

## Keynote

R. Rajwar: Specializing the General

## Session 1: Real-time Data

1. M. Karpathiotakis, I. Alagiannis, T. Heinis, M. Branco, A. Ailamaki: Virtualizing Data with ViDa: Running Queries to Build a Database
2. O. Kennedy, L. Ziarek: Just-in-time Data Structures
3. T. Johnson, V. Shkapenyuk: Data Stream Warehousing in Tidalrace
4. R. Fernandez, P. Pietzuch, J. Koshy, J. Kreps, D. Lin, N. Narkhede, J. Rao, C. Riccomini, G. Wang: Liquid: Unifying Nearline and Offline Big Data Integration

## Session 2: Column Stores

1. Marcel Kornacker et al.: Impala
2. Shilpa Lawande et al.: Vertica
3. Vidhya Srinivasan et al.: Redshift

## Session 3: Data-driven Science

1. J. Duggan, M. Brodie: Hephaestus: Virtual Experiments for Data-Intensive Science
2. A. Parameswaran et al.: DataHub: Collaborative Data Science & Dataset Version Management at Scale
3. Y. Diao, A. Roy, T. Bloom: Building Highly-Optimized, Low-Latency Pipelines for Genomic Data Analysis

## Session 4: Data Integration I

1. H. Lee, S. Balakrishnan, A. Halevy, B. Harb, H. Lee, J. Madhavan, A. Rostamizadeh, W. Shen, K. Wilder, F. Wu, C. Yu: Applying Web Tables in Practice
2. T. Rekatsinas, X. Dong, L. Getoor, D. Srivastava: Finding Quality in Quantity: The Challenge of Discovering Valuable Sources for Integration
3. Z. Ives, Z. Yan, N. Zheng, B. Litt, J. Wagenaar: Looking at Everything in Context

## Session 5: Storage Systems

1. I. Oukid, W. Lehner, T. Kissinger, T. Willhalm, P. Bumbulis: Instant Recovery for Main Memory Databases
2. I. Manolescu, F. Bugiotti, D. Bursztyn, A. Deutsch, I. Ileana: Invisible Glue: Scalable Self-Tuning Multi-Stores
3. P. Helland: Immutability Changes Everything

## Session 6: Model Management

1. **H. Zhen, D. Gawlick:** Management of Flexible Schema Data in RDBMSs
2. **H. Mühleisen, M. Kersten, S. Manegold:** Capturing the Laws of (Data) Nature
3. **Y. Katsis, Y. Freund, Y. Papakonstantinou:** Combining Databases and Signal Processing in Plato
4. **D. Crankshaw, P. Bailis, J. Gonzalez, H. Li, Z. Zhang, M. Franklin, M. Jordan:** The Missing Piece in Complex Analytics: Low Latency, Scalable Model Management and Serving with Velox

## Session 7: Cloud

1. **J. Ortiz, M. Balazinska, V. de Almeida:** Changing the Face of Database Services
2. **W. Lang, R. Nehme, I. Rae:** Database Optimization in the Cloud: Where Costs, Partial Results, and Consumer Choice Meet
3. **C. Curion, A. Vulimiri, B. Godfrey, K. Karanasos, G. Varghese:** WANalytics: Analytics for a Geo-Distributed Data-Intensive World
4. **J. Levandoski, D. Lomet, S. Sengupta, R. Stutsman, R. Wang:** High Performance Transactions in Deuteronomy

## Session 8: Data Integration II

1. **Mona Vernon et al.:** Thomson Reuters
2. **Z. Abedjan, J. Morcos, M. Gubanov, I. Ilyas, M. Stonebraker:** XForms: Leveraging the Web for Semantic Transformations
3. **J. Hellerstein, J. Heer, S. Kandel:** Predictive Interaction for Data Transformation

## Session 9: Query Processing (under new assumptions)

1. **J. Fan, J. Patel:** The case against specialized graph analytics engines
2. **R. Halstead, I. Absalyamov, W. Najjar, V. Tsotras:** FPGA-based Multithreading for In-Memory Hash Joins
3. **R. Ramamurthy, A. Mohapatra, R. Kaushik:** Raising Authorization Awareness in a DBMS
4. **A. Crotty, A. Galakatos, K. Dursun, T. Kraska, U. Cetintemel, S. Zdonik:** Tupleware: ``Big'' Data, Big Analytics, Small Clusters

## Session 10: Knowledge Management

1. **F. Suchanek, F. Mahdisoltani, J. Biega:** YAGO3: A Knowledge Base from Multilingual Wikipedias
2. **Y. Amsterdamer, S. Davidson, A. Kukliansky, T. Milo, S. Novgorodov, A. Somech:** Managing General and Individual Knowledge in Crowd Mining Applications
3. **I. Terrizzano, M. Roth, P. Schwarz:** tbd. (IBM)