

Biographic Summary

Joseph Gonzalez is an Associate Professor at UC Berkeley and a founding member of the Berkeley Sky and RISE labs where he studies the design of next generation cloud systems and systems for high-performance machine learning. The RISE Lab is an NSF Expedition center and both the Sky and RISE Labs are backed by a consortium of leading international industrial sponsors. Gonzalez's research addresses problems in data systems, neural network design, compilers and distributed systems for large scale machine learning, natural language processing computer vision, robotics, autonomous driving, and graph analytics. Gonzalez has published over 100 papers in top-tier venues spanning machine learning (JMLR, ICML, AISTATS, UAI, CVPR, ECCV, ICCV, RAL, ICRA) and computer systems (VLDB, OSDI, NSDI, EuroSys, MLSys, SoCC). He has made significant contributions to the design of efficient neural networks to reduce computation and enable meta-learning. His work on both graph systems and prediction serving systems helped define their respective disciplines. Gonzalez also led the development and teaches the large upper-division data science class (Data100). Outside of Berkeley, Gonzalez is co-founder and VP of product at Aqueduct Inc. Prior to joining Berkeley, Gonzalez co-founded Turi Inc (formerly GraphLab) based on his thesis work and created the GraphX project (now part of Apache Spark). Gonzalez's innovative work has earned him significant recognition, including the Okawa Research Grant, the NSF Expedition Award, and the NSF Early CAREER Award.

For more information visit <http://eecs.berkeley.edu/~jegonzal>.

Employment and Professional Experience

Nov 2020 — Present	Co-founder, Aqueduct
July 2021 — Present	Associate Professor in the EECS Department
Jan 2016 — July 2021	Assistant Professor in the EECS Department
May 2013 — July 2016	Co-founder at Turi Inc., Acquired by Apple Inc.
Jan 2013 — July 2015	Post-doc in the UC Berkeley AMPLab
June 2011 — Aug 2011	Yahoo! Research Intern
June 2007 — Aug 2007	AT&T Labs Intern

Education

Ph.D., Machine Learning: December 2012. Machine Learning Department, School of Computer Science at Carnegie Mellon University.

Title: *“Parallel and Distributed Algorithms and Systems for Probabilistic Reasoning.”*

Thesis Advisor: Carlos Guestrin

M.S., Machine Learning: December 2009. Machine Learning Department, School of Computer Science at Carnegie Mellon University.

B.S. with Honors, Computer Science: June 2006. California Institute of Technology.

Awards

- **Sigmod Systems Award [2022]:** Award for work on Apache Spark: an innovative, widely-used, open-source, unified data processing system encompassing relational, streaming, and machine-learning workloads.
- **NSF Early CAREER Award [2019]:** Funding awarded for research into the design of models and system for efficient prediction. Details: "*CAREER: Towards the Design of Models and Systems for Efficient Prediction*" (Award number 1846431).
- **Okawa Foundation Research Grant Award [2018]:** Funding award for research into the design of systems for machine learning life-cycle management.
- **Nominated for ACM Dissertation Award [2013]:** My thesis was nominated by CMU for the ACM Dissertation Award.
- **AT&T Labs Fellowship (2007):** Graduate research stipend for academic achievement as an under-represented minority.
- **NSF Graduate Research Fellowship (2007):** Graduate research stipend for 3 years.
- **NASA Space Act Award (2005):** Awarded for a sizeable contribution to space exploration.
- **NASA Inventions and Contributions Board Award (2005):** Awarded for the development of an innovative new technology that has made a contribution to space exploration.
- **Caltech Presidential Award (2002-2006):** I was awarded tuition for research and academic achievements.

Peer Reviewed Publications

- [1] Yonghao Zhuang, Lianmin Zheng, Zhuohan Li, Eric Xing, Qirong Ho, Joseph E. Gonzalez, Ion Stoica, Hao Zhang, and Hexu Zhao. "On Optimizing the Communication of Model Parallelism". In: *Proceedings of Machine Learning and Systems*. 2023. URL: https://proceedings.mlsys.org/paper_files/paper/2023/hash/d0b9a3081f811b2a307c38ad457a487c-Abstract-mlsys2023.html.
- [2] Ionel Gog, Sukrit Kalra, Peter Schafhalter, Joseph E. Gonzalez, and Ion Stoica. "D3: A Dynamic Deadline-Driven Approach for Building Autonomous Vehicles". In: *Proceedings of the Seventeenth European Conference on Computer Systems*. EuroSys '22. Rennes, France: Association for Computing Machinery, 2022, pp. 453–471. ISBN: 9781450391627. DOI: 10.1145/3492321.3519576. URL: <https://doi.org/10.1145/3492321.3519576>.
- [3] Wenshuo Guo, Kirthivasan Kandasamy, Joseph Gonzalez, Michael Jordan, and Ion Stoica. "Learning Competitive Equilibria in Exchange Economies with Bandit Feedback". In: *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*. Ed. by Gustau Camps-Valls, Francisco J. R. Ruiz, and Isabel Valera. Vol. 151. Proceedings of Machine Learning Research. PMLR, Mar. 2022, pp. 6200–6224. URL: <https://proceedings.mlr.press/v151/guo22a.html>.
- [4] Paras Jain, Safeen Huda, Martin Maas, Joseph E. Gonzalez, Ion Stoical, and Azalia Mirhoseini. "Learning to Design Accurate Deep Learning Accelerators with Inaccurate Multipliers". In: *2022 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. 2022, pp. 184–189. DOI: 10.23919/DATE54114.2022.9774607. URL: https://www.parasjain.com/projects/21autoapprox_date/paper.pdf.
- [5] Xiaoxuan Liu, Lianmin Zheng, Dequan Wang, Yukuo Cen, Weize Chen, Xu Han, Jianfei Chen, Zhiyuan Liu, Jie Tang, Joseph E. Gonzalez, Michael Mahoney, and Alvin Cheung. "GACT: Activation Compressed Training for Generic Network Architectures". In: *Proceedings of the 39th International Conference on Machine Learning*. Ed. by Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato. Vol. 162. Proceedings of Machine Learning Research. PMLR, July 2022, pp. 14139–14152. URL: <https://proceedings.mlr.press/v162/liu22v.html>.

- [6] Shishir G. Patil, Paras Jain, Prabal Dutta, Ion Stoica, and Joseph Gonzalez. “POET: Training Neural Networks on Tiny Devices with Integrated Rematerialization and Paging”. In: *Proceedings of the 39th International Conference on Machine Learning*. Ed. by Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato. Vol. 162. Proceedings of Machine Learning Research. PMLR, July 2022, pp. 17573–17583. URL: <https://proceedings.mlr.press/v162/patil22b.html>.
- [7] Suzanne Petryk, Lisa Dunlap, Keyan Nasseri, Joseph Gonzalez, Trevor Darrell, and Anna Rohrbach. “On Guiding Visual Attention With Language Specification”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2022, pp. 18092–18102.
- [8] Albert Wilcox, Ashwin Balakrishna, Brijen Thananjeyan, Joseph E. Gonzalez, and Ken Goldberg. “LS3: Latent Space Safe Sets for Long-Horizon Visuomotor Control of Sparse Reward Iterative Tasks”. In: *Proceedings of the 5th Conference on Robot Learning*. Ed. by Aleksandra Faust, David Hsu, and Gerhard Neumann. Vol. 164. Proceedings of Machine Learning Research. PMLR, Nov. 2022, pp. 959–969. URL: <https://proceedings.mlr.press/v164/wilcox22a.html>.
- [9] Bichen Wu, Ruizhe Cheng, Peizhao Zhang, Tianren Gao, Joseph E. Gonzalez, and Peter Vajda. “Data Efficient Language-Supervised Zero-Shot Recognition with Optimal Transport Distillation”. In: *International Conference on Learning Representations*. 2022. URL: <https://openreview.net/forum?id=G89-1yZLFHk>.
- [10] Tianjun Zhang, Tongzheng Ren, Mengjiao Yang, Joseph Gonzalez, Dale Schuurmans, and Bo Dai. “Making Linear MDPs Practical via Contrastive Representation Learning”. In: *Proceedings of the 39th International Conference on Machine Learning*. Ed. by Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato. Vol. 162. Proceedings of Machine Learning Research. PMLR, July 2022, pp. 26447–26466. URL: <https://proceedings.mlr.press/v162/zhang22x.html>.
- [11] Zhengming Zhang, Ashwinee Panda, Linyue Song, Yaoqing Yang, Michael Mahoney, Prateek Mittal, Ramchandran Kannan, and Joseph Gonzalez. “Neurotoxin: Durable Backdoors in Federated Learning”. In: *Proceedings of the 39th International Conference on Machine Learning*. Ed. by Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato. Vol. 162. Proceedings of Machine Learning Research. PMLR, July 2022, pp. 26429–26446. URL: <https://proceedings.mlr.press/v162/zhang22w.html>.
- [12] Lianmin Zheng, Zhuohan Li, Hao Zhang, Yonghao Zhuang, Zhifeng Chen, Yanping Huang, Yida Wang, Yuanzhong Xu, Danyang Zhuo, Eric P. Xing, Joseph E. Gonzalez, and Ion Stoica. “Alpa: Automating Inter- and Intra-Operator Parallelism for Distributed Deep Learning”. In: *16th USENIX Symposium on Operating Systems Design and Implementation (OSDI 22)*. Carlsbad, CA: USENIX Association, July 2022, pp. 559–578. ISBN: 978-1-939133-28-1. URL: <https://www.usenix.org/conference/osdi22/presentation/zheng-lianmin>.
- [13] Raghav Anand, Jeffrey Ichnowski, Chenggang Wu, Joseph M. Hellerstein, Joseph E. Gonzalez, and Ken Goldberg. “Serverless Multi-Query Motion Planning for Fog Robotics”. In: *2021 IEEE International Conference on Robotics and Automation (ICRA)*. May 2021, pp. 7457–7463. DOI: 10.1109/ICRA48506.2021.9561571. URL: <https://goldberg.berkeley.edu/pubs/ICRA21-ichnowski-serverless-motion-planning-submitted.pdf>.
- [14] Robert Avram, Jeffrey Olgin, Alvin Wan, Zeeshan Ahmed, Louis Verreault-Julien, Sean Abreau, Derek Wan, Joseph E. Gonzalez, Derek So, Krishan Soni, and Geoffrey Tison. “CATHAI: FULLY AUTOMATED CORONARY ANGIOGRAPHY INTERPRETATION AND STENOSIS DETECTION USING A DEEP LEARNING-BASED ALGORITHMIC PIPELINE”. In: *Journal of the American College of Cardiology* 77.18_Supplement_1 (2021), pp. 3244–3244. DOI: 10.1016/S0735-1097(21)04598-8. eprint: <https://www.jacc.org/doi/pdf/10.1016/S0735-1097%2821%2904598-8>. URL: <https://www.jacc.org/doi/abs/10.1016/S0735-1097%2821%2904598-8>.

