



Proceedings of the VLDB Endowment

Volume 19, No. 12 – August 2026

Editors in Chief:

Yanlei Diao and Xiaokui Xiao

Associate Editors:

Alkis Polyzotis, Alkis Simitsis, Andrew Pavlo, Angela Bonifati, Anh Dinh, Antonis Deligiannakis, Ashraf Aboulnaga, Badrish Chandramouli, Bin Cui, Bogdan Cautis, Ce Zhang, Elena Ferrari, Evaggelia Pitoura, Fatemeh Nargesian, Guoliang Li, Hong Cheng, Immanuel Trummer, Jana Giceva, Jennie Rogers, Jian Pei, Jianliang Xu, Karima Echihabi, Katja Hose, Khuzaima Daudjee, Kunsoo Park, Kyuseok Shim, Laks Lakshmanan, Lei Cao, Lei Chen, Li Xiong, Mauro Sozio, Meihui Zhang, Melanie Herschel, Michael Böhlen, Nan Tang, Niv Dayan, Norman May, Raul Castro Fernandez, Raymond Chi-Wing Wong, Reynold Cheng, Ruben Mayer, Semih Salihoglu, Shimin Chen, Sibor Wang, Sourav S Bhowmick, Spyros Blanas, Steven Whang, Sudipto Das, Umar Farooq Minhas, Wei Wang, Wook-Shin Han, Yannis Velegrakis, Yanyan Shen, Yi Chen, Yuncheng Wu, Zhifeng Bao, Zi Huang

Publication Editors:

Andrea Mauri, Giovanni Simonini, Siqiang Luo, Stefan Halfpap, Subhadeep Sarkar

PVLDB – Proceedings of the VLDB Endowment

Volume 19, No. 12, August 2026.

All papers published in this issue will be presented at the 52nd International Conference on Very Large Data Bases, Boston, MA, USA, 2026.

Copyright 2026 VLDB Endowment

This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org.

Volume 19, Number 12, August 2026

Pages i – xx and 3834 - 4988

ISSN 2150-8097

Available at: <http://www.pvldb.org> and <https://dl.acm.org/journal/pvldb>

TABLE OF CONTENTS

Front Matter

Copyright Notice	i
Table of Contents	ii
PVLDB Organization and Review Board – Vol. 19	xiii

Industrial Papers

Overlay Bitmap Encoding for Efficient Consumption of Apache Parquet Files	3834
<i>Güneş Aluç, Daniel J. Farrar, Manoj Kumar Sharma, Weibing Zhang</i>	
Enhancing Database Write Performance with the Write Operation Pushdown Framework.....	3847
<i>Hao Wu, Lutong Zhang, Xinying Zheng, Cheng Zhu, Bo Shang, Lixun Cao, Xian Liu, Lei Zhang, Cheng Chen</i>	
K2+: A Multi-Region OLTP Database with Disaggregated Storage for High Availability and Strong Consistency	3860
<i>Haoze Song, Guoli Wei, Xusheng Chen, Yazhi Feng, Xieyun Fang, Yong Men, Hao Feng, Yongkun Li, Heming Cui</i>	
Clock2Q+: A Simple and Efficient Replacement Algorithm for Metadata Cache in VMware vSAN	3874
<i>Yiyan Zhai, Bintang Dwi Marthen, Sarath Balivada, Vamsi Sudhakar Bojji, Eric Knaft, Jitender Rohilla, Jiaqi Zuo, Quanxing Liu, Maxime Austruy, Junlong Gao, Wenguang Wang, Juncheng Yang</i>	
ZipFlow: a Compiler-based Framework to Unleash Compressed Data Movement for Modern GPUs.....	3888
<i>Gwangoo Yeo, Zhiyang Shen, Wei Cui, Matteo Interlandi, Rathijit Sen, Bailu Ding, Qi Chen, Minsoo Rhu</i>	
RayBooster: A Ray Tracing Engine to Accelerate SedonaDB	3902
<i>Liang Geng, Rubao Lee, Dewey Dunnington, Feng Zhang, Jia Yu, Xiaodong Zhang</i>	
Highly Scalable SQL Azure MVCC Version Cleaner	3916
<i>Jithendra Gandikota, Prashanth Purnananda, Chaitanya Sreenivas Ravella, Adrian-Leonard Radu, Hanuma Kodavalla, Max Garton, Allison Ohara, Wayne Chen, Raghavendra Tk</i>	
DBAgent: An RL-Based Agent for Autonomous Database Operations and Maintenance	3928
<i>Xu Chen, Junming Chen, Shuncheng Liu, Yuyang Xia, Zibo Liang, Xiaoming Qin, Kai Zeng, Han Su, Kai Zheng</i>	
Incremental Query Optimizer Statistics in Amazon Redshift	3941
<i>Pascal Pfeil, Zhe Heng Eng, Magnus Müller, George Erickson, Roger Kim, Mohammed Al-Kateb, Majid Saeedan, Dominik Horn, Orestis Polychroniou, Mengchu Cai, Tim Kraska</i>	
TEngineDB-V: An OLAP-Native Vector Search System for Large- <i>k</i> Workloads at Tencent.....	3955
<i>Xufei Wu, Pengcheng Zhang, Yitong Song, Xiaobo Zhang, Anqi Liang, Kai Wang, Jijun Du, Yidi Xiong, Guangxu Cheng, Zhe Chen, Peng Chen, Guoliang Li, Xuanhe Zhou, Fan Wu</i>	
Hermes at Scale: Powering Distributed Queries with a Unified Memory Fabric	3969
<i>Tim Gubner, Matheus Nerone, Diego Tomé, Vitor Oliveira, Rune Humborstad, Manyi Lu</i>	
Tuning the Lookahead Distance for PostgreSQL Asynchronous IO	3982
<i>Wentao Wu, Jiasheng Hu, Manoj Syamala, Andres Freund, Vivek Narasayya</i>	
FastCompose: Eliminating Compilation Cold Starts in Query Execution with Composition.....	3996
<i>Venkatraman Govindaraju, Yankun Shen, Kiran Chinta, Shermal Fernando, Pascal Pfeil, Orestis Polychroniou, Naresh Chainani</i>	

Reaching the Pinnacle of TPC-DS: Co-design of Architecture, Executor, and Storage in TDSQL.....	4009
<i>Yiteng Chu, Jie Jiang, Yuxing Chen, Feng Zhang, Yuxin Tang, Anqun Pan, Xin Wu, Xiang Hu, Hailin Lei, Qian Zhang, Ziyang Lin, Xiaoliang Yan, Xiaojian Fan, Liang Zhang, Bin Qu, Wen Zhang, Mengbang Xing, Lixiong Zheng, Wei Lu, Yunpeng Chai, Xiaoyong Du</i>	
AutoLiquid: Autonomic Data Layout Optimization for the Databricks Lakehouse.....	4023
<i>Supun Nakandala, Naga Raju Bhanoori, Yunjia Zhang, Parimarjan Negi, Ankur Sharma, Eric Liang, Cindy Jiang, Sirui Sun, Terry Kim, Mostafa Mokhtar, Vijayan Prabhakaran, Bart Samwel, Sunitha Beeram, Siddharth Taneja, Amir Hormati, Amit Shukla, Michalis Petropoulos, Reynold Xin</i>	
No Silver Bullet: Boosting GaussDB Performance on the 30TB TPC-H Workload	4036
<i>Tim Zeyl, Jason Lam, Shu Lin, Reza Pournaghi, Qi Cheng, Calvin Wong, Kaixiang Du, Yuliang He, Yang Sun, Weicheng Wang, Paul Lee, Chen Ruo, Yang Xinyi, Li Qunan, Wang Junjie, Hu Dongxing, Chong Chen, Per-Åke Larson</i>	
Evaluating the Practical Effectiveness of LLM-Driven Index Tuning on Microsoft SQL Server	4049
<i>Xiaoying Wang, Wentao Wu, Vivek Narasayya, Surajit Chaudhuri</i>	
ScaleSense: Cost-Intelligent Scaling Framework via Learned Resource Estimation in Alibaba AnalyticDB	4063
<i>Yifan Wu, Yuhang Li, Zhenhua Wang, Ke Chen, Lidan Shou, Zonghao Chen, Liang Lin, Huan Li, Gang Chen</i>	
CoeusBI: A Comprehensive Interactive Business Intelligence System Powered by LLMs at Baidu	4076
<i>Jinqing Lian, Chaofan Li, Yingxia Shao, Ming Wang, Yang Dong, Xinyi Liu, Wei Zhang, Chaoxian Gui, Tianqi Wan, Ming Dong</i>	
OceanBase Bacchus: a High-Performance and Scalable Cloud-Native Shared Storage Architecture for Multi-Cloud	4089
<i>Quanqing Xu, Mingqiang Zhuang, Chuanhui Yang, Quanwei Wan, Fusheng Han, Fanyu Kong, Hao Liu, Hu Xu, Junyu Ye</i>	
Virtualizing Recursion: Just-In-Time Graph Analytics in a Hyperscale Relational Warehouse	4103
<i>Jason Arnold, Knut Stolze</i>	
TQP++: Bridging ML Compilers and Analytical Query Processing on GPUs.....	4116
<i>Wei Cui, Peng Cheng, Carlo Curino, Rathijit Sen, Matteo Interlandi</i>	
The Live Database: Firestore’s Scalable and Consistent Realtime Queries	4130
<i>Matt Bell, David Gay, Per Jacobsson, Ram Kesavan, Namita Lal, Xijiao Li, Hui Wu</i>	
Benchmarking Native In-Database TPCx-AI at 100 Terabytes in Ocident Hyperscale Data Warehouse.....	4143
<i>Jason Arnold, Neesh Dahiya, Knut Stolze</i>	
A Decade of Apache Spark Structured Streaming: How We Evolved The Architecture To Meet Real-World Needs.....	4156
<i>Anish Shrigondekar, Boyang Jerry Peng, Siying Dong, Jungteak Lim, Huanli Wang, Bo Gao, Eric Marnadi, Jason Teoh, B. Micheal Okutubo, Livia Zhu, Yuchen Liu, Chloe Xia, Abhay Bothra, Sagar Mittal, Karthik Ramasamy, Michael Armburst, Fatih Emekci, Shan Shan Huang, Shrikanth Shankar, Indrajit Roy</i>	
Real-time SQL Plan Management in Oracle	4169
<i>Sunil Chakkappen, Mohamed Ziauddin, Hong Su, Shreya Kunjibettu, Nigel Bayliss</i>	
SQL-Native Vector Search at Billion Scale in Presto.....	4182

