Multimodal Learning with Graphs

Yasha Ektefaie, George Dasoulas, Ayush Noori, Maha Farhat, and Marinka Zitnik

Vamsi

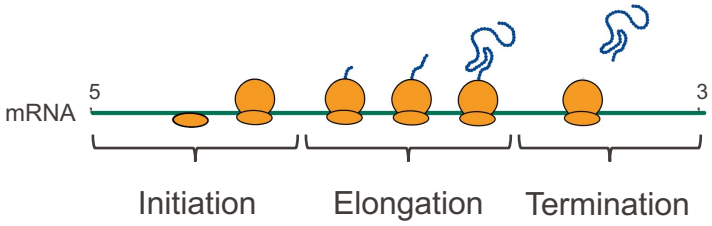

- **Background:** Bachelor's in Computer Science
- 3rd Year EDCB Ph.D. student in the LTS2 laboratory for signal processing (Prof. Pierre Vandergheynst)
- Working on building explainable machine learning models to study translation elongation

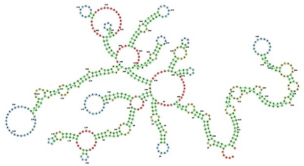

mRNA Secondary Structure

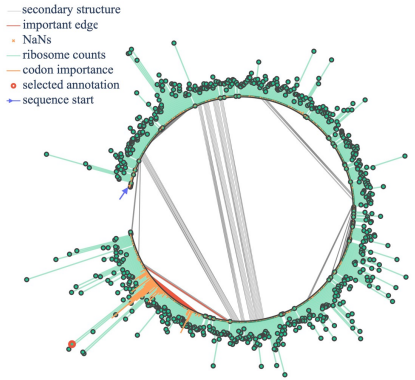

Vamsi Nallapareddy

What is multimodal learning?

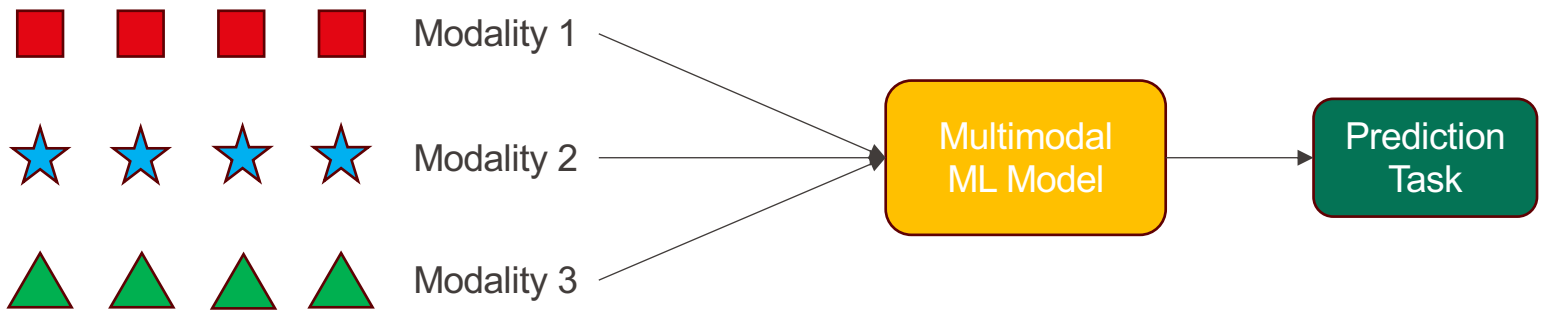

Issues with multimodal learning

- **Modality Collapse:** The model tends to only look at a subset of modalities that are most useful in training whilst the rest could be potentially more useful in implementation.
- **Missing Modalities:** Some samples don't have information regarding all the modalities that the model uses.
- **Complex Inter-Modality Relationships:** The different modalities might have a different relationship with each other, so a simple fusion model might not be enough to extract all the information.

Issues with multimodal learning

- **Modality Collapse:** The model tends to only look at a subset of modalities that are most useful in training whilst the rest could be potentially more useful in implementation.
- **Missing Modalities:** Some samples don't have information regarding all the modalities that the model uses.
- **Complex Inter-Modality Relationships:** The different modalities might have a different relationship with each other, so a simple fusion model might not be enough to extract all the information.

Issues with multimodal learning

- **Modality Collapse:** The model tends to only look at a subset of modalities that are most useful in training whilst the rest could be potentially more useful in implementation.
- **Missing Modalities:** Some samples don't have information regarding all the modalities that the model uses.
- **Complex Inter-Modality Relationships:** The different modalities might have a different relationship with each other, so a simple fusion model might not be enough to extract all the information.

Multimodal Graph Learning (MGL)

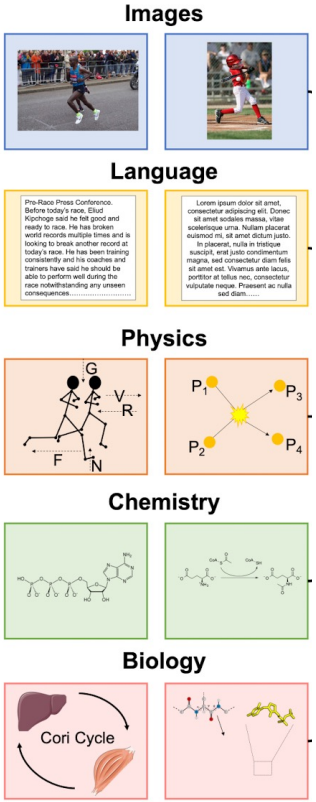

Vamsi Nallapareddy

Multimodal Graph Learning (MGL)

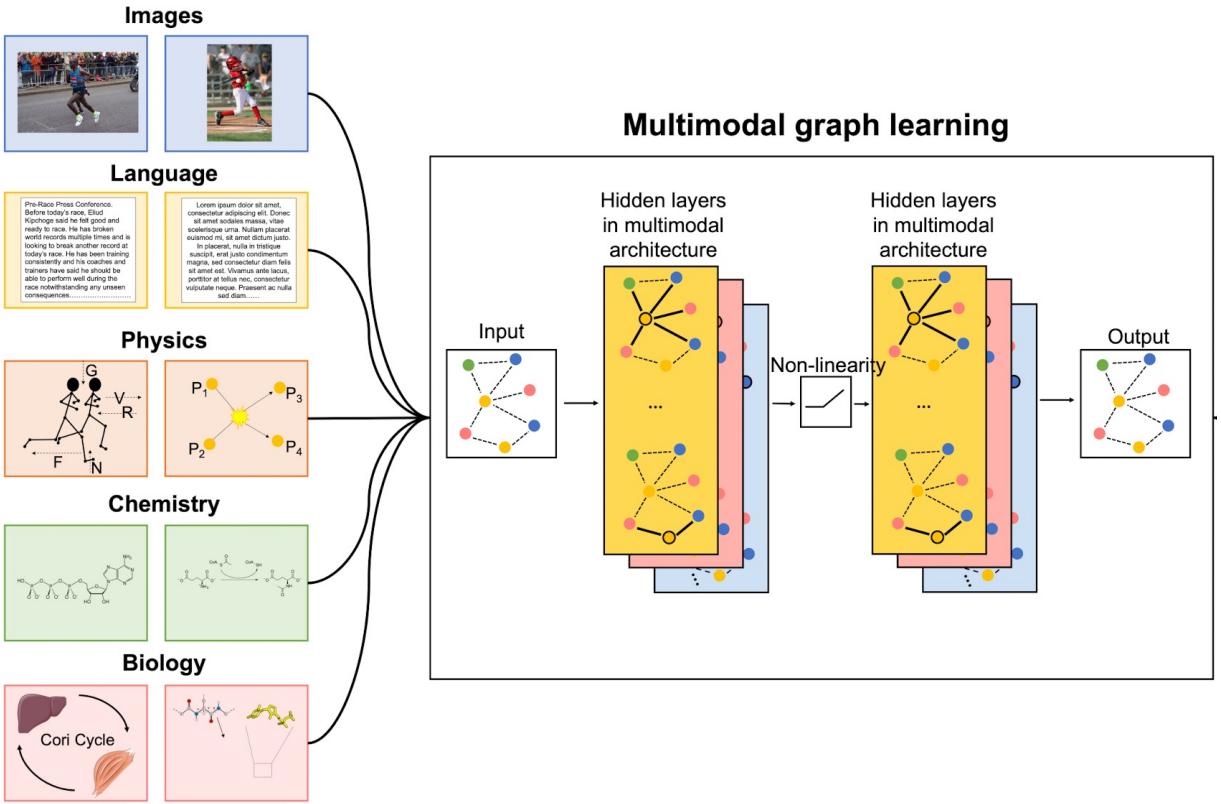

Multimodal Graph Learning (MGL)