- [15] Jianfei Chen, Lianmin Zheng, Zhewei Yao, Dequan Wang, Ion Stoica, Michael Mahoney, and Joseph Gonzalez. “ActNN: Reducing Training Memory Footprint via 2-Bit Activation Compressed Training”. In: *Proceedings of the 38th International Conference on Machine Learning*. Ed. by Marina Meila and Tong Zhang. Vol. 139. Proceedings of Machine Learning Research. PMLR, July 2021, pp. 1803–1813. URL: <https://proceedings.mlr.press/v139/chen21z.html>.
- [16] Kaiyuan Eric Chen, Yafei Liang, Nikhil Jha, Jeffrey Ichnowski, Michael Danielczuk, Joseph Gonzalez, John Kubiawicz, and Ken Goldberg. “FogROS: An Adaptive Framework for Automating Fog Robotics Deployment”. In: *2021 IEEE 17th International Conference on Automation Science and Engineering (CASE)*. Aug. 2021, pp. 2035–2042. DOI: 10.1109/CASE49439.2021.9551628. URL: <https://arxiv.org/abs/2108.11355>.
- [17] Ruizhe Cheng, Bichen Wu, Peizhao Zhang, Peter Vajda, and Joseph E. Gonzalez. “Data-Efficient Language-Supervised Zero-Shot Learning With Self-Distillation”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*. June 2021, pp. 3119–3124.
- [18] Xiaoliang Dai, Alvin Wan, Peizhao Zhang, Bichen Wu, Zijian He, Zhen Wei, Kan Chen, Yuandong Tian, Matthew Yu, Peter Vajda, and Joseph E. Gonzalez. “FBNetV3: Joint Architecture-Recipe Search Using Predictor Pretraining”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2021, pp. 16276–16285.
- [19] Lisa Dunlap, Kirthevasan Kandasamy, Ujval Misra, Richard Liaw, Michael Jordan, Ion Stoica, and Joseph E. Gonzalez. “Elastic Hyperparameter Tuning on the Cloud”. In: *Proceedings of the ACM Symposium on Cloud Computing*. SoCC ’21. Seattle, WA, USA: Association for Computing Machinery, 2021, pp. 33–46. ISBN: 9781450386388. DOI: 10.1145/3472883.3486989. URL: <https://doi.org/10.1145/3472883.3486989>.
- [20] Sayna Ebrahimi, Suzanne Petryk, Akash Gokul, William Gan, Joseph E. Gonzalez, Marcus Rohrbach, and Trevor Darrell. “Remembering for the right reasons: Explanations reduce catastrophic forgetting”. In: *Applied AI Letters* 2.4 (2021), e44. DOI: <https://doi.org/10.1002/ail2.44>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/ail2.44>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/ail2.44>.
- [21] Aditya Ganapathi, Priya Sundaresan, Brijen Thananjeyan, Ashwin Balakrishna, Daniel Seita, Jennifer Grannen, Minho Hwang, Ryan Hoque, Joseph E. Gonzalez, Nawid Jamali, Katsu Yamane, Soshi Iba, and Ken Goldberg. “Learning Dense Visual Correspondences in Simulation to Smooth and Fold Real Fabrics”. In: *2021 IEEE International Conference on Robotics and Automation (ICRA)*. May 2021, pp. 11515–11522. DOI: 10.1109/ICRA48506.2021.9561980. URL: <https://arxiv.org/abs/2003.12698>.
- [22] Ionel Gog, Sukrit Kalra, Peter Schafhalter, Matthew A. Wright, Joseph E. Gonzalez, and Ion Stoica. “Pylot: A Modular Platform for Exploring Latency-Accuracy Tradeoffs in Autonomous Vehicles”. In: *2021 IEEE International Conference on Robotics and Automation (ICRA)*. May 2021, pp. 8806–8813. DOI: 10.1109/ICRA48506.2021.9561747. URL: <https://arxiv.org/abs/2104.07830>.
- [23] J. Weston Hughes, Jeffrey E. Olgin, Robert Avram, Sean A. Abreau, Taylor Sittler, Kaahan Radia, Henry Hsia, Tomos Walters, Byron Lee, Joseph E. Gonzalez, and Geoffrey H. Tison. “Performance of a Convolutional Neural Network and Explainability Technique for 12-Lead Electrocardiogram Interpretation”. In: *JAMA Cardiology* 6.11 (Nov. 2021), pp. 1285–1295. ISSN: 2380-6583. DOI: 10.1001/jamacardio.2021.2746. eprint: https://jamanetwork.com/journals/jamacardiology/articlepdf/2782549/jamacardiology_hughes_2021_oi_210051_1635348688.2775.pdf. URL: <https://doi.org/10.1001/jamacardio.2021.2746>.
- [24] Jeffrey Ichnowski, Paras Jain, Bartolomeo Stellato, Goran Banjac, Michael Luo, Francesco Borrelli, Joseph E Gonzalez, Ion Stoica, and Ken Goldberg. “Accelerating Quadratic Optimization with Reinforcement Learning”. In: *Advances in Neural Information Processing Systems*. Ed. by M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan. Vol. 34. Curran Asso-

- ciates, Inc., 2021, pp. 21043–21055. URL: <https://proceedings.neurips.cc/paper/2021/file/afdec7005cc9f14302cd0474fd0f3c96-Paper.pdf>.
- [25] Anand Padmanabha Iyer, Qifan Pu, Kishan Patel, Joseph E. Gonzalez, and Ion Stoica. “TEGRA: Efficient Ad-Hoc Analytics on Evolving Graphs”. In: *18th USENIX Symposium on Networked Systems Design and Implementation (NSDI 21)*. USENIX Association, Apr. 2021, pp. 337–355. ISBN: 978-1-939133-21-2. URL: <https://www.usenix.org/conference/nsdi21/presentation/iyer>.
- [26] Eric Liang, Zhanghao Wu, Michael Luo, Sven Mika, Joseph E Gonzalez, and Ion Stoica. “RLlib Flow: Distributed Reinforcement Learning is a Dataflow Problem”. In: *Advances in Neural Information Processing Systems*. Ed. by M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan. Vol. 34. Curran Associates, Inc., 2021, pp. 5506–5517. URL: <https://proceedings.neurips.cc/paper/2021/file/2bce32ed409f5ebcee2a7b417ad9beed-Paper.pdf>.
- [27] Ujval Misra, Richard Liaw, Lisa Dunlap, Romil Bhardwaj, Kirthevasan Kandasamy, Joseph E. Gonzalez, Ion Stoica, and Alexey Tumanov. “RubberBand: Cloud-Based Hyperparameter Tuning”. In: *Proceedings of the Sixteenth European Conference on Computer Systems*. EuroSys ’21. Online Event, United Kingdom: Association for Computing Machinery, 2021, pp. 327–342. ISBN: 9781450383349. DOI: 10.1145/3447786.3456245. URL: <https://doi.org/10.1145/3447786.3456245>.
- [28] Charles Packer, Pieter Abbeel, and Joseph E Gonzalez. “Hindsight Task Relabelling: Experience Replay for Sparse Reward Meta-RL”. In: *Advances in Neural Information Processing Systems*. Ed. by M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan. Vol. 34. Curran Associates, Inc., 2021, pp. 2466–2477. URL: <https://proceedings.neurips.cc/paper/2021/file/1454ca2270599546dfcd2a3700e4d2f1-Paper.pdf>.
- [29] Samuel Paradis, Minho Hwang, Brijen Thananjeyan, Jeffrey Ichnowski, Daniel Seita, Danyal Fer, Thomas Low, Joseph E. Gonzalez, and Ken Goldberg. “Intermittent Visual Servoing: Efficiently Learning Policies Robust to Instrument Changes for High-precision Surgical Manipulation”. In: *2021 IEEE International Conference on Robotics and Automation (ICRA)*. May 2021, pp. 7166–7173. DOI: 10.1109/ICRA48506.2021.9561070. URL: <https://arxiv.org/abs/2011.06163>.
- [30] Nathan Pemberton, Johann Schleier-Smith, and Joseph E. Gonzalez. “The RESTless Cloud”. In: *Proceedings of the Workshop on Hot Topics in Operating Systems*. HotOS ’21. Ann Arbor, Michigan: Association for Computing Machinery, 2021, pp. 49–57. ISBN: 9781450384384. DOI: 10.1145/3458336.3465280. URL: <https://doi.org/10.1145/3458336.3465280>.
- [31] Devin Petersohn, Dixin Tang, Rehan Durrani, Areg Melik-Adamyanyan, Joseph E. Gonzalez, Anthony D. Joseph, and Aditya G. Parameswaran. “Flexible Rule-Based Decomposition and Metadata Independence in Modin: A Parallel Dataframe System”. In: *Proc. VLDB Endow.* 15.3 (Nov. 2021), pp. 739–751. ISSN: 2150-8097. DOI: 10.14778/3494124.3494152. URL: <https://doi.org/10.14778/3494124.3494152>.
- [32] Johann Schleier-Smith, Vikram Sreekanti, Anurag Khandelwal, Joao Carreira, Neeraja J. Yadwadkar, Raluca Ada Popa, Joseph E. Gonzalez, Ion Stoica, and David A. Patterson. “What Serverless Computing is and Should Become: The next Phase of Cloud Computing”. In: *Commun. ACM* 64.5 (Apr. 2021), pp. 76–84. ISSN: 0001-0782. DOI: 10.1145/3406011. URL: <https://doi.org/10.1145/3406011>.
- [33] Brijen Thananjeyan, Ashwin Balakrishna, Suraj Nair, Michael Luo, Krishnan Srinivasan, Minho Hwang, Joseph E. Gonzalez, Julian Ibarz, Chelsea Finn, and Ken Goldberg. “Recovery RL: Safe Reinforcement Learning With Learned Recovery Zones”. In: *IEEE Robotics and Automation Letters* 6.3 (July 2021), pp. 4915–4922. ISSN: 2377-3766. DOI: 10.1109/LRA.2021.3070252. URL: <https://arxiv.org/abs/2010.15920>.
- [34] Brijen Thananjeyan, Kirthevasan Kandasamy, Ion Stoica, Michael Jordan, Ken Goldberg, and Joseph Gonzalez. “Resource Allocation in Multi-armed Bandit Exploration: Overcoming Sublinear Scaling with Adaptive Parallelism”. In: *Proceedings of the 38th International Conference on Machine Learning*. Ed. by Marina Meila and Tong Zhang. Vol. 139. Proceedings of Machine Learn-

