Guest Editors' Notes

Welcome to the March 2026 issue of the ACM SIGMOD Record, a special issue presenting papers and technical perspectives covering a selection of the best conference papers from 2025 in data management. These papers have been given the 2025 ACM SIGMOD Research Highlight Award. This is an award for the database community to showcase a set of research projects that exemplify core database research. In particular, these projects address an important problem, the papers represent a definitive milestone in solving the problem, and have the potential of significant impact. This award also aims to make the selected works widely known in the database community, to our industry partners, and to the broader ACM community.

The award committee and editorial board included Carsten Binnig, Rada Chirkova, Alfons Kemper, Samuel Madden, Julia Stoyanovich, and Ke Yi (chair). We solicited articles from SIGMOD, PODS, VLDB, ICDE, EDBT, and ICDT, published in 2025, as well as from community nominations. We used a careful review process, in which each nominated paper was discussed in a virtual meeting. Papers with conflict of interest were discussed in the absence of the conflicting committee members. This year, ten articles were finally selected as 2025 Research Highlight Award winners.

The authors of each article worked closely with an associate editor to rewrite the article into a compact 8-page format and improved it to appeal to the broad data management community. In addition, each research highlight in the March 2026 issue is accompanied by a one-page technical perspective written by an expert on the topic presented in the article. The technical perspective provides the reader with an overview of the background, the motivation, and the key innovation of the featured research highlight, as well as its scientific and practical significance.

2025 ACM SIGMOD Research Highlight Award Winners

1. A theoretical breakthrough on the join order problem together with practical improvements ("DPconv: Super-Polynomially Faster Join Ordering");
2. A technique to vectorize multiple instances of a vertex-centric algorithm ("Automating Vectorized Distributed Graph Computation");
3. An optimal algorithm for computing conjunctive queries in terms of the output size ("Output-Optimal Algorithms for Join-Aggregate Queries");
4. "Output-sensitive Conjunctive Query Evaluation", which has overlapping results with the previous paper (a combined highlight paper is written by the authors of the two papers jointly);
5. A method to analyze collections of documents while preserving strong privacy guarantees ("Differentially Private Substring and Document Counting");
6. A dynamic range filter that simultaneously supports updates, variable-length range queries, and variable-length keys ("Diva: Dynamic Range Filter for Var-Length Keys and Queries");
7. A generic and self-decoding framework for easy integration of data formats and data system ("AnyBlox: A Framework for Self-Decoding Datasets");

8. A new method for optimizing UDFs in SQL databases by combining UDF outlining with selective inlining ("The Key to Effective UDF Optimization: Before Inlining, First Perform Outlining");
9. A framework for reusing intermediate results and memory allocations in multi-backend machine-learning systems ("MEMPHIS: Holistic Lineage-based Reuse and Memory Management for Multi-backend ML Systems");
10. A new programming language for relational data that abandons the traditional sublanguage paradigm and instead constructs a language from a minimal core that can evolve through libraries ("Rel: A Programming Language for Relational Data").

On behalf of the SIGMOD Record Editorial Board, we hope that you enjoy reading the March 2026 issue of the SIGMOD Record!

Carsten Binnig, Rada Chirkova, Alfons Kemper, Samuel Madden, Julia Stoyanovich, Ke Yi

March 2026