


Welcome to TTIC!

Avrim Blum

Chief Academic Officer




What is TTIC?

A COMPUTER SCIENCE GRADUATE INSTITUTE

Affiliated with the
University of Chicago

Supported by an
endowment provided by
TTI-Japan and Toyota




A COMPUTER SCIENCE GRADUATE INSTITUTE

What does that mean?

- We are a single-department university with one program: a PhD in Computer Science (with a Master's within the PhD). Currently ≈ 40 PhD students.
- Tenure-track faculty who conduct research, advise students, teach classes. Research focused on Machine Learning, Theoretical Computer Science, NLP, Speech, Computer Vision, Robotics, and Computational Biology.
- Closely affiliated with the University of Chicago. Students can (and do) take classes across and can be co-advised across.
- Also have 3-year Research Assistant Professors (RAPs) who add to the research environment.

A research-focused mini-university of Machine Learning, AI, Theory, and applications, closely affiliated with the University of Chicago.

Where are we?

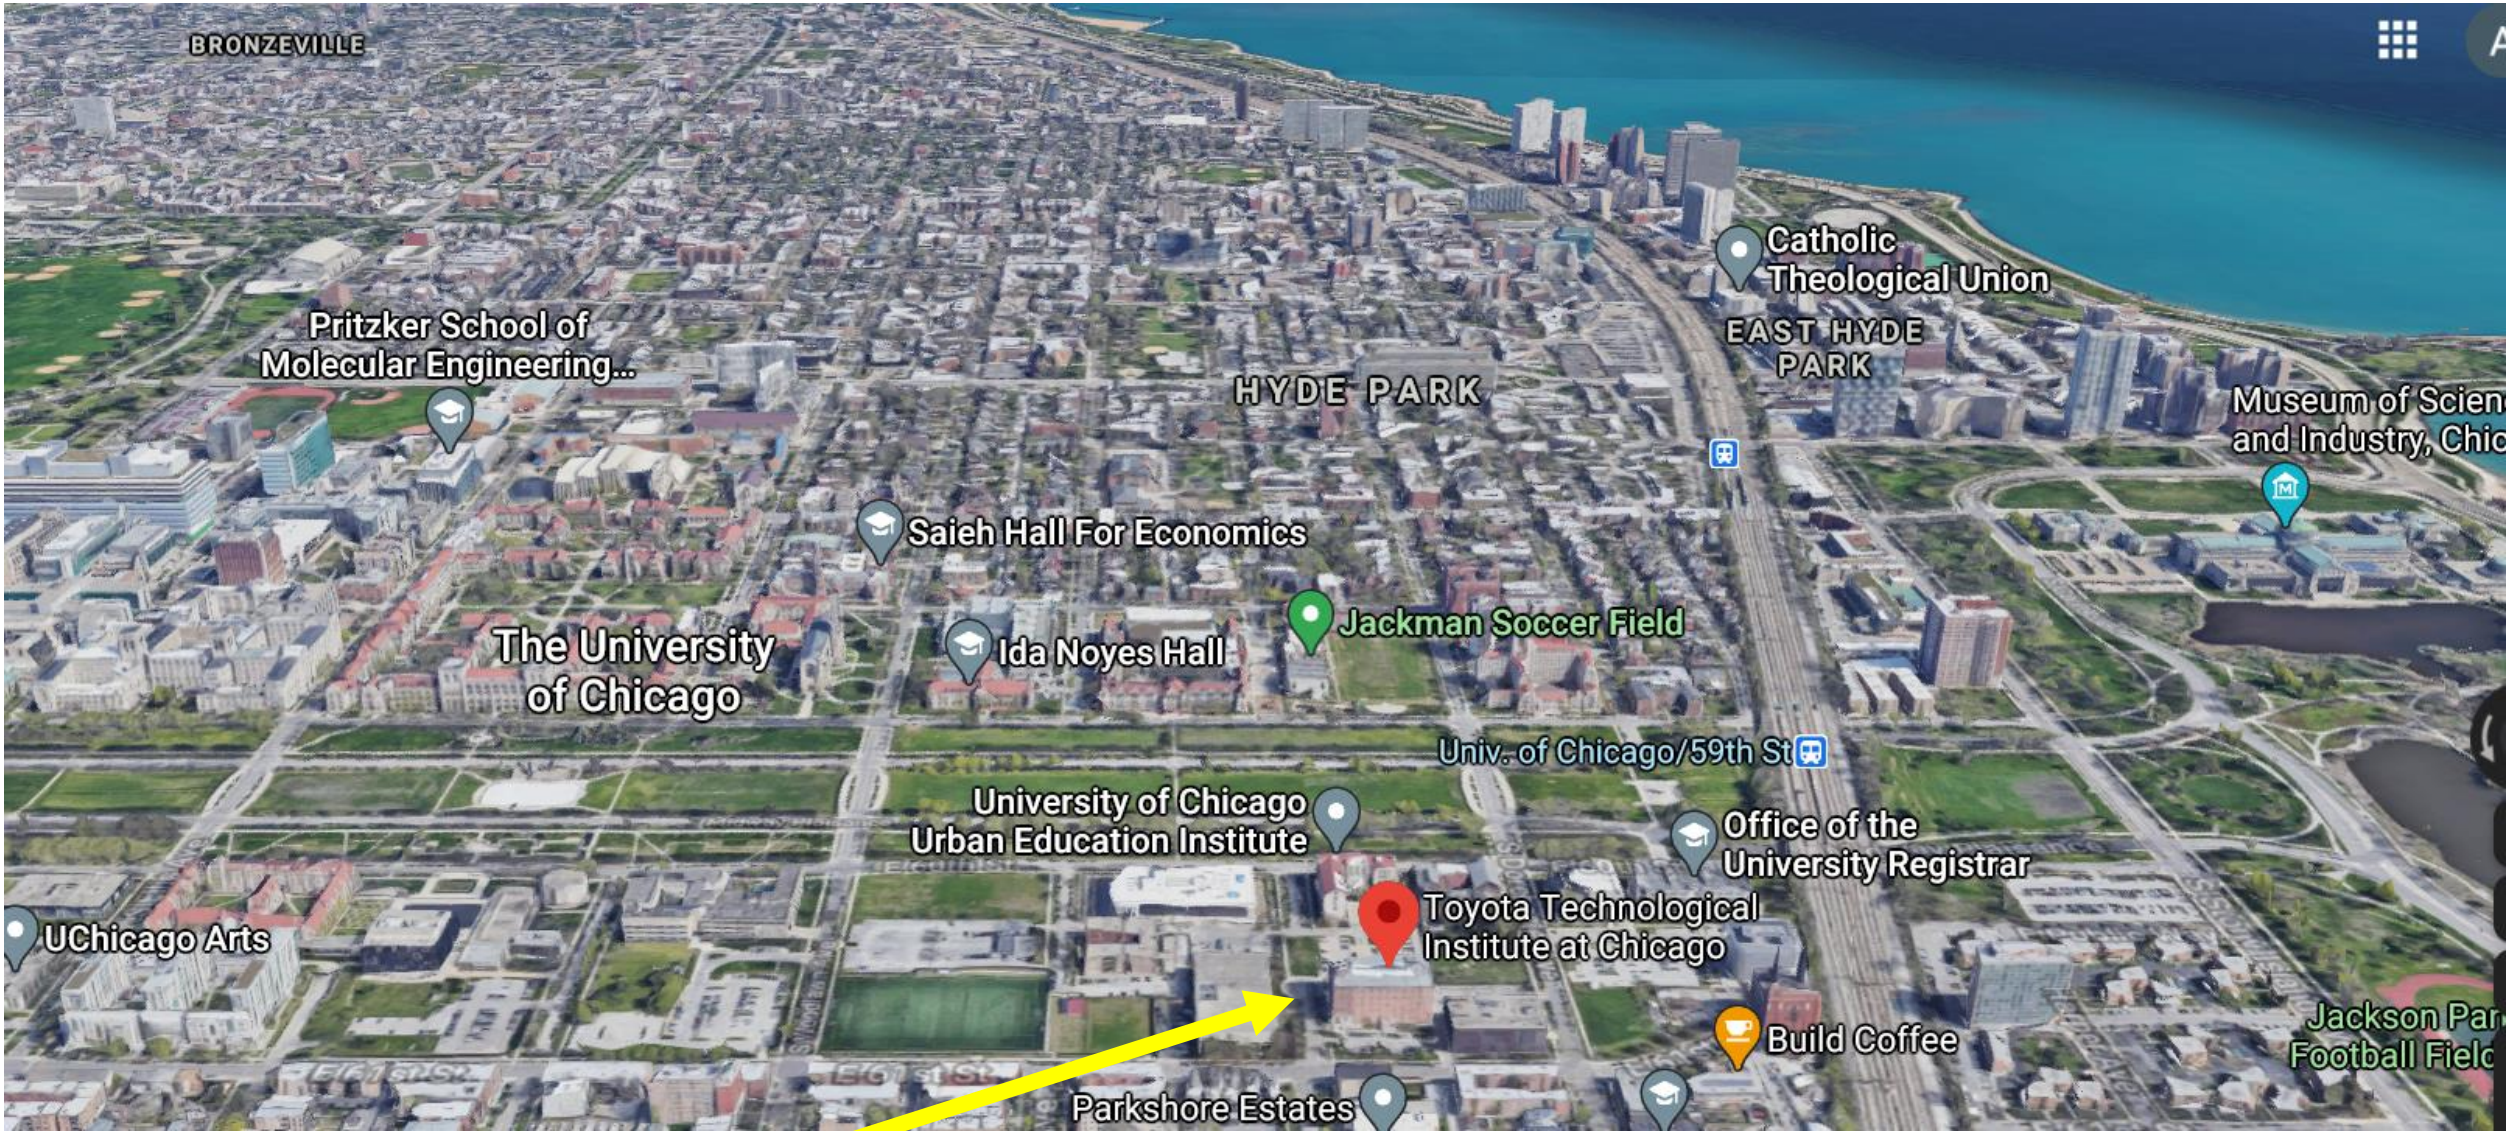


Chicago



Campus of the University of Chicago







Here

Our building







What makes TTIC great (for students)?

- Top faculty, comparable to the best universities in computer science in the world.
- Well-funded. Every student gets a competitive stipend, funded by internal TTIC funds or faculty research grants. Plus equipment funds, travel funds. Students only TA one or two quarters out of their entire PhD.
- Personal attention. 4:1 student-to-faculty ratio. Plus mechanisms to ensure every student has the best shot possible and no student “falls through the cracks”.
- Caring, engaged staff with TTIC’s goals of excellence in mind.
- Highly active research environment: colloquium speakers, reading groups, workshops, Research Assistant Professors, visiting/adjoint faculty, postdocs, and more

What makes TTIC great (for faculty)?

- Fantastic faculty colleagues and a friendly atmosphere, plus excellent folks nearby at UChicago, Northwestern, UIC, and more. [See [RAP testimonials](#)]
- Strong TTIC PhD students, with the ability to also seamlessly work with UChicago PhD, Masters, and Undergraduate students.
- Teaching load is only 1 course/year (tenure-track faculty) or 0 courses/year (RAPs)
- Well-funded, with significant internal research funds in addition to external grants
- Caring, engaged staff with TTIC's goals of excellence in mind.
- Highly active research environment: colloquium speakers, reading groups, workshops, Research Assistant Professors, visiting/adjoint faculty, postdocs, and more



