Dnotitia Paper Accepted to VLDB 2025, World's Top Database Conference

SEOUL, SOUTH KOREA, July 17, 2025 /EINPresswire.com/ -- Dnotitia Inc. (Dnotitia), a fast-rising startup in long-term memory AI and semiconductor-integrated solutions, announced that its collaborative research on vector database optimization has been officially accepted to VLDB 2025, one of the world's most prestigious academic conferences on data systems and infrastructure.

VLDB (Very Large Data Bases) is a top-tier international forum with over 50 years of history, covering the full spectrum of data infrastructure, ranging from data management systems and distributed architecture to large-scale experimentation and high-performance applications. It serves as a key platform where researchers and practitioners from around the world share the latest advances that are shaping the future of data science and AI technologies.

The conference addresses a wide range of core technologies driving AI era, including vector search, distributed processing, and large-scale data infrastructure. Each year, more than 600 papers are submitted by researchers from global tech leaders such as Google, Meta, Microsoft, and Amazon, as well as top universities including MIT, Stanford, and UC Berkeley. With an acceptance rate of around 15%, only a small number of submissions are selected following a rigorous peer-review process based on technical excellence and research contribution. VLDB 2025 will take place in London, UK, from September 1 to 5.

The accepted paper, titled "Turbocharging Vector Databases using Modern SSDs," was led by Professor Sang-Won Lee's research team at Seoul National University(SNU), and conducted as part of the industry-academia collaboration with Dnotitia's Chief Data Officer Hongchan Noh and Professor Jaeyoung Do of SNU.

The presents SSD-based optimization strategies aimed at improving vector search performance, a critical capability in retrieval-augmented generation (RAG) and semantic search. To address the limitations of traditional disk-based vector databases, such as slow retrieval speed and inefficient disk utilization, the authors leveraged the parallel processing capabilities of SSDs, applied insertion reordering techniques, and reorganized data layouts to improve storage efficiency. These optimizations led to significant performance improvements.

The techniques were implemented in pgvector, an open-source vector search extension for PostgreSQL, and evaluated through real-world experiments. The results showed up to 8.5x faster query throughput, over 90% reduction in index build time, and threefold improvement in cache efficiency.

This research is closely aligned with Dnotitia's ongoing development of its Vector Data Processing Unit (VDPU) and large-scale RAG systems. Rather than remaining at the theoretical level, the work demonstrates clear potential for real-world product integration, reinforcing Dnotitia's identity as a research-driven, technology-first AI company.

Dnotitia plans to build on the results of this research by further advancing its work under the national R&D project "Development of Vector Databases for Long-Term Memory in Large-Scale AI Models," funded by Korea's Ministry of Science and ICT and the Institute of Information & Communications Technology Planning & Evaluation (IITP) as part of the Software and Computing Technology Development Program. The company also aims to strengthen its academic collaboration with Seoul National University and continue developing SSD-optimized RAG systems and scalable AI infrastructure technologies.

In August, Dnotitia will present a live demonstration at the upcoming Future of Memory and Storage (FMS) 2025, showcasing SSD performance under real-world vector search workload, going beyond conventional disk benchmarks.

This press release can be viewed online at: https://www.einpresswire.com/article/831418356