

Beyond Incrementalism: How to Change the World Through Data Systems Research

Viktor Leis
Technical University of Munich
Germany
leis@in.tum.de

ABSTRACT

In this panel we argue that traditional data systems research remains highly in today's world. It aims to spark dialogue on choosing impactful research problems, influencing industry, and improving data system technology.

PVLDB Reference Format:

Viktor Leis. Beyond Incrementalism: How to Change the World Through Data Systems Research. PVLDB, 18(12): 5508 - 5509, 2025.
doi:10.14778/3750601.3760515

1 MOTIVATION

Historically, the data management community formed around topics such as storage structures, query processing, query languages, query optimization, and transactions. Today, however, only a small fraction (probably less than 10%) of research papers at VLDB and SIGMOD address these traditional systems topics. In fact, some in our community question their continued relevance. As Eugene Wu put it in a recent blog post:

"Continuing to improve RDBMS technology is helpful, but not necessarily competitive with industry, and the gains are increasingly marginal – we are polishing a round ball."

<https://wp.sigmod.org/?p=3801>

So: is traditional systems work still relevant? We argue emphatically *yes*—now more than ever. Modern data infrastructure is undergoing foundational change: from novel hardware and cloud-native architecture to the rise of AI, open data lakes, and modular data systems. These shifts challenge long-held assumptions and demand rethinking many core techniques. While industry innovates, it is often constrained by short-term incentives that favor incrementalism over bold, transformative ideas.

Academia has a unique role to play. Over the decades, academic research has had a huge impact—not only through ideas, but through full systems that changed the world: Ingres, PostgreSQL, Spark, DuckDB, and many others.

The goal of this panel is to reignite a focus on ambitious, high-impact data systems research and to inspire the next generation of researchers to think beyond incremental improvements.

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.
Proceedings of the VLDB Endowment, Vol. 18, No. 12 ISSN 2150-8097.
doi:10.14778/3750601.3760515

2 STRUCTURE AND GOALS

We will host a lively, moderated panel discussion featuring leading data systems researchers that have achieved real-world impact. The format will emphasize interactive discussion and audience engagement.

Goals

- Highlight high-impact data systems efforts—past and present
- Provide concrete guidance for choosing research topics with real-world relevance
- Discuss how to extend impact beyond papers (via open-source, startups, collaboration, etc.)
- Explore systemic factors (e.g., incentives, community norms) that shape research trajectories

Example Discussion Topics

- What are the most important data systems problems currently being ignored by academia?
- How can academic research influence industry at scale?
- How should researchers choose topics that matter?
- What's the role of open-source and startups in amplifying research impact?
- What lessons can we learn from impactful systems projects?
- How can publication and evaluation norms better support impactful research?

3 PANELISTS

Anastasia Ailamaki is a Professor of Computer and Communication Sciences at the École Polytechnique Fédérale de Lausanne (EPFL), a visiting researcher at Google, and the co-founder and Chair of the Board of Directors of RAW Labs SA, a Swiss company developing systems to analyze heterogeneous big data from multiple sources efficiently. Her research interests are in data-intensive systems and applications, and in particular: (a) in strengthening the interaction between the database software and emerging hardware and I/O devices, and (b) in automating data management to support computationally-demanding, data-intensive scientific applications. **Peter Boncz** leads the Database Architectures group at CWI and is professor at VU University Amsterdam. His academic career started with the open-source column-store MonetDB becoming the outcome of his PhD. He has a track record in bridging the gap between academia and commercial application, founding multiple startups. In 2008 he co-founded Vectorwise around the analytical database system by the same name, which pioneered vectorized query execution and lightweight data compression; which have been adopted broadly in analytical database systems. In recent years

he has collaborated closely with both Databricks and MotherDuck, a startup that is connecting DuckDB to the cloud.

Badrish Chandramouli is a partner research manager in the data systems group at Microsoft Research. His interests include key-value stores, caching, indexing, distributed systems, stream processing, and vector databases. He has worked on a wide range of systems such as Trill (a streaming engine), FASTER (a key-value storage library), Mison (a JSON parser), ALEX (a learned index), and most recently, Garnet (a Redis compatible remote cache-store). These efforts have led to numerous publications and deep impact across both Microsoft products and open source. Badrish's research has won honors such as best paper awards at the IEEE ICDE conference and the ACM SIGMOD's DBTest workshop, the ACM SIGMOD

research highlight award, and has been invited as the best papers at VLDB and ACM SIGMOD's DaMoN workshop. He is a member of ACM and IEEE, and regularly serves in associate editor, meta-reviewer, program committee, and organizer roles for database conferences. Learn more about his research at <https://badrish.net/>.

Andy Pavlo is an Associate Professor with Indefinite Tenure of Databaseology in the Computer Science Department at Carnegie Mellon University. He was the co-founder of the OtterTune start-up and record label, but then it died. He is currently the co-founder of "SO-YOU-DONT-HAVE-TO INCORPORATED"); DROP TABLE companies; -". His goal is to achieve a database-centric lifestyle.