

Jeonghun Kim

ML Engineer / AI Researcher

portfolio.jhkim3ispark.work



Software - Software Engineer, Compiler (Kernel Optimization)

bandlagom0927@gmail.com

+82 10-9559-9471 · Seoul

github.com/thuygom



Cover Letter

Dear FuriosaAI Hiring Team,

I am applying for the Software - Software Engineer, Compiler (Kernel Optimization) position (Job ID 4007040201). I bring experience translating a measured bottleneck into an implementation change and validating the result on fixed inputs. My work includes speech-processing optimization and task-specific model evaluation, with performance and output quality checked separately. Before ISPARK, I served one year as an information-systems NCO in the Republic of Korea Air Force (Dec 2021–Dec 2022), maintaining Linux servers through shell scripting and log analysis.

I redesigned long-audio processing and reduced server processing time by 36.9%, from 122.6 to 77.4 seconds on a 30-minute recording, using Silero voice activity detection and faster-whisper batching at batch size 8. The same workload had also exposed a delivery problem: the server finished after the synchronous HTTP request timed out. I separated live sessions from recoverable asynchronous jobs, with task polling and persisted progress. I treated recognition quality, execution time, and delivery as separate engineering problems, allowing each change to address a specific cause.

I improved a four-billion-parameter Text-to-SQL model from 26.6% to 76.6% normalized database-result agreement on a 64-question internal CRM development benchmark, using teacher-student training and rank-16 LoRA. The model matched 49 of 64 reference results, compared with 17 for the base model; scoring ignores row order and normalizes numbers and dates. I also verified the serving interface: missing schema input caused invented table names despite unchanged weights. I documented the required schema inputs alongside the adapter so the trained behavior could be reproduced in serving.

At FuriosaAI, I can contribute representative workloads, reference checks, and focused performance investigations. I would reproduce a bottleneck, examine the operation's constraints, and validate numerical behavior across input variations as I develop an optimization.

Sincerely,

Jeonghun Kim