MONASH University

Generative Deep Learning

Giới thiệu mô hình tự sinh dữ liệu hiện đại

Phùng Quốc Định, PhD

Email: dinh.phung@monash.edu
dinh@trustingsocial.com
Published under: Dinh Phung

GROUP OF EIGHT AUSTRALIA

Outline

- What is a generative model?
- Why do we need 'generative' and approaches
- Why generative deep learning now?
  - Deep learning in 2-cent
  - Two motivations for deep generative models
    - Curse of dimensionality
    - Reversing a deep NN
  - Generative Adversarial Networks
- What's next?

Outline

- What is a generative model?
- Why do we need ‘generative’ and approaches
- Why generative deep learning now?
  - Deep learning in 2-cent
  - Two motivations for deep generative models
    - Curse of dimensionality
    - Reversing a deep NN
  - Generative Adversarial Networks
- What’s next?

What is a generative model?

BEACH

What is a generative model?

BEACH

y : class, label, response, outcome, ..

x : feature, covariate, ..

What is a generative model?

BEACH

y : class, label, response, outcome, ..

x : feature, covariate, ..

Cúm

{Đau đầu, sổ mũi, …}

What is a generative model?

BEACH

y : class, label, response, outcome, ..

x : feature, covariate, ..

Naïve Bayes Model

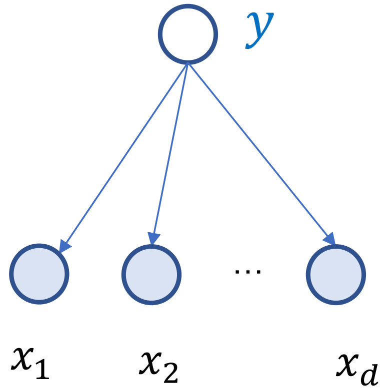

Cúm

{Đau đầu, sổ mũi, …}

MONASH University

What is a generative model?

BEACH

y : class, label, response, outcome, ..

x : feature, covariate, ..

Naïve Bayes Model

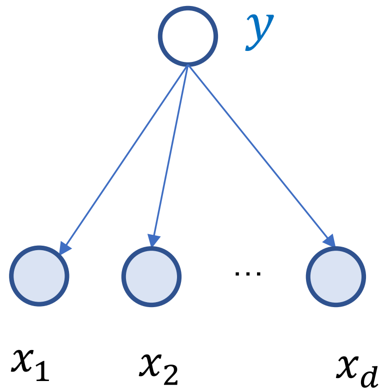

$$P(x, y) = P(y) \prod_{i=1}^{d} P(x \mid y)$$

Cúm

{Đau đầu, sổ mũi, ...}

- Generative model: models joint distribution for all involved: $p(x, y)$

What is a generative model?

BEACH

y : class, label, response, outcome, ..

x : feature, covariate, ..

Naïve Bayes Model

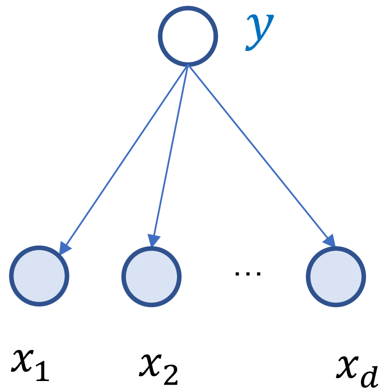

$$P(x, y) = P(y) \prod_{i=1}^{d} P(x \mid y)$$

Cúm

{Đau đầu, sổ mũi, …}

- Generative model: models joint distribution for all involved: $p(x, y)$
- Discriminative model: only concerns with prediction $p(y \mid x)$

Human or Machine?

1 "A shallow magnitude 4.7 earthquake was reported Monday morning five miles from Westwood, California, according to the U.S. Geological Survey. The temblor occurred at 6:25 a.m. Pacific time at a depth of 5.0 miles."

Human

Computer

"Vào sáng thứ 2, một trận địa chấn nhẹ 4,7 độ cách Weswood, California năm dặm đã được ghi nhận theo the cục khảo sát địa chất. Các trận động đất xảy ra lúc 06:25 sáng theo giờ Thái Bình Dương ở độ sâu 5,0 dặm.

Human or Machine?

1 "A shallow magnitude 4.7 earthquake was reported Monday morning five miles from Westwood, California, according to the U.S. Geological Survey. The temblor occurred at 6:25 a.m. Pacific time at a depth of 5.0 miles."

✕ Human

✔ Computer

> This excerpt of an initial report about a March 2014 earthquake was written by an algorithm.

"Vào sáng thứ 2, một trận địa chấn nhẹ 4,7 độ cách Weswood, California năm dặm đã được ghi nhận theo the cục khảo sát địa chất. Các trận động đất xảy ra lúc 06:25 sáng theo giờ Thái Bình Dương ở độ sâu 5,0 dặm.

Human or Machine?

2 "Apple's holiday earnings for 2014 were record shattering. The company earned an $18 billion profit on $74.6 billion in revenue. That profit was more than any company had ever earned in history."

Human

Computer

"Thu nhập kỳ nghỉ năm 2014 của Apple đã phá vỡ kỷ lục. Công ty đã kiếm được 18 tỷ lợi nhuận trên 74,6 tỷ doanh thu. Lợi nhuận đó nhiều hơn bất kỳ công ty nào từng kiếm được trong lịch sử".

Human or Machine?

2 "Apple's holiday earnings for 2014 were record shattering. The company earned an $18 billion profit on $74.6 billion in revenue. That profit was more than any company had ever earned in history."

Human

This was an excerpt from an article on Business Insider.

Computer

"Thu nhập kỳ nghỉ năm 2014 của Apple đã phá vỡ kỷ lục. Công ty đã kiếm được 18 tỷ lợi nhuận trên 74,6 tỷ doanh thu. Lợi nhuận đó nhiều hơn bất kỳ công ty nào từng kiếm được trong lịch sử".


Which face is NOT generated by AI?

Which face is NOT generated by AI?

Which face is NOT generated by AI?

 

Which face is NOT generated by AI?

http://www.whichfaceisreal.com

Outline

- What is a generative model?
- Why do we need 'generative' and approaches
- Why generative deep learning now?
  - Deep learning in 2-cent
  - Two motivations for deep generative models
    - Curse of dimensionality
    - Reversing a deep NN
  - Generative Adversarial Networks
- What's next?

Why generative models?

- Prediction = core intelligence, but what else is fundamental to intelligence?

  It is more than just input to output:

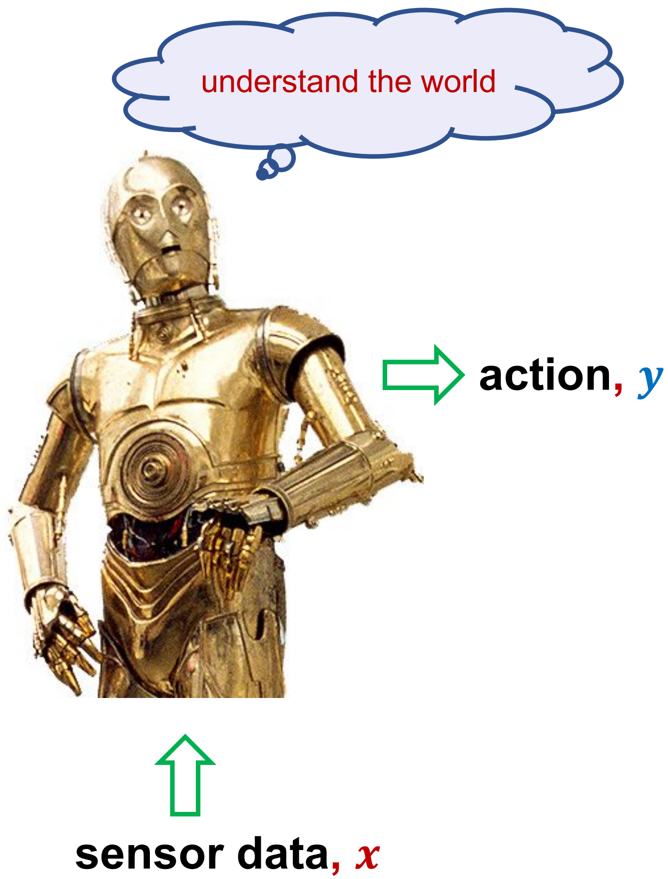

" … the ability to draw deep understanding and insights from data!"

MONASH University

Generative models

- **Generative models** learn how data is generated to re-generate new data.
  - Synthesize poem with Markov process
- Texts, speech, videos, interactions
- Simulations
  - Robot learning without damage
  - Learning tasks with few or expensive labels
- Create new art, music, or fashion
- Synthetic biology
- Many more applications …

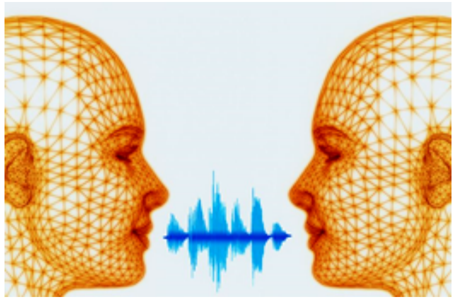

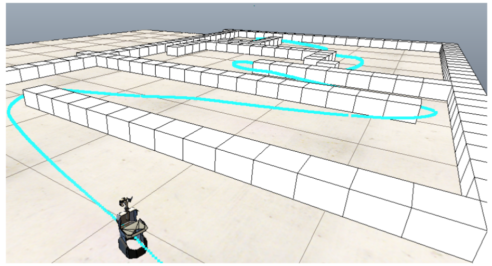

MONASH University

Generative models

- Explain cause for actions beyond associating inputs to outputs.
  - Explainable AI, patients need to know why they will be treated this way or that way!

