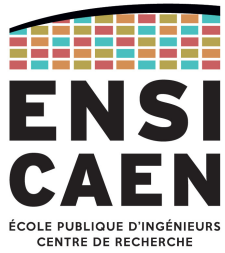

Graph Classification

Septempber 3, 2019
Salerno

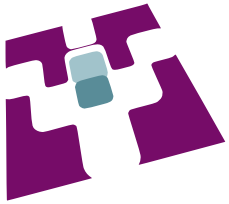

Pros

- Interest: provide a compact encoding of both :
  - a decomposition of objects into meaningful sub-parts,
  - the relationships between these sub-parts.
- Applications:
  - Image processing: segmentation, boundary detection,
  - Pattern Recognition: printed characters, documents, objects (buildings, brain structures), faces, gestures, molecules,. . . ,
  - Image registration,
  - Understanding of structured scenes.
  - . . .

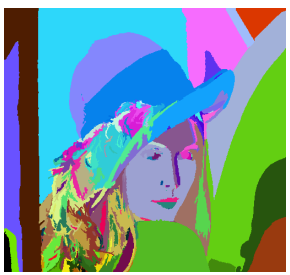
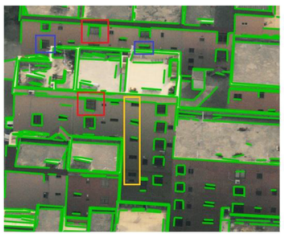
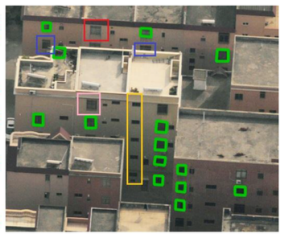
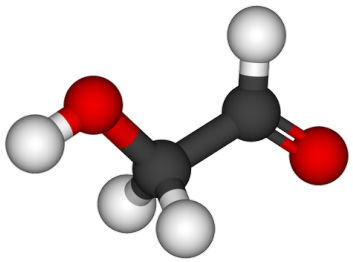

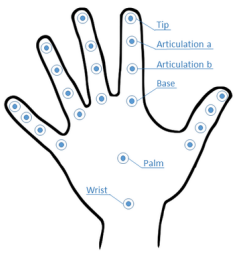

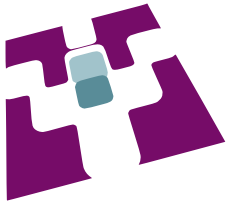

Cons

- Even simple questions are difficult:
  - Are this two graphs the same ?
    - $\mathcal{NP}$-intermediate
  - Is this graph a sub-part of this graph ?
    - $\mathcal{NP}$-complete
  - What is the distance between these two graphs ?
    - $\mathcal{NP}$-hard (for the usual distance)
  - What is the mean/median of a set of graphs ?
    - $\mathcal{NP}$-hard (for the usual distance)
- Pattern Recognition implies:
  1. Metrics,
  2. Non exponential execution times.

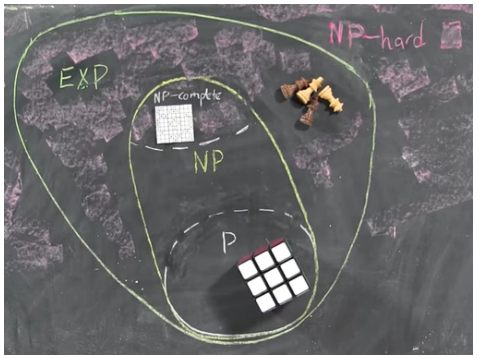

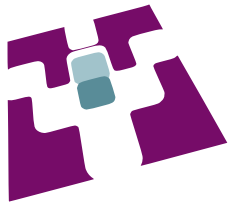

Graph space as an Orbifold

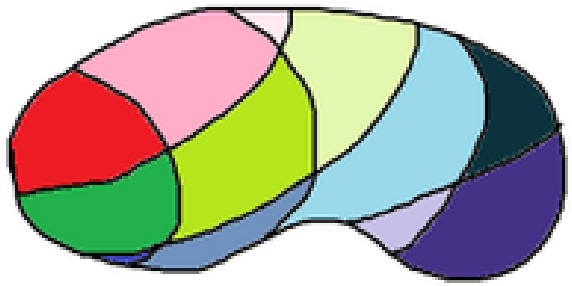

[Jain, 2016, Jain and Wysotzki, 2004, Jain, 2014].

