

# AAAI 2020 (第三十四届 AAAI 人工智能大会) 见闻

潘旻琦

1. AAAI 2020 第一天参加了图网络 Tutorial、微分深度学习 Tutorial, 总结如下:

- (a) 图网络有三种滤波: 谱滤波 ChebNet [Defferrard et al., 2016]、GraphSage [Hamilton et al., 2017]、GAT [Veličković et al., 2017]、MPNN [Gilmer et al., 2017], 空域滤波, 以及两者兼备的 GCN [Kipf and Welling, 2016]
- (b) 基于 RNN 思想的图网络有: GNN [Scarselli et al., 2008]、GGs-NN [Li et al., 2015]
- (c) 图网络的池化层有: gPool [Gao and Ji, 2019]、DiffPool [Ying et al., 2018]、Eigenpooling [Ma et al., 2019b]
- (d) 攻击图网络的方法有很多, 例如先开采干净的图再通过迁移学习元学习攻击 GNN 的 PA-GNN [Tang et al., 2019]
- (e) 图网络的应用: 药物相互作用预测 [Ma et al., 2019a]、药物推荐 [Shang et al., 2019]
- (f) Graph2Seq [Xu et al., 2018] 把 Seq2Seq 推广到了图的嵌入, 该课题的后续工作出现了新课题 (询问生成 [Chen et al., 2019c]) 和新模型 (Graph2Graph [Mohammadshahi and Henderson, 2019] 和图匹配 [Li et al., 2019])
- (g) ResNet [He et al., 2016] 可进化到 ODE-Net [Chen et al., 2018d]、RNN 可进化到微分方程 RNN [Pearlmutter, 1995]、Normalizing Flows [Rezende and Mohamed, 2015] 可进化到微分方程 Flow [Chen et al., 2018d]、NICE [Dinh et al., 2014] 可进化到微分 NICE [Dinh et al., 2016]
- (h) 把 Flow 用在 GNN 上生成新图的方法有: NICE [Dinh et al., 2014]、Real NVP [Dinh et al., 2016]、GraphNVP [Madhawa et al., 2019]、Glow [Kingma and Dhariwal, 2018]
- (i) 复杂图上的神经动力学 NDCN [Zang and Wang, 2019] 统一了连续时间网络动态预测问题、结构化序列预测问题、半监督节点回归/分类问题这三个问题

- (j) 臧承熙等人发现了一个构造样本数据背后的概率分布的用微分方程描述的动力系统的定理 [Zang et al., 2019]
2. AAAI 2020 第二天参加了 XAI Tutorial、首届图深度学习 Workshop, 总结如下:
- (a) 基于注意力机制的 XAI [Choi et al., 2016]
  - (b) 基于原型层的 XAI [Li et al., 2018]
  - (c) XAI 图像识别网络 ProtoPNet [Chen et al., 2018a]
  - (d) 几个归因 XAI [Guidotti et al., 2018] 方法: 积分梯度 [Sundararajan et al., 2017a] [Sundararajan et al., 2017b]、切除法 [Baehrens et al., 2010]、Score 反向传播 LRP [Lapuschkin et al., 2016] Guided BackProp [Springenberg et al., 2014] DeepLift [Shrikumar et al., 2017]
  - (e) 几个计算机视觉领域的 XAI: 网络解剖可解释单元 [Bau et al., 2017]、视觉说明 [Hendricks et al., 2016]、不确定性图 [Kendall and Gal, 2017]、显著性图 [Adebayo et al., 2018]
  - (f) 基于夏普利值 [Shapley, 1953] [Aumann and Shapley, 2015] 的 XAI: [Merrick and Taly, 2019] [Lundberg and Lee, 2017] [Chen et al., 2018c] [Datta et al., 2016] [Kononenko et al., 2010] [Chalkiadakis et al., 2011]
  - (g) 基于因果推断的 XAI [Kononenko et al., 2010] [Datta et al., 2016]
  - (h) 机器人方面的 XAI [Rosenthal et al., 2016] [Brooks et al., 2010] [Sheh, 2017]
  - (i) 基于知识库的迁移学习 XAI [Chen et al., 2018b] [Lecue et al., 2019]
  - (j) 其他个体预测解释 XAI 方法: [Ribeiro et al., 2016] [Ribeiro et al., 2018] [Koh and Liang, 2017] [Kim et al., 2016]
  - (k) 臧承熙等人昨天介绍的把层数看作连续时间用稠密层数神经网络实现端到端的微分方程组学习的对复杂图的连续时间非线性动态的研究获得该 Workshop 最佳论文 [Zang and Wang, 2019]
  - (l) 东京大学又提出一个稠密层数图网络 GDE [Poli et al., 2019]
3. AAAI 2020 第三天参加了 ML Interpretability、Fair Allocation I、ML Causal Learning and Bayes Nets、Robotics and Planning 四个分会场, 总结如下:
- (a) Jilles Vreeken 和他的学生提出一种可解释的数据分解 [Dalleiger and Vreeken, 2020], 用以发现数据集的组成部分
  - (b) 可视化深层网络 [Qi et al., 2019b] 用到了昨天 XAI Tutorial 提到的积分梯度
  - (c) [Tomsett et al., 2019] 用到了昨天 XAI Tutorial 提到的显著性图, 他们提出了显著性指标的健康健全性检查, 研究各种各样的显著性图到底哪个好