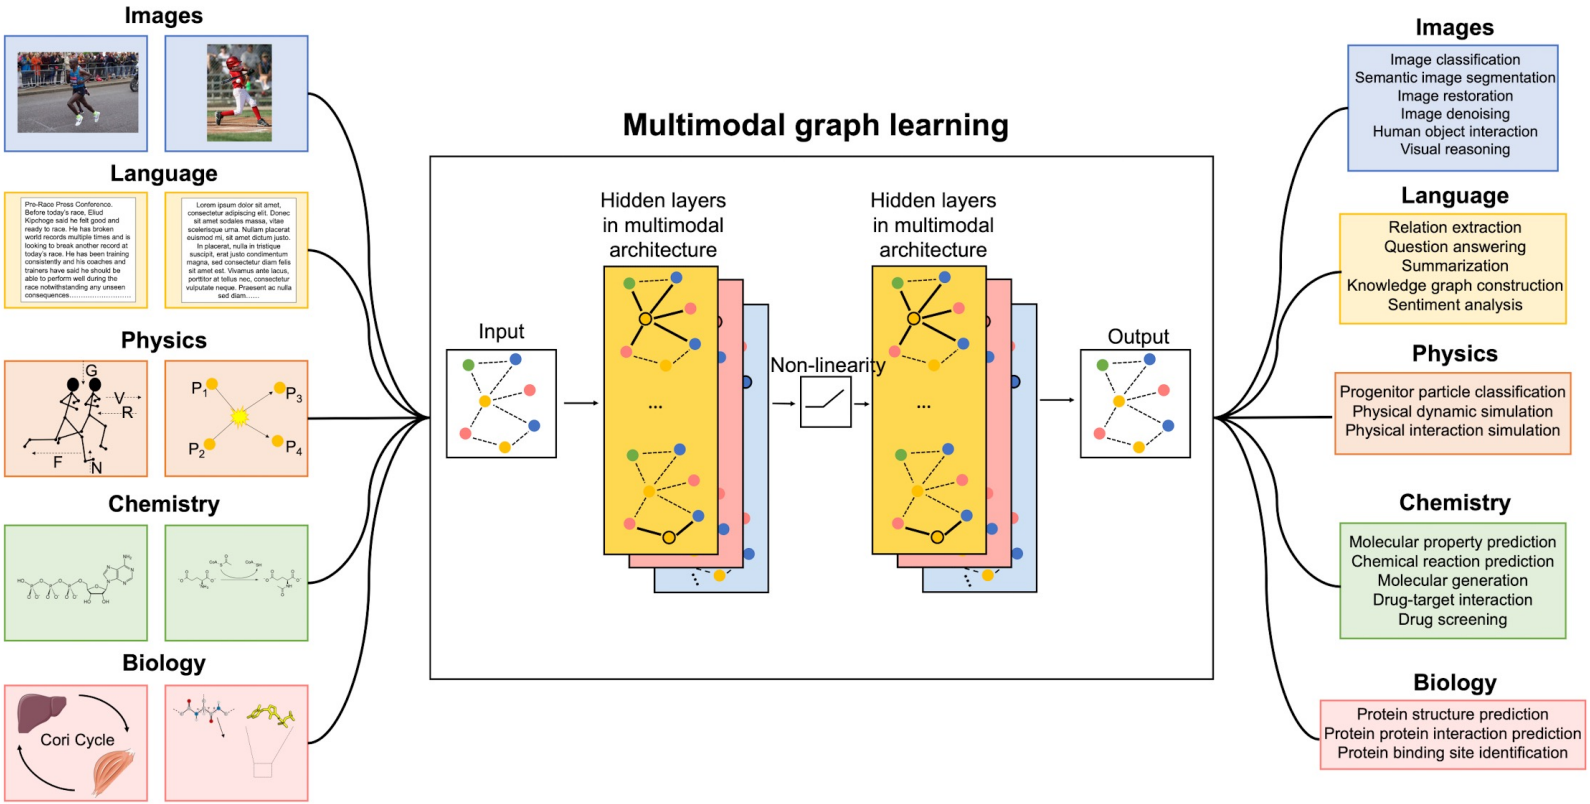

MGL Blueprint

- Simple single-modal
- All-in-one model
- MGL Blueprint

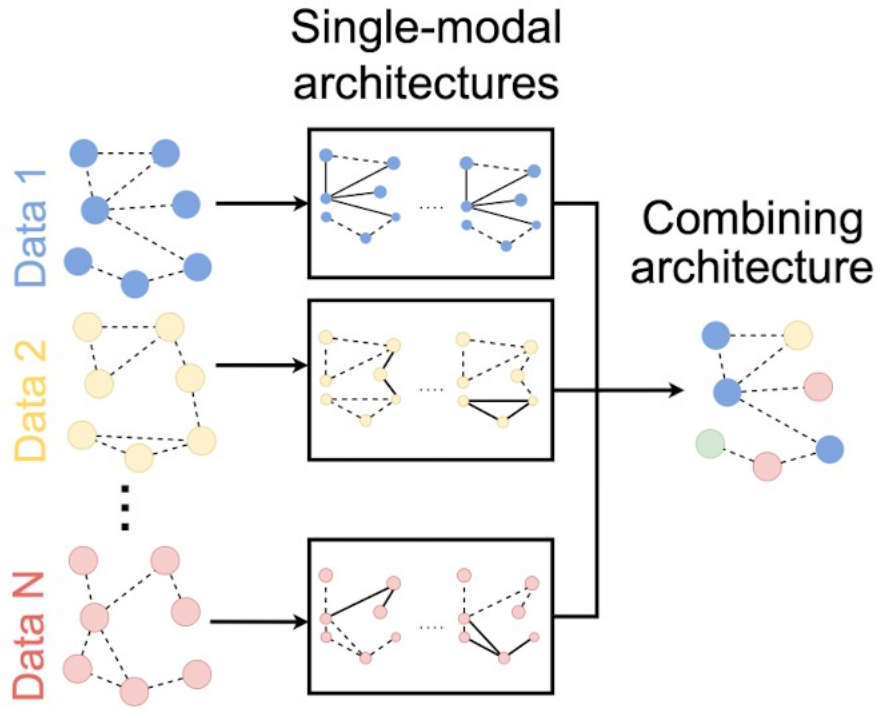

MGL Blueprint

- Simple single-modal
- All-in-one model
- MGL Blueprint

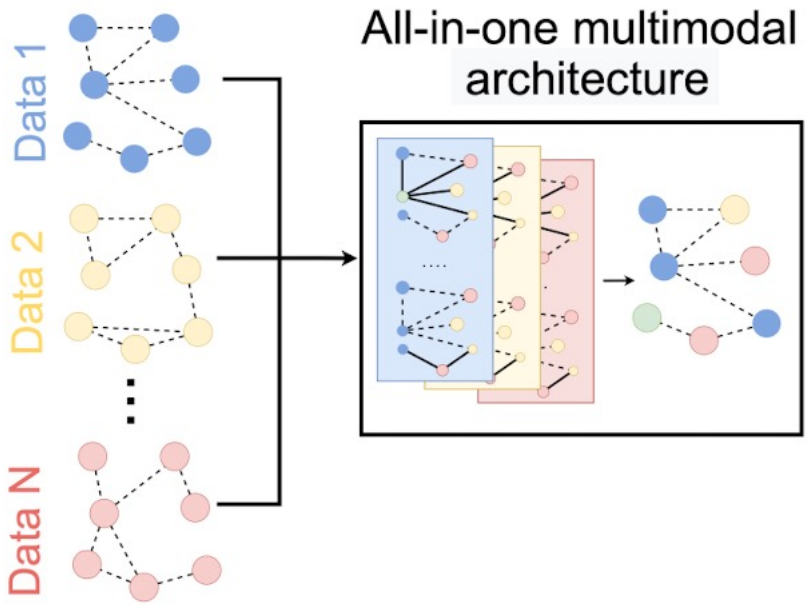

MGL Blueprint

- Simple single-modal
- All-in-one model
- MGL Blueprint
  - Structure Learning
  - Learning on Structure

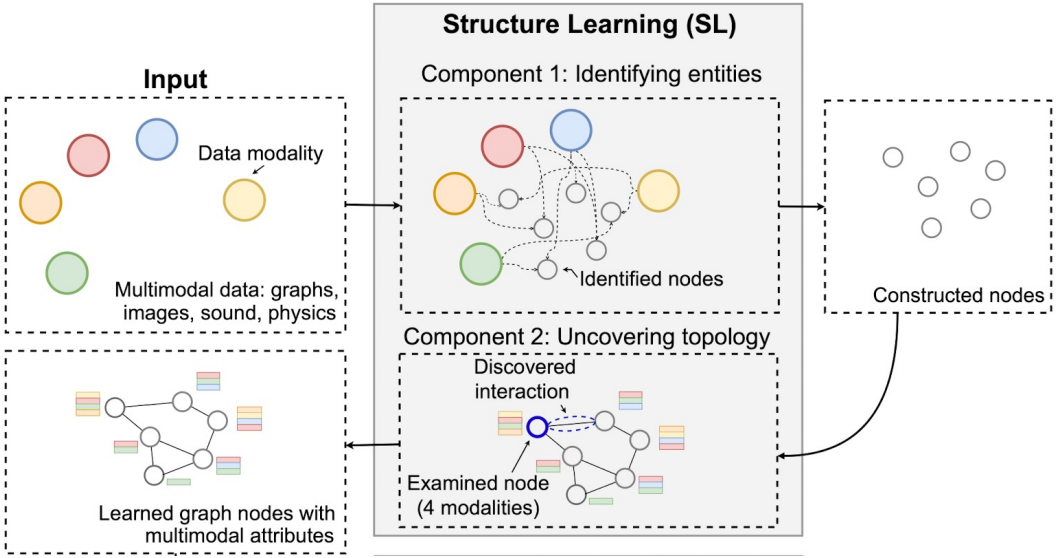

MGL Blueprint

- Simple single-modal
- All-in-one model
- MGL Blueprint
  - Structure Learning
  - Learning on Structure

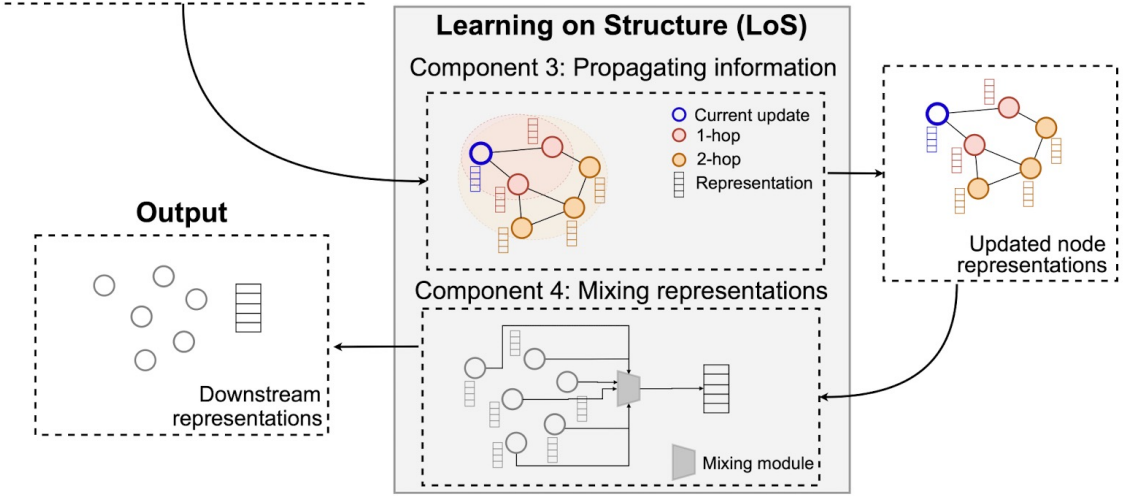

MGL Blueprint

- Simple single-modal
- All-in-one model
- MGL Blueprint
  - Structure Learning
  - Learning on Structure

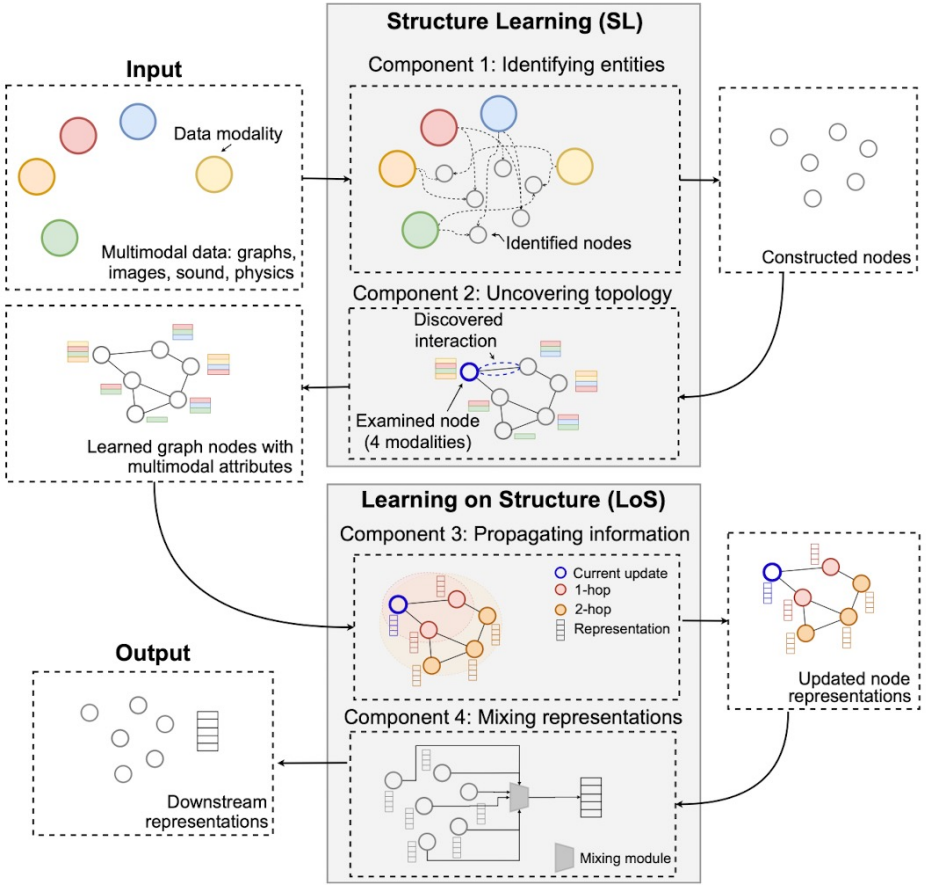

MGL on Images Use-Cases

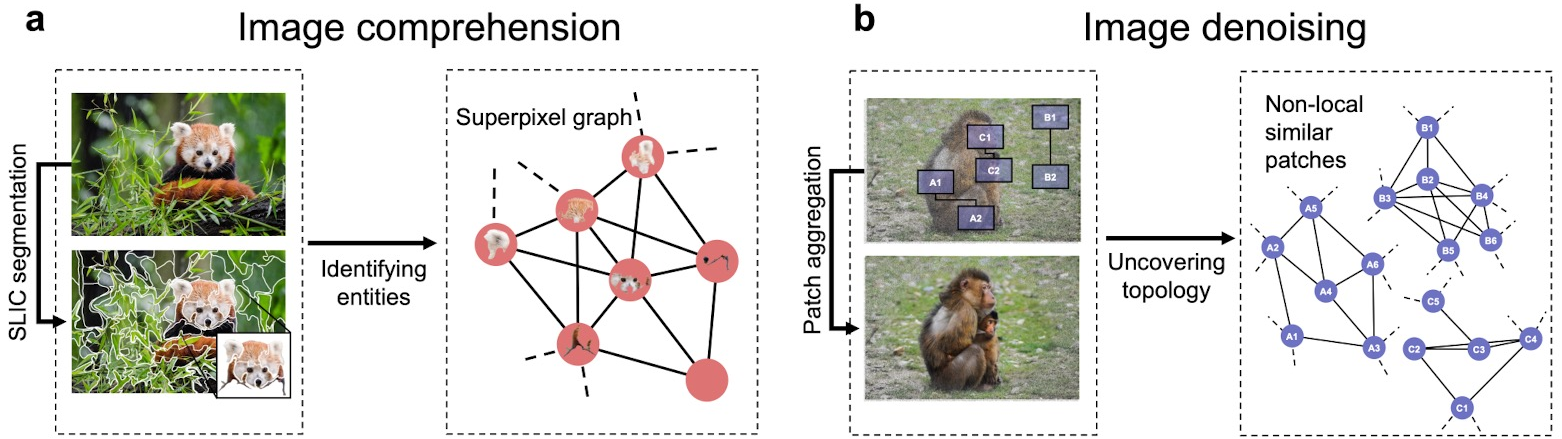

MGL on Images Use-Cases

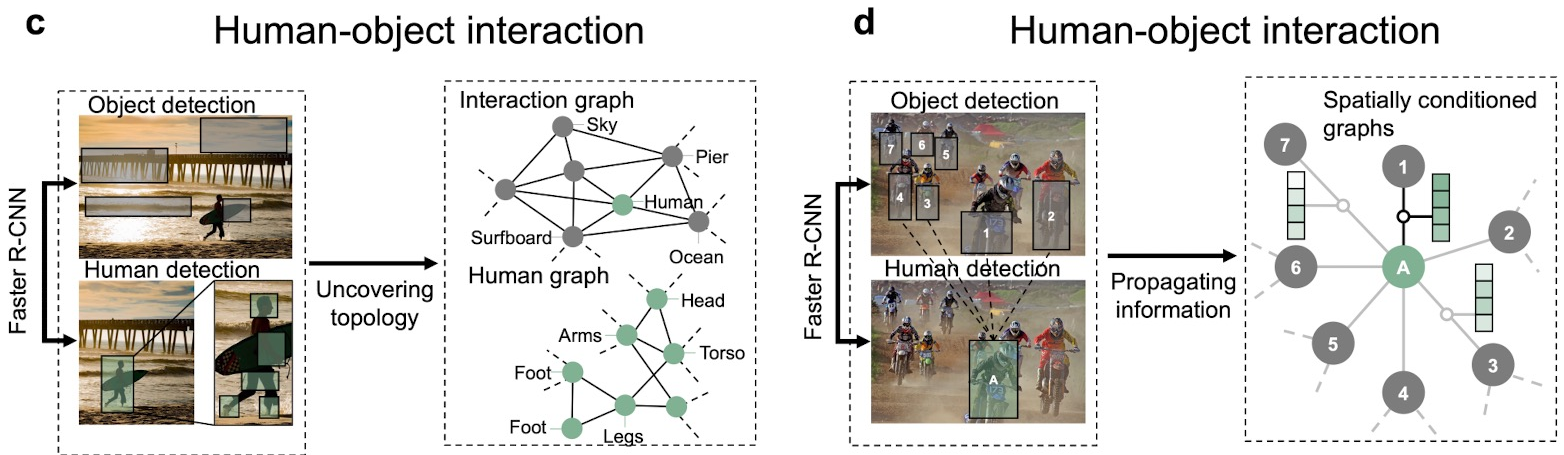

Aspect Based Sentiment Analysis (ABSA)

Tasks:

- Find aspects (start, end)
- For each of the aspects find their sentiment polarity (negative, positive)

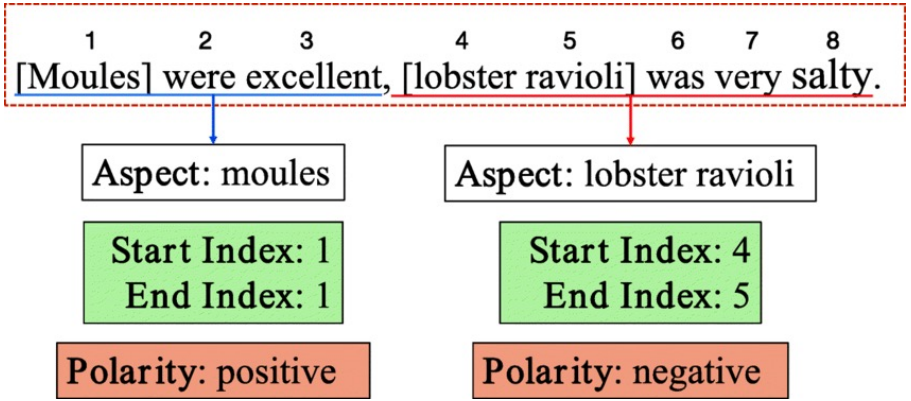

Zhou, Jie & Huang, Xiangji & Hu, Qinmin & He, Liang. (2020). Is position important? deep multi-task learning for aspect-based sentiment analysis. Applied Intelligence. 50. 10.1007/s10489-020-01760-x.

