人工智能的下一个十年

Jie Tang

Computer Science

Tsinghua University

PPT可以在主页下载: http://keg.cs.tsinghua.edu.cn/jietang/ 或者： Google “Jie Tang”

人工智能的第三次浪潮

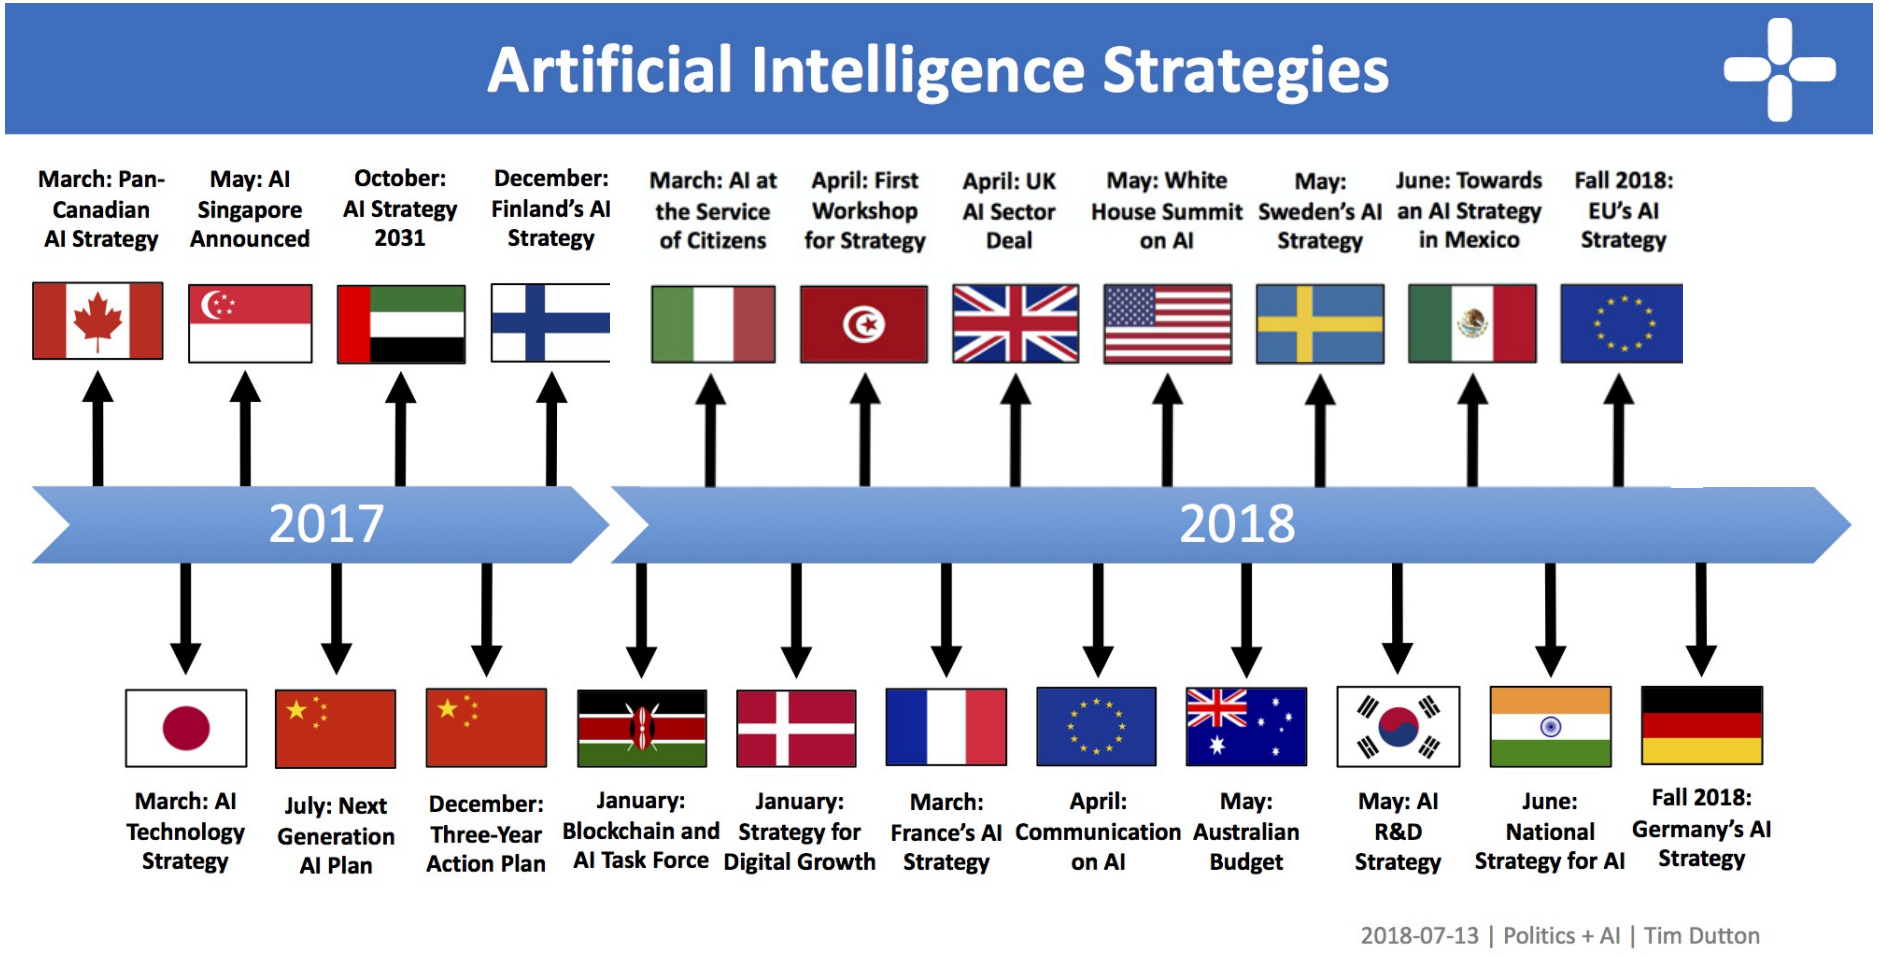
Artificial Intelligence Strategies
March: Pan-Canadian AI Strategy
May: AI Singapore Announced
October: AI Strategy 2031
December: Finland's AI Strategy
March: AI at the Service of Citizens
April: First Workshop for Strategy
April: UK AI Sector Deal
May: White House Summit on AI
May: Sweden's AI Strategy
June: Towards an AI Strategy in Mexico
Fall 2018: EU's AI Strategy
2017
2018
March: AI Technology Strategy
July: Next Generation AI Plan
December: Three-Year Action Plan
January: Blockchain and AI Task Force
January: Strategy for Digital Growth
March: France's AI Strategy
April: Communication on AI
May: Australian Budget
May: AI R&D Strategy
June: National Strategy for AI
Fall 2018: Germany's AI Strategy
2018-07-13 | Politics + AI | Tim Dutton

人工智能近10年

SEARCH

AI趋势：从感知到认知

- 从感知到认知

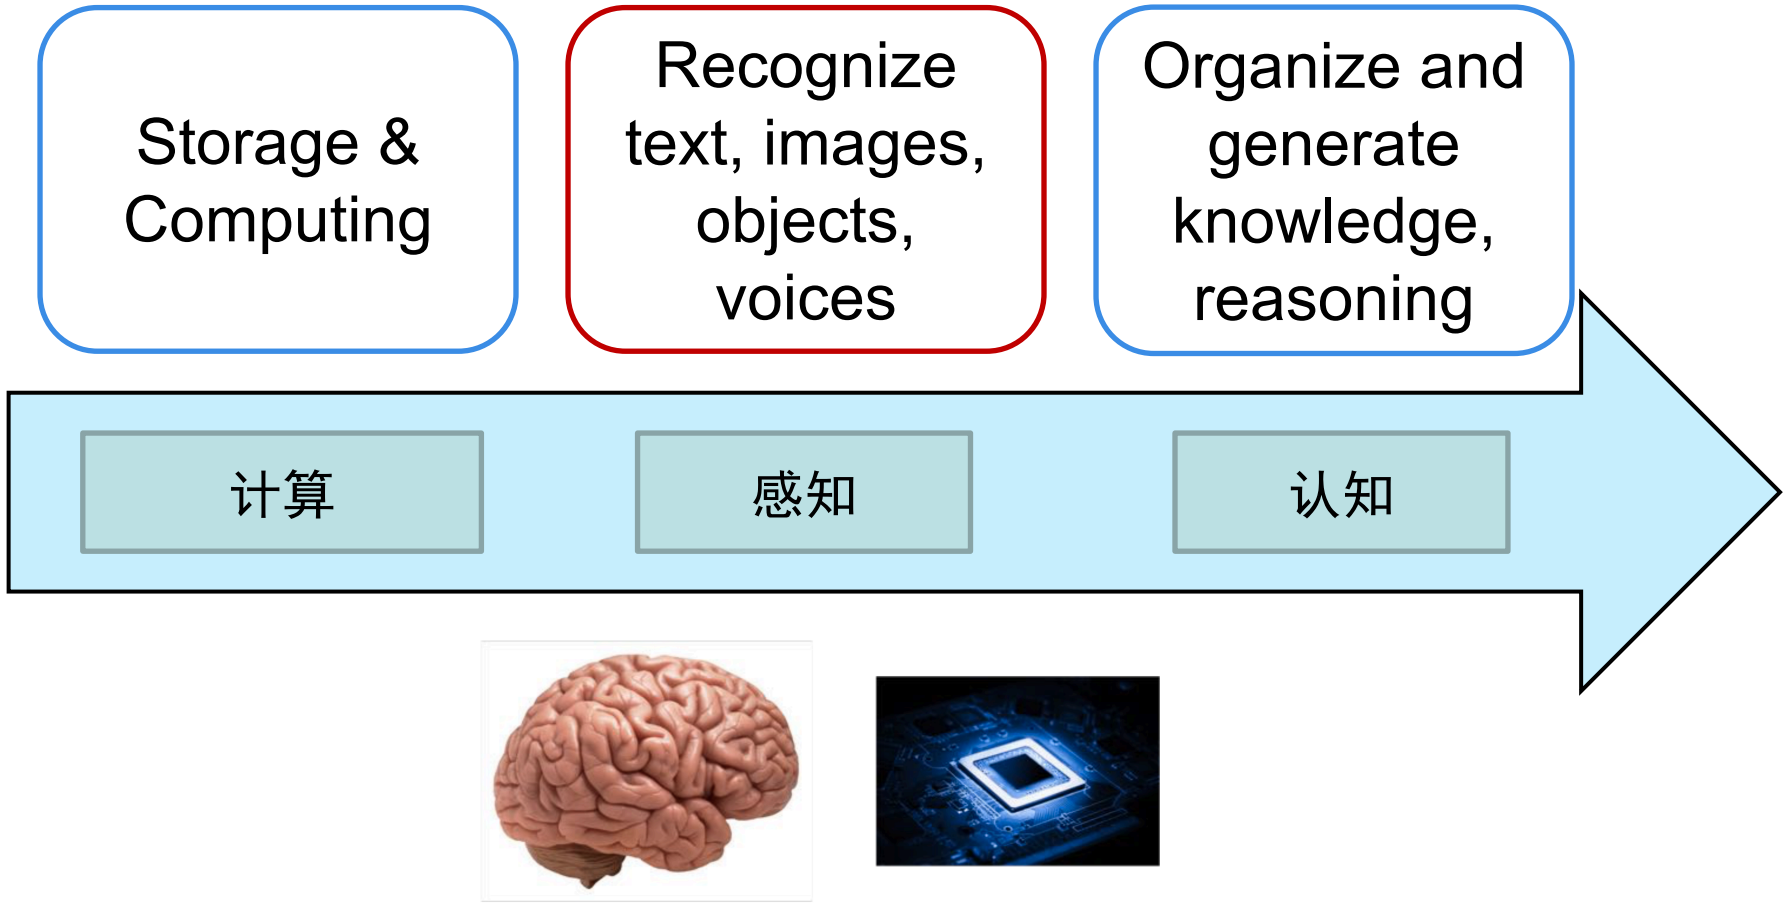

Stochastic vs Deterministic

Uncertainty!

