Appendices

A   Benchmarks

LRU Training Benchmarks (Mean ± Min/Max)

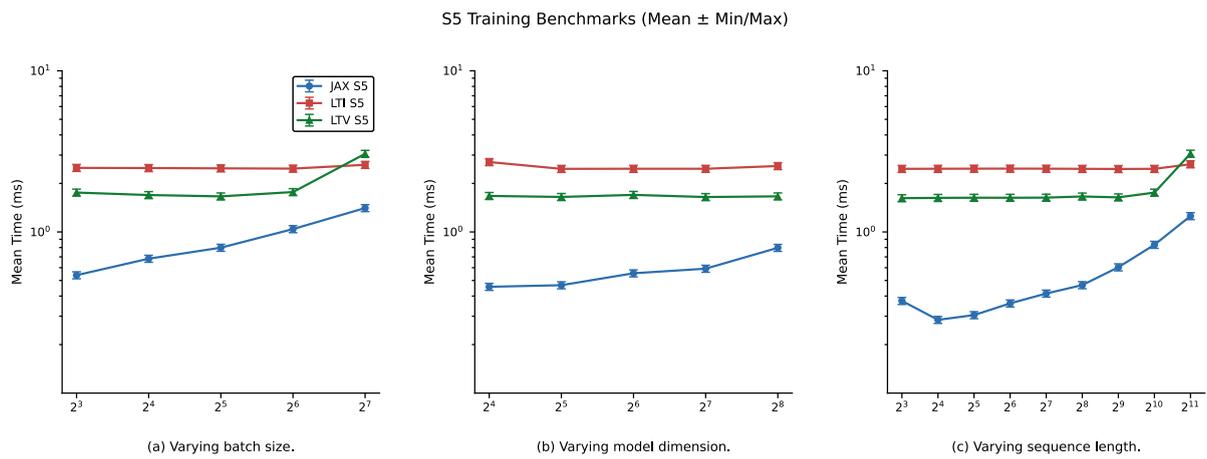

(a) Varying batch size.   (b) Varying model dimension.   (c) Varying sequence length.
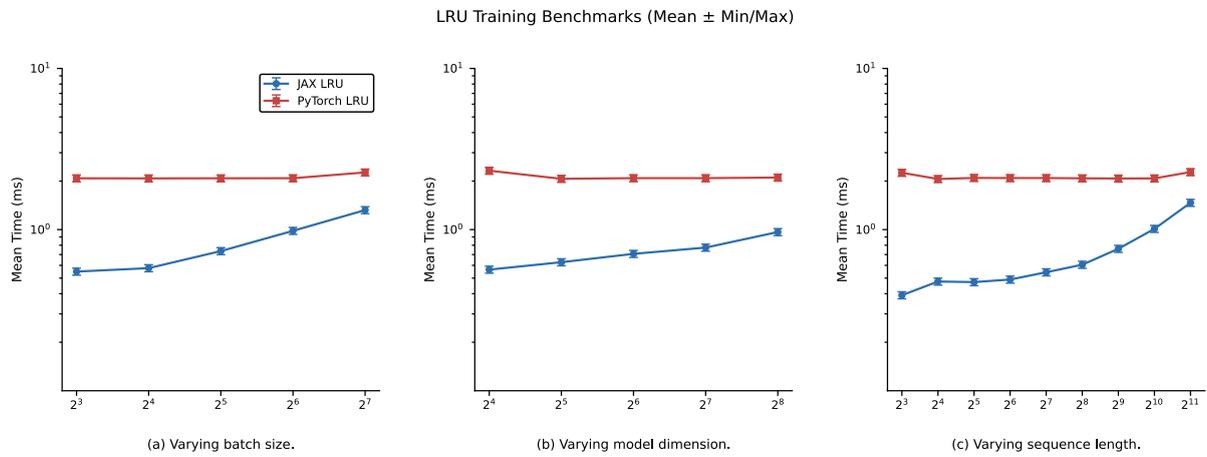
Figure 3: LRU Training Benchmarks.

S5 Training Benchmarks (Mean ± Min/Max)

(a) Varying batch size.   (b) Varying model dimension.   (c) Varying sequence length.

Figure 4: S5 Training Benchmarks.