Orbifolds theory: basics

- Let $\mathcal{T}$ be the manifold of $n \times n$ matrices.
- One graph may have multiple matrix representations

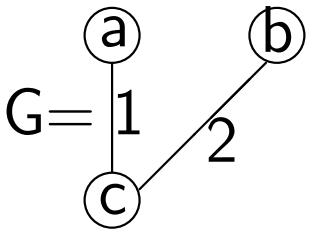

	1	2	3
1	a	0	1
2	0	b	2
3	1	2	c

x

	1	2	3
1	c	1	2
2	1	a	0
3	2	0	b

y

	1	2	3
1	b	2	0
2	2	c	1
3	0	1	a

z...

...

- Let $\mathcal{P}$ denote the set of $n \times n$ permutation matrices:

$$\forall (x, y) \in \mathcal{T}\ x \sim y \Leftrightarrow \exists P \in \mathcal{P} | x = P^t y P$$

- $\mathcal{T}/\sim$ is called an orbifold
- A graph G is encoded in $\mathcal{T}/\sim$ by $[X_G] = \{x, y, z, \dots\}$.

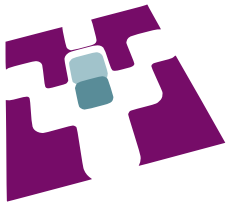

Orbifolds theory: basics

- Let X and Y denote two elements of $\mathcal{T}/\sim$:

$$\begin{cases} < X, Y > & = & \max\{< x, y > \ x \in X, y \in Y\} \\ \delta(X, Y) & = & \sqrt{\|X\|^2 + \|Y\|^2 - 2 < X, Y >} \\ & = & \min_{x \in X, y \in Y} \|x - y\| \end{cases}$$

  This metric is called the graph alignment metric.
- Using real attributes, the space $(\mathcal{G}, \delta)$ is:
  - A complete metric space,
  - a geodesic space,
  - locally compact,
  - every closed bounded subset of $(\mathcal{G}, \delta)$ is compact.

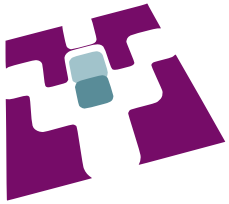

Orbifold theory: Applications

- Computation of the sample mean [Jain, 2016]
  - The sample mean of a set of graphs always exists.
  - Under some conditions, the set of sample means reduces to a singleton.
- Central clustering algorithms [Jain and Wysotzki, 2004],
- Generalized linear classifiers [Jain, 2014]

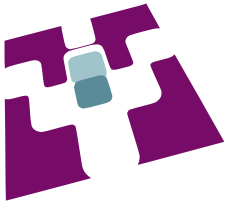

Orbifold theory: limits

- The enumeration of $[X]$ requires $n!$ computations.
- Graph metric is induced by graph representation: Both concepts can not be distinguished.

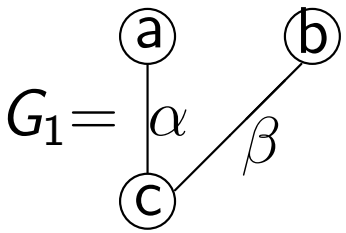

a	0	α
0	b	β
α	β	c

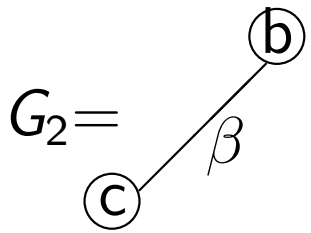

0	0	0
0	b	β
0	β	c

- We have:

$$\delta(G_1, G_2) = \sqrt{a^2 + \alpha^2}$$

We should search for a more flexible metric.

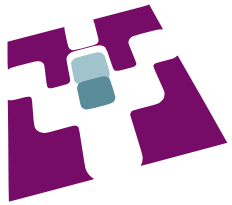

Graph Edit distance

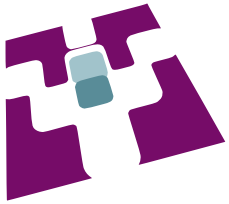

Edit Paths

Definition (Edit path)

Given two graphs G_1 and G_2 an **edit path** between G_1 and G_2 is a sequence of node or edge removal, insertion or substitution which transforms G_1 into G_2.

