


STA414/STA2104

Statistical Methods for Machine Learning II



January 2022

David Duvenaud & Michal Malyska

Motivating Questions:

- How to fill in missing parts, conditioned on a description?



“zebras roaming in the field”



“a girl hugging a corgi on a pedestal”

GLIDE: Towards Photorealistic Image Generation and Editing with Text-Guided Diffusion Models

Motivating Questions:

- How to generate images from text?



“a high-quality oil painting
of a psychedelic hamster
dragon”



“robots meditating in a
vipassana retreat”

Motivating Questions:

- Which measurement would give the most useful information about a patient?
- How can I handle missing data?
- How can I figure out how good a player is from their wins and losses?
- How to combine all information automatically without designing a new algorithm for each type?

Answer:

- Fit joint distributions over pixels / words / measurements
 - Answer queries by marginalizing and conditioning
- Only need to:
 - Specify meaningful models (graphical models)
 - Integrate over states (approximate inference)
 - Fit lots of parameters (gradients and neural nets)

Lecture structure:

- **First 2 hours:** Going through concepts, mostly matching course notes. Medium pace.
- **Third hour:** 1 hour tutorial on same content. Worked examples. Slow pace. Lots of time for questions. Feel free to skip / leave.

Today

- Course information and overview
 - Expectations, course structure, evaluations
 - Learning objectives for course
- Overview of probabilistic machine learning
 - Examples
 - Tools of the trade
- Graphical model notation

Learning Outcomes: Today

- Know what topics are and aren't in the course.
- An idea of if you have the background + how hard the material will be.
- What you should be able to do with this knowledge.
- Basics of representing conditional independence of probability distributions with graphs.

Scope of course

- Designing, fitting, and interpreting parametric probabilistic models.
 - Conditioning, marginalizing, Normalized versus unnormalized distributions, Graphical models
 - Neural nets, gradient-based optimization, automatic differentiation
 - Approximate inference, sampling, variational inference
- A bit of decision theory
- Standard software tools: numerics, autodiff, git

Evaluation plan (tentative)

- **Assignment 1:** Brute force Bayes, Naive Bayes image completion. Due around Feb 6th
- **Assignment 2:** Exact inference and MCMC in univariate + Trueskill. Due around Feb 20th
- **Midterm:** 20%: Graphical models, conditioning, sampling, fitting distributions. Feb 28 and 29
- **Assignment 3:** Stochastic variational inference in Trueskill. Due around March 20th
- **Assignment 4:** Fitting neural net generative models (variational autoencoders and maybe one other). Due around April 3rd.
- **Final Exam:** 30% Some math about distributions, conceptual questions about fitting models,