Mamba Training Benchmarks (Mean ± Min/Max)

(a) Varying batch size.   (b) Varying model dimension.   (c) Varying sequence length.
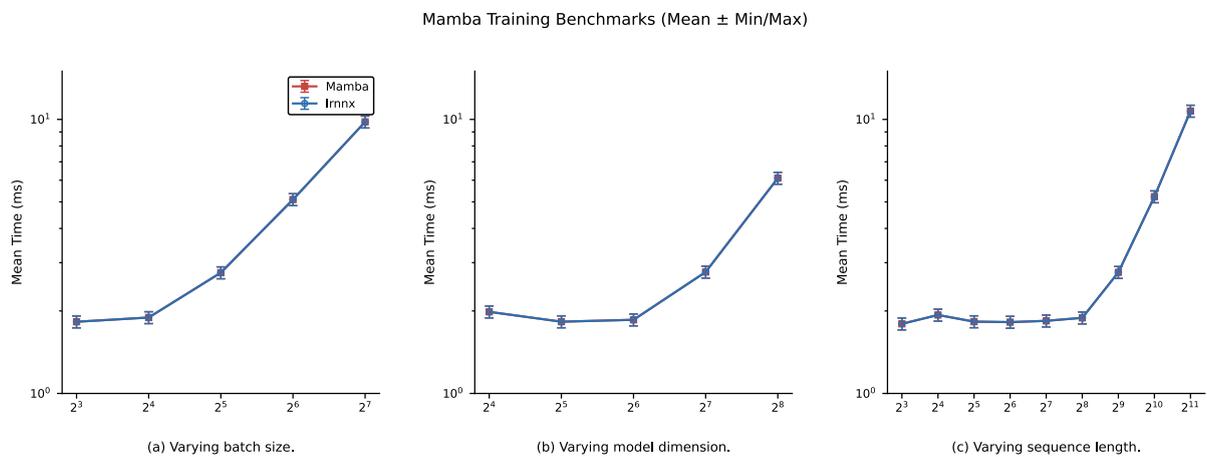
Figure 5: Mamba Training Benchmarks.

LRU Inference Benchmarks (Mean ± Min/Max)

(a) Varying batch size.

(b) Varying model dimension.

(c) Varying sequence length.
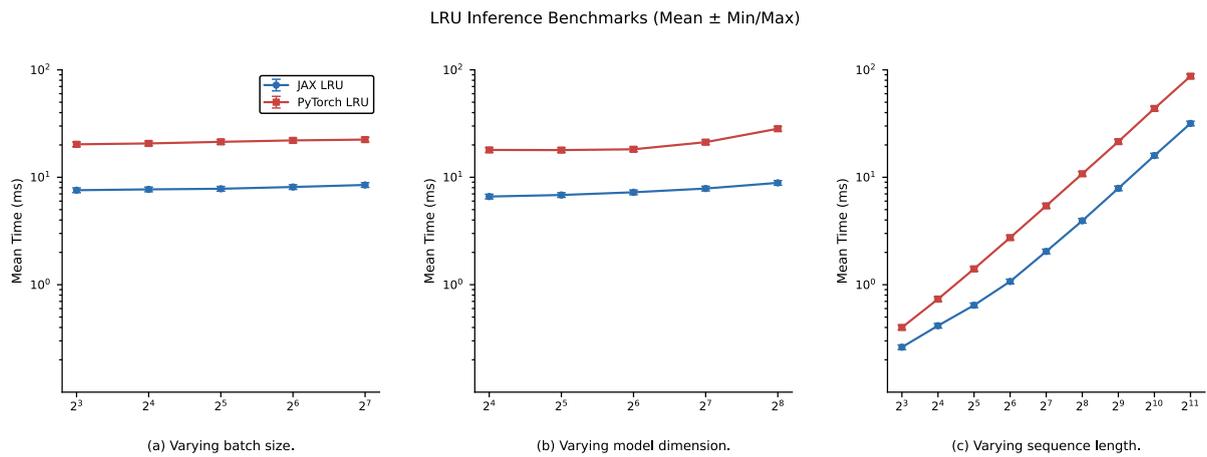
Figure 6: LRU Inference Benchmarks.

