Deep learning

Introduction

Hamid Beigy

Sharif University of Technology

October 30, 2021

Table of contents

Course Information

Course Information

1. Course name : Deep learning
2. The objective of deep learning is moving Machine Learning closer to one of its original goals: Artificial Intelligence.
3. Instructor : Hamid Beigy
   Email : beigy@sharif.edu
4. Class Link: https://vc.sharif.edu/ch/beigy
5. Course Website: http://ce.sharif.edu/courses/00-01/1/ce719-1/
   http://sharif.edu/~beigy/14001-40719.html
6. Lectures: Sat-Mon (10:30-12:30)
7. TAs : Fariba Lotfi Email: flotfi@ce.sharif.edu
   Aryan Sadeghi Email: aryansadeghi1374@gmail.com

Course evaluation

- Evaluation:

Mid-term exam	20%	1400-09-24	
Final exam	25%	1400-11-04	
Practical Assignments	35%		
Quiz	15%		
Paper	5%	1400-09-24	Hard deadline for selection

Main reference

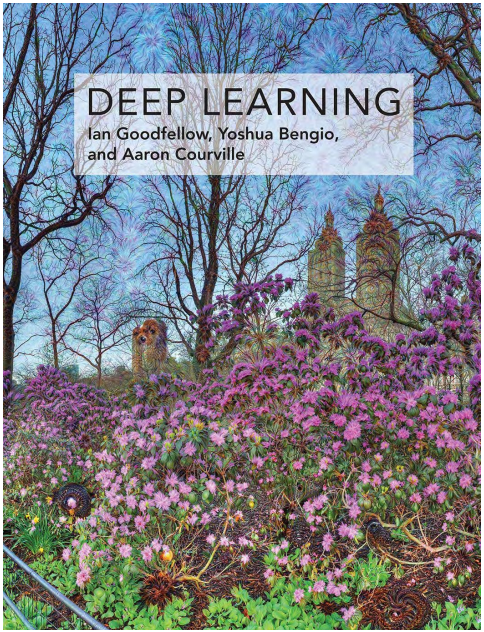

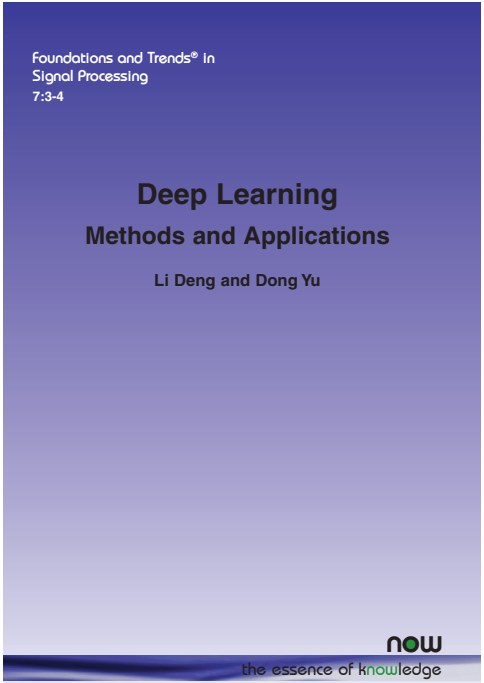

1. IEEE Trans on Pattern Analysis and Machine Intelligence
2. Journal of Machine Learning Research
3. Pattern Recognition
4. Machine Learning
5. Neural Networks
6. Neural Computation
7. Neurocomputing
8. IEEE Trans. on Neural Networks and Learning Systems
9. Annuals of Statistics
10. Journal of the American Statistical Association
11. Pattern Recognition Letters
12. Artificial Intelligence
13. Data Mining and Knowledge Discovery
14. IEEE Transaction on Cybernetics (SMC-B)
15. IEEE Transaction on Knowledge and Data Engineering
16. Knowledge and Information Systems

Relevant conferences

1. Neural Information Processing Systems (NIPS)
2. International Conference on Machine Learning (ICML)
3. European Conference on Machine Learning (ECML)
4. Asian Conference on Machine Learning (ACML)
5. Conference on Learning Theory (COLT)
6. Algorithmic Learning Theory (ALT)
7. Conference on Uncertainty in Artificial Intelligence (UAI)
8. Practice of Knowledge Discovery in Databases (PKDD)
9. International Joint Conference on Artificial Intelligence (IJCAI)
10. IEEE International Conference on Data Mining series (ICDM)