- (d) Exploration-Exploitation 考虑的是在线情况下的期望奖励最大化, 而 [Faury et al., 2019] 考虑的是离线的评价历史政策并提出改进, 他们使用分布鲁棒性优化 (DRO) 解决反事实风险最小化 (CRM) 问题
  - (e) CO 矿业学院的 Peng Gao 团队在做 SLAM 的闭环检测问题, 方法是用图表示地标与背景视觉信息及其之间的距离, 然后做图匹配
  - (f) Hinton 团队提出了 2019 版堆叠式胶囊自动编码器 [Kosiorrek et al., 2019] 用来解决 CNN 不面向对象、无坐标系标架、不能处理旋转、缩放等视角变换等问题, 而且用到了 Set Transformer [Lee et al., 2018]
  - (g) 杨立昆提出了深度学习三大挑战——自监督学习、学习推理、规划复杂的行动序列; 自监督学习在计算机视觉领域的图片完形填空上效果不太好 [Pathak et al., 2016], 视频完形填空尝试过对抗训练 [Mathieu et al., 2015], 他认为应该用潜变量 EBM 实现自监督学习, 而且自监督学习是深度学习的未来; 他认为第二个挑战可以把逻辑推理用向量、连续函数来表示, 然后用潜变量 EBM 能量最小化来解决; 第三个挑战他也没有想到办法
  - (h) Bengio 研究了如何实现第二系统认知, 与杨立昆提出的第二挑战类似; 反观现状, 深度学习已经可以一层一层提高认知能力 [Pascanu et al., 2013] [Montufar et al., 2014], 但是深度学习要达到更高的智能必须实现分布外 (Out of Distribution) 泛化与迁移 [Lake and Baroni, 2017], 目前的研究在这方面不太成功 [Bahdanau et al., 2018]; 他认为注意力机制是解决问题的关键 [Bahdanau et al., 2014], 并且还要发现因果结构并加以利用 [Schölkopf et al., 2012] [Bengio et al., 2019][Ke et al., 2019][Goyal et al., 2019]
4. AAAI 2020 第四天参加了 Vision Object Detection II、Logic and KR Inference and Reasoning II、ML Unsupervised and Semi-Supervised Learning Clustering III、ML Domain Adaptation Adversarial Learning Transfer Learning Semi-Supervised Learning、Vision Deep Learning Methods 五个分会场, 总结如下:
- (a) 瑞士洛桑的 Aude G. Billard 团队让机器人手臂学习接住空中抛物, 提到了手臂可行姿势的概率分布 [Huang and Roumeliotis, 2013]、学习手臂稳定姿势的流形 [Li et al., 2014], 后来有学生实现了两个手臂一起接物 [Mirrazavi Salehian et al., 2016], ICRA 2020 将展示他们的两只手抓物体的同时保持身体平衡的工作
  - (b) [Zhang et al., 2019] 以空间语义网络调制改进联合分割
  - (c) F<sup>3</sup>Net 考察了图像不同区域的显著物体检测学习的难度不一样的问题 [Wei et al., 2019]