MGL on ABSA

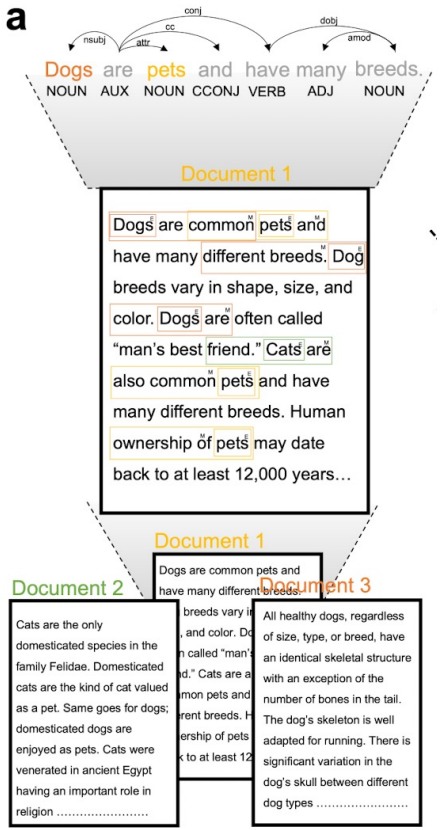
a
conj
cc
dobj
nsubj
attr
amod
Dogs are pets and have many breeds.
NOUN AUX NOUN CCONJ VERB ADJ NOUN
Document 1
Dogs are common pets and have many different breeds. Dog breeds vary in shape, size, and color. Dogs are often called "man's best friend." Cats are also common pets and have many different breeds. Human ownership of pets may date back to at least 12,000 years…
Document 1
Dogs are common pets and have many different breeds.
Document 2
Cats are the only domesticated species in the family Felidae. Domesticated cats are the kind of cat valued as a pet. Same goes for dogs; domesticated dogs are enjoyed as pets. Cats were venerated in ancient Egypt having an important role in religion ……
Document 3
All healthy dogs, regardless of size, type, or breed, have an identical skeletal structure with an exception of the number of bones in the tail. The dog's skeleton is well adapted for running. There is significant variation in the dog's skull between different dog types ……

MGL on ABSA

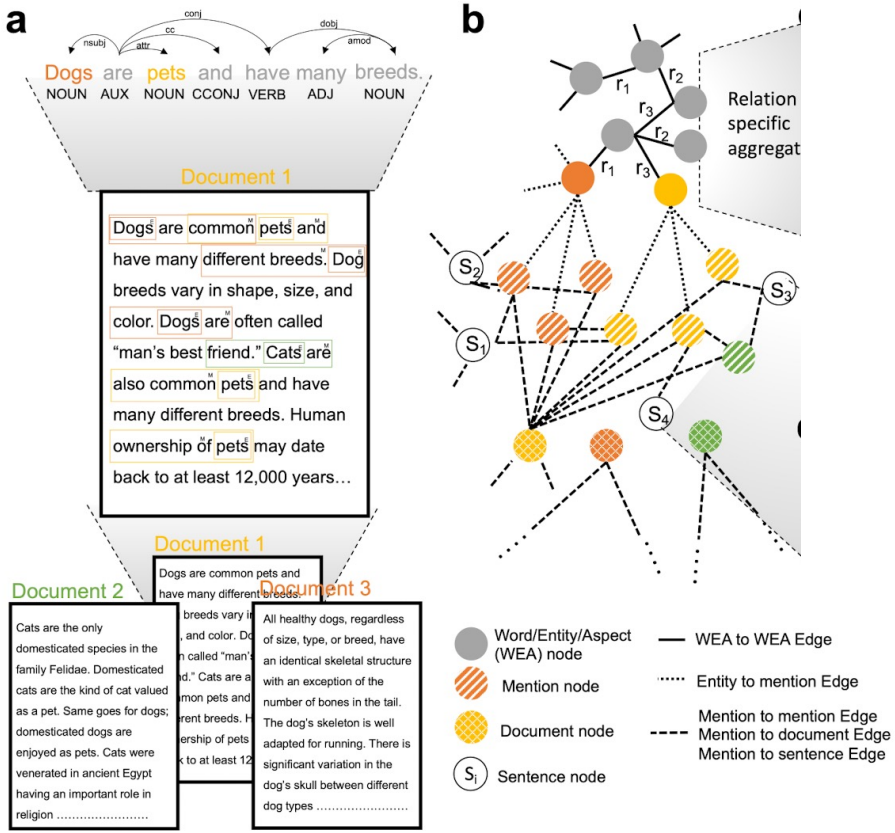
a
b
Dogs are pets and have many breeds.
NOUN AUX NOUN CCONJ VERB ADJ NOUN
Document 1
Dogs are common pets and have many different breeds. Dog breeds vary in shape, size, and color. Dogs are often called "man's best friend." Cats are also common pets and have many different breeds. Human ownership of pets may date back to at least 12,000 years…
Document 2
Cats are the only domesticated species in the family Felidae. Domesticated cats are the kind of cat valued as a pet. Same goes for dogs; domesticated dogs are enjoyed as pets. Cats were venerated in ancient Egypt having an important role in religion ……
Document 3
All healthy dogs, regardless of size, type, or breed, have an identical skeletal structure with an exception of the number of bones in the tail. The dog's skeleton is well adapted for running. There is significant variation in the dog's skull between different dog types ……
Relation specific aggregat
Word/Entity/Aspect (WEA) node
Mention node
Document node
Sentence node
WEA to WEA Edge
Entity to mention Edge
Mention to mention Edge
Mention to document Edge
Mention to sentence Edge

MGL on ABSA

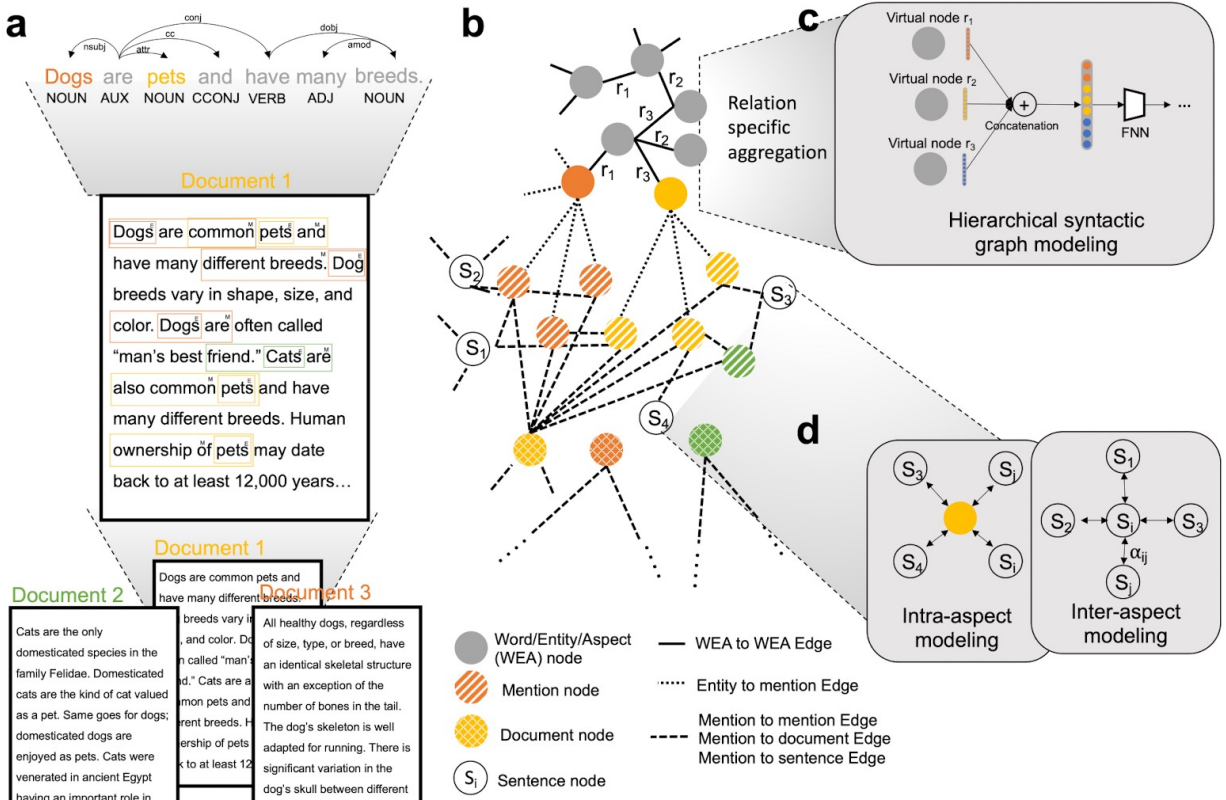
a
conj
cc
dobj
nsubj
attr
amod
Dogs are pets and have many breeds.
NOUN AUX NOUN CCONJ VERB ADJ NOUN
Document 1
Dogs are common pets and have many different breeds. Dog breeds vary in shape, size, and color. Dogs are often called "man's best friend." Cats are also common pets and have many different breeds. Human ownership of pets may date back to at least 12,000 years…
Document 1
Dogs are common pets and have many different breeds.
Document 2
Cats are the only domesticated species in the family Felidae. Domesticated cats are the kind of cat valued as a pet. Same goes for dogs; domesticated dogs are enjoyed as pets. Cats were venerated in ancient Egypt having an important role in religion ……
Document 3
All healthy dogs, regardless of size, type, or breed, have an identical skeletal structure with an exception of the number of bones in the tail. The dog's skeleton is well adapted for running. There is significant variation in the dog's skull between different dog types ……
b
Relation specific aggregation
c
Virtual node r1
Virtual node r2
Virtual node r3
Concatenation
FNN
Hierarchical syntactic graph modeling
d
Intra-aspect modeling
Inter-aspect modeling
Word/Entity/Aspect (WEA) node
Mention node
Document node
Sentence node
WEA to WEA Edge
Entity to mention Edge
Mention to mention Edge
Mention to document Edge
Mention to sentence Edge

MGL in Natural Sciences

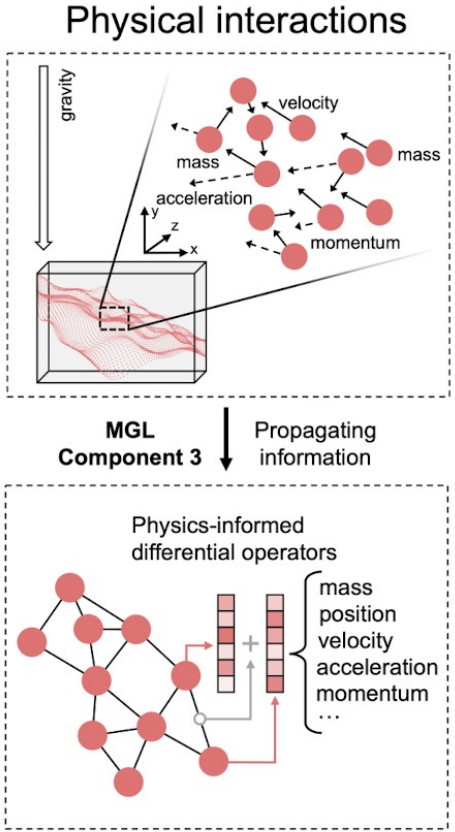

- The usual goal is to understand more about the underlying mechanics of these physical processes.
- To model the graphs, the experimental data + physical constraints are used.

Vamsi Nallapareddy

MGL in Natural Sciences

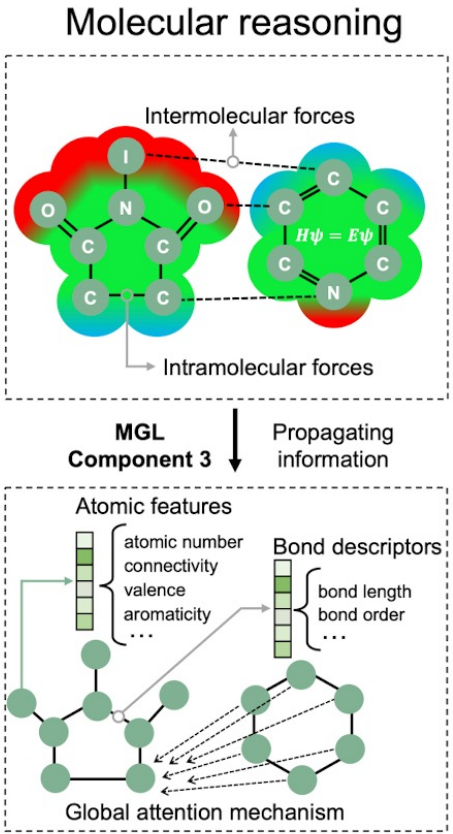

- The graphs are generated using atoms as nodes and chemical bonds as the edges.
- The general goal is to predict chemical properties of these different compounds.