Zhichen Xu, Ge Gao, Ke Wang, Aakash Deep, Junjie Qi, Jort Gemmeke, Amit Dutta, Xiaoxuan Meng, Trang Nguyen, Jiayin Yan, Xiao Du, Vivek Gaur, Kaushik Ravichandran, Vishal Gandhi

SERON: Smart Query Router for Multi-Primary Cloud-Native Databases with Shared Storage..... 4195
Yuwei Huang, Guoliang Li, Jinyu Zhang

A Unified Bandwidth Orchestration Framework for Hierarchical Data Storage Systems..... 4209
Ji Zhang, Li Liu, André Brinkmann, Giovanni Cherubini, Zhu Xiangyu, Shai Bergman, Ke Zhou

IMLane: Composable Framework for Efficient AI Function Execution in Database Engine..... 4223
Chenyang Zhang, Linjun Lu, Qingfeng Pan, Chen Xu, Xianzhong Cao, Quanqing Xu, Chuanhui Yang

SOS: A High-Performance Distributed Key-Value Store for Large-Scale Online Services..... 4237
Yingxin Li, Kai Liu, Han Xie

A Tree-Structured Two-Phase Commit Framework for OceanBase: Optimizing Scalability and Consistency
..... 4250
Quanqing Xu, Chen Qian, Chuanhui Yang, Fanyu Kong, Guixiang Liu, Fusheng Han, Zixiang Zhai

Private Data Collections: Enabling Data Privacy in Permissioned Blockchains 4263
Senthilnathan Natarajan, Manish Sethi, Dave Enyeart, Yacov Manevich, Artem Barger

OmniTable: A Unified Wide-Table System for Petabyte-Scale LLM Data Curation and Exploration 4276
Yuzhuo Fu, Xiangchun Wang, Chao Huang, Liyi Wang, Binwei Zeng, Yuhao Wang, Taotao Nie, Dongke Hu, Wang Hong, Jiayi Wang, Wenwen Cui, Zhuyan Zhou, Yushun Guo, Yuhao Xing, Jiaxin Lian, Peng Lin, Qing Cui, Wenhui Shi

The Evolution of Storage in Apache AsterixDB 4290
Murtadha Al Hubail, Ali Alsuliman, Wail Alkowiileet, Michael Blow, Michael Carey, Peeyush Gupta, Ian Maxon, Daniel Nagy, Ritik Raj, Shahrzad Shirazi, Hussain Towailab

Towards Industrial-Scale Parametric Query Optimization 4303
Songsong Mo, Quanqing Xu, Xuchen Ding, Yue Zhao, Zhifeng Bao, Chuanhui Yang, Gao Cong

BiLink: Bidirectional Meta-paths for Link Discovery in Billion-Scale Heterogeneous Graphs 4316
Jinquan Hang, Zhiqing Hong, Xinyue Feng, Desheng Zhang, Haotian Wang, Guang Wang

Bridging NL2SQL for CAN Signal Analytics at NIO: From Externalized Schemas to CTE Pipelines 4330
Wei Zhou, Haoyu Zhao, Ruiguang Zhong, Shaoli Yu, Xuanhe Zhou, Xue Yang, Xiaosong Jia, Guoliang Li, Shengzhong Liu, Kang Zhang, Jie Wu, Fan Wu

VikingMem: A Memory Base Management System for Stateful LLM-based Applications 4344
Jiajie Fu, Junwen Chen, Mengzhao Wang, Aoxiang He, Maojia Sheng, Xiangyu Ke, Yifan Zhu, Yunjun Gao

Ultron: History-Based Query Optimization at Databricks..... 4358
Supun Nakandala, RK Korlapati, Eric Liang, Yunjia Zhang, Yuhao Zhang, Andrey Gubichev, Kelvin Jiang, Yannis Sismanis, Mostafa Mokhtar, Sid Taneja, Vuk Ercegovic, Amir Hormati, Amit Shukla, Michalis Petropoulos, Reynold Xin, Ryan Marcus

One Ring to Shuffle Them All: Scalable Intra-Process Data Redistribution with Ring-Buffer Shuffle in Redpanda Osla..... 4372
Adam Szymański, Tyler Akidau

Lakebase: Serverless Postgres over Open Lake Storage 4385

Jasraj Dange, Andrei Dragus, Ali Ghodsi, Haoyu Huang, Yihe Huang, Stas Kelvich, Heikki Linnakangas, Hans Norheim, Ippokratis Pandis, Nikita Shamgunov, Em Sharnoff, John Spray, Zhou Sun, Reynold Xin, Matei Zaharia

From Presto to Prestissimo: A Velox-Powered Modernization Journey 4399
Amit Dutta, Pedro Pedreira, Xiaoxuan Meng, Masha Basmanova, Orri Erling, Sergey Pershin, Jialiang Tan, Ge Gao, Ke Wang, Zac Wen, Krishna Pai, Xiao Du, Natasha Sehgal, Heidi Han, Artem Selishchev, Ke Wang, Sebastiano Peluso, Chandrashekhar Singh, Sreeni Viswanandha, Aditi Pandit, Deepak Majeti, Tim Meehan, Christian Zentgraf, Aakash Deep, Rohit Jain, Sridhar Anumandla, Kaushik Ravichandran, Vivek Gaur

Why We Created Yet Another Memory Framework: Understanding MGA’s Role in Next-Gen Database Systems 4413
Vikramraj Sitpal, Pei Li, Shubham Kumar, Somansh Reddy Satish, Ravi Thammaiah, Nagarajan Muthukrishnan

SalesforceDB: A Cloud-native Multi-tenant OLTP Database 4426
Atish Agrawal, Vaibhav Arora, Gary Baker, Milena Bergant, Benjamin Busjaeger, Subho Chatterjee, Terry Chong, Roy Chowdhuri, Doug Doole, Haiyan Du, Thomas Fanghaenel, Pat Helland, Jv Jujjuri, Jim Mace, Ashish Mahajan, Jamie Martin, Nicolas Michael, Ayman Mobarak, Steven Raspudic, Randy Spalten, Simon Wong, Nat Wyatt

IORM: Hierarchical I/O Governance for Thousands of Consolidated Databases on Oracle Exadata 4440
Rajarshi Chowdhury, Akshay Shah, Zakaria Alrmailh, Chenhao Guo, Anubhav Singh, Sue Lee

Nova: A Multi-Purpose Vector Engine for Low-Latency, Multi-Tenant, and Cross-Table Hybrid Retrieval 4453
Dechuang Chen, Bing Chen, Jianwei Lu, Zhiwei Wu, Zeyu Yang, Qiang Gu, Fangyuan Zhang, Sibao Wang, Zeyuan Yu, Caihua Yin, Wenchao Zhou, Feifei Li

Narwhal: Breaking the Local Boundary via Disaggregated Memory Scheduling for Alibaba AnalyticDB 4466
Jian Zhou, Jiachi Zhang, Yang Zhang, Yongbo Wang, Ruonan Guo, Yuemin Wu, Zeguang Chang, Shuai Wang, Ye Yin, Wenchao Zhou, Liang Lin

PolarKV: Tier Locally, Serve Globally—A KV Cache over Cloud Memory and Storage 4480
Yingqiang Zhang, Xinjun Yang, Hao Chen, Feifei Li, Shuwen Wang, Chuan Sun, Sai Wu, Ninglong Weng

Demonstrations

QueryBrew: System-Agnostic SQL-to-SQL Query Optimization 4494
Tobias Schmidt, Maximilian Reif, Altan Birlir, Thomas Neumann

FedBridge: A Federated Query Engine over Embedding-Heterogeneous Vector Databases 4498
Yuxiang Wang, Ruixi Hu, Ziyuan He, Zhilin Liang, Xinyu Zhao, Yongxin Tong

HaDA: Empower Database Experts with Data Dependencies 4502
Daniel Lindner, Jonas Baltruschat, Kisung Park, Taehyung Lee, Heesik Shin, Daniel Ritter, Felix Naumann

OMBench: Taming Data Management Requirements in Cloud-Native Applications 4506
Rodrigo Laigner, Xikun Jiang, Boris Döder, Yongluan Zhou

How Out-of-Bounds Are Your Cardinality Estimates? 4510
Mihail Stoian, Tiemo Bang, Hangdong Zhao, Jesús Camacho-Rodríguez, Yuanyuan Tian, Andreas Kipf

AtomSQL: Interactive Disambiguation of NL-to-SQL via User-Guided Atom-Level Alignment 4514

Aritra Mazumder, Parth Desai, Fuheng Zhao, Anna Fariha

Credo: Declarative Control of LLM Pipelines via Beliefs and Policies 4518
Duo Lu, Andrew Crotty, Uğur Çetintemel

