

---

## YanJun (Jane) Qi

### *Curriculum Vitae*

Department of Computer Science,  
University of Virginia,  
Charlottesville, Virginia 22904-4740

yq2h@virginia.edu  
yanjun.research@gmail.com (preferred)  
<https://qiyanjun.github.io/Homepage/>

---

## Professional Experience

Associate Professor (tenured), Department of Computer Science, University of Virginia (UVa), Charlottesville, VA (08/2019–Now)

- Research focused on trustworthy Generative AI and deep learning generalization. Past research bridged ideas from deep learning, reinforcement learning, NLP, health informatics, and secure computation.
- Publication H-index is 51 and has 16.7k citations (by Aug.25) wrt GoogleScholar.
- Graduated 6 PhDs (now 2 as CS university faculty and 4 as industry researchers), 15 MS with thesis and 25+ BS with thesis.

Amazon Scholar, Amazon Web Services (02/2025–Now)

- Idea creation, strategic planning, product launch science supports on agentic GenAI;

Principal Applied Scientists, (also as Senior Manager) Amazon Web Services (01/2022–02/2025, on leave from UVa)

- Idea creation and project scoping, strategic planning, cross-org/cross-functional collaborations, team building/ managing. Also have built multiple high-profile solutions for large AWS customers to speed up their cloud AI adoption;
- Launched >10 AWS Bedrock high-value generative AI product features, including model evaluation, model customization, model distillation, and LLM guardrails;

NIH Data Scholar, National Institute of Aging National Institute of Health (01/2021–12/2021)

- Strategic planning and managing large center grant projects ( 20 science teams)

Affiliated faculty, The Center for Public Health Genomics, School of Medicine, University of Virginia, Charlottesville, VA (09/2014–Now)

Affiliated faculty, School of Data Science, University of Virginia, Charlottesville, VA (09/2015–Now)

Visting Professor, Department of Statistics, Sabbatical at the University of California, Berkeley, CA (01/2020–04/2020)

Assistant Professor, Department of Computer Science, University of Virginia, Charlottesville, VA (08/2013–08/2019)

Senior Researcher, Machine Learning Department, NEC Labs America, Princeton, NJ (06/2013–08/2013)

Research Staff Member, Machine Learning Department, NEC Labs America, Princeton, NJ (02/2010–05/2013)

Education

Ph.D., Computer and Information Science, Language Technologies Institute, School of Computer Science, Carnegie Mellon University, Pittsburgh (May 2008)

- Thesis: *Learning of Protein Interaction Networks*
- Advisors: Prof. Ziv Bar-Joseph and Prof. Judith Klein-Seetharaman

M.S., Computer and Information Science, Language Technologies Institute, School of Computer Science, Carnegie Mellon University, Pittsburgh (May 2003)

B. Eng. and M. Eng. in Computer Science, Department of Computer Science, Tsinghua University, Beijing China (July 2001)

- Special program to get B. Eng. and M. Eng. simultaneously
- B.Eng. graduating with the first class distinction: (top 3 out of 153)

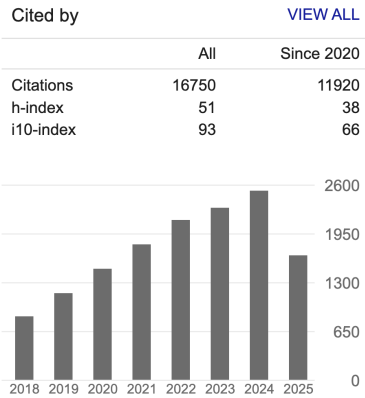
---

Citations Counts

According to Google Scholar

<https://scholar.google.com/citations?user=eXKdSu0AAAAJ&hl=en>

Count statistics On 2025/09/09



Refereed Publications

*Papers under my primary supervision noted with [PS]*

**Conference Papers (Archival Peer Reviewed Proceeding)**

[1] Zhe Wang and **Yanjun Qi**. Augmented adversarial trigger learning. In Luis Chiruzzo, Alan Ritter, and Lu Wang, editors, *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 7068–7100, Albuquerque, New Mexico, April 2025. Association for Computational Linguistics. [PS].

- [2] Zhe Wang, Haozhu Wang, and **Yanjun Qi**. Hierarchical prompt decision transformer: Improving few-shot policy generalization with global and adaptive. In *proceeding of The Web Conference: (WWW25)*. 2025. [\[PS\]](#).
- [3] Jiayu Li, Jennifer Zhu, Fang Liu, and **Yanjun Qi**. AIDE: Attribute-guided Multi-hop data expansion for data scarcity in task-specific fine-tuning. In Georg Rehm and Yunyao Li, editors, *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 1083–1101, Vienna, Austria, July 2025. Association for Computational Linguistics. [\[PS\]](#).
- [4] Jennifer Zhu, Dmitriy Beshpalov, Liwen You, Ninad Kulkarni, and **Yanjun Qi**. TaeBench: Improving quality of toxic adversarial examples. In Weizhu Chen, Yi Yang, Mohammad Kachuee, and Xue-Yong Fu, editors, *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL)*, pages 251–265, Albuquerque, New Mexico, April 2025. Association for Computational Linguistics. [\[PS\]](#).
- [5] Aman Goel, Xian Wu, Zhe Wang, Dmitriy Beshpalov, and **Yanjun Qi**. TurboFuzzLLM: Turbocharging mutation-based fuzzing for effectively jailbreaking large language models in practice. In Weizhu Chen, Yi Yang, Mohammad Kachuee, and Xue-Yong Fu, editors, *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL)*, pages 523–534, Albuquerque, New Mexico, April 2025. Association for Computational Linguistics. [\[PS\]](#).
- [6] Zifan Xu, Haozhu Wang, Dmitriy Beshpalov, Xian Wu, Peter Stone, and **Yanjun Qi**. LaRS: Latent reasoning skills for chain-of-thought reasoning. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen, editors, *Association for Computational Linguistics: EMNLP 2024*, pages 3624–3643, Miami, Florida, USA, November 2024. Association for Computational Linguistics. [\[PS\]](#).
- [7] Tong Wang, Ninad Kulkarni, and **Yanjun Qi**. Less is more for improving automatic evaluation of factual consistency. In *North American Association for Computational Linguistics: NAACL*. Association for Computational Linguistics, 2024. [\[PS\]](#).
- [8] Dmitriy Beshpalov, Sourav Bhabesh, Yi Xiang, Liutong Zhou, and **Yanjun Qi**. Towards building a robust toxicity predictor. In Sunayana Sitaram, Beata Beigman Klebanov, and Jason D Williams, editors, *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 581–598. Association for Computational Linguistics, 2023. [\[PS\]](#).
- [9] Arshdeep Sekhon<sup>(v)</sup>, Hanjie Chen, Aman Shrivastava, Zhe Wang, Yangfeng Ji, and **Yanjun Qi**. Improving interpretability via explicit word interaction graph layer. *Thirty-Seventh AAAI Conference on Artificial Intelligence (AAAI)*, 2023. [\[PS\]](#).
- [10] Zhe Wang, Jake Grigsby, and **Yanjun Qi**. PGrad: Learning principal gradients for domain generalization. In *The Eleventh International Conference on Learning Representations*, 2023. [\[PS\]](#).
- [11] Zhe Wang, Jake Grigsby, Arshdeep Sekhon<sup>(v)</sup>, and **Yanjun Qi**. St-maml : A stochastic-task based method for task-heterogeneous meta-learning. In James Cussens and Kun Zhang, editors, *Proceedings of the Thirty-Eighth Conference on Uncertainty in Artificial Intelligence*, volume 180 of *Proceedings of Machine Learning Research*, pages 2066–2074. PMLR, 01–05 Aug 2022. [\[PS\]](#).
- [12] Arshdeep Sekhon<sup>(v)</sup>, Yangfeng Ji, Matthew Dwyer, and **Yanjun Qi**. White-box testing of NLP models with mask neuron coverage. In *Findings of the Association for*

*Computational Linguistics: NAACL 2022*, pages 1547–1558, Seattle, United States, July 2022. Association for Computational Linguistics. [PS].

