Proceedings of the VLDB Endowment

Volume 15, No. 1 – September 2021

PVLDB – Proceedings of the VLDB Endowment

Volume 15, No. 1, September 2021.

All papers published in this issue will be presented at the 48th International Conference on Very Large Data Bases, Sydney, Australia, 2022.

Copyright 2021 VLDB Endowment

Volume 15, Number 1, September 2021

Pages i – vi and 1 - 140

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. 15 iii

Research Papers

ANN Softmax: Acceleration of Extreme Classification Training 1
Kang Zhao, Liuyihan Song, Yingya Zhang, Pan Pan, Xu Yinghui, Rong Jin

WindTunnel: Towards Differentiable ML Pipelines Beyond a Single Model 11
Gyeong-in Yu, Saeed Amizadeh, Sehoon Kim, Artidoro Pagnoni, Ce Zhang, Byung-gon Chun, Markus Weimer, Matteo Interlandi

DBOS: A DBMS-oriented Operating System 21
Athinagoras Skiadopoulos, Qian Li, Peter Kraft, Kostis Kaffes, Daniel Hong, Shana Mathew, David Bestor, Michael Cafarella, Vijay Gadepally, Goetz Graefe, Jeremy Kepner, Christos Kozyrakis, Tim Kraska, Michael Stonebraker, Lalith Suresh, Matei Zaharia

Deep Indexed Active Learning for Matching Heterogeneous Entity Representations 31
Arjit Jain, Sunita Sarawagi, Prithviraj Sen

A Learned Query Rewrite System using Monte Carlo Tree Search 46
Xuanhe Zhou, Guoliang Li, Chengliang Chai, Jianhua Feng

On Detecting Cherry-picked Generalizations 59
Yin Lin, Brit Youngman, Yuval Moskovitch, H. V. Jagadish, Tova Milo

FACE: A Normalizing Flow based Cardinality Estimator 72
Jiayi Wang, Chengliang Chai, Jiabin Liu, Guoliang Li

Learned Cardinality Estimation: A Design Space Exploration and a Comparative Evaluation 85
Ji Sun, Jintao Zhang, Zhaoyan Sun, Guoliang Li, Nan Tang

DeepEverest: Accelerating Declarative Top-K Queries for Deep Neural Network Interpretation 98
Dong He, Maureen Daum, Walter Cai, Magdalena Balazinska

Cosine: A Cloud-Cost Optimized Self-Designing Key-Value Storage Engine 112
Subarna Chatterjee, Meena Jagadeesan, Wilson Qin, Stratos Idreos

Accelerating Recommendation System Training by Leveraging Popular Choices 127
Muhammad Adnan, Yassaman Ebrahimzadeh Maboud, Divya Mahajan, Prashant Nair