Relevant packages and datasets

1. Packages:
   - Keras https://keras.io
   - TensorFlow http://www.tensorflow.org/
   - Cafe http://caffe.berkeleyvision.org
   - PyTorch https://pytorch.org
2. Datasets:
   - UCI Machine Learning Repository http://archive.ics.uci.edu/ml/
   - MNIST: handwritten digits http://yann.lecun.com/exdb/mnist/
   - 20 newsgroups http://qwone.com/~jason/20Newsgroups/

Introduction

Gartner Hype-Cycle of Emerging Technologies (2016)

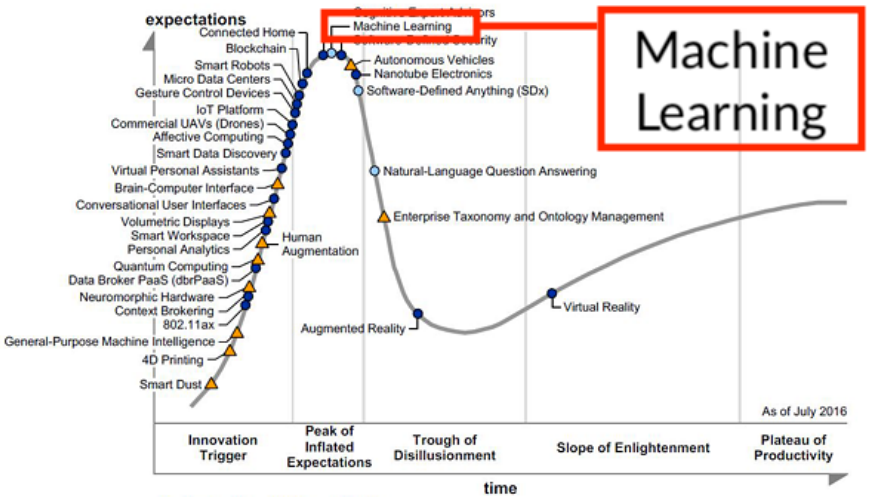

Gartner Hype-Cycle of Emerging Technologies (2017)

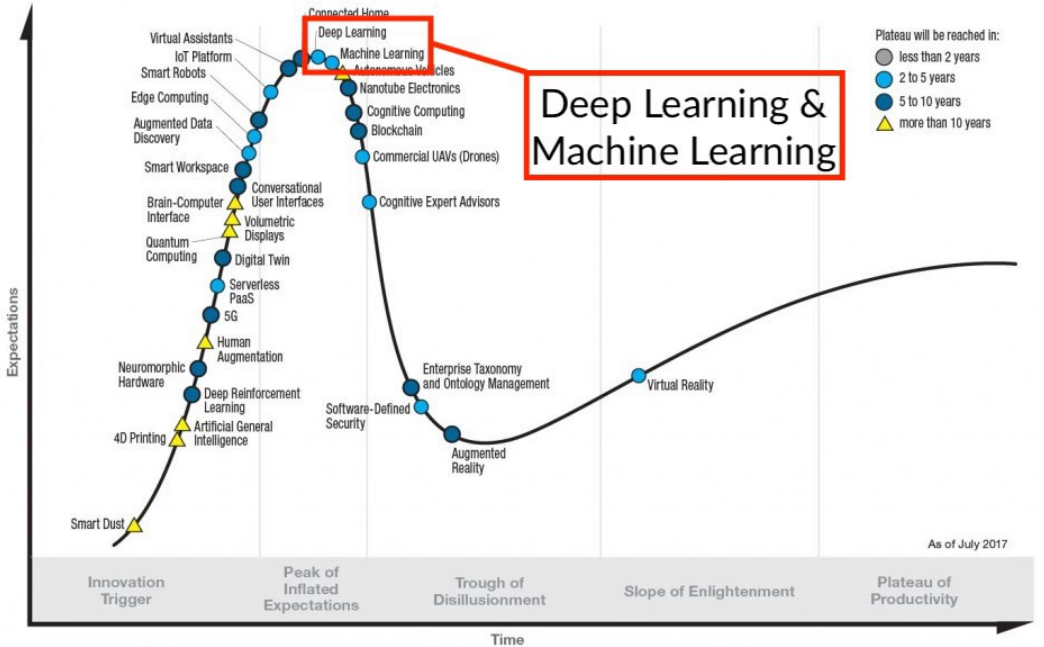

Gartner Hype-Cycle of Emerging Technologies (2018)

