

# Neural Structured Learning Tutorial @ KDD'20

August 25, 2020



TensorFlow



# Organizers



Arjun Gopalan



Da-Cheng Juan



Cesar Ilharco Magalhaes



Chun-Sung Ferng



Allan Heydon



Chun-Ta Lu



Philip Pham



George Yu

All organizers are from Google Research



# Agenda

- Introduction to NSL
- Data Preprocessing
- Regularizing with Natural Graphs (Lab 1)
- Regularizing with Synthesized Graphs (Lab 2)
- Adversarial Learning (Lab 3)
- NSL in TensorFlow Extended (TFX)
- Research and Future Directions
- Conclusion

Tutorial website:

[https://github.com/tensorflow/neural-structured-learning/tree/master/workshops/kdd\\_2020](https://github.com/tensorflow/neural-structured-learning/tree/master/workshops/kdd_2020)



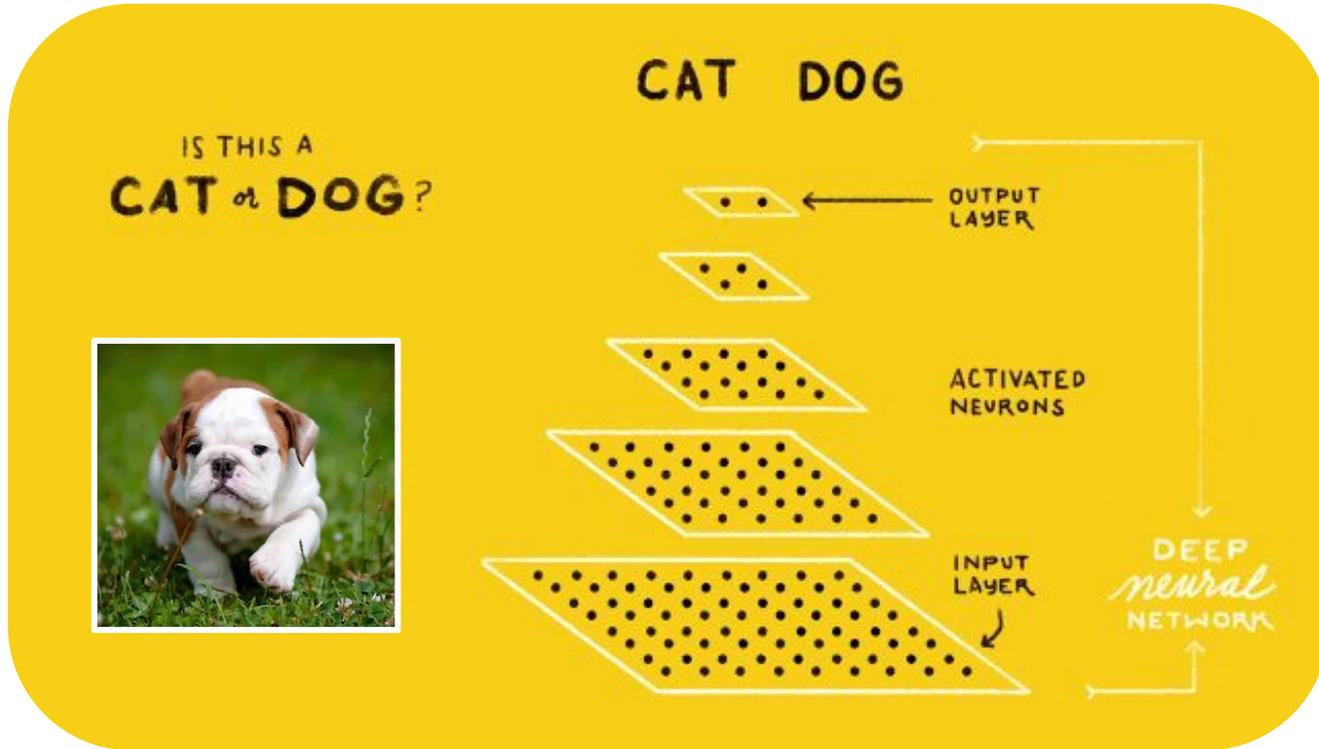
# Introduction to NSL

Da-Cheng Juan



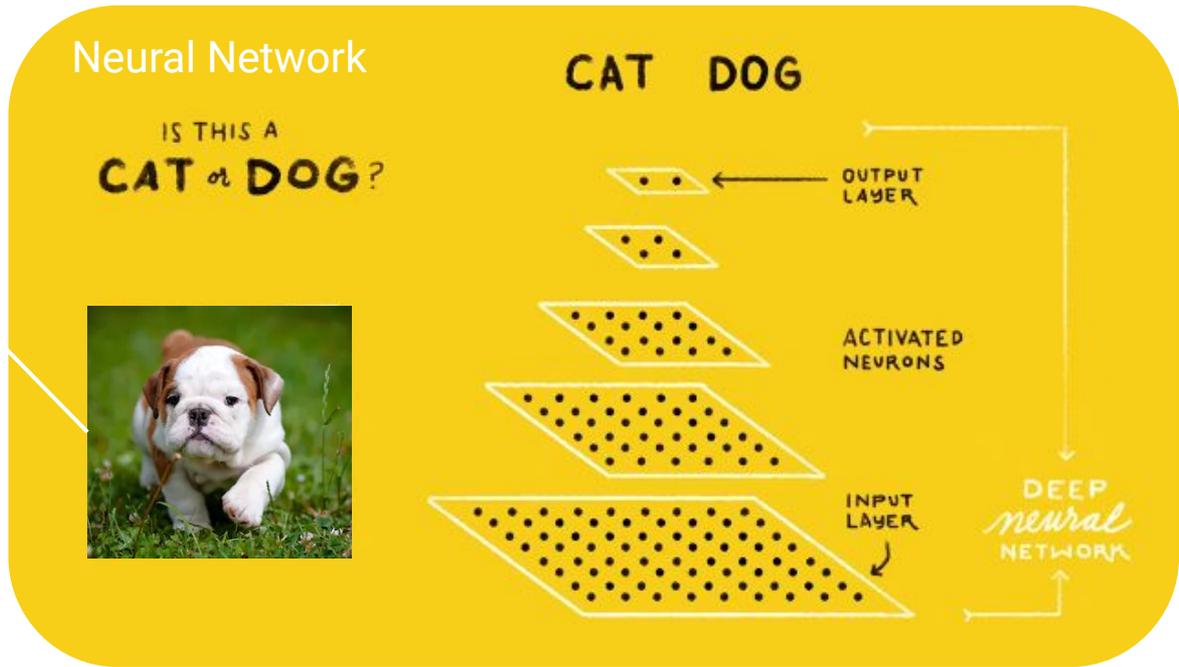
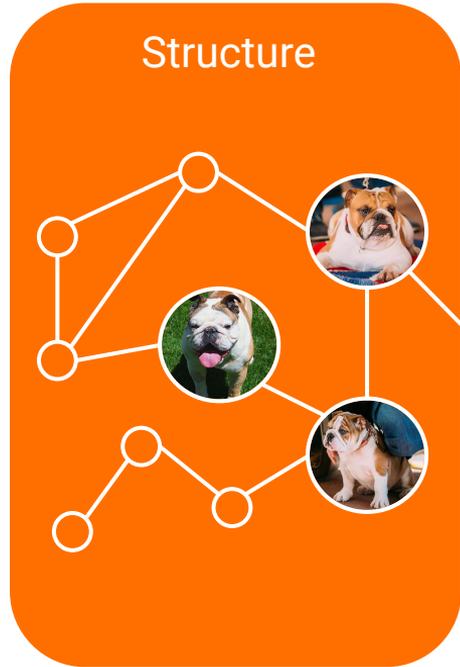


# How a Typical Neural Net Works





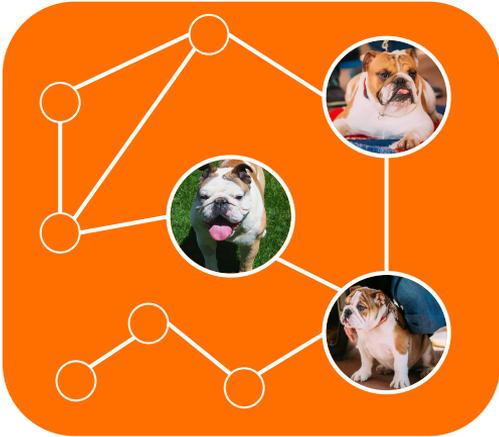
# Neural Structured Learning (NSL)





