Deep Generative Models

Introduction

Hamid Beigy

Sharif University of Technology

March 1, 2024

Table of contents

Introduction

How do you understand complex and unstructured inputs?

Computer vision

Computational speech

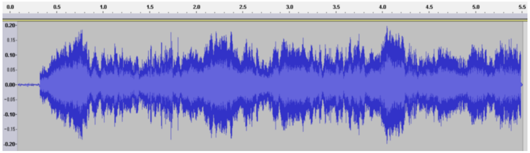

Natural language processing

Robotics

Credit: Aditya Grover

Generative and Discriminative Models

1. **Discriminative modeling** estimates the conditional distribution $p(\mathbf{y} \mid \mathbf{x})$.

2. **Generative modeling** estimates the joint distribution $p(\mathbf{x}, \mathbf{y})$.

Discriminative modeling

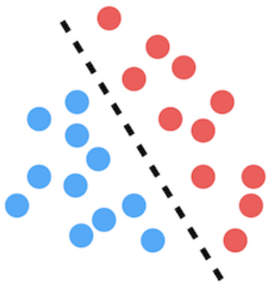

Generative modeling

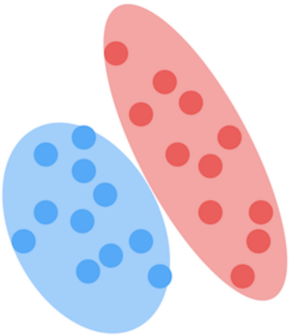

3. Without assuming $\mathbf{y}$, generative models learn $p(\mathbf{x})$ from given data.

4. $p(\mathbf{x})$ enables us to generate new data similar to the training dataset.

1. **A Generative model** (GM) is a probability distribution $p(\mathbf{x})$.
    - A statistical GM is a trainable probabilistic model, $p_\theta(\mathbf{x})$.
    - A deep GM is a statistical generative model parametrized by a neural network.
2. A generative model needs
    - Data ($\mathbf{x}$): Complex, unstructured samples such as images, speech, molecules, text, etc.
    - Prior knowledge: parametric form (e.g., Gaussian, mixture, softmax), loss function (e.g., maximum likelihood, divergence), optimization algorithm, etc.

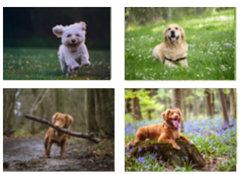

$\mathbf{x}_i \sim P_{\text{data}}$
$i = 1, 2, \ldots, n$

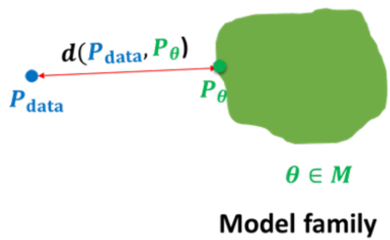

Credit: Aditya Grover

Key Questions

$x_i \sim P_{\text{data}}$
$i = 1, 2, \ldots, n$

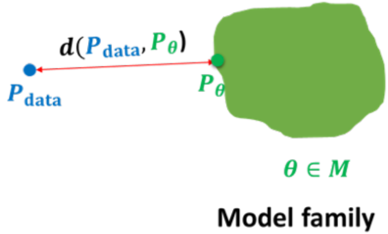

1. **A Representation:** how do we parameterize the joint distribution of many random variables?
2. **A Learning:** what is the right way to compare probability distributions?
3. **A Inference:** how do we invert (or encode) the generation process?