[13] Jin Yong Yoo and **Yanjun Qi**. Towards improving adversarial training of nlp models. *EMNLP*, 2021. [PS].

[14] Jack Lanchantin<sup>(v)</sup>, Tianlu Wang, Vicente Ordonez, and **Yanjun Qi**. General multi-label image classification with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16478–16488, 2021. [PS].

[15] Jack Lanchantin<sup>(v)</sup>, Tom Weingarten, Arshdeep Sekhon<sup>(v)</sup>, Clint Miller, and **Yanjun Qi**. Transfer learning for predicting virus-host protein interactions for novel virus sequences. In *Proceedings of the 12th ACM Conference on Bioinformatics, Computational Biology, and Health Informatics*, BCB ’21, New York, NY, USA, 2021. Association for Computing Machinery. [PS].

[16] Arshdeep Sekhon<sup>(v)</sup>, Beilun Wang<sup>(v)</sup>, Zhe Wang, and **Yanjun Qi**. Beyond data samples: Aligning differential networks estimation with scientific knowledge. *AISTAT*, 2022. [PS].

[17] Paola Cascante-Bonilla, Fuwen Tan, **Yanjun Qi**, and Vicente Ordonez. Curriculum labeling: Self-paced pseudo-labeling for semi-supervised learning. 2021. *Spotlight Talk / Acceptance rate = 22%*; [CS].

[18] John Morris<sup>(v)</sup>, Eli Lifland, Jack Lanchantin<sup>(v)</sup>, Yangfeng Ji, and **Yanjun Qi**. Reevaluating adversarial examples in natural language. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, pages 3829–3839, 2020. *Acceptance Rate = 23%* [PS].

[19] John Morris<sup>(v)</sup>, Eli Lifland, Jin Yong Yoo, Jake Grigsby, Di Jin, and **Yanjun Qi**. Textattack: A framework for adversarial attacks, data augmentation, and adversarial training in nlp. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 119–126, 2020. *Acceptance Rate = 22%* [PS].

[20] Jin Yong Yoo, John Morris<sup>(v)</sup>, Eli Lifland, and **Yanjun Qi**. Searching for a search method: Benchmarking search algorithms for generating nlp adversarial examples. In *Proceedings of the Third BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 323–332, 2020. *Acceptance Rate = 28%* [PS].

[21] Jack Lanchantin<sup>(v)</sup>, Arshdeep Sekhon<sup>(v)</sup>, Ritambhara Singh<sup>(v)</sup>, and **Yanjun Qi**. Neural message passing for multi-label sequence classification. *Proceedings of European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML-PKDD)*, 2019. *Long Talk / Acceptance rate = 20%*; [PS].

[22] Ji Gao<sup>(v)</sup>, Jack Lanchantin<sup>(v)</sup>, Mary Lou Soffa, and **Yanjun Qi**. Black-box generation of adversarial text sequences to evade deep learning classifiers. *IEEE Security and Privacy Symposium (SPW)*, pages 50–56, 2018. *Acceptance rate = 29% (21/73)* ; [PS].

[23] Beilun Wang<sup>(v)</sup>, Arshdeep Sekhon<sup>(v)</sup>, and **Yanjun Qi**. A fast and scalable joint estimator for integrating additional knowledge in learning multiple related sparse gaussian graphical models. In *Proceedings of The 35th International Conference on Machine Learning (ICML)*, 2018. *Acceptance rate = 24.9%*; [PS].

[24] Beilun Wang<sup>(v)</sup>, Arshdeep Sekhon<sup>(v)</sup>, and **Yanjun Qi**. Fast and scalable learning of sparse changes in high-dimensional gaussian graphical model structure. In *Proceedings*

of The 21st International Conference on Artificial Intelligence and Statistics (AISTATS), 2018. Acceptance rate = 28.9% ; [PS].

[25] Weilin Xu<sup>(v)</sup>, David Evans, and **Yanjun Qi**. Feature squeezing: Detecting adversarial examples in deep neural networks. In *Proceedings of the Network and Distributed System Security Symposium*, (NDSS), Feb. 2018. [CS].

[26] I.K. Kim, W. Wang, **Yanjun Qi**, and M. Humphrey. Cloudinsight: Utilizing a council of experts to predict future cloud application workloads. In *IEEE International Conference on Cloud Computing*, 2018. Acceptance rate = 15% ; [AC].

[27] Ritambhara Singh<sup>(v)</sup>, Jack Lanchantin<sup>(v)</sup>, Arshdeep Sekhon<sup>(v)</sup>, and **Yanjun Qi**. Attend and predict: Understanding gene regulation by selective attention on chromatin. In *Proceedings of Advances in Neural Information Processing Systems (NIPS)*, 2017. Acceptance rate = 20.9% (678/3240) ; [PS].

[28] Ritambhara Singh<sup>(v)</sup>, Arshdeep Sekhon<sup>(v)</sup>, Kamran Kowsari, Jack Lanchantin<sup>(v)</sup>, Beilun Wang<sup>(v)</sup>, and **Yanjun Qi**. Gakco: a fast gapped k-mer string kernel using counting. In *Proceedings of European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML-PKDD)*, 2017. Acceptance rate = 20% ; [PS].

[29] Beilun Wang<sup>(v)</sup>, Ji Gao<sup>(v)</sup>, and **Yanjun Qi**. A fast and scalable joint estimators for learning multiple related sparse gaussian graphical models. In *Proceedings of The 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 1168–1177, 2017. Acceptance rate = 28.9% ; [PS].

[30] Beilun Wang<sup>(v)</sup>, Ji Gao<sup>(v)</sup>, and **Yanjun Qi**. A theoretical framework for robustness of (deep) classifiers against adversarial samples. In *the International Conference on Learning Representations (ICLR)*, 2017. Acceptance rate = 32% ; [PS].

[31] Jack Lanchantin<sup>(v)</sup>, Ritambhara Singh<sup>(v)</sup>, Beilun Wang<sup>(v)</sup>, and **Yanjun Qi**. Deep motif dashboard: Visualizing and understanding genomic sequences using deep neural networks. *Pacific Symposium on Biocomputing (PSB)*, 2017. No acceptance rate data available ; [PS].

[32] Andrew Norton<sup>(v)</sup> and **Yanjun Qi**. Adversarial-playground: A visualization suite showing how adversarial examples fool deep learning. In *Visualization for Cyber Security (VizSec), 2017 IEEE Symposium on*, pages 1–4. IEEE, 2017. Acceptance rate = 37% ; [PS].

[33] Zeming Lin<sup>(v)</sup>, Jack Lanchantin<sup>(v)</sup>, and **Yanjun Qi**. MUST-CNN: A multilayer shift-and-stitch deep convolutional architecture for sequence-based protein structure prediction. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence*, (AAAI), Feb. 2016. Acceptance rate = 25% ; [PS].

[34] Ritambhara Singh<sup>(v)</sup> and **Yanjun Qi**. Character based string kernels for bio-entity relation detection. *ACL BioNLP 2016*, page 66, 2016. Acceptance rate = 27% ; [PS].