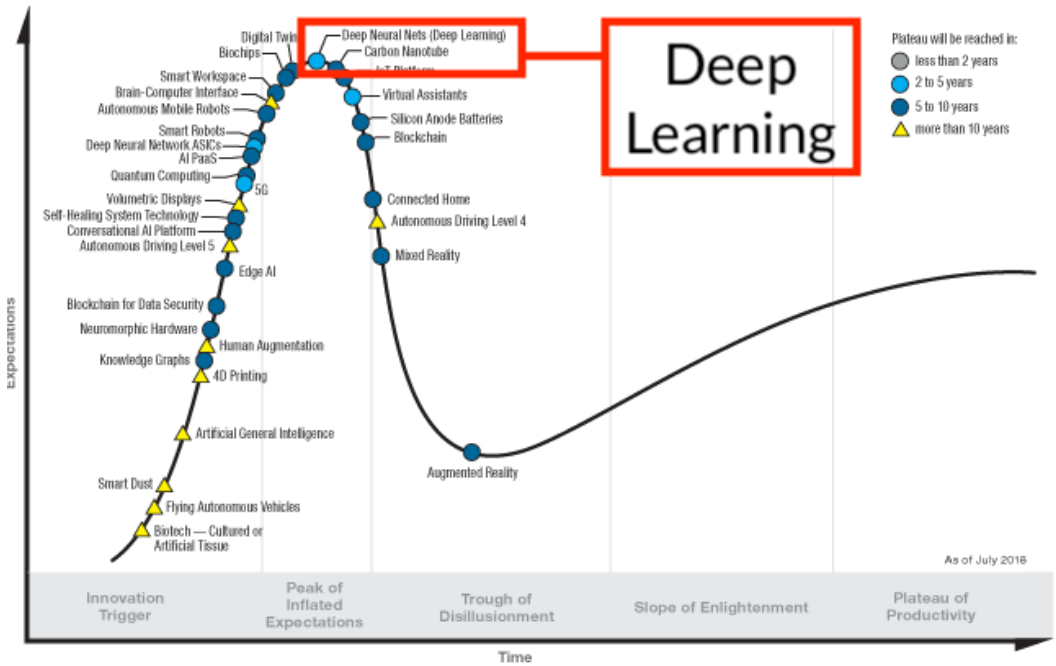

Gartner Hype-Cycle of Emerging Technologies (2019)

