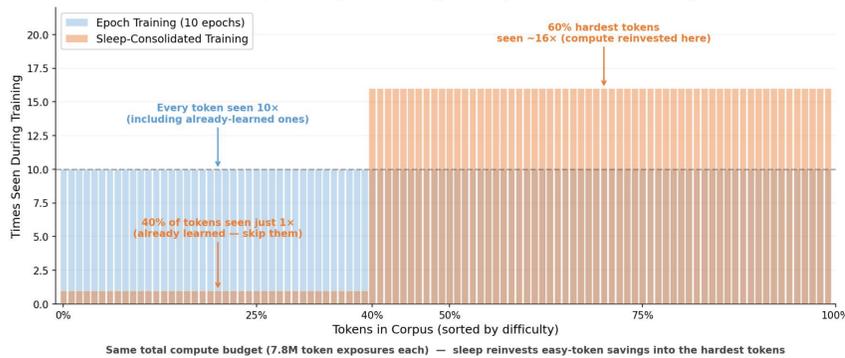


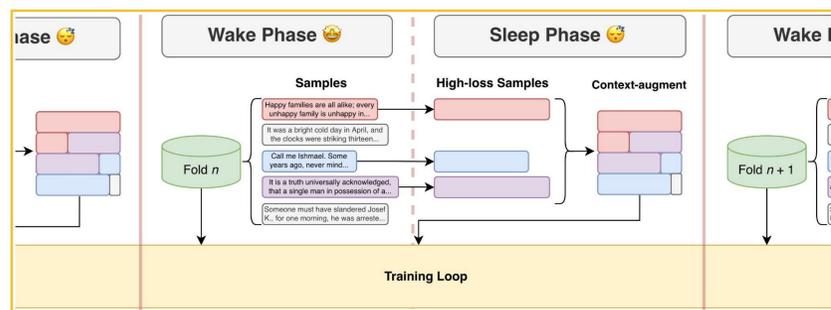
Background

Token Exposure: Epoch Training vs Sleep-Consolidated Learning



- Standard LM training is expensive^[2]
 - Much less data-efficient than humans
 - Requires significant compute
- Epoch training schedules aren't cognitively plausible
 - Humans live through time linearly
 - Replay happens during sleep^[1]

Methods: Training Schedule



Goal: Replace epoch schedule with sleep-like, cognitively-plausible replay schedule

Sleep Mechanism

- Incorporates biologically-inspired replay phase into LLM training
- Model only sees each data point once during wake
- Wake samples selected for sleep phase based on difficulty
- High-loss samples are "recontextualized" to encourage abstraction

Linguistic Evaluation

- BLiMP**: Grammar and linguistic understanding
- WUGS**: Derivational morphology
- Entity Tracking**: Discourse entity states

Compare Sleep LMs to Baseline (Epoch-schedule LMs)

Methods: Dataset & Model

Dataset: BabyLM 100M Corpus

- Approximates linguistic input a child receives by age 13
- Combination of sources from 2 domains: transcribed speech & child-directed text

Model: RoBERTa-Small

- 8 hidden layers, 8 attention heads
- Hidden size: 256, intermediate size: 2048
- Trained on masked language modeling (MLM)

Baseline PPL: 36.63

Hyperparameter Sweep

Best PPL: 109.133

- Gridsearch over # phases, ratio, contextualize, steps
- Fewer phases = Better
 - Longer wake phases
 - Less overfitting(?)
- Replay Ratio and Contextualization do not have significant corr.