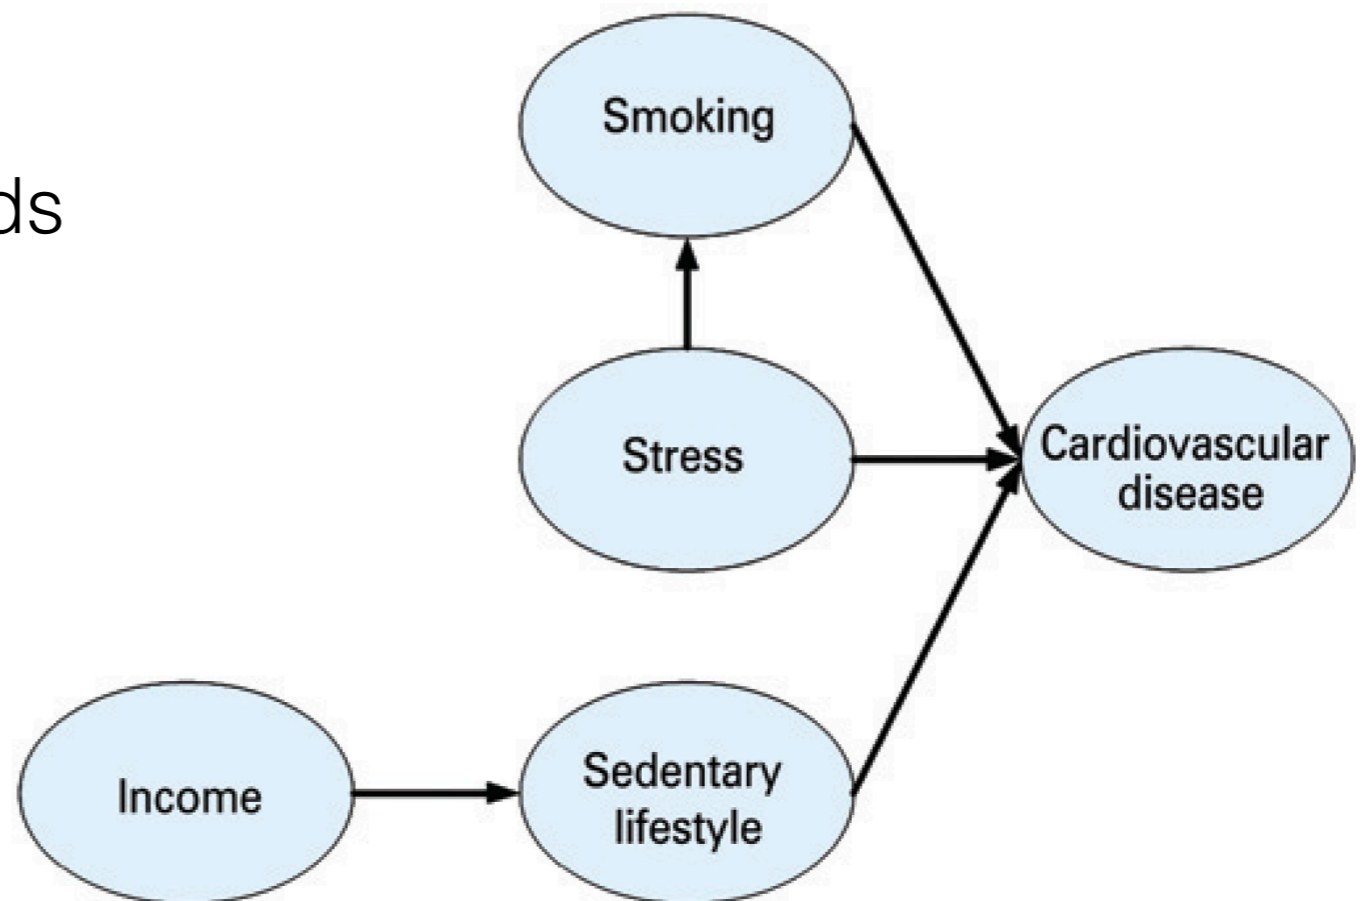
Tools of the trade: Probability

- Probabilities represent uncertainty about a fixed but unknown quantity, conditioned on some information
- Inference and prediction is easy!:
"Just write down the joint probability of everything, and integrate out everything you don't know." - MacKay
- No need to pretend to identify parameters, except for computational efficiency




Tools of the trade: Graphical Models

- Large joint distributions are expensive to reason about
- Conditional independence avoids combinatorial explosions when summing
- Can reason about conditional independence using graphs
- Fun algorithms
- Old-fashioned, since often simpler to assume everything is connected




$$p(i, sm, st, se, c) = p(i)p(se | i)p(s)p(sm | s)p(c | sm, se, st)$$




[1]

Gaussian mixture model




[2]

Linear dynamical system




[3]

Hidden Markov model




[4]

Switching LDS




[5]

Mixture of Experts




[2]

Driven LDS




[6]

IO-HMM




[7]

Factorial HMM



[8,9]

Canonical correlations analysis



[10]

admixture / LDA / NMF

[1] Palmer, Wipf, Kreutz-Delgado, and Rao. Variational EM algorithms for non-Gaussian latent variable models. NIPS 2005.

[2] Ghahramani and Beal. Propagation algorithms for variational Bayesian learning. NIPS 2001.

[3] Beal. Variational algorithms for approximate Bayesian inference, Ch. 3. U of London Ph.D. Thesis 2003.

[4] Ghahramani and Hinton. Variational learning for switching state-space models. Neural Computation 2000.

[5] Jordan and Jacobs. Hierarchical Mixtures of Experts and the EM algorithm. Neural Computation 1994.

[6] Bengio and Frasconi. An Input Output HMM Architecture. NIPS 1995.

[7] Ghahramani and Jordan. Factorial Hidden Markov Models. Machine Learning 1997.

[8] Bach and Jordan. A probabilistic interpretation of Canonical Correlation Analysis. Tech. Report 2005.

[9] Archambeau and Bach. Sparse probabilistic projections. NIPS 2008.

[10] Hoffman, Bach, Blei. Online learning for Latent Dirichlet Allocation. NIPS 2010.

Courtesy of Matthew Johnson

Tools of the trade: Neural Networks


- Not profound or especially mysterious: Just a large nonlinear parametric functions.
- Can fit anything if we overparameterize enough and use gradients.
- Main issues: Overfitting, non-differentiable objectives, hard to debug
- Show autograd demo