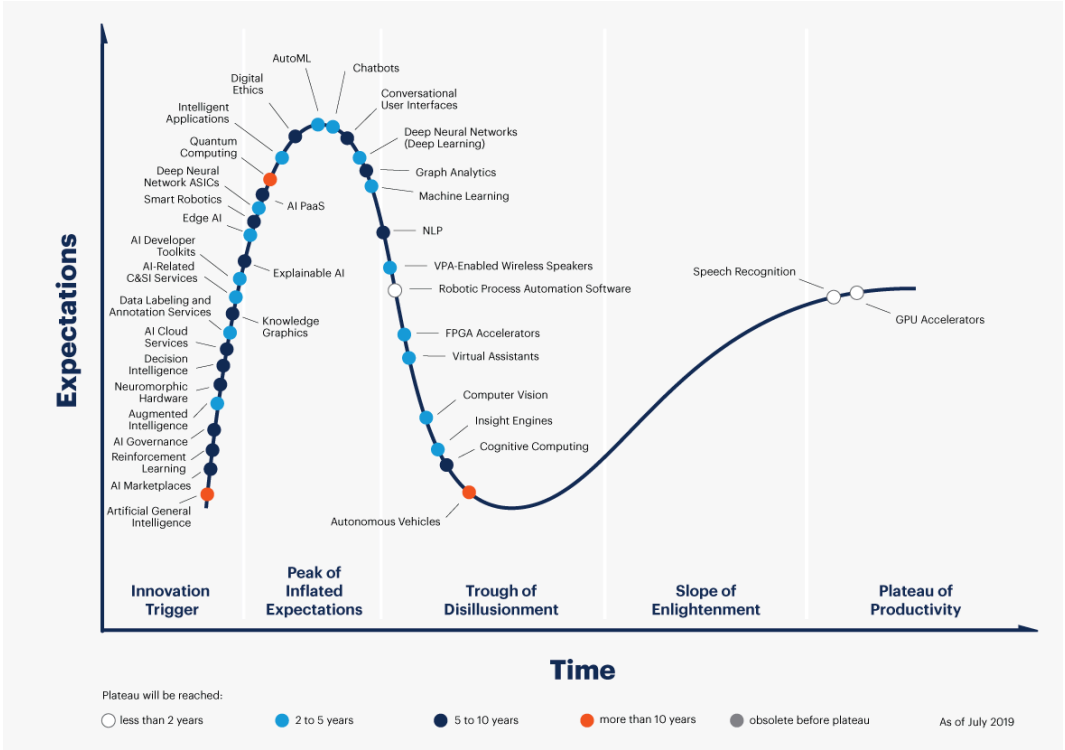

Gartner Hype-Cycle of Emerging Technologies (2020)

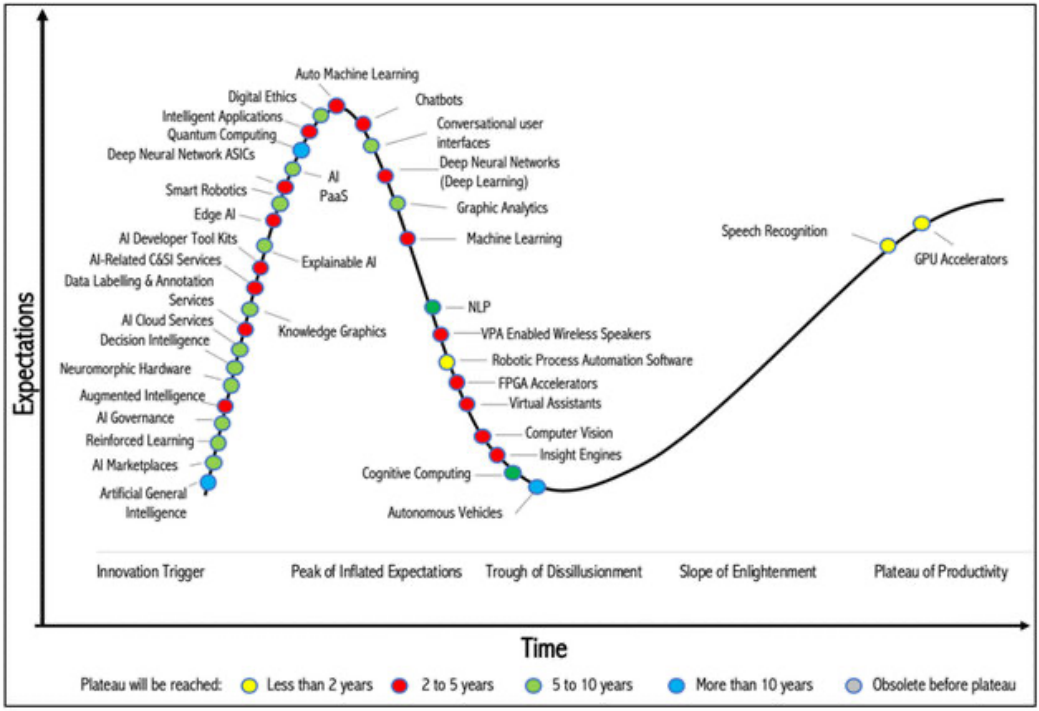

Deep learning has various closely related definitions or high-level descriptions.

Definition (Deep learning)

A sub-field of machine learning that is based on

- learning several levels of representations, corresponding to a hierarchy of features or factors or concepts,
- where
  - higher-level concepts are defined from lower-level ones, and
  - the same lower-level concepts can help to define many higher-level concepts.

Definition (Deep learning)

- Deep learning is part of a broader family of machine learning methods based on learning representations.
- An observation (e.g., an image) can be represented in many ways (e.g., a vector of pixels), but some representations make it easier to learn tasks of interest (e.g., is this the image of a human face?) from examples, and research in this area attempts to define what makes better representations and how to learn them.