SpatialSQL: A Multi-Agent System for Interactive and Observable Spatial Text-to-SQL..... 4522
Yang Song, Yu Sun, Xieyang Wang, Qian Tao, Tengfei Xu, Jianqiu Xu, Shuyue Wei, Yongxin Tong

DataMagic: Transforming Tabular Data into Data Insight Video..... 4526
Yupeng Xie, Chen Ma, Zhenyang Wang, Liangwei Wang, Jiayi Zhu, Chuxuan Zeng, Zhouan Shen, Boyan Li, Yuyu Luo

QuWARTS: Query Workload Aware Relational Table Synthesis from Unstructured Text..... 4530
Aritra Mazumder, Whanhee Cho, Anna Fariha

Hamilton – Interactive Ontology Learning 4534
Lukas Laskowski, Felix Draxler, Michael Hladik, Fabian Panse, Jan Portisch, Padhraic Smyth, Felix Naumann

SemPiper: Interactive Code Synthesis for Semantic Operators in Machine Learning Pipelines 4538
Olga Ovcharenko, Luciano Duarte, Sebastian Schelter

GRAIL: AI translation for scientists application workflow on satellite data 4542
Zhuocheng Shang, Ahmed Eldawy

GraphAlg Playground: An Online Platform for Learning and Experimenting with the GraphAlg Language 4546
Daan De Graaf, Robert Brijder, Soham Chakraborty, George Fletcher, Bram Van De Wall, Nikolay Yakovets

GraphContainer: A Unified Platform for Comparing and Debugging Graph RAG Methods 4550
Seonho An, Chaejeong Hyun, Min-Soo Kim

A Demo of Interactive Thematic Data Collection on the Live Web 4554
Michael West, Eduard Dragut

CypherLens: An Interactive Demo for Evaluating and Diagnosing NL-to-Cypher Systems..... 4558
Zahra Mohammadi, Ali Ansari, Longin Latecki, Eduard Dragut

Guixu: Valuation-Driven Data Discovery for Autonomous AI Agents with On-Chain Attestation..... 4562
Yifan Wu, Yuchen Peng, Jiaqi Chai, Yufei Qian, Xilin Li, Ke Chen, Lidan Shou

Qr-Hint: Formally Verified and AI-Explained SQL Tutoring 4566
Zhouyi Jiang, Yihao Hu, Amir Gilad, Sudeepa Roy, Jun Yang

DataMosaic: An Interactive Demonstration of Constraint-Driven Document-to-Database Construction 4570
Zhengxuan Zhang, Zhuowen Liang, Haixun Wang, Nan Tang

Data Lake for AI-enhanced Quantum Error Correction 4574
Aadi Patwardhan, Max De Groot, Andrei Iliescu, Aditya Shankar, Yuandou Wang, Wenbo Sun, Floris Geerts, Rihan Hai

LayoutPilot: A Lakehouse Physical Design Advisor 4578
Guanli Liu, Andreas Kipf, Renata Borovica-Gajic

Text-to-SQL Evaluation Toolkit..... 4582

Oktie Hassanzadeh, Yotam Perlitz, Nhan Pham, Tanvi Kaple, Karolina Żróbek, Long Vu, Michael Glass, Dharmashankar Subramanian, Mohammadreza Pourreza, Davood Rafiei

Dominance-Based Data Reduction for Package Queries.....	4586
<i>Vasileios Vittis, Azza Abouzied, Peter J. Haas, Alexandra Meliou</i>	
PGShield: Maintaining Path Constraints in Property Graphs	4590
<i>Christopher Spinrath, Angela Bonifati, Rachid Echahed</i>	
Verified LLM-Based Query Rewriting for Microsoft SQL Server	4594
<i>Kukjin Lee, Anshuman Dutt, Vivek Narasayya, Surajit Chaudhuri</i>	
A Demonstration of KAFY: An Extensible and Scalable Transformers-Based System for Trajectory Data Analysis.....	4598
<i>Youssef Hussein, Mohamed F. Mokbel</i>	
Demonstrating GenDB: Instance-Optimized and Customized Query Processing Code Generation via LLM Agents.....	4602
<i>Jiale Lao, Immanuel Trummer</i>	
MemLens: A Value-Aware Memory Management System with Interactive Analytics for LLM-based Agents	4606
<i>Shuyue Wei, Chang Liu, Zimu Zhou, Yongxin Tong, Lizhen Cui</i>	
Demonstrating EarthLake: A Model Lake System for Earth Observation Foundation Model Management	4610
<i>Binger Chen, Haralampos Gavrilidis, Luca Gaedicke, Tacettin Emre Bök, Matthias Boehm, Ziawasch Abedjan, Begüm Demir, Volker Markl</i>	
DeQL Studio: Declarative Decision-Making over Relational Data	4614
<i>Matteo Brucato, Fjodor Kholodkov, Soren Little, Jakob Mayer, Duc Nguyen</i>	
Minimal Data Cleaning for Model Training by MinPrep.....	4618
<i>Cheng Zhen, Prayoga Fnup, Nischal Aryal, Arash Termehchy, Alireza Aghasi</i>	
A Demonstration of Continuous Lifelong Conflict-Aware AGV Routing with Kinematic Constraints.....	4622
<i>Ruizhong Wu, Tianqi Zhang, Yunjie Huang, Lei Li</i>	
FilterPilot: An Interactive Assistant for Adapting Filtering Predicate to Table Content.....	4626
<i>Shiwen Wu, Yuqi Wang, Yue Pang, Xiaofang Zhou</i>	
GraphAgent: An Effective Knowledge-Guided GNN Model Selection System	4630
<i>Mingtao Zhang, Haoyang Li, Yuming Xu, Nicole Hu, Peng Cheng, Chen Jason Zhang, Qing Li, Lei Chen</i>	
CADENZA in Action: Breaking the Monolith with Intent-Dependent Plan Spaces for Semantic Queries .	4634
<i>Jaehyun Ha, Yongjoo Park, Wook-Shin Han</i>	
MCAD: Multivariate Correlation Anomaly Data Generator	4638
<i>Christos C. Papadopoulos, Alkis Simitsis, Torben Bach Pedersen</i>	
Carnot: Interpretable, Interactive, and Optimized Execution of Deep Research Queries.....	4642
<i>Matthew Russo, Yash Agarwal, Tianyu Li, Zhuohan Gu, Michael Cafarella, Omar Khattab, Tim Kraska, Samuel Madden</i>	
NiceT: Named Entity Cleaning and Enhancement with Human-in-the-loop.....	4646

Garima Gaur, Ioana Manolescu, María Mellado Tenorio, Madhulika Mohanty, Moritz Sommer

A Portable Middleware for Plan-Based Adaptive Query Processing..... 4650
Pei Mu, Bitu Asoodeh, Anderson Chaves Carniel, Amir Shaikhha

Policy-Aware Federated Query Orchestration Across Energy Data Spaces and Edge AI Services..... 4654
Vasile Ofrim, Mihai Daian, Dragos Lazea, Liana Todorean, Anca Hangan, Tudor Cioara

Into the CUDA Multiverse of Runtime Compilation: Exploring Kernel Fusion Strategies for GPU Database Systems 4658
Maha Alwahibi, Bachana Jangebashvili, Maximilian E. Schüle

CoBLOC - Continuous Block Level Fairness on Data Streams 4662
Subhadeep Ghosh, Zhihui Du, Angela Bonifati, Manish Kumar, David Bader, Senjuti Basu Roy

Demonstrating Samsara: A Super Optimizer for Multimodal Stream Processing 4666
Alessandro Ferri, Uélison Jean Lopes Dos Santos, Szilard Nistor, Pratyush Agnihotri, Manisha Luthra, Carsten Binnig

GDPView: An Interactive Interface that Unifies Variants of Deletion Propagation (DP) 4670
Yihan Luo, Neha Makhija, Wolfgang Gatterbauer

Exploring the Benefits of Just-in-time Model Replacement 4674
Nils Strassenburg, Boris Glavic, Tilmann Rabl

SuperDRI: Graph-Guided Autonomous Troubleshooting for Cloud Databases 4678
Yiwen Zhu, Joyce Cahoon, Qiushi Bai, Anna Pavlenko, Divya Vermareddy, Meina Wang, Mathieu Demarne, Wenjing Wang, Swati Bararia, Hannah Lerner, Katherine Lin, Subru Krishnan, Miso Cilimdžić

MAPPipe: A System for Bridging Efficiency and Quality in Data Preprocessing via Knowledge-Augmented Structural Pruning..... 4682
Xiaoou Ding, Haifeng Cheng, Yanshuo Liu, Chen Wang, Hongzhi Wang

LabelScope: Context-Aware Discovery of Label Outliers in Radio Astronomy Signals..... 4686
Rabeya Hossain, Brian Mason, Ryan Loomis, Jeff M. Phillips, El Kindi Rezig

SubCure: A System for Stress-Testing Causal Conclusions Using Cardinality Repairs 4690
Shira Raviv, Ben Hechler, Yarden Gabbay, Haoquan Guan, Shaull Almagor, El Kindi Rezig, Brit Youngmann, Babak Salimi

Demonstrating Indexing for Near-Sorted Data 4694
Aneesh Raman, Philip Chindris, Anwesha Saha, Teona Bagashvili, Manos Athanassoulis

cuRPQ+: A System for Interactive Path-Aware Querying Beyond Plain CRPQs 4698
Seohyeon Kim, Sungwoo Park, Min-Soo Kim

