. [No language left behind: Scaling human-centered machine translation](#). *arXiv preprint arXiv:2207.04672*.
- Fangxiaoyu Feng, Yinfei Yang, Daniel Cer, Naveen Ariavazhagan, and Wei Wang. 2022. [Language-agnostic BERT sentence embedding](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 878–891, Dublin, Ireland. Association for Computational Linguistics.
- Francisco Guzmán, Peng-Jen Chen, Myle Ott, Juan Pino, Guillaume Lample, Philipp Koehn, Vishrav Chaudhary, and Marc’Aurelio Ranzato. 2019. [The FLORES evaluation datasets for low-resource machine translation: Nepali–English and Sinhala–English](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6098–6111, Hong Kong, China. Association for Computational Linguistics.
- Katharina Hämmerl, Jindřich Libovický, and Alexander Fraser. 2024. [Understanding cross-lingual alignment—a survey](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 10922–10943, Bangkok, Thailand. Association for Computational Linguistics.
- John Hewitt, John Thickstun, Christopher Manning, and Percy Liang. 2023. [Backpack language models](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9103–9125, Toronto, Canada. Association for Computational Linguistics.
- Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, and 1 others. 2022. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*.
- Masoud Jalili Sabet, Philipp Dufter, François Yvon, and Hinrich Schütze. 2020. [Simalign: High quality word alignments without parallel training data using static and contextualized embeddings](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1627–1643, Online. Association for Computational Linguistics.
- Chong Li, Shaonan Wang, Jiajun Zhang, and Chengqing Zong. 2024. [Improving in-context learning of multilingual generative language models with cross-lingual alignment](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 8058–8076, Mexico City, Mexico. Association for Computational Linguistics.
- Xi Victoria Lin, Todor Mihaylov, Mikel Artetxe, Tianlu Wang, Shuohui Chen, and others. 2021. [Few-shot learning with multilingual language models](#). *arXiv preprint arXiv:2112.10668*.
- Danni Liu and Jan Niehues. 2025. [Middle-layer representation alignment for cross-lingual transfer in fine-tuned llms](#). *Preprint*, arXiv:2502.14830. ACL 2025.
- Daniel Loureiro and Alípio Jorge. 2019. [Language modelling makes sense: Propagating representations through WordNet for full-coverage word sense disambiguation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5682–5691, Florence, Italy. Association for Computational Linguistics.
- Niklas Muennighoff, Thomas Wang, Lintang Sutawika, Adam Roberts, Stella Biderman, Teven Le Scao, M Saiful Bari, Sheng Shen, Zheng Xin Yong, Hailey Schoelkopf, Xiangru Tang, Dragomir Radev, Alham Fikri Aji, Khalid Almubarak, Samuel Albanie, Zaid Alyafeai, Albert Webson, Edward Raff, and Colin Raffel. 2023. [Crosslingual generalization through multitask finetuning](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15991–16111, Toronto, Canada. Association for Computational Linguistics.
- Roberto Navigli and Simone Paolo Ponzetto. 2012. [Babelnet: The automatic construction, evaluation and application of a wide-coverage multilingual semantic network](#). *Artificial Intelligence*, 193:217–250.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Bianca Scarlini, Tommaso Pasini, and Roberto Navigli. 2020. [SenseBERT: Context-enhanced sense embeddings for multilingual word sense disambiguation](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8758–8765.
- Anirudh Sundar, Sinead Williamson, Katherine Metcalf, Barry-John Theobald, Skyler Seto, and Masha Fedzechkina. 2025. [Steering into new embedding spaces: Analyzing cross-lingual alignment induced by model interventions in multilingual language models](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2375–2401, Vienna, Austria. Association for Computational Linguistics.

Jörg Tiedemann. 2012. Parallel data, tools and interfaces in OPUS. In *Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC’12)*, Istanbul, Turkey. European Language Resources Association (ELRA).

Hetong Wang, Pasquale Minervini, and Edoardo Ponti. 2024. [Probing the emergence of cross-lingual alignment during LLM training](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 12159–12173, Bangkok, Thailand. Association for Computational Linguistics.

Weixuan Wang, Minghao Wu, Barry Haddow, and Alexandra Birch. 2025. [Bridging the language gaps in large language models with inference-time cross-lingual intervention](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5418–5433, Vienna, Austria. Association for Computational Linguistics.

Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. [mt5: A massively multilingual pre-trained text-to-text transformer](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online. Association for Computational Linguistics.

A Base Backpack Model Configurations

The base English Backpack models were initialized with GPT-2 backbones of varying sizes, following the configurations summarized in Table 10.

All models use 16 sense embeddings per token and a shared vocabulary of 50,264 tokens. Hyperparameters such as dropout and layer normalization follow the original GPT-2 settings. Each variant was trained on OpenWebText prior to cross-lingual adaptation.

B Scheduler Details

Loss weights and InfoNCE temperatures are controlled by a cosine-based interpolation schedule parameterized by normalized training progress $p \in [0, 1]$. Let a denote the proportion of steps allocated to the alignment phase and z the point at which the polish phase begins. Each weight w and temperature τ is updated as a smooth function of p according to

$$\text{interp}(a, b, t) = b + \frac{1}{2}(a-b)[1 + \cos(\pi t)], t \in [0, 1].$$

Alignment phase ($p \leq a$).

$$w_{\text{sns}}, w_{\text{ctx}} \text{ high}, \quad w_{\text{lm}} \approx 0,$$

with weights interpolated from their alignment values toward mid-phase values. Temperatures τ_{sns} and τ_{ctx} decay gradually from their initial settings.

Joint phase ($a < p < z$).

$$w_{\text{sns}}, w_{\text{ctx}} \text{ decreasing}, \quad w_{\text{lm}} \text{ increasing},$$

following cosine interpolation between mid- and tail-phase targets. Temperatures continue to decrease until reaching their final values.

Polish phase ($p \geq z$).

$$w_{\text{sns}}, w_{\text{ctx}} \text{ low}, \quad w_{\text{lm}} \text{ high}.$$

At this stage, training proceeds with high emphasis on the language modeling objective to improve fluency in the target language.

This schedule provides a smooth curriculum across objectives, avoiding discontinuities in optimization dynamics while emphasizing different learning signals at appropriate stages of adaptation.

We initially experiment with 20%-50%-30% alignment/joint/polish phase boundaries, but have found that the LM loss still reduces considerably. We then shifted to a 20%-30%-50% schedule and have kept this for all our experiments to maintain consistency.

C SENSIA Adaptation Setup

We adapt the pretrained English Backpack model (GPT-2 SMALL backbone) to each target language using parallel data from the OPUS corpus. Training is performed for 150K steps with a learning rate of 5×10^{-5} , batch size of 64, warmup ratio of 0.1, gradient clipping at 1.0, and label smoothing of 0.05. All runs use mixed-precision (FP16) optimization and the same random seed (1234).

Each adaptation follows the three-phase schedule described in Section 2.1: an *alignment* phase (20% of training steps), a *joint* phase (30%), and a *polish* phase (50%). Loss weights are progressively annealed between $(w_{\text{sns}}, w_{\text{ctx}}, w_{\text{lm}}) = (0.54, 0.44, 0.02)$ in the alignment phase, $(0.40, 0.40, 0.20)$ in the joint phase, and $(0.15, 0.15, 0.70)$ in the polish phase. Both sense and context temperatures are fixed at $\tau_{\text{sns}} = 0.05$ and $\tau_{\text{ctx}} = 0.07$, and the sense pool temperature is 0.7. During the polish phase, the sense vectors and sense-weighting network are frozen to prevent representational drift while the transformer and LM head continue to update.

Model	Layers	Hidden	Heads	Senses	Positions	Params (M)
<i>Small</i>	12	768	12	16	512	124
<i>Medium</i>	24	1024	16	16	512	345
<i>Large</i>	36	1280	20	16	512	762

Table 10: Configuration of base English Backpack models used for cross-lingual adaptation. Each model employs GELU activations, dropout of 0.1 across attention, residual, and embedding layers, and a sense intermediate scaling factor of 4.

Evaluation is conducted every 2K steps on the FLORES-200 dev set, tracking recall@1, perplexity, and sense-weight entropy. All hyperparameters are held constant across Estonian, Indonesian, Swahili, and Turkish, differing only in dataset configuration (eng-tgt) and OPUS data splits.

