

Biographic Summary

Joseph Gonzalez is an assistant professor at UC Berkeley and a co-director of the UC Berkeley RISE Lab. His research addresses the design of algorithms, abstractions, and systems for scalable machine learning and data science. Joseph holds a PhD in Machine Learning from CMU where he created the GraphLab open-source graph processing system and a collection of algorithms for graphical model inference. Building on his Thesis work, Joseph co-founded of Turi Inc. (formerly GraphLab Inc.) which was acquired in July 2016 by Apple Inc.

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

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

- **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.
- **Presidential Award (2002-2006):** I was awarded tuition for research and academic achievements.

Peer Reviewed Publications

- [1] 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>.
- [2] 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 ’20. Heraklion, Greece: Association for Computing Machinery, 2020. ISBN: 9781450368827. DOI: 10.1145/3342195.3387552. URL: <https://doi.org/10.1145/3342195.3387552>.
- [3] 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*. 2020, pp. 497–511. URL: <https://arxiv.org/abs/1910.02653>.
- [4] 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_files/icml/2020/6626-Paper.pdf.
- [5] 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>.
- [6] 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_files/icml/2020/5927-Paper.pdf.
- [7] 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*. 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>.
- [8] 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 ’20. Heraklion, Greece: Association for Computing Machinery, 2020. ISBN: 9781450368827. DOI: 10.1145/3342195.3387535. URL: <https://doi.org/10.1145/3342195.3387535>.
- [9] 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>.
- [10] 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>.
- [11] 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. DOI: 10.1109/LRA.2020.2976272. URL: <https://arxiv.org/abs/1905.13402>.
- [12] 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 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2020. URL: <https://arxiv.org/abs/2004.05565>.
- [13] 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_files/icml/2020/2957-Paper.pdf.
- [14] 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>.
- [15] 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>.
- [16] 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>.
- [17] 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. DOI: 10.1145/3357223.3362719. URL: <https://doi.org/10.1145/3357223.3362719>.
- [18] 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 2019 Workshop on Network Meets AI & ML, NetAI at SIGCOMM 2019, Beijing, China, August 23, 2019*. 2019, pp. 44–49. DOI: 10.1145/3341216.3342212. URL: <https://doi.org/10.1145/3341216.3342212>.
- [19] 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. DOI: 10.1109/ICRA.2019.8793690. URL: <https://doi.org/10.1109/ICRA.2019.8793690>.
- [20] 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>.
- [21] 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>.
- [22] 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>.
- [23] *ANODEV2: A Coupled Neural ODE Evolution Framework*. 2019. URL: <https://arxiv.org/abs/1906.04596>.

- [24] *Helen: Maliciously Secure Cooperative Learning for Linear Models*. IEEE Computer Society, 2019. URL: https://people.eecs.berkeley.edu/~wzheng/helen_ieeesp.pdf.
- [25] 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>.
- [26] 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>.
- [27] 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>.
- [28] 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>.
- [29] 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_Virtualization%20\(8\).pdf](http://learningsys.org/nips18/assets/papers/102CameraReadySubmissionGPU_Virtualization%20(8).pdf).
- [30] 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>.
- [31] 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>.
- [32] 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>.
- [33] 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>.
- [34] 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>.
- [35] 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.
- [36] 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. DOI: 10.1145/3127479.3131614. URL: <http://doi.acm.org/10.1145/3127479.3131614>.

- [37] 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>.
- [38] Ankur Dave, Alekh Jindal, Li Erran Li, Reynold Xin, Joseph E. Gonzalez, and Matei Zaharia. “Graph-Frames: An Integrated API for Mixing Graph and Relational Queries.” In: *SIGMOD Grades Workshop*. 2016.
- [39] 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.
- [40] 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.
- [41] 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.
- [42] Veronika Strnadova-Neeley, Aydin Buluc, Jarrod Chapman, John Gilbert, Joseph E. Gonzalez, and Leonid Olikier. “Efficient Data Reduction for Large-Scale Genetic Mapping”. In: *ACM Conference on Bioinformatics, Computational Biology, and Health Informatics (BCB '15)*. 2015.
- [43] 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.
- [44] David Bader, Aydin Buluç, 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.
- [45] 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.
- [46] 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.
- [47] 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.
- [48] 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. DOI: 10.1109/HPEC.2013.6670338. URL: <https://doi.org/10.1109/HPEC.2013.6670338>.
- [49] 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>.
- [50] 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>.

- [51] 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>.
- [52] 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_scalable_inference_in_latent_variable_models.pdf.
- [53] 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>.
- [54] 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>.
- [55] 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>.
- [56] 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>.
- [57] 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>.
- [58] 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 Publications

- [1] 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>.
- [2] 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: *CoRR* (2020). arXiv: 2006.07357 [cs.DC]. URL: <https://arxiv.org/abs/2006.07357>.
- [3] 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>.
- [4] 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>.
- [5] 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: *CoRR* (2020). arXiv: 2006.11436 [cs.CV]. URL: <https://arxiv.org/abs/2006.11436>.
- [6] 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>.

