# Curriculum Vitae

**Ernest K. Ryu**

310-925-2429
eryu@math.ucla.edu
https://ernestryu.com/

## Academic employment

**24–** (Current) Assistant Professor
Applied Mathematics, UCLA.

**20–24** Assistant Professor
Mathematical Sciences, Seoul National University.

**16–19** Assistant Adjunct Professor
Applied Mathematics, UCLA.

## Education

**10–16** Stanford University, Ph.D. in Applied Mathematics.

**10–16** Stanford University, M.S. in Statistics.

**06–10** California Institute of Technology, B.S. with honors in Physics and Electrical Engineering.

## Research interests

- Optimization
- Reinforcement learning theory
- Deep learning theory
- Multimodal LLMs

## Awards

- Sloan Research Fellowship, 2025
- INFORMS Optimization Society Prize for Young Researchers, 2024.
- Education Prize, College of Natural Science, SNU, 2024.
- KSIAM Outstanding Young Investigator Award, 2023.
- Excellent Research Award, College of Natural Science, SNU, 2023.
- Samsung Science & Technology Foundation grant, 2021.
- Gene Golub Best Thesis Award, 2016.
- Simons Math+X Graduate Fellowship, 2014–16.
- DOE Office of Science Graduate Fellowship (SCGF), 2010–13.
- NASA Tech Brief Award, 2011.
- Caltech Upper Class Merit Award, 2008.

## Selected publications

- LoRA Training Provably Converges to a Low-Rank Global Minimum or It Fails Loudly (But it Probably Won't Fail). J. Kim, J. Kim, and E. K. Ryu, *International Conference on Machine Learning* (Oral, top 120/12107=1.0% of papers), 2025.
- Computer-Assisted Design of Accelerated Composite Optimization Methods: OptISTA. U. Jang, S. Das Gupta, and E. K. Ryu, *Mathematical Programming*, 2025.
- Optimization Algorithm Design via Electric Circuits. S. P. Boyd, T. Parshakova, E. K. Ryu, and J. J. Suh, *Neural Information Processing Systems* (Spotlight, top 326/15671=2.1% of papers), 2024.

- Optimal Acceleration for Minimax and Fixed-Point Problems is Not Unique. T. Yoon, J. Kim, J. J. Suh, E. K. Ryu, *International Conference on Machine Learning* (Spotlight, top (144+191)/9473=3.5% of papers), 2024.
- LoRA Training in the NTK Regime has No Spurious Local Minima. U. Jang, J. D. Lee, and E. K. Ryu, *International Conference on Machine Learning* (Oral, top 144/9473=1.5% of papers), 2024.
- Image clustering conditioned on text criteria. S. Kwon, J. Park, M. Kim, J. Cho, E. K. Ryu, and K. Lee, *International Conference on Learning Representations*, 2024.
- Branch-and-bound performance estimation programming: A unified methodology for constructing optimal optimization methods. S. Das Gupta, B. P. G. Van Parys, and E. K. Ryu, *Mathematical Programming*, 2024.
- Time-reversed dissipation induces duality between minimizing gradient norm and function value. J. Kim, A. Ozdaglar, C. Park, and E. K. Ryu, *Neural Information Processing Systems*, 2023.
- Exact optimal accelerated complexity for fixed-point iterations. J. Park and E. K. Ryu, *International Conference on Machine Learning* (long presentation, top 118/5630=2% of papers), 2022.
- Continuous-time analysis of AGM via conservation laws in dilated coordinate systems. J. J. Suh, G. Roh, and E. K. Ryu, *International Conference on Machine Learning* (long presentation, top 118/5630=2% of papers), 2022.
- Accelerated algorithms for smooth convex-concave minimax problems with $\mathcal{O}(1/k^2)$ rate on squared gradient norm, T. Yoon and E. K. Ryu, *International Conference on Machine Learning* (long presentation, top 166/5513=3% of papers), 2021.
- Plug-and-play methods provably converge with properly trained denoisers. E. K. Ryu, J. Liu, S. Wang, X. Chen, Z. Wang, and W. Yin, *International Conference on Machine Learning*, 2019.