MGL in Natural Sciences

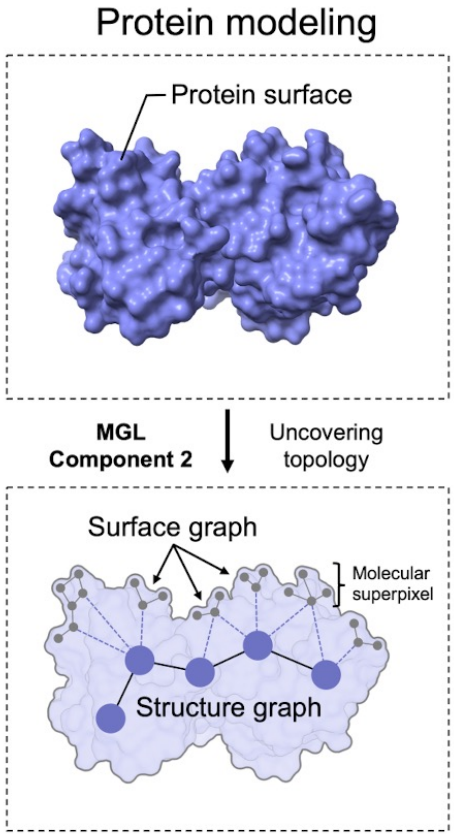

- The graphs are generated usually using protein 3D structure or even the protein surface information by assigning surface vertices.
- The usual tasks in this domain would be to understand protein-protein interactions or even protein-ligand interactions.

Maria Boulougouri

Multimodal Learning with graphs

About me

- Background: Molecular Biology
- 4th Year EDCB Ph.D. student in the LTS2 laboratory (Prof. Pierre Vandergheynst)
- Previous work:
  - Set representations and GNNs in chemistry
  - Explainable ML for single cell omics
- Current work
  - Latent graph learning with gene expression data

ARTICLES
https://doi.org/10.1038/s41592-021-01255-8

nature methods

SpaGCN: Integrating gene expression, spatial location and histology to identify spatial domains and spatially variable genes by graph convolutional network

Jian Hu[1]✉, Xiangjie Li[2], Kyle Coleman[1], Amelia Schroeder[1], Nan Ma[3], David J. Irwin[4], Edward B. Lee[5], Russell T. Shinohara[1] and Mingyao Li[1]✉

Maria Boulougouuri

Spatially resolved transcriptomics

- Gene expression profiling with spatial information to understand context

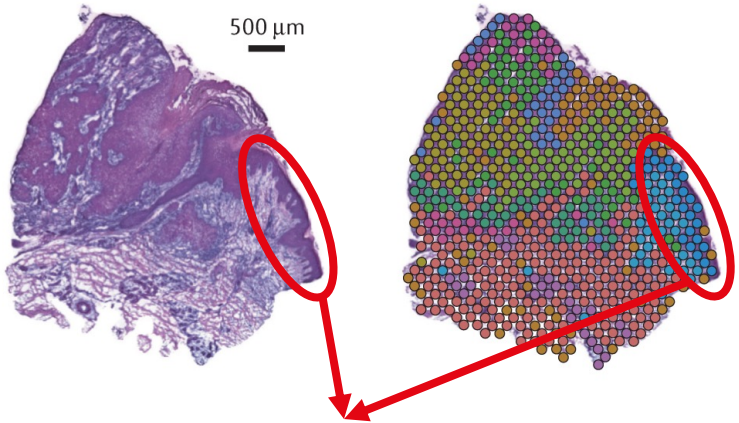

- Identify spatial regions that are coherent in gene expression *and* histology
- Identify spatially variable genes

→ Existing methods don't incorporate spatial information

SpaGCN: Integrating gene expression, spatial location and histology to identify spatial domains and spatially variable genes by graph convolutional network

Longo, S.K., Guo, M.G., Ji, A.L. *et al.* Integrating single-cell and spatial transcriptomics to elucidate intercellular tissue dynamics. *Nat Rev Genet* **22**, 627–644 (2021). https://doi.org/10.1038/s41576-021-00370-8

SpaGCN

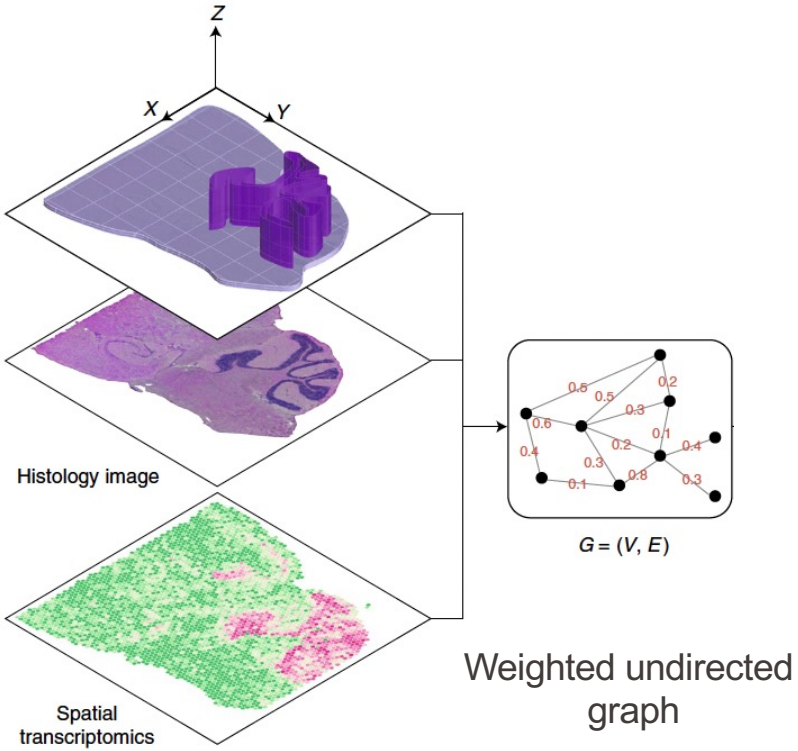

SpaGCN: Integrating gene expression, spatial location and histology to identify spatial domains and spatially variable genes by graph convolutional network

SpaGCN

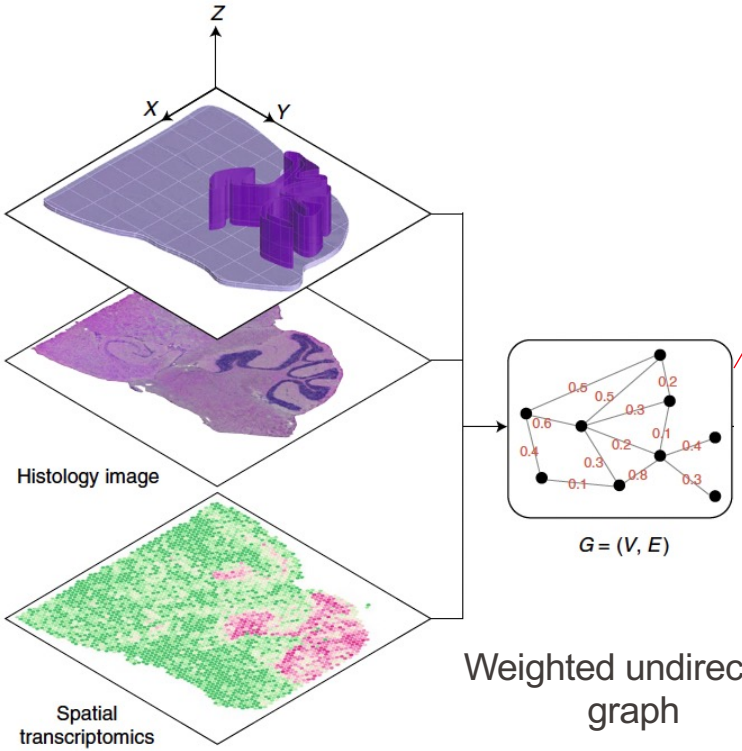

V represents a spot (instead of mRNA)
→ Segmentation-free approach

Weighted undirected graph

SpaGCN: Integrating gene expression, spatial location and histology to identify spatial domains and spatially variable genes by graph convolutional network

SpaGCN

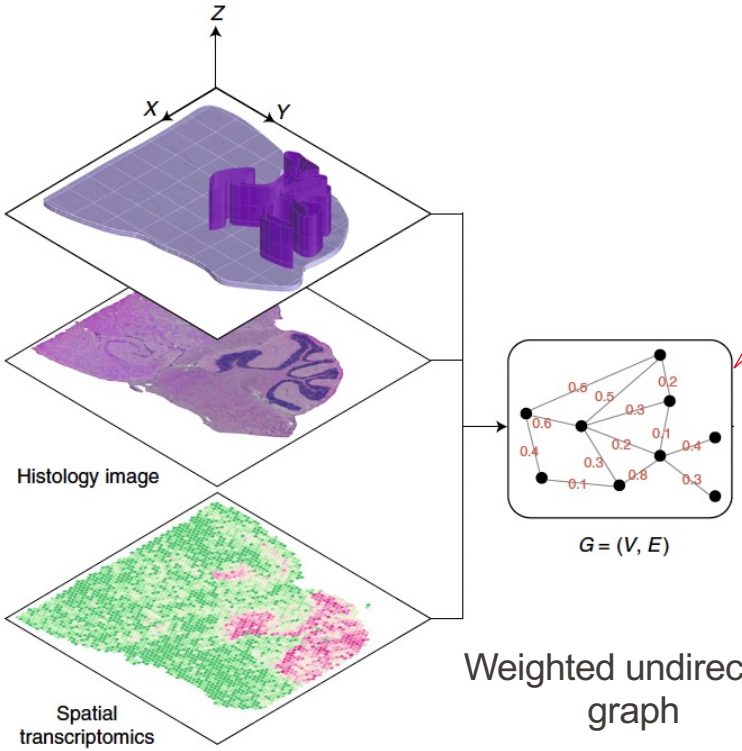

Weighted undirected graph

V represents a spot (instead of mRNA)
→ Segmentation-free approach

E reflects:
- Physical distance on slide
- Histological similarity

SpaGCN: Integrating gene expression, spatial location and histology to identify spatial domains and spatially variable genes by graph convolutional network

Maria Boulougouuri

SpaGCN

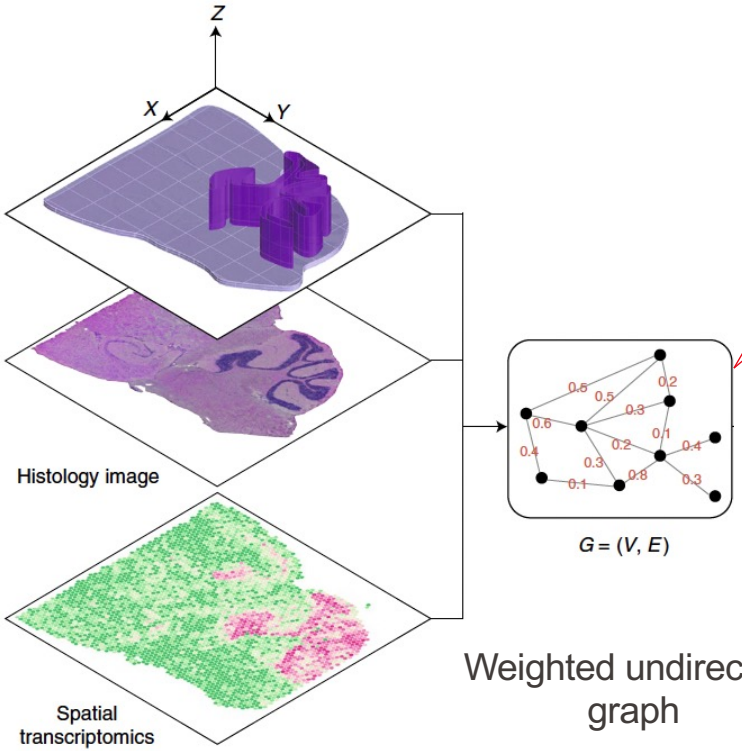

Weighted undirected graph

V represents a spot (instead of mRNA)
→ Segmentation-free approach

E reflects:
- Physical distance on slide
- Histological similarity

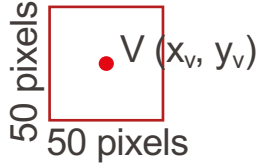

$$z_v = \frac{r_v \times V_r + g_v \times V_g + b_v \times V_b}{V_r + V_g + V_b}$$

→ Mean of RGB values (r_v, g_v, b_v)

SpaGCN: Integrating gene expression, spatial location and histology to identify spatial domains and spatially variable genes by graph convolutional network

Maria Boulougouuri

SpaGCN

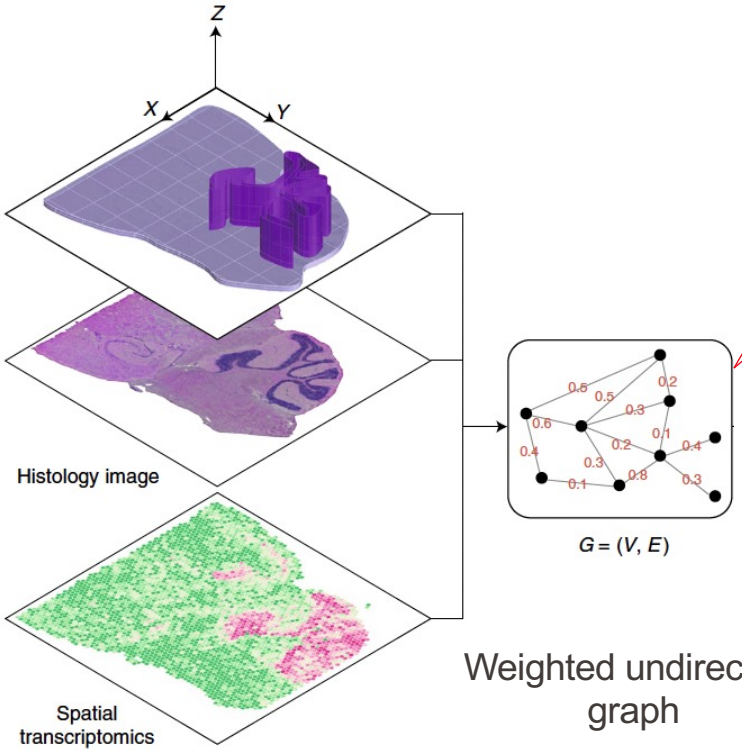

Weighted undirected graph

V represents a spot (instead of mRNA)
→ Segmentation-free approach

E reflects:
- Physical distance on slide
- Histological similarity

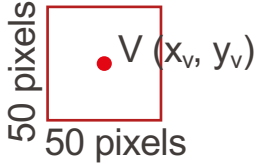

$$z_v^* = \frac{z_v - \mu_z}{\sigma_z} \times \max\left(\sigma_x, \sigma_y\right) \times s$$