[35] Weilin Xu<sup>(v)</sup>, **Yanjun Qi**, and D. Evans. Automatically evading classifiers. In *Proceedings of the Network and Distributed System Security Symposium*, (NDSS), Feb. 2016. Acceptance rate = 15% ; [CS].

[36] I.K. Kim, W. Wang, **Yanjun Qi**, and M. Humphrey. Empirical evaluation of workload forecasting techniques for predictive cloud resource scaling. In *9th IEEE International Conference on Cloud Computing*, 2016. Acceptance rate = 15% ; [AC].

- [37] S.M. Preum<sup>(v)</sup>, J. Stankovic, , and **Yanjun Qi**. MAPer: A multi-scale adaptive personalized model for temporal human behavior prediction. In *Proceedings of the 24th ACM International Conference on Information and Knowledge Management*, (CIKM), Oct. 2015. Acceptance rate = 18% (87/484); [CS].
- [38] J. Gong, J. Lach, **Yanjun Qi**, and MD. Goldman. Causal analysis of inertial body sensors for enhancing gait assessment separability towards multiple sclerosis diagnosis. In *Proceedings of IEEE Body Sensor Network*, (BSN), Jun. 2015. No acceptance rate data available; [CS].
- [39] K. Wang, **Yanjun Qi**, J. Fox, MR. Stan, and K. Skadron. Association rule mining with the micron automata processor. In *29th IEEE International Parallel & Distributed Processing Symposium (IPDPS)*, May. 2015. Acceptance rate = 21.8%, 108 out of 496; [AC].
- [40] O. Tastan, **Yanjun Qi**, J. Carbonell, and J. Klein-Seetharaman. Refining literature curated hiv-1, human protein-protein interactions using expert opinions. In *Proceedings of Pacific Symposium on Biocomputing (PSB)*, volume 20, Jan. 2015. No acceptance rate data available; [CS].
- [41] J. Gong, P. Asare, J. Lach, and **Yanjun Qi**. Piecewise linear dynamical model for actions clustering from inertial body sensors with considerations of human factors. In *Proceedings of 9th International Conference on Body Area Networks (BodyNets)*, Sep. 2014. Acceptance rate = 35% ; [PS].
- [42] **Yanjun Qi**, S. Das<sup>(s)</sup>, J. Weston, and R. Collobort. A deep learning framework for character-based information extraction. In *Proceedings of the European Conference on Information Retrieval (ECIR)*, Apr. 2014. No acceptance rate data available; [PS].
- [43] S. Das<sup>(s)</sup>, **Yanjun Qi**, P. Mitra, and L. Giles. Extracting metadata from academic homepages using labeled features. In *Proceedings of SIAM International Conference on Data Mining (SDM)*, May. 2014. Acceptance rate = 29% (120/389); [CS].
- [44] Y. He<sup>(s)</sup>, K. Kavukcuoglu, Y. Wang, A. Szlam, and **Yanjun Qi**. Unsupervised feature learning by deep sparse coding. In *Proceedings of SIAM International Conference on Data Mining (SDM)*, May. 2014. Acceptance rate = 15% (60/389); [PS].
- [45] I.K. Kim, J. Steele, **Yanjun Qi**, and M. Humphrey. Comprehensive elastic resource management to ensure predictable performance for scientific applications on public iaas clouds. In *IEEE/ACM 7th International Conference on Utility and Cloud Computing (UCC)*, Dec. 2014. No acceptance rate data available; [AC].
- [46] R. Min, S. Chowdhury<sup>(s)</sup>, **Yanjun Qi**, A. Stewart, and R. Ostroff. An integrated approach to blood-based cancer diagnosis and biomarker discovery. In *Proceedings of Pacific Symposium on Biocomputing (PSB)*, volume 19, 2014. No acceptance rate data available; [CS].
- [47] **Yanjun Qi** and P. Laquerre. Retrieving medical records with sennamed: NEC labs america at TREC 2012 medical record track. In *NIST Special Publication for the Text Retrieval Conference (TREC)*, 2012. Ranked 2nd/82 auto runs, 3rd/all 88 runs; [PS].
- [48] Y. He<sup>(s)</sup>, **Yanjun Qi**, K. Kavukcuoglu, and H. Park. Learning the dependency structure of latent factors. In *Proceedings of Advances in Neural Information Processing Systems (NIPS)*, 2012. Acceptance rate = 25% (370/1467); [PS].
- [49] D. Bespalov<sup>(s)</sup>, **Yanjun Qi**, B. Bai, and A. Shokoufandeh. Sentiment classification with supervised sequence encoder. In *Proceedings of European Conference on Machine*

*Learning and Principles and Practice of Knowledge Discovery in Databases (ECML-PKDD)*, 2012. Acceptance rate = 23% (/443); [PS].

[50] D. Bespalov<sup>(s)</sup>, **Yanjun Qi**, B. Bai, and A. Shokoufandeh. Large-scale image classification using supervised spatial encoder. In *Proceedings of 21st conference of the International Association for Pattern Recognition (ICPR)*, 2012. No acceptance rate data available; [PS].

[51] I. Nouretdinov, A. Gammerman, , **Yanjun Qi**, and J. Klein-Seetharaman. Determining confidence of predicted interactions between hiv-1 and human proteins using conformal method. In *Proceedings of Pacific Symposium on Biocomputing*, volume 17, 2012. No acceptance rate data available; [AC].

[52] D. Bespalov<sup>(s)</sup>, B. Bai, **Yanjun Qi**, and A. Shokoufandeh. Sentiment classification based on supervised latent n-gram analysis. In *Proceedings of the 20th ACM international conference on Information and knowledge management*, pages 375–382. ACM, 2011. Acceptance rate = 15% (137/917); [CS].

[53] X. Chen<sup>(s)</sup>, **Yanjun Qi**, B. Bai, Q. Lin, and J.G. Carbonell. Sparse latent semantic analysis. In *Proceedings of SIAM International Conference on Data Mining (SDM)*, 2011. Acceptance rate = 25% (86/343); [CS].

[54] X. Ning<sup>(s)</sup> and **Yanjun Qi**. Semi-supervised convolution graph kernels for relation extraction. In *Proceedings of SIAM International Conference on Data Mining (SDM)*, 2011. Acceptance rate = 25% (86/343); [PS].

[55] X. Chen<sup>(s)</sup>, B. Bai, **Yanjun Qi**, Q. Lin, and J. Carbonell. Learning preferences with millions of parameters by enforcing sparsity. In *Proceedings of IEEE International Conference on Data Mining (ICDM)*, pages 779–784, 2010. Acceptance rate = 19% (155/797); [PS].

[56] **Yanjun Qi**, O. Tastan, J.G. Carbonell, J. Klein-Seetharaman, and J. Weston. Semi-supervised multi-task learning for predicting interactions between HIV-1 and human proteins. In *Proceedings of the 9th European Conference on Computational Biology (ECCB)*, pages i645–i654, 2010. Acceptance rate = 17% (36/215); [PS].

[57] P. Kuksa<sup>(s)</sup> and **Yanjun Qi**. Semi-supervised bio-named entity recognition with word-codebook learning. In *Proceedings of SIAM International Conference on Data Mining (SDM)*, 2010. Acceptance rate = 23% (82/351); [PS].

[58] P. Kuksa<sup>(s)</sup>, **Yanjun Qi**, B. Bai, R. Collobert, J. Weston, V. Pavlovic, and X. Ning<sup>(s)</sup>. Semi-supervised abstraction-augmented string kernel for multi-level bio-relation extraction. In *Proceedings of European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML-PKDD)*, 2010. Acceptance rate = 17% (110/658); [PS].