Source: xkcd

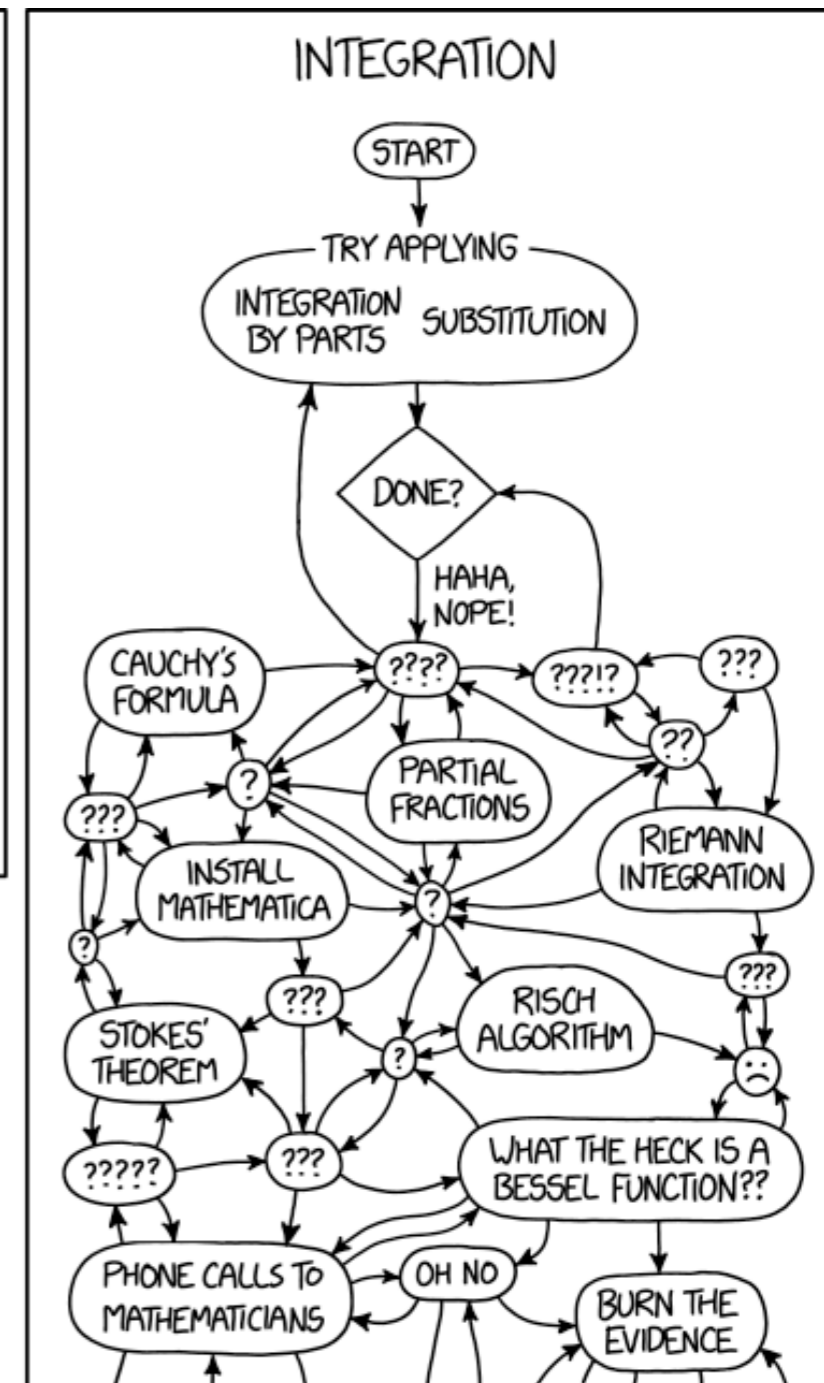
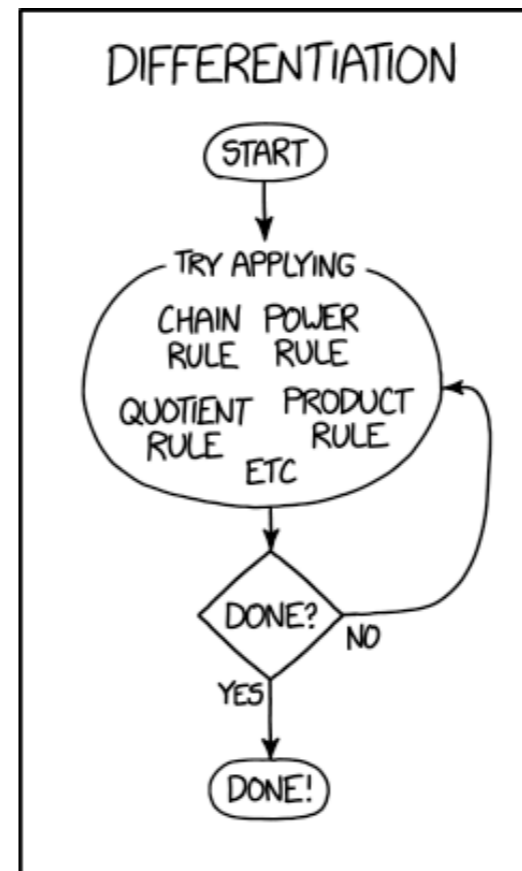
Tools of the trade: Gradient-based optimization

- Unconstrained, high-dimensional, stochastic, first-order gradient descent is surprisingly applicable
- Hinton: SGD "works much better than anyone had any right to expect".
- More parameters -> more progress before getting stuck.



Tools of the trade: Automatic Differentiation

- Reverse-mode grads have same asymptotic time cost as original function
- Biggest change in last 10 years of ML practice
- Vector-Jacobian products are cheap



Source: xkcd


Tools of the trade: Approximate Inference

- Have $p(h, d) = p(h)p(d | h)$

want samples from



$$p(h | d) = \frac{p(h)p(d | h)}{\sum_{h'} p(h')p(d | h')}$$

- Gradient based methods:
 - Variational inference
 - MCMC




What can you build with these tools?

- Naive Bayes, Mixture of Gaussians, Logistic Regression, Bayesian Linear Regression, Hidden Markov Models, Factor Analysis
- Neural network classifiers, LSTMs, RNNs, Transformers, Convnets, Neural ODEs, Deep Equilibrium models
- Variational Autoencoders, Generative Adversarial Networks, Normalizing Flows, Diffusion models





What can do with these models?

- Extend existing models.
 - E.g. What if we know the age of only some of our users?
- Sanity check data
 - E.g. Which piece of data in this form is most surprising?



What's not in scope?

- Statistical Learning Theory. See e.g. Dan Roy or Murat Erdogdu's grad courses.
- Fancy neural network architectures. See e.g. Roger Grosse's or Jimmy Ba's courses
- Logic-based AI, reasoning, discrete search. See e.g. Sheila McLiraith or Faheim Baccus.
- Discrete Nonparametrics, e.g. Indian Buffet processes, Dirichlet processes. But will cover Gaussian processes!



Tools of the trade: Julia and Python

- Python: Allowed, but initially can't use frameworks' network layers, initializers, or optimizers.
 - Suggested: Jax, PyTorch
 - Gotchas: need to learn both Python and a framework on top. Bad error messages.
- Julia: Also a big community, decent autodiff
- Hard mode: Dex (github.com/google-research/dex-lang)
 - Barely working, ultra-experimental, dependently typed functional numerical language with state and algebraic effects. Use at your own risk!



Tools of the trade: Git



- Version control is table stakes for industry, collaboration, your own sanity.
- Github demos add a lot to a resume.
- Interface is awful, imo



GitHub

Intimidated?

- Machine learning involves math, but it's the same math over and over.
- HW0 will evaluate class's current level of knowledge. The student body is always right.
- Will provide starter code / skeleton for at least most of the assignments.
- We want you to succeed!

Extra Resources

- No required textbook. All tested material in lecture notes on website.
- David MacKay (2003) *Information Theory, Inference, and Learning Algorithms*. Great intro, dated on methods.
- Christopher M. Bishop (2006) *Pattern Recognition and Machine Learning*. Great intro, dated on methods.
- Kevin Murphy (2012), *Machine Learning: A Probabilistic Perspective*. Up-to-date, encyclopaedic.
- Trevor Hastie, Robert Tibshirani, Jerome Friedman (2009) *The Elements of Statistical Learning*
- Deep Learning (2016) Goodfellow, Bengio, Courville.