TTIC Alumni

- Faculty at Cornell, UIUC, Tsinghua University, U Edinburgh, Yale, U Waterloo
- Postdocs at top universities/labs (MIT, IAS, ETH, Brown, Caltech, Columbia, IBM)
- Researchers at top labs and startups (Google, Microsoft, Meta, Amazon, Adobe, Tesla, Genentech, Toyota, Zendar, Waymo)
- Two graduates named Sloan Research Fellows. 2020 International Society of Computational Biology (ISCB) Overton Prize.

Recent Student Awards

- Best Paper Award at the 34th Annual Conference on Learning Theory (COLT 2021) and *both* Best Student Paper awards at COLT 2019.
- Best Short Paper award at the 4th Workshop on Computational Models of Reference, Anaphora, and Coreference, 2021.
- 2019, 2020, 2021 Google PhD Fellowships
- 2017, 2019, 2021 NSF Fellowships
- 1st Place at Max Planck Institute for Intelligent Systems Real Robot Challenge. EUR 3500 cash prize 😊

TTIC Tenure-Track Faculty



[Avrim Blum](#)




[Julia Chuzhoy](#)



[Kevin Gimpel](#)
(on leave)



[Shiry Ginosar](#)
(joining Fall 2024)



[Zhiyuan Li](#)
(joined Fall 2023)



[Karen Livescu](#)



[Yury Makarychev](#)



[David McAllester](#)



[Greg Shakhnarovich](#)



[Nathan Srebro](#)



[Madhur Tulsiani](#)



[Matthew Turk](#)



[Matthew Walter](#)



[Jinbo Xu](#)

TTIC Research Assistant Professors (RAPs)