- Understand and predict the world - how it evolves!
  - Self-driving needs the ability to predict what happens next on the street
  - Robots need the ability to foresee the environments
  - Fintech AI needs to ‘simulate’ the future economy, world events, etc.
  - Biologists wish to understand the evolution from data

[Fellow’14, Mohamed ’15, etc]

Generative models

- Explain cause for actions beyond associating inputs to outputs.
  - Explainable AI, patients need to know why they will be treated this way or that way!

- Understand and predict the world - how it evolves!
  - Self-driving needs the ability to predict what happens next on the street
  - Robots need the ability to foresee the environments
  - Fintech AI needs to ‘simulate’ the future economy, world events, etc.
  - Biologists wish to understand the evolution from data

- Imagine and generate rich plans for the future.
- Recognize objects in the world and their factors of variation.
- Establish concepts useful for reasoning and decision making.
- Detect novelty, surprising events in the world.
- Test our ability to manipulate high-dimensional distributions

[Fellow’14, Mohamed ’15, etc]

Outline

- What is a generative model?
- Why do we need 'generative' and approaches
- Why generative deep learning now?
  - Deep learning in 2-cent
  - Two motivations for deep generative models
    - Curse of dimensionality
    - Reversing a deep NN
  - Generative Adversarial Networks
- What's next?

Approaches to generative models

- Modeling goal: "recovers intrinsic, nonlinear and low-dimensional structures from high-dimensional data"
- Approaches
  - Latent variable modelling

    Learning low-dimensional manifolds
  - Probabilistic graphical models
  - Learning relations, correlation structures, etc.
  - … generative deep learning

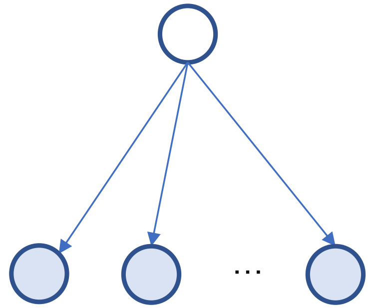

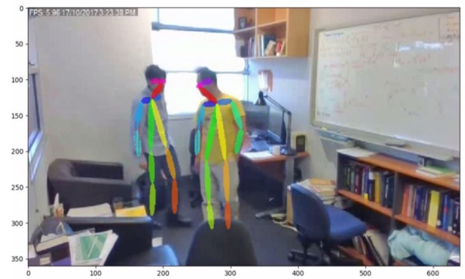

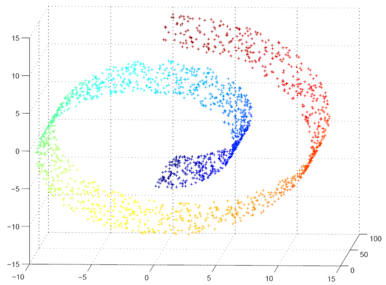

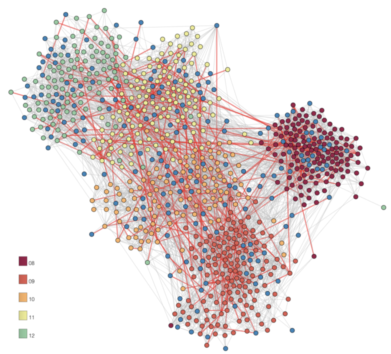

Some work we've done …

- Nested Markov Graphical Models
  - Hierarchical HMM, Coxian HMM, Semi-HMM, Switching HMM, etc.
  - Point out the connection between HHMM and PCFG.
  - Applied to surveillance, activity/abnormality detection and NLP.
    - AIJ'17, NIPS'09, AIJ'09, NIPS'08, CVPR'05, CVPR'05, AAAI'04, etc.

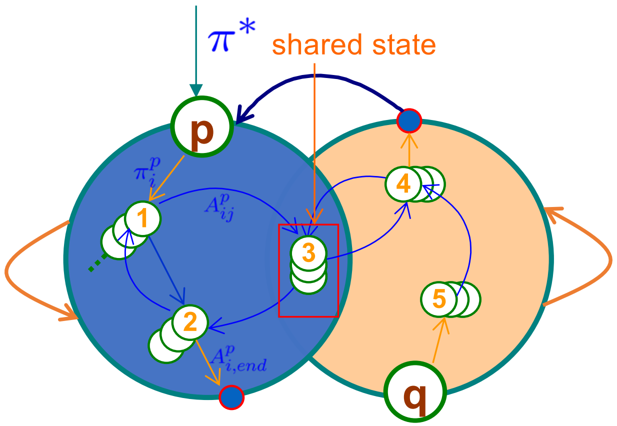

Nested Markov Models

MONASH University

Some work we've done …

- Generative undirected graphical models
  - Hierarchical Conditional Random Fields
  - Preference Networks
  - Restricted Boltzmann Machines (RBM)
  - Applications: recommender system, latent structures learning, choice selection, etc.
    - UAI'17, Info Science'16, AAAI'15, JBI'15, KAIS'15, ACML'11, '12, '13, etc.
    - UAI'09 Best Paper Award Runner-Up.

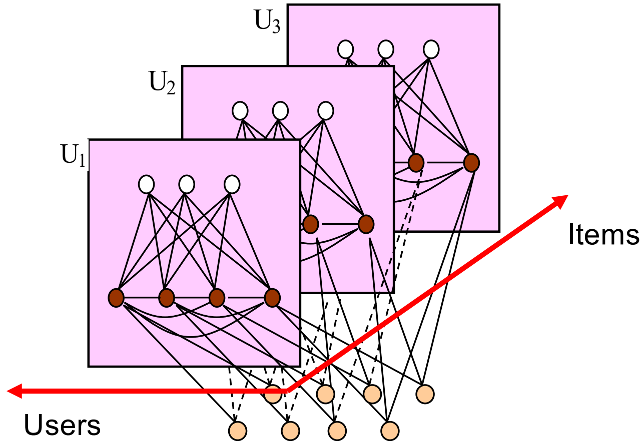

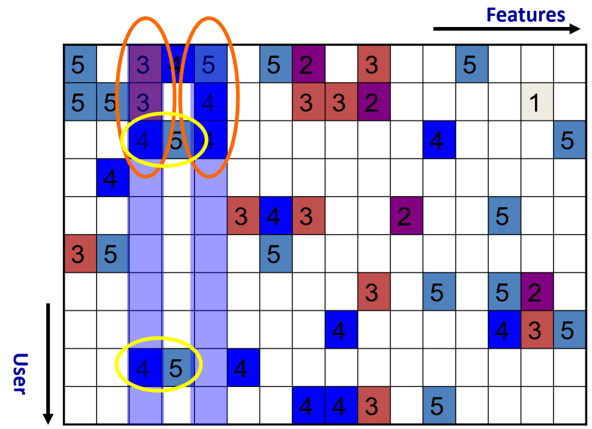

Some work we've done …

- Bayesian topic models
  - Emphasis on nonparametric models for big data and growing complexity.
  - Multilevel Clustering (MC2), Context-Sensitive DPM (CSDP), Hierarchical Beta Processes, etc.
    - ICML'17, IJCAI'17, ICDM'17, UAI'16, ICML'14, SDM'14, ICML'13, SDM'12, UAI'12, etc.

- Numerous ARC-funded and external grants.
- Several applications
  - Surveillance and security, activity recognition and inference, human dynamics, interpreting EMRs, healthcare, suicide prediction, sentiment analysis, burst detection, etc..

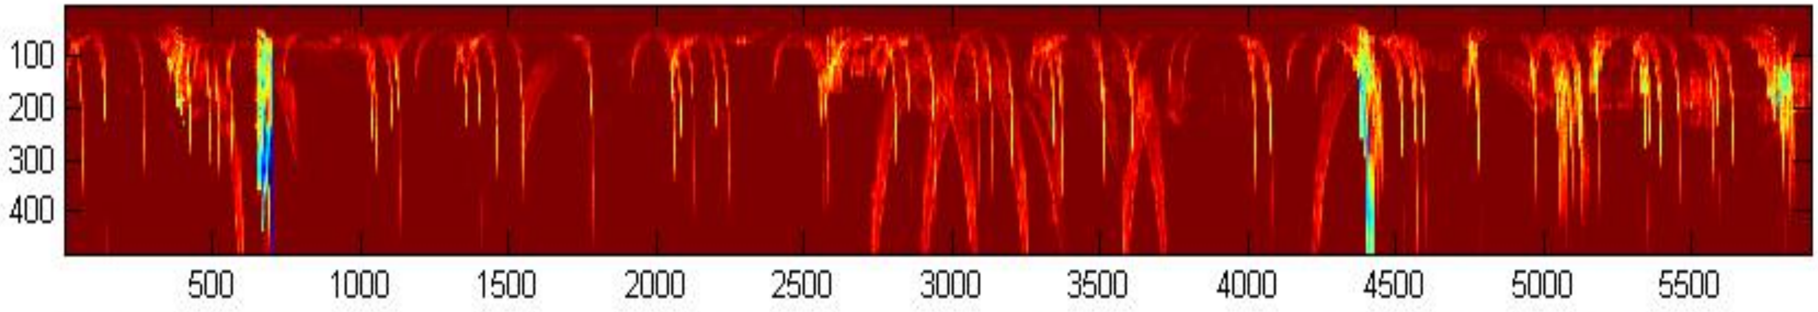

how many movement patterns?