ReStore-in-Action: Adaptive Multi-Tier Storage Management via Intelligent Data Migration Policy 4702
Shadman Saqib Eusuf, Tianru Zhang, Teona Bagashvili, Manos Athanassoulis, Tarikul Islam Papon

PROXAI: Interactive Provenance-Aware Debugging of Machine Learning Pipelines 4706
Pasquale Leonardo Lazzaro, Elia Guglielmi, Paolo Missier, Riccardo Torlone

Painless and Efficient Scaling of Pandas Programs: Demo 4710
Bhushan Pal Singh, Priyesh Kumar, Chiranmoy Bhattacharya, Utkarsh Shetye, Manish Kumar, S. Sudarshan

Persona-Conditioned Query Generation for Selectivity-Controlled Workloads	4714
<i>Angjela Davitkova, Sebastian Michel</i>	
Demonstration of BobFlow: Human-Agent Collaboration Using Dataflows on Data Science Tasks	4718
<i>Jiadong Bai, Yicong Huang, Chen Li</i>	
Demonstrating TACT: Tunable Accuracy-Cost Trade-offs for Semantic Image Filtering	4722
<i>Tharushi Jayasekara, Immanuel Trummer</i>	
RSR-core: A High-Performance Engine for Low-Bit Matrix-Vector Multiplication.....	4726
<i>Mohsen Dehghankar, Abolfazl Asudeh</i>	
TexBench: A Unified Benchmarking Suite for Shifting Workloads	4730
<i>Abhishek Chanda, Shubham Kaushik, Artem Lavrov, Subhadeep Sarkar</i>	
ARGO: An Interactive Data Governance System for Machine Learning via Hierarchical Reinforcement Learning.....	4734
<i>Shuang Hao, Jun Wei, Yingqi Mu, Diwen Cheng</i>	
Demonstrating TOFFEE: A Learned System for Synthesizing Data Agent Trajectories at Scale	4738
<i>Ziting Wang, Yin Li, Zuhao Yang, Xiuchang Li, Jiale Bai, Gao Cong</i>	
ImputePilot: A Graphical Model Selection Toolkit for Time Series Imputation	4742
<i>Yuanyuan Yao, Zhexin Jin, Lu Chen, Anthony K. H. Tung, Mourad Khayati</i>	
LLM-CER: An Interactive System for In-context Clustering-based Entity Resolution with Large Language Models	4746
<i>Haoyu Wang, Haitong Tang, Jiajie Fu, Arijit Khan, Sharad Mehrotra, Xiangyu Ke, Yunjun Gao</i>	
ECO-Hadoop: Energy Consumption Optimization System for Apache Hadoop.....	4750
<i>Zemin Chao, Xiao Chen, Jianhe Cen, Yumeng Fu, Hongzhi Wang</i>	
QueryCraft: A Natural Language-Driven NoSQL Database Querying System Powered by Large Language Models	4754
<i>Jinwei Lu, Zhiqian Qin, Chen Zhang, Haodi Zhang, Yuanfeng Song, Raymond Chi-Wing Wong</i>	
ChatQPT: Towards Conversing with Relational Query Engines.....	4758
<i>Hui Li, Sourav S Bhowmick, Baochao Xu, Zhenning Shi, Xiyue Gao, Jingjing Li</i>	
LIMA: Denial Constraint Discovery in Large Dynamic Databases	4762
<i>Albert Martin, Eduardo C. De Almeida, Oscar Romero, Anna Queralt</i>	
DA-Studio: An Agentic System for End-to-End Data Analysis.....	4766
<i>Yizhe Liu, Shaolei Zhang, Ju Fan</i>	
Provleipsis4j: Querying Future Graphs in Neo4j	4770
<i>Evangelos Iliadis, Christos Gkartzios, Evaggelia Pitoura</i>	
GALACTICA: An Interactive System for Local and Global Counterfactual Explanations for Time-series Clustering.....	4774
<i>Christos Fragkathoulas, Eleni Psaroudaki, Themis Palpanas, Evaggelia Pitoura</i>	
The Case for Multi-Version Experimental Evaluation (MVEE)	4778
<i>Simon Jörz, Felix Schuhknecht</i>	

iPDB: SQL with ML and LLM Predicates (Towards a Database Engine for AI)	4782
<i>Udesh Kumarasinghe, Tyler Liu, Ahmed Mahmood, Chunwei Liu, Walid G. Aref</i>	
Tursio for Credit Unions: Structured Data Search with Automated Context Graphs	4786
<i>Shivani Tripathi, Ravi Shetye, Shi Qiao, Alekh Jindal</i>	
Cedar: A Columnar LSM-Engine for Temporal Property Graphs	4790
<i>Yang Wang, Xuelian Lin, Jinghe Song, Jianyong Zhu, Yu Zhao, Shuai Ma</i>	
MoDora: A Multimodal Document AI Assistant Harness	4794
<i>Yukai Wu, Bangrui Xu, Shaoli Yu, Zirui Tang, Xuanhe Zhou, Yeye He, Bin Wang, Conghui He, Guoliang Li, Fan Wu</i>	
ARCADE: A Real-Time Data System for Hybrid and Continuous Query Processing across Diverse Data Modalities	4798
<i>Jingyi Yang, Songsong Mo, Jiachen Shi, Zihao Yu, Kunhao Shi, Gao Cong, Xuchen Ding</i>	
Q-ACER: Query Aggregate Constraint Efficient Repair System	4802
<i>Vaishnavi Deshpande, Seokki Lee, Shatha Algarni, Boris Glavic, Adriane Chapman</i>	
privJedAI: A Unified Open-Source Pipeline for Privacy-Preserving Record Linkage	4806
<i>Lefteris Stetsikas, Dimitrios Karapiperis, George Papadakis, Manolis Koubarakis</i>	
Demonstration of Doc*: Access Control for Information-Theoretically Secure Data	4810
<i>Komal Kumari, Yin Li, Sharad Mehrotra, Shantanu Sharma, Venkata Sreeramakavacham</i>	
Hypothesis Exploration with LLM-driven Intelligent Orchestration System	4814
<i>Bruno G. Tavares Dos Santos, Vicente Nejar De Almeida, Eduardo Ribeiro, Sihem Amer-Yahia, João Luiz Dhl Comba</i>	
Enzyme Demo: Incremental View Maintenance for Data Engineering	4818
<i>Yuhong Chen, Ritwik Yadav, Min Yang, Supun Abeysinghe, Manuel Ung, Jeffrey Helt, Michael Armbrust, Shrikanth Shankar</i>	
LORY-J: Location-aware Join Discovery from Data Lakes	4822
<i>Lu Xing, Mir Mahathir Mohammad, Walid G. Aref, El Kindi Rezig</i>	
QFusion: A Demonstration of Boundary-Aware Fusion Planning and Execution for Large-Scale QUBO Optimization	4826
<i>Hanwen Liu, Ibrahim Sabek</i>	
TurboLynx in Action: A Schemaless Graph Engine for General-Purpose Analytics	4830
<i>Taesung Lee, Jaehyun Ha, Byungchul Tak, Wook-Shin Han</i>	
SAGA: Synthetic Agentic Graph Architecture for Temporal Benchmark Generation	4834
<i>Jiacheng Ding, Xiaofei Zhang</i>	
ATLAS: Adaptive Text-to-SQL with Lifecycle-Aware Self-Maintaining Context	4838
<i>Qing Zhang, Shijing Hu, Zhihui Lu</i>	
DBCooker: A Database Kernel-Specialized Coding Agent Harness	4842
<i>Wei Zhou, Xuanhe Zhou, Qikang He, Guoliang Li, Bingsheng He, Chao Peng, Yun Lin, Fan Wu</i>	
ParSEval: Interactive Counterexample-driven Evaluation for Text-to-SQL	4846
<i>Chunyu Chen, Zhengjie Miao, Yong Zhang, Jiannan Wang</i>	

DialectEvo: Bridging SQL Dialects with Execution-Guided Knowledge Refinement..... 4850
Jianling Gao, Chongyang Tao, Liu Yang, Xiya Jiang, Shuai Ma

TrustBCI: A Trustworthy Platform for Collaborative Brain-Computer Interface Data Exchange 4854
Jianyu Yang, Jiannan Wang, Guoliang Li

Hard to Read, Hard to Maintain? Refactoring Complex SQL into Wvlet Pipelines 4858
Kaina Anderson, Yohanes Yohanie Fridelin Panduman, Makoto Onizuka, Taro L. Saito

Tutorials

Hardware Acceleration for Spatial Databases: Building Next-Generation Query Engines Using Ray Tracing Cores..... 4862
Rubao Lee, Liang Geng, Xiaodong Zhang

Bridging LLMs and Database Systems: A Deep Dive into Enhanced Relational Operators..... 4867
Tianjing Zeng, Yin Lin, Rong Zhu, Yunxiang Su, Zhongjun Ding, Bolin Ding, Jingren Zhou

Are We There Yet? A Grand Tour of ETA Queries 4872
Areeg Mostafa, Mohamed Hemdan, Mohamed F. Mokbel

From Human-Graph to Agent-Graph Interaction: State of the Art and Future Directions..... 4878
Sourav S Bhowmick, Byron Choi, Arijit Khan

Instance-Optimal Acyclic Joins: From Theory to Systems..... 4884
Paraschos Koutris, Stijn Vansummeren, Qichen Wang, Yisu Remy Wang, Xiangyao Yu

Data Management for Agentic Memory: Foundations, Systems, and Challenges 4888
Guoliang Li, Jiaqi Tian, Xuanhe Zhou