人工智能

博弈对策

无人驾驶

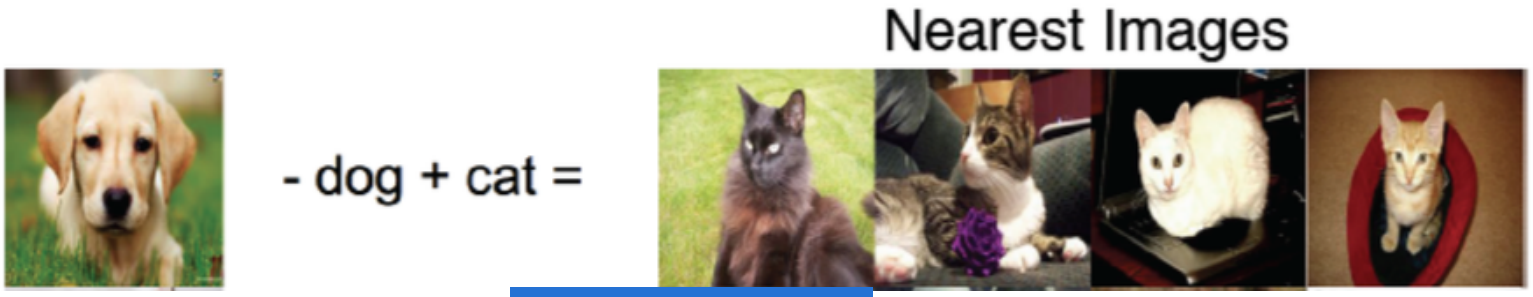

图像识别

回顾机器学习

Classification Models

Decision tree
Bayesian Classifier
Perceptron
Neural Networks

Maximal Margin

SVM
(Vapnik)




Reinforcement Learning

Q-Learning
Policy Gradient
Actor-Critic
TRPO


Sequential Learning

HMM, MEMM, CRF,
Voted Perceptron




Deep Learning

DNN, CNN, ResNet




Deep RL

DQN, VIN, A3C


Graphical Models

Factor Graph,
Exponential Model,
Topic Model (PLSI, LDA)





Recurrent Nets

RNN, LSTM, GRU


Deep Generative Models

DBN, AutoEncoder,
VAE, GAN, WGAN


For more from http://aminer.cn/

DL算法

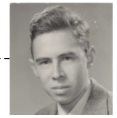

Frank Rosenblatt
Cornell University
psychologist

Geoffery Hinton
University of Toronto
deep learning

Yann Lecun
New York University
deep learning

Perceptron(1958)

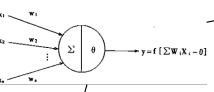

BPNN/MLP(1986)

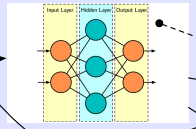

Neocognitron(1980)
[convolution & pooling]

Hopfield Network(1982)
[recurrent & feedback]

LeNet/CNN(1998)

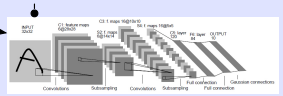

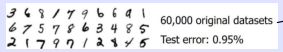

AlexNet(2012)

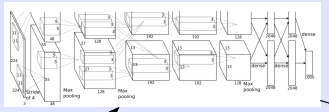

Relu, dropout & bigger

Stacked Capsule Autoencoders(2019)

Capsule Nets(2017)

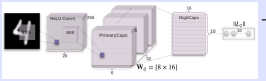

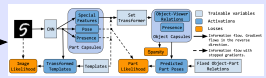

Kaiming He (何恺明)
MSRA => FAIR
computer vision

MoCo(2019)

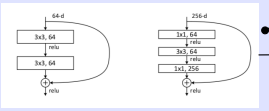

constrastive learning

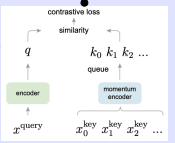

ResNet(2016)

VGG(2014)
GoogLeNet(2015)

DenseNet(2017)
SENet(2018)

NAS(2018)

EfficientNet (2019)

RegNet (2020)

RBM(1986/2006)

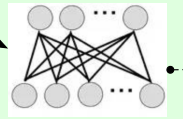

stack

Deep Belief Nets(2006)

AutoEncoder(1989/2006)
Denosing Autoencoder(2008)

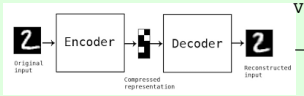

Variational Inference

VAE(2013)

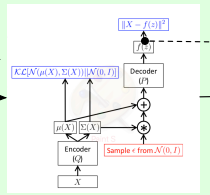

beta-VAE(2018)
VQ-VAE-2(2019)

Max Welling
University of Amsterdam
statistical learning

RNN/LSTM(1997)

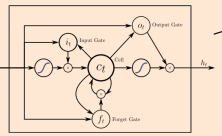

RNN in Speech Recognition(2013)

GAN(2014)

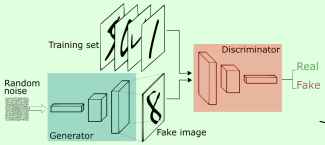

Ian Goodfellow
Google Brain
deep adversarial learning

DCGAN(2014)
WGAN(2017)
BigGAN(2018)

styleGAN(2018)
sinGAN(2019)

IRGAN(2017)
SeqGAN(2017)

ELECTRA(2019)

Christopher Manning
Stanford University
natural languge process

Neural Probabilistic Language Model(2003)

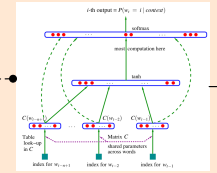

word2Vec(2013)

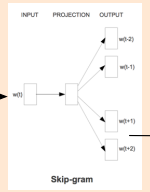

Seq2Seq(2014)

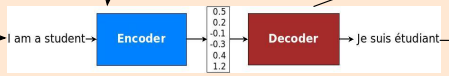

Character CNN(2015)
Transformer(2017)

BERT(2018)

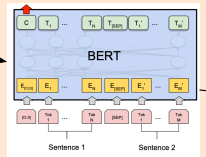

XLNet(2019)
Roberta(2019)
Albert(2019)
T5(2019)

Jürgen Schmidhuber
IDSIA
Universal AI

Yoshua Bengio
University of Montreal
Deep learning

Deep Q-learning(2013)

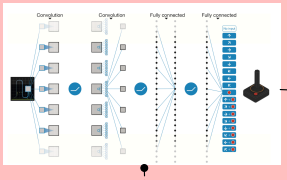

DDPG(2015)

Double DQN(2015)
Dueling Net(2016)

MCTS

AlphaGo(2016)

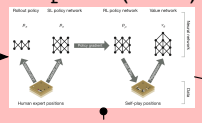

Alpha Zero(2017)

A3C(2016)

TD3(2018)
PPO(2018)

SAC(2019)

Sergey Levine
UC Berkeley
Robotics

David Silver
DeepMind
Reinforcement learning

预训练模型BERT

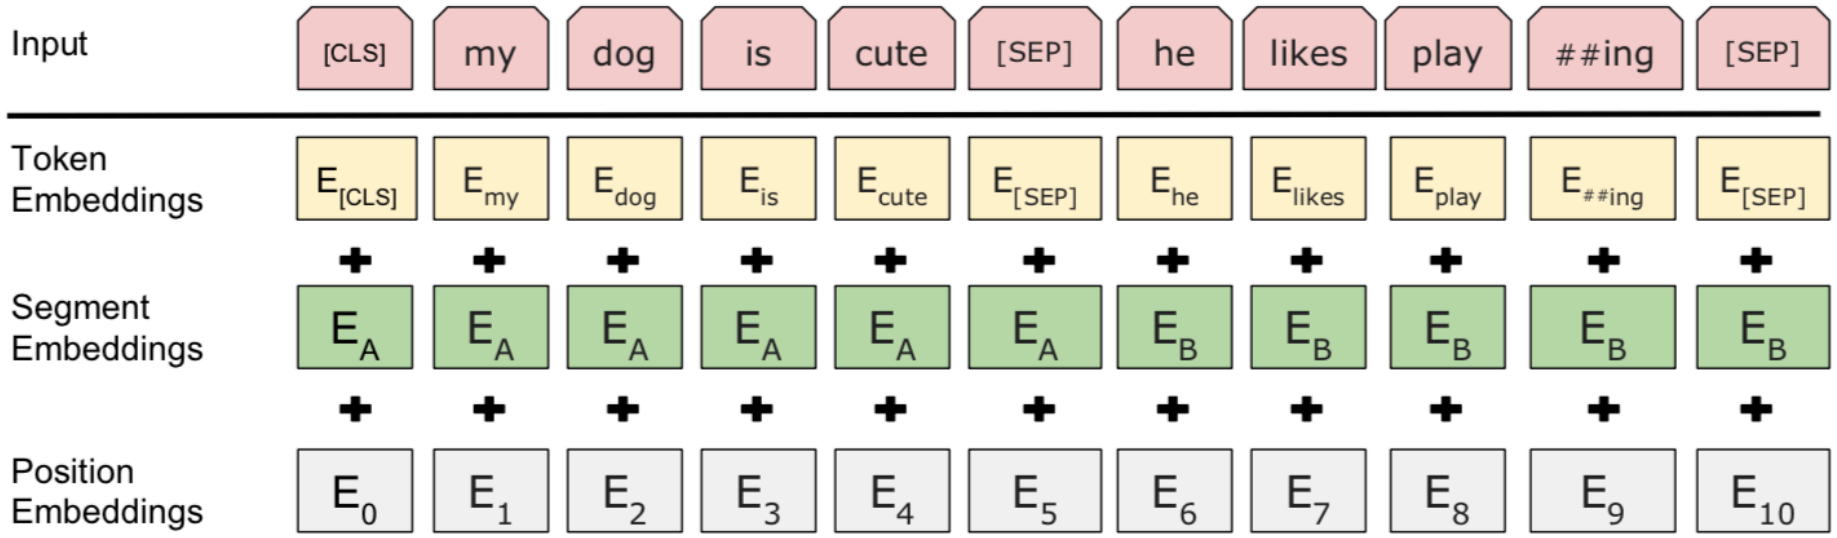

- 预训练Pre-train
- 微调Fine tune
- Beat all state-of-the-arts on 11 NLP tasks in 2018

