

KEVIN YANG
CURRICULUM VITAE

425-786-7398 yangk@berkeley.edu

EDUCATION

- Ph.D. in Artificial Intelligence** 2019-
UC Berkeley
Advisor: Prof. Dan Klein
- M.Eng., Computer Science** 2018-2019
MIT
Advisor: Prof. Regina Barzilay
Thesis: *Are Learned Molecular Representations Ready for Prime Time?*
- B.S, Computer Science and Mathematics** 2015-2019
MIT
Double major; GPA 5.0/5.0

AWARDS AND HONORS

- NSF Graduate Fellowship 2019
Putnam Top 200 2016-2018
International Linguistics Olympiad Gold Medal (5th place) 2015
USA Math Olympiad Honorable Mention (14th place) 2013-2015

RESEARCH INTERESTS

I am broadly interested in controllable generation, and also in designing more modular and/or language-controllable agents. I have worked primarily in the NLP and biochemistry domains, and am also interested in game-playing agents.

PUBLICATIONS

Learning Space Partitions for Path Planning

Kevin Yang, Tianjun Zhang, Chris Cummins, Brandon Cui et al.

NeurIPS 2021

We propose a path planning method inspired by a theoretical analysis of search space partitioning, and show strong performance on difficult multimodal, long-horizon path planning problems.

FUDGE: Controlled Text Generation with Future Discriminators

Kevin Yang, Dan Klein

NAACL 2021

We propose a simple, flexible, and highly effective method for controlling generation toward desired attributes using lightweight classifiers.

A Streaming Approach for Efficient Batched Beam Search

Kevin Yang, Violet Yao, John DeNero, Dan Klein

EMNLP 2020

We propose an efficient batching strategy for variable-length decoding on GPU architectures, demonstrating substantial speedups over existing fixed-width and variable-width beam searches.

Improving Molecular Design by Stochastic Iterative Target Augmentation

Kevin Yang, Wengong Jin, Kyle Swanson, Regina Barzilay, Tommi Jaakkola
ICML 2020

We use a simple and theoretically motivated self-training approach guided by an external property predictor to substantially improve over state-of-the-art approaches in molecular design.

Uncertainty Quantification Using Neural Networks for Molecular Property Prediction

Lior Hirschfeld, Kyle Swanson, Kevin Yang, Regina Barzilay, Tommi Jaakkola
JCIM 2020

We comprehensively evaluate and compare several approaches for uncertainty estimation in neural models on molecular property prediction tasks.

A Deep Learning Approach to Antibiotic Discovery

Jonathan Stokes, Kevin Yang, Kyle Swanson, Wengong Jin et al.
Cell 2020

We use computational property prediction models to screen drug databases for potential antibiotic activity, and discover previously unknown antibiotics with novel mechanisms of action which are effective even against bacteria which are resistant to commonly used antibiotics.

Analyzing Learned Molecular Representations for Property Prediction

Kevin Yang, Kyle Swanson, Wengong Jin, Connor Coley et al.
JCIM 2019

We introduce a new variant of message-passing neural networks, demonstrating consistently strong performance that significantly improves over existing baselines on many datasets. We also carefully benchmark models on both public and proprietary industry datasets.

Learning Multimodal Graph-to-Graph Translation for Molecular Optimization

Wengong Jin, Kevin Yang, Regina Barzilay, Tommi Jaakkola
ICLR 2019

We introduce an encoder-decoder architecture for molecular optimization that operates directly on the molecular graph, substantially outperforming string-based baselines as well as pre-existing state of the art.

PROFESSIONAL ACTIVITIES

Research Intern, Microsoft Research (Semantic Machines) 2021

12-week summer internship working on multiturn semantic parsing with a particular focus on data privacy.

Teaching Assistant, MIT 2018

Course: Introduction to Inference

Research Intern, ASAPP 2018

12-week summer internship researching and productionizing natural language processing models on difficult real-world text datasets.

Equities Research Intern, DE Shaw 2017

11-week summer internship in equities research group. I developed a model predicting existence of hidden liquidity in equities exchanges, with strong results on real-world data.

Software Engineering Intern, Google 2016

12-week summer internship focusing on improving search ad efficiency and optimization.

LANGUAGES

Proficient in Python and Java; experienced in C, R, HTML, Javascript. Also proficient in PyTorch. Natively fluent in Mandarin Chinese, advanced Japanese (JLPT N3), advanced Spanish.

REFERENCES

Dan Klein, Professor
Computer Science Division, UC Berkeley
510-642-3214 klein@berkeley.edu

Regina Barzilay, Professor
CSAIL, MIT
617-258-5706 regina@csail.mit.edu