Advances of Query Processing in Vector Databases 4893
Jiadong Xie, Yingfan Liu, Jeffrey Xu Yu

Graph-Based Retrieval-Augmented Generation: Applications, Challenges, Solutions, and Opportunities 4900
Angela Bonifati, Yixiang Fang, Sibor Wang, Yingli Zhou

How Can Quantum Computing and Databases Meet “in Practice”? From Algorithms to Systems 4906
Hanwen Liu, Yiran Li, Ibrahim Sabek, Xuanhe Zhou, Hans-Arno Jacobsen

Interoperability in Healthcare: A Primer and New Frontiers..... 4912
Andra Ionescu, Paris Carbone, Sebastiaan Meijer, Christos Koutras, Juliana Freire, Katja Hose, Jayanth Raghothama

Graph Foundation Models: State of the Art and Future Directions 4918
Alexander Zhou, Haoyang Li, Lei Chen

CRDTs and Databases: A Hands-On Tutorial..... 4922
Paulo Sérgio Almeida, Nuno Faria, José Pereira

Private LLM Inference with Homomorphic Encryption..... 4927
Lawrence Lim, Divyakant Agrawal, Amr El Abbadi

An Extended Tutorial and Vocabulary for Relational Language Design in an Era of AI-Assisted Query Generation	4932
<i>Wolfgang Gatterbauer</i>	

From Big Time Series Forecasting to Foundation Models for Structured Data	4939
<i>Xiuyan Zhang, Abdul Fatir Ansari, Christos Faloutsos, George Karypis, Yuyang Wang</i>	

Panels

Publish or Ship? The Research-to-Production Gap in Data Systems	4944
<i>Subarna Chatterjee</i>	

Data Agents: Rethinking Data Systems in the AI Agent Era	4946
<i>Guoliang Li, Yuyu Luo</i>	

HW-SW Co-Design for Databases in the Age of AI-for-Systems	4950
<i>Monica Chiosa</i>	

Endowment Awards

The Dataflow Model Revisited	4953
<i>Tyler Akidau, Rafael Fernández-Moctezuma, Reuven Lax, Daniel Mills</i>	

Future-Proof Data Systems	4965
<i>Jana Giceva</i>	

Semantic Data Systems: From Data Management to Data Understanding	4973
<i>Juliana Freire</i>	

Keynotes

The Three Golden Ages of Database Engineering: From SIGMOD'85 to the Agentic Era	4981
<i>Reynold Xin</i>	

Efficient and Reliable Systems for Building Foundation Models at Scale	4982
<i>Jingren Zhou</i>	

Special Event

Return to the Boston Area After 50+ Years: What can we learn from VLDB 1975?	4983
<i>Philip A. Bernstein, Lei Chen, David Dewitt, Sebastian Link, Volker Markl, Renée J. Miller, C. Mohan, Michael Stonebraker, Divesh Srivastava</i>	

PVLDB ORGANIZATION AND REVIEW BOARD - Vol. 19

Editors in Chief of PVLDB

Yanlei Diao (Ecole Polytechnique and Amazon Web Services)
Xiaokui Xiao (National University of Singapore)

Associate Editors of PVLDB

Alkis Polyzotis (Databricks)
Alkis Simitsis (Athena Research Center)
Andrew Pavlo (Carnegie Mellon University)
Angela Bonifati (Univ. of Lyon)
Anh Dinh (Deakin University)
Antonis Deligiannakis (Technical University of Crete)
Ashraf Aboulmaga (University of Texas at Arlington)
Badrish Chandramouli (Microsoft Research)
Bin Cui (Peking University)
Bogdan Cautis (University of Paris-Saclay)
Ce Zhang (University of Chicago)
Elena Ferrari (University of Insubria, Varese)
Evaggelia Pitoura (Univ. of Ioannina)
Fatemeh Nargesian (University of Rochester)
Guoliang Li (Tsinghua University)
Hong Cheng (Chinese University of Hong Kong)
Immanuel Trummer (Cornell University)
Jana Giceva (TU Munich)
Jennie Rogers (Northwestern University)
Jian Pei (Duke University)
Jianliang Xu (Hong Kong Baptist University)
Karima Echihiabi (Mohammed VI Polytechnic University)
Katja Hose (TU Wien)
Khuzaima Daudjee (University of Waterloo)
Kunsoo Park (Seoul National University)
Kyuseok Shim (Seoul National University)
Laks Lakshmanan (The University of British Columbia)
Lei Cao (University of Arizona)
Lei Chen (Hong Kong University of Science and Technology)
Li Xiong (Emory University)
Mauro Sozio (Luiss University of Rome)
Meihui Zhang (Beijing Institute of Technology)
Melanie Herschel (Nanyang Technological University)
Michael Böhlen (University of Zurich)
Nan Tang (HKUST (GZ))
Niv Dayan (University of Toronto)
Norman May (SAP SE)
Raul Castro Fernandez (The University of Chicago)
Raymond Chi-Wing Wong (Hong Kong University of Science and Technology)
Reynold Cheng (The University of Hong Kong, China)
Ruben Mayer (University of Bayreuth)
Semih Salihoglu (University of Waterloo)
Shimin Chen (Chinese Academy of Sciences)
Sibo Wang (The Chinese University of Hong Kong)
Sourav S Bhowmick (Nanyang Technological University)
Spyros Blanas (The Ohio State University)
Steven Whang (KAIST)

Sudipto Das (Amazon Web Services)
Umar Farooq Minhas (Apple)
Wei Wang (ByteDance)
Wook-Shin Han (POSTECH)
Yannis Velegrakis (Utrecht University)
Yanyan Shen (Shanghai Jiao Tong University)
Yi Chen (NJIT)
Yuncheng Wu (Renmin University of China)
Zhifeng Bao (The University of Queensland)
Zi Huang (University of Queensland)

Publication Editors

Andrea Mauri (Lyon 1 University)
Giovanni Simonini (University of Modena and Reggio Emilia)
Siqiang Luo (Nanyang Technological University)
Stefan Halfpap (TU Berlin)
Subhadeep Sarkar (Brandeis University)

PVLDB Managing Editor

Jun Yang (Duke University)

PVLDB Advisory Board

Alexandra Meliou (University of Massachusetts Amherst)
Divesh Srivastava (AT&T Labs - Research)
Fatma Ozcan (Google)
Felix Naumann (HPI)
Katja Hose (Technische Universität Wien)
Lei Chen (Hong Kong University of Science and Technology)
Matthias Boehm (Technische Universität Berlin)
Meihui Zhang (Beijing Institute of Technology)
Nesime Tatbul (Intel Labs and MIT)
Sourav S. Bhowmick (Nanyang Technological University)
Themis Palpanas (Universite Paris Cite)
Torsten Grust (University of Tuebingen)
Vanessa Braganholo (Universidade Federal Fluminense)
Wolfgang Lehner (TU Dresden)
Xiaokui Xiao (National University of Singapore)
Xin Luna Dong (Meta)
Yanlei Diao (Ecole Polytechnique, Amazon Web Services)

Review Board

Abhishek Roy (Microsoft)
Ahmed Aly (Google)
Ahmed Eldawy (University of California, Riverside)
Ahmed Mahmood (Google)
Alberto Lerner (Computing Flows GmbH)
Alberto Sonnino (MystenLabs & University College London)
Alexander Artikis (University of Piraeus)
Alexander Böhm (SAP SE)
Alexander Krause (TU Dresden)
Amarnadh Sai Eluri (Google)
Amelie Marian (Rutgers University)
Amine Mhedhbi (Polytechnique Montréal)
Amir Shaikhha (University of Edinburgh)
Amr Magdy (University of California Riverside)
Andrea Mauri (Lyon 1 University)
Andreas Züfle (Emory University)
Antonio Boffa (EPFL)
Arash Termehchy (Oregon State University)
Arijit Khan (Bowling Green State University)
Aristides Gionis (KTH Royal Institute of Technology)
Arkaprava Saha (University of Tsukuba)
Arnab Phani (TU Berlin)
Arnd Christian König (Microsoft)
Avinash Kumar (Google)
Baihua Zheng (Singapore Management University)
Berthold Reinwald (IBM Research - Almaden)
Bin Yang (Aalborg University)
Bo Tang (Southern University of Science and Technology)
Bo Zhao (Aalto University)
Bobbi W Yogatama (NVIDIA)
Bojan Karlaš (Harvard University)
Boris Glavic (University of Illinois Chicago)
Brian Kroth (Microsoft)
Brit Youngmann (Technion)
Calisto Zuzarte (IBM)
Carlo Curino (Microsoft)
Chang Ge (University of Minnesota)
Chao Zhang (University of Waterloo)
Chen Zhang (The Hong Kong Polytechnic University)
Cheng Long (Nanyang Technological University)
Chenghao Lyu (Amazon)
Chengkai Li (The University of Texas at Arlington)
Chengliang Chai (Beijing Institute of Technology)
Chenhao Ma (The Chinese University of Hong Kong, Shenzhen)
Christan Grant (University of Florida)
Christos Doukeridis (University of Piraeus)
Christos Koutras (New York University)
Chrysanthi Kosyfaki (The University of Hong Kong)
Chuan Lei (Oracle)
Chunwei Liu (MIT CSAIL)
Cindy Chen (University of Massachusetts Lowell)
Cong Yue (National University of Singapore)
Daniel Kang (UIUC)
Daniel Ritter (SAP)
Davood Rafiei (University of Alberta)
Demetris Zeinalipour (University of Cyprus)
Dimitrios Skoutas (Athena Research Center)