https://arxiv.org/pdf/1810.04805.pdf

Video-to-Video Synthesis

- The best video synthesis performance

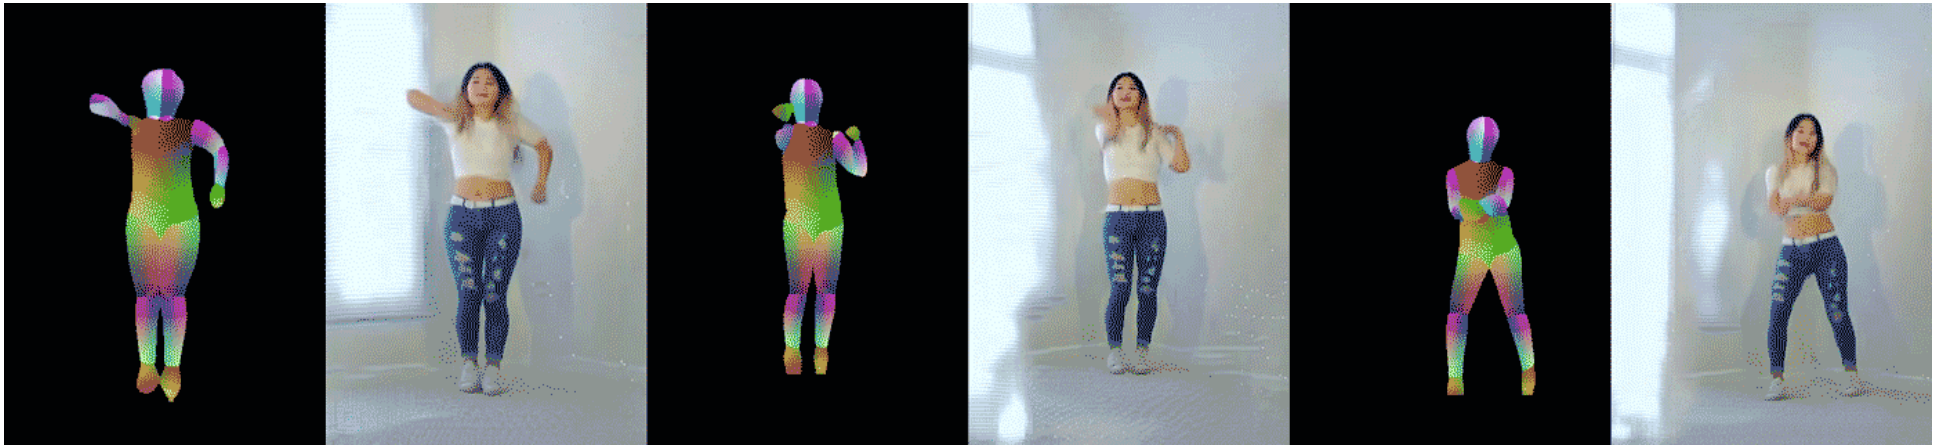

https://arxiv.org/abs/1808.06601

自监督学习— MoCo

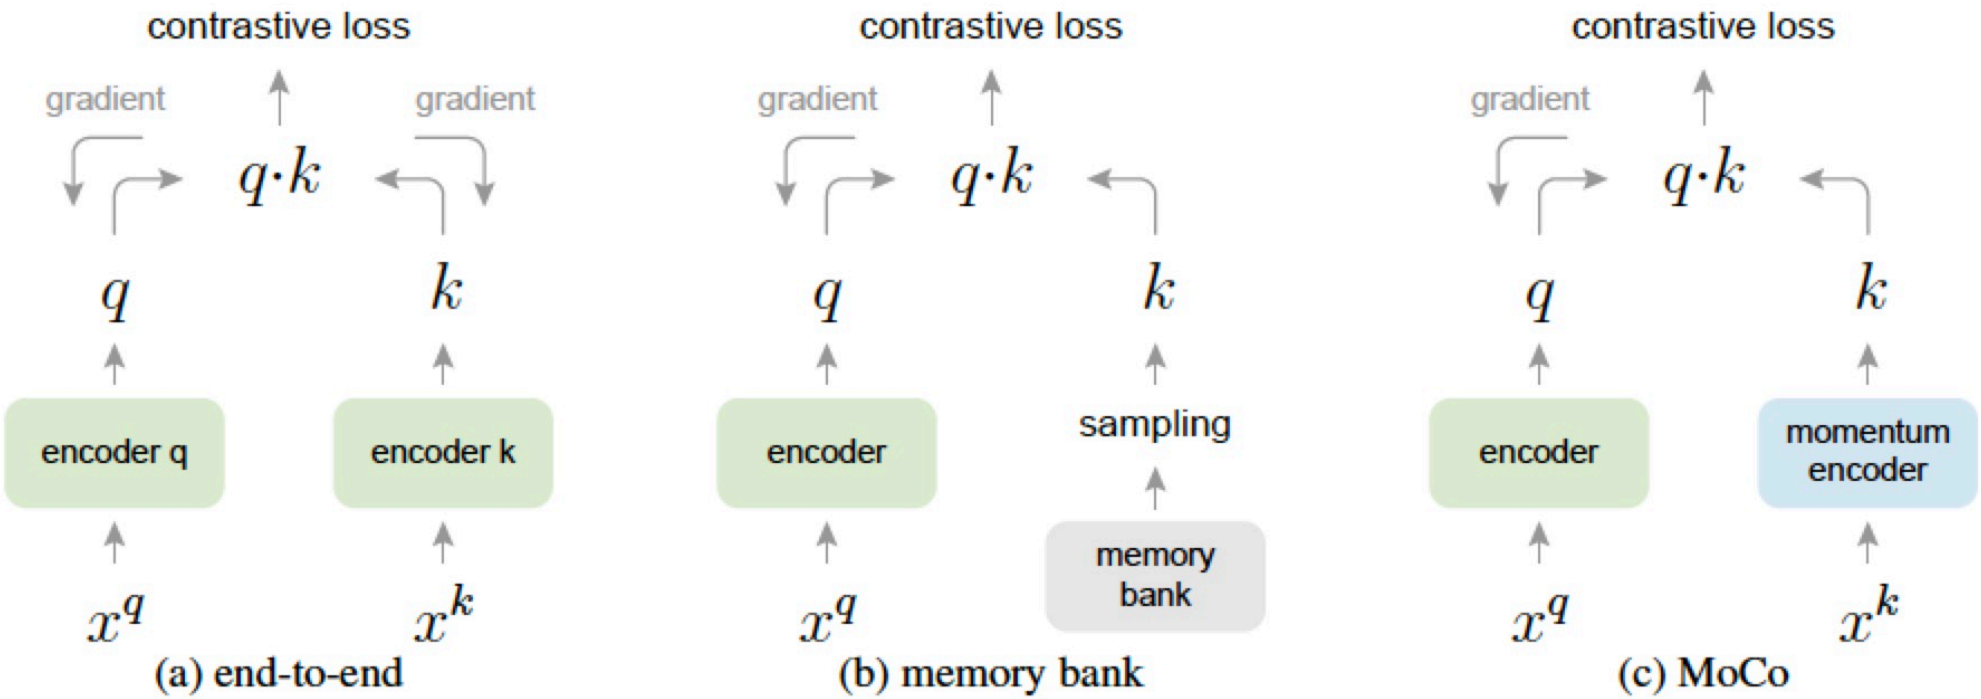

- 无需标记样本，即可学习图形表示
- Momentum contrastive learning
- 效果甚至超过有监督学习结果

https://arxiv.org/abs/1911.05722

面向图数据的自监督学习

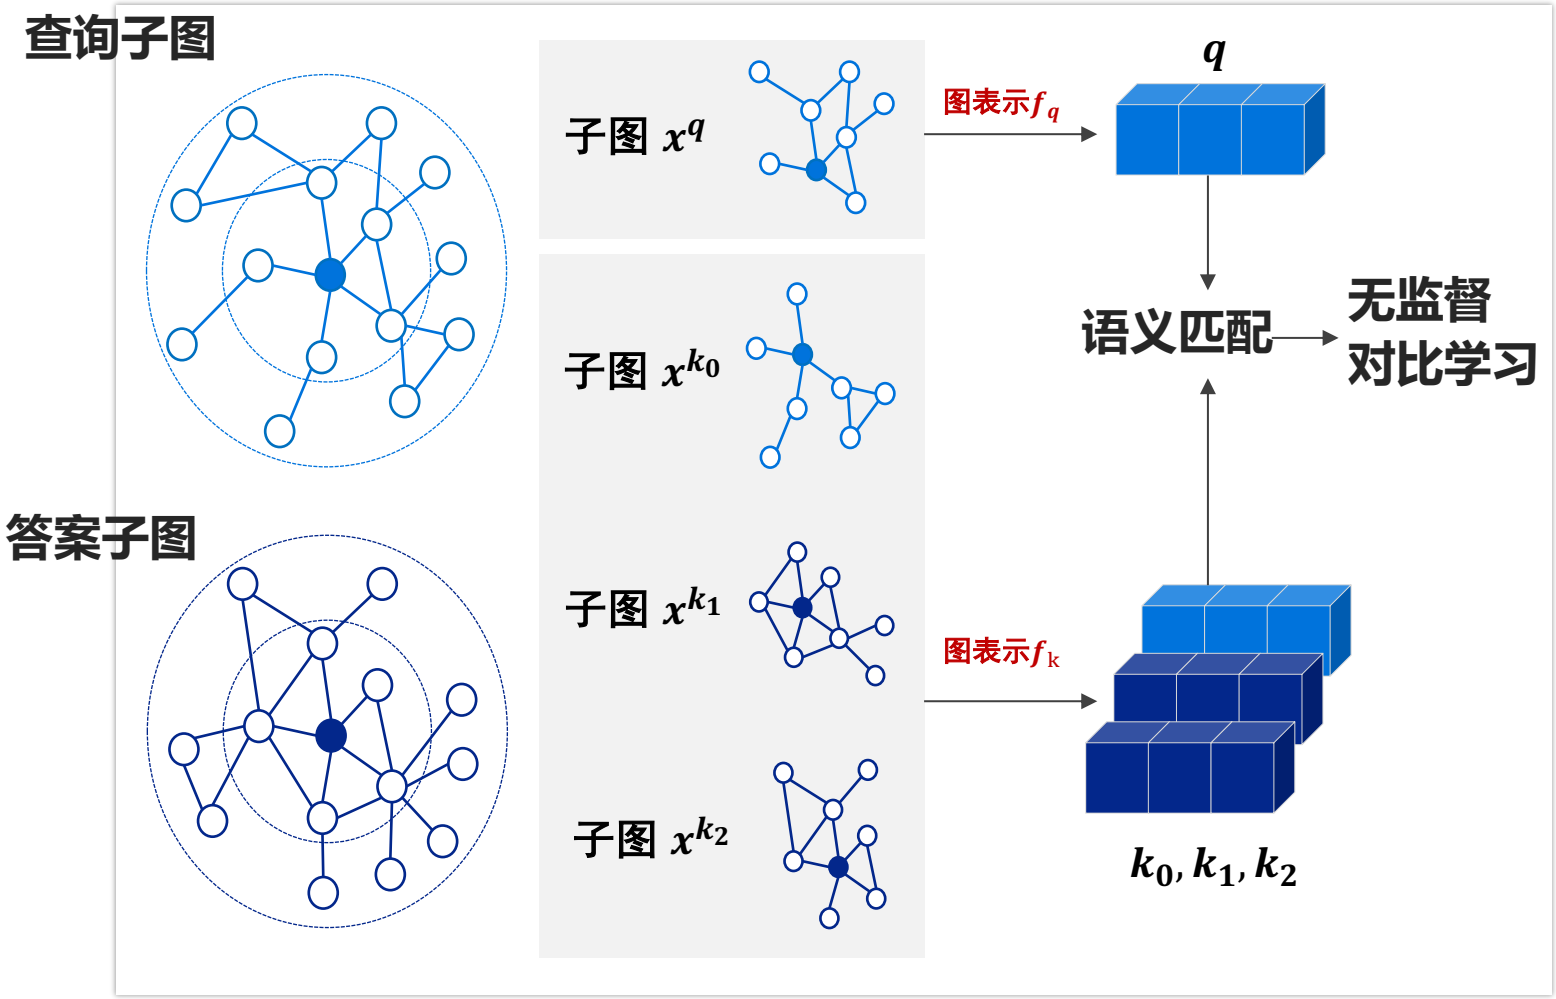

1. J. Qiu, Q. Chen, Y. Dong, J. Zhang, H. Yang, M. Ding, K. Wang, and J. Tang. GCC: Graph Contrastive Coding for Graph Neural Network Pre-Training. *KDD'20*.

人工智能未来...