An Example

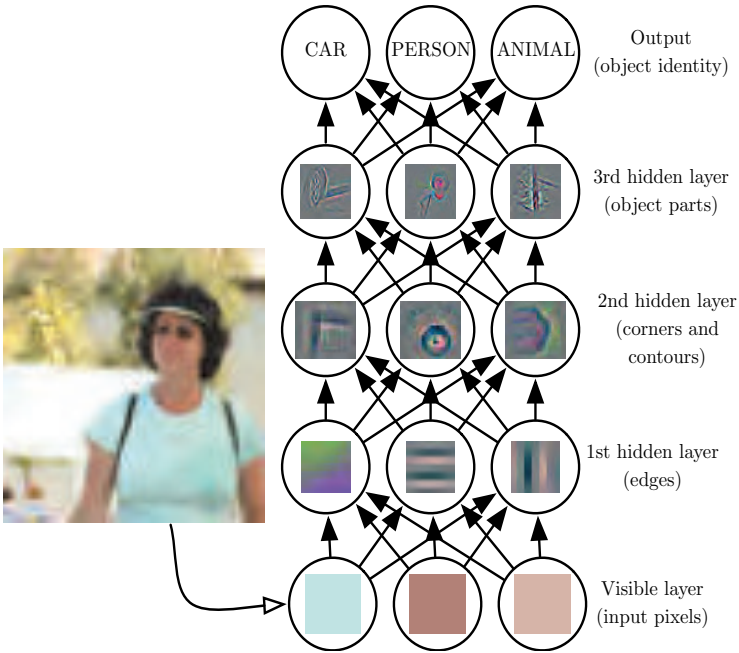
CAR
PERSON
ANIMAL
Output
(object identity)
3rd hidden layer
(object parts)
2nd hidden layer
(corners and
contours)
1st hidden layer
(edges)
Visible layer
(input pixels)

An Example

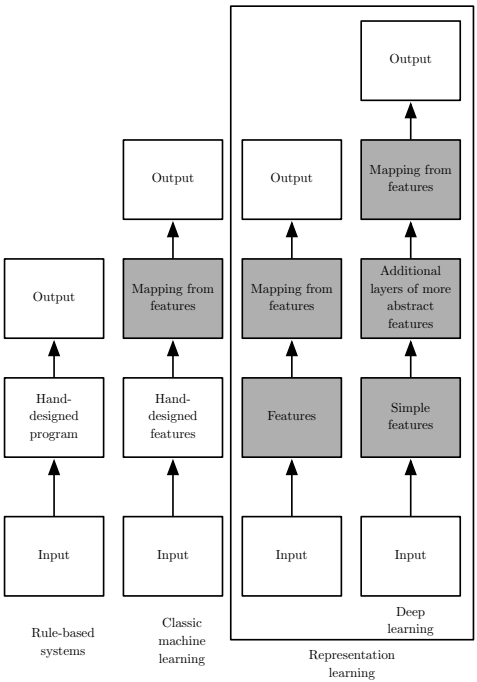
Output
Output
Output
Output
Mapping from features
Mapping from features
Mapping from features
Additional layers of more abstract features
Hand-designed program
Hand-designed features
Features
Simple features
Input
Input
Input
Input
Rule-based systems
Classic machine learning
Representation learning
Deep learning

What is deep learning?

Common among the various high-level descriptions of deep learning are two key aspects:

1. Models consisting of multiple layers/stages of nonlinear information processing
2. Methods for supervised or unsupervised learning of feature representation at successively higher, more abstract layers.

Deep learning is in the intersections among the research areas of

1. Neural networks
2. Artificial intelligence
3. Graphical modeling
4. Optimization
5. Pattern recognition
6. Signal processing.