- ing Research. PMLR, July 2021, pp. 10236–10246. URL: <https://proceedings.mlr.press/v139/thananjeyan21a.html>.
- [35] Vainavi Viswanath, Jennifer Grannen, Priya Sundaresan, Brijen Thananjeyan, Ashwin Balakrishna, Ellen Novoseller, Jeffrey Ichnowski, Michael Laskey, Joseph E. Gonzalez, and Ken Goldberg. “Disentangling Dense Multi-Cable Knots”. In: *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2021, pp. 3731–3738. DOI: 10.1109/IR051168.2021.9636397. URL: <https://arxiv.org/abs/2106.02252>.
- [36] Alvin Wan, Lisa Dunlap, Daniel Ho, Jihan Yin, Scott Lee, Suzanne Petryk, Sarah Adel Bargal, and Joseph E. Gonzalez. “NBDT: Neural-Backed Decision Tree”. In: *International Conference on Learning Representations*. 2021. URL: <https://openreview.net/forum?id=mCLVeEplNE>.
- [37] Guanhua Wang, Zhuang Liu, Brandon Hsieh, Siyuan Zhuang, Joseph Gonzalez, Trevor Darrell, and Ion Stoica. “sensAI: ConvNets Decomposition via Class Parallelism for Fast Inference on Live Data”. In: *Proceedings of Machine Learning and Systems*. Ed. by A. Smola, A. Dimakis, and I. Stoica. Vol. 3. 2021, pp. 664–679. URL: <https://proceedings.mlsys.org/paper/2021/file/c4ca4238a0b923820dcc509a6f75849b-Paper.pdf>.
- [38] Xin Wang, Thomas E. Huang, Benlin Liu, Fisher Yu, Xiaolong Wang, Joseph E. Gonzalez, and Trevor Darrell. “Robust Object Detection via Instance-Level Temporal Cycle Confusion”. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. Oct. 2021, pp. 9143–9152.
- [39] Bichen Wu, Chenfeng Xu, Xiaoliang Dai, Alvin Wan, Peizhao Zhang, Zhicheng Yan, Masayoshi Tomizuka, Joseph E. Gonzalez, Kurt Keutzer, and Peter Vajda. “Visual Transformers: Where Do Transformers Really Belong in Vision Models?” In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. Oct. 2021, pp. 599–609.
- [40] Zhanghao Wu, Paras Jain, Matthew Wright, Azalia Mirhoseini, Joseph E Gonzalez, and Ion Stoica. “Representing Long-Range Context for Graph Neural Networks with Global Attention”. In: *Advances in Neural Information Processing Systems*. Ed. by M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan. Vol. 34. Curran Associates, Inc., 2021, pp. 13266–13279. URL: <https://proceedings.neurips.cc/paper/2021/file/6e67691b60ed3e4a55935261314dd534-Paper.pdf>.
- [41] Kevin Yang, Tianjun Zhang, Chris Cummins, Brandon Cui, Benoit Steiner, Linnan Wang, Joseph E Gonzalez, Dan Klein, and Yuandong Tian. “Learning Space Partitions for Path Planning”. In: *Advances in Neural Information Processing Systems*. Ed. by M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan. Vol. 34. Curran Associates, Inc., 2021, pp. 378–391. URL: <https://proceedings.neurips.cc/paper/2021/file/03a3655fff3e9bdea48de9f49e938e32-Paper.pdf>.
- [42] Yaoqing Yang, Liam Hodgkinson, Ryan Theisen, Joe Zou, Joseph E Gonzalez, Kannan Ramchandran, and Michael W Mahoney. “Taxonomizing local versus global structure in neural network loss landscapes”. In: *Advances in Neural Information Processing Systems*. Ed. by M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan. Vol. 34. Curran Associates, Inc., 2021, pp. 18722–18733. URL: <https://proceedings.neurips.cc/paper/2021/file/9b72e31dac81715466cd580a448cf823-Paper.pdf>.
- [43] Tianjun Zhang, Paria Rashidinejad, Jiantao Jiao, Yuandong Tian, Joseph E Gonzalez, and Stuart Russell. “MADE: Exploration via Maximizing Deviation from Explored Regions”. In: *Advances in Neural Information Processing Systems*. Ed. by M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan. Vol. 34. Curran Associates, Inc., 2021, pp. 9663–9680. URL: <https://proceedings.neurips.cc/paper/2021/file/5011bf6d8a37692913fce3a15a51f070-Paper.pdf>.
- [44] Tianjun Zhang, Huazhe Xu, Xiaolong Wang, Yi Wu, Kurt Keutzer, Joseph E Gonzalez, and Yuandong Tian. “NovelD: A Simple yet Effective Exploration Criterion”. In: *Advances in Neural Information Processing Systems*. Ed. by M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and

- J. Wortman Vaughan. Vol. 34. Curran Associates, Inc., 2021, pp. 25217–25230. URL: <https://proceedings.neurips.cc/paper/2021/file/d428d070622e0f4363fcea11f4a3576-Paper.pdf>.
- [45] Zhengming Zhang, Yaoqing Yang, Zhewei Yao, Yujun Yan, Joseph E. Gonzalez, Kannan Ramchandran, and Michael W. Mahoney. “Improving Semi-supervised Federated Learning by Reducing the Gradient Diversity of Models”. In: *2021 IEEE International Conference on Big Data (Big Data)*. Dec. 2021, pp. 1214–1225. DOI: 10.1109/BigData52589.2021.9671693.
- [46] Lianmin Zheng, Ruochen Liu, Junru Shao, Tianqi Chen, Joseph E. Gonzalez, Ion Stoica, and Ameer Haj Ali. “TenSet: A Large-scale Program Performance Dataset for Learned Tensor Compilers”. In: *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*. 2021. URL: <https://openreview.net/forum?id=aIfp8kLuv9>.
- [47] Ashwin Balakrishna, Brijen Thananjeyan, Jonathan Lee, Felix Li, Arsh Zahed, Joseph E. Gonzalez, and Ken Goldberg. “On-Policy Robot Imitation Learning from a Converging Supervisor”. In: *Proceedings of the Conference on Robot Learning*. Ed. by Leslie Pack Kaelbling, Danica Kragic, and Komei Sugiura. Vol. 100. Proceedings of Machine Learning Research. PMLR, Oct. 2020, pp. 24–41. URL: <http://proceedings.mlr.press/v100/balakrishna20a.html>.
- [48] Jianfei Chen, Yu Gai, Zhewei Yao, Michael W. Mahoney, and Joseph E. Gonzalez. “A Statistical Framework for Low-bitwidth Training of Deep Neural Networks”. In: *Advances in Neural Information Processing Systems*. Ed. by H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin. Vol. 33. Curran Associates, Inc., 2020, pp. 883–894. URL: <https://proceedings.neurips.cc/paper/2020/file/099fe6b0b444c23836c4a5d07346082b-Paper.pdf>.
- [49] Daniel Crankshaw, Gur-Eyal Sela, Corey Zumar, Xiangxi Mo, Joseph E. Gonzalez, Ion Stoica, and Alexey Tumanov. “InferLine: ML Inference Pipeline Composition Framework”. In: *Proceedings of the ACM Symposium on Cloud Computing*. SoCC ’20. Association for Computing Machinery, Nov. 2020. URL: <http://arxiv.org/abs/1812.01776>.
- [50] Ankur Dave, Chester Leung, Raluca Ada Popa, Joseph E. Gonzalez, and Ion Stoica. “Oblivious Cooperative Analytics Using Hardware Enclaves”. In: *Proceedings of the Fifteenth European Conference on Computer Systems (EuroSys)*. EuroSys ’20. Heraklion, Greece: Association for Computing Machinery, 2020. ISBN: 9781450368827. URL: <https://doi.org/10.1145/3342195.3387552>.
- [51] Rolando Garcia, Eric Liu, Vikram Sreekanti, Bobby Yan, Anusha Dandamudi, Joseph E. Gonzalez, Joseph M. Hellerstein, and Koushik Sen. “Hindsight Logging for Model Training”. In: *Proc. VLDB Endow.* 14.4 (Dec. 2020), pp. 682–693. ISSN: 2150-8097. DOI: 10.14778/3436905.3436925. URL: <https://doi.org/10.14778/3436905.3436925>.
- [52] Jeffrey Ichnowski, William Lee, Victor Murta, Samuel Paradis, Ron Alterovitz, Joseph E. Gonzalez, Ion Stoica, and Ken Goldberg. “Fog Robotics Algorithms for Distributed Motion Planning Using Lambda Serverless Computing”. In: *2020 IEEE International Conference on Robotics and Automation, ICRA 2020, Paris, France, May 31 - August 31, 2020*. IEEE, 2020, pp. 4232–4238. DOI: 10.1109/ICRA40945.2020.9196651. URL: <https://doi.org/10.1109/ICRA40945.2020.9196651>.
- [53] Paras Jain, Ajay Jain, Aniruddha Nrusimha, Amir Gholami, Pieter Abbeel, Kurt Keutzer, Ion Stoica, and Joseph E. Gonzalez. “Breaking the Memory Wall with Optimal Tensor Rematerialization”. In: *Proceedings of Machine Learning and Systems*. 2020, pp. 497–511. URL: <https://arxiv.org/abs/1910.02653>.
- [54] Samvit Jain, Xun Zhang, Yuhao Zhou, Ganesh Ananthanarayanan, Junchen Jiang, Yuanhao Shu, Paramvir Bahl, and Joseph Gonzalez. “Spatula: Efficient cross-camera video analytics on large camera networks”. In: *5th IEEE/ACM Symposium on Edge Computing, SEC 2020, San Jose, CA, USA, November 12-14, 2020*. IEEE, 2020, pp. 110–124. DOI: 10.1109/SEC50012.2020.00016. URL: <https://doi.org/10.1109/SEC50012.2020.00016>.
- [55] Zhuohan Li, Eric Wallace, Sheng Shen, Kevin Lin, Kurt Keutzer, Dan Klein, and Joseph E. Gonzalez. “Train Big, Then Compress: Rethinking Model Size for Efficient Training and Inference of Transformers”. In: *Proceedings of the International Conference on Machine Learning (ICML)*. ICML’20.

- July 2020. URL: <https://proceedings.icml.cc/static/paper%5C%5Ffiles/icml/2020/6626-Paper.pdf>.
- [56] Xiayue Charles Lin, Joseph E. Gonzalez, and Joseph M. Hellerstein. “Serverless Boom or Bust? An Analysis of Economic Incentives”. In: *12th USENIX Workshop on Hot Topics in Cloud Computing (HotCloud)*. USENIX Association, July 2020. URL: <https://www.usenix.org/conference/hotcloud20/presentation/lin>.
- [57] Mong H. Ng, Kaahan Radia, Jianfei Chen, Dequan Wang, Ionel Gog, and Joseph E. Gonzalez. “BEV-Seg: Bird’s Eye View Semantic Segmentation Using Geometry and Semantic Point Cloud”. In: *Proceedings of the Workshop in Scalability for Autonomous Driving at CVPR’20*. 2020. arXiv: 2006.11436. URL: <https://arxiv.org/abs/2006.11436>.
- [58] Devin Petersohn, William W. Ma, Doris Jung Lin Lee, Stephen Macke, Doris Xin, Xiangxi Mo, Joseph E. Gonzalez, Joseph M. Hellerstein, Anthony D. Joseph, and Aditya G. Parameswaran. “Towards Scalable Dataframe Systems”. In: *Proceedings of Very Large Data Bases (PVLDB)*. Vol. 13. 2020. URL: <http://www.vldb.org/pvldb/vol13/p2033-petersohn.pdf>.
- [59] Daniel Rothchild, Ashwinee Panda, Enayat Ullah, Nikita Ivkin, Ion Stoica, Vladimir Braverman, Joseph E. Gonzalez, and Raman Arora. “FetchSGD: Communication-Efficient Federated Learning with Sketching”. In: *Proceedings of the International Conference on Machine Learning (ICML)*. ICML’20. July 2020. URL: <https://proceedings.icml.cc/static/paper%5C%5Ffiles/icml/2020/5927-Paper.pdf>.
- [60] Vidit Saxena, Joakim Jalden, and Joseph E. Gonzalez. “Thompson Sampling for Linearly Constrained Bandits”. In: *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics (AISTATS)*. Ed. by Silvia Chiappa and Roberto Calandra. Vol. 108. Proceedings of Machine Learning Research. Online: PMLR, Aug. 2020, pp. 1999–2009. URL: <http://proceedings.mlr.press/v108/saxena20a.html>.
- [61] Vikram Sreekanti, Chenggang Wu, Saurav Chhatrapati, Joseph E. Gonzalez, Joseph M. Hellerstein, and Jose M. Faleiro. “A Fault-Tolerance Shim for Serverless Computing”. In: *Proceedings of the Fifteenth European Conference on Computer Systems (EuroSys)*. EuroSys ’20. Heraklion, Greece: Association for Computing Machinery, 2020. ISBN: 9781450368827. URL: <https://doi.org/10.1145/3342195.3387535>.
- [62] Vikram Sreekanti, Chenggang Wu, Xiayue Charles Lin, Johann Schleier-Smith, Jose M. Faleiro, Joseph E. Gonzalez, Joseph M. Hellerstein, and Alexey Tumanov. “Cloudburst: Stateful Functions-as-a-Service”. In: *Proceedings of Very Large Data Bases (PVLDB)*. Vol. 13. 2020. URL: <http://www.vldb.org/pvldb/vol13/p2438-sreekanti.pdf>.
- [63] Priya Sundaresan, Jennifer Grannen, Brijen Thananjeyan, Ashwin Balakrishna, Michael Laskey, Kevin Stone, Joseph E. Gonzalez, and Ken Goldberg. “Learning Rope Manipulation Policies Using Dense Object Descriptors Trained on Synthetic Depth Data”. In: *2020 IEEE International Conference on Robotics and Automation, ICRA 2020, Paris, France, May 31 - August 31, 2020*. IEEE, 2020, pp. 9411–9418. DOI: 10.1109/ICRA40945.2020.9197121. URL: <https://doi.org/10.1109/ICRA40945.2020.9197121>.
- [64] Ajay Kumar Tanwani, Raghav Anand, Joseph E. Gonzalez, and Ken Goldberg. “RILaaS: Robot Inference and Learning as a Service”. In: *IEEE Robotics and Automation Letters* 5.3 (2020), pp. 4423–4430. URL: <https://ieeexplore.ieee.org/document/9103220>.
- [65] Brijen Thananjeyan, Ashwin Balakrishna, Ugo Rosolia, Joseph E. Gonzalez, Aaron D. Ames, and Ken Goldberg. “ABC-LMPC: Safe Sample-Based Learning MPC for Stochastic Nonlinear Dynamical Systems with Adjustable Boundary Conditions”. In: *Proceedings of the Int. Workshop on the Algorithmic Foundations of Robotics (WAFR)*. 2020. URL: <https://arxiv.org/abs/2003.01410>.
- [66] Brijen Thananjeyan, Ashwin Balakrishna, Ugo Rosolia, Felix Li, Rowan McAllister, Joseph E. Gonzalez, Sergey Levine, Francesco Borrelli, and Ken Goldberg. “Safety Augmented Value Estimation

