Where Does Academic Database Research Go From Here?

Eugene Wu
ewu@cs.columbia.edu
Columbia University
New York, NY, USA

Raul Castro Fernandez
raulcf@uchicago.edu
The University of Chicago
Chicago, IL, USA

ABSTRACT

An open forum to discuss and debate the future of database research in the context of industry, other research communities, and AI.

PVLDB Reference Format:
Eugene Wu and Raul Castro Fernandez. Where Does Academic Database Research Go From Here?. PVLDB, 18(12): 5510 - 5511, 2025.
doi:10.14778/3750601.3760516

1 INTRODUCTION

We can informally define an academic community as a group that jointly contributes to advancing knowledge within a specific discipline, with shared research questions and methodologies. Communities are brought together due to the importance and impact – broadly defined – of the questions that the communities pursue.

The world, of course, is not sentimental. The fortunes of a community ebb and wane based on its competitive advantage as compared to the rest of the world. Historically, alchemy faded in prominence and was replaced by more effective methods, namely chemistry and the scientific method. Similarly, humoralism was replaced by evidence-based medicine, and behaviorism was replaced by cognitive psychology and neuroscience. In engineering, cybernetics grew out of war efforts and appeared poised to take over the world, but ultimately evolved into the pragmatic discipline of control theory. In computer science, "machine learning" research as repeatedly evolved from the 50s (logic-based), the 70s (expert systems), to today's deep neural networks and LLMs that have taken center stage due to the massive potential that automated intelligence can bring.

Closer to our hearts, database systems have grown in prominence over the past 50 years due to the demands of accounting, logistics, the Internet, and digitization. Yet, the data management success does not guarantee future success; that is up to us—the community—to chart our place in the AI era.

A recent SIGMOD blog post [7] reflected on the state of academic research from this perspective of competitive advantage as compared to startups, industry, and other research communities. It argued that much of the success of the data management community is tied to the commercial success of relational database management systems (RDBMS). Yet, this success has transformed RDBMSes into a commodity that sits on the plateau of the technology S-curve, and continued investment provides marginal gains that cannot compete with industry. At the same time, the article asked what advantage academic database research has in AI.

Proceedings of the VLDB Endowment, Vol. 18, No. 12 ISSN 2150-8097.
doi:10.14778/3750601.3760516

If this argument is true, it shakes the foundation of our academic community. We must answer what our role is and how the problems we solve are relevant not only to the researchers associated with our community but also to the outside world. On the other hand, many members of the community have objected to this characterization, and argue that our community is healthy, relevant, and simply poorly marketed to the outside world.

This naturally sets the stage for a debate, and it is in this context that we organized a panel with members of the database community to debate where we should be going to maintain a competitive edge at SIGMOD 2025. The panel was largely well-received. We aimed to foster a bottom-up conversation on the future of academic database research in the face of major technological and economic shifts—especially the rise of AI and its implications for talent, funding, and institutional relevance. The discussion centered on our community's comparative advantage: what unique contributions can academic researchers make in a world where industry investment in AI is surging (e.g., over $100B in VC funding in 2024 alone), while public research funding faces new constraints? The panel generated a wide range of reactions. Many attendees appreciated the opportunity to name shared concerns and think collectively about our future. Others saw the panel as part of an overly self-critical tradition. A more detailed summary of SIGMOD's panel can be found in the SIGMOD Blog [8] and an extended version of this document is in arXiv, where we keep an up-to-date version [9].

The strong and varied response to the SIGMOD panel made it clear that our community is eager to grapple with the deep questions facing academic database research. Yet it also revealed the need for continued dialogue—dialogue that not only surfaces concerns, but channels them into constructive next steps. With that in mind, we bring the conversation to VLDB, not as a repetition, but as a continuation and evolution. The context of VLDB, with its international scope and long-term perspective, invites us to ask: What will define our community's relevance over the next fifty years?[1] This panel features a different set of voices—including leaders from diverse institutional and intellectual backgrounds—and takes a more forward-looking stance. We focus on identifying our comparative advantage and outlining tangible interventions that can help steer our field toward long-term impact in the age of AI and beyond.