[59] **Yanjun Qi**, P. Kuksa<sup>(s)</sup>, R. Collobert, K. Sadamasa, K. Kavukcuoglu, and J. Weston. Semi-supervised sequence labeling with self-learned features. In *Proceedings of the 9th IEEE International Conference on Data Mining (ICDM)*, pages 428–437, 2009. Acceptance rate = 9% (70/786); [CS].

[60] **Yanjun Qi**, R. Collobert, P. Kuksa<sup>(s)</sup>, K. Kavukcuoglu, and J. Weston. Combining labeled and unlabeled data with word-class distribution learning. In *Proceeding of the 18th ACM International Conference on Information and Knowledge Management (CIKM)*, pages 1737–1740. ACM, 2009. Acceptance rate = 20% (294/847); [CS].

- [61] O. Tastan, **Yanjun Qi**, J.G. Carbonell, and J. Klein-Seetharaman. Prediction of interactions between HIV-1 and human proteins by information integration. In *Proceedings of Pacific Symposium on Biocomputing*, volume 516, 2009. No acceptance rate data available; [CS].
- [62] B. Bai, J. Weston, D. Grangier, R. Collobert, K. Sadamasa, **Yanjun Qi**, C. Cortes, and M. Mohri. Polynomial semantic indexing. In *Proceedings of Advances in Neural Information Processing Systems (NIPS)*, pages 64–72, 2009. Acceptance rate = 23% (263/1105); [AC].
- [63] B. Bai, J. Weston, D. Grangier, R. Collobert, K. Sadamasa, **Yanjun Qi**, O. Chapelle, and K. Weinberger. Supervised semantic indexing. In *Proceeding of the 18th ACM International Conference on Information and Knowledge Management (CIKM)*, pages 187–196. ACM, 2009. Acceptance rate = 15% (123/847); [AC].
- [64] **Yanjun Qi**, F. Balem, C. Faloutsos, J. Klein-Seetharaman, and Z. Bar-Joseph. Protein complex identification by supervised graph local clustering. In *Proceedings of the 16th Annual International Conference Intelligent Systems for Molecular Biology (ISMB)*, pages i250–i258, 2008. Acceptance rate = 17% (49/292); [AC].
- [65] **Yanjun Qi**, J. Klein-Seetharaman, and Z. Bar-Joseph. Random forest similarity for protein-protein interaction prediction from multiple sources. In *Proceedings of Pacific Symposium on Biocomputing*, volume 10, pages 531–542, 2005. No acceptance rate data available; [AC].
- [66] **Yanjun Qi**, A. Hauptmann, and T. Liu. Supervised classification for video shot segmentation. In *Proceedings of IEEE Conference on Multimedia & Expro (ICME)*, volume 2, pages 689–692, 2003. No acceptance rate data available; [AC].
- [67] A. Hauptmann, R. Yan, **Yanjun Qi**, R. Jin, M. Christel, M. Derthick, MY Chen, R. Baron, WH Lin, and T.D. Ng. Video classification and retrieval with the informedia digital video library system. In *NIST Special Publication for the Eleventh Text Retrieval Conference*, pages 119–127, 2003. No acceptance rate data available; [AC].
- [68] A. Hauptmann, R. Jin, N. Papernick, D. Ng, **Yanjun Qi**, R. Houghton, and S. Thornton. Video retrieval with the informedia digital video library system. In *NIST Special Publication for the Tenth Text Retrieval Conference*, pages 94–102, 2002. No acceptance rate data available; [AC].
- [69] R. Jin, **Yanjun Qi**, and A. Hauptmann. A probabilistic model for camera zoom detection. In *Proceedings of the 16th International Conference on Pattern Recognition, ICPR '02*, pages 30859–30865, Washington, DC, USA, 2002. No acceptance rate data available; [AC].
- [70] **Yanjun Qi**, Y. Zhong, and S. Yang. Design and implementation of STB sopca software modules. In *Proceeding of International Society for Optical Engineering (SPIE) for Digital Cinema and Microdisplays*, volume 4207-18, pages 70–76, 2000.

#### Journal Papers (Full-Length Peer Reviewed)

- [71] Xi Fang, Weijie Xu, Fiona Anting Tan, Jiani Zhang, Ziqing Hu, **Yanjun Qi**, Scott Nickleach, Diego Socolinsky, Srinivasan Sengamedu, and Christos Faloutsos. Large language models (LLMs) on tabular data: Prediction, generation, and understanding – a survey. *Transactions on Machine Learning Research (TMLR)*, 2024.
- [72] Zichen Wang, Ryan Brand, Jared Adolf-Bryfogle, Jasleen Grewal, **Yanjun Qi**, Steven A. Combs, Nataliya Golovach, Rebecca Alford, Huzefa Rangwala, and Peter M.



Clark. EGGNet, a generalizable geometric deep learning framework for protein complex pose scoring. 9(7):7471–7479. [PS].

[73] Jack Lanchantin<sup>(v)</sup> and Yanjun Qi. Graph convolutional networks for epigenetic state prediction using both sequence and 3D genome data. *Bioinformatics*, 36(S2):i659–i667, 12 2020. *Impact Factor: 7.307; Acceptance Rate = 16% [PS]*.

[74] Derrick Blakely<sup>(v)</sup>, Eamon Collins, Ritambhara Singh<sup>(v)</sup>, Andrew Norton<sup>(v)</sup>, Jack Lanchantin<sup>(v)</sup>, and Yanjun Qi. FastSK: fast sequence analysis with gapped string kernels. *Bioinformatics*, 36(S2):i857–i865, 12 2020. *Impact Factor: 7.307; Acceptance Rate = 16% [PS]*.

[75] Arshdeep Sekhon<sup>(v)</sup>, Ritambhara Singh<sup>(v)</sup>, and Yanjun Qi. Deepdiff: Deep-learning for predicting differential gene expression from histone modifications. *Bioinformatics*, 34(17):i891–i900, 2018. *Impact Factor: 7.307; Acceptance Rate = 15% (40/280) [PS]*.

[76] Travers Ching, Daniel S Himmelstein, Brett K Beaulieu-Jones, Alexandr A Kalinin, Brian T Do, Gregory P Way, Enrico Ferrero, Paul-Michael Agapow, Michael Zietz, Michael M Hoffman, others, Yanjun Qi, et al. Opportunities and obstacles for deep learning in biology and medicine. *Journal of The Royal Society Interface*, 15(141), 2018. *Impact Factor 3.917; [AC]*.

[77] Beilun Wang<sup>(v)</sup>, Ritambhara Singh<sup>(v)</sup>, and Yanjun Qi. A constrained L1 minimization approach for estimating multiple sparse gaussian or nonparanormal graphical models. *Machine Learning*, 106:1381–1417, 2017. *Impact Factor: 1.848; [PS]*.

[78] Ritambhara Singh<sup>(v)</sup>, Jack Lanchantin<sup>(v)</sup>, Gabriel Robins, and Yanjun Qi. DeepChrome: deep-learning for predicting gene expression from histone modifications. *Bioinformatics*, 32(17):i639–i648, 2016. *Impact Factor: 7.307; [PS]*.

[79] Feiyu Xiong<sup>(v)</sup>, Moshe Kam, Leonid Hrebien, Beilun Wang<sup>(v)</sup>, and Yanjun Qi. Kernelized information-theoretic metric learning for cancer diagnosis using high-dimensional molecular profiling data. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 10(4):38, 2016. *Impact Factor: 1.892; [CS]*.

[80] Ritambhara Singh<sup>(v)</sup>, Jack Lanchantin<sup>(v)</sup>, Gabriel Robins, and Yanjun Qi. Transfer string kernel for cross-context dna-protein binding prediction. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 2016. *Impact Factor: 1.955; [PS]*.