- From Demonstrations (SAVED): Safe Deep Model-Based RL for Sparse Cost Robotic Tasks”. In: *IEEE Robotics Autom. Lett.* 5.2 (2020), pp. 3612–3619. URL: <https://arxiv.org/abs/1905.13402>.
- [67] Alvin Wan, Xiaoliang Dai, Peizhao Zhang, Zijian He, Yuandong Tian, Saining Xie, Bichen Wu, Matthew Yu, Tao Xu, Kan Chen, Peter Vajda, and Joseph E. Gonzalez. “FBNetV2: Differentiable Neural Architecture Search for Spatial and Channel Dimensions”. In: *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2020. URL: <https://arxiv.org/abs/2004.05565>.
- [68] Xin Wang, Thomas E. Huang, Trevor Darrell, Joseph E. Gonzalez, and Fisher Yu. “Frustratingly Simple Few-Shot Object Detection”. In: *Proceedings of the International Conference on Machine Learning (ICML)*. ICML’20. July 2020. URL: <https://proceedings.icml.cc/static/paper%5C%5Ffiles/icml/2020/2957-Paper.pdf>.
- [69] Yaoqing Yang, Rajiv Khanna, Yaodong Yu, Amir Gholami, Kurt Keutzer, Joseph E. Gonzalez, Kannan Ramchandran, and Michael W. Mahoney. “Boundary thickness and robustness in learning models”. In: *Advances in Neural Information Processing Systems*. Ed. by H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin. Vol. 33. Curran Associates, Inc., 2020, pp. 6223–6234. URL: <https://proceedings.neurips.cc/paper/2020/file/44e76e99b5e194377e955b13fb12f630-Paper.pdf>.
- [70] Lianmin Zheng, Chengfan Jia, Minmin Sun, Zhao Wu, Cody Hao Yu, Ameer Haj-Ali, Yida Wang, Jun Yang, Danyang Zhuo, Koushik Sen, Joseph E. Gonzalez, and Ion Stoica. “Ansor: Generating High-Performance Tensor Programs for Deep Learning”. In: *14th USENIX Symposium on Operating Systems Design and Implementation, OSDI 2020, Virtual Event, November 4-6, 2020*. USENIX Association, 2020, pp. 863–879. URL: <https://www.usenix.org/conference/osdi20/presentation/zheng>.
- [71] Joseph M. Hellerstein, Jose M. Faleiro, Joseph E. Gonzalez, Johann Schleier-Smith, Vikram Sreekanti, Alexey Tumanov, and Chenggang Wu. “Serverless Computing: One Step Forward, Two Steps Back”. In: *Conference on Innovative Data Systems Research (CIDR ’19)*. Jan. 2019. URL: <https://arxiv.org/abs/1812.03651>.
- [72] Samvit Jain, Ganesh Ananthanarayanan, Junchen Jiang, Yuanchao Shu, and Joseph E. Gonzalez. “Scaling Video Analytics Systems to Large Camera Deployments”. In: *HotMobile ’19, Proceedings of the 20th International Workshop on Mobile Computing Systems and Applications*. Feb. 2019. URL: <https://arxiv.org/abs/1809.02318>.
- [73] Samvit Jain, Xin Wang, and Joseph Gonzalez. “Accel: A Corrective Fusion Network for Efficient Semantic Segmentation on Video”. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2019. URL: <http://arxiv.org/abs/1807.06667>.
- [74] Richard Liaw, Romil Bhardwaj, Lisa Dunlap, Yitian Zou, Joseph E. Gonzalez, Ion Stoica, and Alexey Tumanov. “HyperSched: Dynamic Resource Reallocation for Model Development on a Deadline”. In: *Proceedings of the ACM Symposium on Cloud Computing*. SoCC ’19. Santa Cruz, CA, USA: Association for Computing Machinery, 2019, pp. 61–73. ISBN: 9781450369732. URL: <https://doi.org/10.1145/3357223.3362719>.
- [75] Vidit Saxena, Joakim Jaldén, Joseph E. Gonzalez, Mats Bengtsson, Hugo M. Tullberg, and Ion Stoica. “Contextual Multi-Armed Bandits for Link Adaptation in Cellular Networks”. In: *Proceedings of the Workshop on Network Meets AI (NetAI) at SIGCOMM*. 2019, pp. 44–49. URL: <https://doi.org/10.1145/3341216.3342212>.
- [76] Ajay Kumar Tanwani, Nitesh Mor, John Kubiawicz, Joseph E. Gonzalez, and Ken Goldberg. “A Fog Robotics Approach to Deep Robot Learning: Application to Object Recognition and Grasp Planning in Surface Decluttering”. In: *International Conference on Robotics and Automation, ICRA 2019, Montreal, QC, Canada, May 20-24, 2019*. 2019, pp. 4559–4566. URL: <https://doi.org/10.1109/ICRA.2019.8793690>.

- [77] Xin Wang, Fisher Yu, Lisa Dunlap, Yi-An Ma, Ruth Wang, Azalia Mirhoseini, Trevor Darrell, and Joseph E. Gonzalez. “Deep Mixture of Experts via Shallow Embedding”. In: *Proceedings of the Thirty-Fifth Conference on Uncertainty in Artificial Intelligence, UAI 2019, Tel Aviv, Israel, July 22-25, 2019*. 2019, p. 192. URL: <http://auai.org/uai2019/proceedings/papers/192.pdf>.
- [78] Xin Wang, Fisher Yu, Ruth Wang, Trevor Darrell, and Joseph E. Gonzalez. “TAFE-Net: Task-Aware Feature Embeddings for Low Shot Learning”. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2019. URL: <https://arxiv.org/abs/1904.05967>.
- [79] Zuxuan Wu, Xin Wang, Joseph E. Gonzalez, Tom Goldstein, and Larry S. Davis. “ACE: Adapting to Changing Environments for Semantic Segmentation”. In: *International Conference in Computer Vision (ICCV)*. Oct. 2019. URL: <http://arxiv.org/abs/1904.06268>.
- [80] Tianjun Zhang, Zhewei Yao, Amir Gholami, Kurt Keutzer, Joseph E. Gonzalez, George Biros, and Michael W. Mahoney. “ANODEV2: A Coupled Neural ODE Evolution Framework”. In: *Neural Information Processing Systems (NeurIPS)*. 2019. URL: <https://arxiv.org/abs/1906.04596>.
- [81] Wenting Zheng, Raluca Ada Popa, Joseph E. Gonzalez, and Ion Stoica. “Helen: Maliciously Secure Coopetitive Learning for Linear Models”. In: *IEEE Symposium on Security and Privacy (Oakland)*. IEEE Computer Society, 2019. URL: <https://people.eecs.berkeley.edu/~wzheng/helen%5C%5Fieesp.pdf>.
- [82] Rolando Garcia, Vikram Sreekanti, Neeraja Yadwadkar, Daniel Crankshaw, Joseph E. Gonzalez, and Joseph M. Hellerstein. “Context: The Missing Piece in the Machine Learning Lifecycle”. In: *Proceedings of the KDD Workshop on Common Model Infrastructure (CMI)*. Aug. 2018. URL: <http://www.vikrams.io/papers/flor-cmi18.pdf>.
- [83] Samvit Jain and Joseph E. Gonzalez. “Fast Semantic Segmentation on Video Using Block Motion-Based Feature Interpolation”. In: *The Third International Workshop on Video Segmentation (IWVS)*. Mar. 2018. URL: <https://arxiv.org/abs/1803.07742>.
- [84] Eric Liang, Richard Liaw, Robert Nishihara, Philipp Moritz, Roy Fox, Joseph Gonzalez, Ken Goldberg, and Ion Stoica. “Ray RLLib: A Composable and Scalable Reinforcement Learning Library”. In: *Proceedings of the 35th International Conference on Machine Learning. ICML '18*. ACM, July 2018. URL: <https://arxiv.org/abs/1712.09381>.
- [85] Richard Liaw, Eric Liang, Robert Nishihara, Philipp Moritz, Joseph E. Gonzalez, and Ion Stoica. “Tune: A Research Platform for Distributed Model Selection and Training”. In: *Proceedings of the ICML Workshop on AutoML*. 2018. URL: <https://arxiv.org/abs/1807.05118>.
- [86] Xiangxi Mo, Paras Jain, Ajay Jain, Alexey Tumanov, Joseph E. Gonzalez, and Ion Stoica. “A Case for Dynamic GPU Inference Multitenancy and Scheduling”. In: *Proceedings of the Learning Systems Workshop at NIPS 2018*. Dec. 2018. URL: [http://learningsys.org/nips18/assets/papers/102CameraReadySubmissionGPU%5C%5FVirtualization%5C%20\(8\).pdf](http://learningsys.org/nips18/assets/papers/102CameraReadySubmissionGPU%5C%5FVirtualization%5C%20(8).pdf).
- [87] Xin Wang, Yujia Luo, Dan Crankshaw, Alexey Tumanov, Fisher Yu, and Joseph E. Gonzalez. “IDK Cascades: Fast Deep Learning by Learning not to Overthink”. In: *Conference on Uncertainty in Artificial Intelligence (UAI)*. July 2018. URL: <https://arxiv.org/abs/1706.00885>.
- [88] Xin Wang, Fisher Yu, Zi-Yi Dou, and Joseph E. Gonzalez. “SkipNet: Learning Dynamic Routing in Convolutional Networks”. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. July 2018. URL: <https://arxiv.org/abs/1711.09485>.
- [89] Bichen Wu, Alvin Wan, Xiangyu Yue, Peter Jin, Sicheng Zhao, Noah Golmant, Amir Gholaminejad, Joseph E. Gonzalez, and Kurt Keutzer. “Shift: A Zero FLOP, Zero Parameter Alternative to Spatial Convolutions”. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2018. URL: <https://arxiv.org/abs/1711.08141>.
- [90] Francois W. Belletti, Evan R. Sparks, Michael J. Franklin, Alexandre M. Bayen, and Joseph E. Gonzalez. “Random Projection Design for Scalable Implicit Smoothing of Randomly Observed Stochastic Processes”. In: *Artificial Intelligence and Statistics (AISTATS '17)*. July 2017. URL: <http://proceedings.mlr.press/v54/belletti17a/belletti17a.pdf>.