## Students

- Griffin Pinney (Ph.D. 2025–)
- Xinjie He (Ph.D. 2025–)
- Donghwan Rho (Ph.D. 2022–)
- Uijeong Jang (Ph.D. 2022–)
- Jaesung Park (Ph.D. 2023–)
- Kibeom Myoung (M.S. 2023–)
- Hyunsik Chae (M.S. 2023–)
- Junsu Kim (Undergraduate intern 2024–)

## Past students

- Se-Hyeon Kwon (Ph.D. 2020–2024. Current: Samsung Research.)
- Jooyoung Choi (Ph.D. 2021–2024. Current: KRAFTON AI Research.)
- Jisun Park (Ph.D. 2020–2024. Current: Princeton Postdoc.)
- TaeHo Yoon (Ph.D. 2020–2024. Current: Johns Hopkins Postdoc.)
- Jaewook Suh (Ph.D. 2020–2024. Current: Rice Postdoc.)
- Jaeyeon Kim (Undergraduate intern 2022–2024. Current: Harvard CS Ph.D.)
- Jinhee Paeng (Undergraduate intern 2022–2024. Current: Stanford CME PhD.)
- Chanwoo Park (Undergraduate intern 2020-2021. Current: MIT EECS Ph.D.)
- Soheun Yi (Undergraduate intern 2022. Current: CMU Statistics Ph.D.)

## Published books

1. *Large-Scale Convex Optimization via Monotone Operators*, W. Yin and E. K. Ryu, Cambridge University Press, 2022.

## Full list of publications

1. Clip-low increases entropy and clip-high decreases entropy in reinforcement learning of large language models. J. R. Park, J. Kim, G. Kim, J. Jo, S. Choi, J. Cho, and E. K. Ryu, *Under review*, 2025.
2. PinTok: Tokenizers deserve dedicated pinned CPU-compute and memory. S. Choi, M. Chin, and E. K. Ryu, *Under review*, 2025.
3. Why policy gradient algorithms work for undiscounted total-reward MDPs. J. Lee and E. K. Ryu, *Under review*, 2025.
4. STORK: Faster Diffusion and Flow Matching Sampling by Resolving both Stiffness and Structure-Dependence. Z. Tan, W. Wang, A. L. Bertozzi, and E. K. Ryu, *Under review*, 2025.
5. Sharpness-aware minimization can hallucinate minimizers, C. Park, U. Jang, E. K. Ryu, and I. Yang, *Under review*, 2025.
6. Decomposing complex visual comprehension into atomic visual skills for vision language models. H. Chae, S. Yoon, J. Park, C. Y. Chun, Y. Cho, M. Cai, Y. J. Lee, E. K. Ryu, *Under review*, 2025.
7. Finite-Time Bounds for Average-Reward Fitted Q-Iteration. J. Lee and E. K. Ryu, *Neural Information Processing Systems*, 2025.
8. LoRA Training Provably Converges to a Low-Rank Global Minimum or It Fails Loudly (But it Probably Won't Fail). J. Kim, J. Kim, and E. K. Ryu, *International Conference on Machine Learning* (Oral, top 120/12107=1.0% of papers), 2025.
9. Optimal Non-Asymptotic Rates of Value Iteration for Average-Reward Markov Decision Processes. J. Lee and E. K. Ryu, *International Conference on Learning Representations*, 2025.
10. Encryption-Friendly LLM Architecture. D. Rho, T. Kim, M. Park, J. W. Kim, H. Chae, E. K. Ryu, and J. H. Cheon, *International Conference on Learning Representations*, 2025.
11. Computer-Assisted Design of Accelerated Composite Optimization Methods: OptISTA. U. Jang, S. Das Gupta, and E. K. Ryu, *Mathematical Programming*, 2025.
12. Convergence analyses of Davis–Yin splitting via scaled relative graphs. J. Lee, S. Yi, and E. K. Ryu, *SIAM Journal on Optimization*, 2025.
13. Accelerated minimax algorithms flock together. T. Yoon and E. K. Ryu, *SIAM Journal on Optimization*, 2025.
14. Coordinate-Update Algorithms can Efficiently Detect Infeasible Optimization Problems. J. Paeng, J. Park, and E. K. Ryu, *Journal of Mathematical Analysis and Applications*, 2025.
15. Task Diversity Shortens the ICL Plateau. J. Kim, S. Kwon, J. Y. Choi, J. Park, J. Cho, J. D. Lee, and E. K. Ryu, *Transactions on Machine Learning Research*, 2024.
16. Deflated Dynamics Value Iteration. J. Lee, A. Rakhsha, E. K. Ryu, and A.-M. Farahmand, *Transactions on Machine Learning Research, 2025.*
17. Scaled Relative Graph of Normal Matrices. X. Huang, E. K. Ryu, and W. Yin, *Journal of Convex Analysis*, 2025.
18. Optimization Algorithm Design via Electric Circuits. S. P. Boyd, T. Parshakova, E. K. Ryu, and J. J. Suh, *Neural Information Processing Systems* (Spotlight, top 326/15671=2.1% of papers), 2024.
19. Gradient-free Decoder Inversion in Latent Diffusion Models. S. Hong, S. Y. Jeon, K. Lee, E. K. Ryu, and S. Y. Chun, *Neural Information Processing Systems*, 2024.
20. LoRA Training in the NTK Regime has No Spurious Local Minima. U. Jang, J. D. Lee, and E. K. Ryu, *International Conference on Machine Learning* (Oral, top 144/9473=1.5% of papers), 2024.
21. Optimal Acceleration for Minimax and Fixed-Point Problems is Not Unique. T. Yoon, J. Kim, J. J. Suh, E. K. Ryu, *International Conference on Machine Learning* (Spotlight, top (144+191)/9473=3.5% of papers), 2024.