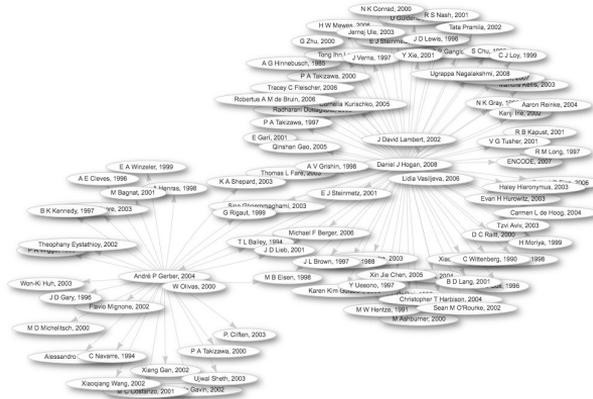
# Structure Among Samples

## Co-Occurrence Graph



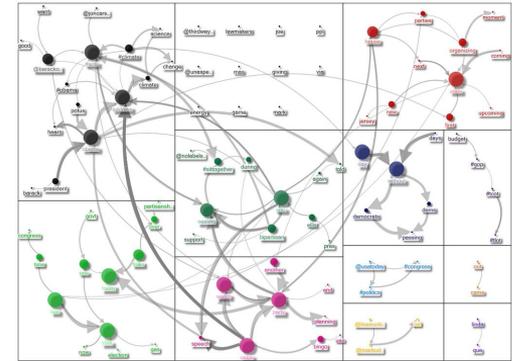
[Source: graph concept is from Juan et al., WSDM'20. Original images are from pixabay.com]

## Citation Graph



[Source: [https://commons.wikimedia.org/wiki/File:Partial\\_citation\\_graph\\_for\\_%2A\\_screen\\_for\\_RNA-binding\\_proteins\\_in\\_yeast\\_indicates\\_dual\\_functions\\_for\\_many\\_enzymes%22\\_as\\_of\\_April\\_12\\_2017.png](https://commons.wikimedia.org/wiki/File:Partial_citation_graph_for_%2A_screen_for_RNA-binding_proteins_in_yeast_indicates_dual_functions_for_many_enzymes%22_as_of_April_12_2017.png)]

## Text Graph



[Source: copied without modification from [https://www.flickr.com/photos/marc\\_smith/6705382867/sizes/l/](https://www.flickr.com/photos/marc_smith/6705382867/sizes/l/)]



# NSL: Advantages of Learning with Structure

➔ **Less Labeled Data Required**

➔ **Robust Model**





# Scenario II: Model Robustness Required

Example task: Image Classification



$x$

“panda”

+ .007 ×



$\text{sign}(\nabla_x J(\theta, x, y))$

“nematode”

=



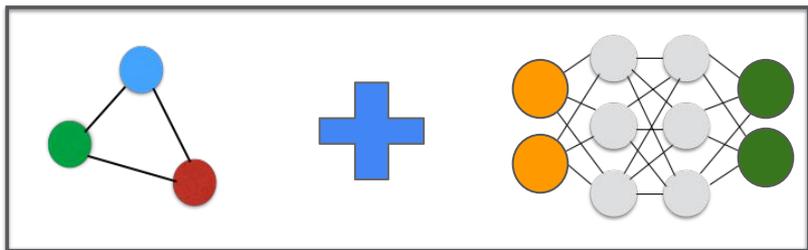
$x + \epsilon \text{sign}(\nabla_x J(\theta, x, y))$   
“gibbon”

[Source: Goodfellow, et al., ICLR'15]



# NSL: Neural Graph Learning

## Graph + Neural Net

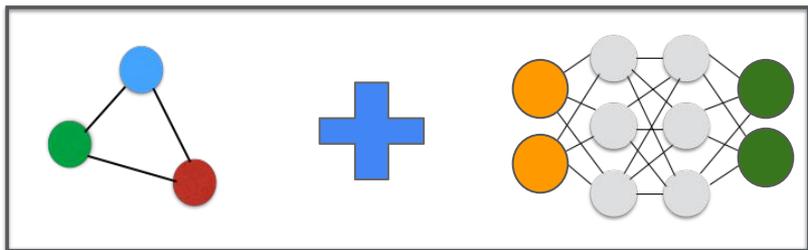


- **Jointly** optimizes both features & structured signals for better models

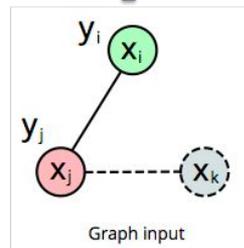
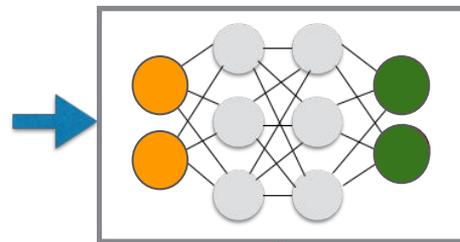
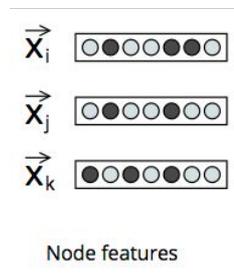