S5 Inference Benchmarks (Mean ± Min/Max)

(a) Varying batch size.

(b) Varying model dimension.

(c) Varying sequence length.
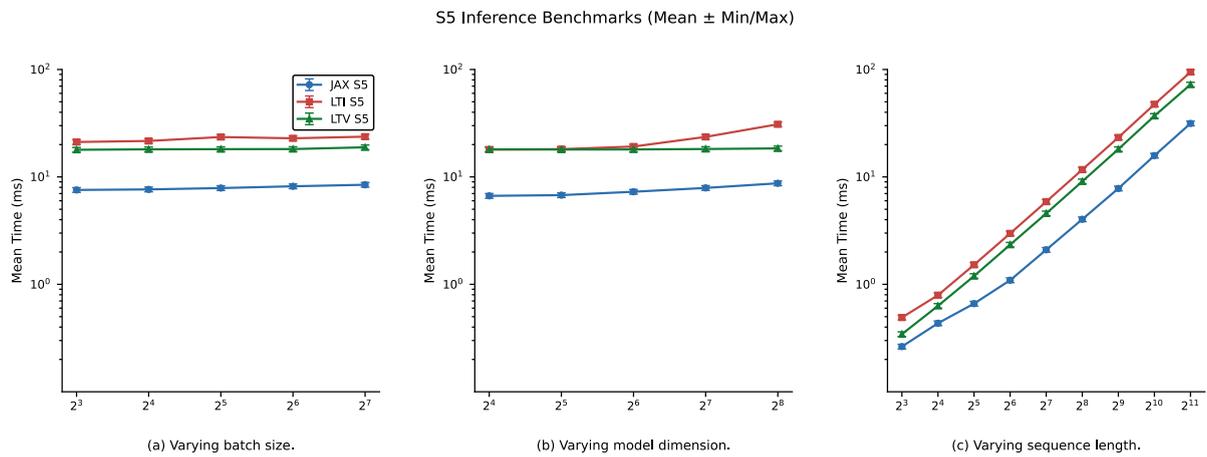
Figure 7: S5 Inference Benchmarks.

Mamba Inference Benchmarks (Mean ± Min/Max)

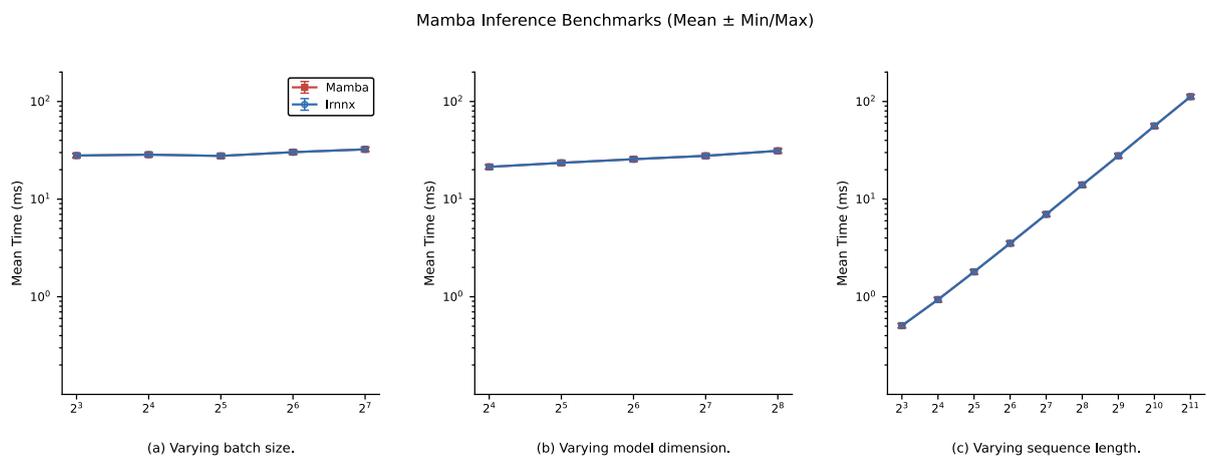(a) Varying batch size.

(b) Varying model dimension.

(c) Varying sequence length.

Figure 8: Mamba Inference Benchmarks.