MONASH University

Some recent work too …

- Deep abnormality detection
  - *Ideas*: automatically learn multilevel features, then build ‘universal’ detection system.
  - Detection at semantic level!
- Deep Boltzmann Machines
  - Extremely hard to train
  - Develop new exponential family representation (prep. ICML’19); develop new method to train (AAAI’19)
    - SOTA in clustering, classification, reconstruction tasks!
- Other methods underway.

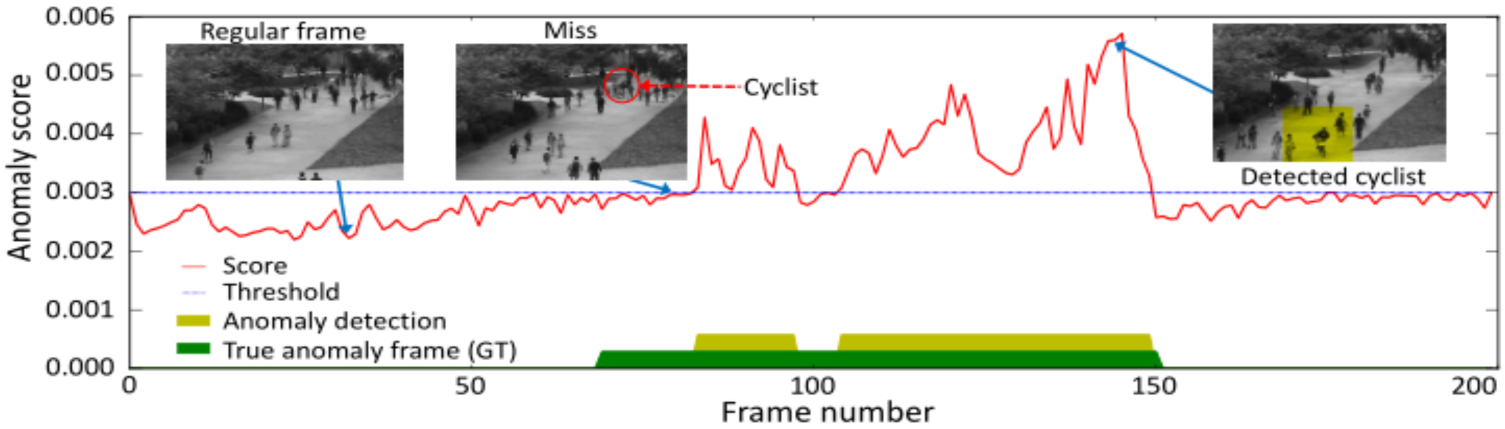

ACML’18, Best Application Paper Award, PAKDD’17

MONASH University

What is common among these approaches?

- Assume a true, but unknown, data distribution: $P(x)$
- Data: $\mathrm{D} = \{x_1, x_2, \dots, x_n\}$ where $x_i \sim P(x)$.
- Goal: estimate density $\tilde{P}(x)$ from data D such that $\tilde{P}(x)$ is as 'close' to $P(x)$ as possible.

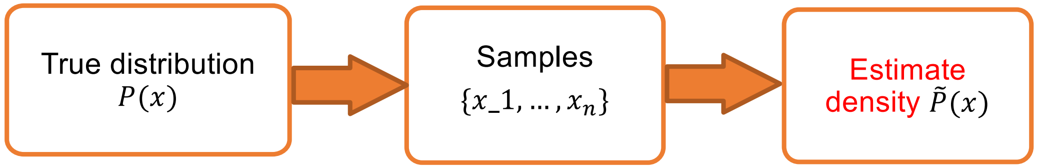

- This is incredibly a useful and important task
  - E.g., to detect abnormality, need to evaluate likelihood of unseen data!
- It is also a very challenging task, e.g., sampling data is hard
- But, what if we're only interested in generating samples?

MONASH University

Outline

- What is a generative model?
- Why do we need ‘generative’ and approaches
- Why generative deep learning now?
  - Deep learning
  - Two motivations for deep generative models
    - Curse of dimensionality
    - Reversing a deep NN
  - Generative Adversarial Networks
- What’s next?

Deep Learning

"Deep Learning: machine learning algorithms based on learning multiple levels of representation and abstraction" - Yoshua Bengio

output y

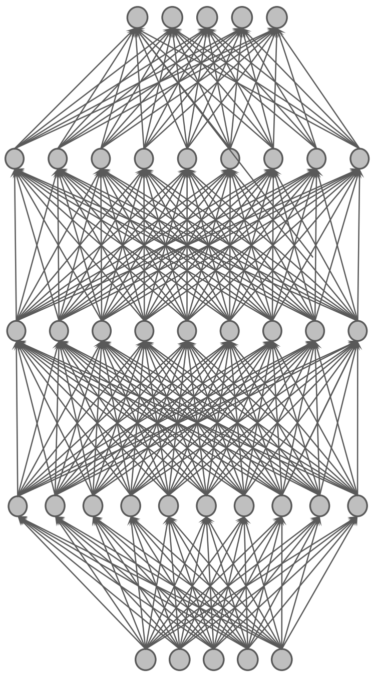

Input x

Hidden layers

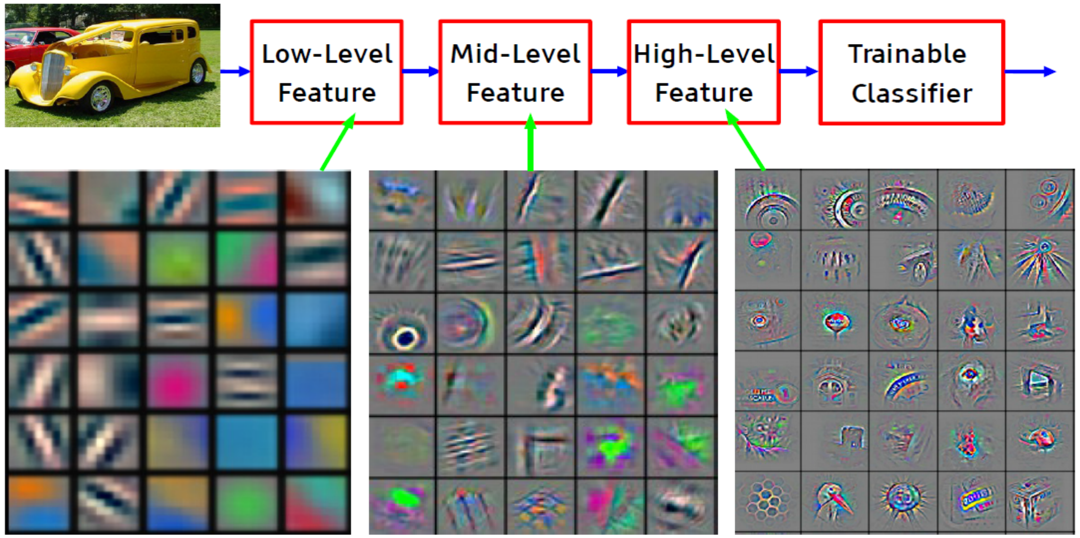

(Image courtesy of Yann Lecun)

Deep learning là những kỹ thuật máy học dựa trên việc học nhiều lớp biểu diễn để mô tả dữ liệu.

Yann LeCun
Director of AI Research, Facebook

Yoshua Bengio
Head, Montreal Institute for Learning Algorithms

Geoff Hinton
Google Brain Toronto, University of Torronto

Fathers of the Deep Learning Revolution Receive ACM A.M. Turing Award

Bengio, Hinton and LeCun Ushered in Major Breakthroughs in Artificial Intelligence

New York, NY, March 27, 2019 – ACM, the Association for Computing Machinery, today named Yoshua Bengio, Geoffrey Hinton, and Yann LeCun recipients of the 2018 ACM A.M. Turing Award for conceptual and engineering breakthroughs that have made deep neural networks a critical component of computing. Bengio is Professor at the University of Montreal and Scientific Director at Mila, Quebec's Artificial Intelligence Institute; Hinton is VP and Engineering Fellow of Google, Chief Scientific Adviser of The Vector Institute, and University Professor Emeritus at the University of Toronto; and LeCun is Professor at New York University and VP and Chief AI Scientist, Facebook.

