

# YAQI DUAN

<https://duanyq22.github.io/> | [yaqi.duan@stern.nyu.edu](mailto:yaqi.duan@stern.nyu.edu)  
KMC 8-54, 44 West 4th Street, New York, NY 10012

## ACADEMIC POSITIONS

---

|   |                              |
|---|------------------------------|
| <b>New York University, Stern School of Business</b><br>Assistant Professor in the Department of Technology, Operations, and Statistics | New York, NY<br>2023 –       |
| <b>Massachusetts Institute of Technology</b><br>Postdoc, hosted by Professor Martin J. Wainwright                                       | Cambridge, MA<br>2022 – 2023 |

## EDUCATION

---

|   |                               |
|---|-------------------------------|
| <b>Princeton University</b><br>Ph.D. in Operations Research and Financial Engineering | Princeton, NJ<br>2017 – 2022  |
| <b>Peking University</b><br>B.S. in Mathematics                                       | Beijing, China<br>2013 – 2017 |

## PUBLICATIONS AND PREPRINTS

---

( $\alpha$ - $\beta$ : author names are sorted alphabetically; \*,<sup>†</sup>: equal contribution.)

### Journal publications and preprints

- PILAF: Optimal human preference sampling for reward modeling.  
Feng, Y., Kwiatkowski<sup>†</sup>, A., Zheng, K.<sup>†</sup>, **Duan, Y.\***, Kempe, J.\*.  
*arXiv:2502.04270*.
- Localized exploration in contextual dynamic pricing achieves dimension-free regret.  
Chai, J., **Duan, Y.**, Fan, J., Wang, K. ( $\alpha$ - $\beta$ )  
*arXiv:2412.19252*.
- Taming “data-hungry” reinforcement learning? Stability in continuous state-action spaces.  
**Duan, Y.**, Wainwright, M. J.  
*arXiv:2401.05233*.  
Preliminary version appeared in *Conference on Neural Information Processing Systems (NeurIPS)* 2024.
- Policy evaluation from a single path: Multi-step methods, mixing and mis-specification.  
**Duan, Y.**, Wainwright, M. J.  
*arXiv:2211.03899*.  
Preliminary version appeared in *Annual Learning for Dynamics & Control Conference (L4DC)* 2023.
- Optimal policy evaluation using kernel-based temporal difference methods.  
**Duan, Y.**, Wang, M., Wainwright, M. J.  
*Annals of Statistics (AoS)*, 52 (5), 1927-1952, 2024.  
*2023 IMS Lawrence D. Brown Ph.D. Student Award*

- Adaptive and robust multi-task learning.  
**Duan, Y.**, Wang, K. ( $\alpha$ - $\beta$ )  
*Annals of Statistics (AoS)*, 51(5): 2015-2039, 2023.
- Learning good state and action representations for Markov decision process via tensor decomposition.  
Ni, C., **Duan, Y.**, Dahleh, M., Wang, M., Zhang, A.  
*Journal of Machine Learning Research (JMLR)*, 24(115):1–53, 2023.  
Preliminary version appeared in *IEEE International Symposium on Information Theory (ISIT)* 2021.
- Adaptive low-nonnegative-rank approximation for state aggregation of Markov chains.  
**Duan, Y.**, Wang, M., Wen, Z., Yuan, Y.  
*SIAM Journal on Matrix Analysis and Applications (SIMAX)*, 41(1):pp. 244-278, 2020.

### Conference publications

- Taming “data-hungry” reinforcement learning? Stability in continuous state-action spaces.  
**Duan, Y.**, Wainwright, M. J.  
*Conference on Neural Information Processing Systems (NeurIPS)* 2024.
- A finite-sample analysis of multi-step temporal difference estimates.  
**Duan, Y.**, Wainwright, M. J.  
*Annual Learning for Dynamics & Control Conference (LADC)* 2023.
- Near-optimal offline reinforcement learning with linear representation: leveraging variance information with pessimism.  
Yin, M., **Duan, Y.**, Wang, M., Wang, Y.  
*International Conference on Learning Representations (ICLR)* 2022.
- Risk bounds and Rademacher complexity in batch reinforcement learning.  
**Duan, Y.**, Jin, C., Li, Z. ( $\alpha$ - $\beta$ )  
*International Conference on Machine Learning (ICML)* 2021.
- Bootstrapping statistical inference for off-policy evaluation.  
Hao, B., Ji, X., **Duan, Y.**, Lu, H., Szepesvári, C., Wang, M.  
*International Conference on Machine Learning (ICML)* 2021.
- Sparse feature selection makes reinforcement learning more sample efficient.  
Hao, B., **Duan, Y.**, Lattimore, T., Szepesvári, C., Wang, M.  
*International Conference on Machine Learning (ICML)* 2021.
- Learning good state and action representations via tractable tensor decomposition.  
Ni, C., Zhang, A., **Duan, Y.**, Wang, M.  
*IEEE International Symposium on Information Theory (ISIT)* 2021.
- Minimax-optimal off-policy evaluation with linear function approximation.  
**Duan, Y.**, Wang, M.  
*International Conference on Machine Learning (ICML)* 2020.

- State aggregation learning from Markov transition data.  
**Duan, Y.**, Ke, Z., Wang, M.  
*Conference on Neural Information Processing Systems (NeurIPS) 2019.*
- Learning low-dimensional state embeddings and metastable clusters from time series data.  
Sun, Y., **Duan, Y.**, Gong, H., Wang, M.  
*Conference on Neural Information Processing Systems (NeurIPS) 2019.*

## GRANTS AND FUNDING

---

- NSF grant DMS- 2413812 2024-2027

## SELECTED AWARDS AND HONORS

---

- IMS Lawrence D. Brown Ph.D. Student Award, *Institute of Mathematical Statistics* 2023
- EECS Rising Star, *Massachusetts Institute of Technology* 2021
- Gordon Y. S. Wu Fellowship in Engineering, *Princeton University* 2017 – 2021

## INVITED TALKS

---

- The 2025 Joint Statistical Meetings Aug. 2025
- The 2025 INFORMS Applied Probability Conference June. 2025
- Statistical Learning and Optimization Workshop, Columbia University Apr. 2025
- CILVR Seminar, New York University Feb. 2025
- RL Theory Seminar (virtual) Dec. 2024
- Department of Statistics, Rutgers University Oct. 2024
- S. S. Wilks Memorial Seminar in Statistics, Princeton University Sept. 2024
- Math & Data (MaD) Seminar, New York University Feb. 2024
- The 2023 INFORMS Annual Meeting Oct. 2023
- The 2023 Joint Statistical Meetings Aug. 2023

## PROFESSIONAL SERVICES

---

### *Journal review*

Annals of Statistics (AoS), Journal of the Royal Statistical Society: Series B (JRSSB),  
Journal of the American Statistical Association (JASA), Operations Research (OR),  
Journal of Machine Learning Research (JMLR).

## TEACHING EXPERIENCES

---

### *New York University*

STAT-UB 103 - *Statistics for Business Control, Regression & Forecasting Models*: Spring 2024, 2025  
STAT-UB 003 - *Regression and Forecasting Models*: Spring 2024