- [91] Daniel Crankshaw, Xin Wang, Guilio Zhou, Michael J. Franklin, Joseph E. Gonzalez, and Ion Stoica. “Clipper: A Low-Latency Online Prediction Serving System”. In: *14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17)*. Boston, MA: USENIX Association, 2017, pp. 613–627. ISBN: 978-1-931971-37-9. URL: <https://www.usenix.org/conference/nsdi17/technical-sessions/presentation/crankshaw>.
- [92] Joseph M. Hellerstein, Vikram Sreekanti, Joseph E. Gonzalez, Sudhansku Arora, Arka Bhattacharyya, Shirshanka Das, Akon Dey, Mark Donsky, Gabriel Fierro, Sreyashi Nag, Krishna Ramachandran, Chang She, Eric Sun, Carl Steinbach, and Venkat Subramanian. “Establishing Common Ground with Data Context”. In: *Conference on Innovative Data Systems Research (CIDR '17)*. 2017.
- [93] Neeraja J. Yadwadkar, Bharath Hariharan, Joseph E. Gonzalez, Burton Smith, and Randy H. Katz. “Selecting the Best VM Across Multiple Public Clouds: A Data-driven Performance Modeling Approach”. In: *Proceedings of the 2017 Symposium on Cloud Computing*. SoCC '17. Santa Clara, California: ACM, Sept. 2017, pp. 452–465. ISBN: 978-1-4503-5028-0. URL: <https://doi.acm.org/10.1145/3127479.3131614>.
- [94] Wenting Zheng, Ankur Dave, Jethro G. Beekman, Raluca Ada Popa, Joseph E. Gonzalez, and Ion Stoica. “Opaque: An Oblivious and Encrypted Distributed Analytics Platform”. In: *14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17)*. Boston, MA: USENIX Association, 2017, pp. 283–298. ISBN: 978-1-931971-37-9. URL: <https://www.usenix.org/conference/nsdi17/technical-sessions/presentation/zheng>.
- [95] Ankur Dave, Alekh Jindal, Li Erran Li, Reynold Xin, Joseph E. Gonzalez, and Matei Zaharia. “GraphFrames: An Integrated API for Mixing Graph and Relational Queries.” In: *SIGMOD Grades Workshop*. 2016.
- [96] Neeraja J. Yadwadkar, Bharath Hariharan, Joseph E. Gonzalez, and Randy Katz. “Multi-Task Learning for Straggler Avoiding Predictive Job Scheduling”. In: *Journal of Machine Learning Research (JMLR '16)*. 2016.
- [97] Daniel Crankshaw, Peter Bailis, Joseph E. Gonzalez, Haoyuan Li, Zhao Zhang, Michael J. Franklin, Ali Ghodsi, and Michael I. Jordan. “The Missing Piece in Complex Analytics: Low Latency, Scalable Model Management and Serving with Velox”. In: *Conference on Innovative Data Systems Research (CIDR '15)*. 2015.
- [98] Daniel Crankshaw, Xin Wang, Joseph E. Gonzalez, and Michael J. Franklin. “Scalable Training and Serving of Personalized Models”. In: *Proceedings of the Learning Systems Workshop at NIPS 2015*. 2015.
- [99] Veronika Strnadova-Neeley, Aydin Buluc, Jarrod Chapman, John Gilbert, Joseph E. Gonzalez, and Leonid Oliker. “Efficient Data Reduction for Large-Scale Genetic Mapping”. In: *ACM Conference on Bioinformatics, Computational Biology, and Health Informatics (BCB '15)*. 2015.
- [100] Neeraja J. Yadwadkar, Bharath Hariharan, Joseph E. Gonzalez, and Randy Katz. “Faster Jobs in Distributed Data Processing using Multi-Task Learning”. In: *SIAM International Conference on Data Mining (SDM '15)*. 2015.
- [101] David Bader, Aydin Buluc, John Gilbert, Joseph E. Gonzalez, Jeremy Kepner, and Timothy Mattson. “The Graph BLAS effort and its implications for Exascale”. In: *SIAM Workshop on Exascale Applied Mathematics Challenges and Opportunities (EX14)*. 2014.
- [102] Joseph E. Gonzalez, Reynold S. Xin, Ankur Dave, Daniel Crankshaw, Michael J. Franklin, and Ion Stoica. “GraphX: Graph Processing in a Distributed Dataflow Framework”. In: *11th USENIX Symposium on Operating Systems Design and Implementation (OSDI 14)*. 2014, pp. 599–613.
- [103] Xinghao Pan, Stefanie Jegelka, Joseph E. Gonzalez, Joseph K. Bradley, and Michael I. Jordan. “Parallel Double Greedy Submodular Maximization”. In: *Neural Information Processing Systems (NIPS '14)*. 2014.

- [104] Veronika Strnadova, Aydin Buluc, Leonid Olikier, Joseph E. Gonzalez, Stefanie Jegelka, Jarrod Chapman, and John Gilbert. “Fast Clustering Methods for Genetic Mapping in Plants”. In: *16th SIAM Conference on Parallel Processing for Scientific Computing*. 2014.
- [105] T. Mattson, D. Bader, J. Berry, A. Buluc, J. Dongarra, C. Faloutsos, J. Feo, J. Gilbert, J. Gonzalez, B. Hendrickson, J. Kepner, C. Leiserson, A. Lumsdaine, D. Padua, S. Poole, S. Reinhardt, M. Stonebraker, S. Wallach, and A. Yoo. “Standards for graph algorithm primitives”. In: *2013 IEEE High Performance Extreme Computing Conference (HPEC)*. Sept. 2013, pp. 1–2. URL: <https://doi.org/10.1109/HPEC.2013.6670338>.
- [106] Xinghao Pan, Joseph E. Gonzalez, Stefanie Jegelka, Tamara Broderick, and Michael I. Jordan. “Optimistic Concurrency Control for Distributed Unsupervised Learning”. In: *NIPS ’13*. 2013. URL: <https://arxiv.org/abs/1307.8049>.
- [107] Evan Sparks, Ameet Talwalkar, Virginia Smith, Xinghao Pan, Joseph E. Gonzalez, Tim Kraska, Michael I. Jordan, and Michael J. Franklin. “MLI: An API for Distributed Machine Learning”. In: *International Conference on Data Mining (ICDM)*. IEEE. Dec. 2013. URL: <https://ieeexplore.ieee.org/abstract/document/6729619>.
- [108] Reynold Xin, Joseph E. Gonzalez, Michael Franklin, and Ion Stoica. “GraphX: A Resilient Distributed Graph System on Spark”. In: *SIGMOD Grades Workshop*. 2013. URL: <https://dl.acm.org/citation.cfm?id=2484427>.
- [109] Amr Ahmed, Mohamed Aly, Joseph Gonzalez, Shравan Narayanamurthy, and Alex Smola. “Scalable Inference in Latent Variable Models”. In: *Conference on Web Search and Data Mining (WSDM)*. 2012. URL: <http://www.cs.cmu.edu/~jegonzal/papers/ahmed%5C%5Fscalable%5C%5Finference%5C%5Fin%5C%5Flatent%5C%5Fvariable%5C%5Fmodels.pdf>.
- [110] Joseph E. Gonzalez, Yucheng Low, Haijie Gu, Danny Bickson, and Carlos Guestrin. “PowerGraph: Distributed Graph-Parallel Computation on Natural Graphs”. In: *OSDI ’12*. 2012. URL: <https://www.usenix.org/system/files/conference/osdi12/osdi12-final-167.pdf>.
- [111] Yucheng Low, Joseph E. Gonzalez, Aapo Kyrola, Danny Bickson, Carlos Guestrin, and Joseph M. Hellerstein. “Distributed GraphLab: A Framework for Machine Learning and Data Mining in the Cloud.” In: *Proceedings of Very Large Data Bases (PVLDB)*. Aug. 2012. URL: <https://arxiv.org/abs/1204.6078>.
- [112] Joseph E. Gonzalez, Yucheng Low, Arthur Gretton, and Carlos Guestrin. “Parallel Gibbs Sampling: From Colored Fields to Thin Junction Trees”. In: *Artificial Intelligence and Statistics (AISTATS)*. May 2011. URL: <http://proceedings.mlr.press/v15/gonzalez11a.html>.
- [113] Yucheng Low, Joseph E. Gonzalez, Aapo Kyrola, Daniel Bickson, Carlos Guestrin, and Joseph M. Hellerstein. “GraphLab: A New Parallel Framework for Machine Learning”. In: *Conference on Uncertainty in Artificial Intelligence (UAI)*. 2010. URL: <https://arxiv.org/abs/1006.4990>.
- [114] Joseph E. Gonzalez, Yucheng Low, and Carlos Guestrin. “Residual Splash for Optimally Parallelizing Belief Propagation”. In: *Artificial Intelligence and Statistics (AISTATS)*. Apr. 2009. URL: <http://proceedings.mlr.press/v5/gonzalez09a.html>.
- [115] Joseph E. Gonzalez, Yucheng Low, Carlos Guestrin, and David O’Hallaron. “Distributed Parallel Inference on Large Factor Graphs”. In: *Conference on Uncertainty in Artificial Intelligence (UAI)*. July 2009. URL: <https://arxiv.org/pdf/1205.2645.pdf>.

Preprint arXiv Technical Reports

- [1] Lisa Dunlap, Alyssa Umno, Han Zhang, Jiezhi Yang, Joseph E. Gonzalez, and Trevor Darrell. *Diversify Your Vision Datasets with Automatic Diffusion-Based Augmentation*. 2023. arXiv: 2305.16289 [cs.CV]. URL: <https://arxiv.org/abs/2305.16289>.