MONASH University

Outline

- What is a generative model?
- Why do we need ‘generative’ and approaches
- Why generative deep learning now?
  - Deep learning in 2-cent
  - Two motivations for deep generative models
    - Curse of dimensionality
    - Reversing a deep NN
  - Generative Adversarial Networks
- What’s next?

Two motivations for Deep GM

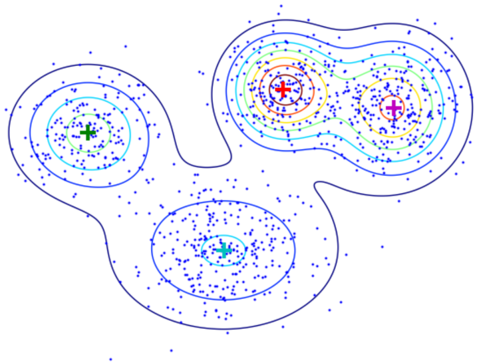

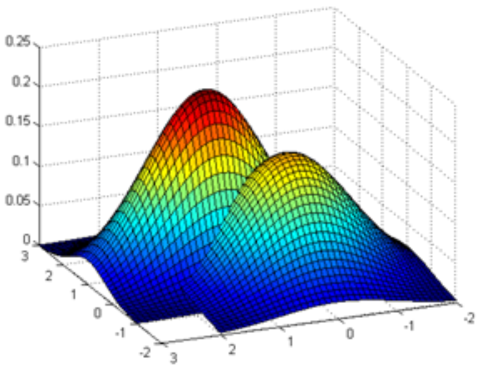

- Realistic images/videos?
- Realistic speeches?
- Realistic languages, texts, dialogs, translation?
- Modern AI systems: deal with really high-dimensional objects!

#1: generating *high-dimensional* objects !

MONASH University

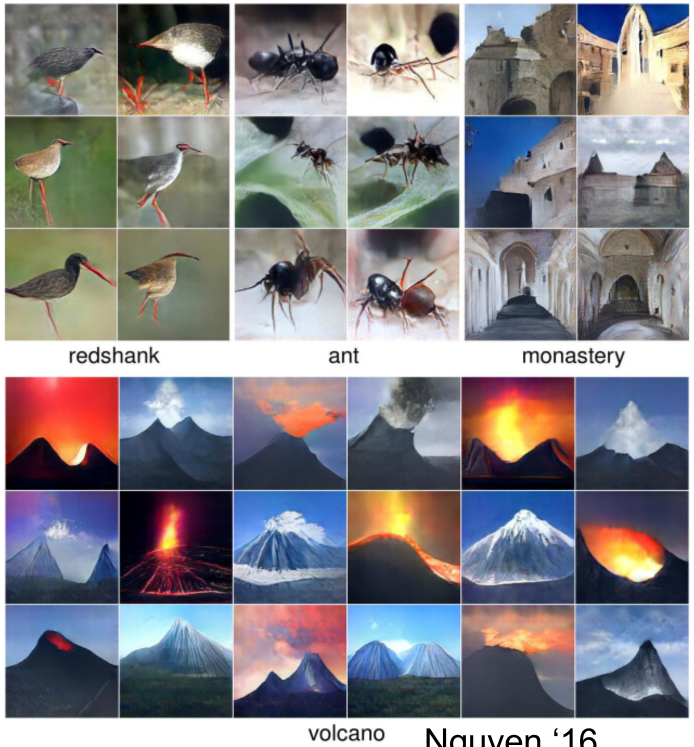

Nguyen '16

Vinyals '15

Human: *i am seeing an error related to vpn*
Machine: *what is the error message that you are getting when connecting to vpn using network connect ?*
Human: *connection refused or something like that*
Machine: *may i know the version of network connect you connect ?*
Human: *i am not sure i know that*

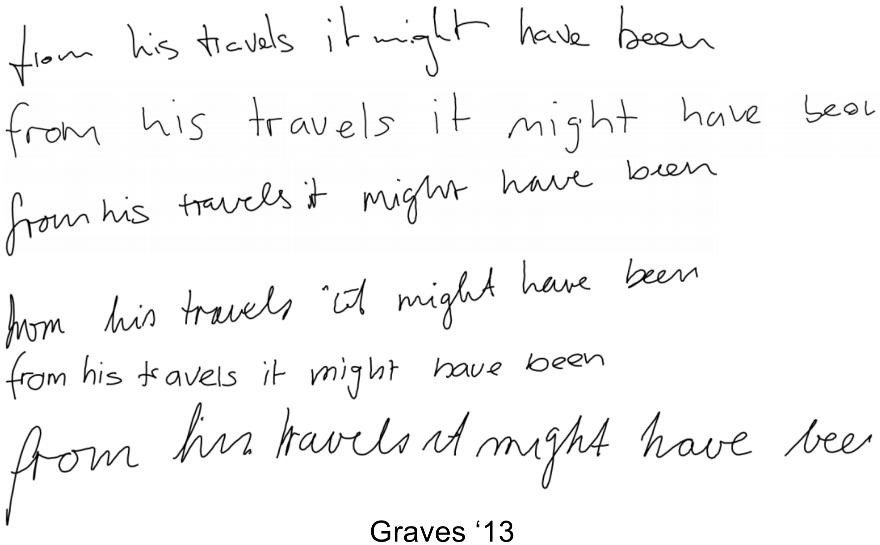

Graves '13

Karras '17

Deep Dream …



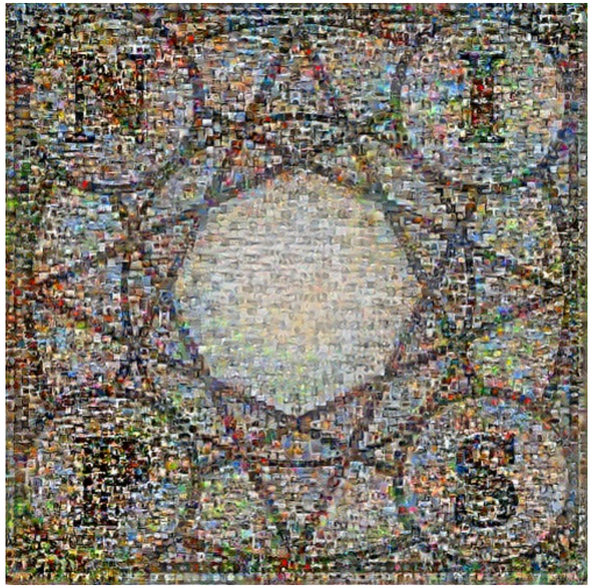






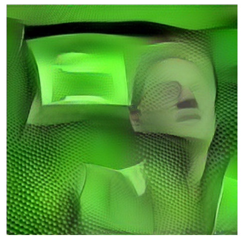

Deep Fake …

› ARTIFICIAL INTELLIGENCE

Seeing Isn't Believing: This New AI System Can Create "Deep Fake" Videos

Sophisticated image processing technology threatens to swamp the internet with next-generation fake news.