- (d) 哈佛的刘小泮等人考察了机器人视觉的语义分割的交叉熵损失函数不能区分物体类别重要性的问题
  - (e) HyGNN [Luo et al., 2020] 用 GNN 做人群计数
  - (f) 复旦的 Zhaolong Zhang 等人用 GCN 做零样本学习图像检索
  - (g) [Chen et al., 2020] 以更细粒度的加权改进多视图聚类
  - (h) 西安交大的 Xiaoyu Tao 等人研究了两目标持续学习，学新的不能忘旧的
  - (i) HoMM [Chen et al., 2019a] 提出无监督域适配的高阶统计对齐
  - (j) [Pareja et al., 2019] 用 RNN 实现 GCN 动态化，动态进化 GCN 的参数
  - (k) [Huang et al., 2019] 让不同的网络层共享一个注意力模块
  - (l) Henry Kautz 与昨天杨立昆、Bengio 的研究兴趣类似，谈了融合符号推理与神经网络的几种构想：符号到网络到符号、网络嵌入符号、网络  $\cup$  编译了的符号、网络  $\rightarrow$  符号、符号嵌入网络
5. AAAI 2020 第五天参加了 Vision Object Detection III、ML Neural Nets Theory Models and Algorithms IV 两个分会和海报展会，总结如下：
- (a) [Chen et al., 2019b] 以互信息最大化估计特征之间的关系来改进显著性检测
  - (b) [Kim et al., 2019] 号称第一次用脉冲神经网络做目标检测
  - (c) TX 农工的 Hongyang Gao 提出了新的 ReLU 激活函数解决神经元死亡问题
  - (d) VA 理工的 Zhou Zhou、Klan Hamedani 分享了两个 Reservoir Computing 方面的课题
  - (e) GraLSP [Jin et al., 2019] 通过随机 Anonymous Walk 将局部结构模式显式地纳入 GNN 邻域聚合中
  - (f) 浙大的 Zhou Sheng 的 DGE 在做图嵌入的时候把通用性和个性化这两个概念提出来研究
  - (g) 中科院的迟程一次性分享了两篇论文，一篇是在现有基于区域的检测网络之上引入遮挡处理能力的在拥挤的场景中检测行人的 PedHunter [Chi et al., 2019a]，另一篇 JointDet [Chi et al., 2019b] 探索了头部检测和人体检测两个问题的固有相关性
  - (h) 微观经济学家 Susan Athey 介绍了消费者购物时的选择与反事实推断
  - (i) SONY 的北野宏明介绍 RoboCup 竞赛时提到 Hector SLAM——ROS 社区中使用最广泛的开源 SLAM
  - (j) [Wang et al., 2019] 用图网络做协同过滤，但其实是个二分图，没用到太多的图网络

- (k) Graph-Hist [Magelinski et al., 2019] 用可微的软直方图层把 GCN 和 CNN 连起来
  - (l) 东北大学的王莅尘分享了半监督多标签之间相关性的融合, 他说 Label Correlation Graph 百试不爽
6. AAAI 2020 第六天参加了 ML Probabilistic Methods II、NLP Semantics and Summarization IV、NLP Relational Learning III 三个分会场, 总结如下:
- (a) UCLA 的 [Nijkamp et al., 2019] 发现用 CNN 做能量函数的 EBM 的最大似然学习用 MCMC 跑, 跑太久会出现过饱和, 前人的工作虽然图漂亮但没有收敛到真正的分布
  - (b) MIT 的 [Tan et al., 2019] 提出一种深度马尔可夫模型的分解推理方法做多模态时间序列完形填空
  - (c) 曼海姆的 Anne Lauscher 认为词向量嵌入的性别和种族偏向有隐性偏向和显性偏向之分 [Lauscher et al., 2019], 提出了可以针对两种偏向的新的去偏向框架, 他的实验还尝试了德西意俄克土六种语言之间的迁移学习, 非常酷
  - (d) 甲骨文的 Daniel Peterson 等人研究了怎么对动词进行聚类 and 词义归纳, 前人的做法是纯无监督的聚类, 他们的做法是对一小部分样本进行人工监督加标签, 这一点小改动居然导致结果出现十几个百分点的提升
  - (e) 清华的 [Qi et al., 2019a] 预测 BabelNet 中的那些不同语言中的同义词集的词原, 构造一个跨语言的词原知识库
  - (f) 腾讯的 [Jin et al., ] 提出了完全依存森林数据结构; 前人的工作只用了树, 他们的工作则更进一步用了森林, 这个森林包含了全部可能的词与词的依存关系、全部的句法分析的信息, 他们再把这个森林编码成三维的 tensor, 再拿 CNN 从这些 tensor 上获取特征, 最后接了一个逻辑回归上实现关系抽取

## 参考文献

- [Adebayo et al., 2018] Adebayo, J., Gilmer, J., Muelly, M., Goodfellow, I., Hardt, M., and Kim, B. (2018). Sanity checks for saliency maps. In *Advances in Neural Information Processing Systems*, pages 9505–9515.
- [Aumann and Shapley, 2015] Aumann, R. and Shapley, L. (2015). *Values of Non-Atomic Games*. Princeton Legacy Library. Princeton University Press.
- [Baehrens et al., 2010] Baehrens, D., Schroeter, T., Harmeling, S., Kawanabe, M., Hansen, K., and MÅžller, K.-R. (2010). How to explain individual classification decisions. *Journal of Machine Learning Research*, 11(Jun):1803–1831.

