

## A Note from the PC Chairs

Welcome to SIGMOD 2025! This year's conference experience includes some substantive changes to the program. These changes were prominently discussed in the Call for Papers posted in November of 2023. In this note we provide background on the most significant changes: **selected research talks** and **“must-see” plenary poster sessions**.

### **Background: Growth and Goal-Setting**

SIGMOD is growing — fast! For the first time this year, we received over 1,000 submissions, and accepted 250 papers. By comparison, just 3 years ago SIGMOD accepted 151 papers from 514 submissions. In business terms, the submissions have a Compound Annual Growth Rate (CAGR) of 25%—something like double the revenue growth of leading tech companies in that period<sup>1</sup>.

This growth reflects a vibrant, highly relevant community. It also prompts us to adapt our processes to accommodate.

When contemplating change, it's important to return to first principles: What is the core purpose of an in-person scientific conference? We believe it is to *come together for shared experiences that foster a vibrant scientific community*. That includes exposure to diverse streams of research, meeting emerging and established thinkers, and engaging in face-to-face discussions in both small groups and large communal settings.

### **Keeping People Together**

We might continue with the traditional conference format, where all papers are presented in talks. However, in 2025 and beyond this would mean *eight or more concurrent activities* on the program.

Spreading our community that thin trades away some of the key benefits of holding an in-person conference. Trends suggest this will only get worse in future years.

So we made the decision to try a new process this year, with fewer concurrent sessions and 15-minute talks. Our hope is that the program will bring the community together in larger groups for a more engaging and thoughtful conference experience.

### **Plenary Poster Sessions and Videos**

We are conscious that limiting the number of speaking slots is a trade-off. There is excellent work accepted to SIGMOD 2025 that will not be presented in a talk session.

---

<sup>1</sup> FinanceCharts, <https://www.financecharts.com/stocks/AAPL/summary/revenue-ttm-cagr>, <https://www.financecharts.com/stocks/AMZN/summary/revenue-ttm-cagr>, <https://www.financecharts.com/stocks/GOOG/summary/revenue-ttm-cagr>, <https://www.financecharts.com/stocks/MSFT/summary/revenue-ttm-cagr>, <https://www.financecharts.com/stocks/META/summary/revenue-ttm-cagr>, viewed June 20, 2025.

To address this reality, our attention and value system shifts to poster sessions, as in many other scientific communities. The *SIGMOD '25 poster sessions will be plenary*, meaning nothing will be happening concurrent with the poster sessions. This “must-see” experience should offer undistracted opportunities for like-minded researchers to find each other and review work quite deeply. It also gives every researcher with an accepted paper a chance to gain visibility, and engage in active dialogue with interested colleagues.

Additionally, modern times do bring some benefits: the ability to host a generous number of digital artifacts. Authors of all papers have the opportunity to upload a full video presentation on their paper to the ACM digital library as part of the SIGMOD proceedings and their journal article.

### **Why use random sampling?**

Many scientific conferences assign speaking slots to papers or speakers that were deemed to be among the best of those accepted. In consultation with the SIGMOD Executive Committee, we decided that 2025 was not the year to create additional exclusive tiers among accepted papers, beyond the tradition of Award Papers. Hence *receiving a speaking slot at SIGMOD 2025 is not a differentiating feature*, and should not be used to judge the relative merit of papers at the conference.

### **What was the sampling procedure?**

This year’s program has room for 80 talks from the 250 papers that were accepted. To get to our goal of creating a cohesive program that reflects the variety of research in our community, we adopted a randomized method of talk selection used in many SIGMOD papers: stratified sampling.

#### *The Strata*

On submission, the authors of each paper chose a primary topic area, selecting from a two-level hierarchy: a broad first-level category (e.g., “Models and Languages”) and optionally a more specific second-level category (e.g., “Spatial and Temporal Data Management”). Once the PC review process was complete, we assessed the number of accepted papers in each category. When a second-level category contained a substantial number of papers (more than 10), it was treated as its own stratum. Otherwise, it was grouped with related papers under the broader first-level category. This process resulted in 14 distinct strata, listed below in descending order of size.

- Models and Languages -> Graphs, social networks, web data, and semantic web
- Data Management Systems
- Modern AI & Data Management
- Data Management Systems -> Query processing and optimization
- Models and Languages
- Data Governance -> Data security, privacy, and access control

- Modern AI & Data Management -> Data and metadata management for machine learning pipelines
- Data Management Systems -> Database systems on emerging hardware
- Data Management Systems -> Storage, indexing, and physical database design
- Data Governance
- Data Governance -> Data quality, data cleaning
- Models and Languages -> Spatial and temporal data management
- Human-Centric Data Management
- Data-Driven Applications

### *The Speaking Slots*

Award papers all received a speaking slot. Then, one paper was randomly sampled from each stratum to ensure topic coverage. Among the remaining slots, sampling was performed with weights proportional to the number of papers within each stratum. This approach allowed us to balance intellectual diversity (covering the full range of topics) with representation of community interest (via weights across strata) and fairness in the process (randomization without intervention).

We then contacted the authors of the selected papers to confirm their attendance and willingness to present at the conference. If any declined the speaking slot, we moved on to the next randomly sampled paper. A few weeks later, we reached out again to reconfirm their attendance, ensuring that the speaking slots were reserved only for those who committed to attending.

We then created the program by grouping the selected papers thematically into 16 different sessions, giving each session its own representative name as is currently published online.

### **How will we evaluate the experiment?**

As scientists and engineers, we are conscious of the benefit of being open-minded about trying new things, while also collecting feedback and iteratively improving over time. Together with the Executive Committee (EC), we are committed to evaluating both the strengths and limitations of this year's process and program. To that end, the EC will be releasing a survey at the end of the conference to assess how these changes have affected the experience of all participants including authors, presenters, and attendees.

We encourage you to take a few minutes to thoughtfully complete the survey when it is launched, and share your feedback. Your input is essential in shaping the future of SIGMOD conferences going forward.