Siddharth Bhandari



Anand Bhattad



Sam Buchanan



Emily Diana



Jungo Kasai



Hongyuan Mei



Theodor Misiakiewicz



Derek Reiman



Liren Shan



Saeed Sharifi-Malvajerdi



Ohad Trabelsi



Ali Vakilian



Santhoshini Velusamy

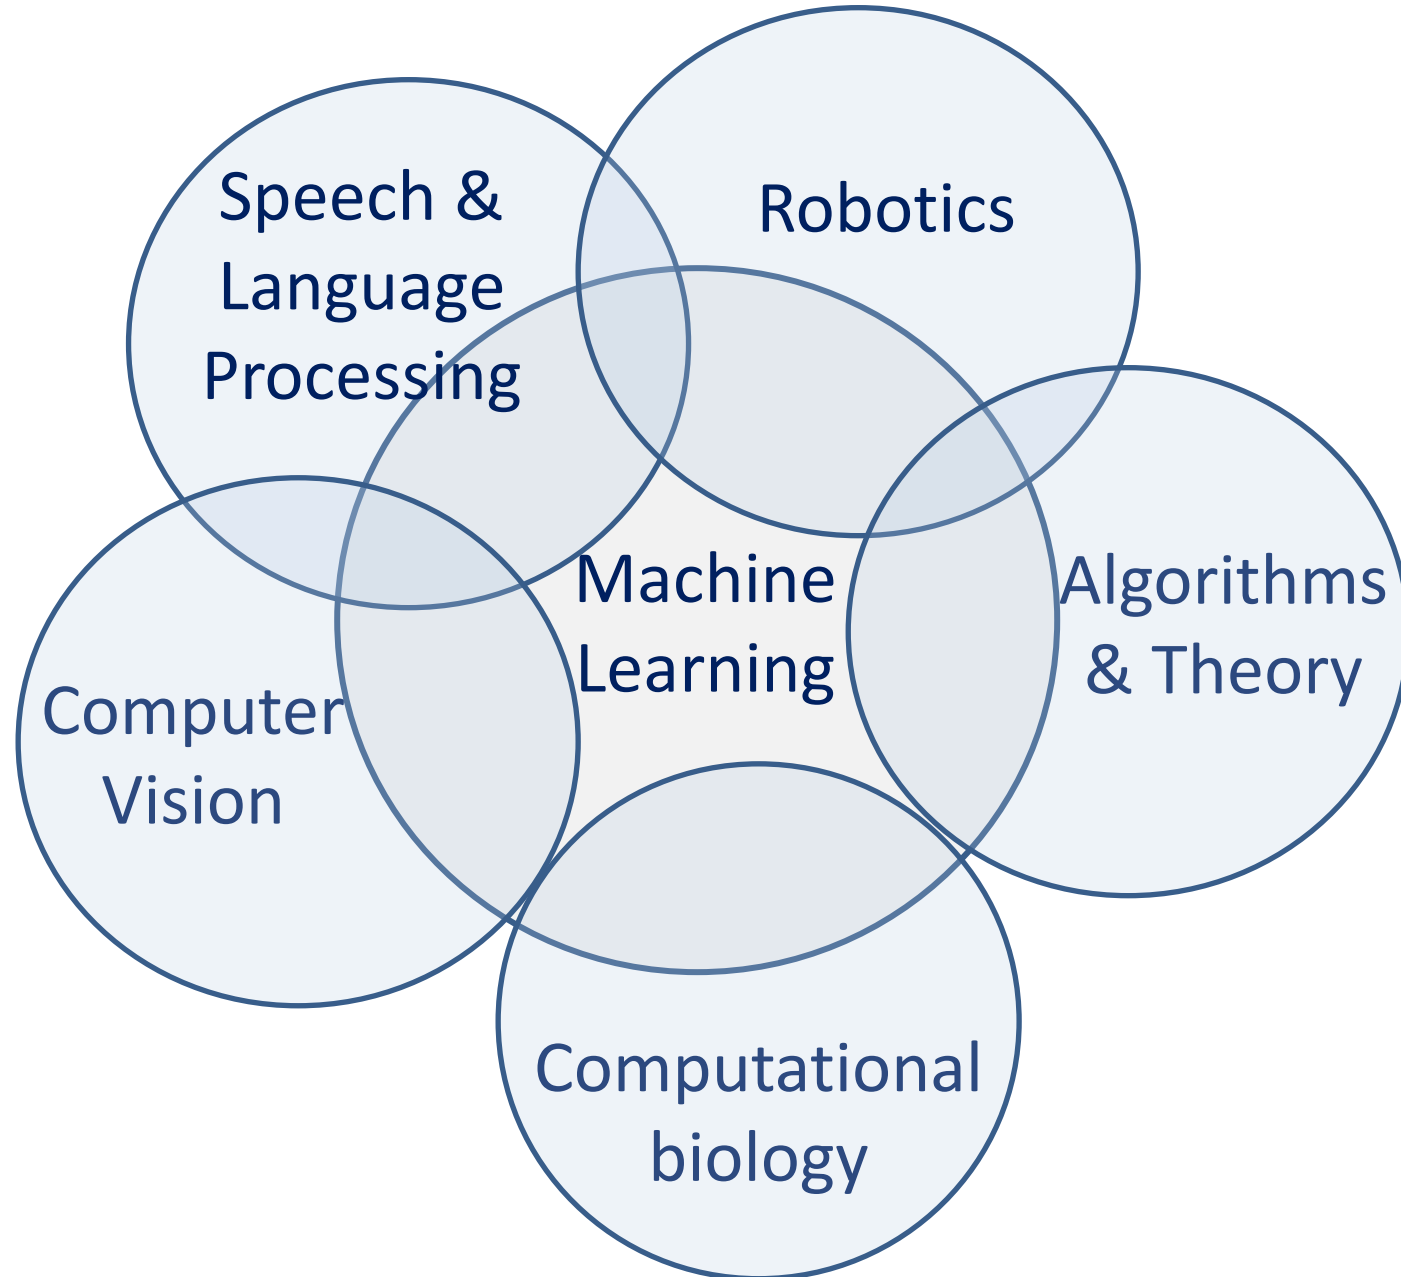


Lingxiao Wang



Jiawei (Joe) Zhou

Research Directions




Avrim Blum



- Mathematical analysis of machine learning (Machine Learning Theory)
- Fairness, adversarial robustness
- Game-theoretic / incentive-aware algorithms

