



education

phd | machine learning

uc berkeley | '17-present
research: interpretable ml
advisor: bin yu
gpa: 3.95

ms | data science

uc berkeley | '17-present
research: statistical biology

bs | cs & math

university of virginia | '14-'17
double major: cs, math
concentration: statistics

skills

machine learning

pytorch • scikit-learn • tensorflow
aws ec2 • s3 • sagemaker
keras • mllib • caffe2

languages

experienced
python • java • matlab
proficient
r • c/c++ • web basics
human languages
english • spanish • hindi

awards

berkeley grad slam semifinalist '19
outstanding teaching award '18
uva rader research award '17
uva undergrad symposium winner '17
raven honor society '16-'17
icpc regional qualification '14-'16
1st place microsoft code jam '16
3rd place google games uva '17
2nd place apt puzzle competition '17
intermediate honors '16
dean's list '14-'17

funding awards

pdsoros fellowship finalist '19
ircn workshop travel award '19
vidya shelat fund award '16
rodman scholarship '14-'17

experience

berkeley | ml research (bin yu lab ☺)

fall '17 - present

- developed interpretation methods for ml models (e.g. neural nets)
- created interpretable models in medicine, biology, and computer vision

amazon | research internship (pietro perona lab ☺)

summer '20

- testing for bias with causal matching using GANs
- interpreting semantic directions in generative models

response4life | volunteer data scientist

spring '20

- helped develop, integrate, and deploy models to forecast covid-19 severity

pacmed ai | interpretable ml internship

summer '19

- developed techniques to interpret machine-learning models for healthcare
- integrated interpretability techniques for predicting icu re-admission

facebook | computer vision internship

summer '17

- investigated unsupervised deep learning for segmentation of satellite imagery
- implemented crfs for segmentation post-processing

uva | ml research (yanjun qi lab ☺)

fall '16 - spring '17

- developed novel weighted- ℓ_1 , multi-task gaussian graphical model
- analyzed large-scale functional brain connectivity with graphical models

hhmi | ml research (srini turaga lab ☺)

summer '15, winter '15, summer '16

- extended cnns and watershed algorithms for neural image segmentation
- implemented distributed random forests for image segmentation

uva | comp. neuroscience research (william levy lab ☺)

fall '14 - fall '16

- developed detailed biophysical models of neural computation
- analyzed energy efficiency, noise, and variability in stochastic neurons

hhmi scientific computing | comp. neuroscience research

summer '14

- analyzed backpropagating action potentials via biophysical simulations

research innovations inc. | web dev + android internship

summer '13 - spring '14

- developed web/mobile app for task coordination with qr codes

coursework

computation

machine learning
computer vision
structure learning
algorithms
artificial intelligence
deep learning
learning theory
ai in graphics
cs theory
data structures
software dev. I & II
information retrieval
computer architecture

stat/math

statistical models
probability
statistics
optimization
linear algebra
info theory
real analysis
linear models
stochastic processes
chaos theory I & II
multivariate calculus
discrete mathematics
differential equations
abstract algebra

neuroscience

neural coding
neural network models
neurobiology
visual neuroscience
cognitive science

teaching

berkeley | student instructor

summer 2018

cs 189/289: machine learning
(lectures to class of 80+ students)

fall 2019

cs 188: artificial intelligence

papers

interpretability

- interpretations are useful: penalizing explanations to align neural networks with prior knowledge: rieger, singh, murdoch, & yu, *icml 2020*
- transformation importance with applications to cosmology: singh*, ha*, lanusse, boehm, liu & yu, *iclr 2020 workshop (spotlight talk)*
- hierarchical interpretations for neural network predictions: singh*, murdoch*, & yu, *iclr 2019*
- interpretable machine learning: definitions, methods, and applications: murdoch*, singh*, kumbier, abbasi-asl, & yu, *pnas 2019*
- disentangled attribution curves for interpreting random forests and boosted trees: devlin, singh, & yu *arXiv 2019*

interpretable data science projects

- curating a covid-19 data repository and forecasting county-level death counts in the united states: altieri et al. *harvard data science review 2020 (accepted)*
- an interpretable clinical-decision rule for intra-abdominal injury *in prep*
- predicting successful clathrin-coated pits in clathrin-mediated endocytosis via auxilin *in prep*

statistical neuroscience

- large scale image segmentation with structured-loss-based deep learning for connectome reconstruction: funke et al. *tpami 2018*
- a weighted- ℓ_1 , multi-task graphical model with applications to heterogeneous brain connectivity: singh, wang, & qi, *neurips 2017 amlcd workshop*
- linearized synaptic integration at no extra cost: morel, singh, & levy, *journal of computational neuroscience 2018*
- a consensus layer V pyramidal neuron can sustain interpulse-interval coding: singh & levy *plos one 2017*

machine learning misc

- revisiting complexity and the bias-variance tradeoff: dwivedi*, singh*, yu, & wainwright *arXiv 2020*
- benchmarking bias in facial recognition via matching using neural networks: singh, balakrishnan, & perona 2020 *in submission*

selected talks

- interpreting ml models: *uc berkeley bair seminar, 2020*
- uncovering brain connections underlying autism via graphical models: *tom tom founder's machine learning conference, 2017*

projects / activities

notes, blog, & slides '14-'20
imodels package '18-
covid-19 forecasting '20
hummingbird tracking '18
news balancer django app '17
madison house volunteering '15-'16
java mini-games '14-'16

academic service

neurips ml4h reviewer '20
basis education volunteering '19-'20
bair undergrad mentoring '18-'20