My Origin Story

- [https://
bayes.wustl.edu/etj/
prob/book.pdf](https://bayes.wustl.edu/etj/prob/book.pdf)
- Derives probability
from scratch
- Part manifesto

Probability Theory: The Logic of Science

by

E. T. Jaynes

Wayman Crow Professor of Physics

Washington University

St. Louis, MO 63130, U. S. A.

ML as a bag of tricks

Special cases:

- K-means
- Kernel Density Estimation
- Support Vector Machines
- Boosting
- Random Forests
- K-Nearest Neighbors

Extensible family:

- Mixture of Gaussians
- Latent variable models
- Gaussian processes
- Deep neural nets
- Bayesian neural nets
- Attention-based models

Regularization as a bag of tricks

Fast special cases:

- Early stopping
- Ensembling
- L2 Regularization
- Gradient noise
- Dropout
- Expectation-Maximization

Extensible family:

- Stochastic variational inference

AI as a bag of tricks

Russel and Norvig's
parts of AI:

- Machine learning
- Natural language processing
- Knowledge representation
- Automated reasoning
- Computer vision
- Robotics

Extensible family:

- Deep probabilistic latent-variable models + decision theory
- a.k.a. Model-based Reinforcement learning

Stats vs Machine Learning



- Statistician: Look at dataset, consider the problem, design an interpretable model
 - Want guarantees, few assumptions, explanations
- ML: Mostly only predictions. Must handle new data automatically.
 - No way around making assumptions. Just make model big enough, hopefully it includes something close to the truth.
 - Model needs to have a million parameters somewhere, reality is messy.
 - Can't use guarantees or bounds in practice, so empirically choose model details
- Probabilistic ML: Distinguish model from fitting algorithm

Statistical Learning


vs


Deep learning

- Conceptually, a lot going on, mathematically and algorithmically simpler




Examples











Probabilistic graphical models

- + structured representations
- + priors and uncertainty
- + data and computational efficiency
- rigid assumptions may not fit
- feature engineering
- top-down inference

Deep learning

- neural net “goo”
- difficult parameterization
- can require lots of data
- + flexible
- + feature learning
- + recognition networks

Deep generative models


- Model distributions implicitly by a variable pushed through a deep net:

$$y = f_{\theta}(x)$$

- Approximate intractable distribution by a tractable distribution parameterized by a deep net:

$$p(y|x) = \mathcal{N}(y|\mu = f_{\theta}(x), \Sigma = g_{\theta}(x))$$


- Optimize all parameters using stochastic gradient descent




$$c_i \sim \text{Cat}(\pi)$$




$$z_i \sim N(\mu_i, \Sigma_i)$$

$$x_i \sim N(f_\theta(z_i), g_\theta(z_i))$$






Modeling idea: graphical models on latent variables,
neural network models for observations



Composing graphical models with neural networks for structured representations and fast inference. Johnson, Duvenaud, Wiltchko, Datta, Adams, NIPS 2016






data space






latent space

unsupervised
learning



supervised
learning



Courtesy of Matthew Johnson

Types of Learning

- **Supervised Learning:** Given input-output pairs (x,y) the goal is to predict correct output y given a new input x .
- **Unsupervised Learning:** Given unlabeled data instances $x_1, x_2, x_3 \dots$ build a model of x , which can be used for making predictions, decisions.
- **Semi-supervised Learning:** We are given only a limited amount of (x,y) pairs, but lots of unlabeled x 's.
- All just special cases of estimating distributions from data: $p(y|x), p(x), p(x, y)$.

Image Infill

- Just sampling from $p(\text{missing pixels} \mid \text{remaining})$
- <https://www.youtube.com/watch?v=9V7rNoLVmSs>

Pluralistic Image Completion

NANYANG TECHNOLOGICAL UNIVERSITY SINGAPORE

Options: CelebA-HQ-random

Input: Original/Output Score: 0.023



StyleGAN2

- "Just" a big GAN with some training tricks + data preprocessing.
- Representation ends up being intuitive.
- <https://www.youtube.com/watch?v=c-NJtV9Jvp0&feature=youtu.be>

Analyzing and Improving the Image Quality of StyleGAN

Tero Karras
NVIDIA

Samuli Laine
NVIDIA

Miika Aittala
NVIDIA

Janne Hellsten
NVIDIA

Jaakko Lehtinen
NVIDIA and Aalto University

Timo Aila
NVIDIA



occluded

completions

original



Pixel Recurrent Neural Networks (2015)

Aaron van den Oord, Nal Kalchbrenner, Koray Kavukcuoglu

Image to Text



LZ
a car is parked in
the middle of nowhere .



a wooden table and chairs
arranged in a room .



there is a cat sitting on a shelf .




a ferry boat on a marina
with a group of people .



a little boy with a bunch
of friends on the street .

Text to Image






Captions

This is a large brown bird with a bright green head, yellow bill and orange feet.




With long brown upper converts and giant white wings, the grey breasted bird flies through the air.

A grey bird with a small head and short beak with lighter grey wing bars and a bright yellow belly.

Keypoints




Samples



A white large bird with orange legs and gray secondaries and primaries, and a short yellow bill.