$$\mathcal{L}_{S'}(\text{STOP}) = (\text{stop}, \text{CERTIFICATE})$$






2020 National Academy of
Sciences Michael and Sheila
Held Prize

Julia Chuzhoy



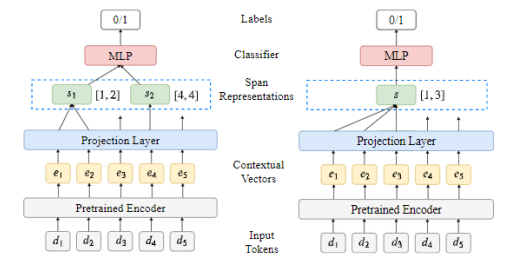
- Approximation algorithms for combinatorial optimization problems
- Network design and routing, dynamic graph algorithms (e.g., APSP), flow and cut problems, other graph algorithms
- Limits of approximation



Kevin Gimpel



- Natural Language Processing
- Automatically generating paraphrases, summarizing text
- Systems that can understand text the way people do and answer questions about it
- Methods for representing text and world knowledge



(a) Probing model for two-span tasks. This model can be used to decide whether two spans (here [1, 2] and [4, 4]) are coreferent or not.


(b) Probing model for single-span tasks. This model can be used to decide whether a span (here [1, 3]) refers to a constituent or not.

Shiry Ginosar


(joining Fall 2024)



- Computer perception, including vision, audio, and other sensory modalities
- Artificial social intelligence, including social learning, Human-Human, and Human-AI interaction
- Gesture recognition, learning to respond like a good listener



Zhiyuan Li



- Machine Learning Theory
- Analysis of Deep Learning
- Computational and sample efficiency of optimization methods


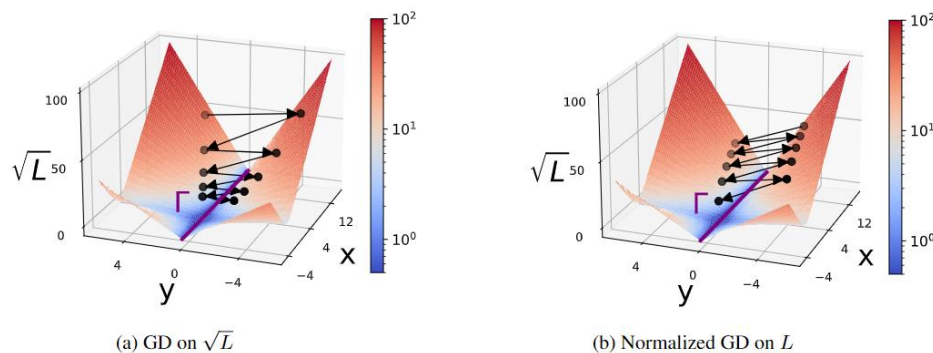
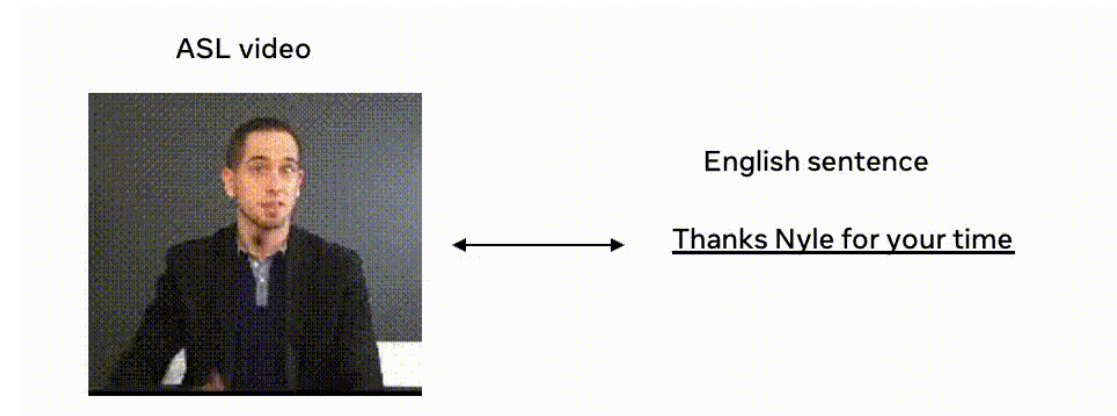