SEARCH

第三代人工智能

符号AI

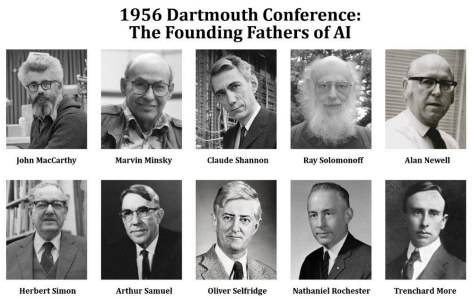

符号模型/规则模型/感知机

第一代

认知智能

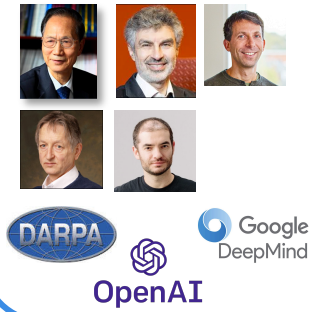

张钹院士2016年提出第三代人工智能雏形，DARPA 2018年发布AI Next计划。核心思路是推进统计与知识推理融合的计算；与脑认知的融合

第三代

第二代

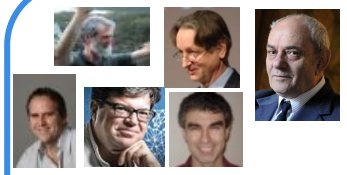

大数据驱动的统计学习方法初步实现了针对文本、图像、语音等的感知与识别

感知智能

目前急需的是高质量超大规模知识图谱（AI的基础设施）以及对知识的理解能力（面向认知的深度学习算法）

脑认知

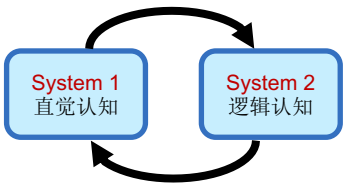

超越深度学习

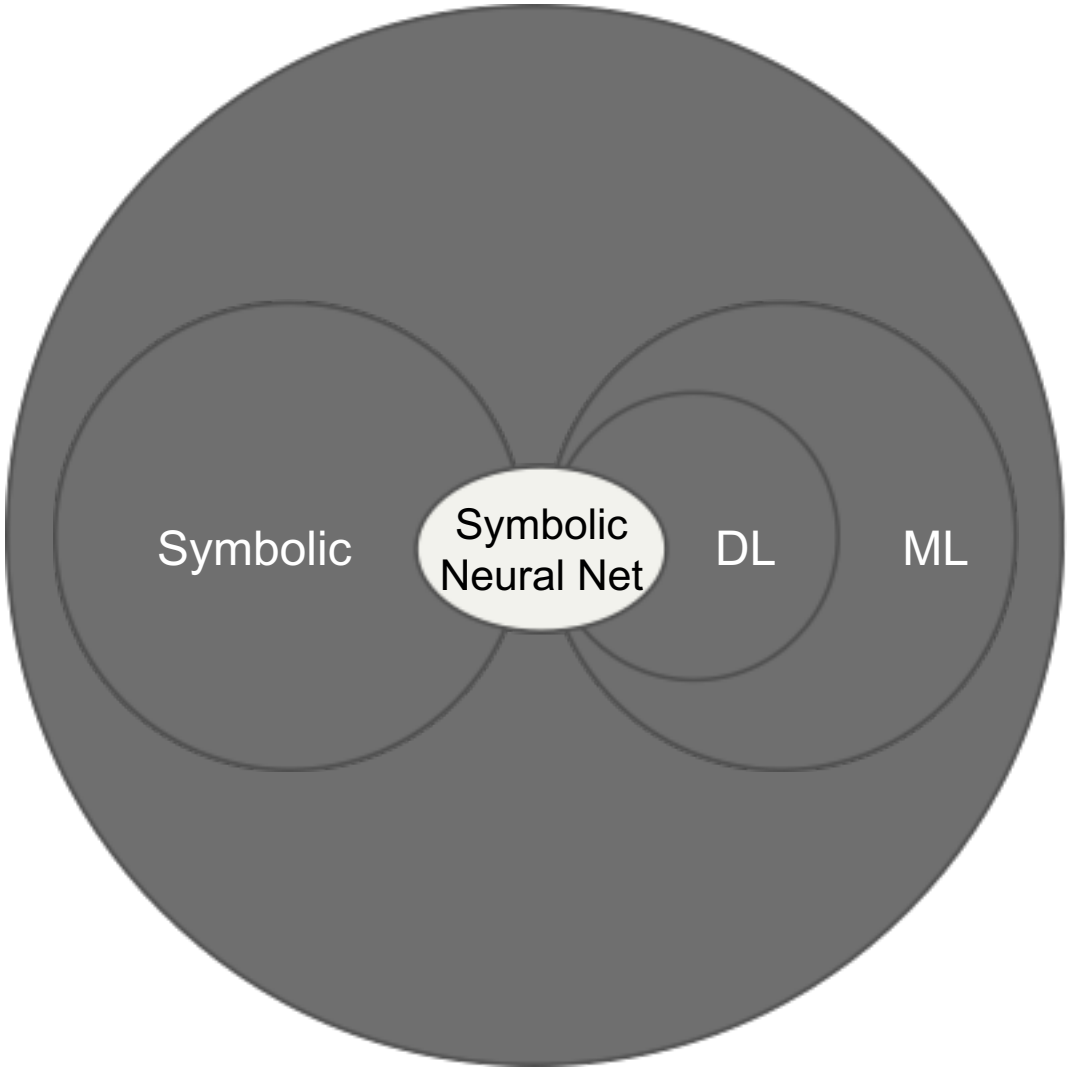

算法是核心，计算、数据是基础

人工智能的核心：算法

算法

实现核心：机器通过**算法**实现人工智能

突破方向：**认知智能**是下一个突破方向

突破途径：人工智能突破主要通过**算法性能的提升**，主要有工程学法和模拟法

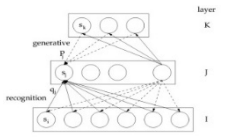

方法一：工程学方法

（Engineering Approach）

采用传统的编程技术，利用大量数据处理经验改进提升算法性能。

方法二：模拟法

（Modeling Approach）

模仿人类或其他生物所用的方法或机理，提升算法性能，例如遗传算法和神经网络。

基础条件：计算和知识

计算能力

- **现状：**使用GPU并行计算神经网络
- **作用：**提升运算速度，降低计算成本
- **未来：**量子计算、速度更快的芯片

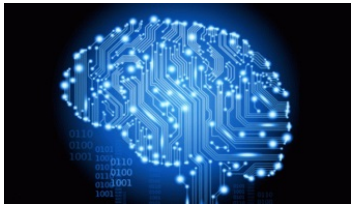

知识

- **现状：**互联网发展积累了片段化知识
- **作用：**训练机器，提升算法性能
- **未来：**面向全世界的常识知识图谱

场景驱动、算法为核心、知识为基础

认知图谱 (Cognitive Graph)

—知识图谱, 认知推理, 逻辑表达

认知图谱：算法与认知的结合

Question: Who is the director of the 2003 film which has scenes in it filmed at the Quality Cafe in Los Angeles?

Quality Café

The Quality Cafe is a now-defunct diner in Los Angeles, California. The restaurant has appeared as a location featured in a number of Hollywood films, including Old School, Gone in 60 Seconds, ...

Los Angeles

Los Angeles is the most populous city in California, the second most populous city in the United States, after New York City, and the third most populous city in North America.

Alessandro Moschitti

Alessandro Moschitti is a professor of the CS Department of the University of Trento, Italy. He is currently a Principal Research Scientist of the Qatar Computing Research Institute (QCRI)

WIKIPEDIA
The Free Encyclopedia

Old School

Old School is a 2003 American comedy film released by Dream Works Pictures and The Montecito Picture Company and directed by Todd Phillips.

Todd Phillips

Todd Phillips is an American director, producer, screenwriter, and actor. He is best known for writing and directing films, including Road Trip (2000), Old School (2003), Starsky & Hutch (2004), and The Hangover Trilogy.

Tsinghua University

Tsinghua University is a major research university in Beijing and dedicated to academic excellence and global development. Tsinghua is perennially ranked as one of the top academic institutions in China, Asia, and worldwide...