- [Bahdanau et al., 2014] Bahdanau, D., Cho, K., and Bengio, Y. (2014). Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- [Bahdanau et al., 2018] Bahdanau, D., Murty, S., Noukhovitch, M., Nguyen, T. H., de Vries, H., and Courville, A. C. (2018). Systematic generalization: What is required and can it be learned? *CoRR*, abs/1811.12889.
- [Bau et al., 2017] Bau, D., Zhou, B., Khosla, A., Oliva, A., and Torralba, A. (2017). Network dissection: Quantifying interpretability of deep visual representations. In *Computer Vision and Pattern Recognition*.
- [Bengio et al., 2019] Bengio, Y., Deleu, T., Rahaman, N., Ke, R., Lachapelle, S., Bilaniuk, O., Goyal, A., and Pal, C. (2019). A meta-transfer objective for learning to disentangle causal mechanisms. *arXiv preprint arXiv:1901.10912*.
- [Brooks et al., 2010] Brooks, D., Shultz, A., Desai, M., Kovac, P., and Yanco, H. A. (2010). Towards state summarization for autonomous robots. In *2010 AAAI Fall Symposium Series*.
- [Chalkiadakis et al., 2011] Chalkiadakis, G., Elkind, E., and Wooldridge, M. (2011). Computational aspects of cooperative game theory. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 5(6):1–168.
- [Chen et al., 2019a] Chen, C., Fu, Z., Chen, Z., Jin, S., Cheng, Z., Jin, X., and Hua, X.-S. (2019a). Homm: Higher-order moment matching for unsupervised domain adaptation. *arXiv preprint arXiv:1912.11976*.
- [Chen et al., 2018a] Chen, C., Li, O., Barnett, A., Su, J., and Rudin, C. (2018a). This looks like that: deep learning for interpretable image recognition. *CoRR*, abs/1806.10574.
- [Chen et al., 2019b] Chen, C., Sun, X., Hua, Y., Dong, J., and Xu, H. (2019b). Learning deep relations to promote saliency detection. In *34th AAAI Conference on Artificial Intelligence (AAAI-20): Proceedings*, Proceedings of the AAAI Conference on Artificial Intelligence. Association for the Advancement of Artificial Intelligence (AAAI).
- [Chen et al., 2018b] Chen, J., Lécué, F., Pan, J. Z., Horrocks, I., and Chen, H. (2018b). Knowledge-based transfer learning explanation. In *Sixteenth International Conference on Principles of Knowledge Representation and Reasoning*.

- [Chen et al., 2018c] Chen, J., Song, L., Wainwright, M. J., and Jordan, M. I. (2018c). L-shapley and c-shapley: Efficient model interpretation for structured data. *arXiv preprint arXiv:1808.02610*.
- [Chen et al., 2020] Chen, M.-S., Huang, L., Wang, C.-D., and Huang, D. (2020). Multi-view clustering in latent embedding space. In *Proc. of AAAI Conference on Artificial Intelligence*.
- [Chen et al., 2018d] Chen, T. Q., Rubanova, Y., Bettencourt, J., and Duvenaud, D. (2018d). Neural ordinary differential equations. *CoRR*, abs/1806.07366.
- [Chen et al., 2019c] Chen, Y., Wu, L., and Zaki, M. J. (2019c). Reinforcement learning based graph-to-sequence model for natural question generation. *arXiv preprint arXiv:1908.04942*.
- [Chi et al., 2019a] Chi, C., Zhang, S., Xing, J., Lei, Z., Li, S. Z., and Zou, X. (2019a). Pedhunter: Occlusion robust pedestrian detector in crowded scenes. *arXiv preprint arXiv:1909.06826*.
- [Chi et al., 2019b] Chi, C., Zhang, S., Xing, J., Lei, Z., Li, S. Z., and Zou, X. (2019b). Relational learning for joint head and human detection. *arXiv preprint arXiv:1909.10674*.
- [Choi et al., 2016] Choi, E., Bahadori, M. T., Sun, J., Kulas, J., Schuetz, A., and Stewart, W. (2016). Retain: An interpretable predictive model for healthcare using reverse time attention mechanism. In *Advances in Neural Information Processing Systems*, pages 3504–3512.
- [Dalleiger and Vreeken, 2020] Dalleiger, S. and Vreeken, J. (2020). Explainable data decompositions.
- [Datta et al., 2016] Datta, A., Sen, S., and Zick, Y. (2016). Algorithmic transparency via quantitative input influence: Theory and experiments with learning systems. In *2016 IEEE Symposium on Security and Privacy (SP)*, pages 598–617.
- [Datta et al., 2016] Datta, A., Sen, S., and Zick, Y. (2016). Algorithmic transparency via quantitative input influence: Theory and experiments with learning systems. In *2016 IEEE symposium on security and privacy (SP)*, pages 598–617. IEEE.
- [Defferrard et al., 2016] Defferrard, M., Bresson, X., and Vandergheynst, P. (2016). Convolutional neural networks on graphs with fast localized spectral filtering. In