D Dataset Construction

Parallel training data are built automatically using a unified pipeline that filters and merges sentence pairs from multiple OPUS corpora. For each target language, we run the same data-building script with language-specific corpus selections. The pipeline downloads OPUS data, removes duplicates, applies sentence-length and ratio filters, computes sentence similarity using LaBSE (Feng et al., 2022) embeddings, and outputs a final merged dataset capped at two million aligned pairs. The pipeline uses OpusTools⁶ to retrieve data from online sources.

Sentence pairs are retained only if they satisfy a length ratio constraint (typically $0.6 < \frac{|x_{\text{src}}|}{|x_{\text{tgt}}|} < 1.7$) and exceed a LaBSE cosine similarity threshold (≥ 0.80 – 0.82 depending on the language). Each corpus contributes up to one million pairs before sampling to reach the target size of two million examples. The process automatically cleans temporary files after successful export and logs the filtering statistics for reproducibility.

Corpora used to build each language pair include:

- **Estonian (en-et):** Tatoeba, WikiMatrix, CCAligned, OpenSubtitles, TED2020, GlobalVoices.
- **Indonesian (en-id):** JW300, WikiMatrix, TED2020, GlobalVoices, QED, CCAligned.
- **Swahili (en-sw):** JW300, Tatoeba, GlobalVoices, QED, WikiMatrix, CCAligned, TED2013, CCMatrix.

- **Turkish (en-tr):** SETimes, OpenSubtitles, WikiMatrix, TED2020, GlobalVoices, CCAligned.

All datasets are built with a batch size of 2048, a maximum of one to three million pairs per corpus before filtering, and final sampling to two million sentence pairs per language. Length ratio, similarity thresholds, and corpus composition are held constant across languages except for Swahili, which uses a wider length ratio range (0.6–1.7) and slightly lower similarity cutoff (0.80) to accommodate noisier data sources.

E Benchmarking Implementation Details

Evaluation procedure. For all benchmarks, model predictions are scored using token-level log-likelihoods from the language modeling head. Given a context C and a candidate continuation A , we compute its conditional probability $p(A | C) = \exp\left(\frac{1}{|A|} \sum_{t=1}^{|A|} \log p(a_t | a_{<t}, C)\right)$, where log-probabilities are averaged over all tokens.

Following prior work, two complementary scoring schemes are used:

1. a conditional likelihood (*cond*) method, which evaluates $p(A | C)$ directly, and
2. a pointwise mutual information (PMI) formulation, $p(A | C) - p(A)$, which corrects for unconditional continuation bias.

For each example, the model selects the candidate with the highest final score as the prediction. Accuracy is the proportion of correctly selected completions, and perplexity (PPL) is the exponentiated negative mean log-likelihood per token.

Scoring parameters. We perform a grid search for our models over interpolation weights λ and α that combine PMI and conditional scores into a unified objective:

$$s(A | C) = (1 - \lambda) \text{PMI}(A, C) + \lambda \log p(A | C) + \alpha \text{len}(A),$$

⁶<https://github.com/Helsinki-NLP/OpusTools>

Language	XCOPA		XStoryCloze		Belebele	
	λ	α	λ	α	λ	α
Estonian	<i>cond</i>	0.10	—	—	1.00	0.20
Indonesian	0.50	0.10	<i>cond</i>	1.00	1.00	0.20
Swahili	0.60	0.70	<i>cond</i>	0.70	0.90	0.90
Turkish	0.40	0.40	—	—	0.40	0.50

Table 11: Optimal λ and α parameters per language and task. “*cond*” indicates the use of conditional scoring only (no PMI term).

where the final term penalizes length bias in generation. Grid search was performed on the XCOPA validation set over $\lambda \in \{0.0, 0.1, \dots, 1.0\}$ and $\alpha \in \{0.0, 0.1, \dots, 1.0\}$, selecting parameters that maximized validation accuracy per language and task. The optimal parameters used for all reported results are summarized in Table 11.

Baseline results. Scores for larger multilingual baselines and previously published models are sourced directly from prior work for transparency and consistency:

- **XCOPA:** BLOOMZ-7B1 and XGLM-7B5 from the BLOOMZ paper; Goldfish from the Goldfish paper.
- **XStoryCloze:** BLOOMZ-7B1 and XGLM-7B5 from the BLOOMZ paper; Goldfish from the Goldfish paper.
- **Belebele:** XGLM-7B5 and Goldfish from the Goldfish paper. BLOOMZ-7B1 results were obtained by running the released Goldfish evaluation scripts on BLOOMZ to ensure consistent scoring across models.