Replay Strategy

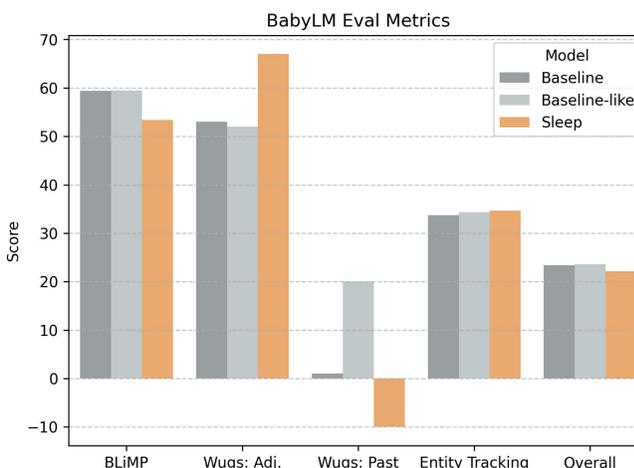
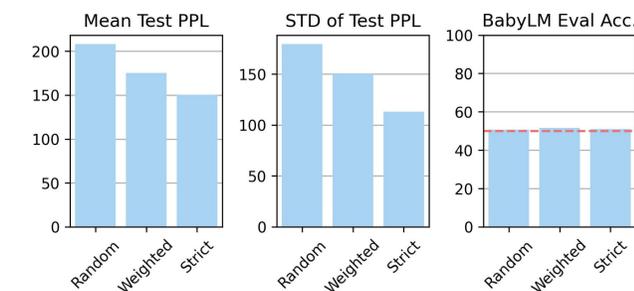
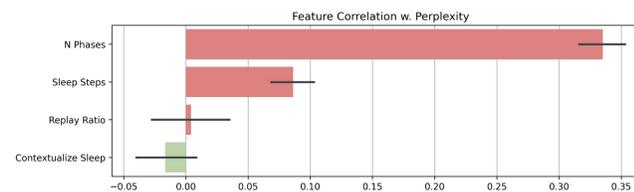
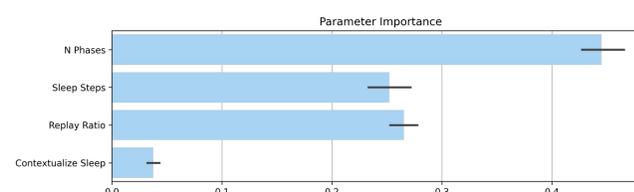
Best PPL: 150.21

- Tested selection of samples for buffer
- Strict strategy was best
 - Mean and STD of PPL reduced
- Similar in BabyLM eval

BabyLM Evaluation

- Sleep LM > Baseline in Adj. Nominalization
 - More human-like than baseline
- Linguistic understanding (BLiMP) is still poorer than baselines
- Entity tracking results are inconclusive

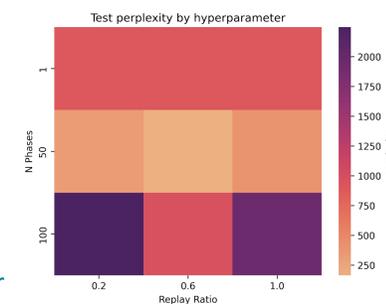
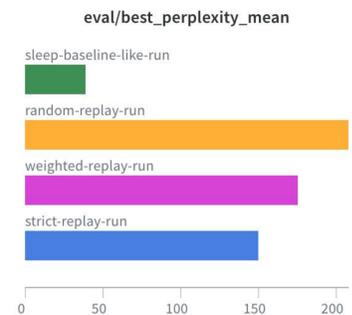
Results



Conclusions

Currently, standard epoch training outperforms sleep mechanism in both perplexity and BabyLM evaluation metrics.

- Selective replay reduces wasted compute on already-mastered data
 - Strict replay = more consistent PPL
- "Catastrophic forgetting" may occur
 - PPL increases drastically
 - Plasticity decay may further close the gap between biological and LM language acquisition
- Promising results in hyperparameter sweeps
 - perplexity seems to be lower when diverging from epoch schedule



Limitations & Future Work

- Still underperforming epoch-training models
 - More rigorous/comprehensive hyperparameter tuning
- Replay selection strategy - loss might be encoding noise
 - Try other ways of selecting useful samples for learning
- Implement plasticity decay: weight downscaling + hierarchical layer freezing
 - Could help model retain information from earlier folds
- Extend to larger-scale models and datasets beyond BabyLM
 - Looking at PPL pattern with other datasets

Acknowledgements

We would like to thank our mentor, Dr. Alex Warstadt, for his feedback and guidance. This work also used resources available through the National Research Platform (NRP) at the University of California, San Diego. NRP has been developed, and is supported in part, by funding from National Science Foundation, from awards 1730158, 1540112, 1541349, 1826967, 2112167, 2100237, and 2120019, as well as additional funding from community partners.

References

- [1] Rasch and Born. 2013. "About Sleep's Role in Memory."
- [2] Warstadt and Bowman. 2024. "What Artificial Neural Networks Can Tell Us About Human Language Acquisition."



Scan Me!