算法：BIDAF, BERT, XLNet

- 目标：理解整个文档，而不仅仅是局部片段
- 但仍然缺乏在知识层面上的推理能力

BiDAF

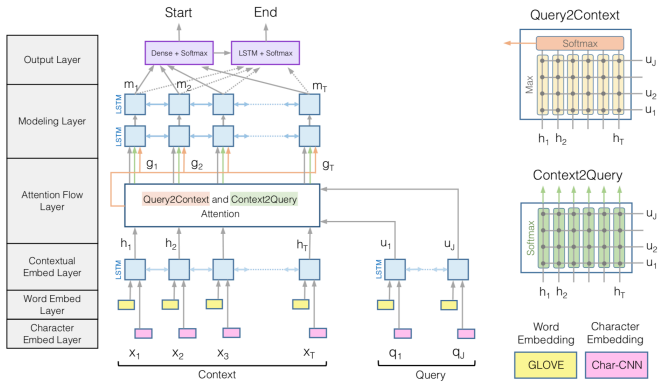

BERT

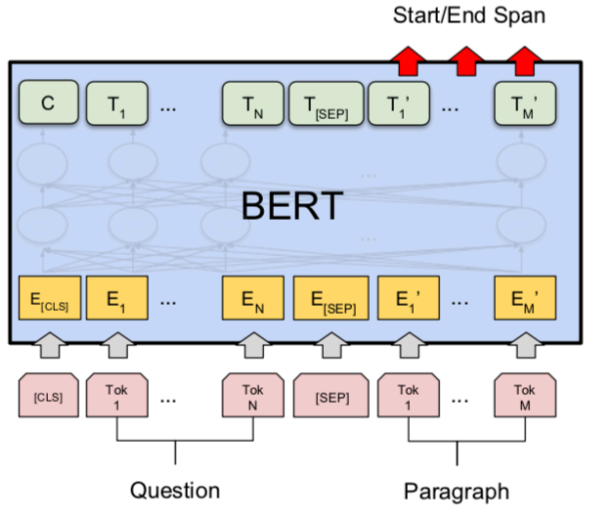

XLNet

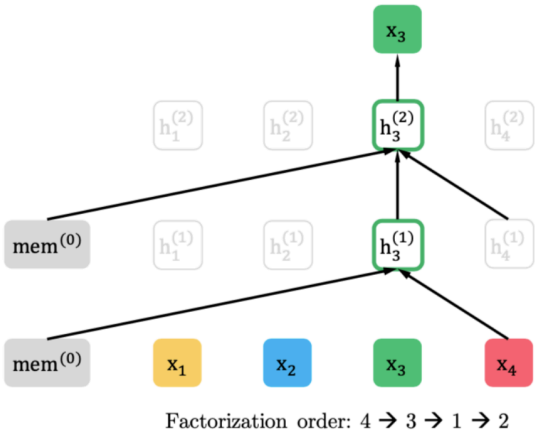

挑战：可解释性

- 大部分阅读理解方法都只能看做**黑盒**：
  - 输入：问题和文档
  - 输出：答案文本块（在文档中的起止位置）

- 如何让用户可以验证答案的对错：
  - 推理路径或者子图
  - 每个推理节点上的支撑事实
  - 用于对比的其他可能答案和推理路径

认知图谱：知识表示，推理和决策

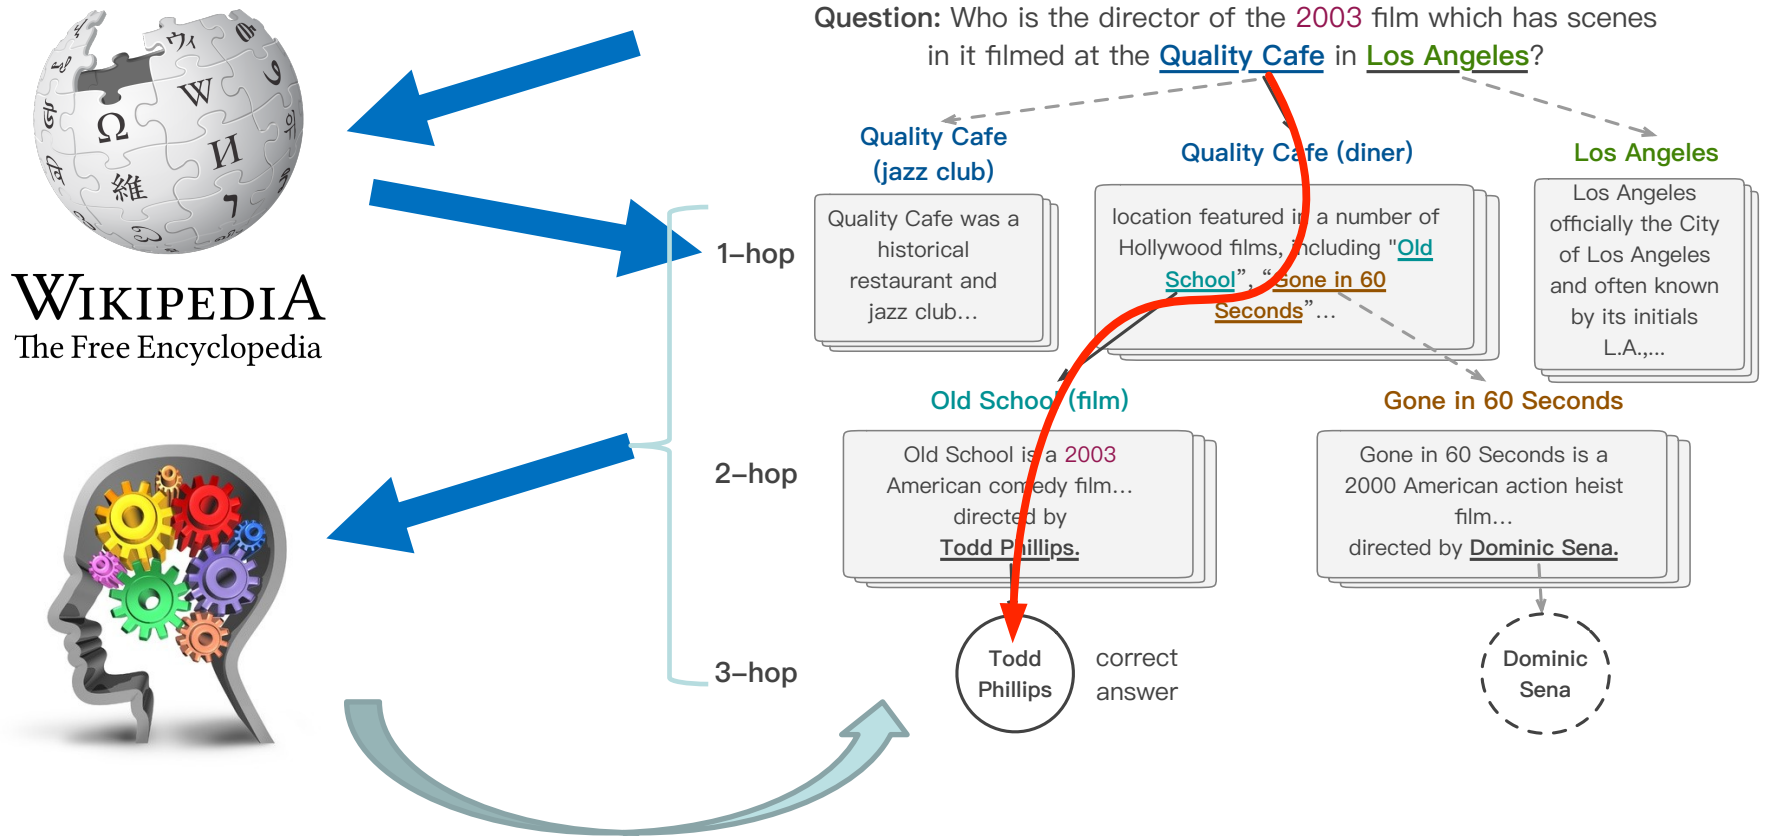

和认知科学的结合

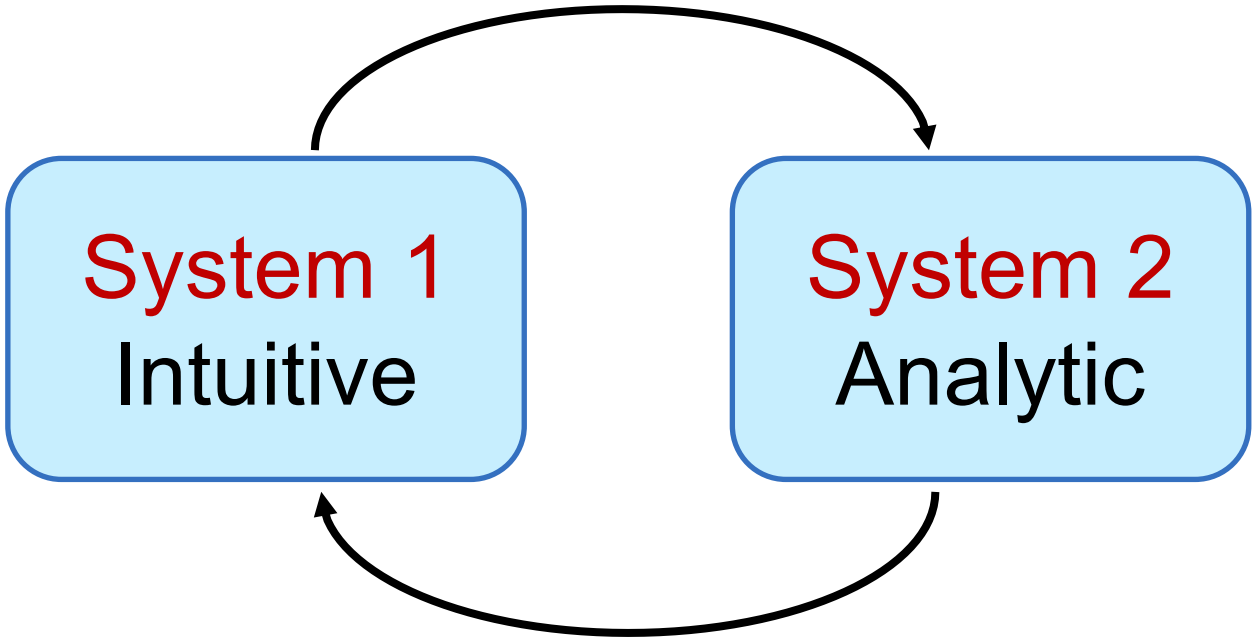

Dual Process Theory (Cognitive Science)

SYSTEM 1 VS. SYSTEM 2 COGNITION

2 systems (and categories of cognitive tasks):

Manipulates high-level / semantic concepts, which can be recombined combinatorially

System 1

- Intuitive, fast, **UNCONSCIOUS**, non-linguistic, habitual
- Current DL

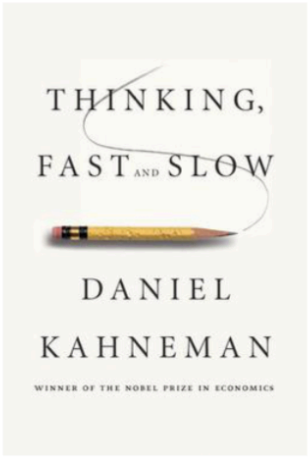

System 2

- Slow, logical, sequential, **CONSCIOUS**, linguistic, algorithmic, planning, reasoning
- Future DL

Mila

3

1. From Bengio's NIPS'2019 Keynote

认知图谱的推理模型

- System 1:
  - Knowledge expansion by association in text when reading
- System 2:
  - Decision making w/ all the information

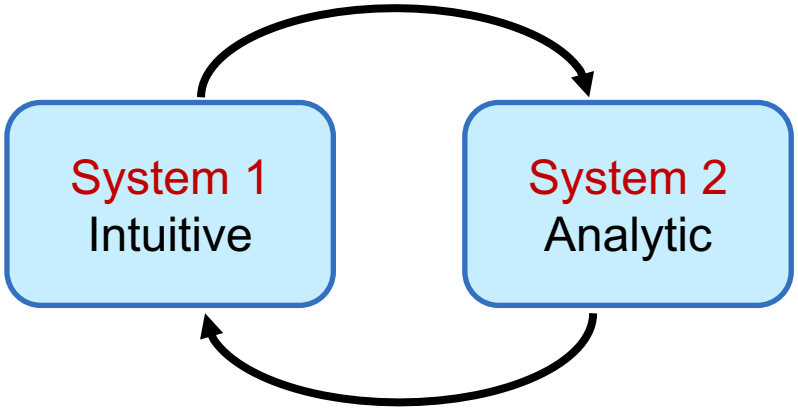

认知图谱的推理模型

- An iterative framework corresponding to dual process theory
- System 1
  - extract entities to build the cognitive graph
  - generate semantic vectors for each node
- System 2
  - Do reasoning based on semantic vectors and graph
  - Feed clues to System 1 to extract next-hop entities

