

Stanford



Matei Zaharia

Assistant Professor of Computer Science and, by courtesy, of Electrical Engineering

Bio

BIO

Homepage: <https://cs.stanford.edu/~matei/>

ACADEMIC APPOINTMENTS

- Assistant Professor, Computer Science
- Assistant Professor (By courtesy), Electrical Engineering

LINKS

- Matei Zaharia's Homepage: <https://cs.stanford.edu/~matei/>

Teaching

COURSES

2019-20

- Principles of Data-Intensive Systems: CS 245 (Win)
- Value of Data and AI: CS 320 (Win)

2018-19

- Cloud Computing Technology: CS 349D (Aut)
- Principles of Data-Intensive Systems: CS 245 (Spr)

2017-18

- Cloud Computing Technology: CS 349D (Aut)
- Parallel Computing: CS 149 (Win)

2016-17

- Parallel Computing: CS 149 (Win)
- Project in Mining Massive Data Sets: CS 341 (Spr)

STANFORD ADVISEES

Doctoral Dissertation Reader (AC)

Henry Corrigan-Gibbs, Edward Gan, Zhihao Jia, Yawen Wang, Stephen Yang

Postdoctoral Faculty Sponsor

Fiodar Kazhamiaka

Doctoral Dissertation Advisor (AC)

Firas Abuzaid, Shoumik Palkar

Orals Evaluator

Benjamin Braun, Henry Corrigan-Gibbs, Stephen Yang

Master's Program Advisor

Nikhil Athreya, Advay Pal

Doctoral Dissertation Co-Advisor (AC)

James Thomas

Doctoral (Program)

Cody Coleman, Peter Kraft, Deepak Narayanan, Deepti Raghavan, Pratiksha Thaker

Publications

PUBLICATIONS

- **To Index or Not to Index: Optimizing Exact Maximum Inner Product Search**
Abuzaid, F., Sethi, G., Bailis, P., Zaharia, M., IEEE
IEEE.2019: 1250–61
- **Structured Streaming: A Declarative API for Real-Time Applications in Apache Spark**
Armbrust, M., Das, T., Torres, J., Yavuz, B., Zhu, S., Xin, R., Ghodsi, A., Stoica, I., Zaharia, M., Das, G., Jermaine, C., Bernstein, P., Eldawy, et al
ASSOC COMPUTING MACHINERY.2018: 601–13
- **MISTIQUE: A System to Store and Query Model Intermediates for Model Diagnosis**
Vartak, M., da Trindade, J. F., Madden, S., Zaharia, M., Das, G., Jermaine, C., Bernstein, P., Eldawy, A.
ASSOC COMPUTING MACHINERY.2018: 1285–1300
- **Splinter: Practical Private Queries on Public Data**
Wang, F., Yun, C., Goldwasser, S., Vaikuntanathan, V., Zaharia, M., USENIX Assoc
USENIX ASSOC.2017: 299–313
- **DIY Hosting for Online Privacy**
Palkar, S., Zaharia, M., Assoc Comp Machinery
ASSOC COMPUTING MACHINERY.2017: 1–7
- **Making Caches Work for Graph Analytics**
Zhang, Y., Kiriansky, V., Mendis, C., Amarasinghe, S., Zaharia, M., Nie, J. Y., Obradovic, Z., Suzumura, T., Ghosh, R., Nambiar, R., Wang, C., Zang, H., BaezaYates, et al
IEEE.2017: 293–302
- **Apache Spark: A Unified Engine for Big Data Processing** *COMMUNICATIONS OF THE ACM*
Zaharia, M., Xin, R. S., Wendell, P., Das, T., Armbrust, M., Dave, A., Meng, X., Rosen, J., Venkataraman, S., Franklin, M. J., Ghodsi, A., Gonzalez, J., Shenker, et al
2016; 59 (11): 56-65
- **MLlib: Machine Learning in Apache Spark** *JOURNAL OF MACHINE LEARNING RESEARCH*
Meng, X., Bradley, J., Yavuz, B., Sparks, E., Venkataraman, S., Liu, D., Freeman, J., Tsai, D. B., Amde, M., Owen, S., Xin, D., Xin, R., Franklin, et al
2016; 17