# NSL: Neural Graph Learning

## Graph + Neural Net



- **Jointly** optimizes both features & structured signals for better models



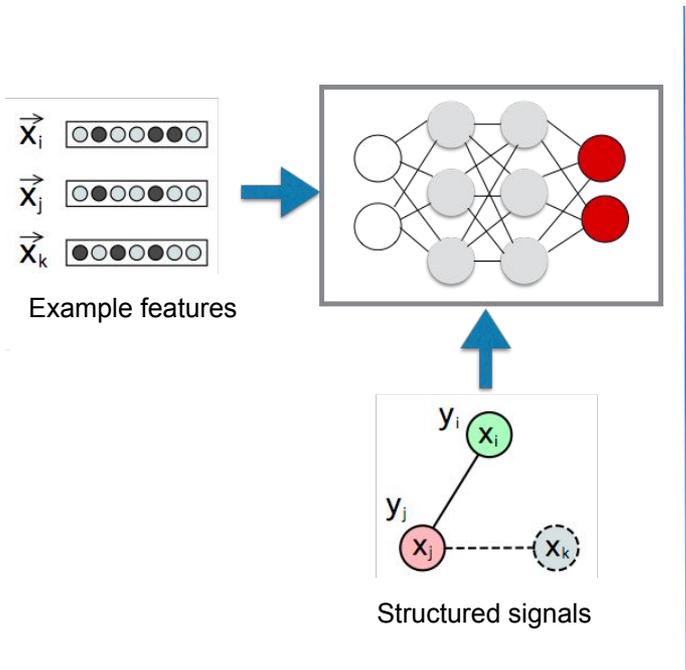
## Neural Graph Machines (NGM)

[Paper](#): Bui, Ravi & Ramavajjala [WSDM'18]



# NSL: Neural Graph Learning

Joint optimization with label and structured signals:



$$\text{Optimize: } \text{loss} = \sum_{i=1}^B \mathcal{L}(y_i, \hat{y}_i) + \alpha \sum_{i=1}^B \mathcal{L}_{\mathcal{N}}(y_i, x_i, \mathcal{N}(x_i))$$

$$x_i \rightarrow f(\cdot) \rightarrow y_i$$



Supervised Loss

$$\sum_{i=1}^B \mathcal{E}(y_i, g_{\theta}(x_i))$$

$g_{\theta}(x_i)$ : NN output for input  $x_i$   
 $\mathcal{E}(\cdot)$ : Loss function

Examples: L2 (for regression)  
 Cross-Entropy (for classification)

Neighbor Loss

$$\sum_{x_j \in \mathcal{N}(x_i)} w_{ij} \cdot \mathcal{D}(h_{\theta}(x_i), h_{\theta}(x_j))$$

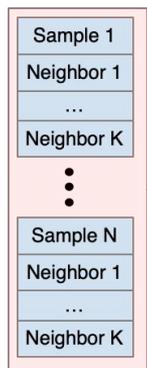
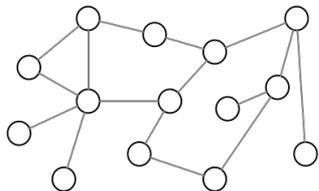
$h_{\theta}(\cdot)$ : Target hidden layer  
 $\mathcal{D}(\cdot)$ : Distance metric

Examples: L1, L2, ...



# NSL: Neural Graph Learning in Practice

Training samples with labels



Batch of labeled samples with neighbors

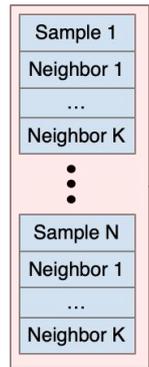
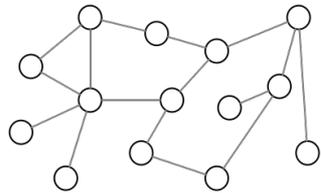
Structured signals (e.g., graphs)

[Source: Juan, et al., WSDM'20]



# NSL: Neural Graph Learning in Practice

Training samples with labels



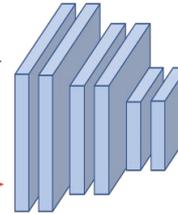
Batch of labeled samples with neighbors

Input Layer

Sample Features

Neighbor Features

Neural Net

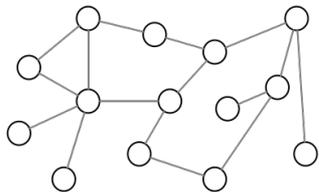


Structured signals (e.g., graphs)