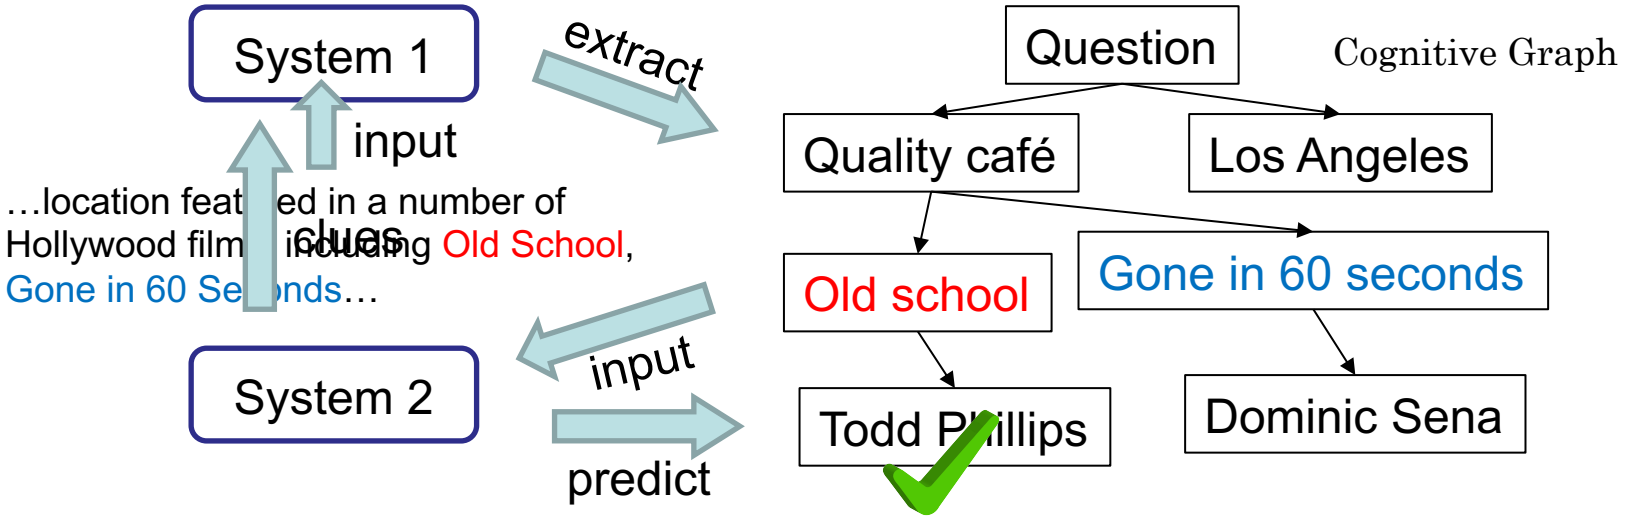

认知图谱的推理模型

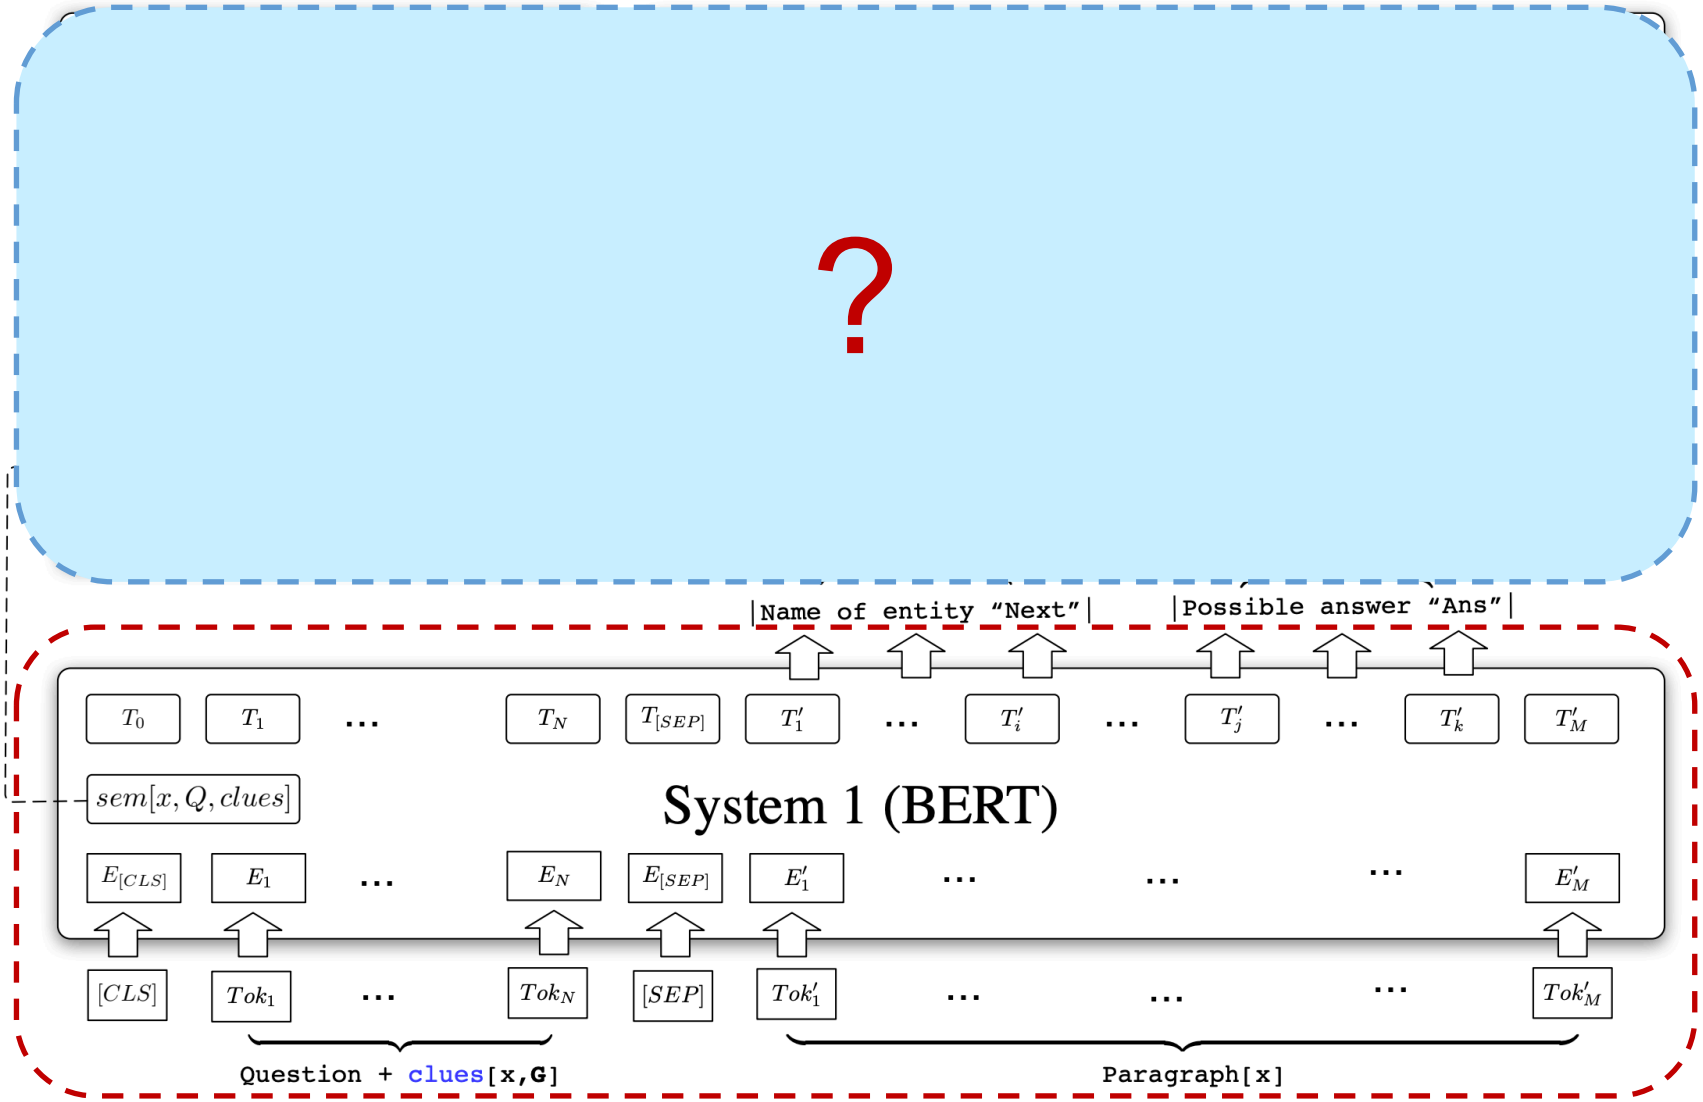

1. M. Ding, C. Zhou, Q. Chen, H. Yang, and J. Tang. Cognitive Graph for Multi-Hop Reading Comprehension at Scale. ACL'19.

认知图谱的推理模型

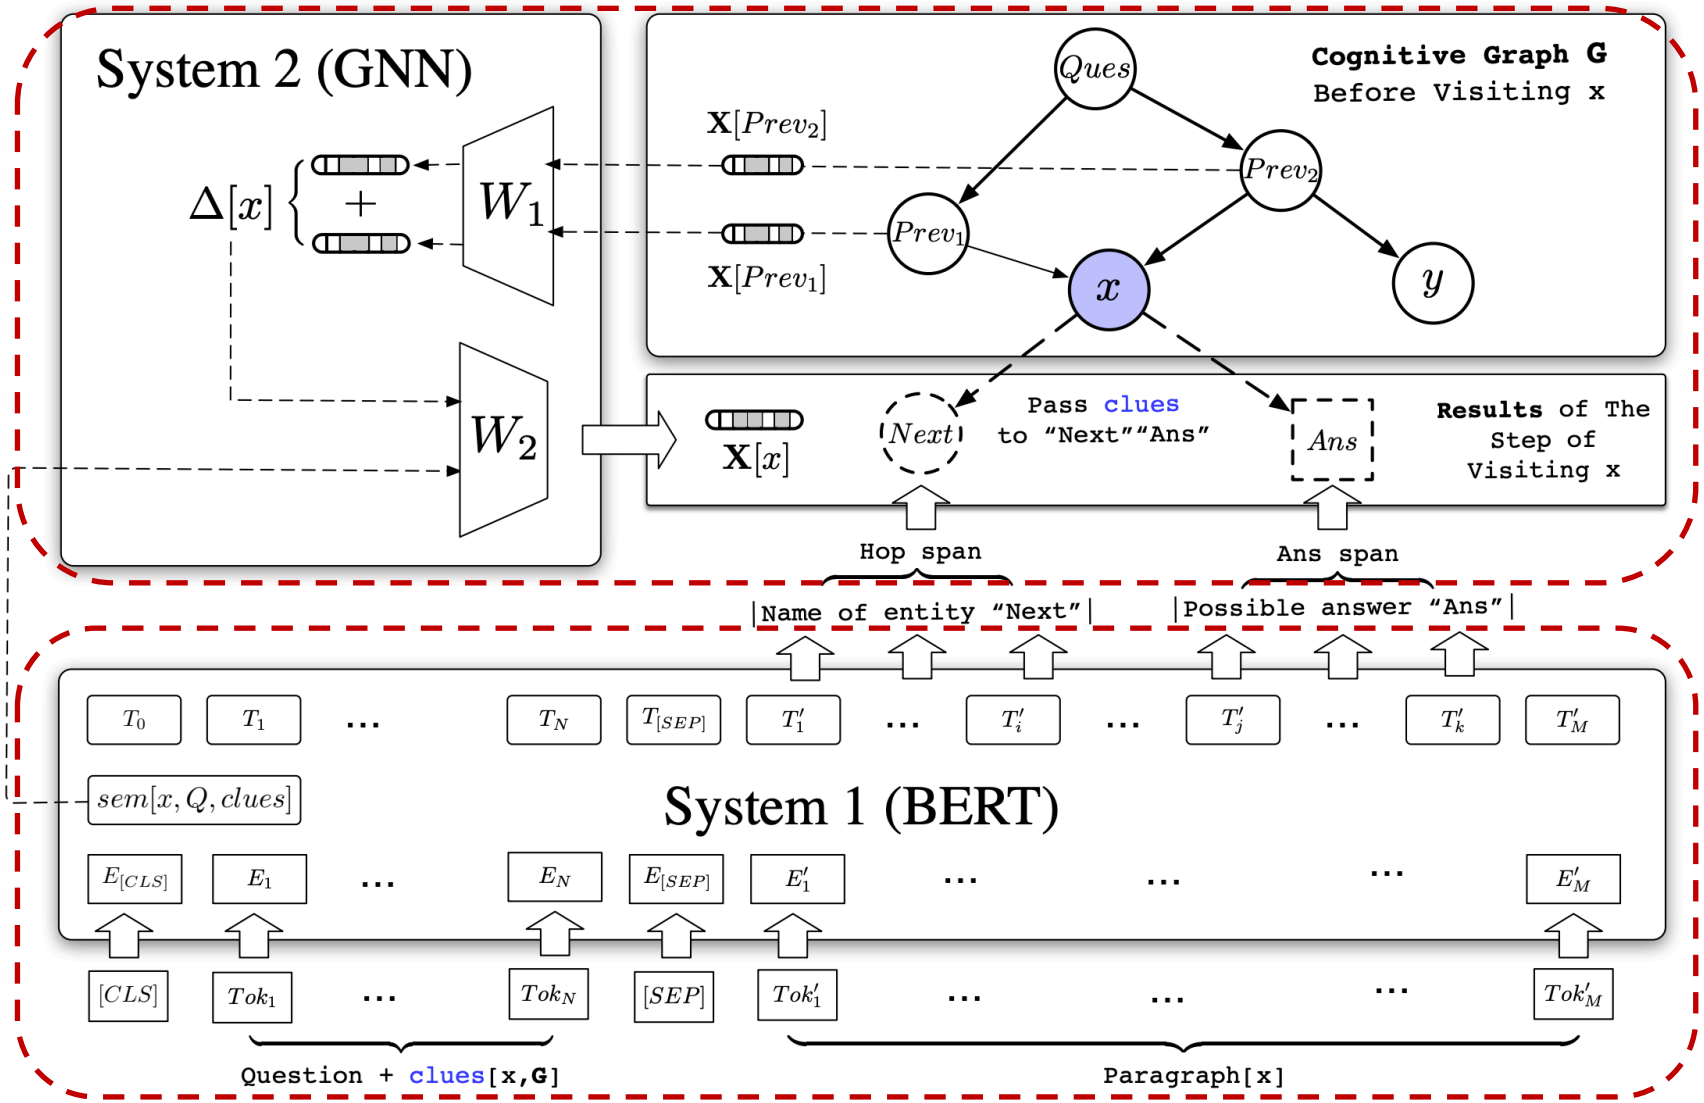

1. M. Ding, C. Zhou, Q. Chen, H. Yang, and J. Tang. Cognitive Graph for Multi-Hop Reading Comprehension at Scale. ACL'19.