[81] Sarah Mohamed, Nick Janus, and Yanjun Qi. SCODE: a cytoscape app for supervised complex detection in protein-protein interaction graphs. *F1000Research*, 5, 2016. *RG Impact Fator: 1.13; [PS]*.

[82] J. Gong, P. Asare, Yanjun Qi, and J. Lach. Piecewise linear dynamical model for actions clustering from inertial body sensors with considerations of human factors. *IEEE Transactions on Affective Computing*, 2016. *Impact Factor 3.149; [AC]*.

[83] Ritambhara Singh<sup>(v)</sup>, C. Kucsu, A. Quinlan, Yanjun Qi, and M. Adli. Cas9-chromatin binding information enables more accurate crispr off-target prediction. *Nucleic Acids Research*, page gkv575, Oct. 2015. *Impact Factor 8.808; [CS]*.

[84] Yanjun Qi, M. Oja, J. Weston, and W.S. Noble. A unified multitask architecture for predicting local protein properties. *PLoS ONE*, 7(3):e32235, 2012. *Impact Factor 4.411; [PS]*.

- [85] **Yanjun Qi**, O. Tastan, J.G. Carbonell, J. Klein-Seetharaman, and J. Weston. Semi-supervised multi-task learning for predicting interactions between HIV-1 and human proteins. *Bioinformatics (Oxford Journal)*, 26(18):i645–i655, 2010. *Impact Factor 5.468*; [PS].
- [86] B. Bai, J. Weston, D. Grangier, R. Collobert, K. Sadamasu, **Yanjun Qi**, O. Chapelle, and K. Weinberger. Learning to rank with (a lot of) word features. *Information Retrieval*, 13(3):291–314, 2010. *Impact Factor: 0.917*; [AC].
- [87] **Yanjun Qi**, H.K. Dhiman, N. Bhola, I. Budyak, S. Kar, D. Man, A. Dutta, K. Tirupula, B.I. Carr, J. Grandis, et al. Systematic prediction of human membrane receptor interactions. *Proteomics*, 9(23):5243–5255, 2009. *Impact Factor 5.479*; [AC].
- [88] **Yanjun Qi**, F. Balem, C. Faloutsos, J. Klein-Seetharaman, and Z. Bar-Joseph. Protein complex identification by supervised graph local clustering. *Bioinformatics (Oxford Journal)*, 24(13):i250–i258, 2008. *Impact Factor 5.468*; [AC].
- [89] L. Peña-Castillo, M. Tasan, C.L. Myers, H. Lee, T. Joshi, C. Zhang, Y. Guan, M. Leone, A. Pagnani, W.K. Kim, **Yanjun Qi**, others, and F.P. Roth. A critical assessment of Mus musculus gene function prediction using integrated genomic evidence. *Genome Biology*, 9(Suppl 1):S2, 2008. *Impact Factor 6.15*; [AC].
- [90] H. Nozawa, G. Howell, S. Suzuki, Q. Zhang, **Yanjun Qi**, J. Klein-Seetharaman, A. Wells, J.R. Grandis, and S.M. Thomas. Combined Inhibition of PLC $\gamma$ -1 and c-Src Abrogates Epidermal Growth Factor Receptor-Mediated Head and Neck Squamous Cell Carcinoma Invasion. *Clinical Cancer Research*, 14(13):4336–4350, 2008. *Impact Factor 6.488*; [AC].
- [91] **Yanjun Qi**, J. Klein-Seetharaman, and Z. Bar-Joseph. A mixture of feature experts approach for protein-protein interaction prediction. *BMC Bioinformatics*, 8(Suppl 10):S6, 2007. *Impact Factor 3.78*; [AC].
- [92] **Yanjun Qi**, Z. Bar-Joseph, and J. Klein-Seetharaman. Evaluation of different biological data and computational classification methods for use in protein interaction prediction. *Proteins: Structure, Function, and Bioinformatics*, 63(3):490–500, 2006. *Impact Factor 3.39*; [AC].
- [93] A.G. Hauptmann, J. Gao, R. Yan, **Yanjun Qi**, J. Yang, and H.D. Wactlar. Automated analysis of nursing home observations. *IEEE Pervasive Computing, Special Issue on Pervasive Computing for Successful Aging*, 3(2):15–21, 2004. *Impact Factor: 1.545*; [AC].

#### Book & Chapters & Dissertation

- [94] **Yanjun Qi**. Random forest for bioinformatics. In C. Zhang and Y. Ma, editors, *Ensemble Machine Learning: Methods and Applications*, pages 307–322. Springer, 2012.
- [95] **Yanjun Qi** and W. Noble. Protein interaction networks: Protein domain interaction and protein function prediction. In *Handbook of Computational Statistics: Statistical Bioinformatics*. Springer, 2011.
- [96] **Yanjun Qi**. *Learning of Protein-Protein Interaction Networks*. PhD thesis, Carnegie Mellon University, May 2008.

#### Workshop Papers (Archival Peer Reviewed Proceeding)

- [97] Wanyu Du, Song Feng, James Gung, Lijia Sun, Yi Zhang, Saab Mansour, and **Yanjun Qi**. DFLOW: Diverse dialogue flow simulation with large language models. In Ehsan

Kamalloo, Nicolas Gontier, Xing Han Lu, Nouha Dziri, Shikhar Murty, and Alexandre Lacoste, editors, *Proceedings of the 1st Workshop for Research on Agent Language Models (REALM 2025)*, pages 17–32, Vienna, Austria, July 2025. Association for Computational Linguistics. [\[PS\]](#).

[98] Daniel Schwarz, Dmitriy Beshpalov, Zhe Wang, Ninad Kulkarni, and **Yanjun Qi**. Graph of attacks with pruning: Optimizing stealthy jailbreak prompt. generation for enhanced LLM content moderation. In Agostina Calabrese, Christine de Kock, Debora Nozza, Flor Miriam Plaza-del Arco, Zeerak Talat, and Francielle Vargas, editors, *Proceedings of the The 9th Workshop on Online Abuse and Harms (WOAH)*, pages 482–489, Vienna, Austria, August 2025. Association for Computational Linguistics. [\[PS\]](#).

[99] Vanamala Venkataswamy, Jake Grigsby, Andrew Grimshaw, and **Yanjun Qi**. Launchpad: Learning to schedule using offline and online RL methods. In Dalibor Klusáček, Julita Corbalán, and Gonzalo P. Rodrigo, editors, *Job Scheduling Strategies for Parallel Processing*, pages 60–83, Cham, 2024. Springer Nature Switzerland. [\[PS\]](#).

[100] Mihai Christodorescu, Ryan Craven, et al, **Yanjun Qi**, Khawaja Shams, and Matthew Turek. Securing the future of genai: Policy and technology. May 2024.

[101] Hanyu Liu, Chengyuan Cai, and **Yanjun Qi**. Expanding scope: Adapting english adversarial attacks to chinese. In Anaelia Ovalle, Kai-Wei Chang, Ninareh Mehrabi, Yada Pruksachatkun, Aram Galystan, Jwala Dhamala, Apurv Verma, Trista Cao, Anoop Kumar, and Rahul Gupta, editors, *Proceedings of the 3rd Workshop on Trustworthy Natural Language Processing (TrustNLP 2023)*, pages 276–286. Association for Computational Linguistics. [\[PS\]](#).

[102] Aman Shrivastava, **Yanjun Qi**, and Vicente Ordonez. Estimating and maximizing mutual information for knowledge distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshop on Robustness*, pages 48–57, 2022.