PVLDB ORGANIZATION AND REVIEW BOARD - Vol. 15

Review Board

Abolfazl Asudeh (University of Michifan)
Aécio Santos (New York University)
Ahmed Eldawy (University of California, Riverside)
Alexander Hall (RelationalAI)
Alexander J Ratner (University of Washington)
Aline Bessa (New York University)
Alkis Simitsis (Athena Research Center)
Altigran da Silva (Universidade Federal do Amazonas)
AnHai Doan (University of Wisconsin-Madison)
Anna Fariha (Microsoft)
Anton Dignös (Free University of Bozen-Bolzano)
Antonio Cavalcante Araujo Neto (University of Alberta)
Arijit Khan (Nanyang Technological University)
Arvind Arasu (Microsoft)
Babak Salimi (University of California, San Diego)
Bailu Ding (Microsoft Research)
Bertram Ludaescher (University of Illinois)
Bolong Zheng (Huazhong University of Science and Technology)
Brandon Haynes (Gray Systems Lab, Microsoft)
Byron Choi (Hong Kong Baptist University)
Carlo Curino (Microsoft -- GSL)
Carlos Scheidegger (The University of Arizona)
Carsten Binnig (TU Darmstadt)
Ce Zhang (ETH)
Cheng Long (Nanyang Technological University)
Chengfei Liu (Swinburne University of Technology)
Chuan Lei (Instacart)
Chunbin Lin (Amazon AWS)
Curtis Dyreson (Utah State University)
Dan Kifer (Pennsylva State University)
Dana M Van Aken (Carnegie Mellon University)
Daniel Deutch (Tel Aviv University)
Daniel Oliveira (UFF, Brazil)
David Koop (Northern Illinois University)
Davide Mottin (Aarhus University)
Dong Xie (Penn State University)
Eduardo Ogasawara (CEFET-RJ)
Eleni Tzirita Zacharatou (TU Berlin)
Fabio Porto (LNCC)
Faisal Nawab (University of California at Irvine)
Fan Zhang (Guangzhou University)
Fatemeh Nargesian (University of Rochester)
Fei Chiang (McMaster University)
Florin Rusu (UC Merced)
Floris Geerts (University of Antwerp)
Fotis Psallidas (Microsoft)
George Fletcher (Eindhoven University of Technology)
George Papadakis (University of Athens)
Gerhard Weikum (Max-Planck-Institut fur Informatik)
Germain Forestier (University of Haute Alsace)
Guoliang Li (Tsinghua University)
Haipeng Dai (Nanjing University)
Harish Doraiswamy (Microsoft Research India)
Heiko Mueller (DeepReason.ai)
Herodotos Herodotou (Cyprus University of Technology)
Holger Pirk (Imperial College)
Hongzhi Yin (The University of Queensland)
Huiping Cao (New Mexico State University)
Immanuel Trummer (Cornell)
Ioana Manolescu (INRIA and Institut Polytechnique de Paris)
Ippokratis Pandis (Amazon)
Ishtiyaque Ahmad (University of California, Santa Barbara)
Jae-Gil Lee (KAIST)
Jana Giceva (TU Munich)
Jeffrey Xu Yu (Chinese University of Hong Kong)
Jens Teubner (TU Dortmund University)
Jia Zou (Arizona State University)
Jian Pei (Simon Fraser University)
Jianguo Wang (Purdue University)
Jiannan Wang (Simon Fraser University)
Jianxin Li (Deakin University)
Jianye Yang (Central South University)
Jiwon Seo (Hanyang University)
Johannes Gehrke (Microsoft)
Jorge Arnulfo Quiane Ruiz (TU Berlin)
Joseph Near (University of Vermont)
Junhu Wang (Griffith University)
Kaiping Zheng (National University of Singapore)
Kangfei Zhao (The Chinese University of Hong Kong)
Karima Echihabi (Mohammed VI Polytechnic University)
Katja Hose (Aalborg University)
Kenneth A Ross (Columbia University)
Kostas Zoumpatianos (Snowflake Computing)
Lei Zou (Peking University)
Leopoldo Bertossi (Universidad Adolfo Ibanez)
Li Xiong (Emory University)
Lianke Qin (University of California, Santa Barbara)
Lijun Chang (The University of Sydney)
Lin Ma (Carnegie Mellon University)
Long Yuan (Nanjing University of Science and Technology)
Lu Qin (UTS)
Luciano Barbosa (Universidade Federal de Pernambuco)
Marcelo Arenas (Universidad Catolica & IMFD)
Maria Luisa Sapino (U. Torino)
Matteo Lissandrini (Aalborg University)
Matthias Boehm (Graz University of Technology)
Matthias Renz (University of Kiel)
Max Heimel (Snowflake)
Maximilian Schleich (University of Washington)
Meihui Zhang (Beijing Institute of Technology)
Melanie Herschel (Universität Stuttgart)
Michael Abebe (University of Waterloo)
Min Xie (Instacart)
Mirella M Moro (Universidade Federal de Minas Gerais)
Mohamed Sarwat (Arizona State University)
Mohammad Dashti (MongoDB)
Mohammad Javad Amiri (University of Pennsylvania)
Mohammad Sadoghi (University of California, Davis)
Muhammad Aamir Cheema (Monash University)