- Lee, D. D., Sugiyama, M., von Luxburg, U., Guyon, I., and Garnett, R., editors, *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 3837–3845.
- [Dinh et al., 2014] Dinh, L., Krueger, D., and Bengio, Y. (2014). Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*.
- [Dinh et al., 2016] Dinh, L., Sohl-Dickstein, J., and Bengio, S. (2016). Density estimation using real NVP. *CoRR*, abs/1605.08803.
- [Faury et al., 2019] Faury, L., Tanielian, U., Vasile, F., Smirnova, E., and Dohmatob, E. (2019). Distributionally robust counterfactual risk minimization. *arXiv preprint arXiv:1906.06211*.
- [Gao and Ji, 2019] Gao, H. and Ji, S. (2019). Graph u-nets. *CoRR*, abs/1905.05178.
- [Gilmer et al., 2017] Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., and Dahl, G. E. (2017). Neural message passing for quantum chemistry. *CoRR*, abs/1704.01212.
- [Goyal et al., 2019] Goyal, A., Lamb, A., Hoffmann, J., Sodhani, S., Levine, S., Bengio, Y., and Schölkopf, B. (2019). Recurrent independent mechanisms. *ArXiv*, abs/1909.10893.
- [Guidotti et al., 2018] Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., and Pedreschi, D. (2018). A survey of methods for explaining black box models. *ACM computing surveys (CSUR)*, 51(5):1–42.
- [Hamilton et al., 2017] Hamilton, W., Ying, Z., and Leskovec, J. (2017). Inductive representation learning on large graphs. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R., editors, *Advances in Neural Information Processing Systems 30*, pages 1024–1034. Curran Associates, Inc.
- [He et al., 2016] He, K., Zhang, X., Ren, S., and Sun, J. (2016). Identity mappings in deep residual networks. In Leibe, B., Matas, J., Sebe, N., and Welling, M., editors, *Computer Vision – ECCV 2016*, pages 630–645, Cham. Springer International Publishing.
- [Hendricks et al., 2016] Hendricks, L. A., Akata, Z., Rohrbach, M., Donahue, J., Schiele, B., and Darrell, T. (2016). Generating visual explanations. In *European Conference on Computer Vision*, pages 3–19. Springer.