Dixin Tang (The University of Texas, Austin)
Domagoj Vrgoč (PUC)
Dominik Durner (CedarDB)
Donatella Firmani (Sapienza University)
Donato Tiano (Università degli Studi di Modena e Reggio Emilia)
Dong Deng (Alibaba Cloud Computing)
Dong Xie (Penn State University)
Dumitrel Loghin (OKX)
Eduard Dragut (Temple Univ.)
Egemen Tanin (University of Melbourne)
El Kindi Rezig (University of Utah)
Eleni Tzirita Zacharatou (Hasso Plattner Institute (HPI))
Eliana Pastor (Politecnico di Torino)
Elisa Bertino (Purdue University)
Enzo Veltri (Università della Basilicata)
Ergute Bao (Mohamed bin Zayed University of Artificial Intelligence)
Eser Kandogan (Megagon Labs)
Essam Mansour (Concordia University)
Farnoush Banaei-Kashani (University of Colorado Denver)
Felix M Schuhknecht (Johannes Gutenberg University Mainz)
Florian Hahn (University of Twente)
Florin Rusu (UC Merced)
Francesco Gullo (University of L'Aquila)
Fuad Jamour (Amazon Web Services)
Gao Cong (Nanyang Technological University)
Garima Gaur (Inria Saclay, Ecole Polytechnique)
Gengrui Zhang (Concordia University)
George Christodoulou (TU Delft)
George Papadakis (University of Athens)
George Papastefanatos (ATHENA Research Center)
Georgios Siachamis (Inria)
Georgios R Theodorakis (Nvidia)
Giorgio Vinciguerra (Università di Pisa)
Giuseppina Andresini (University of Bari Aldo Moro)
Goce Trajcevski (Iowa State University)
Grigorios Loukides (King's College London)
Guido Moerkotte (University of Mannheim)
Guillaume Lachaud (Ecole Polytechnique)
Haibo Hu (Hong Kong Polytechnic University)
Haipeng Dai (Nanjing University)
Hannes Voigt (Neo4j)
Hantian Zhang (Google)
Hanzhi Wang (University of Copenhagen)
Haralampos Gavrilidis (BIFOLD & Technische Universität Berlin)
Haridimos Kondylakis (FORTH-ICS & Computer Science Department, University of Crete)
Harry Kai-Ho Chan (The University of Sheffield)
Hazar Harmouch (University of Amsterdam)
Herald Kllapi (Google Inc)
Hongzhi Wang (Harbin Institute of Technology)
Hua Lu (Aalborg University)
Huanchen Zhang (Tsinghua University)
Huiping Chen (University of Birmingham)
Hyungsoo Jung (Seoul National University)
Hyunjoon Kim (Hanyang University)
Ibrahim Sabek (University of Southern California)

Igor Zablotchi (Mysten Labs)
 Ilia Petrov (Reutlingen University)
 Ilias Azizi (Mohammed VI Polytechnic University)
 Ilie Sarpe (KTH Royal Institute of Technology)
 Indrakshi Ray (Colorado State University)
 Ioannis Demertzis (UCSC)
 Jaeyoung Do (Seoul National University)
 Jan Hidders (Birkbeck, University of London)
 Jeeta Ann Chacko (Technical University of Munich)
 Jeffrey Xu Yu (The Hong Kong University of Science and Technology (Guangzhou))
 Jelle Hellings (McMaster University)
 Jesús Camacho-Rodríguez (Microsoft)
 Jia Li (Hong Kong University of Science and Technology)
 Jia Zou (Arizona State University)
 Jialin Ding (Princeton University)
 Jianguo Wang (Purdue University)
 Jiannan Wang (Tsinghua University)
 Jianxin Li (Edith Cowan University)
 Jiawei Jiang (Wuhan University)
 Jieming Shi (The Hong Kong Polytechnic University)
 Jignesh Patel (Carnegie Mellon University)
 Jin Wang (Arizona State University)
 Jiwon Seo (Seoul National University)
 Johannes Pietrzyk (TU Dresden)
 Johes Bater (Tufts University)
 John Paparrizos (The Ohio State University, USA & Aristotle University of Thessaloniki, Greece)
 Jonathan Fürst (ZHAW Zurich University of Applied Sciences)
 Jonghyeok Park (Korea University)
 Joseph Near (University of Vermont)
 Ju Fan (Renmin University of China)
 Junghoon Kim (UNIST)
 Junhao Gan (University of Melbourne)
 K. Selçuk Candan (Arizona State University)
 Kaisong Huang (University of Calgary)
 Kai-Uwe Sattler (TU Ilmenau)
 Kanchan Chowdhury (Marquette University)
 Karim Benouaret (Univ Lyon)
 Katsiaryna Milylenka (Zalando SE)
 Kaustubh Beedkar (Indian Institute of Technology Delhi)
 Keke HUANG (The University of British Columbia)
 Kenneth A Ross (Columbia University)
 KI HYUN TAE (Samsung Research)
 Kostas Stefanidis (Tampere University)
 Kurt Stockinger (Zurich University of Applied Sciences)
 Kwanghyun Park (Yonsei University)
 Kyoungmin Kim (EPFL)
 Larissa Capobianco Shimomura (Hasselt University)
 Le Gruenwald (The University of Oklahoma)
 Lefteris Kokoris Kogias (MystenLabs)
 Lei Zou (Peking University)
 Leong Hou U (University of Macau)
 Leopoldo Bertossi (Carleton University)
 Liang Liang (EPFL)
 Lingyang Chu (McMaster University)
 Loredana Caruccio (University of Salerno)
 Lu Chen (Zhejiang University)
 Luca Gagliardelli (eCampus University)
 Lucas Braun (Oracle Labs)
 Lukasz Golab (University of Waterloo)
 Madhulika Mohanty (Inria & Institut Polytechnique de Paris)
 Magnus Mueller (AWS)
 Mahdi Esmailoghli (HU Berlin)
 Makoto Onizuka (Osaka University)
 Man Lung Yiu (Hong Kong Polytechnic University)
 Manuel Rigger (National University of Singapore)
 Marco Patella (University of Bologna)
 Markos Markakis (Massachusetts Institute of Technology)
 Martin Prammer (Carnegie Mellon University)
 Matteo Ceccarello (University of Padova)
 Matteo Interlandi (Microsoft)
 Matthaïos Olma (MongoDB)
 Matthew J Perron (Microsoft)
 Matthias Boehm (Technische Universität Berlin)
 Matthias Weidlich (Humboldt-Universität zu Berlin)
 Maximilian E Schüle (University of Bamberg)
 Meichun Hsu (Oracle America)
 Miao Qiao (The University of Auckland)
 Miao Yu (Meta)
 Michael J Cahill (University of Sydney)
 Michael Grossniklaus (University of Konstanz)
 Michael Gubanov (Florida State University)
 Michael J Mior (Rochester Institute of Technology)
 Mohamed Eltabakh (Qatar Foundation)
 Mohamed Mokbel (University of Minnesota - Twin Cities)
 Mohammad Sadoghi (University of California, Davis)
 Mohammad Javad Amiri (Stony Brook University)
 Mohammed J Saeed (Apple)
 Mostafa Milani (The University of Western Ontario)
 Nick Koudas (University of Toronto)
 Nikolay Yakovets (TU Eindhoven)
 Oktie Hassanzadeh (IBM Research)
 Oliver A Kennedy (University at Buffalo, SUNY)
 Oscar Romero (Universitat Politècnica de Catalunya)
 Panagiotis Bouros (Johannes Gutenberg University Mainz)
 Panagiotis Karras (University of Copenhagen)
 Paolo Garza (Politecnico di Torino)
 Paolo Merialdo (Università degli Studi Roma Tre)
 Paolo Missier (University of Birmingham)
 Paolo Papotti (EURECOM)
 Patrick Damme (Technische Universität Berlin)
 Patrick Schäfer (Humboldt-Universität zu Berlin)
 Paul Boniol (Inria, Ecole normale supérieure)
 Peizhi Wu (University of Pennsylvania)
 Peng Peng (Hunan University)
 Periklis Chrysogelos (Oracle)
 Peter M. Fischer (University of Augsburg)
 Pierangela Samarati (Università degli Studi di Milano)
 Prashant Pandey (Northeastern University)
 Protiva Rahman (University of Florida)
 Qingqing Ye (Hong Kong Polytechnic University)
 Qitong Wang (Harvard University)
 Qizhen Zhang (University of Toronto)
 Rada Chirkova (NC State University)
 Rainer Gemulla (Universität Mannheim)