认知图谱的推理效果

- HotpotQA is a dataset with leaderboard similar to SQuAD
- CogQA ranked 1st from 21, Feb to 15, May (nearly 3 month)

	Model	Ans				Sup				Joint			
		EM	F_1	Prec	Recall	EM	F_1	Prec	Recall	EM	F_1	Prec	Recall
Dev	Yang et al. (2018)	23.9	32.9	34.9	33.9	5.1	40.9	47.2	40.8	2.5	17.2	20.4	17.8
	Yang et al. (2018)-IR	24.6	34.0	35.7	34.8	10.9	49.3	52.5	52.1	5.2	21.1	22.7	23.2
	BERT	22.7	31.6	33.4	31.9	6.5	42.4	54.6	38.7	3.1	17.8	24.3	16.2
	CogQA-sys1	33.6	45.0	47.6	45.4	**23.7**	58.3	**67.3**	56.2	12.3	32.5	39.0	31.8
	CogQA-onlyR	34.6	46.2	48.8	46.7	14.7	48.2	56.4	47.7	8.3	29.9	36.2	30.1
	CogQA-onlyQ	30.7	40.4	42.9	40.7	23.4	49.9	56.5	48.5	**12.4**	30.1	35.2	29.9
	CogQA	**37.6**	**49.4**	**52.2**	**49.9**	23.1	**58.5**	64.3	**59.7**	12.2	**35.3**	**40.3**	**36.5**
Test	Yang et al. (2018)	24.0	32.9	-	-	3.86	37.7	-	-	1.9	16.2	-	-
	QFE	28.7	38.1	-	-	14.2	44.4	-	-	8.7	23.1	-	-
	DecompRC	30.0	40.7	-	-	N/A	N/A	-	-	N/A	N/A	-	-
	MultiQA	30.7	40.2	-	-	N/A	N/A	-	-	N/A	N/A	-	-
	GRN	27.3	36.5	-	-	12.2	48.8	-	-	7.4	23.6	-	-
	CogQA	**37.1**	**48.9**	-	-	**22.8**	**57.7**	-	-	**12.4**	**34.9**	-	-

Table 1: Results on HotpotQA (fullwiki setting). The test set is not public. The maintainer of HotpotQA only offers EM and F_1 for every submission. N/A means the model cannot find supporting facts.

** Code available at https://github.com/THUDM/CogQA

认知图谱的推理效果

CogQA Performs much better on question with more hops!

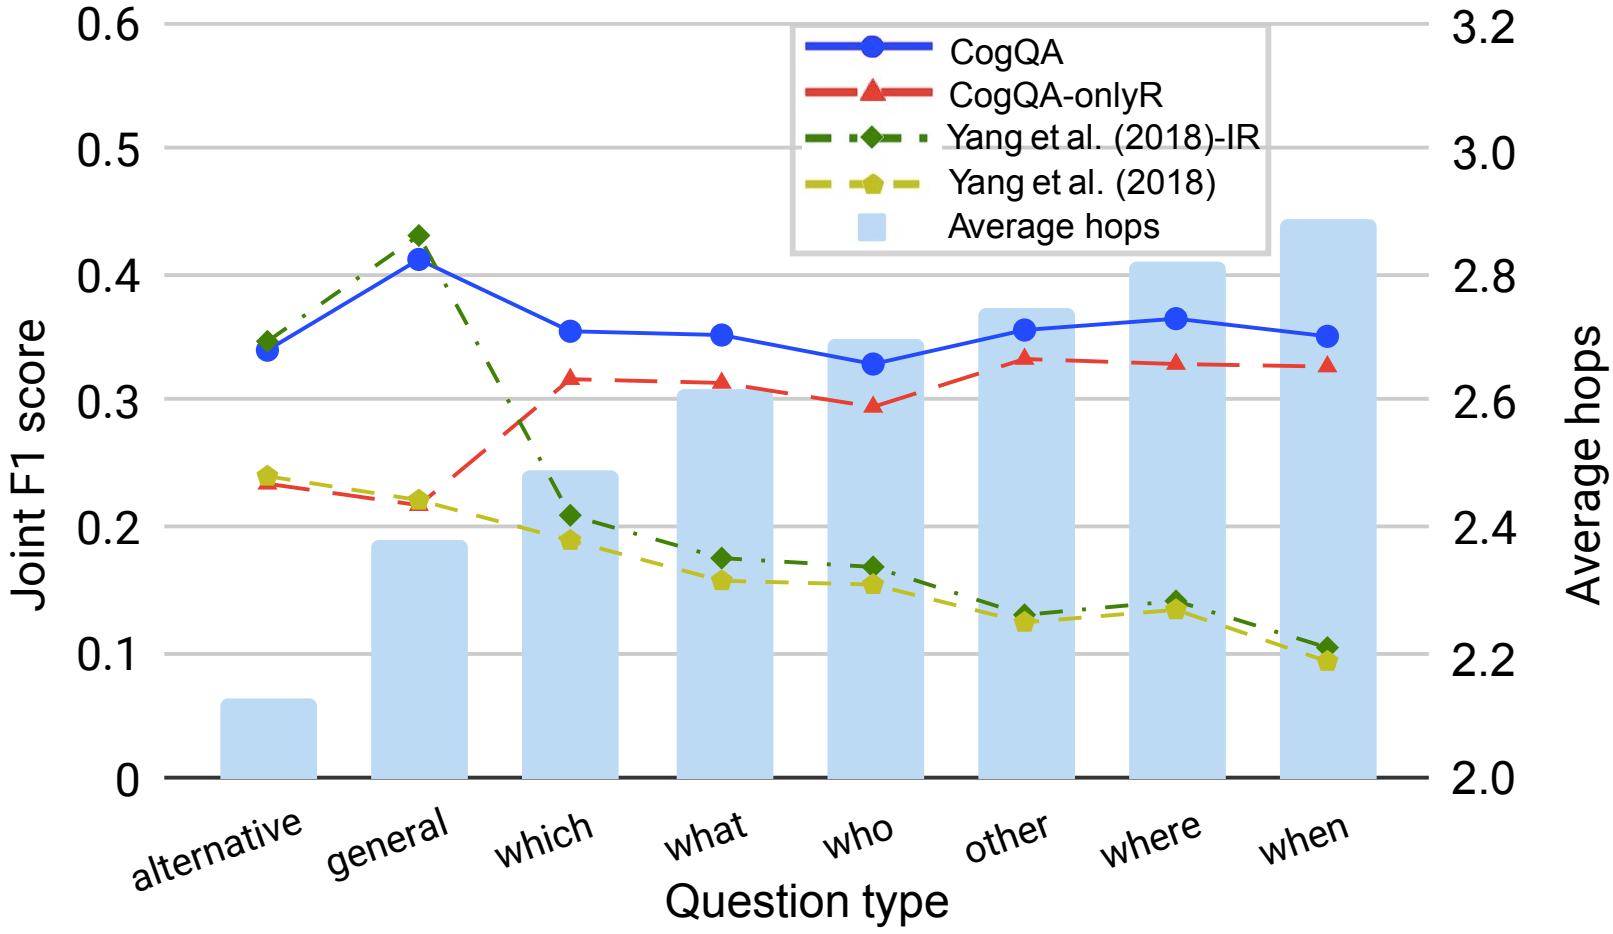

认知图谱的推理效果

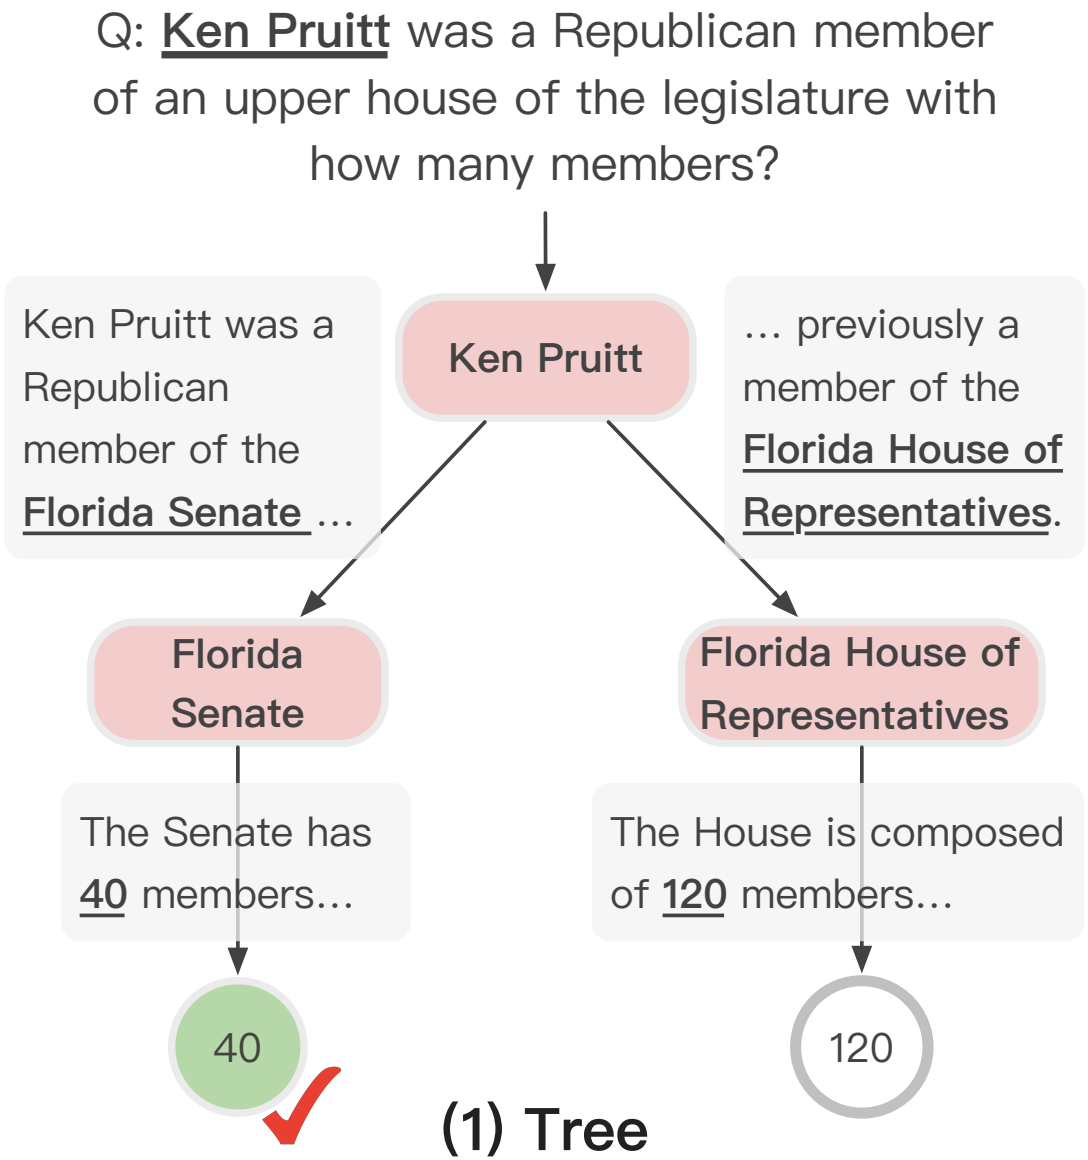

(1) Tree

- Tree-shape Cognitive Graph
- Users can verify the answer by comparing it with another possible reasoning chain.
- "*Upper House*" in the question is similar to "*Senate*" not "*House of Representative*"