Credit: Aditya Grover

Deep Generative Models

Progress in Deep Generative Models

Credit: Ian Goodfellow, 2019

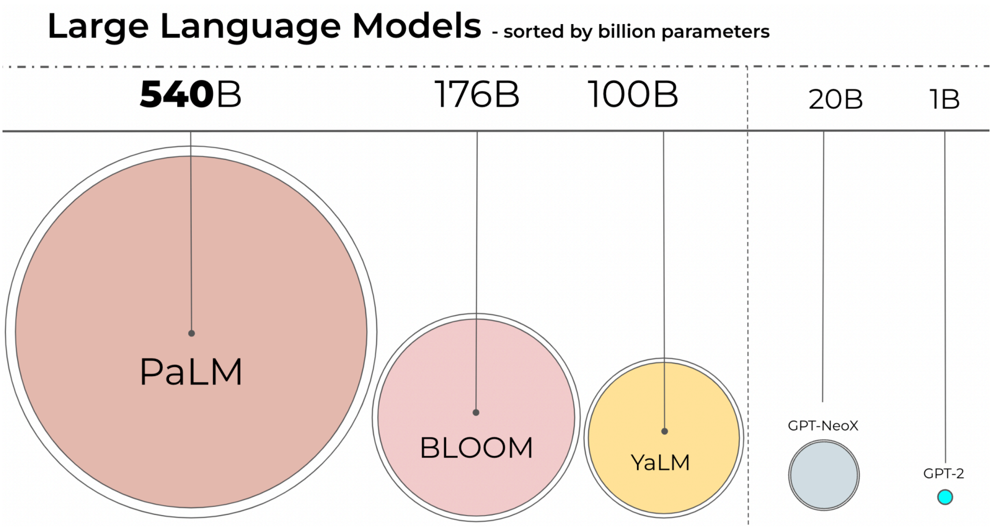
Large Language Models - sorted by billion parameters
540B
176B
100B
20B
1B
PaLM
BLOOM
YaLM
GPT-NeoX
GPT-2

Sample Applications of Generative Models

Image generation

Text prompt: an armchair in the shape of an avocado ...

Language Generation

Text prompt: to get an A+ in deep generative modes, students should have to

Completion

To get an A+ in deep generative models, students have to be willing to work with problems that are a whole lot more interesting than, say, the ones that most students work on in class. If you're a great student, the question above can be avoided and you'll be able to do great work, but if you're not, you will need to go beyond the basics before getting good.

Now to be clear, this advice is not just for the deep-learning crowd; it is good advice for any student who is taking his or her first course in machine learning.

The key point is that if you have a deep, deep brain of a computer scientist, that's just as important to you.

Credit: Aditya Grover

Applications of Deep Generative Models

Generating Data

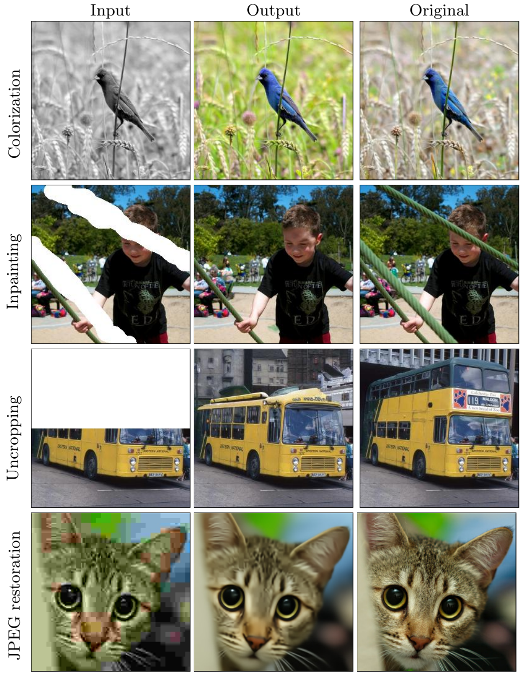

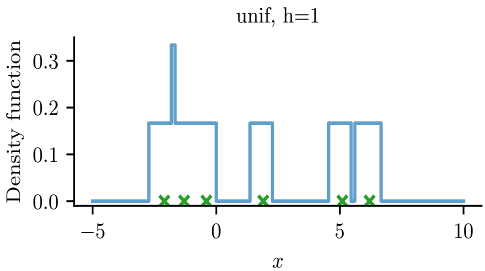
unif, h=1
Density function
0.0
0.1
0.2
0.3
−5
0
5
10
x

(a)

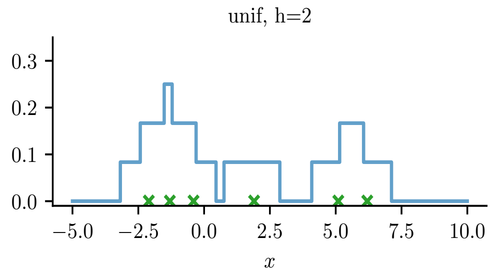
unif, h=2
0.0
0.1
0.2
0.3
−5.0
−2.5
0.0
2.5
5.0
7.5
10.0
x

(b)

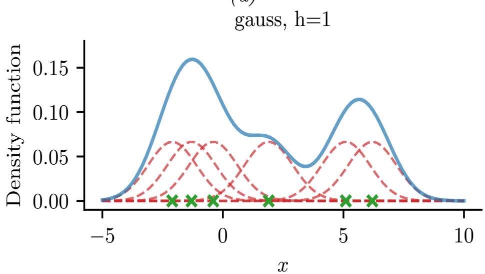
gauss, h=1
Density function
0.00
0.05
0.10
0.15
−5
0
5
10
x