We obtain GPT-2 + FT and MCL scores via our own implementation of these methods. Both GPT-2 + FT and MCL use the same data and token budget as SENSIA-adapted models during training to keep comparisons data-matched. During decoding, we also use the same (optimal) hyperparameters used by the SENSIA-adapted models to ensure comparability.

Reproducibility. All SENSIA evaluations were run with fixed random seeds and deterministic decoding order. Scoring outputs, predictions, and selected hyperparameters are included with the released evaluation materials to facilitate replication.

F Generation Examples

We test with short translated prompts (for example, “*President Barack Obama said*”) to qualitatively assess fluency. The adapted model produces fluent, well-structured continuations in all four target languages:

- **Estonian:** President Barack Obama ütles, et ta on kõige kallim, kes on kunagi olnud.
“*President Barack Obama said he is the most expensive person who has ever lived.*”
- **Indonesian:** Presiden Barack Obama mengatakan bahwa kami tidak akan mengambil keuntungan dari kemampuan kami untuk mengumpulkan informasi...
“*President Barack Obama said that we will not take advantage of our ability to collect information...*”
- **Swahili:** Rais Barack Obama alisema kuwa Marekani inaamini kuwepo kwa makubaliano ya Paris...
“*President Barack Obama said the United States believes in the Paris Agreement on climate change.*”
- **Turkish:** Başkan Barack Obama şöyle dedi: Bunu yapmak için çok çalıştım.
“*President Barack Obama said: I worked very hard to do this.*”

G Additional Analysis and Ablation Details

In this section we provide the exact formulations and practical protocols for the analyses and empirical ablations in Sections 4 and 5. We follow the notation from the main methodology section: each token is represented by a mixture of K latent sense vectors $s_{t,1}, \dots, s_{t,K}$ with mixture weights $\pi_{t,k}$, and sentence-level vectors are ℓ_2 -normalized before contrastive training.

G.1 Sense-Space Analyses (Section 4)

Data and Languages. All analyses use the FLORES-200 devtest split (Guzmán et al., 2019; Costa-Jussà et al., 2022), which provides balanced English–target pairs for each evaluation language: Indonesian (ind), Swahili (swh), Estonian (est), and Turkish (tur). Each English–target pair was drawn from the same dataset used for tracking recall@1 and perplexity during training, ensuring

consistency between intrinsic and extrinsic evaluations.

Word Alignment. Word-level alignments were obtained using SIMALIGN (Jalili Sabet et al., 2020) with the multilingual MINILM encoder in *itermax* mode, which iteratively refines bidirectional attention for high-precision alignments. We retained only one-to-one alignments to ensure semantic comparability between tokens. Multiword expressions and one-to-many alignments were excluded, as were rare or function words that did not appear in both languages.

Sense Extraction. For each aligned word pair $(w_{\text{en}}, w_{\text{tgt}})$, we extracted the model’s pre-context sense embeddings from the Backpack sense module. When a target-language word was tokenized into multiple subwords under the shared BPE vocabulary, we pooled the corresponding sense vectors by averaging their activations before normalization. To ensure stable estimates, we filtered to include only content words (nouns, verbs, adjectives) appearing at least five times in the corpus. All sense vectors were ℓ_2 -normalized and mean-centered before computing correlations or alignments.

Sense topology alignment. For each aligned word pair $(w_{\text{src}}, w_{\text{tgt}})$ from FLORES-200 obtained with SIMALIGN, we extract the *pre-context* sense vectors from the Backpack sense module. This yields matrices $S_{\text{src}}, S_{\text{tgt}} \in \mathbb{R}^{K \times d}$, where each row corresponds to one of the K senses of the word in the given language. Because a target-language word may span multiple GPT-2 BPE tokens, we first pool sense vectors over all subword positions belonging to the word span (simple average over positions), obtaining $V_{\text{src}}, V_{\text{tgt}} \in \mathbb{R}^{K \times d}$.⁷

We then mean-center and row-normalize the sense matrices:

$$\tilde{V} = \text{normalize}\left(V - \frac{1}{K} \sum_{k=1}^K V_k\right), \quad (4)$$

where V_k is the k -th row of V and $\text{normalize}(\cdot)$ applies row-wise ℓ_2 -normalization. All sense vectors are thus mean-centered and unit-norm before we compare geometries.