[103] Jake Grigsby and **Yanjun Qi**. A closer look at advantage-filtered behavioral cloning in high-noise datasets, 2023. [\[PS\]](#).

[104] Jake Grigsby, Zhe Wang, Nam Nguyen, and **Yanjun Qi**. Long-range transformers for dynamic spatiotemporal forecasting. *Proceeding of the 9th SIGKDD International Workshop on Mining and Learning from Time Series (MILETS23) + arXiv preprint arXiv:2109.12218*, 2021.

[105] Sanchit Sinha, Hanjie Chen, Arshdeep Sekhon<sup>(v)</sup>, Yangfeng Ji, and **Yanjun Qi**. Perturbing inputs for fragile interpretations in deep natural language processing. In Jasmijn Bastings, Yonatan Belinkov, Emmanuel Dupoux, Mario Giulianelli, Dieuwke Hupkes, Yuval Pinter, and Hassan Sajjad, editors, *Proceedings of the Fourth BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 420–434. Association for Computational Linguistics. [\[PS\]](#).

[106] Vanamala Venkataswamy, Jake Grigsby, Andrew Grimshaw, and **Yanjun Qi**. Rare: Renewable energy aware resource management in datacenters. In *Job Scheduling Strategies for Parallel Processing: 25th International Workshop, JSSPP 2022, Virtual Event, June 3, 2022, Revised Selected Papers*, pages 108–130. Springer, 2022.

[107] Jake Grigsby, Jin Yong Yoo, and **Yanjun Qi**. Towards automatic actor-critic solutions to continuous control. *Deep Reinforcement Learning Workshop at NeurIPS*, 2021.

[108] Jake Grigsby and **Yanjun Qi**. Measuring visual generalization in continuous control from pixels. *Deep Reinforcement Learning Workshop at NeurIPS*, 2020.

- [109] Sanchit Sinha, Hanjie Chen, Arshdeep Sekhon and Yangfeng Ji, and **Yan-jun Qi**. Fragile interpretation in deep natural language processing, 2021. *EMNLP Blackbox NLP*, [PS].
- [110] Arshdeep Sekhon<sup>(v)</sup>, Zhe Wang, and **Yan-jun Qi**. Relate and predict: Structure-aware prediction with jointly optimized neural dependency graph. *ICML Graph Relational Learning Workshop*, 2020. [PS].
- [111] John Morris<sup>(v)</sup>, Jin Yong Yoo, and **Yan-jun Qi**. Textattack: Lessons learned in designing python frameworks for nlp. In *Proceedings of Second Workshop for NLP Open Source Software (NLP-OSS)*, pages 126–131, 2020. [PS].
- [112] Arshdeep Sekhon<sup>(v)</sup>, Beilun Wang<sup>(v)</sup>, Zhe Wang, and **Yan-jun Qi**. Beyond data samples: Aligning differential networks estimation with scientific knowledge. *AISTAT*, 2022. [PS].
- [113] Jack Lanchantin<sup>(v)</sup>, Tom Weingarten, Arshdeep Sekhon<sup>(v)</sup>, Clint Miller, and **Yan-jun Qi**. Transfer learning for predicting virus-host protein interactions for novel virus sequences. In *Proceedings of the 12th ACM Conference on Bioinformatics, Computational Biology, and Health Informatics*, BCB '21, New York, NY, USA, 2021. Association for Computing Machinery. [PS].
- [114] Jack Lanchantin<sup>(v)</sup>, Arshdeep Sekhon<sup>(v)</sup>, Ritambhara Singh<sup>(v)</sup>, and **Yan-jun Qi**. Prototype matching networks for large-scale multi-label genomic sequence classification. *NeurIPS Learning Meaningful Representations of Life Workshop (LMRL)*, 2019. [PS].
- [115] Jack Lanchantin<sup>(v)</sup>, Ritambhara Singh<sup>(v)</sup>, and **Yan-jun Qi**. Memory matching networks for genomic sequence classification. *International Conference on Learning Representations (ICLR) Workshop Track*, 2017. Acceptance rate = 32% ; [PS].
- [116] Chandan Singh, Beilun Wang<sup>(v)</sup>, and **Yan-jun Qi**. A constrained, weighted-l1 minimization approach for joint discovery of heterogeneous neural connectivity graphs. *NIPS 2017 workshop for Advances in Modeling and Learning Interactions from Complex Data*, 2017. Acceptance rate = 40% (17/42); [PS].
- [117] Ji Gao<sup>(v)</sup>, Beilun Wang<sup>(v)</sup>, and **Yan-jun Qi**. DeepCloak: Masking dnn models for robustness against adversarial samples. *International Conference on Learning Representations (ICLR) Workshop Track*, 2017. Acceptance rate = 32% ; [PS].
- [118] Ritambhara Singh<sup>(v)</sup>, G. Robins, and **Yan-jun Qi**. Transfer string kernel for cross-context transcription factor binding prediction. In *Proceedings of the 14th International Workshop on Data Mining in Bioinformatics at 21st ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, (BIOKDD), Oct. 2015. No acceptance rate data available; [PS].
- [119] Weilin Xu<sup>(v)</sup>, **Yan-jun Qi**, and D. Evans. Automatically evading classifiers. Posters of the 6th IEEE Symposium on Security and Privacy, May. 2015. No acceptance rate data available; [CS].
- [120] J. Gong, J. Lach, **Yan-jun Qi**, R.S. Zee, S.J. Seaman, and N.S. Schenkman. Motion assessment for robotic surgery education using inertial body sensors. In *Proceedings of the Wireless Health Annual Conference*, Oct. 2014. Acceptance rate = 25% ; Best poster award; [CS].
- [121] Feiyu Xiong<sup>(v)</sup>, M. Kam, L. Hrebien, and **Yan-jun Qi**. Ranking with distance metric learning for biomedical severity detection. In *Proceedings of SIAM International Confer-*

ence on Data Mining (SDM), 3rd Workshop on Data Mining for Medicine and Healthcare (DMMH), May. 2014. No acceptance rate data available; [PS].

[122] Y. He<sup>(s)</sup>, K. Kavukcuoglu, Y. Wang, A. Szlam, and **YanJun Qi**. Unsupervised feature learning by deep sparse coding. International Conference on Learning Representations (ICLR) Workshop, May. 2014. No acceptance rate data available; [PS].

[123] R. Min, X. Ning<sup>(s)</sup>, **YanJun Qi**, C. Cheng, I. Durdanovic, and M. Gerstein. Ensemble learning based sparse high-order boltzmann machine for unsupervised feature interaction identification. NIPS Workshop on Challenges in Learning Hierarchical Models: Transfer Learning and Optimization, Dec. 2014. Acceptance rate = 25%; [CS].

[124] Y. He<sup>(s)</sup>, **YanJun Qi**, K. Kavukcuoglu, and H. Park. Structured latent factor analysis. NIPS Workshop on Challenges in Learning Hierarchical Models: Transfer Learning and Optimization, 2011. No acceptance rate data available; [CS].

[125] X. Chen<sup>(s)</sup>, **YanJun Qi**, B. Bai, Q. Lin, and J.G. Carbonell. Sparse latent semantic analysis. NIPS Workshop on Practical Application of Sparse Modeling: Open Issues and New Directions, 2010. No acceptance rate data available; [CS].

[126] **YanJun Qi**, O. Tastan, J. Carbonell, J. Klein-Seetharaman, and J. Weston. Semi-supervised embedding for predicting protein-protein interactions from multiple sources. The Learning Workshop in Clearwater, 2009. No acceptance rate data available; [AC].