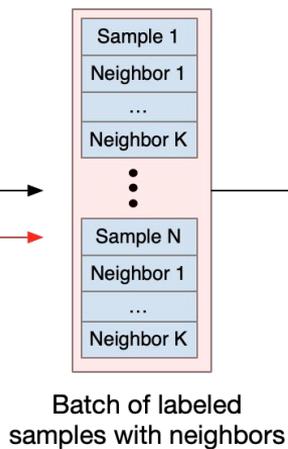


# NSL: Neural Graph Learning in Practice

Training samples with labels



Structured signals (e.g., graphs)

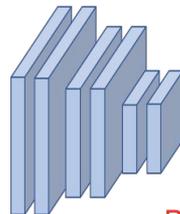


Input Layer

Sample Features

Neighbor Features

Neural Net



Sample Embedding

$\phi(x_u)$



Sample Labels

Supervised Loss

Regularization

$\phi(x_v)$

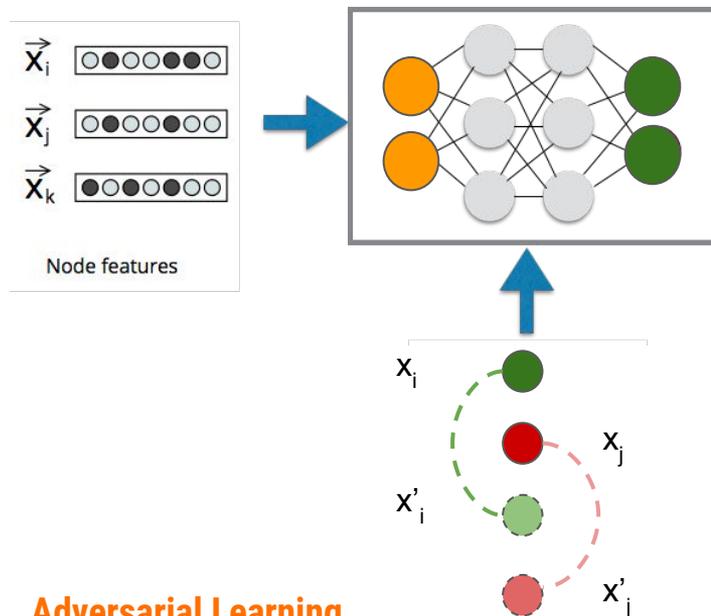
Neighbor Embedding



# NSL: Adversarial Learning

## Adversarial + Neural Net

- **Jointly** optimize features from original and “adversarial” examples for more robust models



**Adversarial Learning**

[Paper](#): Goodfellow, et al. [ICLR'15]



# Libraries, Tools and Trainers

## Standalone Tool

`build_graph`

`pack_nbrs`

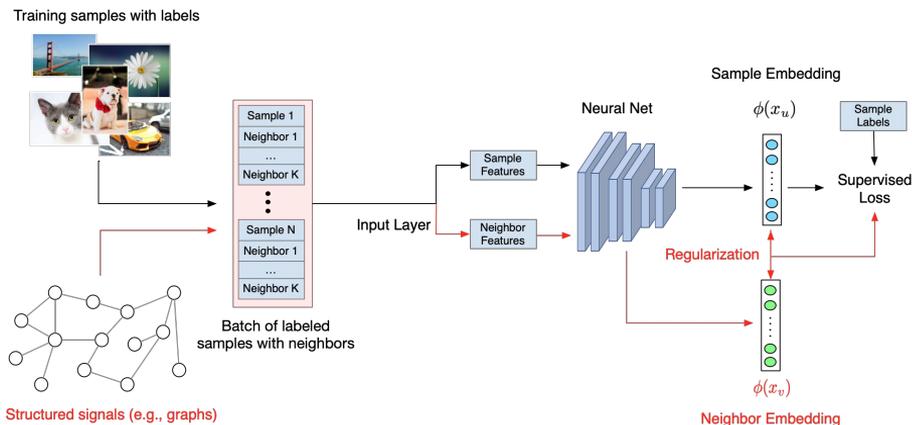
## Graph Functions

`read_tsv_graph`

`write_tsv_graph`

`add_edge`

`add_undirected_edges`



## Lib

`unpack_neighbor_features`

`gen_adv_neighbor`

`replicate_embeddings`

## Keras

`GraphRegularization`

`AdversarialRegularization`

`layers`

Web: [tensorflow.org/neural\\_structured\\_learning](https://tensorflow.org/neural_structured_learning)