BY GLENN MCDONALD · PUBLISHED ON 09/28/2018 · 3:56 PM EDT

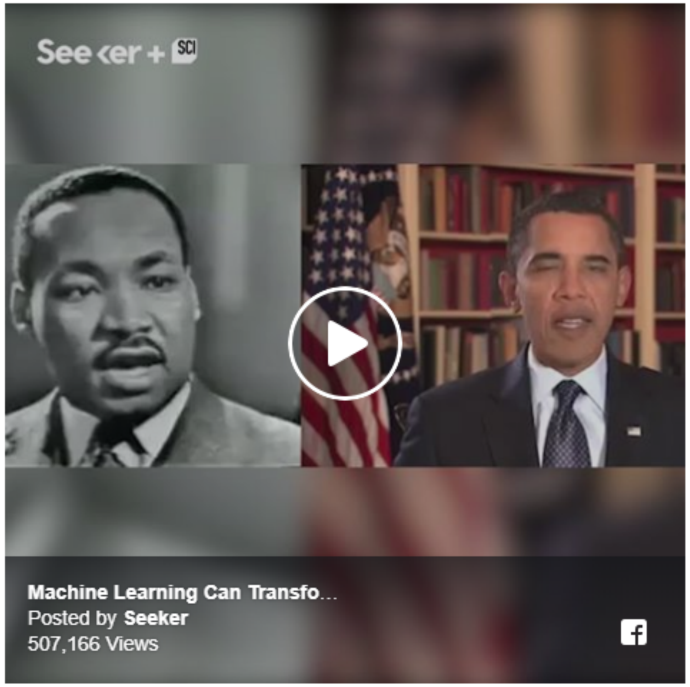

MONASH University

Two motivations for Deep GM

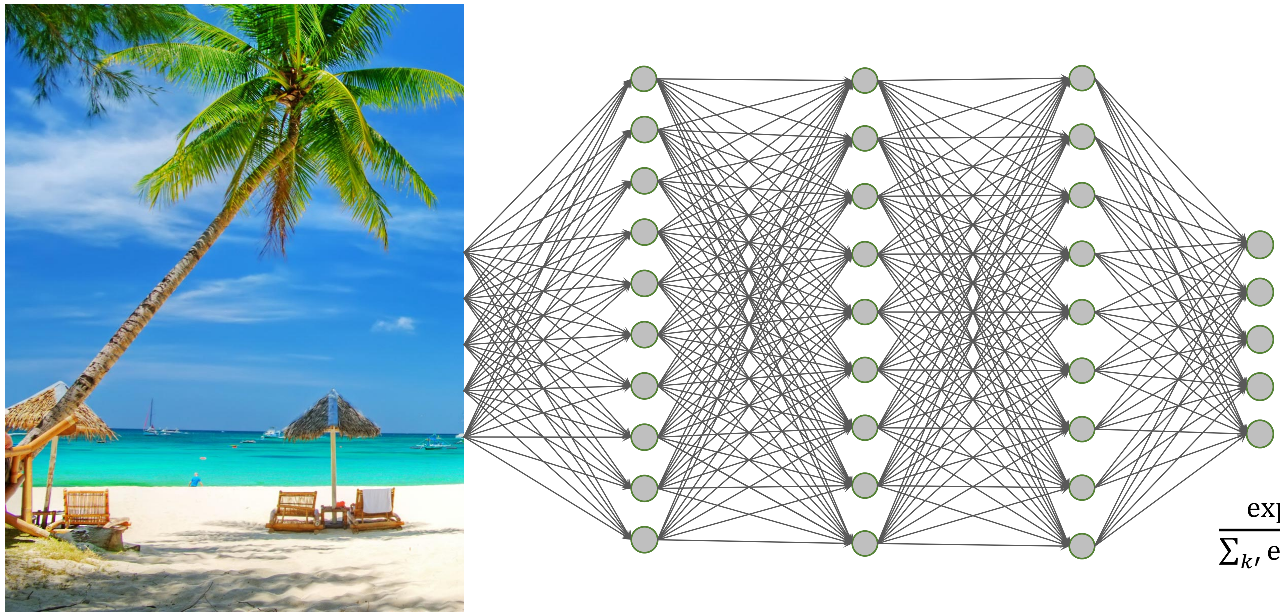

$\Rightarrow$ BEACH

$$\frac{\exp(h^T w_k)}{\sum_{k'} \exp(h^T w_{k'})}$$

#2: can we reverse a deep neural network?

Two motivations for Deep GM

But, can we also build model to ‘generate’ instead of ‘recognize’?

BEACH

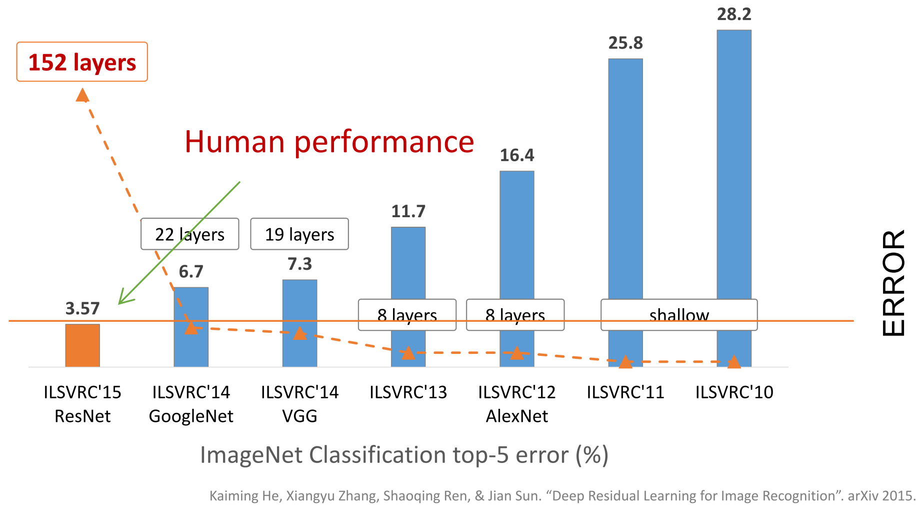

This is the current key success of deep learning!

- Motivation 2: can we reverse a deep neural network?

Two motivations for Deep GM

BEACH

#2: can we reverse a deep neural network?

But, what is a more fundamental notion when learning to generate data?

- To generate, the model must ‘understand the world’:
  - Capture properties of the data, context, knowledge
  - Internalize and build data representation
  - Capture conditional structures, contextual
  - ‘Imagine’ novel structures

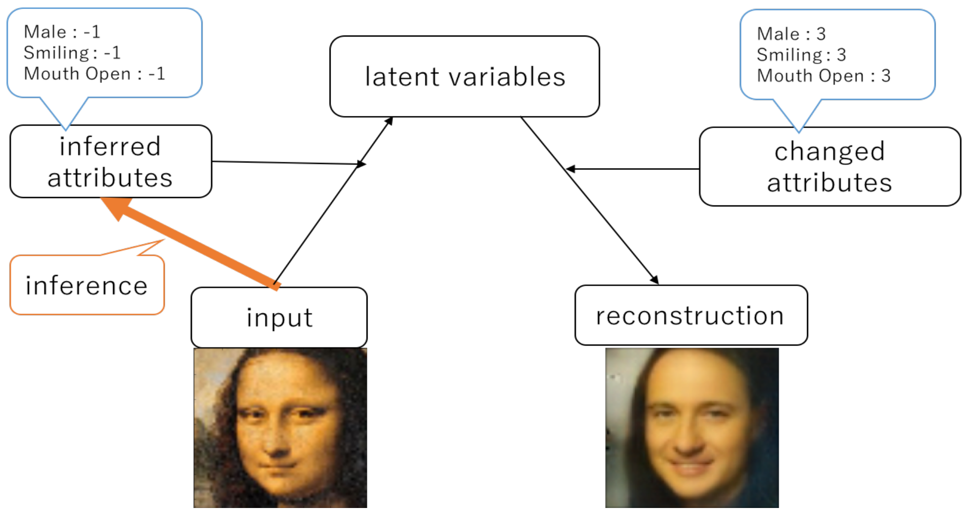

Sohn et al ’15

Explicit vs. Implicit density

Explicit density estimation

True distribution $P(x)$ → Samples $\{x\_1, \dots, x_n\}$ → Estimate density $\tilde{P}(x)$

- Goal: estimate density $\tilde{P}(x)$ such that $\tilde{P}(x)$ is as 'close' to $P(x)$ as possible.

Bayesian networks
Markov random fields
Restricted Boltzmann machines
Deep Boltzmann machines
Generative graphical models
…

Implicit density estimation

True distribution $P(x)$ → Samples $\{x\_1, \dots, x_n\}$ → Estimate generator $G(z)$

- Goal: estimate generator $G(z)$ such that samples generated by $G(z)$ is 'indistinguishable' from samples generated by $P(x)$ where z comes from some arbitrary spaces, like random noise.

Variational Autoencoders
Generative Adversarial Nets

[Goodfellow '16]

But, what if we're only interested in generating samples?

Outline

- What is a generative model?
- Why do we need ‘generative’ and approaches
- Why generative deep learning now?
  - Deep learning in 2-cent
  - Two motivations for deep generative models
    - Curse of dimensionality
    - Reversing a deep NN
  - Generative Adversarial Networks
- What’s next?

Generative Adversarial Networks (GAN)

- GAN estimates a generative model via an *adversarial process* by simultaneously training two models:
  - A *generator* G that induces distribution as close to the true data distribution as possible – this is the generator!
  - A *discriminator* D that estimates the probability that a sample came from the true data distribution rather than from G.
- Output from GAN is the generator G, <u>not</u> an explicit density.

Minimax game theory analogy

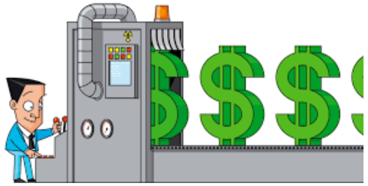

produce counterfeit money after seeing the real ones

knowing what real money looks like, attempt to detect the counterfeit ones

MONASH University

Generative Adversarial Networks (GAN)

How to formulate this?

- Introduce a noise variable $\boldsymbol{z} \sim p_{\boldsymbol{z}}$
  - This is the prior distribution; think of it as the resource used to make money.
- A generator G that maps $\boldsymbol{z}$ to $\widetilde{\boldsymbol{x}}$ in data space:

$$\widetilde{\boldsymbol{x}} = G(\boldsymbol{z})$$

  - Parameterize G by a deep neural network with parameter $\boldsymbol{\theta}_g$.
- A discriminator $D(\boldsymbol{x})$ maps $\boldsymbol{x}$ to a probability in $[0, 1]$ representing the probability that $\boldsymbol{x}$ comes from the real data rather than G.
  - Parameterize $D(\boldsymbol{x})$ by another deep neural network $\boldsymbol{\theta}_d$.