- What is the comparative advantage of the academic database community over the next decades?
- What is needed to keep it relevant?
- What specific interventions can we implement?
- What is the academic database community's place in the current AI revolution?

[1] VLDB'25 will be the 51st installment of the conference!

A full version of this document that combines information about the SIGMOD panel, its report, and hot takes from community members can be found at https://arxiv.org/abs/2504.08948.

1.1 Panel Composition and Format

For VLDB, we have a different mix of panelists. Rather than concentrating on senior academic researchers, we reached out to senior students, junior faculty, and senior industry researchers. We believe that this diverse mix of perspectives will unlock new insights and answers to the questions above.

The panelists include:

- Shreya Shankar, 5th-year PhD student in EECS at UC Berkeley.
- Natacha Crooks, Assistant Professor of EECS at UC Berkeley.
- Jiannan Wang, Associate Professor at Simon Fraser University
- Gustavo Alonso, Professor of Computer Science at ETH Zurich.
- Divesh Srivastava, head of Database Research at AT&T, President of the VLDB endowment, Co-chair of the ACM Publications Board, ACM Fellow

We will request each panelist to present a 1-minute thesis, followed by a 1-hour moderated panel discussion with interactive inputs from the audience.

Moderators Eugene Wu and Raul Castro Fernandez will divide tasks to incorporate an interactive element, allowing the audience to participate actively in the conversation.

1.2 Self-Reflection Through the Ages

The data management community has a history of self-reflection. Roughly every 5 years, a select group of senior academics convenes a closed-door meeting to discuss trends, opportunities, and challenges for the database community and issues a report [1–3, 5]. In contrast, this panel serves as an open forum for the database community to engage in the discussion. Members of the community have spoken out on threats to the community, including Dewitt's 1995 keynote [4] and Stonebraker's 2018 Ten Fears talk [6].

REFERENCES

[1] Daniel Abadi, Rakesh Agrawal, Anastasia Ailamaki, Magdalena Balazinska, Philip A Bernstein, Michael J Carey, Surajit Chaudhuri, Jeffrey Dean, AnHai Doan, Michael J Franklin, et al. 2016. The Beckman report on database research. *Commun. ACM* 59, 2 (2016), 92–99.

[2] Daniel Abadi, Anastasia Ailamaki, David Andersen, Peter Bailis, Magdalena Balazinska, Philip Bernstein, Peter Boncz, Surajit Chaudhuri, Alvin Cheung, AnHai Doan, et al. 2020. The Seattle report on database research. *ACM Sigmod Record* 48, 4 (2020), 44–53.

[3] Phil Bernstein, Michael Brodie, Stefano Ceri, David DeWitt, Mike Franklin, Hector Garcia-Molina, Jim Gray, Jerry Held, Joe Hellerstein, HV Jagadish, et al. 1998. The asilomar report on database research. *ACM Sigmod record* 27, 4 (1998), 74–80.

[4] J. David Dewitt. 1995. Database Systems: Road Kill on the Information Superhighway? VLDB Keynote.

[5] Avi Silberschatz and Stan Zdonik. 1996. Strategic directions in database systems—breaking out of the box. *ACM Computing Surveys (CSUR)* 28, 4 (1996), 764–778.

[6] Michael Stonebraker. 2018. My 10 Fears about the Future of the DBMS Field. https://www.jfsowa.com/ikl/Stonebraker.pdf.

[7] Eugene Wu. 2024. Where Does Database Research Go From Here? https://wp.sigmod.org/?p=3801.

[8] Eugene Wu. 2025. Panel Summary: Where Does Academic Database Research Go From Here? SIGMOD Blog, ACM. https://wp.sigmod.org/?p=3828 Posted June 26, 2025; updated July 4, 2025.

[9] Eugene Wu and Raul Castro Fernandez. 2025. Where Does Academic Database Research Go From Here? arXiv:2504.08948 [cs.DB] https://arxiv.org/abs/2504.08948