→ Rescaled according to st.dev. and scaling factor **s** (can be adjusted to increase importance of histology)

SpaGCN: Integrating gene expression, spatial location and histology to identify spatial domains and spatially variable genes by graph convolutional network

Maria Boulougouuri

SpaGCN

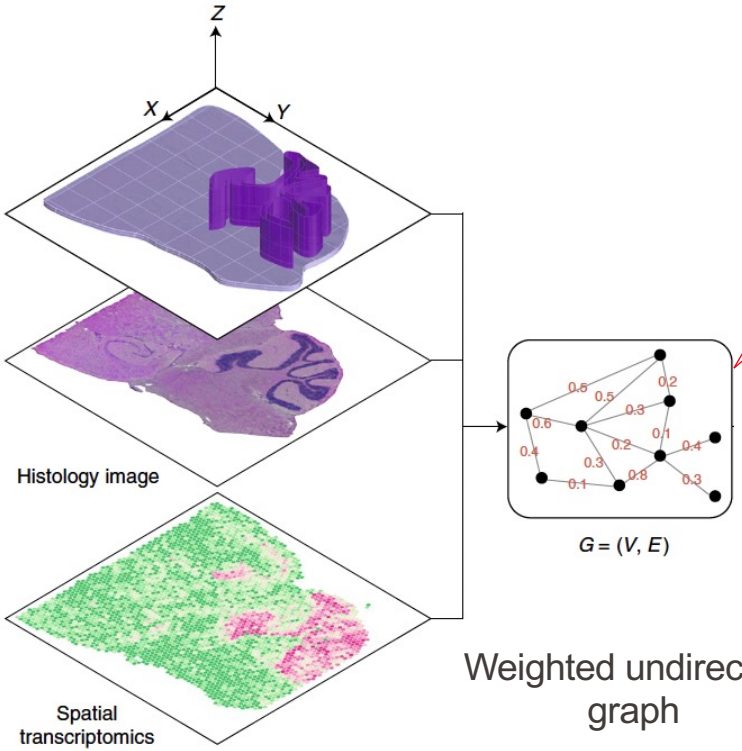

Weighted undirected graph

V represents a spot (instead of mRNA)
→ Segmentation-free approach

E reflects:
- Physical distance on slide
- Histological similarity

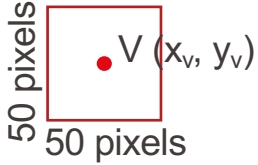

$$d(u,v) = \sqrt{(x_u - x_v)^2 + (y_u - y_v)^2 + (z_u^* - z_v^*)^2}$$

→ Distance in 3D Euclidean space

SpaGCN: Integrating gene expression, spatial location and histology to identify spatial domains and spatially variable genes by graph convolutional network

Maria Boulougouuri

SpaGCN

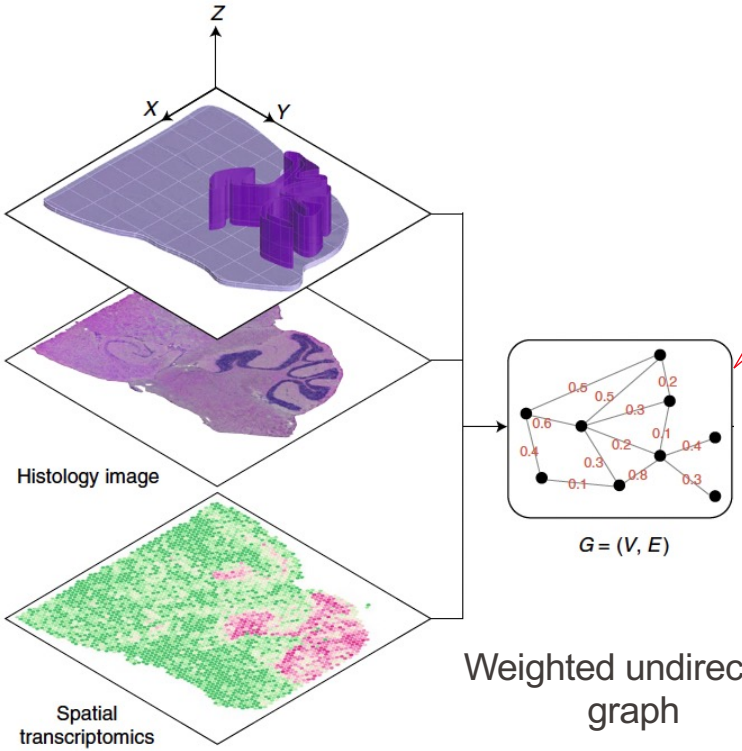

Weighted undirected graph

V represents a spot (instead of mRNA)
→ Segmentation-free approach

E reflects:
- Physical distance on slide
- Histological similarity

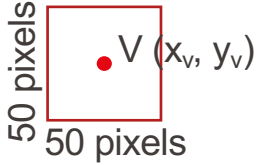

$$w(u,v) = \exp\left(-\frac{d(u,v)^2}{2l^2}\right)$$

→ Characteristic length scale **l** (can be adjusted to increase neighbour contribution to gene expression aggregation)

SpaGCN: Integrating gene expression, spatial location and histology to identify spatial domains and spatially variable genes by graph convolutional network

SpaGCN

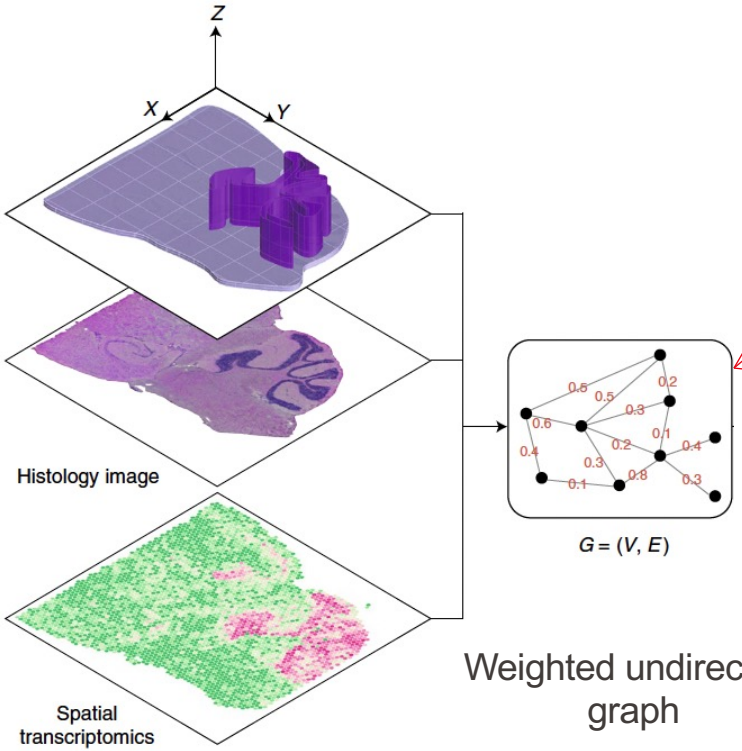

Weighted undirected graph

V represents a spot (instead of mRNA)
→ Segmentation-free approach

E reflects:
- Physical distance on slide
- Histological similarity

Feature matrix:
- 50 PCs of gene expression matrix

SpaGCN: Integrating gene expression, spatial location and histology to identify spatial domains and spatially variable genes by graph convolutional network

SpaGCN

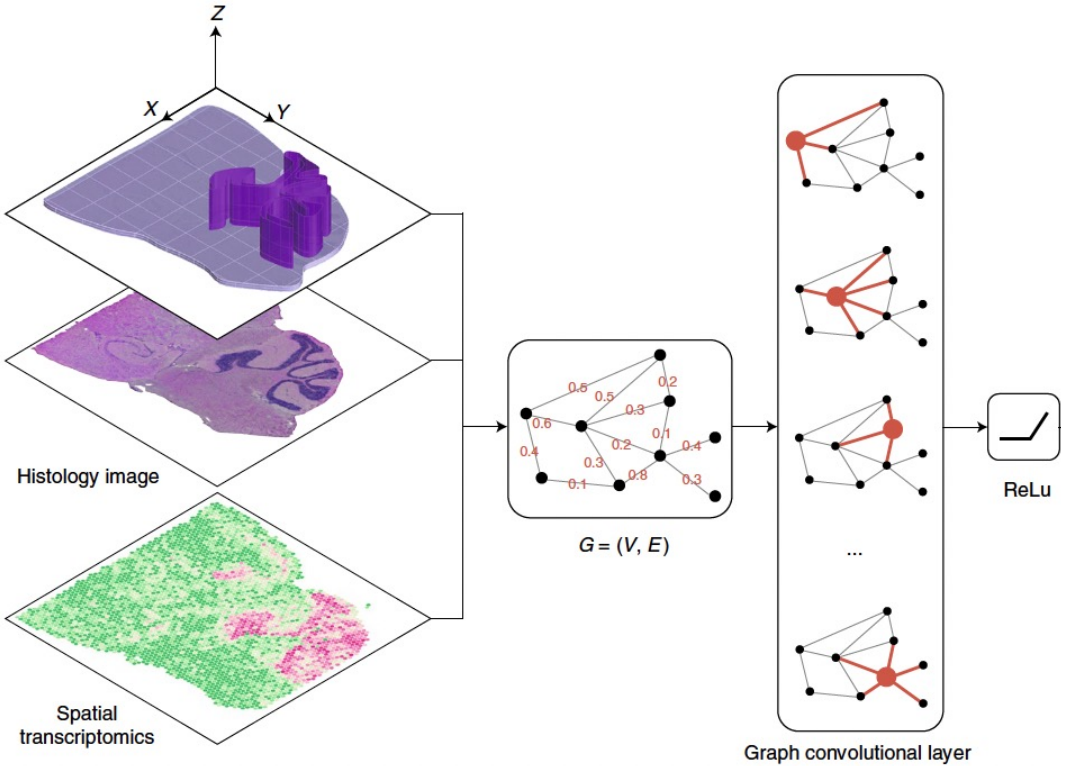

GCN
→ information aggregation from the neighborhood

SpaGCN: Integrating gene expression, spatial location and histology to identify spatial domains and spatially variable genes by graph convolutional network

Maria Boulougouuri

SpaGCN

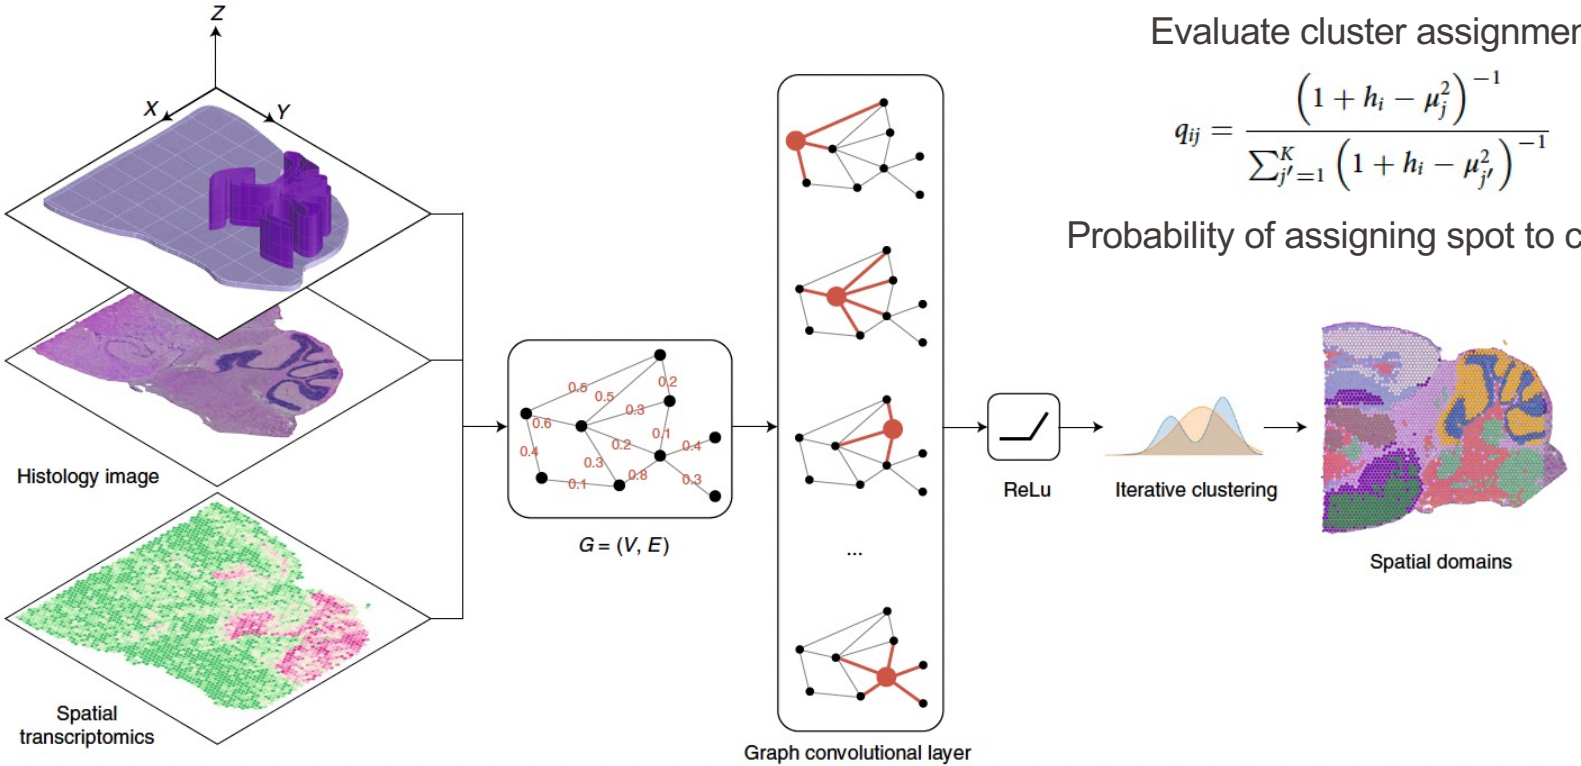

Evaluate cluster assignment:

$$q_{ij} = \frac{\left(1 + h_i - \mu_j^2\right)^{-1}}{\sum_{j'=1}^{K} \left(1 + h_i - \mu_{j'}^2\right)^{-1}}$$