- Parallel Multiscale Autoregressive Density Estimation. Reed et al., 2017


Text to Image




“a boat in the canals of venice”



“a painting of a fox in the style of starry night”



“a red cube on top of a blue cube”




“a stained glass window of a panda eating bamboo”



“a crayon drawing of a space elevator”



“a futuristic city in synthwave style”



“a pixel art corgi pizza”



“a fog rolling into new york”

- GLIDE, 2021

Sequential Data: Video

- Stochastic Video Generation with a Learned Prior. Emily Denton, Rob Fergus

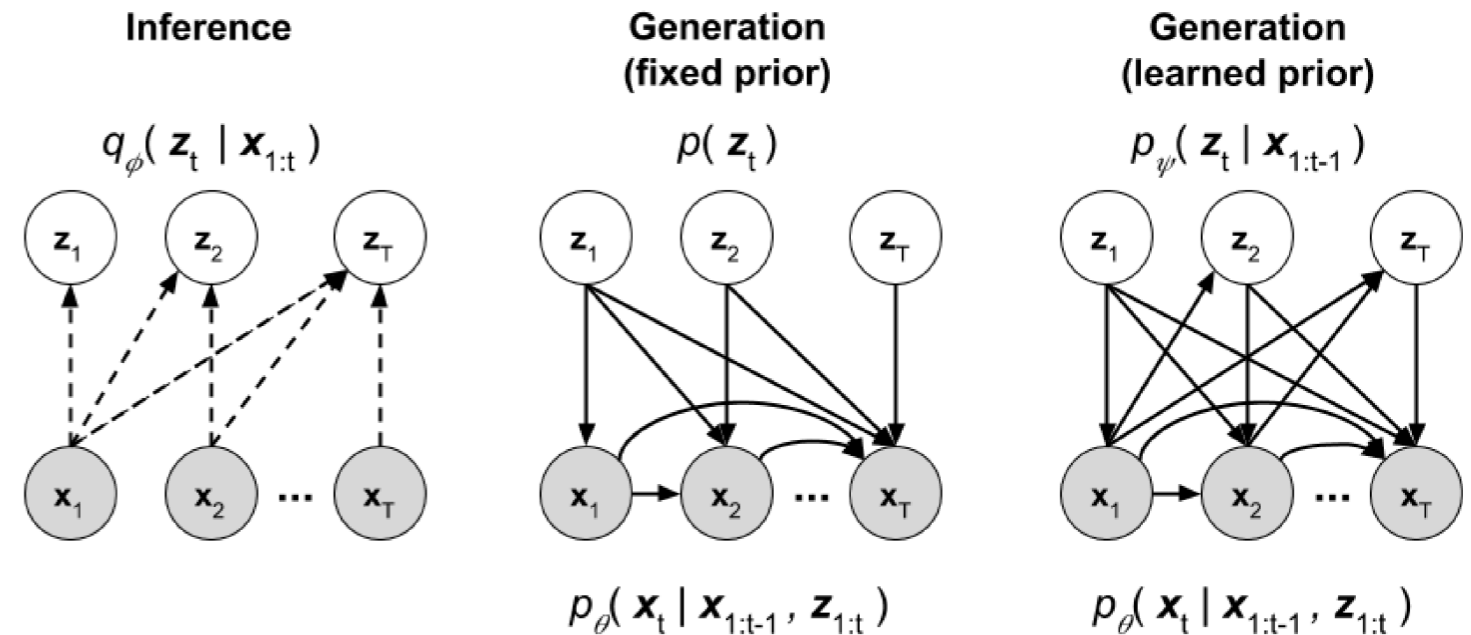

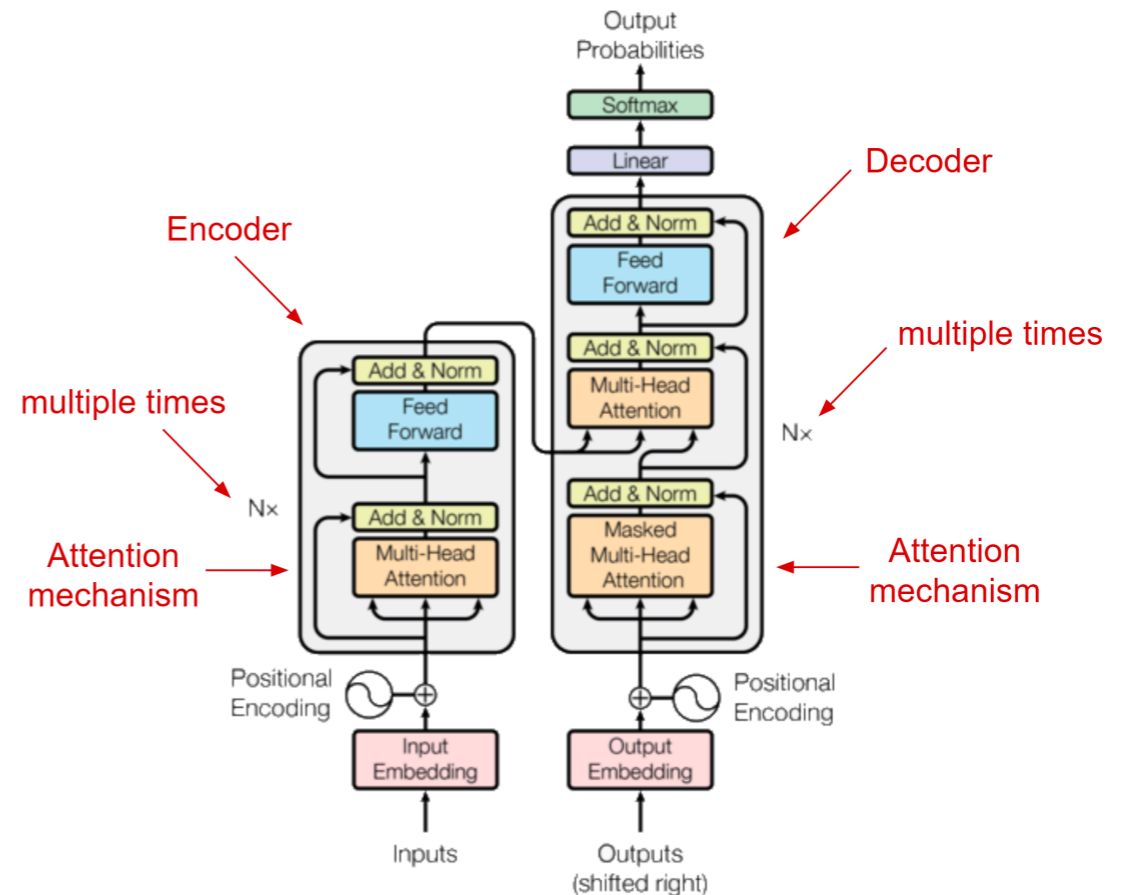


Figure 1. Inference (left) and generation in the SVG-FP (middle) and SVG-LP models (right).