- [7] 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: *CoRR* abs/2001.04592 (2020). arXiv: 2001.04592. URL: <https://arxiv.org/abs/2001.04592>.
- [8] 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: *CoRR* abs/2003.01835 (2020). arXiv: 2003.01835. URL: <https://arxiv.org/abs/2003.01835>.
- [9] Alvin Wan, Lisa Dunlap, Daniel Ho, Jihan Yin, Scott Lee, Henry Jin, Suzanne Petryk, Sarah Adel Bargal, and Joseph E. Gonzalez. “NBDT: Neural-Backed Decision Trees”. In: *CoRR* (2020). arXiv: 2004.00221 [cs.CV]. URL: <https://arxiv.org/abs/2004.00221>.
- [10] 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>.
- [11] 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: *CoRR* (2020). arXiv: 2007.05086 [cs.LG]. URL: <https://arxiv.org/abs/2007.05086>.
- [12] 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>.
- [13] 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: *CoRR* (2020). arXiv: 2006.06762 [cs.LG]. URL: <https://arxiv.org/abs/2006.06762>.
- [14] 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>.
- [15] Daniel Crankshaw, Gur-Eyal Sela, Corey Zumar, Xiangxi Mo, Joseph E. Gonzalez, Ion Stoica, and Alexey Tumanov. “InferLine: ML Inference Pipeline Composition Framework”. In: *CoRR* abs/1812.01776 (Nov. 2018). arXiv: 1812.01776. URL: <http://arxiv.org/abs/1812.01776>.
- [16] 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>.
- [17] 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>.
- [18] 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>.
- [19] 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* (2018). arXiv: 1812.00497 [cs.LG]. URL: <https://arxiv.org/abs/1812.00497>.
- [20] 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>.

- [21] 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>.
- [22] 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>.
- [23] 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>.

Technical Reports and Invited Publications

- [1] 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>.
- [2] 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>.
- [3] 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. DOI: 10.1145/3190574. URL: <http://doi.acm.org/10.1145/3190574>.
- [4] 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>.
- [5] 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>.
- [6] 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. DOI: 10.1145/2934664. URL: <http://doi.acm.org/10.1145/2934664>.
- [7] 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. DOI: 10.1145/2567948.2580059. URL: <http://doi.acm.org/10.1145/2567948.2580059>.

Selected Invited Talks

- [2018] **Data Eng Conference Keynote** *Managing the Machine Learning Lifecycle*
- [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*

Teaching

Advising

- **Alexander Ku (MS)** [2017-2018] Studied neural network techniques for epigenetic profiling.
- **Corey Zumar (MS)** [2017-2018] Studied resource provisioning for pipeline prediction serving.
- **Samvit Jain (MS)** [2017-Present] Studying methods to accelerate neural network inference on video data.

- **Richard Liaw (PhD) [2017-Present]** Studying the design of systems abstractions for reinforcement learning.
- **Vlad Feinberg (PhD) [2017-Present]** Studying model based variance reduction methods for reinforcement learning.
- **Francois Belletti (PhD) [2015-2018]** Studying algorithms and systems for irregular time-series analysis.
- **Xin Wang (PhD) [2015-Present]** Studying new algorithms for anytime predictions and developing new Deep Learning models for code analysis.
- **Neeraja Yadwadkar (PhD) [2015-2018]** Studying the use of model based techniques for system management.
- **Daniel Crankshaw (PhD) [2015-Present]** Studying the design of systems for managing the life cycle of machine learning models.

Teaching

- **Principles and Techniques of Data Science [Data100]:** Joseph developed and taught 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 600 students a semester. For more information visit <http://ds100.org>.
- **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.

Workshop Organizer

- **[2018] NextProf Nexus Workshop** I helped organize and review applicants for 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 help launch and co-organize the first SOSP learning systems workshop. (<http://learningsys.org/sosp17/>)
- **[2017] NIPS ML Systems Workshop** I help co-organize the NIPS learning systems workshop. (<http://learningsys.org/nips17/>)
- **[2016] ICML ML Systems Workshop** I help co-organize 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>

Grant and Funding

- **[2018]** NSF National Science Foundation, 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.
- **[2017-2018]** Raised gift funding for RISELab from Alibaba, Amazon, Ant Financial ARM, Capital One, Ericsson, Facebook, GE, Google, Huawei, IBM, Intel, Microsoft, Scotiabank, Splunk, and VMware.
- **[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

- **[2019] EuroSys Review Committee**
- **[2019] SOSP Review Committee**
- **[2018] OSDI Extended Review Committee**
- **[2018] ICML Area Chair**
- **[2017] Sigmod PC Member**
- **[2017] AISTats Area Chair**
- **[2016] VLDB PC Member (Best Reviewer Award)**
- **[2017] EuroSys PC Member**
- **[2016] ICML Reviewer**
- **[2016] AISTATS Reviewer**
- **[2014] OSDI Extended 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

- **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

- **Clipper Prediction Serving System** A system for serving predictions across a wide range of machine learning frameworks. <https://clipper.ai>
- **GraphX** GraphX is the 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