Rana Alotaibi (King Abdulaziz City for Science and Technology (KACST))
 Rana Shahout (Harvard)
 Renata Borovica-Gajic (University of Melbourne)
 Renchi Yang (Hong Kong Baptist University)
 Ria Borromeo (University of Philippines)
 Riccardo Torlone (Roma Tre University)
 Rihan Hai (TU Delft)
 Ritesh Sarkhel (Amazon)
 Roe Shraga (WPI)
 Romila Pradhan (Purdue University)
 Ruiyuan Li (Chongqing University)
 Ryan Marcus (University of Pennsylvania)
 Ryan McKenna (University of Massachusetts Amherst)
 Sai Wu (Zhejiang University)
 Sara Foresti (Universita' degli Studi di Milano)
 Sarah Kleest-Meißner (Hasselt University)
 Sebastian Link (University of Auckland)
 Sebastian Michel (RPTU Kaiserslautern-Landau)
 Sepanta Zeighami (University of California Berkeley)
 Shantanu Sharma (New Jersey Institute of Technology)
 Shaoxu Song (Tsinghua University)
 Sharad Mehrotra (U.C. Irvine)
 Shen Liang (Université Paris Cité)
 Shihabur Chowdhury (Apple)
 Shimin Di (The Hong Kong University of Science and Technology)
 Silviu Maniu (Université Grenoble Alpes)
 Siqiang Luo (Nanyang Technological University)
 Søren Keiser Jensen (Aalborg University)
 Soror Sahri (Université Paris Cité)
 Stavros Maroulis (ATHENA Research Center)
 Stefania Dumbrava (ENSIIE, INRIA, IRIF, Télécom-SudParis)
 Stefano Marchesin (Università di Padova)
 Stefano Paraboschi (Università degli Studi di Bergamo)
 Steffen Zeuch (TU Berlin)
 Surajit Chaudhuri (Microsoft)
 Suyash Gupta (University of Oregon)
 Sven Groppe (Universität zu Lübeck/Institute of Information Systems (IFIS))
 Tanzima Hashem (Bangladesh University of Engineering and Technology)
 Tarique Ashraf Siddiqui (Microsoft Corporation)
 Tenindra Abeywickrama (RIKEN Center for Computational Science)
 Theodoros Chondrogiannis (Norwegian University of Science and Technology (NTNU))
 Thomas Neumann (TUM)
 Tianyi Li (Aalborg University)
 Tianzheng Wang (Simon Fraser University)
 Tieying Zhang (Bytedance)
 Ting Yu (MBZUAI)
 Tingjian Ge (University of Massachusetts, Lowell)
 Tong Chen (The University of Queensland)
 Torben Bach Pedersen (Aalborg University)
 Tristan Allard (Univ Rennes, CNRS, IRISA)
 Tsz Nam Chan (Shenzhen University)
 Ulf Leser (Humboldt-Universität zu Berlin)
 Utku Sirin (Harvard University)
 Venkata Vamsikrishna Meduri (IBM Research-Almaden)

Verena Kantere (University of Ottawa)
 Vikram Nathan (Amazon, Inc.)
 Viktor Leis (Technische Universität München)
 Vivek Narasayya (Microsoft)
 Vivek Shah (Samsung)
 Wang-Chien Lee (Pennsylvania State University, USA)
 Wei Dong (Nanyang Technological University)
 WEI LU (Renmin University of China)
 Weiguo Zheng (Fudan University)
 Wenqi FAN (The Hong Kong Polytechnic University)
 Wentao Wu (Microsoft Research)
 Wentao Zhang (Peking University)
 William Schultz (MongoDB)
 Xiang Lian (Kent State University)
 Xiangyao Yu (University of Wisconsin-Madison)
 Xiangyu Ke (Zhejiang University)
 Xiao Qin (Snowflake)
 Xiao Yan (Wuhan University)
 Xiaofei Zhang (University of Memphis)
 Xiaohui Yu (York University)
 Xiaou Ding (Harbin Institute of Technology)
 Xiaoying Wang (Microsoft)
 Xike Xie (University of Science and Technology of China)
 Xin Huang (Hong Kong Baptist University)
 Xuhao Chen (MSU)
 Xun Jian (Northwestern Polytechnical University)
 Xupeng Miao (Purdue University)
 Yael Amsterdamer (Bar-Ilan University)
 Yan Huang (University of North Texas)
 Yang Cao (Institute of Science Tokyo)
 Yang Wang (The Ohio State University)
 Yao Lu (NUS)
 Yesdaulet Izenov (Nazarbayev University)
 Yeye He (Microsoft Research)
 Yifan Wang (University of Hawaii)
 Yihao Ang (National University of Singapore)
 Yin Yang (Hamad bin Khalifa University)
 Yinan Li (Microsoft Research)
 Yinghui Wu (Case Western Reserve University)
 Yingtai Xiao (TikTok)
 Yiwen Zhu (Microsoft)
 Yixiang Fang (School of Data Science, The Chinese University of Hong Kong, Shenzhen)
 Yongjoo Park (UIUC)
 Yongluan Zhou (University of Copenhagen)
 Yongxin Tong (Beihang University)
 Yu Yang (City University of Hong Kong)
 Yuan Qiu (Southeast University)
 Yujie Hui (Google)
 Yuke Wang (Rice University)
 Yuyu Luo (HKUST (GZ))
 Zeyu Ding (Binghamton University)
 Zhao Cao (Remin University of China)
 Zhaojing Luo (Beijing Institute of Technology)
 Zhengjie Miao (Simon Fraser University)
 Zhewei Wei (Renmin University of China)
 Zhiwei Fan (Meta)
 Zhongle Xie (Zhejiang University)
 Zhuoyue Zhao (University at Buffalo)
 Zoi Kaoudi (IT University of Copenhagen)

Light Load Review Board

Abdulkhakim Qahtan (Utrecht University)
Abolfazl Asudeh (University of Illinois Chicago)
Ahmed El-Roby (Carleton University)
Alberto Abelló (Universitat Politècnica de Catalunya)
Alex Conway (Cornell Tech)
Alexander Erben (Technical University of Munich)
Alexandros Labrinidis (University of Pittsburgh)
Anupam Sanghi (TU Darmstadt)
Ariful Azad (Texas A&M University)
Babak Salimi (University of California at San Diego)
Baotong Lu (Microsoft Research)
Bettina Kemme (McGill University)
C. Seshadhri (UCSC)
Curtis Dyreson (Utah State University)
Danfeng Zhang (Duke University)
Daniel Kocher (University of Salzburg)
Davide Mottin (Aarhus University)
Fabian Panse (University of Augsburg)
Felix Naumann (Hasso Plattner Institute, University of Potsdam)
Haitao Yuan (Nanyang Technological University)
Hangdong Zhao (Microsoft)
Hideyuki Kawashima (Keio University)
Hongyan Chang (NUS)
Hongzhi Yin (The University of Queensland)
Ingo Müller (Google)
Jiaheng Lu (University of Helsinki)
Jinho Lee (Seoul National University)
Juan Sequeda (data.world)
Katia Papakonstantinou (Athens University of Economics and Business)
Kexin Rong (Georgia Institute of Technology)
Konstantinos Mamouras (Rice University)
Kyriakos Mouratidis (Singapore Management University)
Lipeng Wan (Georgia State University)
Manolis Terrovitis (Athena Research Center)
Matteo Brucato (Microsoft Research)
Mayuresh Kunjir (Amazon AWS)
Michele Linardi (CYU)
Mo Sha (Alibaba Group)
Murat Demirbas (MongoDB Research)
Myeongjae Jeon (POSTECH)
Nikos Giatrakos (Technical University of Crete)
Odysseas Papapetrou (TU Eindhoven)
Panagiotis Liakos (Athens University of Economics and Business)
Panagiotis Liakos (University of Athens)
Panos Vassiliadis (University of Ioannina)
Paris Carbone (KTH Royal Institute of Technology)
Paul Groth (University of Amsterdam)
Pierre Bourhis (CNRS)
Quang-Trung Ta (National University of Singapore)
Rathijit Sen (Microsoft)
Remy Wang (University of California, Los Angeles)
Renzo Angles (Universidad de Talca)
Seokki Lee (University of Cincinnati)
Sergi Nadal (Universitat Politècnica de Catalunya)
Sivaprasad Sudhir (MIT)
Song Jiang (University of Texas, Arlington)

Sonia Bergamaschi (Università di Modena e Reggio Emilia)
Subarna Chatterjee (Datastax)
Subhadeep Sarkar (Brandeis University)
Sujaya Maiyya (University of Waterloo)
Supawit Chockchowwat (University of Illinois Urbana-Champaign)
Theodore Johnson (AT&T Labs - Research)
Thomas Fahringer (University of Innsbruck)
Thorsten Papenbrock (Philipps University of Marburg)
Vassilis J. Tsotras (UC Riverside)
Venkatesh Emani (Microsoft)
Wenjia He (University of Michigan)
Wolfgang Gatterbauer (Northeastern University)
Xingguang Chen (National University of Singapore)
Yogesh Simmhan (Indian Institute of Science (IISc), Bangalore)
Yuxin Tang (Rice University)
Zhe Jiang (University of Florida)
Ziawasch Abedjan (BIFOLD/TU Berlin)

LETTER FROM THE EDITORS IN CHIEF

It is our pleasure to present the twelfth issue of PVLDB Volume 19. While the previous eleven issues covered research track papers that have been accepted to PVLDB and will be presented at the VLDB 2026 Conference, Issue 12 includes all other contributions, including the VLDB Endowment Award talks, Keynote talks, VLDB 2026 Special Event Panels, Research Panels, as well as peer-reviewed Industrial papers, Demonstrations, Tutorials, and Workshops, which are also part of the VLDB 2026 program.

The **Industrial Track** of VLDB 2026 bridges the gap between innovative data systems research and practical deployment, covering the full spectrum of commercial and industrial-strength (including open-source) data management systems and solutions, as well as novel applications and firsthand experiential reports from tackling challenging industry problems. Of particular note, over thirty percent of the submissions explored the integration of artificial intelligence with database systems, while the remainder addressed more conventional database topics. A total of 121 papers were submitted, of which 49 were accepted following a rigorous review process that ensured that each manuscript was evaluated by at least three knowledgeable reviewers and included a paper revision phase. Prior to final acceptance, the Program Committee carefully verified that all requested revisions had been adequately addressed. The resulting program features eight thematic sessions, supplemented by dedicated poster sessions.