- [2] Kevin Lin, Kyle Lo, Joseph E. Gonzalez, and Dan Klein. *Decomposing Complex Queries for Tip-of-the-tongue Retrieval*. 2023. arXiv: 2305.15053 [cs.CL]. URL: <https://arxiv.org/abs/2305.15053>.
- [3] Shishir G. Patil, Tianjun Zhang, Xin Wang, and Joseph E. Gonzalez. *Gorilla: Large Language Model Connected with Massive APIs*. 2023. arXiv: 2305.15334 [cs.CL]. URL: <https://arxiv.org/abs/2305.15334>.
- [4] Suzanne Petryk, Spencer Whitehead, Joseph E. Gonzalez, Trevor Darrell, Anna Rohrbach, and Marcus Rohrbach. *Simple Token-Level Confidence Improves Caption Correctness*. 2023. arXiv: 2305.07021 [cs.CV]. URL: <https://arxiv.org/abs/2305.07021>.
- [5] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric. P Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. *Judging LLM-as-a-judge with MT-Bench and Chatbot Arena*. 2023. arXiv: 2306.05685 [cs.CL]. URL: <https://arxiv.org/abs/2204.05149>.
- [6] Sarah Chasins, Alvin Cheung, Natacha Crooks, Ali Ghodsi, Ken Goldberg, Joseph E. Gonzalez, Joseph M. Hellerstein, Michael I. Jordan, Anthony D. Joseph, Michael W. Mahoney, Aditya Parameswaran, David Patterson, Raluca Ada Popa, Koushik Sen, Scott Shenker, Dawn Song, and Ion Stoica. *The Sky Above The Clouds*. 2022. DOI: 10.48550/ARXIV.2205.07147. URL: <https://arxiv.org/abs/2205.07147>.
- [7] Jeffrey Ichnowski, Kaiyuan Chen, Karthik Dharmarajan, Simeon Adebola, Michael Danielczuk, Victor Mayoral-Vilches, Hugo Zhan, Derek Xu, Ramtin Ghassemi, John Kubiawicz, Ion Stoica, Joseph Gonzalez, and Ken Goldberg. *FogROS 2: An Adaptive and Extensible Platform for Cloud and Fog Robotics Using ROS 2*. 2022. DOI: 10.48550/ARXIV.2205.09778. URL: <https://arxiv.org/abs/2205.09778>.
- [8] Brijen Thananjeyan, Justin Kerr, Huang Huang, Joseph E. Gonzalez, and Ken Goldberg. *All You Need is LUV: Unsupervised Collection of Labeled Images using Invisible UV Fluorescent Indicators*. 2022. DOI: 10.48550/ARXIV.2203.04566. URL: <https://sites.google.com/berkeley.edu/luv>.
- [9] Spencer Whitehead, Suzanne Petryk, Vedaad Shakib, Joseph Gonzalez, Trevor Darrell, Anna Rohrbach, and Marcus Rohrbach. *Reliable Visual Question Answering: Abstain Rather Than Answer Incorrectly*. 2022. DOI: 10.48550/ARXIV.2204.13631. URL: <https://arxiv.org/abs/2204.13631>.
- [10] Yaoqing Yang, Ryan Theisen, Liam Hodgkinson, Joseph E. Gonzalez, Kannan Ramchandran, Charles H. Martin, and Michael W. Mahoney. *Evaluating natural language processing models with generalization metrics that do not need access to any training or testing data*. 2022. DOI: 10.48550/ARXIV.2202.02842. URL: <https://arxiv.org/abs/2202.02842>.
- [11] Yu Gai, Paras Jain, Wendi Zhang, Joseph E. Gonzalez, Dawn Song, and Ion Stoica. *Grounded Graph Decoding Improves Compositional Generalization in Question Answering*. 2021. DOI: 10.48550/ARXIV.2111.03642. URL: <https://arxiv.org/abs/2111.03642>.
- [12] Wenshuo Guo, Kirthevasan Kandasamy, Joseph E Gonzalez, Michael I. Jordan, and Ion Stoica. *Learning Competitive Equilibria in Exchange Economies with Bandit Feedback*. 2021. DOI: 10.48550/ARXIV.2106.06616. URL: <https://arxiv.org/abs/2106.06616>.
- [13] Alan Pham, Eunice Chan, Vikranth Srivatsa, Dhruva Ghosh, Yaoqing Yang, Yaodong Yu, Ruiqi Zhong, Joseph E. Gonzalez, and Jacob Steinhardt. *The Effect of Model Size on Worst-Group Generalization*. 2021. DOI: 10.48550/ARXIV.2112.04094. URL: <https://arxiv.org/abs/2112.04094>.
- [14] Nicholas Rhinehart, Jeff He, Charles Packer, Matthew A. Wright, Rowan McAllister, Joseph E. Gonzalez, and Sergey Levine. *Contingencies from Observations: Tractable Contingency Planning with Learned Behavior Models*. 2021. DOI: 10.48550/ARXIV.2104.10558. URL: <https://arxiv.org/abs/2104.10558>.
- [15] Daniel Rothchild, Alex Tamkin, Julie Yu, Ujval Misra, and Joseph Gonzalez. *C5T5: Controllable Generation of Organic Molecules with Transformers*. 2021. DOI: 10.48550/ARXIV.2108.10307. URL: <https://arxiv.org/abs/2108.10307>.

- [16] Priya Sundaresan, Jennifer Grannen, Brijen Thananjeyan, Ashwin Balakrishna, Jeffrey Ichnowski, Ellen Novoseller, Minh Hwang, Michael Laskey, Joseph E. Gonzalez, and Ken Goldberg. *Untangling Dense Non-Planar Knots by Learning Manipulation Features and Recovery Policies*. 2021. DOI: 10.48550/ARXIV.2107.08942. URL: <https://arxiv.org/abs/2107.08942>.
- [17] Brijen Thananjeyan, Kirthevasan Kandasamy, Ion Stoica, Michael I. Jordan, Ken Goldberg, and Joseph E. Gonzalez. *PAC Best Arm Identification Under a Deadline*. 2021. DOI: 10.48550/ARXIV.2106.03221. URL: <https://arxiv.org/abs/2106.03221>.
- [18] Matthew A. Wright and Joseph E. Gonzalez. *Transformers are Deep Infinite-Dimensional Non-Mercer Binary Kernel Machines*. 2021. DOI: 10.48550/ARXIV.2106.01506. URL: <https://arxiv.org/abs/2106.01506>.
- [19] Doris Xin, Devin Petersohn, Dixin Tang, Yifan Wu, Joseph E. Gonzalez, Joseph M. Hellerstein, Anthony D. Joseph, and Aditya G. Parameswaran. *Enhancing the Interactivity of Dataframe Queries by Leveraging Think Time*. 2021. DOI: 10.48550/ARXIV.2103.02145. URL: <https://arxiv.org/abs/2103.02145>.
- [20] Tianjun Zhang, Benjamin Eysenbach, Ruslan Salakhutdinov, Sergey Levine, and Joseph E. Gonzalez. *C-Planning: An Automatic Curriculum for Learning Goal-Reaching Tasks*. 2021. DOI: 10.48550/ARXIV.2110.12080. URL: <https://arxiv.org/abs/2110.12080>.
- [21] Xiaoliang Dai, Alvin Wan, Peizhao Zhang, Bichen Wu, Zijian He, Zhen Wei, Kan Chen, Yuandong Tian, Matthew Yu, Peter Vajda, and Joseph E. Gonzalez. “FBNetV3: Joint Architecture-Recipe Search using Neural Acquisition Function”. In: *CoRR* abs/2006.02049 (2020). arXiv: 2006.02049. URL: <https://arxiv.org/abs/2006.02049>.
- [22] Paras Jain, Ajay Jain, Tianjun Zhang, Pieter Abbeel, Joseph E. Gonzalez, and Ion Stoica. “Contrastive Code Representation Learning”. In: *CoRR* (2020). arXiv: 2007.04973 [cs.LG]. URL: <https://arxiv.org/abs/2007.04973>.
- [23] Kirthevasan Kandasamy, Joseph E. Gonzalez, Michael I. Jordan, and Ion Stoica. “Mechanism Design with Bandit Feedback”. In: *CoRR* abs/2004.08924 (2020). arXiv: 2004.08924. URL: <https://arxiv.org/abs/2004.08924>.
- [24] Vikram Sreekanti, Harikaran Subbaraj, Chenggang Wu, Joseph E. Gonzalez, and Joseph M. Hellerstein. “Optimizing Prediction Serving on Low-Latency Serverless Dataflow”. In: *CoRR* abs/2007.05832 (2020). arXiv: 2007.05832. URL: <https://arxiv.org/abs/2007.05832>.
- [25] Alvin Wan, Daniel Ho, Younjin Song, Henk Tillman, Sarah Adel Bargal, and Joseph E. Gonzalez. “SegNBDT: Visual Decision Rules for Segmentation”. In: *CoRR* (2020). arXiv: 2006.06868 [cs.CV]. URL: <https://arxiv.org/abs/2006.06868>.
- [26] Bohan Zhai, Tianren Gao, Flora Xue, Daniel Rothchild, Bichen Wu, Joseph E. Gonzalez, and Kurt Keutzer. “SqueezeWave: Extremely Lightweight Vocoders for On-device Speech Synthesis”. In: *CoRR* abs/2001.05685 (2020). arXiv: 2001.05685. URL: <https://arxiv.org/abs/2001.05685>.
- [27] Tianjun Zhang, Huazhe Xu, Xiaolong Wang, Yi Wu, Kurt Keutzer, Joseph E. Gonzalez, and Yuandong Tian. *BeBold: Exploration Beyond the Boundary of Explored Regions*. 2020. DOI: 10.48550/ARXIV.2012.08621. URL: <https://arxiv.org/abs/2012.08621>.
- [28] Xin Wang, Fisher Yu, Trevor Darrell, and Joseph E. Gonzalez. “Task-Aware Feature Generation for Zero-Shot Compositional Learning”. In: *CoRR* (2019). arXiv: 1906.04854 [cs.CV]. URL: <https://arxiv.org/abs/1906.04854>.
- [29] Vladimir Feinberg, Alvin Wan, Ion Stoica, Michael I. Jordan, Joseph E. Gonzalez, and Sergey Levine. “Model-Based Value Estimation for Efficient Model-Free Reinforcement Learning”. In: *CoRR* abs/1803.00101 (Feb. 2018). arXiv: 1803.00101. URL: <http://arxiv.org/abs/1803.00101>.
- [30] Noah Golmant, Nikita Vemuri, Zhewei Yao, Vladimir Feinberg, Amir Gholami, Kai Rothauge, Michael W. Mahoney, and Joseph Gonzalez. “On the Computational Inefficiency of Large Batch Sizes for Stochastic Gradient Descent”. In: *CoRR* abs/1811.12941 (Nov. 2018). arXiv: 1811.12941. URL: <http://arxiv.org/abs/1811.12941>.

- [31] J. Weston Hughes, Taylor Sittler, Anthony D. Joseph, Jeffrey E. Olgin, Joseph E. Gonzalez, and Geoffrey H. Tison. “Using Multitask Learning to Improve 12-Lead Electrocardiogram Classification”. In: *CoRR* abs/1812.00497 (Dec. 2018). arXiv: 1812.00497. URL: <http://arxiv.org/abs/1812.00497>.
- [32] Sicheng Zhao, Bichen Wu, Joseph Gonzalez, Sanjit A. Seshia, and Kurt Keutzer. “Unsupervised Domain Adaptation: from Simulation Engine to the RealWorld”. In: *CoRR* abs/1803.09180 (Mar. 2018). arXiv: 1803.09180. URL: <http://arxiv.org/abs/1803.09180>.
- [33] Richard Liaw, Sanjay Krishnan, Animesh Garg, Daniel Crankshaw, Joseph E. Gonzalez, and Ken Goldberg. “Composing Meta-Policies for Autonomous Driving Using Hierarchical Deep Reinforcement Learning”. In: *CoRR* abs/1711.01503 (Nov. 2017). arXiv: 1711.01503. URL: <http://arxiv.org/abs/1711.01503>.
- [34] Francois W. Belletti, Evan R. Sparks, Michael J. Franklin, Alexandre M. Bayen, and Joseph E. Gonzalez. “Scalable Linear Causal Inference for Irregularly Sampled Time Series with Long Range Dependencies”. In: *CoRR* abs/1603.03336 (2016). arXiv: 1603.03336. URL: <http://arxiv.org/abs/1603.03336>.
- [35] Joseph E. Gonzalez, Peter Bailis, Michael I. Jordan, Michael J. Franklin, Joseph M. Hellerstein, Ali Ghodsi, and Ion Stoica. “Asynchronous Complex Analytics in a Distributed Dataflow Architecture”. In: *CoRR* abs/1510.07092 (2015). arXiv: 1510.07092. URL: <http://arxiv.org/abs/1510.07092>.