[127] B. Bai, J. Weston, D. Grangier, R. Collobert, K. Sadamasu, **YanJun Qi**, O. Chapelle, and Kilian Weinberger. Learning to rank with low rank. SIGIR workshop: Learning to Rank for Information Retrieval, 2009. No acceptance rate data available; [AC].

[128] **YanJun Qi**, J. Klein-Seetharaman, and Z. Bar-Joseph. A mixture of experts approach for protein-protein interaction prediction. NIPS workshop on Computational Biology and the Analysis of Heterogeneous Data, 2009. No acceptance rate data available; [AC].

[129] H.K. Dhiman, **YanJun Qi**, A. Dutta, K. Tirupula, Z. Bar-Joseph, and J. Klein-Seetharaman. Validation of predicted protein-protein interaction partners of the egfr. 51st Annual Meeting of the Biophysical Society in Baltimore, Maryland, 2007. Selected Oral Talk.

[130] E. Gardner, Y. Tung, **YanJun Qi**, N. Yanamala, K. Tirupula, D. Man, T. Reinhart, Z. Bar-Joseph, and J. Klein-Seetharaman. Rhodopsin interacts with chemokines. 51st Annual Meeting of the Biophysical Society in Baltimore, Maryland, 2007. Selected Oral Talk.

[131] **YanJun Qi**, Z. Bar-Joseph, and J. Klein-Seetharaman. Predicting protein-protein interactions from labeled and unlabeled data. Cold Spring Harbor Meeting of Systems Biology: Genomics Approaches to Transcriptional Regulation, 2004.

## Preprints

[132] Kefan Song, Amir Moeini, Peng Wang, Lei Gong, Rohan Chandra, **YanJun Qi**, and Shangdong Zhang. Reward is enough: Llms are in-context reinforcement learners, 2025.

[133] Aman Goel, Daniel Schwartz, and **YanJun Qi**. Zero-knowledge llm hallucination detection and mitigation through fine-grained cross-model consistency. *arXiv preprint arXiv:2508.14314*, 2025. [PS].

- [134] Xun Xian, Tong Wang, Liwen You, and **Yanjun Qi**. Understanding data poisoning attacks for RAG: Insights and algorithms, 2025. [\[PS\]](#).
- [135] Zhuotong Chen, Fang Liu, Xuan Zhu, **Yanjun Qi**, and Mohammad Ghavamzadeh. Preference optimization via contrastive divergence: Your reward model is secretly an nll estimator, 2025.
- [136] Changhoon Kim and **Yanjun Qi**. A comprehensive survey on concept erasure in text-to-image diffusion models, 2025. [\[PS\]](#).
- 

## Granted Patents

- [137] R. Min and **Yanjun Qi**. Sparse higher-order markov random field, Jun. 2013. [US Patent 9,183,503; Type: Granted.](#)
- [138] D. Besspalov<sup>(s)</sup>, B. Bai, and **Yanjun Qi**. Sentiment classification based on supervised latent n-gram analysis, Oct. 2012. [US Patent App. 13/424,900.](#)
- [139] **Yanjun Qi** and B. Bai. Document classification with weighted supervised n-gram embedding, Nov. 2014. [US Patent 8,892,488; Type: Granted.](#)
- [140] Y. He<sup>(s)</sup>, **Yanjun Qi**, and K. Kavukcuoglu. Latent factor deendency structure determination, Mar. 2015. [US Patent 8,977,579; Type: Granted.](#)
- [141] X. Chen<sup>(s)</sup>, **Yanjun Qi**, and B. Bai. System and methods for finding hidden topics of documents and preference ranking documents, Dec. 2013. [US Patent 8,738,547; Type: Granted.](#)
- [142] **Yanjun Qi**, B. Bai, X. Ning<sup>(s)</sup>, and P. Kuksa<sup>(s)</sup>. Systems and methods for semi-supervised relationship extraction, Apr. 2014. [US Patent 8,874,432; Type: Granted.](#)
- 

## UVA Graduate Students Mentored

*Graduated Ph.D. Mentee from Department of Computer Science (CS), University of Virginia (UVA)*

1. Ritambhara Singh (Ph.D. (female), as advisor 2013-2018)  
*Current Job: Tenured Associate Professor at Brown University Computer Science (Tenure track 19-25)*
2. Beilun Wang (Ph.D., as advisor 2013-2018)  
*Current Job: Associate Professor at School of Computer Science and Engineering of Southeast University in China.*
3. Weilin Xu (Ph.D. , as advisor, co-advisor: Dr. David Evans 2013-2019)  
*Current Job: Research Staff Member in Intel Research.*
4. Jack Lanchantin (Ph.D., as advisor 2015-2021)  
*Current Job: Researcher Staff Member in Meta Research.*
5. Arshdeep Sekhon (Ph.D., as advisor 2016-2022)  
*Current Job: Data Scientist in Microsoft.*
6. Zhe Wang (Ph.D., as advisor 2019-2024)  
*Current Job: Data Scientist in Amazon Web Service AI.*

7. Kefan Song (Ph.D. student, as advisor 2025.06-Now)

*Master Students from Department of Computer Science (CS), University of Virginia (UVA)*

1. Zeming Lin (CS Undergraduate, BSCS Capstone + Master thesis, both as advisor 2014-2016.01)  
*Current Job: Founder of EvolutionaryScale.*
2. Jefferey Yoo (Distinguished BA + Master thesis, both as Advisor 2018-2021)
3. Derrick Blakely (BSCS capstone + Master Thesis, both as Advisor 2017-2019.12)
4. Kevin Ivey (BSCS Capstone + master advisor 09/2019– 12/2023)
5. Sinha, Sanchit (Master student, as thesis advisor)
6. Ji Gao (Master Thesis, as advisor 2015-2018)
7. Nicholas Janus (Master student, as advisor 2013-2015)
8. Sarah Masud Preum (Master student (female), as co-advisor with Dr. Jack Stankovic 2013-2015)
9. Tung Dao (Master student, as co-advisor, with Dr. Jack W. Davidson 2013-2014)

## UVA Undergraduate Research Students Mentored

*Undergraduate Thesis Students from Department of Computer Science (CS), University of Virginia (UVA)*

1. Jake Grigsby (Fourth year undergraduate, as Distinguished BA advisor 2019-2022)
2. John Morris (Fourth year undergraduate students, as Capstone thesis advisor 2019-2020)
3. Hanyu Liu, Chengyuan Cai, Grant Dong, Meghan Anderson, Malcolm Mashig, Nithin Vijayakumar, Tanay Bapat (Third to Fourth year undergraduate students, as research advisor 2020-2022)
4. Dillon Lue (Fourth year undergraduate students, as capstone research advisor 2020-2021)
5. Mint Lin, Tahmid Kazi, Yuchen Sun, Alan Zheng, Soukarya Ghosh (fourth year undergraduate students, as capstone research advisor 2020-2021)
6. Eli Lifeland (Fourth year undergraduate students, as Capstone BS advisor 2019-2020)
7. Eric Wang (Fourth year undergraduate students, as Distinguished BA advisor 2018-2019)
8. Yang Yu (Third year undergraduate students, as Capstone BS advisor 2019)
9. Jennifer Fang (Fourth year undergraduate students, as capstone advisor 2018-2019)
10. Brandon Liu (Fourth year undergraduate students, as Distinguished BA advisor 2028-2019)
11. Eamon Collins (Fourth year undergraduate students, as capstone advisor 2017)
12. Christopher chen (Fourth year undergraduate students, as capstone advisor)
13. Zachary Skemp (Fourth year undergraduate student, as capstone advisor 2017)

