Korean Researchers Profile

Korea-EU Horizon Europe Researchers Networking Forum 2026 on Cluster 4

Full Name Tae-Wan KIM

Position Professor, Seoul National University / Founder & CEO, LLM Core AI, Inc.

Organization Seoul National University

Email taewan@snu.ac.kr

Personal Link(s)

Deep Learning Laboratory, Seoul National University: https://dllab.snu.ac.kr/
LLM Core AI, Inc.: https://llmcore.ai/en
LinkedIn: linkedin.com/in/tae-wan-kim-8285b9145

Your Target Call

Primary Target Call: HORIZON-CL4-2027-04-DATA-03 — New approaches for decentralized, federated and sustainable AI data processing (RIA)
Related strategic alternatives under evaluation: HORIZON-CL4-2027-04-DIGITAL-EMERGING-11 and HORIZON-CL4-2027-04-DATA-09.

Your Expected Contribution to the Project / Consortium

Our team can contribute a cross-layer technology stack for energy-efficient long-context AI, spanning algorithms, runtime software, memory systems, and hardware validation. Our core contributions include TurboQuant-based ultra-low-bit KV-cache compression, BitNet-inspired low-bit and ternary inference, hardware-aware kernels, and memory-centric accelerator co-design. We can provide GPU reference implementations and extend validation to FPGA and ASIC/MPW-oriented prototypes, using reproducible metrics such as memory footprint, bandwidth, latency, quality retention, and tokens per watt. Within a consortium, we would be well suited to lead or co-lead tasks on model and inference efficiency, hardware– software co-design, accelerator architecture, and prototype validation. We can also contribute Korea– Europe semiconductor collaboration links and help connect research results to practical silicon and industrial deployment pathways.

Biography

Prof. Tae-Wan Kim is a Professor at Seoul National University and Founder & CEO of LLM Core AI, Inc., working at the intersection of artificial intelligence, computer architecture, and semiconductor implementation. He received his Ph.D. from Arizona State University and has professional experience spanning academia and industry. His current research focuses on energy-efficient AI systems, especially long-context large language model inference, low-bit and ternary neural networks, KV-cache compression, memory-centric computing, and hardware–software co-design.

His group is developing TURBOECS, a cross-layer approach that combines TurboQuant-based 3-bit KV- cache compression, BitNet-inspired low-bit inference, hardware-aware runtime kernels, and a GPU-to- FPGA-to-ASIC/MPW validation path. The goal is to reduce memory traffic and energy consumption while improving practical AI efficiency measured by metrics such as tokens per watt, latency, bandwidth, and quality retention.

Prof. Kim is currently building Korea–Europe R&D collaborations around Horizon Europe Cluster 4, Chips JU/Xecs, and related semiconductor and AI ecosystems. His collaboration strategy is to connect Korean strengths in

memory and semiconductor execution with European capabilities in AI/HPC systems, microelectronics, research infrastructures, and industrial validation. He is particularly interested in contributing as a technical partner or work-package leader in projects on sustainable AI processing, efficient frontier AI, and next-generation AI hardware.

Research Areas / Expertise

Energy-efficient AI and sustainable AI computing; long-context LLM inference; KV-cache compression; TurboQuant-based ultra-low-bit quantization; BitNet and ternary neural networks; low-bit inference; hardware–software co-design for AI accelerators; memory-centric and near-memory computing; processing-in-memory (PIM); GPU/FPGA/ASIC prototyping; AI semiconductor architecture; heterogeneous computing; and performance/energy benchmarking including tokens per watt, latency, bandwidth, and memory traffic.

Motivation, Expectation and Plan

My main motivation is to convert current Korea–Europe technical discussions into a focused Horizon Europe consortium for a 2027 Cluster 4 proposal. At the forum, I would like to validate the best target-call positioning, identify a credible European coordinator and complementary partners, and define clear work- package responsibilities around efficient AI processing, hardware–software co-design, prototyping, and system-level validation.

I am particularly interested in partners that can provide AI/HPC or industrial use cases, European computing and testbed infrastructure, semiconductor or FPGA/ASIC expertise, sustainability and energy measurement, and strong Horizon Europe coordination experience.

As a concrete follow-up, I propose to agree on one common proof-of-concept using public models and benchmarks, prepare a 1–2 page joint concept note and preliminary work-package map, identify the remaining partner gaps, and schedule technical and proposal-development meetings immediately after the forum. The objective is to move from networking to a technically grounded consortium and proposal plan.