Other Technical Reports and Invited Publications

- [1] David Patterson, Joseph Gonzalez, Urs Hölzle, Quoc Le, Chen Liang, Lluís-Miquel Munguia, Daniel Rothchild, David R. So, Maud Texier, and Jeff Dean. “The Carbon Footprint of Machine Learning Training Will Plateau, Then Shrink”. In: *Computer* 55.7 (2022), pp. 18–28. DOI: 10.1109/MC.2022.3148714. URL: <https://arxiv.org/abs/2204.05149>.
- [2] David E. Culler, Prabal Dutta, Gabe Fierro, Joseph E. Gonzalez, Nathan Pemberton, Johann Schleier-Smith, Kalyanaraman Shankari, Alvin Wan, and Thomas Zachariah. “CoVista: A Unified View on Privacy Sensitive Mobile Contact Tracing”. In: *IEEE Data Eng. Bull.* 43.2 (2020), pp. 83–94. URL: <http://sites.computer.org/debull/A20june/p83.pdf>.
- [3] Eric Jonas, Johann Schleier-Smith, Vikram Sreekanti, Chia-Che Tsai, Anurag Khandelwal, Qifan Pu, Vaishaal Shankar, Joao Menezes Carreira, Karl Krauth, Neeraja Yadwadkar, Joseph E. Gonzalez, Raluca Ada Popa, Ion Stoica, and David A. Patterson. *Cloud Programming Simplified: A Berkeley View on Serverless Computing*. Tech. rep. UCB/EECS-2019-3. EECS Department, University of California, Berkeley, Feb. 2019. URL: <http://www2.eecs.berkeley.edu/Pubs/TechRpts/2019/EECS-2019-3.html>.
- [4] Dan Crankshaw, Joseph E. Gonzalez, and Peter Bailis. “Research for Practice: Prediction-Serving Systems”. In: *Commun. ACM* 61.8 (July 2018), pp. 45–49. ISSN: 0001-0782. URL: <http://doi.acm.org/10.1145/3190574>.
- [5] Ion Stoica, Dawn Song, Raluca Ada Popa, David A. Patterson, Michael W. Mahoney, Randy H. Katz, Anthony D. Joseph, Michael Jordan, Joseph M. Hellerstein, Joseph E. Gonzalez, Ken Goldberg, Ali Ghodsi, David E. Culler, and Pieter Abbeel. *A Berkeley View of Systems Challenges for AI*. Tech. rep. UCB/EECS-2017-159. EECS Department, University of California, Berkeley, Sept. 2017. URL: <http://www2.eecs.berkeley.edu/Pubs/TechRpts/2017/EECS-2017-159.html>.
- [6] Rong Gu, Qianhao Dong, Haoyuan Li, Joseph E. Gonzalez, Zhao Zhang, Shuai Wang, Yihua Huang, Scott Shenker, Ion Stoica, and Patrick P. C. Lee. *DFS-Perf: A Scalable and Unified Benchmarking Framework for Distributed File Systems*. Tech. rep. UCB/EECS-2016-133. EECS Department, University of California, Berkeley, July 2016. URL: <http://www2.eecs.berkeley.edu/Pubs/TechRpts/2016/EECS-2016-133.html>.

- [7] Matei Zaharia, Reynold S. Xin, Patrick Wendell, Tathagata Das, Michael Armbrust, Ankur Dave, Xiangrui Meng, Josh Rosen, Shivaram Venkataraman, Michael J. Franklin, Ali Ghodsi, Joseph E. Gonzalez, Scott Shenker, and Ion Stoica. “Apache Spark: A Unified Engine for Big Data Processing”. In: *Commun. ACM* 59.11 (Sept. 2016), pp. 56–65. ISSN: 0001-0782. URL: <https://doi.acm.org/10.1145/2934664>.
- [8] Joseph E. Gonzalez. “From Graphs to Tables the Design of Scalable Systems for Graph Analytics”. In: *Proceedings of the 23rd International Conference on World Wide Web. WWW '14 Companion*. Seoul, Korea: ACM, 2014, pp. 1149–1150. ISBN: 978-1-4503-2745-9. URL: <https://doi.acm.org/10.1145/2567948.2580059>.

Journal Editor

- [2020] **Editor of the IEEE Data Engineering Bulletin Special Issue: Data Technologies Behind Digital Contact Tracing**
- [2018] **Editor of the IEEE Data Engineering Bulletin Special Issue: Machine Learning Life-cycle Management**

Selected Invited Talks

- [2022] **Invited Speaker at the Cloud Intelligence and AIOps Workshop** *The Evolution of Model Deployment from the Cloud to the Sky*
- [2022] **Invited Speaker at the China Database Committee Conference** *The Co-Evolution of Data and ML Systems from the Cloud to the Sky*
- [2021] **UC Berkeley Office of the CTO Cloud Computing Meetup** *To the Cloud and Beyond A Brief View of Cloud Computing Research at UC Berkeley*
- [2020] **Speaker in the Colloquium Lecture Series hosted by the Department of Computer and Information Science at the University of Pennsylvania** *Advances in Machine Learning Systems Research*
- [2020] **Keynote Speaker at the 14th Annual Machine Learning Symposium at the New York Academy of Sciences** *Advances in AI-Systems Research*
- [2020] **Interviewed on This Week in Machine Learning Podcast**
- [2020] **Keynote at the International Solid-State Circuits Conference Forum on Machine Learning Processors** *Hardware Opportunities in the Machine Learning Lifecycle*
- [2020] **Keynote at the Workshop on MLOps (co-located and MLSys'20):** *Advances in Prediction Serving Systems*
- [2019] **Spoke in a Panel for Hispanic Engineers and Scientists (HES):** *Discussed career planning and my research trajectory to undergrads at UC Berkeley.*
- [2018] **Data Eng Conference Keynote** *Managing the Machine Learning Lifecycle*
- [2018] **ISAT TAMALE Workshop Presentation:** *This invitation only workshop covered the Tools and Algorithms for Machine learning. Gonzalez presented his work on prediction serving and experiment management systems.*
- [2017] **SOSP Workshop Keynote** *Research at the Intersection of AI and Systems*

- [2017] **AAAI Systems Workshop Keynote** *Rise to the Challenges of AI Systems*
- [2016] **ICML Workshop Keynote** *Prediction Serving: What happens after learning?*
- [2015] **NIPS Workshop Keynote** *Intelligent Services: Serving Machine Learning Predictions*
- [2015] **ODSC** *Intelligent Services: Serving Machine Learning Predictions*
- [2015] **CIDR Talk** *The Missing Piece in Complex Analytics: Low Latency, Scalable Model Management and Serving with Velox*
- [2014] **OSDI Conference** *GraphX: Graph Processing in a Distributed Dataflow Framework*
- [2014] **Annual meeting of the International Society for Bayesian Analysis (ISBA)** *Concurrency Control For Scalable Bayesian Inference.*
- [2014] **Tutorial at the International Conference for Machine Learning (ICML)** *Emerging Systems for Large-Scale Machine Learning.*
- [2014] **Session on Graph Algorithms Building Blocks at the International Parallel and Distributed Processing Systems (IPDPS)** *GraphX: Unifying Table and Graph Analytics.*
- [2014] **NetApp ATG University Day** *Large Scale Graph Analytics: Applications and Systems*
- [2014] **Keynote Speaker: Workshop on Big Graph Mining at the International World Wide Web Conference (WWW)** *From Graphs to Tables: The Design of Scalable Systems for Graph Analytics.*
- [2013] **SIAM CSE'13 Minisymposium Frontiers in Large-Scale Graph Analysis** *Large-Scale Graph-Structured Machine Learning: GraphLab in the Cloud and GraphChi in your PC*
- [2012] **OSDI Conference** *PowerGraph: Distributed Graph-Parallel Computation on Natural Graphs*
- [2011] **IDGA Data Center Consolidation Summit** *GraphLab: A New Parallel Framework for Machine Learning*
- [2010] **DARPA Future Ideas Symposium** *Invited speaker at the DARPA future ideas symposium.*
- [2009] **UAI Conference.** *Distributed Parallel Inference on Large Factor Graphs*

Advising

Postdocs

1. Jianfei Chen [2019-2021] *Tsinghua University*
2. Ionel Gog [2018-2021] *Google Brain*
3. Kirthevasan Kandasamy [2019-2021] *UW-Madison*
4. Alexey Tumanov [2016-2019] *Georgia Tech.*
5. Matthew Wright [2019-2021]
6. Yaoqing Yang [2019-2022] *Dartmouth*

PhD Candidates

1. Francois Belletti (PhD) [2015-2018]: **Thesis:** *Alternate Representations for Scalable Analysis and Control of Heterogeneous Time Series*
2. Daniel Crankshaw (PhD) [2015-2019]: **Thesis:** *The Design and Implementation of Low-Latency Prediction Serving Systems*
3. Lisa Dunlap (PhD) [2021-Present]: *Studying how to use language to correct biases in images.*
4. Vlad Feinberg [2017-2018]: *Studying model based variance reduction methods for reinforcement learning.*
5. Yu Gai [2019-Present]: *Studying graph neural networks for question answering.*
6. Paras Jain [2018-Present]: *Studying machine learning techniques applied to program analysis.*
7. Richard Liaw[2017-2020]: *Studying the design of systems for hyperparameter search.*
8. Charles Lin [2019-2020]: *Studying economics of serverless computing.*
9. Kevin Lin [2019-Present]: *Studying efficient training of NLP models.*
10. Charles Packer [2019-Present]: *Studying meta-reinforcement learning.*
11. Shishir Patil [2019-Present]: *Studying machine learning on edge devices.*
12. Nathan Pemberton [2020-2022]: **Thesis:** *The Serverless Datacenter: Hardware and Software Techniques for Resource Disaggregation*
13. Devin Petersohn [2015-2021]: **Thesis:** *Dataframe Systems: Theory, Architecture, and Implementation*
14. Suzanne Petryk [2019-Present]: *Studying explainable computer vision for few shot learning.*
15. Daniel Rothchild [2018-Present]: *Studying federated learning and model design for chemistry.*
16. Peter Schafhalter [2019-Present]: *Studying the design of software systems for autonomous vehicles.*
17. Eyal Sela [2018-Present]: *Studying the design of prediction serving systems.*
18. Vikram Sreekanti [2016-2020]: **Thesis:** *The Design of Stateful Serverless Infrastructure*
19. Brijen Thananjeyan [2018-2022]: **Thesis:** *Safe Reinforcement Learning Using Learned Safe Sets*
20. Alvin Wan [2018-Present]: **Thesis:** *Efficiently Designing Efficient Deep Neural Networks*
21. Xin Wang (PhD) [2016-2020]: **Thesis:** *The Design of Dynamic Neural Networks for Efficient Learning and Inference*
22. Justin Wang [2020-present]: *Learning based approaches to accelerate combinatorial optimization problems.*
23. Sarah Wooders [2020-present]: *Studying the design of feature stores in machine learning.*
24. Neeraja Yadwadkar (PhD) [2015-2018]: **Thesis:** *Machine Learning for Automatic Resource Management in the Datacenter and the Cloud*
25. Tianjun Zhang [2018-Present]: *Studying multi-agent reinforcement learning.*
26. Lianmin Zheng [2019-Present]: *Studying tensor program optimization.*