Probability of assigning spot to centroid

SpaGCN: Integrating gene expression, spatial location and histology to identify spatial domains and spatially variable genes by graph convolutional network

Maria Boulougouuri

SpaGCN

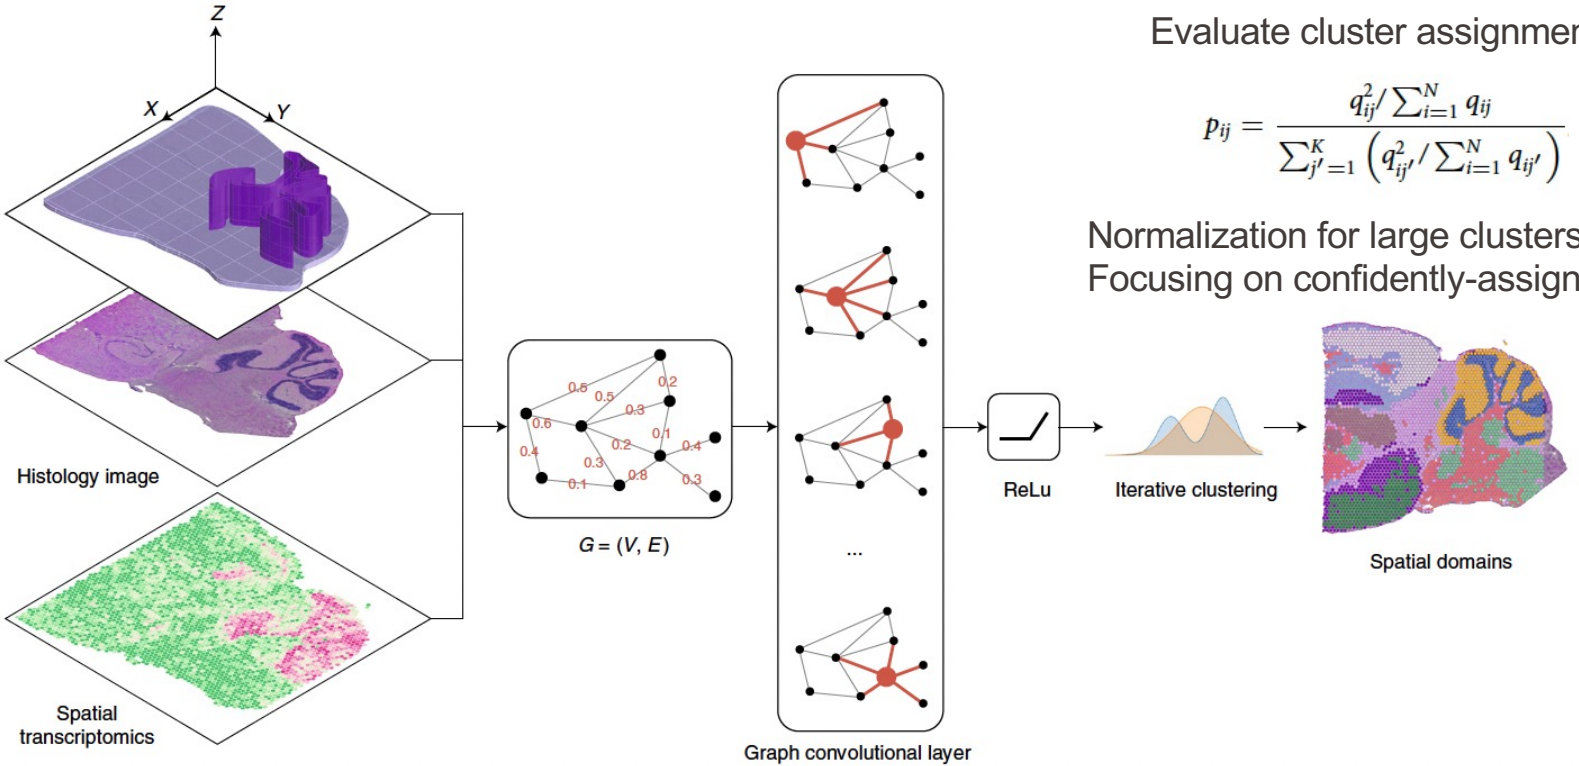

Evaluate cluster assignment:

$$p_{ij} = \frac{q_{ij}^2 / \sum_{i=1}^{N} q_{ij}}{\sum_{j'=1}^{K} \left(q_{ij'}^2 / \sum_{i=1}^{N} q_{ij'} \right)}$$

Normalization for large clusters
Focusing on confidently-assigned spots

SpaGCN: Integrating gene expression, spatial location and histology to identify spatial domains and spatially variable genes by graph convolutional network

Maria Boulougouuri

SpaGCN

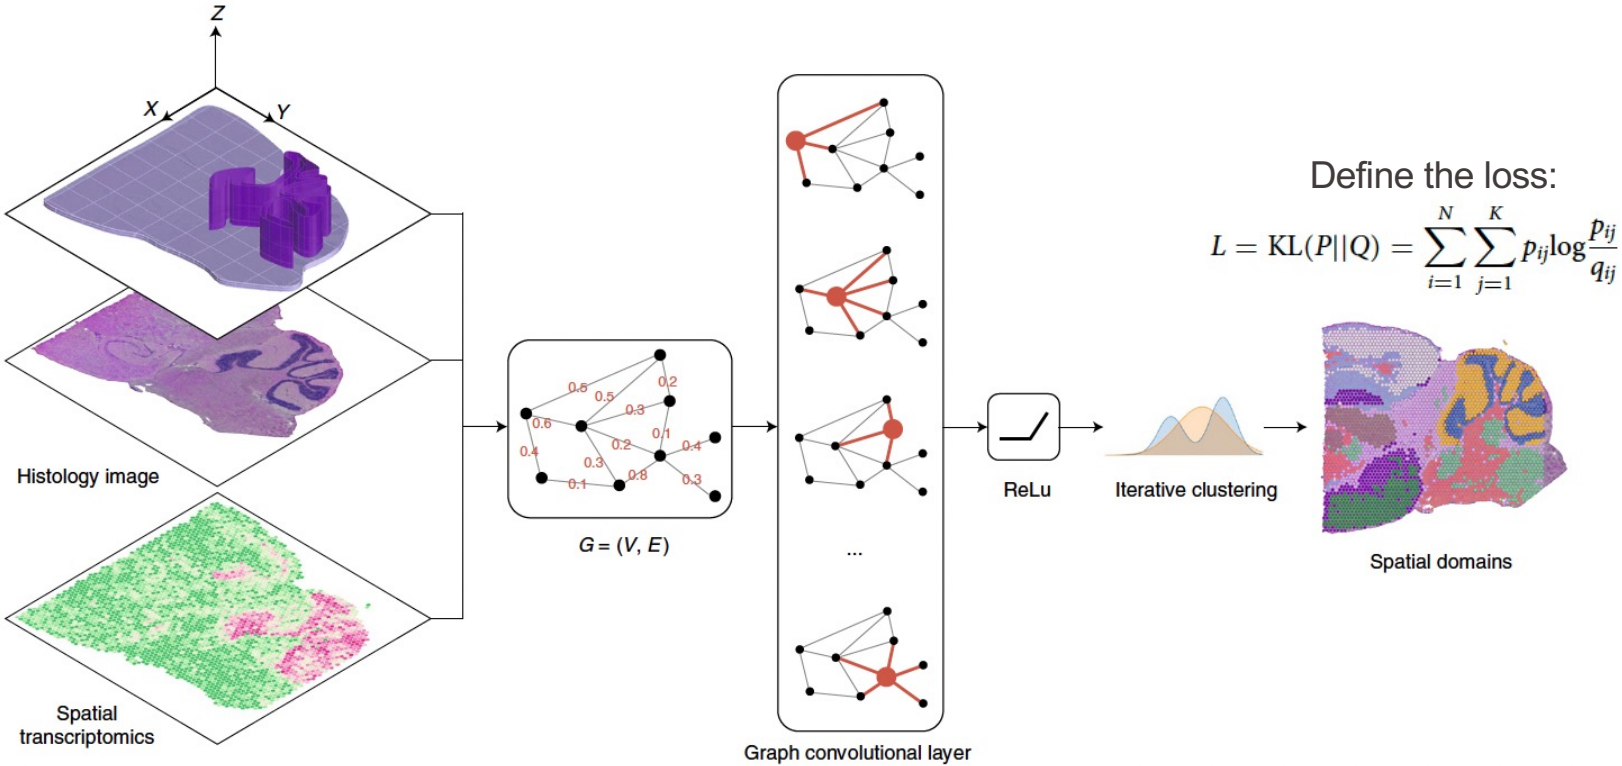

SpaGCN: Integrating gene expression, spatial location and histology to identify spatial domains and spatially variable genes by graph convolutional network

Maria Boulougouuri

SpaGCN

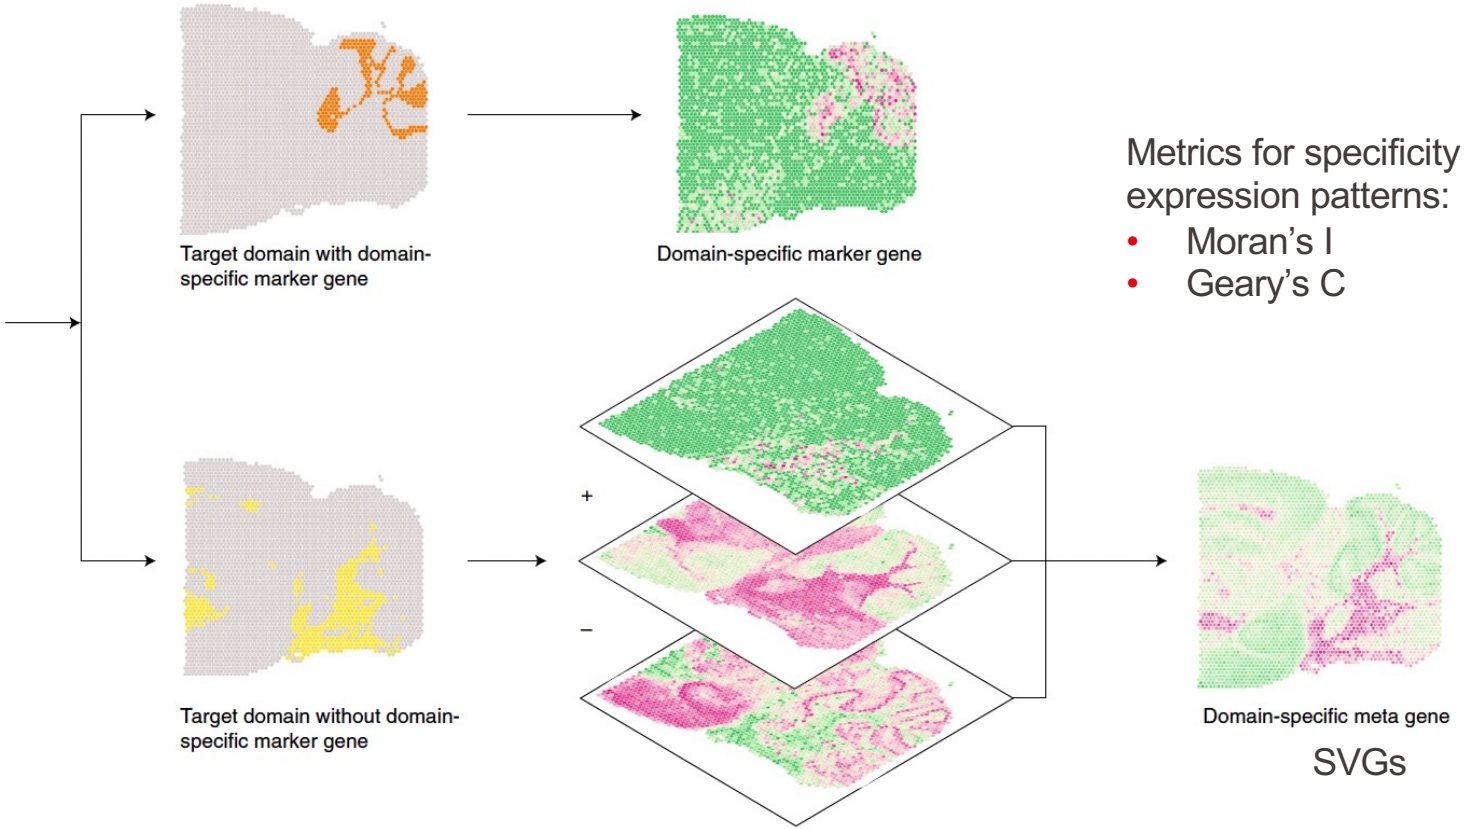

Metrics for specificity of spatial expression patterns:

- Moran's I
- Geary's C

SpaGCN: Integrating gene expression, spatial location and histology to identify spatial domains and spatially variable genes by graph convolutional network

Results: 7 datasets

- Primary pancreatic cancer tissue (human)
- Dorsolateral prefrontal cortex (human)
- Posterior brain (mouse)
- Cortex (mouse)
- Visual cortex (mouse)
- Olfactory bulb (mouse)
- Hypothalamus (mouse)