14. Andrew Norton (Fourth year undergraduate student, as capstone advisor 2016-2017)
  15. Chandan Singh (Fourth year undergraduate student, as capstone advisor 2016-2017)
  16. Muthu Chidambaram (Undergraduate, as independent study advisor 2016-2017)
  17. Sarah Mohamed (Fourth year undergraduate student (female), as thesis advisor 2015-2016)
  18. Sugandha Agrawal (Fourth-year capstone student (female), as thesis advisor 2015-16)
  19. Anat Gilboa (Fourth year undergraduate student (female), as thesis advisor 2014–05/2015,)
  20. Jasdev Singh (Undergraduate thesis advisor 2013–05/2014)
- 

## Invited Talks/Keynotes:

### *Invited Conference Keynotes & Tutorials & External Institutional Talks*

- “Towards Building Robust Deep NLP Systems and better Content Moderation”  
+ *Invited Keynote speech at the NeurIPS workshop on Responsibly Building Next Generation of Multimodal Foundation Models (2024)*  
+ *Invited Keynote at the Sabancı University Science Academy Artificial Learning Summer School (2021)*
- “Deep Learning and 3 Principals for Biomedical Data Analytics”  
+ *Invited talk at NIH Data Science Showcase Seminar Series (2021)*  
+ *Invited Talk at the National Cancer Institute Data Science Seminar Series (2021)*
- “Interpretable Deep Learning for Sequential Data for Gene Regulation”  
+ *Invited Tutorials at the 9th ACM Conference on Bioinformatics, Computational Biology, and Health Informatics (ACM BCB) (2018)*  
+ *Invited Tutorials at UCLA Computational Genomics Winter Institute (2018)*
- “Joint Nets: Extracting Relations from Heterogeneous Samples”, Invited Tutorials  
+ *Invited Tutorials at UCLA Computational Genomics Winter Institute (2019)*
- “Feature Squeezing for Detecting Adversarial Examples in Deep Neural Networks”,  
+ *Distinguished webinar talk at the Institute for Information Infrastructure Protection (I3P) hosted by the George Washington University and SRI International (2018)*  
+ *The ARO Invitational Workshop on Foundations of Autonomous Adaptive Cyber Systems at the George Mason University (2019)*
- “Making Deep Learning Understandable when Analyzing Sequential Data”  
+ *Invited Talk at National Institutes of Health as part of the National Center for Biotechnology Information (NCBI) Scientific Visitors Program (2018)*  
+ *Invited Keynote Speech at the Pacific Northwest National Laboratory, Computing @ PNNL Lecture Series (2017)*
- “Machine Learning for Biomedical Data Analytics”  
+ *Invited Talk at the University of Massachusetts Medical School Molecular Medicine Seminar (Mar. 2013)*  
+ *Invited Talk at Case Western Reserve University School of Engineering (Mar. 2013)*



- “Machine Learning for Computational Biology”,  
+ *Invited Talk in the Seminar Series of Drexel University ECE (2013)*  
+ *Invited Talk in the Seminar Series of New Jersey Institute of Technology CS, Newark NJ (Dec. 2012)*
- “Machine Learning in NEC Labs America”, Invited Tutorial at the San Diego Center for Algae Biotechnology of University of California San Diego (2011)
- “Relational Learning for Protein Interaction Networks Discovery”,  
+ *Harvard University, School of Medicine, Boston, MA (2008)*  
+ *University of Washington, Department of Genome Sciences, Seattle, WA (2008)*  
+ *Siemens Medical Solution, Malvern, PA (2007)*
- “Information Integration for Learning of Protein Interaction Networks”,  
+ *Invited Talk at the Molecular Biophysics and Structural Biology Data and Literature Club, University of Pittsburgh (2006)*

#### *Invited Institutional Talks (Internal)*

- “Interpretable Deep Learning Interpretable”,  
+ *University of Virginia, Blitzarama Seminar hosted by the Center for Dynamics of Healthy Development (Oct. 2018)*
- “Making Deep Learning Interpretable for Analyzing Sequential Data about Gene Regulation”,  
+ *University of Virginia, Genome Sciences Seminar hosted by the Center for Public Health Genomics (Oct. 2018)*
- “Deep Learning: an Overview” ,  
+ *University of Virginia, Human and Machine Intelligence Seminar (2018)*
- “A Constrained, Weighted-L1 Minimization Approach for Joint Discovery of Heterogeneous Neural Connectivity Graphs” ,  
+ *University of Virginia, Neuroscience Graduate Program Retreat (2017)*
- “Machine Learning for ”Big Data Complexity” in Biomedical Data Analytics”,  
+ *University of Virginia, Data Science Institute Datapalooza (2016)*  
+ *University of Virginia, Department of Molecular Physiology and Biological Physics (Jul. 2015)*  
+ *University of Virginia, Genome Sciences Seminar hosted by the Center for Public Health Genomics (Mar. 2014)*  
+ *University of Virginia Department of Computer Science (Apr. 2013)*
- “Machine Learning for Tackling ”Big Data Complexity” in Cancer Data”,  
+ *University of Virginia, Emily Couric Cancer Center (Feb. 2015)*
- “Advice from Early Career Faculty” invited talk at CRA Career Mentoring Workshop (2014)
- “Machine Learning for Modern Health-care”, NEC Labs Pizza Talk Series, NJ (2010)

---

## Teaching Experience

- 08/2025–12/2025 “Machine Learning and Deep Learning” (Senior Undergraduate level), Instructor, Department of Computer Science, University of Virginia
- 01/2025–05/2025 “Generative AI - Overview” (Master + Senior Undergraduate level - Seminar), Instructor, Department of Computer Science, University of Virginia

01/2024–05/2024	“Generative AI - Risks and Benefits” (Graduate Advanced level - Seminar), Instructor, Department of Computer Science, University of Virginia
01/2022–05/2022	“Machine Learning Foundation, Deep Learning, and Good Uses” (Undergraduate Advanced level), Instructor, Department of Computer Science, University of Virginia
08/2020–12/2020	“Machine Learning Foundation, Deep Learning, and Good Use on COVID19” (Undergraduate Advanced level), Instructor, Department of Computer Science, University of Virginia
08/2019–12/2019	“Machine Learning” (Master level), Instructor, Department of Computer Science, University of Virginia
01/2019–05/2019	“Deep Learning Advances Graphs” (PhD Student level), Instructor, Department of Computer Science, University of Virginia
08/2018–12/2018	“Machine Learning” (Undergraduate level), Instructor, Department of Computer Science, University of Virginia
01/2018–05/2018	“Machine Learning” (Undergraduate level), Instructor, Department of Computer Science, University of Virginia
08/2017–12/2017	“Advanced Deep Learning” (PhD Student level), Instructor, Department of Computer Science, University of Virginia
08/2016–12/2016	“Machine Learning” (Master+ Advanced Senior level), Instructor, Department of Computer Science, University of Virginia
08/2015–12/2015	“Machine Learning” (Master level), Instructor, Department of Computer Science, University of Virginia
01/2015–05/2015	“Special topic: Large Scale Machine Learning” (Graduate-level), Instructor, Department of Computer Science, University of Virginia
01/2015–05/2015	“Special topic: Large Scale Machine Learning” (Graduate-level), Instructor, Department of Computer Science, University of Virginia
08/2014–12/2014	“Introduction to Machine Learning and Data Mining” (Undergraduate-JuniorGrad-level), Instructor, Department of Computer Science, University of Virginia
01/2014–05/2014	“Special topic: Machine Learning and Data Mining in Practice” (Graduate-level), Instructor, Department of Computer Science, University of Virginia
08/2013–12/2013	“Special topic: Machine Learning and Data Mining” (Graduate-level), Instructor, Department of Computer Science, University of Virginia

---