⁷In practice, we discard (word, language) pairs where fewer than K senses are available after pooling, which is rare under our Backpack configuration.

The local geometry for a word in each language is summarized by a cosine Gram matrix

$$G_{\text{src}} = \tilde{V}_{\text{src}} \tilde{V}_{\text{src}}^\top, \quad G_{\text{tgt}} = \tilde{V}_{\text{tgt}} \tilde{V}_{\text{tgt}}^\top. \quad (5)$$

We flatten the upper-triangular entries (excluding the diagonal) of each Gram matrix into vectors $g_{\text{src}}, g_{\text{tgt}} \in \mathbb{R}^{K(K-1)/2}$, and compute the Spearman rank correlation

$$\rho(w_{\text{src}}, w_{\text{tgt}}) = \text{Spearman}(g_{\text{src}}, g_{\text{tgt}}). \quad (6)$$

This ρ score measures how similarly the two languages organize the K senses of the aligned word pair. We report the mean ρ over all sampled aligned word pairs per language, together with the corresponding control model (LM-only finetuning with $w_{\text{sns}} = w_{\text{ctx}} = 0, w_{\text{lm}} = 1.0$).

For statistical significance, we run 10,000 paired bootstrap iterations, resampling aligned word pairs with replacement at each iteration and recomputing the average ρ for both the adapted model and the control. The reported ρ values in Table 3 are averages over these iterations, and the differences $\Delta\rho$ are significant at the 95% level ($p < 0.01$).

Orthogonal Procrustes analysis. For the global structure analysis, we focus on *mixture-weighted* sense embeddings. For each aligned word pair, we compute sense mixtures using the model’s sense-weighting network: given sense vectors $s_{t,1}, \dots, s_{t,K}$ and per-sense mixture weights $\pi_{t,1}, \dots, \pi_{t,K}$ for the word,

$$e = \sum_{k=1}^K \pi_{t,k} s_{t,k}, \quad (7)$$

followed by ℓ_2 -normalization of e . We collect $N = 10,000$ aligned word pairs per language, forming matrices $E, T \in \mathbb{R}^{N \times d}$ whose rows are the English and target-language mixture embeddings, respectively.

We then solve the orthogonal Procrustes problem

$$Q^* = \arg \min_{Q: Q^\top Q = I} \|TQ - E\|_F, \quad (8)$$

whose solution is $Q^* = UV^\top$, where $T^\top E = U\Sigma V^\top$ is the singular value decomposition (SVD). After alignment, we compute the average cosine similarity

$$\text{Cosine} = \frac{1}{N} \sum_{i=1}^N \cos(T_i Q^*, E_i), \quad (9)$$

where T_i and E_i are the i -th rows of T and E , respectively. We average this score over five random subsamples of $N = 10,000$ aligned pairs to ensure stability. Table 4 reports the mean cosine for both the adapted model and the LM-only control, along with the difference $\Delta \cos$.

G.2 Empirical Ablation Protocols (Section 5)

Unless otherwise noted, all ablations reuse the SENSIA adaptation setup from Appendix C (learning rate, batch size, phase schedule, loss weights, and temperatures), and the evaluation protocol from Appendix E (conditional/PMI scoring and length normalization). Each ablation is either an inference-time modification of a fixed checkpoint or a separate training run with a single component turned off, and all scores are reported on the same FLORES-200 devtest split or downstream benchmarks as in the main results.

Sense-mixture ablation. We test whether downstream performance depends on the sense-mixture mechanism by modifying the mixture weights $\pi_{t,k}$ at inference time while keeping all model parameters fixed.

For each token t and sense index k , let $\pi_{t,k}$ be the mixture weight produced by the sense-weight network, and let $\text{Top-}k_t$ be the set of indices for the k most probable senses at position t . We compare:

1. **Full mixture** (default SENSIA): use the learned $\pi_{t,k}$ unchanged.
2. **Top- k mixture** (we use $k = 1$):

$$\pi'_{t,k} = \begin{cases} \frac{\pi_{t,k}}{\sum_{j \in \text{Top-}k_t} \pi_{t,j}} & k \in \text{Top-}k_t, \\ 0 & \text{otherwise.} \end{cases} \quad (10)$$

3. **Uniform mixture:**

$$\pi_{t,k}^{\text{uni}} = \frac{1}{K} \quad \forall k \in \{1, \dots, K\}. \quad (11)$$

In all three cases, the token representation is recomputed as $h_t = \sum_k \pi_{t,k} s_{t,k}$ (or with $\pi'_{t,k} / \pi_{t,k}^{\text{uni}}$), and the resulting sequence is fed through the frozen transformer and LM head.