(c)

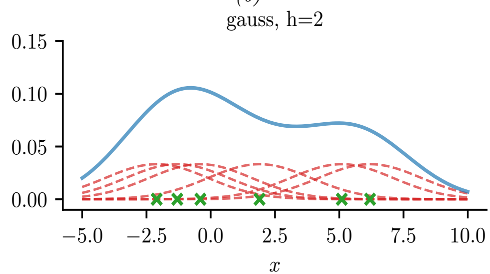
gauss, h=2
0.00
0.05
0.10
0.15
−5.0
−2.5
0.0
2.5
5.0
7.5
10.0
x

(d)

Imputation (Filling in missing values)

Variables

Data sample	A	B	C
1	6	6	NA
2	NA	6	0
3	NA	6	NA
4	10	10	10
5	10	10	10
6	10	10	10
Average	***9***	***8***	***7.5***

Missing values replaced by means

A	B	C
2	6	**7.5**
9	6	0
9	6	**7.5**
10	10	10
10	10	10
10	10	10
9	*8*	*7.5*

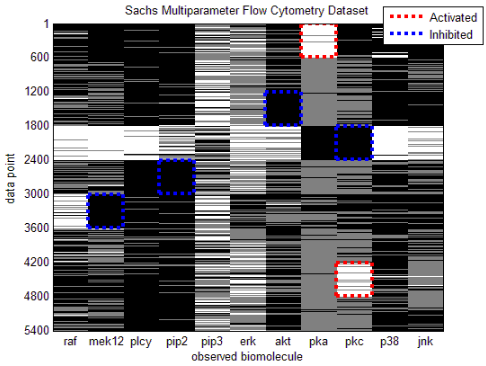
Sachs Multiparameter Flow Cytometry Dataset
Activated
Inhibited
data point
1
600
1200
1800
2400
3000
3600
4200
4800
5400
raf mek12 plcy pip2 pip3 erk akt pka pkc p38 jnk
observed biomolecule

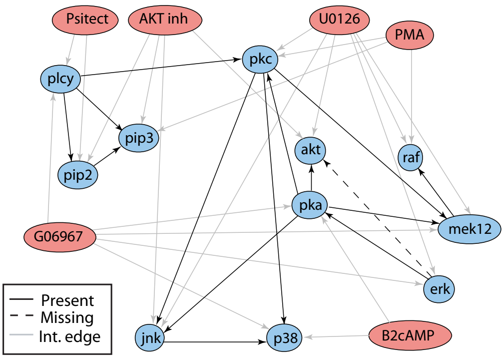
Psitect
AKT inh
U0126
PMA
pkc
plcy
pip3
akt
raf
pip2
pka
mek12
G06967
erk
jnk
p38
B2cAMP
Present
Missing
Int. edge