Maria Boulougouuri

Results: human pancreatic cancer

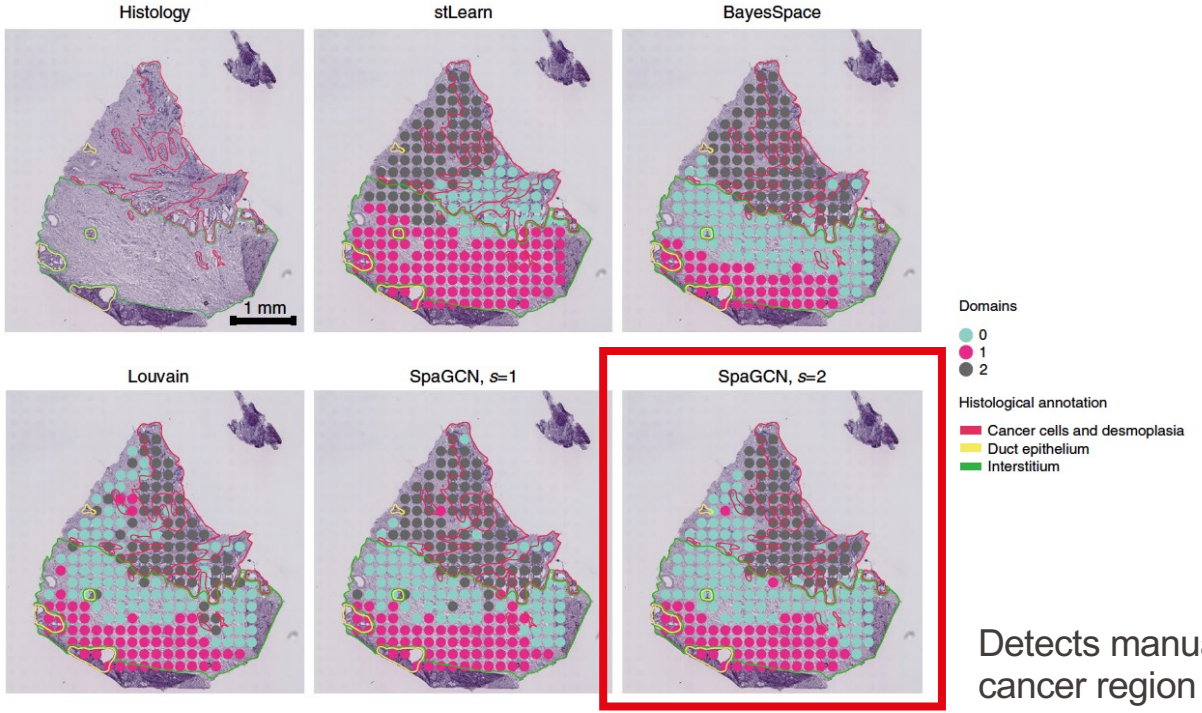

Detects manually annotated cancer region

SpaGCN: Integrating gene expression, spatial location and histology to identify spatial domains and spatially variable genes by graph convolutional network

Results: human pancreatic cancer

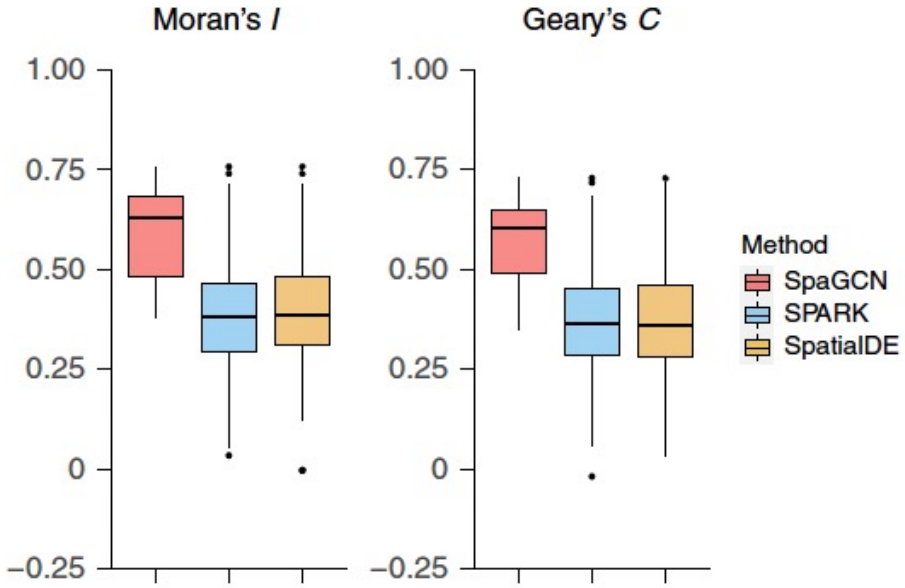

Cancer region SVGs have higher spatial autocorrelation

SpaGCN: Integrating gene expression, spatial location and histology to identify spatial domains and spatially variable genes by graph convolutional network

Conclusions

- Aims:
  - Identification of spatial domains
  - Identification of domain enriched spatially variable genes
- Advantages:
  - Weights of histology can be adjusted (tissue-dependent)
  - Graph weights are updated (technology-dependent)
- Limitations:
  - Gene expression is the main driver
  - Spatial and cell type variation are not distinguished

nature COMMUNICATIONS

ARTICLE

Check for updates

https://doi.org/10.1038/s41467-021-23774-w OPEN

MOGONET integrates multi-omics data using graph convolutional networks allowing patient classification and biomarker identification

Tongxin Wang [1,8], Wei Shao [2,8], Zhi Huang[2,3], Haixu Tang[1], Jie Zhang [4], Zhengming Ding [5✉] & Kun Huang [2,6,7✉]

Maria Boulougouuri

MOGONET

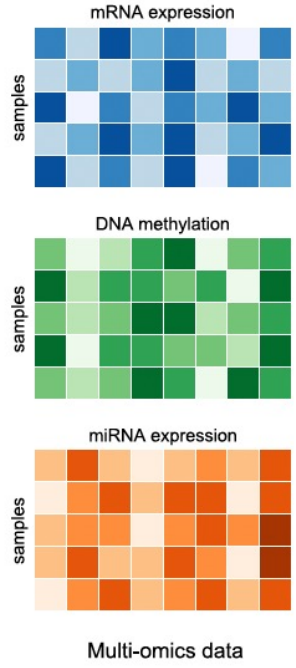

	ROSMAP → 2 classes	LGG → 2 classes	KIPAN → 3 classes	BRCA → 5 classes
mRNA expression	200	2000	2000	1000
DNA methylation	200	2000	2000	1000
miRNA expression	200	548	445	503

MOGONET integrates multi-omics data using graph convolutional networks allowing patient classification and biomarker identification

Maria Boulougouuri

MOGONET

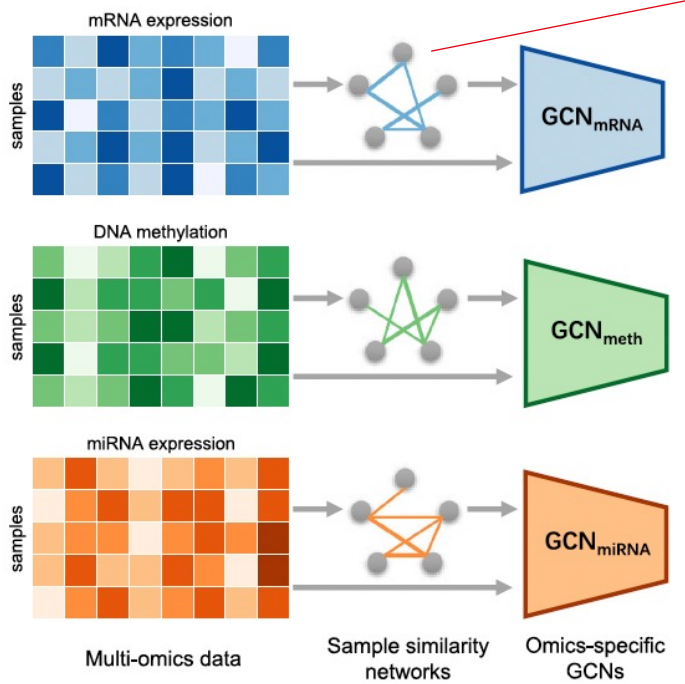

Omics-specific GCN:

- V represents a sample
- E reflects cosine distance of samples
- Threshold ε determined given a parameter **k** {2, 10}
  → represents the average number of edges per node

$$k = \sum_{i,j} I(s(\mathbf{x}_i, \mathbf{x}_j) \geq \epsilon)/n$$

MOGONET integrates multi-omics data using graph convolutional networks allowing patient classification and biomarker identification

Maria Boulougouuri

MOGONET

- Pre-training of omics-specific GCNs
- For every epoch
  - → Fix VCDN, train GCNs
  - → Fix GCNs, train VCDN

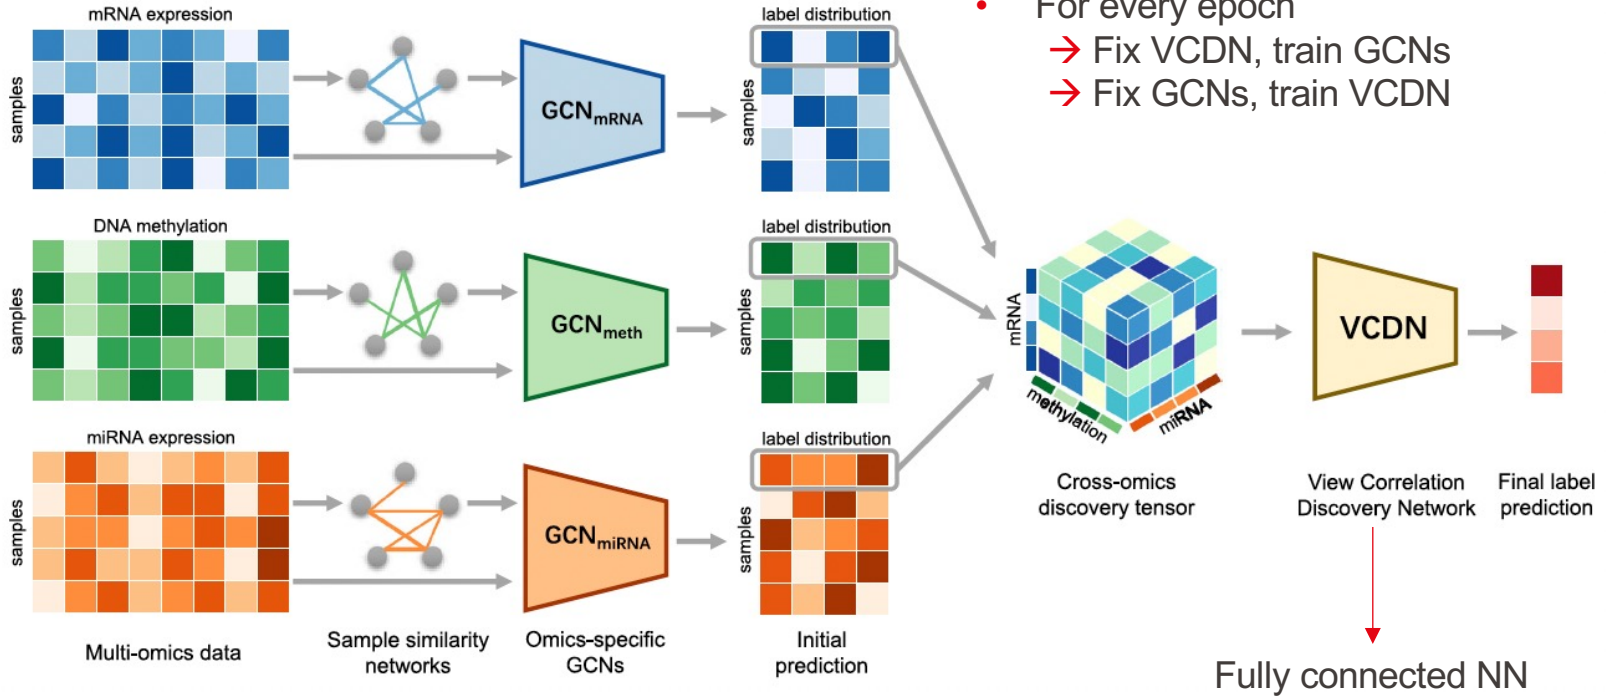

MOGONET integrates multi-omics data using graph convolutional networks allowing patient classification and biomarker identification

Results

Method
KNN
SVM
Lasso
RF
XGBoost
NN
GRridge
block PLSDA
block sPLSDA
NN_NN
NN_VCDN
MOGONET_NN (Ours)
MOGONET (Ours)

→ Concatenation of preprocessed features as input

→ Multi-omics integration methods

→ Internal controls

Results

Method
KNN
SVM
Lasso
RF
XGBoost
NN
GRridge
block PLSDA
block sPLSDA
NN_NN
NN_VCDN
MOGONET_NN (Ours)
MOGONET (Ours)

- NN_NN
  → fully connected NN for omics-specific classification
  → fully connected NN for multi-omics integration
  (vector instead of tensor as input)
- NN_VCDN
  → fully connected NN for omics-specific classification
  → VCDN for multi-omics integration
- MOGONET_NN
  → GCNs for omics-specific classification
  → fully connected NN for multi-omics integration

Results

Table 2 Classification results on ROSMAP dataset.