We evaluate each variant on the FLORES-200 devtest set, computing average token-level cross-entropy

$$\text{CE} = -\frac{1}{T} \sum_{t=1}^T \log p_{\theta}(y_t | y_{<t}, x), \quad (12)$$

where T is the number of target tokens and x is the source sentence. Table 5 reports CE for four target languages, all evaluated with the same adapted small Backpack checkpoint.

Loss and phase ablations. We ablate individual loss terms and training phases using Indonesian as the target language. Recall the total objective

$$\mathcal{L}_{\text{total}} = w_{\text{sns}} \mathcal{L}_{\text{sns}} + w_{\text{ctx}} \mathcal{L}_{\text{ctx}} + w_{\text{lm}} \mathcal{L}_{\text{lm}}, \quad (13)$$

with phase-dependent weights ($w_{\text{sns}}, w_{\text{ctx}}, w_{\text{lm}}$) annealed according to the schedule in Appendix C. For each ablation variant in Table 6, we train a separate model with the following modifications:

- **No Context loss:** set $w_{\text{ctx}} = 0$ at all steps; keep w_{sns} and w_{lm} as in the default schedule.
- **No Sense loss:** set $w_{\text{sns}} = 0$ at all steps; keep w_{ctx} and w_{lm} as in the default schedule.
- **No LM loss:** set $w_{\text{lm}} = 0$ at all steps; keep w_{sns} and w_{ctx} unchanged.
- **No Alignment phase:** skip the alignment phase entirely and start from the joint-phase weights, then proceed to the polish phase as usual.
- **No Joint phase:** run the alignment phase and then jump directly to the polish phase.
- **No Polish phase:** run only the alignment and joint phases (truncating total training steps to match the baseline), with the same per-phase weights.

All other hyperparameters (learning rate, batch size, temperatures, phase boundaries, etc.) follow Appendix C, and we reuse the same evaluation protocol and scoring parameters as the full model. XCOPA, XStoryCloze, Belebele, and FLORES perplexity are evaluated on the same splits as Table 1, producing the numbers in Table 6.

Scaling experiments. For scaling, we extend Backpack to GPT-2 medium (345M) and large (762M) backbones. Pretraining follows the OpenWebText setup of Hewitt et al. (2023), with token budgets chosen according to the Chinchilla scaling law to balance model size and data scale. We then adapt each backbone to Indonesian and Swahili using the same SENSIA schedule as for the small model.

Generation and scoring hyperparameters (cond/PMI interpolation and length normalization) are held fixed across model sizes to isolate the effect of scale, giving the results in Table 7.

Low-resource adaptation. To approximate low-resource conditions, we downsample the Indonesian and Swahili parallel corpora used in the main experiments to a subset of 100k sentence pairs per language. Downsampling is performed by uniformly sampling without replacement from the full OPUS-derived corpus (after all filtering), and we keep the Backpack architecture, SENSIA schedule, loss weights, and hyperparameters identical to the full-data setting.

Each low-resource model is trained from the same English Backpack initialization and evaluated on the same XCOPA, XStoryCloze, and Belebele test splits as the full-data model. We reuse the same decoding/scoring setup (Appendix E), so differences in accuracy and FLORES perplexity in Table 8 are attributable solely to the reduction in parallel data.

Non-Latin script stress test. For the Chinese stress test, we adapt the same English Backpack model to Chinese using English–Chinese parallel data mined from OPUS, but *keep the original English GPT-2 BPE tokenizer unchanged*. Chinese text is thus represented as long sequences of over-segmented subword fragments, intentionally placing the model in a hostile tokenization regime.

We reuse the adaptation hyperparameters, phase schedule, and loss weights from the main experiments, and evaluate on Chinese XCOPA, Chinese XStoryCloze, and the Chinese split of Belebele, comparing directly against the monolingual Chinese Goldfish baseline under the same evaluation protocol and scoring setup. Table 9 reports per-task accuracies for this adversarial setting and supports our conclusion that sense adaptation cannot fully overcome a severely mismatched tokenizer, even though it still yields modest gains on some narrative-style tasks.