Latent Space Arithmetic

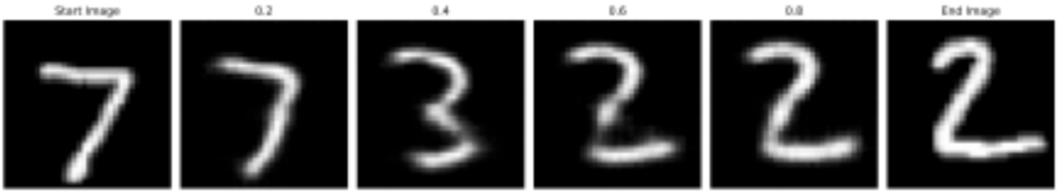

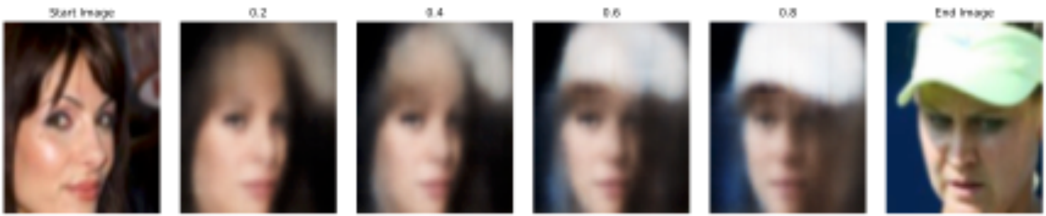

Input pattern

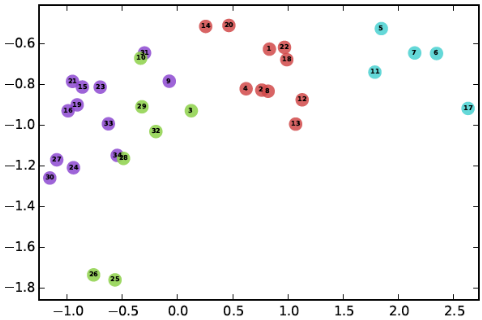

Representation

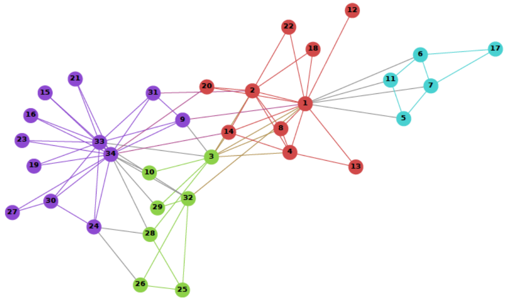

Course Information

Course Information

1. Course name : **Deep Generative Model**
2. Instructor : Hamid Beigy   Email : `beigy@sharif.edu`
3. Class : CE 202
4. Virtual class link: `https://vc.sharif.edu/beigy`
5. Course Website: `http://sharif.edu/~beigy/14022-40957.html`
6. Lectures: Sat-Mon (10:30-12:00)
7. Teaching Assistant : Mohaddeseh Mirbeygi   Email: `m.mirbeygi@gmail.com`

Course overview

Course overview

1. Introduction
2. Structured density
3. Disentangled Representation Learning
4. Generative adversarial network
5. Flow-based models
6. Variational auto-encoder
7. Autoregressive models
8. Energy-based models
9. Diffusion models
10. Hybrid models
11. Evaluation of generative models
12. Differential privacy
13. Causal representation learning
14. Causal generative models
15. Other topics

Course evaluation

- Evaluation:

Mid-term exam	20%	**1403-01-25**
Final exam	20%	
Homeworks	35%	
Quiz	15%	
Paper	10%	**Hard deadline for paper selection: 1403-01-25**
Class activity	5%	

Main references

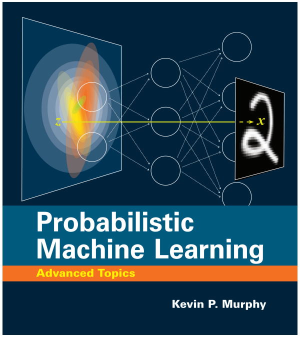

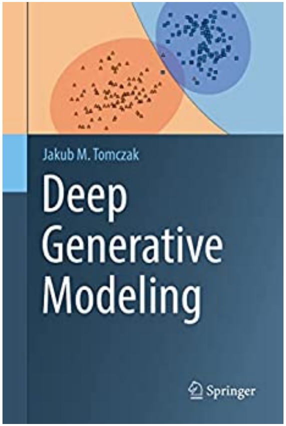

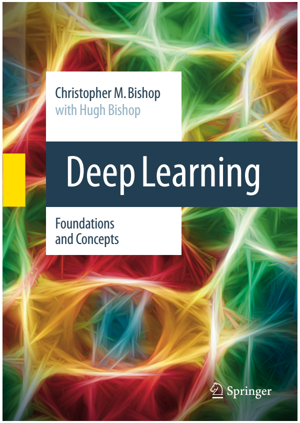

- Bishop, Christopher M. and Hugh Bishop (2024). *Deep Learning: Foundations and Concepts*. Springer.
- Murphy, Kevin P. (2023). *Probabilistic Machine Learning: Advanced Topics*. The MIT Press.
- Tomczak, Jakub M. (2022). *Deep Generative Modeling*. Springer.

Several research papers will be used as references in the class.

References

Reading

1. Chapter 20 of Probabilistic Machine Learning: Advanced Topics (Murphy 2023).
2. Chapter 1 of Deep Generative Modeling (Tomczak 2022).

References

- Bishop, Christopher M. and Hugh Bishop (2024). *Deep Learning: Foundations and Concepts*. Springer.
- Murphy, Kevin P. (2023). *Probabilistic Machine Learning: Advanced Topics*. The MIT Press.
- Tomczak, Jakub M. (2022). *Deep Generative Modeling*. Springer.

Questions?