- [Huang and Roumeliotis, 2013] Huang, G. P. and Roumeliotis, S. I. (2013). Analytically-guided-sampling particle filter applied to range-only target tracking. In *2013 IEEE International Conference on Robotics and Automation*, pages 3168–3175.
- [Huang et al., 2019] Huang, Z., Liang, S., Liang, M., and Yang, H. (2019). Dianet: Dense-and-implicit attention network. *CoRR*, abs/1905.10671.
- [Jin et al., ] Jin, L., Song, L., Zhang, Y., Xu, K., Ma, W.-y., and Yu, D. Relation extraction exploiting full dependency forests.
- [Jin et al., 2019] Jin, Y., Song, G., and Shi, C. (2019). Gralsp: Graph neural networks with local structural patterns. *arXiv preprint arXiv:1911.07675*.
- [Ke et al., 2019] Ke, N. R., Bilaniuk, O., Goyal, A., Bauer, S., Larochelle, H., Pal, C., and Bengio, Y. (2019). Learning neural causal models from unknown interventions. *ArXiv*, abs/1910.01075.
- [Kendall and Gal, 2017] Kendall, A. and Gal, Y. (2017). What uncertainties do we need in bayesian deep learning for computer vision? In *Advances in neural information processing systems*, pages 5574–5584.
- [Kim et al., 2016] Kim, B., Khanna, R., and Koyejo, O. O. (2016). Examples are not enough, learn to criticize! criticism for interpretability. In *Advances in neural information processing systems*, pages 2280–2288.
- [Kim et al., 2019] Kim, S., Park, S., Na, B., and Yoon, S. (2019). Spiking-yolo: Spiking neural network for real-time object detection. *arXiv preprint arXiv:1903.06530*.
- [Kingma and Dhariwal, 2018] Kingma, D. P. and Dhariwal, P. (2018). Glow: Generative flow with invertible 1x1 convolutions. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R., editors, *Advances in Neural Information Processing Systems 31*, pages 10215–10224. Curran Associates, Inc.
- [Kipf and Welling, 2016] Kipf, T. N. and Welling, M. (2016). Semi-supervised classification with graph convolutional networks. *CoRR*, abs/1609.02907.
- [Koh and Liang, 2017] Koh, P. W. and Liang, P. (2017). Understanding black-box predictions via influence functions. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1885–1894. JMLR. org.

- [Kononenko et al., 2010] Kononenko, I. et al. (2010). An efficient explanation of individual classifications using game theory. *Journal of Machine Learning Research*, 11(Jan):1–18.
- [Kosioerek et al., 2019] Kosioerek, A., Sabour, S., Teh, Y. W., and Hinton, G. E. (2019). Stacked capsule autoencoders. In Wallach, H., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E., and Garnett, R., editors, *Advances in Neural Information Processing Systems 32*, pages 15512–15522. Curran Associates, Inc.
- [Lake and Baroni, 2017] Lake, B. M. and Baroni, M. (2017). Still not systematic after all these years: On the compositional skills of sequence-to-sequence recurrent networks. *CoRR*, abs/1711.00350.
- [Lapuschkin et al., 2016] Lapuschkin, S., Binder, A., Montavon, G., Müller, K.-R., and Samek, W. (2016). The lrp toolbox for artificial neural networks. *Journal of Machine Learning Research*, 17(114):1–5.
- [Lauscher et al., 2019] Lauscher, A., Glavas, G., Ponzetto, S. P., and Vulic, I. (2019). A general framework for implicit and explicit debiasing of distributional word vector spaces. *ArXiv*, abs/1909.06092.
- [Lecue et al., 2019] Lecue, F., Chen, J., Pan, J. Z., and Chen, H. (2019). Augmenting transfer learning with semantic reasoning. *arXiv preprint arXiv:1905.13672*.
- [Lee et al., 2018] Lee, J., Lee, Y., Kim, J., Kosioerek, A. R., Choi, S., and Teh, Y. W. (2018). Set transformer. *CoRR*, abs/1810.00825.
- [Li et al., 2014] Li, M., Yin, H., Tahara, K., and Billard, A. (2014). Learning object-level impedance control for robust grasping and dexterous manipulation. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6784–6791.
- [Li et al., 2018] Li, O., Liu, H., Chen, C., and Rudin, C. (2018). Deep learning for case-based reasoning through prototypes: A neural network that explains its predictions. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- [Li et al., 2019] Li, Y., Gu, C., Dullien, T., Vinyals, O., and Kohli, P. (2019). Graph matching networks for learning the similarity of graph structured objects. *arXiv preprint arXiv:1904.12787*.
- [Li et al., 2015] Li, Y., Tarlow, D., Brockschmidt, M., and Zemel, R. (2015). Gated graph sequence neural networks. *arXiv preprint arXiv:1511.05493*.