Success stories

1. Finding nearest images

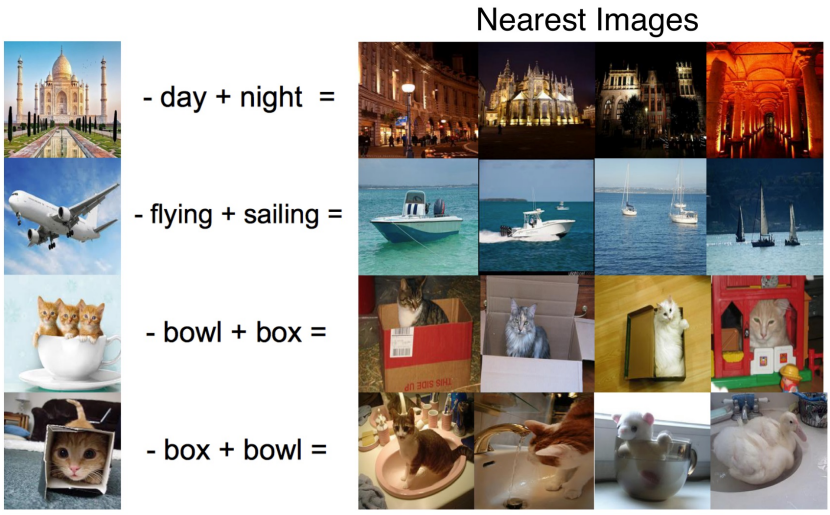

[1]This slide is taken from Prof. Ghodsi's slides.

1. **Word2vec** (Mikolov et al. 2013).
   king – man + woman = queen
2. Google neural machine translation[2]

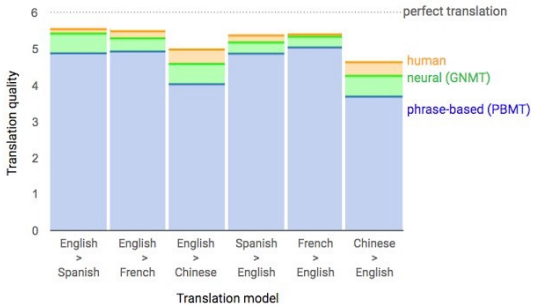

[2]Borrowed from https://blog.statsbot.co/deep-learning-achievements-4c563e034257

1. Wavenet : Generating voice [3]

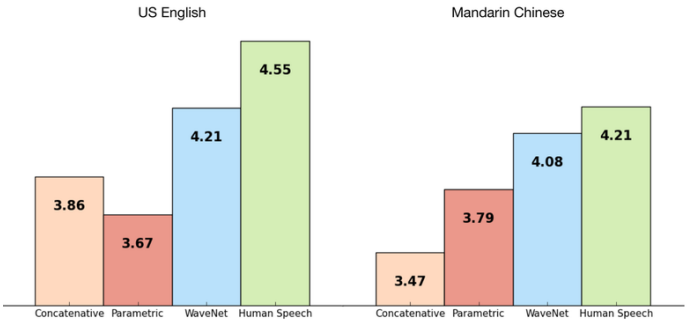

2. Lip Reading

[3] Borrowed from https://blog.statsbot.co/deep-learning-achievements-4c563e034257

1. **LeNet-5**
   LeNet-5 is designed for handwritten and machine-printed character recognition
   Live demo : http://yann.lecun.com/exdb/lenet/index.html
2. **Sentiment Trees**
   Predicting the sentiment of movie reviews.
   Live demo : http://nlp.stanford.edu:8080/sentiment/rntnDemo.html

Success stories of Deep RL

1. TD-Gammon
2. DQN in Atari
3. Deep RL in Robotics
4. Alpha Go and Alpha Zero
5. Dota2 (Video Game)

Outline of course

Outline of course

1. Introduction
2. Review of machine learning and history of deep learning
3. Multi-layer Perceptrons and Backpropagation (MLP)
4. Optimization and Regularization
5. Convolutional networks (CNN)
6. Recurrent networks (RNN)
7. Sum-Product networks (SPN)
8. Dual learning
9. Deep reinforcement learning (Deep RL)
10. Representation learning
11. Deep generative models
12. Graph convolutional networks (GCN)
13. Applications
    - Text mining and natural language processing
    - Computer vision
    - Social networks
14. Advanced topics

References

Readings

1. Chapter 1 of Deep Learning Book (Goodfellow, Bengio, and Courville 2016).
2. Chapter 1 of Deep Learning: Methods and Applications (Deng and Yu 2013).

References i

- Deng, Li and Dong Yu (2013). "Deep Learning: Methods and Applications". In: *Foundations and Trends in Signal Processing* 7.3–4, pp. 197–387.
- Goodfellow, Ian, Yoshua Bengio, and Aaron Courville (2016). *Deep Learning*. MIT Press.
- Mikolov, Tomas et al. (2013). "Distributed Representations of Words and Phrases and their Compositionality". In: *Advances in Neural Information Processing Systems 26*, pp. 3111–3119.

Questions?