Masters Students

1. Raghav Anand [2019-2020]
2. Aman Dhar [2019-2020]
3. Edward Fang [2019-2020]
4. Daniel Ho [2019-2020]
5. Samvit Jain [2017-2019]
6. Sona Jeswani [2019-2020]
7. Alvin Kao [2019-2020]
8. Alexander Ku [2019-2020]
9. Scott Lee [2019-2020]
10. Eric Liu [2019-2020]
11. Ashwinee Panda [2019-2020]
12. Hari Subbaraj [2019-2020]
13. Victor Sun [2019-2020]
14. Nikita Vemuri [2019-2020]
15. Flora Xue [2019-2020]
16. Andy Yan [2019-2020]
17. Corey Zumar [2017-2018]
18. Yika Yujia Luo [2018-2019]

Undergraduate Research Assistants

1. Daiyaan Arfeen
2. Avinash Arjavalingham
3. Michael Chau
4. Saurav Chhatrapati
5. Harrisson Constantinio
6. Anusha Dandamudi
7. Ziyi Dou
8. Rehan Durrani
9. Tianren Gao
10. Noah Golmant
11. Jennifer Grannen
12. Yijie Huang
13. Weston Hughes
14. Yuxin Ji
15. Henry Jin
16. Sun Ah Lee
17. Cathy Li
18. Michael Luo
19. William Ma
20. Sean Meng
21. Simon Mo
22. Mong Ng
23. Aniruddha Nrusimha
24. Aditya Ramkumar
25. Allen Shen
26. Younjin Song
27. Ishaan Srivasta
28. Priya Sundaresan
29. Andrew Tan
30. Julie Yu
31. Anthony Yu
32. Bohan Zhai
33. Ziyun Zhao
34. Kevin Zhu
35. Tyler Zhu
36. Yitian Zou

Teaching

- **Principles and Techniques of Data Science [Data100]:** Joseph developed and teaches the large upper division data science class at UC Berkeley. This class covers topics ranging from data cleaning and visualization to machine learning and hypothesis testing with a focus on real-world data and problems. Students from across campus take this class which routinely has more than 1000 students a semester. For more information visit <http://ds100.org>.
- **Conversations with Thought Leaders in Technology (CS198-100):** Joseph developed and co-hosted this class with Vik Singh external industry leader. This interactive zoom seminar class created in the middle of the COVID-19 pandemic leveraged the virtual format to invite world class technology leaders to speak with students about their careers and the lessons they learned. In this class, we heavily leveraged slack and zoom to have student write and vote on questions that were then used to drive the interactive interview. Speakers in this series included Henry Schuck, Sanjay Poonen, Mamoon Hamid, Archana Agrawal, Hilary Mason, and Nick Caldwell.
- **Machine Learning Systems [CS194-162]:** Joseph developed and teaches the machine learning systems graduate seminar. This graduate seminar covers key topics in machine learning systems research and combines paper reading with structured paper discussions and class projects. For more information visit <https://ucbrise.github.io/cs294-ai-sys-sp22/>.
- **Introduction to Databases [CS186]:** Joseph co-taught this 500+ student class with Joseph Hellerstein and helped to update content and create a series of new lectures covering concepts at the intersection of databases and data science.

Workshops

- [2019] **ACM-IMS Interdisciplinary Summit on the Foundations of Data Science:** I co-organized this event and moderated the panel: Deep Learning, Reinforcement Learning, and Role of Methods in Data Science.
- [2019] **SOSP ML Systems Workshop** I co-organized the SOSP learning systems workshop. (<http://learningsys.org/sosp19/index.html>)
- [2019] **NeurIPS ML Systems Workshop** I co-organized the NeurIPS learning systems workshop. (<http://learningsys.org/neurips19/>)
- [2019] **Workshop on Human In the Loop Learning (HILL)** I co-organized the first workshop on Human in the Loop Learning co-located with ICML
- [2018] **NIPS ML Systems Workshop:** I co-organized the NIPS learning systems workshop. (<http://learningsys.org/nips18/>)
- [2018] **NextProf Nexus Workshop** I co-organized this three day workshop intended to strengthen and diversify the next generation of academic leaders in engineering. (<http://nextprof2018.engin.umich.edu/nexus/>)
- [2017] **SOSP ML Systems Workshop** I launched and co-organized the first SOSP learning systems workshop. (<http://learningsys.org/sosp17/>)
- [2017] **NIPS ML Systems Workshop** I co-organized the NIPS learning systems workshop. (<http://learningsys.org/nips17/>)
- [2016] **ICML ML Systems Workshop** I co-organized the ICML learning systems workshop. (<https://sites.google.com/site/mlsys2016/>)
- [2015] **NIPS Learning Systems Workshop** I help co-organize the NIPS learning systems workshop. (<http://learningsys.org>)

- **[2014] DIMACS Workshop Organizer** I organized the DIMACS workshop on the Systems and Analytics of Big Data (<http://dimacs.rutgers.edu/Workshops/Analytics/>)
- **[2013] NIPS Workshop Organizer** I helped organize the third annual NIPS “Big Learning: Algorithms, Systems, and Tools” workshop. (<http://biglearn.org>)
- **[2012] NIPS Workshop Organizer** I helped organize the second annual NIPS “Big Learning: Algorithms, Systems, and Tools” workshop. (<http://biglearn.org>)
- **[2011] NIPS Workshop Organizer** I organized and led the workshop entitled “Big Learning: Algorithms, Systems, and Tools for Learning at Scale” For more information visit the workshop website <http://biglearn.org>
- **[2009] NIPS Workshop Organizer** I organized and led the first NIPS BigLearn workshop entitled “Large-Scale Machine Learning: Parallelism and Massive Datasets.” For more information visit the workshop website <http://www.select.cs.cmu.edu/meetings/biglearn09>

Funding Activities

- **[2019]** NSF National Science Foundation, CAREER: Towards the Design of Models and Systems for Efficient Prediction. UCB Award ID: 047303-001. PI: Joseph Gonzalez (UC Berkeley). Total: \$450,648.
- **[2018]** NSF National Science Foundation, ECDI: Secure Fog Robotics Using the Global Data Plane. UCB Award ID: 045561-001. Co-PI: Joseph Gonzalez (UC Berkeley). Total: \$2,000,000.
- **[2018]** NSF National Science Foundation, Expedition Award: Secure, Real-Time Decisions on Live Data. UCB Award ID: 044393-001. Co-PI: Joseph Gonzalez (UC Berkeley). Total: \$9,999,999. Project length: 5 years.
- **[2017]** DOD Advanced Research Projects Agency, Hierarchical Learning-Guided Automatic Model Discovery. UCB Award ID: 042431-001. Co-PI: Joseph Gonzalez (UC Berkeley). Total: \$3,811,944. Project length: 4 years.
- **[2017-present]** Raised gift funding for RISELab from Alibaba, Amazon, Ant Financial ARM, Ant, Capital One, Ericsson, Facebook, Futurewei, GE, Google, Hauwei, IBM, Intel, Microsoft, Nvidia, Scotiabank, Splunk, and VMware. Roughly \$5,000,000 annually.
- **[2013]** Raised 6.75M in series A funding for the GraphLab Inc. startup.
- **[2010]** Applied for and was awarded a grant to have early access to the Intel Single-chip Cloud Computer (SCC) as part of the Many-core Applications Research Community.
- **[2008 - 2009] Helped Lead a DARPA Interdisciplinary Sciences and Technology Study (ISAT) Group** to investigate the future of parallel machine learning from an interdisciplinary perspective I also participated in the final Woodshole annual ISAT meeting to prepare a proposal for the DARPA director.
- **[2008]** Applied for and was awarded funding for (BAA 08-34) “Machine Learning and AI in the context of Multicore and Cluster Computing.”

Reviewing

- | | |
|---|--|
| • [2021] OSDI Program Committee | • [2019] SOSP Program Committee |
| • [2021] SysML Program Committee | • [2019] MLSys (formerly SysML) Program Committee |
| • [2020] SysML Conference Workshop Committee | • [2018] ICML Area Chair |
| • [2020] EuroSys Program Committee | • [2018] LXAI Workshop Program Com- |

- mittee
- [2018] MLSys (formerly SysML) Program Committee
- [2018] Sigmod Program Committee
- [2017] AISTats Senior Program Committee
- [2016] VLDB PC Member (Best Reviewer Award)
- [2017] EuroSys PC Member
- [2016] ICML Reviewer
- [2016] AISTATS Reviewer
- [2014] OSDI External Review Committee
- [2014] HotCloud Review Committee
- [2014] ICML Reviewer
- [2013] Transactions on Pattern Analysis and Machine Intelligence
- [2013] NIPS Reviewer
- [2013] Super Computing Reviewer
- [2012] Parallel Computing Reviewer
- [2010] ICML Reviewer
- [2009] JMLR Reviewer
- [2007] JMLR Reviewer
- [2007] IPSN Reviewer

Commercial Involvement

- **Co-founder of Aqueduct Inc. (2020-Present):** I co-founded this company with former students and I am head of product.
- **Advisory Board for AnyScale Inc. (2019-Present):** I am an advisory for AnyScale.
- **Advisory Board for DeepScale Inc. (2017-2019):** I was an advisor for DeepScale and was involved in technical discussions around vision and systems for autonomous driving.
- **Turi Inc. (2013-2016):** I co-founded Turi Inc. (formerly GraphLab Inc.) to commercialize my research and help lead the initial technology roadmap. When Turi was acquired by Apple Inc. in July of 2016 it had grown to over 50 employees and had raised nearly 25M in venture capital.
- **Yahoo! Research (2011):** Developed the next generation of the GraphLab abstraction to enable large-scale machine learning on natural graphs derived from social media and web-content. (*Alex Smola: smola@yahoo-inc.com*)
- **AT&T Labs Research (2007):** Developed models for statistically assessing DSL quality from limited noisy data. (*Steven Phillips: phillips@research.att.com*)
- **Intern at ADAPT (2006)** Worked on an automated AdWords auction agent. I developed and implemented models for assessing word value. (*Alex Bäcker: alex@caltech.edu*)
- **Microsoft Developer Internship (2005):** Worked with MSN Search team developing techniques to use behavioral information to identify search spam. (*Greg Hullender: greghull@windows.microsoft.com*)
- **Caltech Research Fellowship (2004):** Developed a new query-less search technology that uses prior reading interests to identify novel documents. (*Alex Bäcker: alex@caltech.edu*)
- **NASA Jet Propulsion Labs Fellowship (2003):** Developed a new algorithm for efficiently evaluating line-of-sight on digital elevation maps at JPL. (*Robert Chamberlain: rgc@jpl.nasa.gov*)

Publicly Released Software

- **Alpa:** A compiler for distributed deep learning. <https://github.com/alpa-projects/alpa>
- **SkyPlane:** Fast bulk data transfers across regions and clouds. <https://skyplane.org>
- **RLLib:** A distributed reinforcement learning system. <https://docs.ray.io/en/master/rllib.html>
- **Tune:** A system for distributed hyperparameter search and model development. <https://docs.ray.io/en/latest/tune.html>
- **Modin:** A system for scalable execution of Pandas dataframes. <https://modin.readthedocs.io/>

- **Clipper Prediction Serving System** A system for serving predictions across a wide range of machine learning frameworks. <https://clipper.ai>
- **Opaque:** Secure analytics on top of Apache Spark. <https://github.com/mc2-project/opaque>
- **GraphX:** A graph computation framework built into the widely adopted Apache Spark open-source project. <https://spark.apache.org/graphx/>
- **GraphLab/PowerGraph** A sophisticated API for building parallel and distributed machine learning algorithms on top of multicore and cloud architectures. GraphLab generalizes the MapReduce abstraction to support iterative asynchronous computation on graph structured dependent data. <http://graphlab.org>
- **Distributed SplashBP** This library implements the SplashBP algorithm for factor graph inference in the distributed setting using MPI. http://www.select.cs.cmu.edu/code/mpi_splash.tar.gz