- [Lundberg and Lee, 2017] Lundberg, S. M. and Lee, S.-I. (2017). A unified approach to interpreting model predictions. In *Advances in neural information processing systems*, pages 4765–4774.
- [Luo et al., 2020] Luo, A., Yang, F., Li, X., Nie, D., Jiao, Z., Zhou, S., and Cheng, H. (2020). Hybrid graph neural networks for crowd counting. *ArXiv*, abs/2002.00092.
- [Ma et al., 2019a] Ma, T., Shang, J., Xiao, C., and Sun, J. (2019a). Genn: Predicting correlated drug-drug interactions with graph energy neural networks. *arXiv preprint arXiv:1910.02107*.
- [Ma et al., 2019b] Ma, Y., Aggarwal, C., Wang, S., and Tang, J. (2019b). Graph convolutional networks with eigenpooling. In *KDD 2019 - Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pages 723–731. Association for Computing Machinery.
- [Madhawa et al., 2019] Madhawa, K., Ishiguro, K., Nakago, K., and Abe, M. (2019). Graphnvp: An invertible flow model for generating molecular graphs. *arXiv preprint arXiv:1905.11600*.
- [Magelinski et al., 2019] Magelinski, T., Beskow, D., and Carley, K. M. (2019). Graph-hist: Graph classification from latent feature histograms with application to bot detection. *arXiv preprint arXiv:1910.01180*.
- [Mathieu et al., 2015] Mathieu, M., Couprie, C., and LeCun, Y. (2015). Deep multi-scale video prediction beyond mean square error. *arXiv preprint arXiv:1511.05440*.
- [Merrick and Taly, 2019] Merrick, L. and Taly, A. (2019). The explanation game: Explaining machine learning models with cooperative game theory. *arXiv preprint arXiv:1909.08128*.
- [Mirrazavi Salehian et al., 2016] Mirrazavi Salehian, S. S., Figueroa Fernandez, N. B., and Billard, A. (2016). Coordinated multi-arm motion planning: Reaching for moving objects in the face of uncertainty (rss 2016 best student paper award).
- [Mohammadshahi and Henderson, 2019] Mohammadshahi, A. and Henderson, J. (2019). Graph-to-graph transformer for transition-based dependency parsing. *arXiv preprint arXiv:1911.03561*.

- [Montufar et al., 2014] Montufar, G. F., Pascanu, R., Cho, K., and Bengio, Y. (2014). On the number of linear regions of deep neural networks. In Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N. D., and Weinberger, K. Q., editors, *Advances in Neural Information Processing Systems 27*, pages 2924–2932. Curran Associates, Inc.
- [Nijkamp et al., 2019] Nijkamp, E., Hill, M., Han, T., Zhu, S.-C., and Wu, Y. N. (2019). On the anatomy of mcmc-based maximum likelihood learning of energy-based models. *arXiv preprint arXiv:1903.12370*.
- [Pareja et al., 2019] Pareja, A., Domeniconi, G., Chen, J., Ma, T., Suzumura, T., Kanezashi, H., Kaler, T., and Leiserson, C. E. (2019). Evolvegcnn: Evolving graph convolutional networks for dynamic graphs. *CoRR*, abs/1902.10191.
- [Pascanu et al., 2013] Pascanu, R., Gulcehre, C., Cho, K., and Bengio, Y. (2013). How to construct deep recurrent neural networks. *arXiv preprint arXiv:1312.6026*.
- [Pathak et al., 2016] Pathak, D., Krahenbuhl, P., Donahue, J., Darrell, T., and Efros, A. A. (2016). Context encoders: Feature learning by inpainting. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [Pearlmutter, 1995] Pearlmutter, B. A. (1995). Gradient calculations for dynamic recurrent neural networks: a survey. *IEEE Transactions on Neural Networks*, 6(5):1212–1228.
- [Poli et al., 2019] Poli, M., Massaroli, S., Park, J., Yamashita, A., Asama, H., and Park, J. (2019). Graph neural ordinary differential equations. *arXiv preprint arXiv:1911.07532*.
- [Qi et al., 2019a] Qi, F., Chang, L., Sun, M., Ouyang, S., and Liu, Z. (2019a). Towards building a multilingual sememe knowledge base: Predicting sememes for babelnet synsets. *ArXiv*, abs/1912.01795.
- [Qi et al., 2019b] Qi, Z., Khorram, S., and Li, F. (2019b). Visualizing deep networks by optimizing with integrated gradients. *arXiv preprint arXiv:1905.00954*.
- [Rezende and Mohamed, 2015] Rezende, D. J. and Mohamed, S. (2015). Variational inference with normalizing flows. *arXiv preprint arXiv:1505.05770*.
- [Ribeiro et al., 2016] Ribeiro, M. T., Singh, S., and Guestrin, C. (2016). Model-agnostic interpretability of machine learning. *arXiv preprint arXiv:1606.05386*.

