

F.4 Choice of the Teacher Model

Throughout this work we used Gemini 2.5 Flash as the generator model. To demonstrate that Simula is a model-agnostic framework we conducted a targeted ablation using the open-source *Qwen3-Next-80B Instruct* model (Yang et al., 2025). We replicated the downstream experiment on the CTI MCQ dataset, scaling the generated data size from 4k to 84k, while keeping the student model and optimization parameters unchanged.

As shown in Figure 13 below, performance trends remain consistent with the results reported in Section 4.3 using the Gemini model: student models fine-tuned on data generated by the Full Simula system consistently outperform the Baseline and partial ablations across data scales. This provides additional evidence that the observed improvements are due to our reasoning-driven data generation approach, rather than model-specific factors.




Figure 13: **Downstream Performance using Qwen3-Next-80B Instruct** We report mean accuracy with 95% CI for the CTI MCQ dataset for data generated using the Qwen-3 80B instruction-tuned model. We continue to see that the full Simula system version with the critic, provides the best performance compared to the other versions.