Figure 1: SGD+WD optimizes the scale invariant training loss of SIBERT robustly for all initialization scales, and thus for loss scalings and different learning rates (with $\lambda\eta$ fixed). Here the default initialization for parameters in SIBERT encoder is a truncated normal distribution with standard deviation equal to 0.02 (the same as BERT).

Karen Livescu





- Speech and language processing, speech recognition, acoustic+semantic representations of words and text.
- Statistical modeling techniques that can take advantage of both data and prior knowledge.
- Connections with computer vision, including sign language recognition, semantic speech retrieval



Yury Makarychev



- Combinatorial optimization
- Non-worst-case analysis of algorithms
- High-dimensional geometry
- Clustering and dimensionality reduction



David McAllester



- AI: automated reasoning and knowledge representation
- Machine learning
- Natural language processing
- Deep networks and big questions



Greg Shakhnarovich



- Computer vision and machine learning
- Automatic understanding of visual scenes
- Self-supervision
- Vision and language, sign language recognition

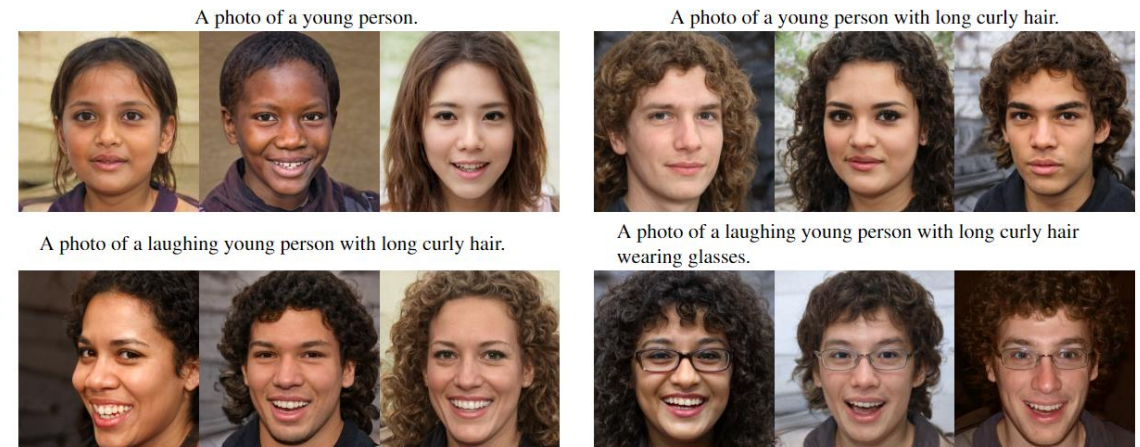



Figure 1. Examples of text-driven face image synthesis by our proposed method *Fast text2StyleGAN*. The text prompts are increasingly more detailed. Each image takes about 0.09s to produce. See https://github.com/duxiaodan/Fast_text2StyleGAN.git