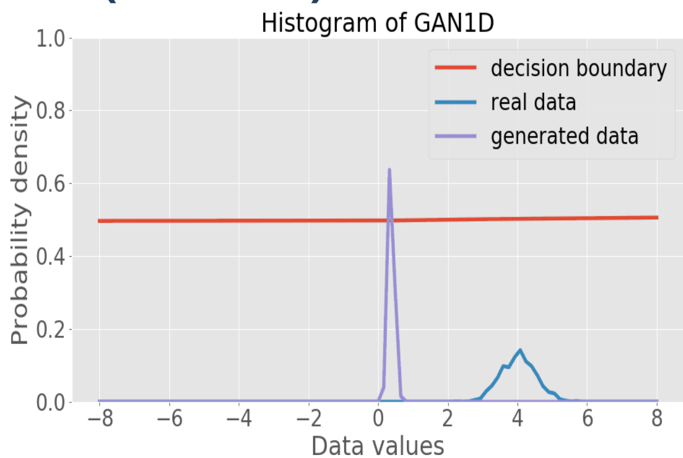

real

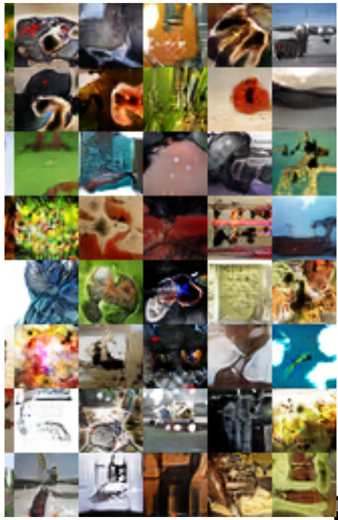

generated

Generative Adversarial Networks (GAN)

training D is to maximize the probability of detecting correct labels

Update D: $\boldsymbol{z}^{(i)} \sim p_z, \boldsymbol{x}^{(i)} \sim p_{data}$

$$\nabla_{\boldsymbol{\theta}_{\mathrm{d}}} \frac{1}{\mathrm{M}} \sum_{i=1}^{M} \left[\log D\left(\boldsymbol{x}^{(i)}\right) + \log\left(1 - D\left(G\left(\boldsymbol{z}^{(i)}\right)\right)\right)\right]$$

$$\min_{G} \max_{\mathrm{D}} J(G, D) = \mathop{\mathbb{E}}_{\boldsymbol{x} \sim p_{data}(\boldsymbol{x})} [\log D(\boldsymbol{x})] + \mathop{\mathbb{E}}_{\boldsymbol{z} \sim p_{\boldsymbol{z}}} [\log(1 - D(G(\boldsymbol{z}))\ ]$$

training G is to minimize the probability of D making mistake, or minimize $1 - D\left(G(\boldsymbol{z})\right)$ to fool the discriminator.

Update G: $\boldsymbol{z}^{(i)} \sim p_z$

$$\nabla_{\boldsymbol{\theta}_{\mathrm{g}}} \frac{1}{\mathrm{M}} \sum_{i=1}^{M} \left[\log\left(1 - D\left(G\left(\boldsymbol{z}^{(i)}\right)\right)\right)\right]$$

Problems with GAN

Two serious technical problems with GAN

- #1: mode collapsing problem
  - Impeach the ability to generate diverse, realistic data/images
- Our solution: Dual GAN (NIPS'17), MGAN (ICLR'18)

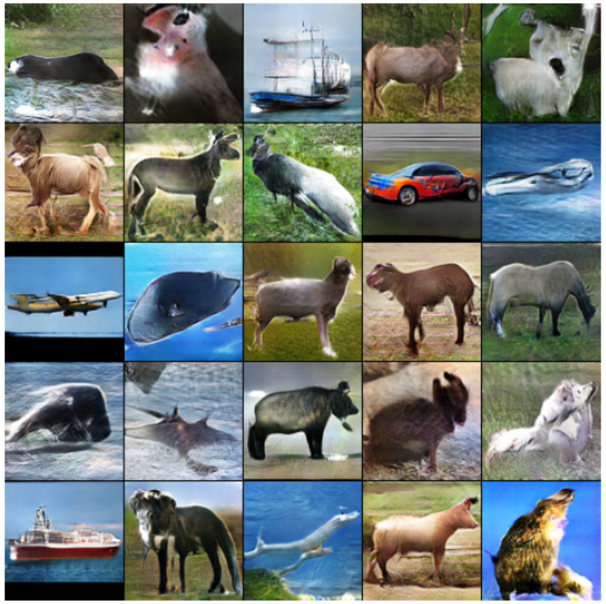

ImageNet
25.78$\pm$0.47
–
–
–
–
–
–
–
7.89
9.18$\pm$0.13
8.25
9.32$\pm$0.10

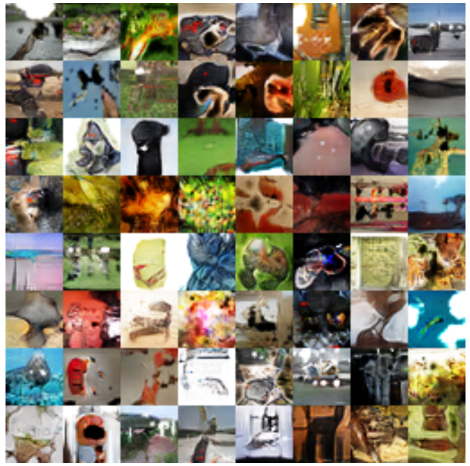

(Original GAN)

Problems with GAN

Two serious technical problems with GAN

- #2: Equilibrium via minimax training is not guaranteed.
  - Training GAN is challenging
- Our solution: Kernelized GAN (arXiv, prep. ICML'19)
  - Exploit duality between f-divergence and loss function.
  - Be able to transform min-max problem into a max-max dual problem, easy to train with theoretical guarantee.

Alternative to GAN?

- GAN is a daring and brilliant idea. But, its original goal is to just generate (high-dimensional) images.
- Is there another way to achieve the same goal?
  - with better geometric interpretation, easier to train and analyse?

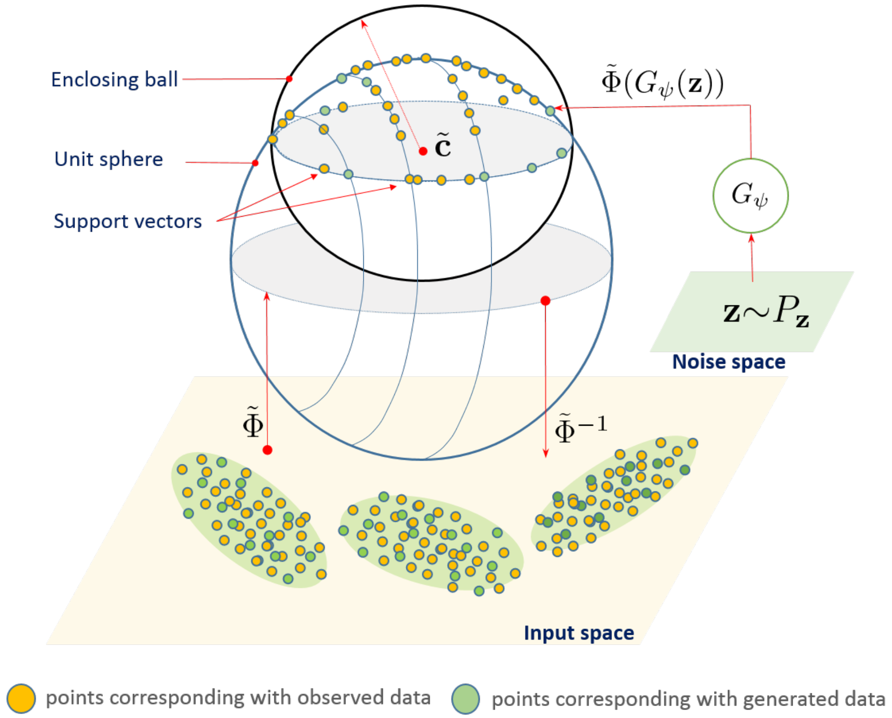

Our solution (on-going):
Generative Enclosing Networks (GEN) [IJCAI'18]

49

Alternative to GAN?