Method	ACC	F1	AUC
KNN	0.657 ± 0.036	0.671 ± 0.044	0.709 ± 0.045
SVM	0.770 ± 0.024	0.778 ± 0.016	0.770 ± 0.026
Lasso	0.694 ± 0.037	0.730 ± 0.033	0.770 ± 0.035
RF	0.726 ± 0.029	0.734 ± 0.021	0.811 ± 0.019
XGBoost	0.760 ± 0.046	0.772 ± 0.045	0.837 ± 0.030
NN	0.755 ± 0.021	0.764 ± 0.021	0.827 ± 0.025
GRridge	0.760 ± 0.034	0.769 ± 0.029	0.841 ± 0.023
block PLSDA	0.742 ± 0.024	0.755 ± 0.023	0.830 ± 0.025
block sPLSDA	0.753 ± 0.033	0.764 ± 0.035	0.838 ± 0.021
NN_NN	0.766 ± 0.023	0.777 ± 0.019	0.819 ± 0.017
NN_VCDN	0.775 ± 0.026	0.790 ± 0.018	0.843 ± 0.021
MOGONET_NN (Ours)	0.804 ± 0.016	0.808 ± 0.010	0.858 ± 0.024
MOGONET (Ours)	0.815 ± 0.023	0.821 ± 0.022	0.874 ± 0.012

Results

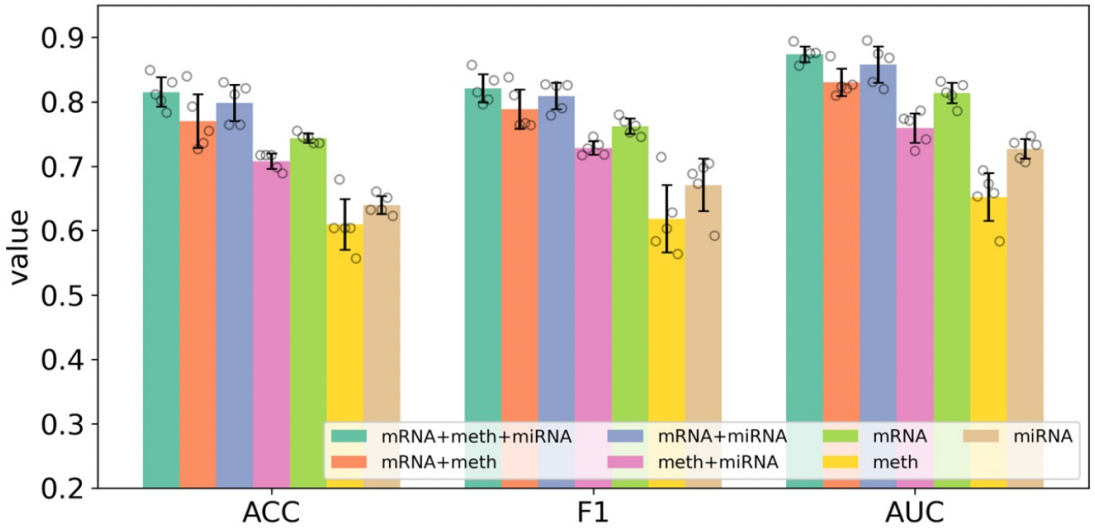
value
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
ACC
F1
AUC
mRNA+meth+miRNA
mRNA+meth
mRNA+miRNA
meth+miRNA
mRNA
meth
miRNA

Results

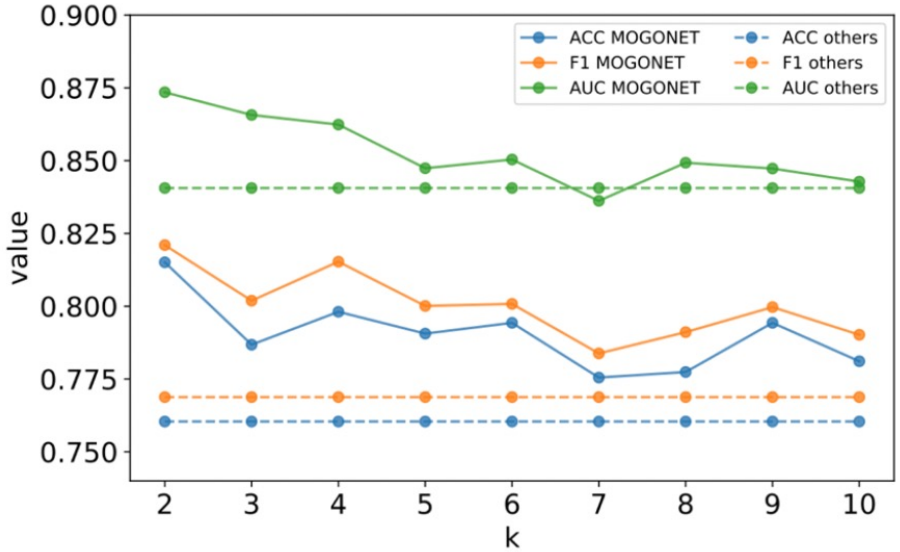
ACC MOGONET
F1 MOGONET
AUC MOGONET
ACC others
F1 others
AUC others
value
k
0.900
0.875
0.850
0.825
0.800
0.775
0.750
2
3
4
5
6
7
8
9
10

Results

Table 3 Classification results on LGG dataset.

Method	ACC	F1	AUC
KNN	0.729 ± 0.034	0.738 ± 0.033	0.799 ± 0.038
SVM	0.754 ± 0.046	0.757 ± 0.050	0.754 ± 0.046
Lasso	0.761 ± 0.018	0.767 ± 0.022	0.823 ± 0.027
RF	0.748 ± 0.012	0.742 ± 0.010	0.823 ± 0.010
XGBoost	0.756 ± 0.040	0.767 ± 0.032	0.840 ± 0.023
NN	0.737 ± 0.023	0.748 ± 0.024	0.810 ± 0.037
GRridge	0.746 ± 0.038	0.756 ± 0.036	0.826 ± 0.044
block PLSDA	0.759 ± 0.025	0.738 ± 0.031	0.825 ± 0.023
block sPLSDA	0.685 ± 0.027	0.662 ± 0.030	0.730 ± 0.026
NN_NN	0.740 ± 0.039	0.756 ± 0.036	0.824 ± 0.036
NN_VCDN	0.740 ± 0.030	0.771 ± 0.021	0.826 ± 0.031
MOGONET_NN (Ours)	0.804 ± 0.025	0.811 ± 0.023	0.832 ± 0.029
MOGONET (Ours)	0.816 ± 0.016	0.814 ± 0.014	0.840 ± 0.027

Results

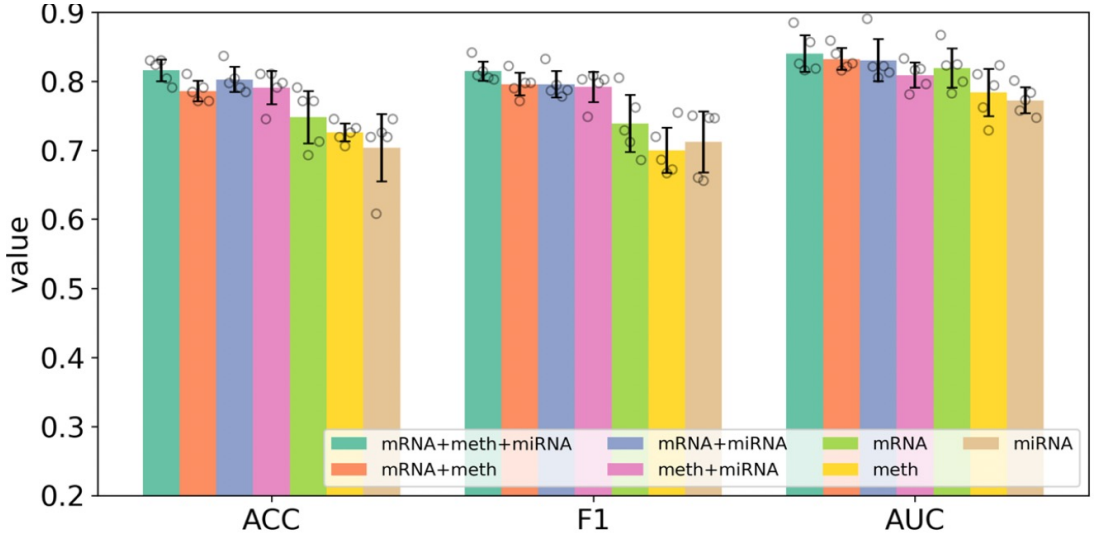

MOGONET integrates multi-omics data using graph convolutional networks allowing patient classification and biomarker identification

Results

Table 4 Classification results on BRCA dataset.

Method	ACC	F1_weighted	F1_macro
KNN	0.742 ± 0.024	0.730 ± 0.023	0.682 ± 0.025
SVM	0.729 ± 0.018	0.702 ± 0.015	0.640 ± 0.017
Lasso	0.732 ± 0.012	0.698 ± 0.015	0.642 ± 0.026
RF	0.754 ± 0.009	0.733 ± 0.010	0.649 ± 0.013
XGBoost	0.781 ± 0.008	0.764 ± 0.010	0.701 ± 0.017
NN	0.754 ± 0.028	0.740 ± 0.034	0.668 ± 0.047
GRridge	0.745 ± 0.016	0.726 ± 0.019	0.656 ± 0.025
block PLSDA	0.642 ± 0.009	0.534 ± 0.014	0.369 ± 0.017
block sPLSDA	0.639 ± 0.008	0.522 ± 0.016	0.351 ± 0.022
NN_NN	0.796 ± 0.012	0.784 ± 0.014	0.723 ± 0.018
NN_VCDN	0.792 ± 0.010	0.781 ± 0.006	0.721 ± 0.018
MOGONET_NN (Ours)	0.805 ± 0.017	0.782 ± 0.030	0.737 ± 0.038
MOGONET (Ours)	0.829 ± 0.018	0.825 ± 0.016	0.774 ± 0.017

Results

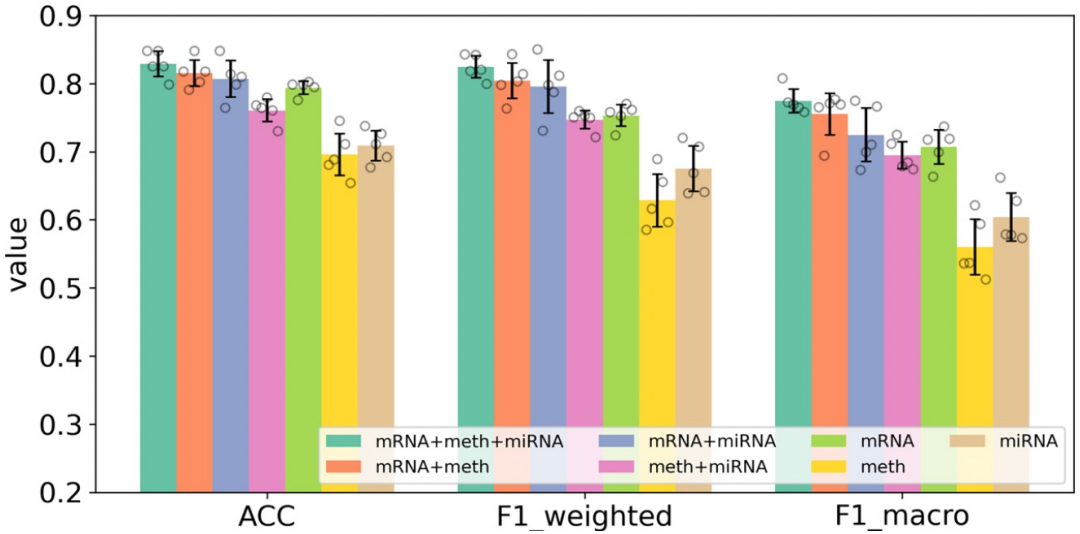

MOGONET integrates multi-omics data using graph convolutional networks allowing patient classification and biomarker identification

Results

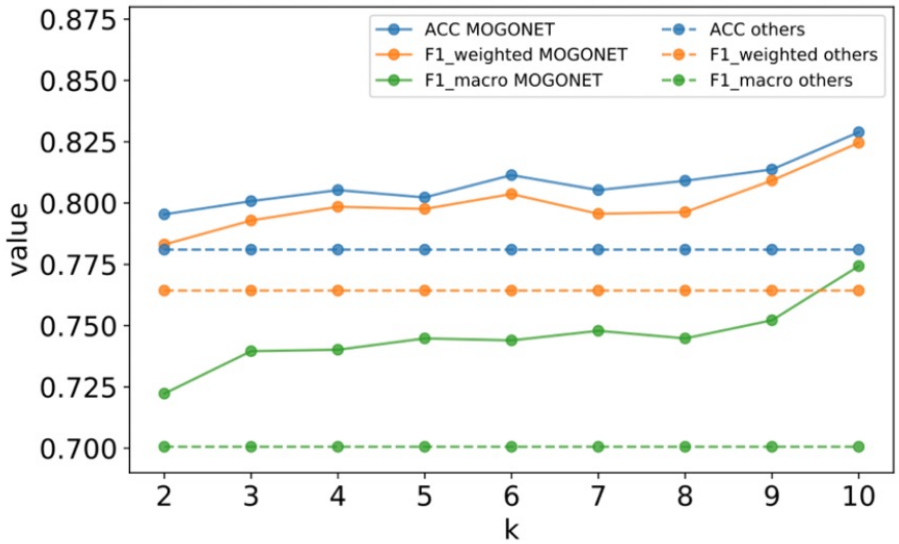
ACC MOGONET
F1_weighted MOGONET
F1_macro MOGONET
ACC others
F1_weighted others
F1_macro others
value
k
0.875
0.850
0.825
0.800
0.775
0.750
0.725
0.700
2
3
4
5
6
7
8
9
10

Conclusions

- Aims:
  - Supervised multi-omics integration method for biomedical classification tasks
  - Demonstrate that both GCNs and VCDN are essential
  - Adversarial attacks for biomarker discovery
- Advantages:
  - GCNs can utilize both the features *and* the geometrical structures of the data
  - Flexibility (number of omics, type, etc.)
  - First method to consider the correlations among different omics data types.
    → less biased toward certain omics data types
- Limitations:
  - Benchmark selection

General conclusions

- Underexplored field in biological applications

- What is the definition of a modality
  - Different views of the same entity (histology + spatial transcriptomics)
  - Different entities (mRNAs + miRNAs)
  - Does base knowledge count (PPI)

- Issues:
  - Lack of data
  - Lack of correlation between different views of the same entity (genes + proteins)
  - Sparsity

Thank you!

Maria Boulougouuri