Nikita Bhutani (Megagon Labs)
Oliver A Kennedy (University at Buffalo, SUNY)
Panos K. Chrysanthis (University of Pittsburgh)
Paolo Missier (Newcastle University)
Parth Nagarkar (NMSU)
Paul Groth (University of Amsterdam)
Peng CHENG (East China Normal University)
Peter Pietzuch (Imperial College London)
Pierangela Samarati (Universita delgi Studi di Milano)
Pinar Karagoz (METU, Turkey)
Pinar Tozun (IT University of Copenhagen)
Prithu Banerjee (UBC)
Raoni Lourenço (New York University)
Raul Castro Fernandez (UChicago)
Ravi Ramamurthy (Microsoft)
Raymond Chi-Wing Wong (Hong Kong University of Science and Technology)
Renata Borovica-Gajic (University of Melbourne)
Reynold Cheng (The University of Hong Kong)
Rui Mao (Shenzhen University)
Ruoming Jin (Kent State University)
Sai Wu (Zhejiang University)
Sainyam Galhotra (University of Chicago)
Sanjay Krishnan (University of Chicago)
Sanjib Kumar Das (Google)
Sayan Ranu (IIT Delhi)
Sebastian Link (University of Auckland)
Semih Salihoglu (University of Waterloo)
Senjuti Basu Roy (New Jersey Institute of Technology)
Sergey Melnik (Google)
Shantanu Sharma (New Jersey Institute of Technology)
Shaoxu Song (Tsinghua University)
Sheng Wang (New York University)
Shimin Chen (Chinese Academy of Sciences)
Shumo Chu (University of California, Santa Barbara)
Shweta Jain (University of Illinois, Urbana-Champaign)
Sibo Wang (The Chinese University of Hong Kong)
Srinivasan Keshav (University of Cambridge)
Steffen Zeuch (DFKI GmbH)
Steven E Whang (KAIST)
Subarna Chatterjee (Harvard University)
Sudip Roy (Google)
Supun C Nakandala (University of California, San Diego)
Tamer Özsu (University of Waterloo)
Tarique A Siddiqui (Microsoft Research)
Thomas Heinis (Imperial College)
Thomas Neumann (TUM)
Tianzheng Wang (Simon Fraser University)
Tien Tuan Anh Dinh (Singapore University of Technology and Design)
Tilmann Rabl (HPI, University of Potsdam)
Ting Yu (Qatar Computing Research Institute)
Torben Bach Pedersen (Aalborg University)
Torsten Grust (Universität Tübingen)
Umar Farooq Minhas (Microsoft Research)
Vasiliki Kalavri (Boston University)
Verena Kantere (National Technical University of Athens)
Victor Zakhary (Oracle)
Vivek Narasayya (Microsoft Research)
Vraj Shah (University of California, San Diego)
Walid G Aref (Purdue)
Wasay Abdul (Harvard)
Wei Wang (Hong Kong University of Science and Technology (Guangzhou))
Wei Lu (Renmin university of china)
Weiren Yu (University of Warwick)
Wen Hua (The University of Queensland)
Wolfgang Lehner (TU Dresden)
Xi He (University of Waterloo)
Xiang Lian (Kent State University)
Xiao Qin (IBM Research)
Xiaofei Zhang (University of Memphis)
Xiaokui Xiao (National University of Singapore)
Xiaolan Wang (Megagon Labs)
Xiaoyang Wang (Zhejiang Gongshang University)
Xin Huang (Hong Kong Baptist University)
Yael Amsterdamer (Bar-Ilan university)
Yanyan Shen (Shanghai Jiao Tong University)
Ye Yuan (Northeastern University)
Yeye He (Microsoft Research)
Yi Chen (NJIT)
Yi Lu (MIT)
Yikai Zhang (Chinese University of Hong Kong)
Yinan Li (Microsoft Research)
Ying Zhang (University of Technology Sydney)
Yongxin Tong (Beihang University)
Yuanyuan Zhu (Wuhan University)
Yue Wang (Shenzhen Institute of Computing Sciences, Shenzhen University)
Yufei Tao (Chinese University of Hong Kong)
Yuliang Li (Megagon Labs)
Yuncheng Wu (National University of Singapore)
Yunjun Gao (Zhejiang University)
Yuval Moskovitch (University of Michigan)
Zhifeng Bao (RMIT University)
Zhongle Xie (Zhejiang University)
Zi Huang (University of Queensland)
Ziawasch Abedjan (Leibniz Universität Hannover)
Zohar Karnin (Amazon)
Zsolt István (IT University of Copenhagen)