认知图谱的推理效果

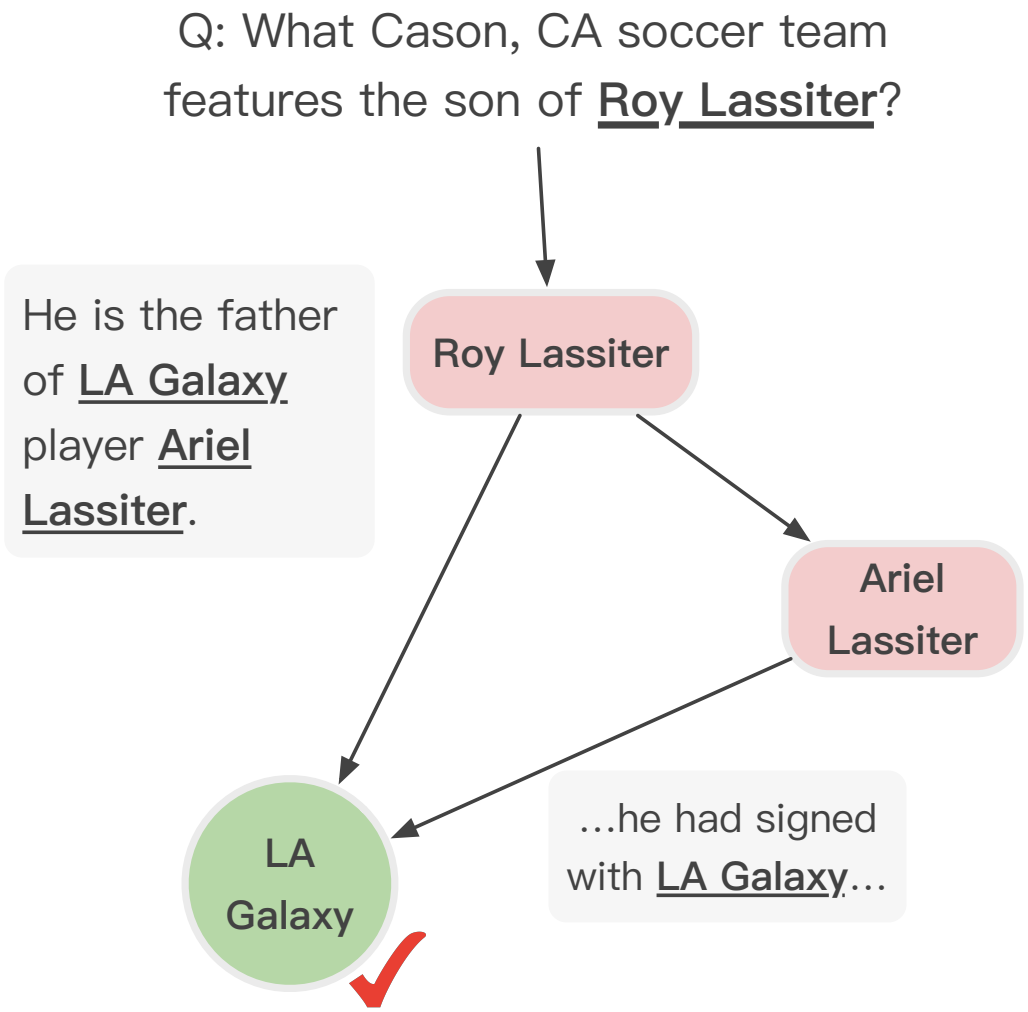

(2) DAG

- DAG-shape Cognitive Graph
- Multiple supporting facts provides richer information, increasing the credibility of the answer.

认知图谱的推理效果

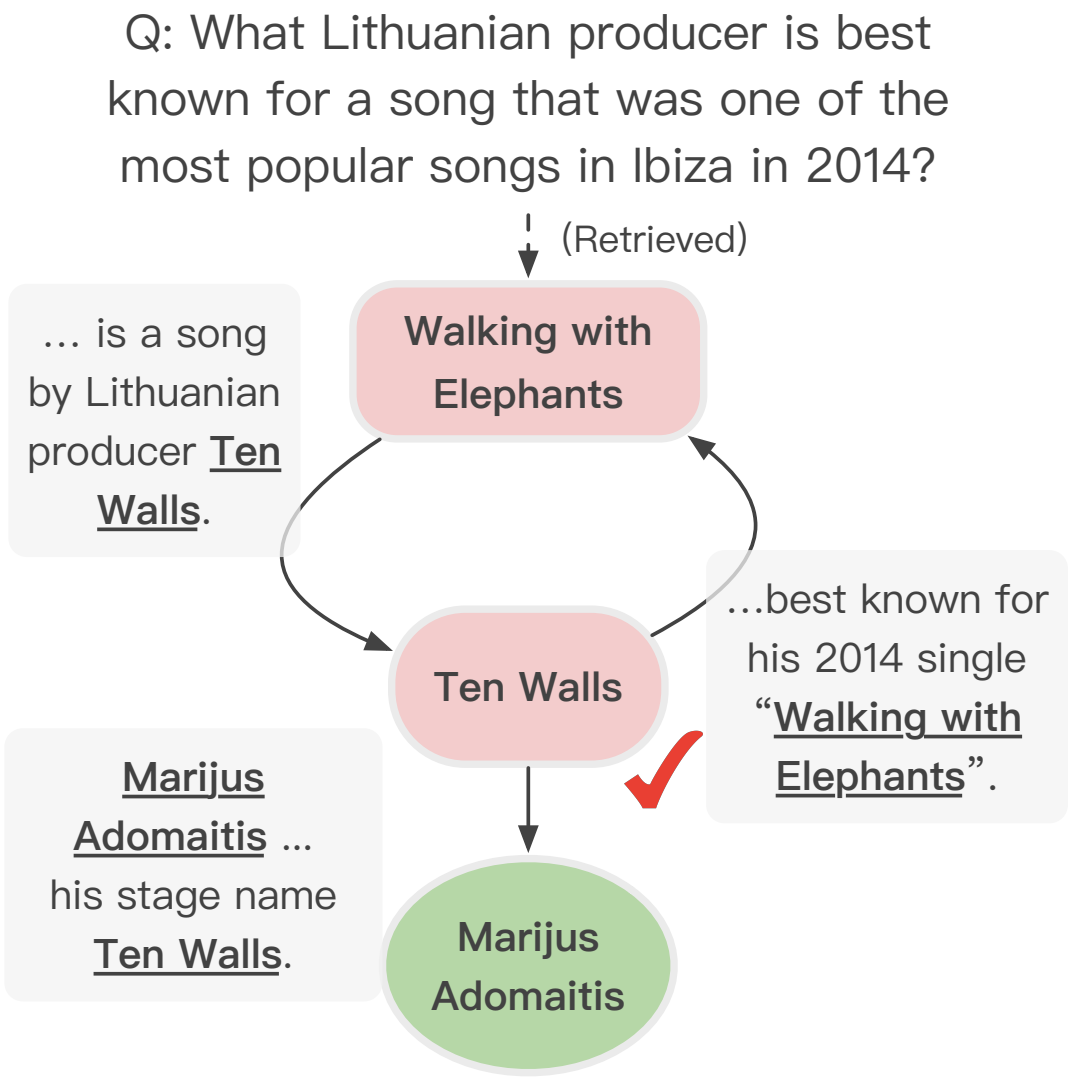

(3) Cyclic Graph

- CogQA gives the answer "Marijus Adomaitis" while the ground truth is "Ten Walls".
- By examining, Ten Walls is just the stage name of Marijus Adomaitis!
- Without cognitive graphs, black-box models cannot achieve it.

Next

挑战与未来(Next 10)

Edward Feigenbaum
Turing Award Winner

认知与推理

—Trillion-scale common-sense knowledge graph

Tim Berners Lee
Turing Award
Winner

Big Data ⇨ Knowledge ⇨ Intelligence

* AI = Knowledge + Intelligence

1. J. Tang, J. Zhang, L. Yao, J. Li, L. Zhang, and Z. Su. ArnetMiner: Extraction and Mining of Academic Social Networks. KDD'08. pp.990-998.

挑战与未来(Next 30)

意识

—让计算机具有自我意识

认知推理 ⇨ 记忆 ⇨ 自我意识

- Next AI = Reasoning + Memory + Consciousness

1. J. Tang, J. Zhang, L. Yao, J. Li, L. Zhang, and Z. Su. ArnetMiner: Extraction and Mining of Academic Social Networks. KDD'08. pp.990-998.

意识：全局工作理论—GWT

所谓的有意识学习，就是有目标的机器学习，给定数据，训练一个模型，通过模型进行分类（决策），所以有意识学习输出是一个短期记忆模型，但这个短期记忆模型比较简单：数据有限、无背景知识。所以无意识就对应着长期记忆模型，长期记忆模型有点类似半监督或者无监督学习模型，或者当下比较流行的预训练和自监督学习；无意识处理对应多种长期记忆，所以无意识可以考虑多个不同的处理器，这些处理器之间可以有链接，也可以没有，很多时候是并行处理，但针对特定任务，比如有意识思考某个问题的时候，形成特定连接，包括无意识处理器（无监督模型）和有意识处理器（有监督模型）之间的连接，这里可以考虑成fine-tune。当然连接权重可以通过外界反馈强化学习来实现。无意识处理器之间的连接以及和有意识处理器之间的连接可以类比为注意力机制。最后值得注意的是长期记忆的构造和实现，人脑记忆保存的是模型图，而不是概念图。每个长期记忆都可能是一个模型，可以生成样本，具体学习方法，可以想象一下是一个层次聚类。通过这样就可以用有监督、无监督、强化、注意力、fine-tune来实现GWT模型。

Related Publications

For more, check http://keg.cs.tsinghua.edu.cn/jietang

- Jiezhong Qiu, Qibin Chen, Yuxiao Dong, Jing Zhang, Hongxia Yang, Ming Ding, Kuansan Wang, and Jie Tang. GCC: Graph Contrastive Coding for Structural Graph Representation Pre-Training. KDD'20.
- Zhen Yang, Ming Ding, Chang Zhou, Hongxia Yang, Jingren Zhou, and Jie Tang. Understanding Negative Sampling in Graph Representation Learning. KDD'20.
- Yukuo Cen, Jianwei Zhang, Xu Zou, Chang Zhou, Hongxia Yang, and Jie Tang. Controllable Multi-Interest Framework for Recommendation. KDD'20.
- Yuxiao Dong, Ziniu Hu, Kuansan Wang, Yizhou Sun and Jie Tang. Heterogeneous Network Representation Learning. IJCAI'20.
- Ming Ding, Chang Zhou, Qibin Chen, Hongxia Yang, and Jie Tang. Cognitive Graph for Multi-Hop Reading Comprehension at Scale. ACL'19.
- Jie Zhang, Yuxiao Dong, Yan Wang, Jie Tang, and Ming Ding. ProNE: Fast and Scalable Network Representation Learning. IJCAI'19.
- Yukuo Cen, Xu Zou, Jianwei Zhang, Hongxia Yang, Jingren Zhou and Jie Tang. Representation Learning for Attributed Multiplex Heterogeneous Network. KDD'19.
- Fanjin Zhang, Xiao Liu, Jie Tang, Yuxiao Dong, Peiran Yao, Jie Zhang, Xiaotao Gu, Yan Wang, Bin Shao, Rui Li, and Kuansan Wang. OAG: Toward Linking Large-scale Heterogeneous Entity Graphs. KDD'19.
- Qibin Chen, Junyang Lin, Yichang Zhang, Hongxia Yang, Jingren Zhou and Jie Tang. Towards Knowledge-Based Personalized Product Description Generation in E-commerce. KDD'19.
- Yifeng Zhao, Xiangwei Wang, Hongxia Yang, Le Song, and Jie Tang. Large Scale Evolving Graphs with Burst Detection. IJCAI'19.
- Yu Han, Jie Tang, and Qian Chen. Network Embedding under Partial Monitoring for Evolving Networks. IJCAI'19.
- Jiezhong Qiu, Yuxiao Dong, Hao Ma, Jian Li, Chi Wang, Kuansan Wang, and Jie Tang. NetSMF: Large-Scale Network Embedding as Sparse Matrix Factorization. WWW'19.
- Jiezhong Qiu, Jian Tang, Hao Ma, Yuxiao Dong, Kuansan Wang, and Jie Tang. DeepInf: Modeling Influence Locality in Large Social Networks. KDD'18.
- Jiezhong Qiu, Yuxiao Dong, Hao Ma, Jian Li, Kuansan Wang, and Jie Tang. Network Embedding as Matrix Factorization: Unifying DeepWalk, LINE, PTE, and node2vec. WSDM'18.
- Jie Tang, Jing Zhang, Limin Yao, Juanzi Li, Li Zhang, and Zhong Su. ArnetMiner: Extraction and Mining of Academic Social Networks. KDD'08.

Thank you !

Jie Tang, KEG, Tsinghua U

Download all data & Codes

http://keg.cs.tsinghua.edu.cn/jietang
https://keg.cs.tsinghua.edu.cn/cogdl/
https://github.com/THUDM