22. Simple Drop-in LoRA Conditioning on Attention Layers will Improve Your Diffusion Model. J. Y. Choi, J. R. Park, I. Park, J. Cho, A. No, and E. K. Ryu, *Transactions on Machine Learning Research*, 2024.
23. Optimal first-order algorithms as a function of inequalities. C. Park and E. K. Ryu, *Journal of Machine Learning Research*, 2024.
24. Image clustering conditioned on text criteria. S. Kwon, J. Park, M. Kim, J. Cho, E. K. Ryu, and K. Lee, *International Conference on Learning Representations*, 2024.
25. Branch-and-bound performance estimation programming: A unified methodology for constructing optimal optimization methods. S. Das Gupta, B. P. G. Van Parys, and E. K. Ryu, *Mathematical Programming*, 2024.
26. Mirror duality in convex optimization. J. Kim, C. Park, J. Diakonikolas, A. Ozdaglar, and E. K. Ryu, *Submitted to SIAM Journal on Optimization*, 2023.
27. Censored sampling of diffusion models using 3 minutes of human feedback. T. Yoon, K. Myoung, K. Lee, J. Cho, A. No, and E. K. Ryu, *Neural Information Processing Systems*, 2023.
28. Time-reversed dissipation induces duality between minimizing gradient norm and function value. J. Kim, A. Ozdaglar, C. Park, and E. K. Ryu, *Neural Information Processing Systems*, 2023.
29. Accelerating value iteration with anchoring. J. Lee and E. K. Ryu, *Neural Information Processing Systems*, 2023.
30. Continuous-time analysis of anchor acceleration. J. J. Suh, J. Park, and E. K. Ryu, *Neural Information Processing Systems*, 2023.
31. Rotation and translation invariant representation learning with implicit neural representations. S. Kwon, J. Y. Choi, and E. K. Ryu, *International Conference on Machine Learning*, 2023.
32. Accelerated infeasibility detection of constrained optimization and fixed-point iterations. J. Park and E. K. Ryu, *International Conference on Machine Learning*, 2023.
33. Factor $\sqrt{2}$-Acceleration of Accelerated Gradient Methods. C. Park, J. Park, and E. K. Ryu, *Applied Mathematics & Optimization*, 2023.
34. Exact optimal accelerated complexity for fixed-point iterations. J. Park and E. K. Ryu, *International Conference on Machine Learning* (long presentation, top 118/5630=2% of papers), 2022.
35. Continuous-time analysis of AGM via conservation laws in dilated coordinate systems. J. J. Suh, G. Roh, and E. K. Ryu, *International Conference on Machine Learning* (long presentation, top 118/5630=2% of papers), 2022.
36. Neural tangent kernel analysis of deep narrow neural networks. J. Lee, J. Y. Choi, E. K. Ryu, and A. No, *International Conference on Machine Learning*, 2022.
37. Robust probabilistic time series forecasting. T. Yoon, Y. Park, E. K. Ryu, and Y. Wang, *International Conference on Artificial Intelligence and Statistics*, 2022.
38. Scaled relative graph: Nonexpansive operators via 2D Euclidean geometry. E. K. Ryu, R. Hannah, and W. Yin, *Mathematical Programming*, 2022.
39. A Geometric Structure of Acceleration and Its Role in Making Gradients Small Fast. J. Lee, C. Park, and E. K. Ryu, *Neural Information Processing Systems*, 2021.
40. Accelerated algorithms for smooth convex-concave minimax problems with $\mathcal{O}(1/k^2)$ rate on squared gradient norm, T. Yoon and E. K. Ryu, *International Conference on Machine Learning* (long presentation, top 166/5513=3% of papers), 2021.
41. WGAN with an infinitely wide generator has no spurious stationary points. A. No, T. Yoon, S. Kwon, and E. K. Ryu, *International Conference on Machine Learning*, 2021.
42. Decentralized proximal gradient algorithms with linear convergence rates. S. A. Alghunaim, E. K. Ryu, K. Yuan, A. H. Sayed, *IEEE Transactions on Automatic Control*, 2021.
43. Tight coefficients of averaged operators via scaled relative graph. X. Huang, E. K. Ryu, and W. Yin, *Journal of Mathematical Analysis and Applications*, 2020.