Sequential Data: Text

- 'Attention Is All You Need'
Vaswani et al., 2017
- Variant of RNNs with attention,
aka key-query layers
- <https://talktotransformer.com/>



Grammar Variational Autoencoder

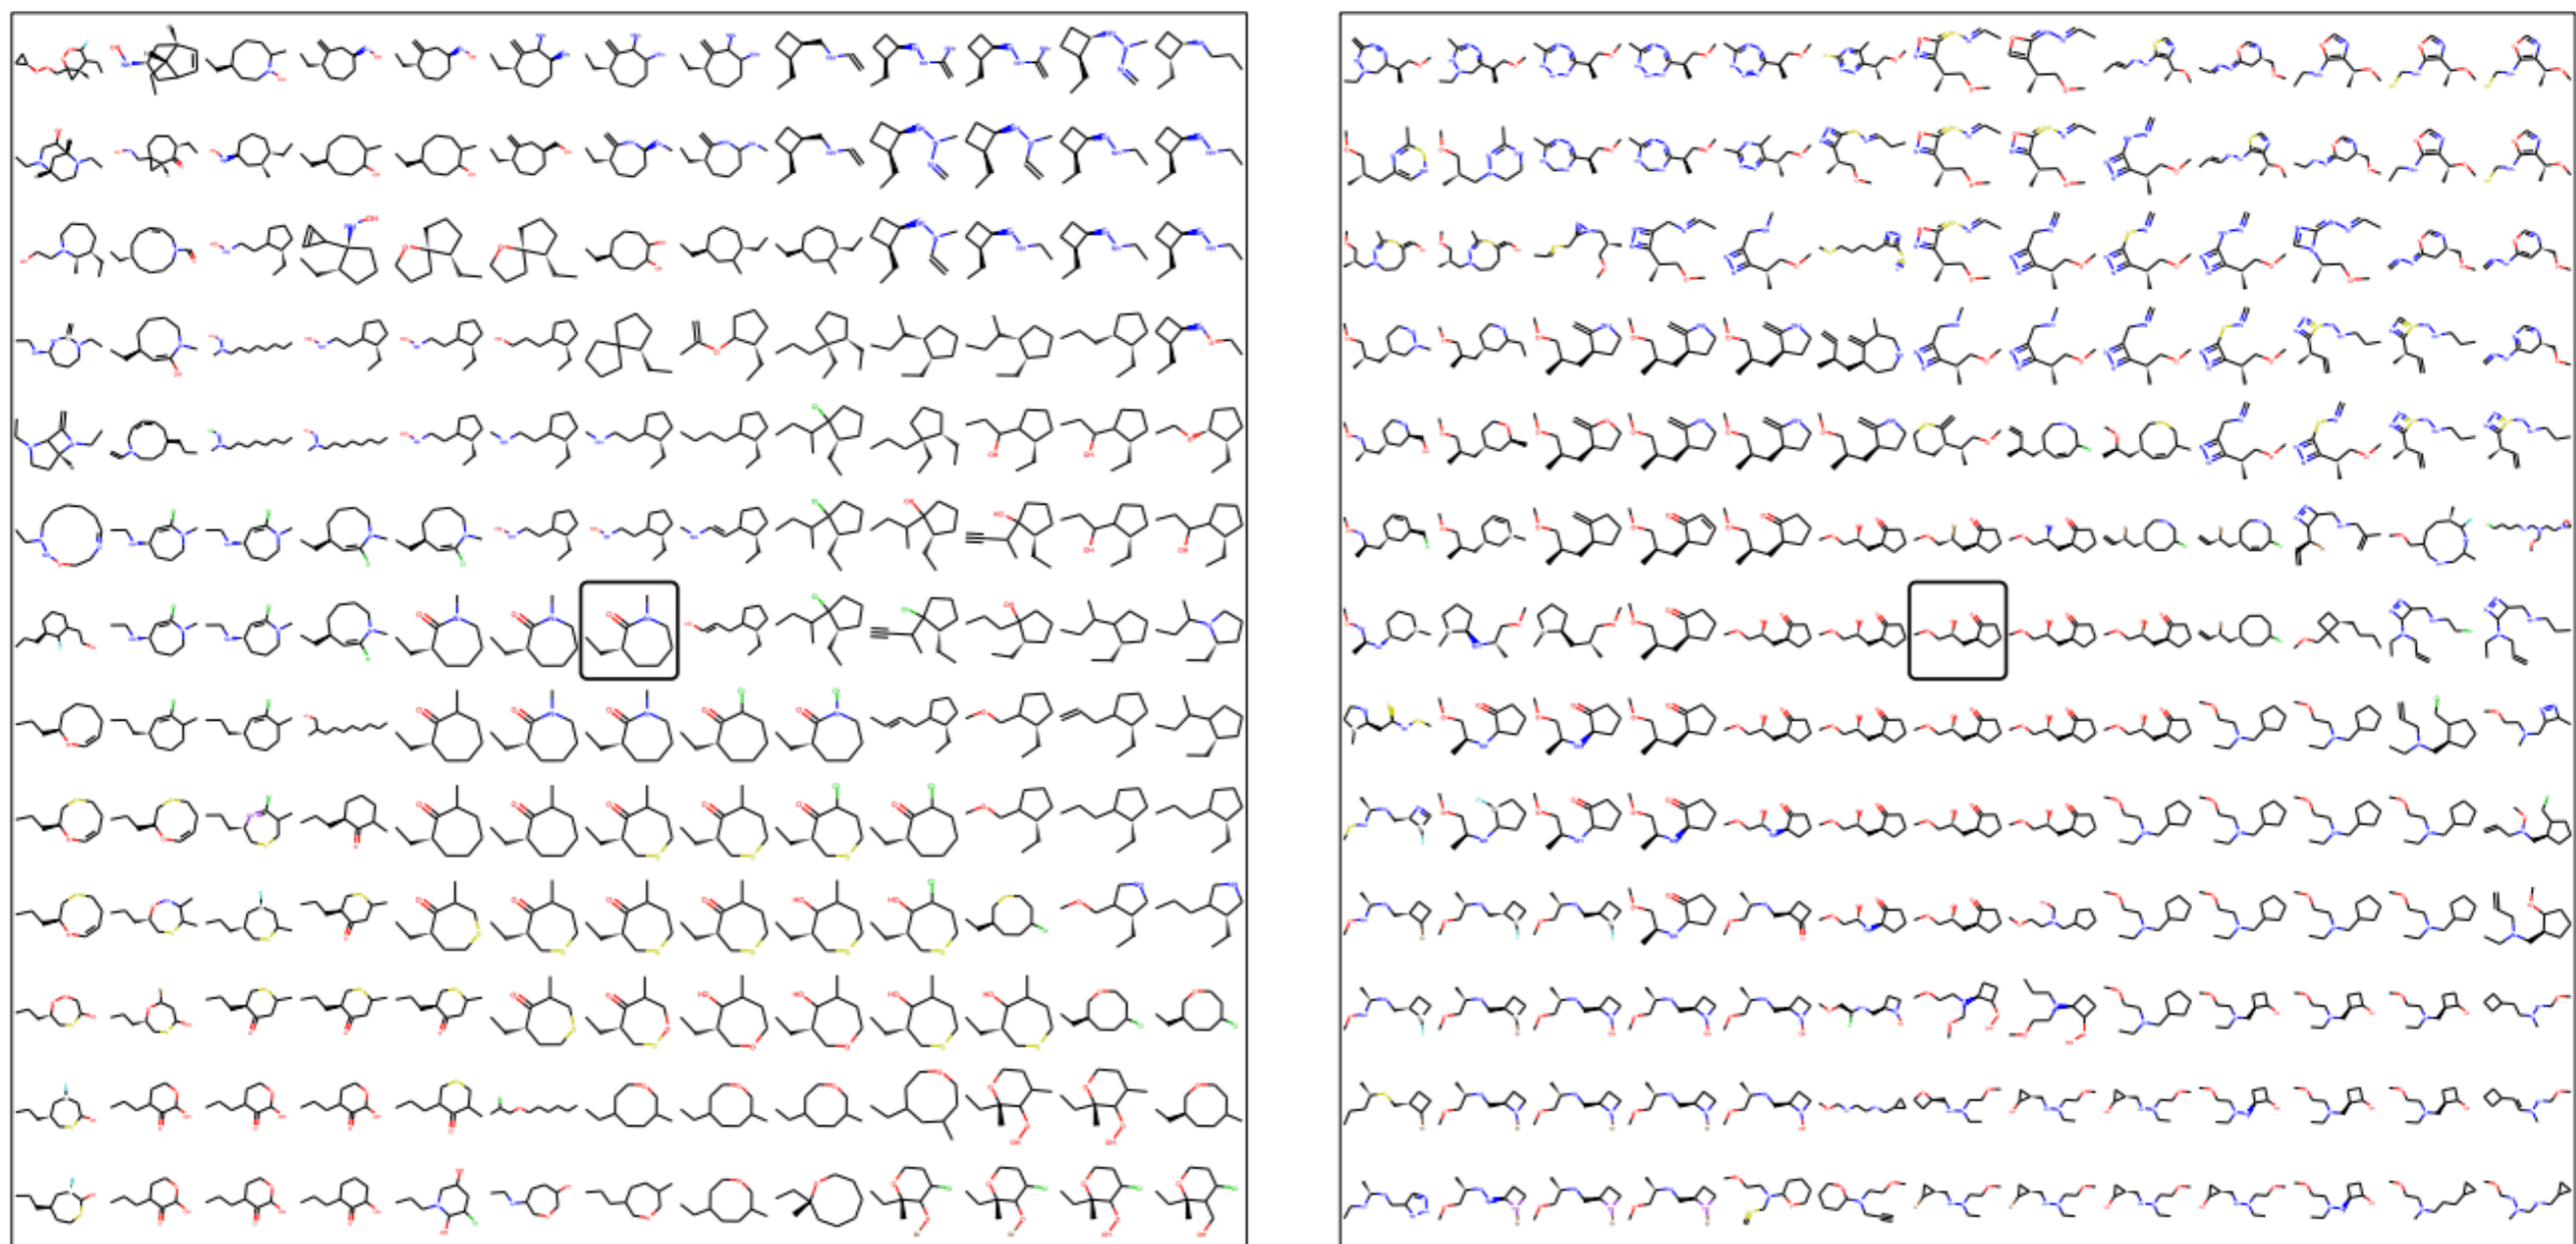


Figure 3. Searching the 56-dimensional latent space of the GVAE, starting at the molecule in the center.