LETTER FROM THE EDITORS IN CHIEF

We are very pleased to present the first issue of PVLDB's Volume 15. The Proceedings of the VLDB present the latest research in the area of database and information system technology. Together with expert boards of associate editors and reviewers, submissions are carefully peer-reviewed, often entering a revision phase, then published in the journal and ultimately presented at the following VLDB conference. We are very grateful to all colleagues who contribute to the success of PVLDB.

Our community recognizes that reproducibility and replicability are essential to move database research forward. Transparency – the availability of code, data, and experimental settings – is a requirement for results to be reproduced and validated. In addition, it enables others to replicate and build upon existing results to advance the state of the art. In 2018, PVLDB established a process for reproducibility evaluation (http://vldb.org/pvldb/reproducibility): authors of accepted papers are invited to submit their code and data, and members of the reproducibility committee assess the reproducibility of the results reported in the paper. For 2022, we introduced a new policy that aims to increase the transparency of papers submitted to PVLDB: authors are now invited to submit supplemental material, such as code, data and other implementation artifacts used to produce the results reported in the paper. Reviewers have access to the supplemental material and consider it in their evaluation of the scientific quality of the contribution. If authors are not able to submit the supplemental material, they must explain why. Authors of accepted papers are 1) expected to include the supplementary material with the camera ready, which will be given an official ACM Artifacts Available badge (https://www.acm.org/publications/policies/artifact-review-and-badging-current), and 2) strongly encouraged to participate in the Reproducibility Evaluation (http://vldb.org/pvldb/reproducibility) and compete for the VLDB Best Reproducible Paper Award.

PVLDB strives to give high-quality and constructive feedback in the form of reviews and meta-reviews. Each paper is evaluated by at least three reviewers and an Associate Editor, who summarizes the reviews and the results of the three-week discussion phase in which reviewers exchange their view of the paper and converge to a joint decision. The availability of supplementary material for accepted papers is evaluated by the PVLDB Reproducibility chairs.

This first edition of PVLDB's Volume 15 includes eleven papers. Zhao et al. propose ANN Softmax, a strategy based on approximate nearest neighbor search to sample classes for "extreme classification". Yu et al. present WindTunnel, a framework that translates operators of a given ML pipeline into differentiable Neural Network (NN) modules, so as to achieve end-to-end training through backpropagation. Skiadopoulos et al. lay out a vision for building a DBMS-oriented operating system (OS). Jain et al. propose DIAL, a scalable active learning approach for entity resolution that jointly learns embeddings to maximize recall for blocking and accuracy for matching blocked pairs. Zhou et al. address the problem of query optimization via rewrite rules. They propose exploring the space of rewrite rule sequences via Monte-Carlo Tree Search and to evaluate alternatives by a learned cost model. Lin et al. formalize the problems of identifying cherry-picked generalizations and quantifying the extent of the generalization and present a framework for detecting and explaining cherry-picked generalizations by refining aggregate queries. Wang et al. present FACE, a new data-driven framework for cardinality estimation. Sun et al. perform a detailed evaluation of different learning-based cardinality estimators. He et al. propose DeepEverest, a system for the efficient execution of interpretation by example queries over the activation values of a deep neural network. Chatterjee et al. propose Cosine, a tool for creating customized cloud-based key-value stores specialized to particular workloads. Finally, Adnan et al. propose a GPU-aware framework to accelerate recommendation system training that explores the semantics of training data and leverages the presence of popular choices to create a hot-embedding aware data layout.

All papers will be presented at the 2022 Conference on Very Large Databases (VLDB 2022) in Sydney, Australia. We hope you enjoy reading this issue and look forward to seeing you in Sydney!

Juliana Freire and Xuemin Lin,
Editors-in-Chief of PVLDB Volume 15
Program Chairs for VLDB 2022