Nati Srebro



- Statistical machine learning algorithms
- Understanding implicit bias of optimization methods
- ML algorithms that satisfy fairness, non-discrimination
- Scalable / distributed learning



Madhur Tulsiani



- Error-correcting codes. E.g., locally-decodable error-correcting codes for high noise rates
- Solvability of constraint satisfaction problems
- Analysis of semidefinite programming techniques
- High-dimensional expanders and their algorithmic applications



Matthew Turk





- Computer vision and machine learning
- Augmented and mixed reality
- Societal/ethical concerns in AI (e.g., facial recognition systems)
- TTIC president



Matthew Walter



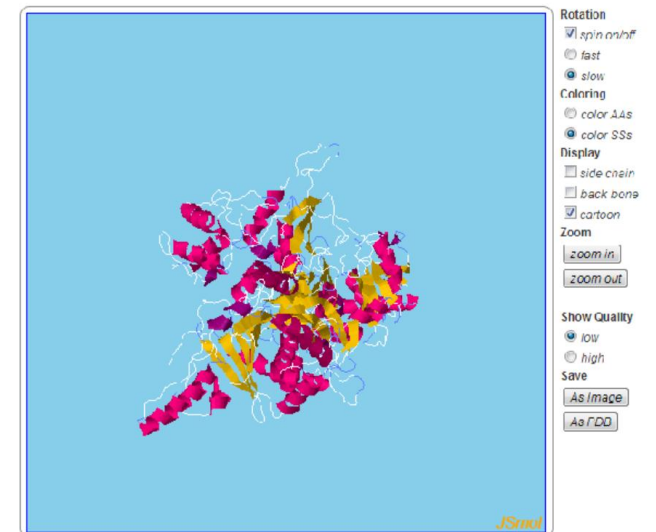
- Developing intelligent, perceptually aware robots that are able to work effectively with and alongside people in unstructured environments
- Robot-human interaction
 - Deciding What to Talk About & How
- Jointly optimizing design and control



Jinbo Xu



- Machine learning and data-driven methods for computational biology
- First successful deep learning algorithm for protein folding / protein structure prediction (RaptorX)
- Convex optimization for protein-protein interaction network alignment



Courses (all PhD-level)

Autumn 2023

[TTIC 31230 - Fundamentals of Deep Learning \(CMSC 31230\)](#) - Instructor: [David McAllester](#)

[TTIC 31080 - Approximation Algorithms \(CMSC 37503\)](#) - Instructor: [Julia Chuzhoy](#)

[TTIC 31150 - Mathematical Toolkit \(CMSC 31150\)](#) - Instructor: [Madhur Tulsiani](#)

[TTIC 31190 - Natural Language Processing](#) - Instructor: [Jiawei Zhou](#) and [Freda Shi](#)

Winter 2024

[TTIC 31010 - Algorithms \(CMSC 37000\)](#) - Instructor: [Yury Makarychev](#)

[TTIC 31020 - Introduction to Machine Learning](#) - Instructor: [Nati Srebro](#)

[TTIC 31070 - Convex Optimization \(...\)](#) - Instructor: [Zhiyuan Li](#)

[TTIC 31050 - Introduction to Bioinformatics and Computational Biology](#) - Instructor: [Jinbo Xu](#)

Spring 2024

[TTIC 31260 - Algorithmic Game Theory](#) - Instructor: [Avrim Blum](#)

[TTIC 31180 - Probabilistic Graphical Models](#) - Instructor: [Matt Walter](#)

Activities

- External talks, including [colloquium](#) and [distinguished lecture series](#)
- Internal [Research@TTIC talk](#) + lunch every Friday
- Reading groups
- [Workshops](#) (usually 4 each summer on ML / AI / Theory topics, plus various during the year)
- Variety of outreach activities



Chicago



Chicago



Chicago



Chicago



More information at <http://www.ttic.edu>