Scientific Data

- Need to marginalize over all the things we don't know
- Graphical models + approximate inference let us combine all sorts of measurements into a single model

Modeling the photometric red clump

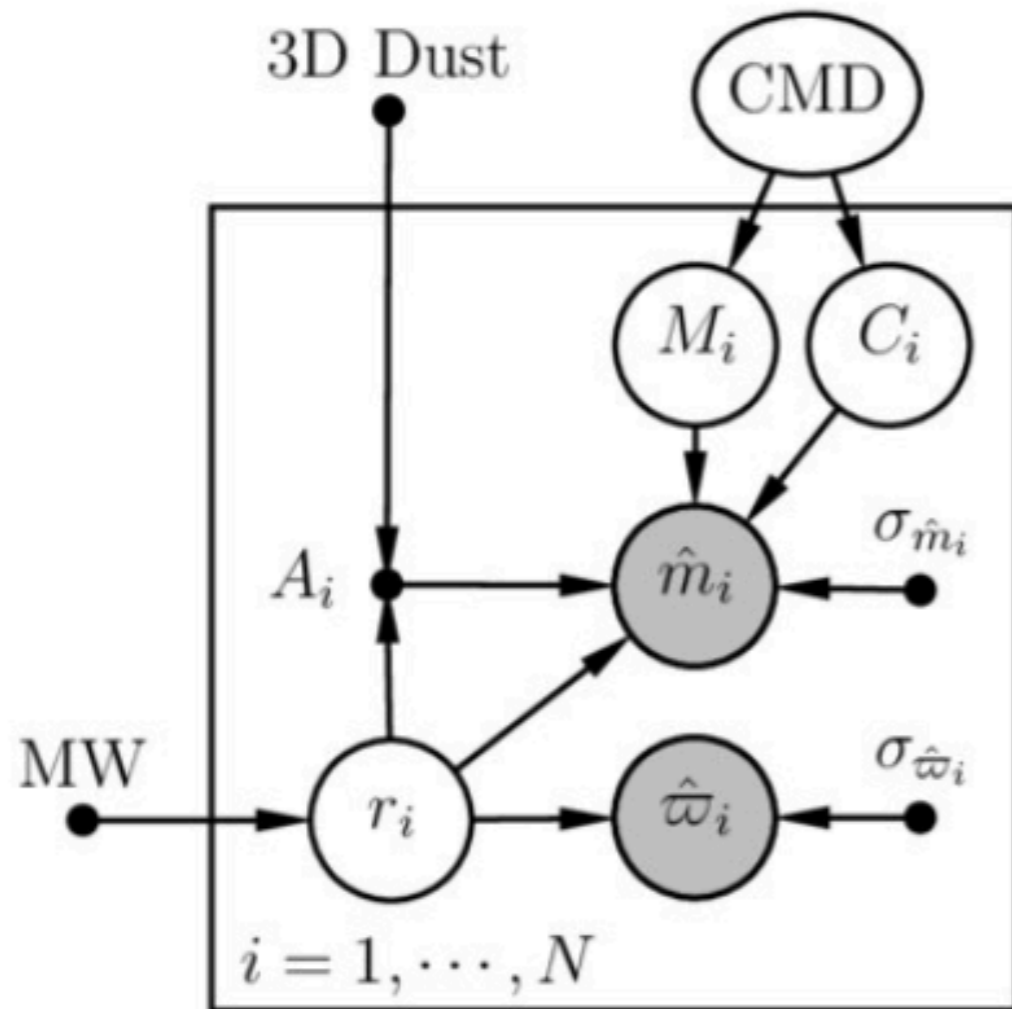
Use existing spectroscopic or astroseismic RC catalogs

Hierarchical probabilistic model: Gaussian for the RC + outliers, marginalizing over dust, parallaxes, observed magnitudes.



Model+MCMC with stan
Sample joint posterior

TGAS PGM and model



Absolute magnitude:

$$M_V = m_V - 5 \log_{10} \left(\frac{d}{10 \text{ pc}} \right)$$

Parallax & magnitude likelihoods:

$$p(\hat{\varpi} | d, \sigma_{\varpi}) = \mathcal{N}(\hat{\varpi} - 1/d; \sigma_{\varpi}^2)$$

$$p(\hat{\vec{m}} | d, \vec{C}, M, \Sigma_{\hat{\vec{m}}})$$

$$= \mathcal{N}(\hat{\vec{m}} - \vec{m}(d, \vec{C}, M); \Sigma_{\hat{\vec{m}}})$$

Posterior distribution now tying all objects together, with CMD

Variational Autoencoders

- One type of probabilistic latent variable model.
- We'll fit these in this class.
- Variations on these models are the current state of the art.



DiffusionCLIP: Text-Guided Diffusion Models for Robust Image Manipulation. 2021

Advantages of probabilistic latent-variable models

- **Data-efficient** - automatic regularization, can take advantage of more information
- **Composable** - e.g. incorporate data corruption model.
- **Handle missing or corrupted data** - no imputation, always integration.
- **Predictive uncertainty** - useful for decision-making.
- **Conditional predictions** - e.g. if brexit happens, the value of the pound will fall
- **Active learning** - What data would be expected to increase our confidence about a prediction?
- **Disadvantages:**
 - intractable integral over latent variables

Reasons to take this course

- Data science jobs.
- Getting into research in ML. (but it's a gold rush)
- Doing research in another area, but being able to build / tweak / question models. (recommended)
- Not being impressed by "it was done with deep learning / reinforcement learning / AI"

Collaboration policy

- Can work with up to 2 others on assignments
- Need to implement / write up your own answers
- Need to list collaborators

Learning Outcomes: Today

- Know what topics are and aren't in the course.
- An idea of if you have the background + how hard the material will be.
- What you should be able to do with this knowledge.
- Next: Graphical model notation.

Questions?