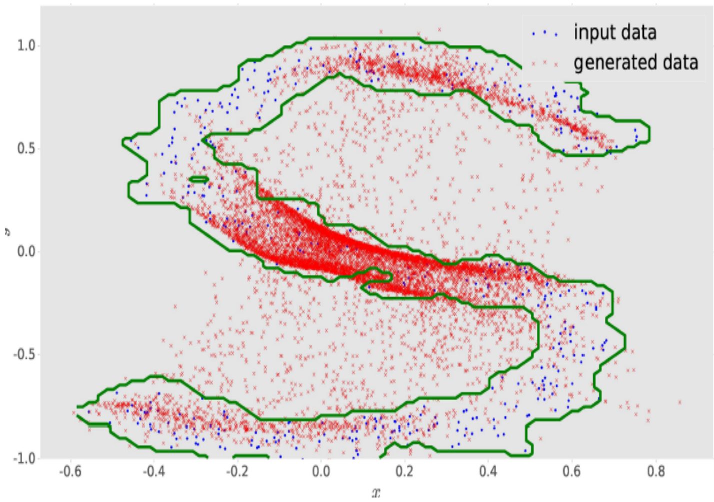

learn and generate data within data manifold

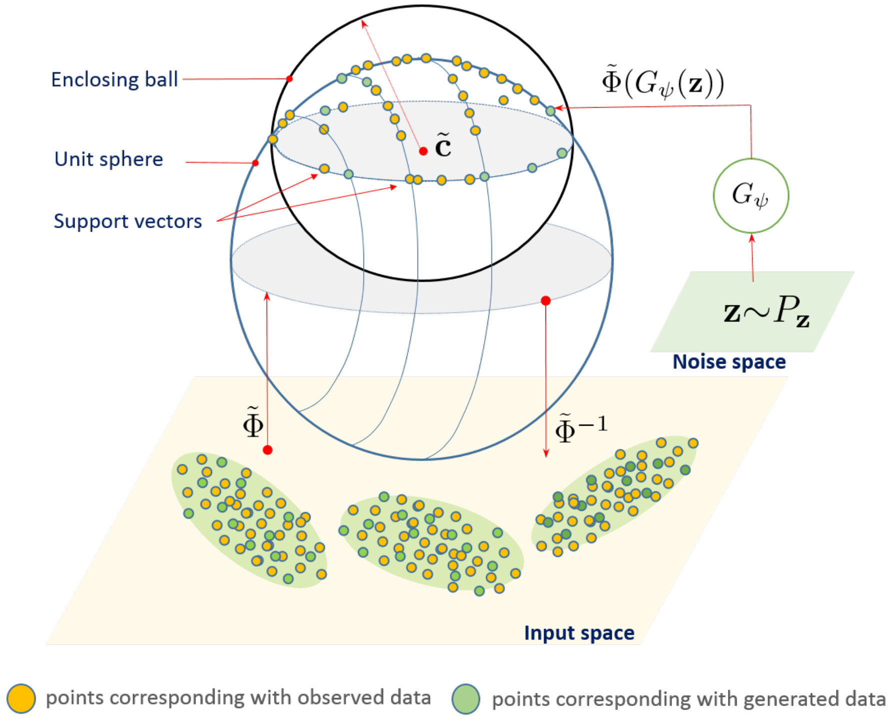

GEN: Generative Enclosing Networks
[IJCAI'18]

50

Outline

- What is a generative model?
- Why do we need ‘generative’ and approaches
- Why generative deep learning now?
  - Deep learning in 2-cent
  - Two motivations for deep generative models
    - Curse of dimensionality
    - Reversing a deep NN
  - Generative Adversarial Networks
- What’s next?

What's next?

- Explosion in computer vision applications, speech and NLP

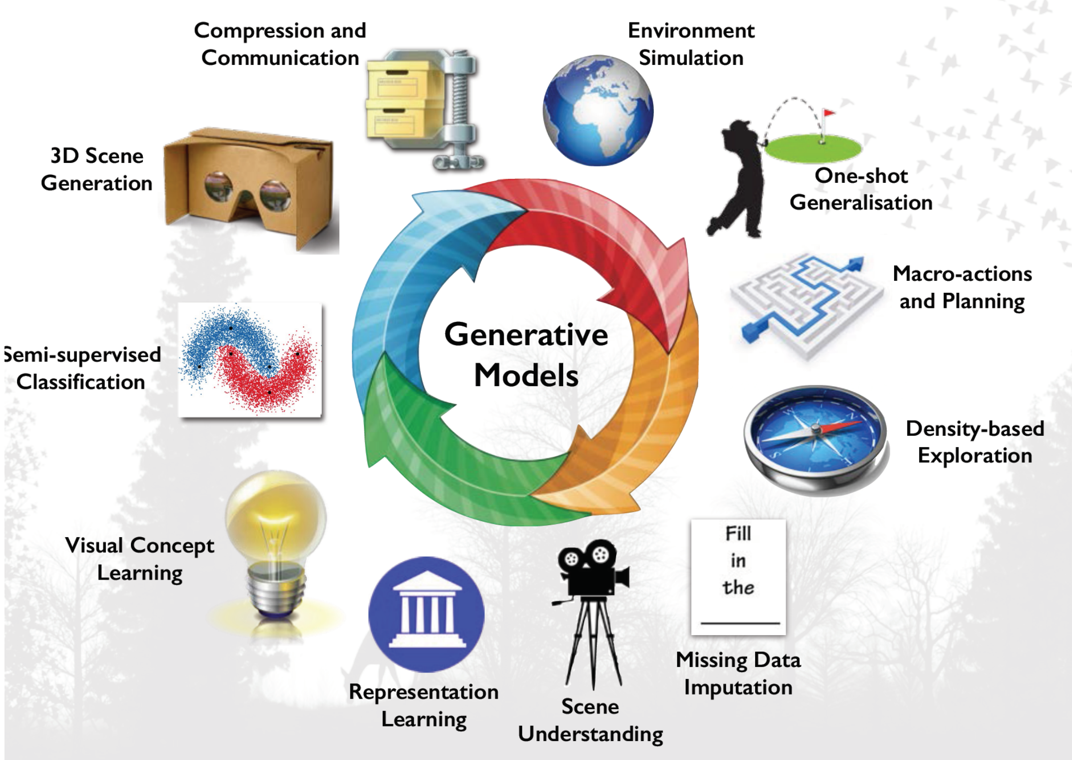

[slide from Mohamed '15]

What’s next?

- Explosion in computer vision applications, speech and NLP

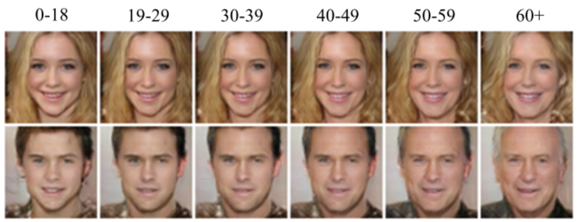

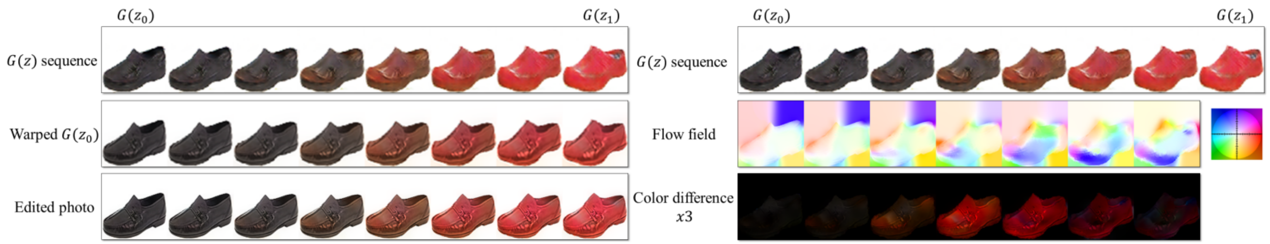

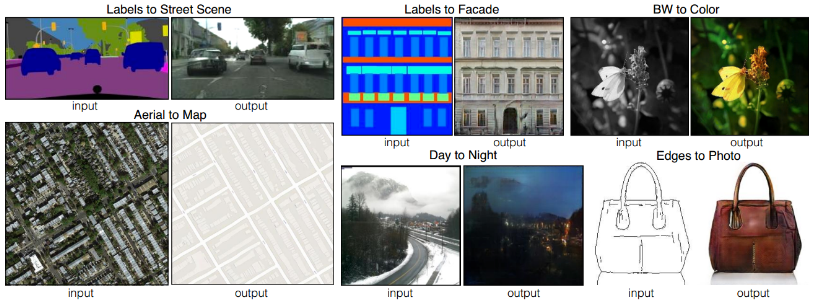

What’s next?

- Explosion in computer vision applications, speech and NLP, and … starting to impact scientific discovery

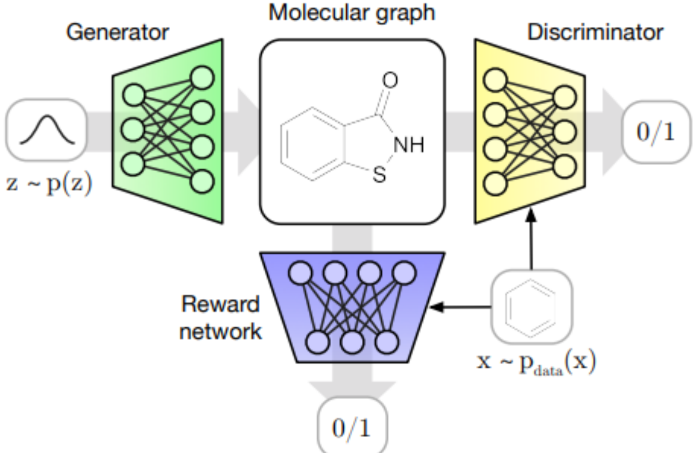

MolGAN: An implicit generative model for small molecular graphs

Nicola De Cao [1] Thomas Kipf [1]

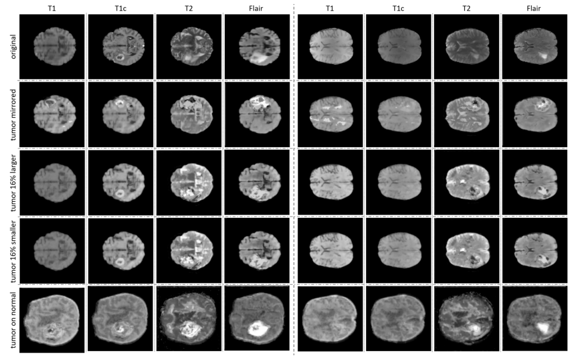

Medical Image Synthesis for Data Augmentation and Anonymization using Generative Adversarial Networks