44. Operator splitting performance estimation: Tight contraction factors and optimal parameter selection. E. K. Ryu, A. Taylor, C. Bergeling, and P. Giselsson. *SIAM Journal on Optimization*, 2020.
45. Splitting with near-circulant linear systems: Applications to total variation CT and PET. E. K. Ryu, S. Ko, and J.-H. Won. *SIAM Journal on Scientific Computing*, 2020.
46. Linear convergence of cyclic SAGA. Y. Park and E. K. Ryu. *Optimization Letters*, 2020.
47. Finding the forward-Douglas–Rachford-forward method. E. K. Ryu and B. C. Vũ, *Journal of Optimization Theory and Applications*, 2020.
48. Uniqueness of DRS as the 2 operator resolvent-splitting and impossibility of 3 operator resolvent-splitting. E. K. Ryu, *Mathematical Programming*, 2020.
49. Plug-and-play methods provably converge with properly trained denoisers. E. K. Ryu, J. Liu, S. Wang, X. Chen, Z. Wang, and W. Yin, *International Conference on Machine Learning*, 2019.
50. Douglas–Rachford splitting for pathological convex optimization. E. K. Ryu, Y. Liu, and W. Yin, *Computational Optimization and Applications*, 2019.
51. A new use of Douglas–Rachford splitting and ADMM for identifying infeasible, unbounded, and pathological conic programs. Y. Liu, E. K. Ryu, and W. Yin, *Mathematical Programming*, 2019.
52. Vector and matrix optimal mass transport: Theory, algorithm, and applications. E. K. Ryu, Y. Chen, W. Li, and S. Osher, *SIAM Journal on Scientific Computing*, 2018.
53. Cosmic divergence, weak cosmic convergence, and fixed points at infinity. E. K. Ryu, *Journal of Fixed Point Theory and its Applications*, 2018.
54. Unbalanced and partial L1 Monge–Kantorovich problem: A scalable parallel first-order method. E. K. Ryu, W. Li, P. Yin, and S. Osher, *Journal of Scientific Computing*, 2018.
55. A parallel method for earth mover's distance. W. Li, E. K. Ryu, S. Osher, W. Yin, and W. Gangbo, *Journal of Scientific Computing*, 2018.
56. Primer on monotone operator methods. E. K. Ryu and S. Boyd, *Applied and Computational Mathematics*, 2016.
57. Risk-constrained Kelly gambling. E. Busseti, E. K. Ryu, and S. Boyd, *Journal of Investing*, 2016.
58. Extensions of Gauss quadrature via linear programming. E. K. Ryu and S. Boyd, *Foundations of Computational Mathematics*, 2015.
59. Computing reaction rates in bio-molecular systems using discrete macro-states. E. Darve and E. K. Ryu. In T. Schlick, editor, *Innovations in Biomolecular Modeling and Simulations*. Royal Society of Chemistry, 2012.
60. Structural characterization of unsaturated phosphatidylcholines using traveling wave ion mobility spectrometry. H. Kim, H. Kim, E. Pang, E. K. Ryu, L. Beegle, J. Loo, W. Goddard, and I. Kanik. *Analytical Chemistry*, 2009.