- [Ribeiro et al., 2018] Ribeiro, M. T., Singh, S., and Guestrin, C. (2018). Anchors: High-precision model-agnostic explanations. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- [Rosenthal et al., 2016] Rosenthal, S., Selvaraj, S. P., and Veloso, M. M. (2016). Verbalization: Narration of autonomous robot experience. In *IJCAI*, pages 862–868.
- [Scarselli et al., 2008] Scarselli, F., Gori, M., Tsoi, A. C., Hagenbuchner, M., and Monfardini, G. (2008). The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80.
- [Schölkopf et al., 2012] Schölkopf, B., Janzing, D., Peters, J., Sgouritsa, E., Zhang, K., and Mooij, J. (2012). On causal and anticausal learning. *arXiv preprint arXiv:1206.6471*.
- [Shang et al., 2019] Shang, J., Ma, T., Xiao, C., and Sun, J. (2019). Pre-training of graph augmented transformers for medication recommendation. *arXiv preprint arXiv:1906.00346*.
- [Shapley, 1953] Shapley, L. S. (1953). A value for n-person games. *Contributions to the Theory of Games*, 2(28):307–317.
- [Sheh, 2017] Sheh, R. K.-M. (2017). ” why did you do that?” explainable intelligent robots. In *Workshops at the Thirty-First AAAI Conference on Artificial Intelligence*.
- [Shrikumar et al., 2017] Shrikumar, A., Greenside, P., and Kundaje, A. (2017). Learning important features through propagating activation differences. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 3145–3153. JMLR. org.
- [Springenberg et al., 2014] Springenberg, J. T., Dosovitskiy, A., Brox, T., and Riedmiller, M. (2014). Striving for simplicity: The all convolutional net. *arXiv preprint arXiv:1412.6806*.
- [Sundararajan et al., 2017a] Sundararajan, M., Taly, A., and Yan, Q. (2017a). Ax- iomatic attribution for deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 3319–3328. JMLR. org.
- [Sundararajan et al., 2017b] Sundararajan, M., Taly, A., and Yan, Q. (2017b). Ax- iomatic attribution for deep networks. *CoRR*, abs/1703.01365.

- [Tan et al., 2019] Tan, Z.-X., Soh, H., and Ong, D. C. (2019). Factorized inference in deep markov models for incomplete multimodal time series. *ArXiv*, abs/1905.13570.
- [Tang et al., 2019] Tang, X., Li, Y., Sun, Y., Yao, H., Mitra, P., and Wang, S. (2019). Robust graph neural network against poisoning attacks via transfer learning. *arXiv preprint arXiv:1908.07558*.
- [Tomsett et al., 2019] Tomsett, R., Harborne, D., Chakraborty, S., Gurram, P., and Preece, A. (2019). Sanity checks for saliency metrics. *arXiv preprint arXiv:1912.01451*.
- [Veličković et al., 2017] Veličković, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., and Bengio, Y. (2017). Graph attention networks. *arXiv preprint arXiv:1710.10903*.
- [Wang et al., 2019] Wang, X., Wang, R., Shi, C., Song, G., and Li, Q. (2019). Multi-component graph convolutional collaborative filtering. *arXiv preprint arXiv:1911.10699*.
- [Wei et al., 2019] Wei, J., Wang, S., and Huang, Q. (2019). F3net: Fusion, feedback and focus for salient object detection. *ArXiv*, abs/1911.11445.
- [Xu et al., 2018] Xu, K., Wu, L., Wang, Z., Feng, Y., Witbrock, M., and Sheini'n, V. (2018). Graph2seq: Graph to sequence learning with attention-based neural networks. *arXiv preprint arXiv:1804.00823*.
- [Ying et al., 2018] Ying, Z., You, J., Morris, C., Ren, X., Hamilton, W., and Leskovec, J. (2018). Hierarchical graph representation learning with differentiable pooling. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R., editors, *Advances in Neural Information Processing Systems 31*, pages 4800–4810. Curran Associates, Inc.
- [Zang et al., 2019] Zang, C., Cui, P., Zhu, W., and Wang, F. (2019). Dynamical origins of distribution functions. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD ' 19*, page 469–478, New York, NY, USA. Association for Computing Machinery.
- [Zang and Wang, 2019] Zang, C. and Wang, F. (2019). Neural dynamics on complex networks. *arXiv preprint arXiv:1908.06491*.
- [Zhang et al., 2019] Zhang, K., Chen, J.-R., Liu, B., and Liu, Q. (2019). Deep object co-segmentation via spatial-semantic network modulation. *ArXiv*, abs/1911.12950.