Hoo-Chang Shin, Neil A Tenenholtz, Jameson K Rogers, Christopher G Schwarz, Matthew L Senjem, Jeffrey L Gunter, Katherine Andriole, Mark Michalski

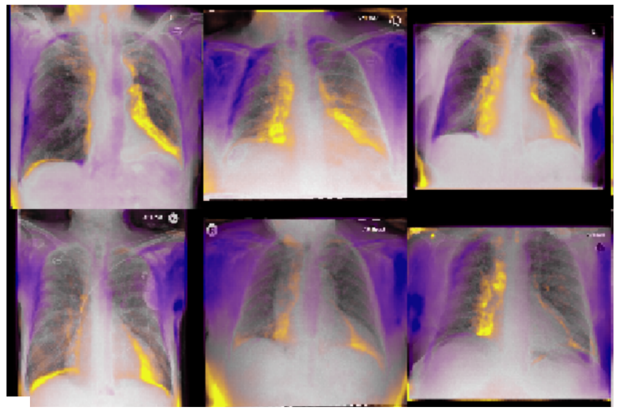

GENERATIVE VISUAL RATIONALES

Jarrel Seah *
Department of Radiology
Alfred Health
Melbourne, VIC 3181, Australia
jarrelscy@gmail.com

Jennifer Tang
Department of Radiology
Royal Melbourne Hospital
Melbourne, VIC 3000, Australia

Andy Kitchen
No affiliation

Jonathan Seah
No affiliation

What's next?

- Explosion in computer vision applications, speech and NLP
- Ripe and exciting time to advance learning theory that wasn't possible before:
  - The concept of adversarial learning is profound and quite paradigm
    - Help us to design next generation of robust and trustworthy machine learning
  - Learning disentangled representation to control and decision making
  - Precision generative models
    - Going beyond faces and entertainment domain to accelerate innovation and scientific discovery
  - Towards strong AI: ability to imagine, create and infer causality !?

Our research

- Deep generative models with a focus on:
  - Solving fundamental problems of Generative Adversarial Networks (GAN)
    - Address mode collapse, better optimization
    - Towards making them useable and stable for real-world applications
    - Papers at NIPS17, ICLR18, IJCAI18, etc.

- Develop theoretical foundations:
  - Deep neural networks, representation learning, Bayesian and inference methods for deep networks, optimization, optimal transport to ML.
  - Papers at: NIPS18 (wsk), ICLR19, AAAI'19, AISTAT'19

- Applying Deep GM to modern AI problems:
  - Human behaviour modelling in finance
  - Medical AI
  - Generative material design
  - Cybersecurity, detecting software vulnerability
  - Automated deep abnormality detection at semantic level
  - Video surveillance, data imputation, tactical behaviours modelling, embed2control
  - Planning and reasoning with Deep GM, and Deep RL

Recent relevant publications

- *Three Player Wasserstein GAN via Amortised Duality,* Nhan Dam, Quan Hoang, Tu Nguyen, Trung Le, Hung Bui **Dinh Phung**, (**IJCAI**) 2019
- *Learning Generative Adversarial Networks from Multiple Data Sources*, Trung Le, Quan Hoang, Tu Nguyen, Hung Bui, **Dinh Phung**, (**IJCAI**) 2019
- *Probabilistic Multilevel Clustering via Composite Transportation Distance* , Viet Huynh, Nhat Ho, **Dinh Phung** and Michael I. Jordan, (**AISTAT**) 2019
- *Robust Anomaly Detection in Videos using Multilevel Representations* , Hung Vu, Tu Dinh Nguyen, Trung Le, Wei Luo and **Dinh Phung**. In In Proceedings of Thirty-third AAAI Conference on Artificial Intelligence ***(AAAI)***, Honolulu, USA, 2019
- *Maximal Divergence Sequential Autoencoder for Binary Software Vulnerability Detection* , Tue Le, Tuan Nguyen, Trung Le, **Dinh Phung**, Paul Montague, Olivier De Vel and Lizhen Qu. In International Conference on Learning Representations ***(ICLR)***, 2019
- *Text Generation with Deep Variational GAN* , Mahmoud Hossam, Trung Le, Michael Papasimeon, Viet Huynh and **Dinh Phung**. In 32nd Neural Information Processing System ***(NIPS)*** Workshop on Bayesian Deep Learning, 2018
- *MGAN: Training Generative Adversarial Nets with Multiple Generators* , Quan Hoang, Tu Dinh Nguyen, Trung Le and **Dinh Phung**. In International Conference on Learning Representations ***(ICLR)***, 2018
- *Dual Discriminator Generative Adversarial Nets* , Tu Dinh Nguyen, Trung Le, Hung Vu and **Dinh Phung**. In Advances in Neural Information Processing Systems 29 ***(NIPS)***, 2017
- Working papers:
  - *KGAN: How to break the minimax game in GAN*, Trung Le, Tu Nguyen, **Dinh Phung**, 2018
  - *Theoretical perspective of deep domain adaptation*, Trung Le, Khanh Nguyen Nhat Ho, **Dinh Phung**, 2019 (NeurIPS, under submission)

MONASH University
57

Acknowledgement and collaborators

Dr Hung Bui Google DeepMind, now Vin AI

A/Prof Long Nguyen Michigan Uni (Ann Arbor)

Dr Nhat Ho UC Berkeley

- My team in the Deep Learning and AI Research Lab at Monash:
  - Trung Le, PhD (postdocs)
  - Ethan Zhao (postdocs)
  - Tu Nguyen, PhD (postdocs)
  - Viet Huynh, PhD (postdocs)
  - Mahmoud Hossam, PhD candidate (Deep GM for sequential structure)
  - Nhan Dam, PhD candidate (Optimal transport theory for Deep GM)
  - Van Nguyen, PhD candidate (Sequential Deep GM and deep domain adaptation)
  - Dai Nguyen, PhD candidate (Representation learning)
  - Quan Hoang, PhD candidate (Deep GM and Self-Attention Model)
  - Tue Le, Master candidate (Maximal Divergence VAE for sequential data)
  - Tuan Nguyen, Master candidate (Deep learning for cybersecurity)
  - ..

MONASH University

58

trustingsocial

"Advancing data science and technology to provide financial access for all"

AI Chief Scientist,
AI Research Lab, Trusting Social, Melbourne, Australia

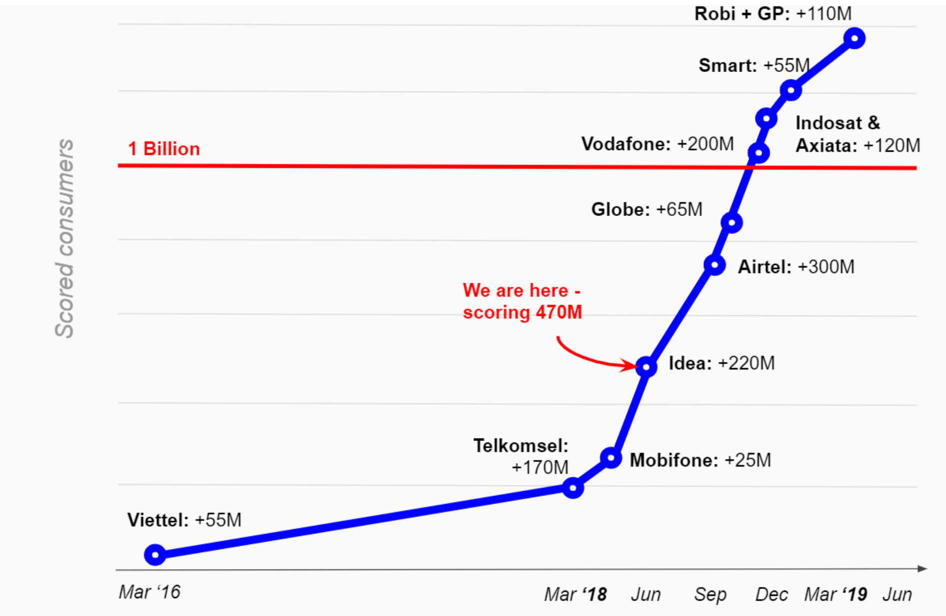

MONASH University

GROUP OF EIGHT AUSTRALIA

Professor of Machine Learning and Data Science
Panel Chair, Master of Artificial Intelligence Program

- **05 markets** covering 2 billion consumers in India, Indonesia, Vietnam, Philippines, Bangladesh
- **200+ team,** 22+ PhDs, 60+ Ms
- **11 offices** in 7 countries
- **Two** research labs in Melbourne and HMC City

Graph analytics
Machine Learning
Deep Learning
Representation Learning and Embeddings
Attention and Transformer Networks
Deep Generative Models and GAN
Transfer Learning
Computer Vision
NLP and Conversational AI

About me ...