The VLDB 2026 **Demonstration Track** showcases innovative data-management systems, tools, and applications through interactive presentations. To support a review process that was both rigorous and scalable, a diverse group of distinguished researchers served as Associate Editors (AEs), overseeing review quality and consistency and facilitating discussions among reviewers. In addition, automated conflict of interest (CoI) checks were conducted using the CLOSET System. Furthermore, the review form was specifically tailored to better align with the unique characteristics of demonstration submissions. The track received a substantial number of submissions, totaling 224, out of which 92 were accepted for presentation, resulting in an acceptance rate of 41%. Each submission was reviewed by three expert reviewers under the responsibility of an AE. The accepted demonstrations were organized into three thematic double sessions: Graph, Systems and Infrastructure; Query Processing, AI and Data Agents; and Data Quality, Analytics and Applications. Together they spanned graph and query infrastructure, agentic and LLM-assisted data access, data preparation and quality, and trustworthy data-intensive applications. The candidates for the Best Demonstration Award were proposed by AEs on the basis of the accepted proposals.

The **Tutorial Track** of VLDB 2026 covers state-of-the-art research, development, and applications in data management or related areas, including interdisciplinary areas. This year, we received 31 submissions, out of which we accepted 15 tutorials. We thank all authors for their contributions to the track, and the program committee for their work in reviewing a growing number of submissions. The Tutorial Program Committee was assembled with the diversity of topics and demographics within our community in mind. Each proposal was assigned to the committee and subsequently reviewed by the Tutorial Program Chairs for a final decision, in coordination with the Program Committee Chairs and the General Chairs. The accepted proposals cover a wide range of exciting topics, including GPU-based query evaluation, human-computer interaction, conflict-free replicated data types, instance-optimal joins, data management in healthcare, quantum databases, homomorphic encryption, LLMs in and around databases, agent/database interactions, and vector databases.

VLDB 2026 will offer a diverse set of **panels** covering Data Agents, HW-SW Co-Design and Research-to-Production Challenges:

- Data Agents: Rethinking Data Systems in the AI Agent Era
 - Moderators: Guoliang Li (Tsinghua University), Yuyu Luo (HKUST (Guangzhou))
 - Panelists: Carsten Binnig (TU Darmstadt), Samuel Madden (MIT), Fatma Özcan (Google), Aditya Parameswaran (UC Berkeley), Nan Tang (HKUST (Guangzhou)), Immanuel Trummer (Cornell University), and Eugene Wu (Columbia University)
- HW-SW Co-Design for Databases in the Age of AI-for-Systems
 - Moderator: Monica Chiosa (MIT)
 - Panelists: Jana Giceva (TU Munich), Tae Jun Ham (Google), Matteo Interlandi (Microsoft Gray Systems Lab), Vasiliki (Vasia) Kalavri (Boston University), and Clemens Lutz (NVIDIA)
- Publish or Ship? The Research-to-Production Gap in Data Systems
 - Moderator: Subarna Chatterjee (IBM WatsonX)
 - Panelists: Natacha Crooks (UC Berkeley), Carlo Curino (Microsoft Gray Systems Lab), Sudipto Das (Databricks), Andreas Kipf (University of Technology Nuremberg), Ryan Marcus (University of Pennsylvania), and Yuanyuan Tian (Microsoft Gray Systems Lab)

VLDB 2026 received 21 Workshop proposals. Like previous years, we decided to accept a larger number of workshops to ensure a good diversity of topics and allow new and revived workshops to establish themselves with sufficiently many submissions. After reviewing the proposals and evaluating their merits based on their history, relevance of topics, and distinctiveness, we accepted 12 proposals directly. For five additional proposals, we identified overlaps in their topics. We therefore encouraged the respective organizers to merge their proposals and organize joint workshops. As a result, we ended up with 14 high-quality workshops (8 full-day and 6 half-day). This year's conference program included the following workshops:

Monday, August 31:

- ADMS: The 17th International Workshop on Accelerating Analytics and Data Management Systems Using Modern Processor and Storage Architectures (Chairs: Rajesh Bordawekar, Tirthankar Lahiri)
- AIDB: The 7th Applied AI for Database Systems and Applications (Chairs: Thaleia Dimitra Doudali, Subru Krishnan, Konstantinos Kanellis)
- NOVAS: Novel Optimization for Visionary AI Systems (Chairs: Gerardo Vitagliano, Chunwei Liu, Liana Patel, Rana Shahout, Andreas Kipf, Paolo Papotti)
- QDB: The 15th International Workshop on Quality in Databases (Chairs: Lorena Etcheverry, Steven Euijong Whang, Cinzia Cappiello)
- TPCTC: The 18th TPC Technology Conference on Performance Evaluation and Benchmarking (Chairs: Raghunath Nambiar, Meikel Poess)

Friday, September 4:

- Agents+Graphs: The 3rd International Workshop on Data Management Opportunities in Bringing Agents with Graph Data (Chairs: Yixiang Fang, Arijit Khan, Tianxing Wu, Da Yan)
- BioDMS: Workshop on Biomedical Data Management Systems (Chairs: Bojan Karlas, Gerardo Vitagliano, Benjamin Gyori, Ulf Leser)
- CDMS: The Fourth International Workshop on Composable Data Management Systems (Chairs: Satya Valluri, Mohamed Zait)
- DASHSys: Workshop on Systems for Data-centric Agents with Human-in-the-loop (Chairs: Yiru Chen, Kun Qian, Estevam Hruschka, Zhengjie Miao, Dan Zhang, Sajjadur Rahman, Anna Fariha, Sainyam Galhotra, Umar Farooq Minhas, Sepanta Zeighami)
- ADS: The Joint Workshop on Agentic Data Systems and Data-Centric AI (Chairs: Guoliang Li, Yuyu Luo, Ju Fan, Xuanhe Zhou)
- QCDKM: Quantum Computing and Data/Knowledge Management (Chairs: Haridimos Kondylakis, Stefanie Scherzinger, Bettina Kemme, Vasilis Efthymiou, Yannis Tzitzikas, Manuel Wimmer, Jose Garcia-Alonso)
- TaDA: The 4th International Workshop on Tabular Data Analysis (Chairs: Vasilis Efthymiou, Oktie Hassanzadeh, Chuan Lei, Kavitha Srinivas)
- SCCP: Symposium on Consistency Checking Principles (Chairs: Cheng Tan, Shuai Mu, Hengfeng Wei, Si Liu)
- VecDB: The 2nd Workshop on Vector Databases (Chairs: Yusuke Matsui, Martin Aumuller, Magdalen Manohar, Harsha Simhadri)

VLDB 2026 features 2 **Keynote** talks:

- The Three Golden Ages of Database Engineering: From SIGMOD '85 to the Agentic Era (Dr. Reynold Xin, Databricks)
- Efficient and Reliable Systems for Building Foundation Models at Scale (Dr. Jingren Zhou, Alibaba Cloud)

In addition, VLDB 2026 hosts a special **“Return to Boston”** event: The first VLDB conference was held in 1975 in the Boston area, which was home to many computer companies at the time. After travelling around the world for 50 years, the VLDB conference returns to Boston in 2026. To mark this momentous occasion, VLDB 2026 is organizing a plenary “Return to Boston” event to celebrate the impact of VLDB on the global data management community. This event will feature three sessions:

- The story of the founding of VLDB in 1975 and what led to its subsequent global footprint, as told by an attendee of VLDB 1975 who is an active member of our community 50+ years later.
- The technical impact of some of the key papers from VLDB 1975 that are relevant even today, as told by several members of our community with interesting perspectives on this topic.
- The pioneering innovations and global influence of the VLDB conferences over the years, as told by recent VLDB endowment presidents from around the world.

Besides the Best Paper and Best Demonstration awards announced during the conference, VLDB 2025 features 3 **VLDB Endowment awards** with invited talks in the program:

- 2026 VLDB Test of Time Award: Tyler Akidau, Robert Bradshaw, Craig Chambers, Slava Chernyak, Rafael J. Fernández-Moctezuma, Reuven Lax, Sam McVeety, Daniel Mills, Frances Perry, Eric Schmidt, and Sam Whittle (Google Research) for their VLDB 2015 paper, “The Dataflow Model: A Practical Approach to Balancing Correctness, Latency, and Cost in Massive-scale, Unbounded, Out-of-order Data Processing.”
- 2026 VLDB Women in Database Research Award: Juliana Freire (New York University) for “pioneering contributions in provenance, reproducibility and urban data science, and for transformative leadership in the data management community.”
- 2026 VLDB Early Career Research Contribution Award: Jana Giceva (Technical University of Munich) for “visionary contributions to building data systems whose core software abstractions decouple logical intent from physical hardware microarchitectures.”

In closing, we wish to express our collective and heartfelt gratitude to all members of the various program committees and to our Proceedings Chairs. Their dedication and tireless efforts throughout the past year have been essential to ensuring the timely publication and high quality of PVLDB Volume 19.

Yanlei Diao and Xiaokui Xiao
Editors-in-Chief of PVLDB Vol. 19
Program Chairs for VLDB 2026

Christian König, Efthymia Tsamoura, and Hannes Voigt
Industrial Program Chairs

Stefania Dumbrava and John Paparrizos
Demonstration Program Chairs

Oliver Kennedy and Wenjie Zhang
Tutorial Program Chairs

Bettina Kemme and Steven Euijong Whang
Panel Program Chairs

Tilmann Rabl and Meihui Zhang
Workshop Program Chairs

Mohamed Eltabakh and Nikos Tziavelis
PhD Workshop Chairs

Divesh Srivastava
VLDB 2026 Special Event Chair

Jiuqi Wei and Zhaojing Luo
Workshop Proceedings Chair

Stefan Halfpap, Siqiang Luo, Andrea Mauri, Subhadeep Sarkar, and Giovanni Simonini
Proceedings Chairs