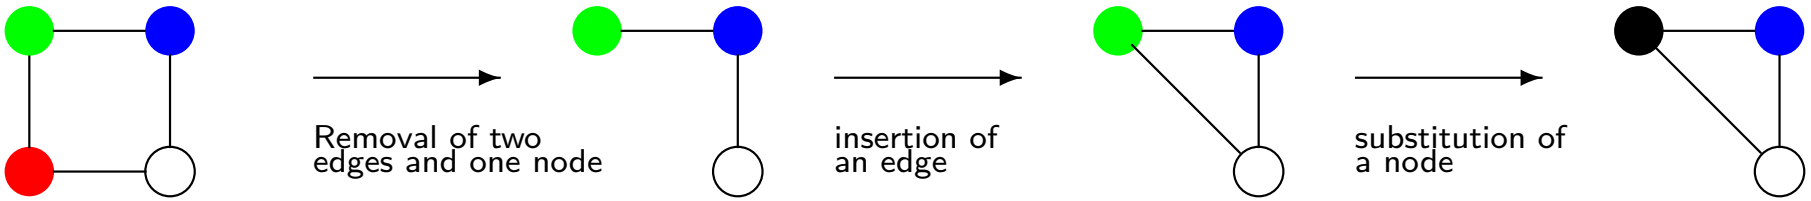

A substitution is denoted $u \rightarrow v$, an insertion $\epsilon \rightarrow v$ and a removal $u \rightarrow \epsilon$.

Alternative edit operations such as merge/split have been also proposed[Ambauen et al., 2003].

Costs

Let $c(.)$ denote the cost of any elementary operation. The cost of an edit path is defined as the sum of the costs of its elementary operations.

- All cost are positive: $c() \geq 0$,
- A node or edge substitution which does not modify a label has a 0 cost: $c(l \rightarrow l) = 0$.

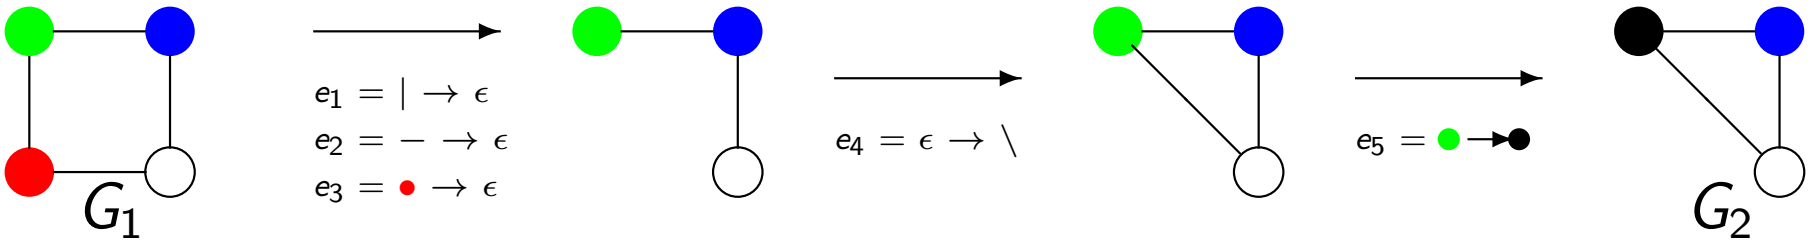

If all costs are equal to 1, the cost of this edit path is equal to 5.

Conversely to graph alignment metric, costs allow to distinguish graph representation and graph metrics.

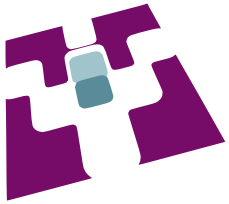

Graph edit distance

Definition (Graph edit distance)

The graph edit distance between G_1 and G_2 is defined as the cost of the less costly path within $\Gamma(G_1, G_2)$. Where $\Gamma(G_1, G_2)$ denotes the set of edit paths between G_1 and G_2.

$$d(G_1, G_2) = \min_{\gamma \in \Gamma(G_1, G_2)} cost(\gamma) = \min_{\gamma \in \Gamma(G_1, G_2)} \sum_{e \in \gamma} c(e)$$

An $\mathcal{NP}$-hard problem.

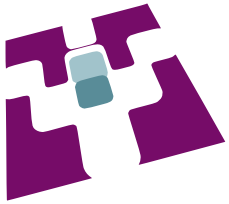

Graph edit distance: Main methods

1. A^* like algorithms [Abu-Aisheh, 2016],

2. Formulation as a quadratic problem [Bougleux et al., 2017] solved by Franck-Wolfe [Frank and Wolfe, 1956] like algorithms (see also:
[Liu and Qiao, 2014, Boria et al., 2018, Daller et al., 2018].

3. Integer Programming [Lerouge et al., 2017, Darwiche, 2018]

4. Fast (and often rough)
approximations [Riesen and Bunke, 2009, Gaüzère et al., 2014a, Carletti et al., 2015, Blumenthal et al., 2018]

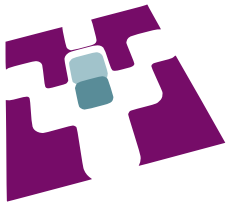

Graph edit distance: Comparison

- Relative comparison:

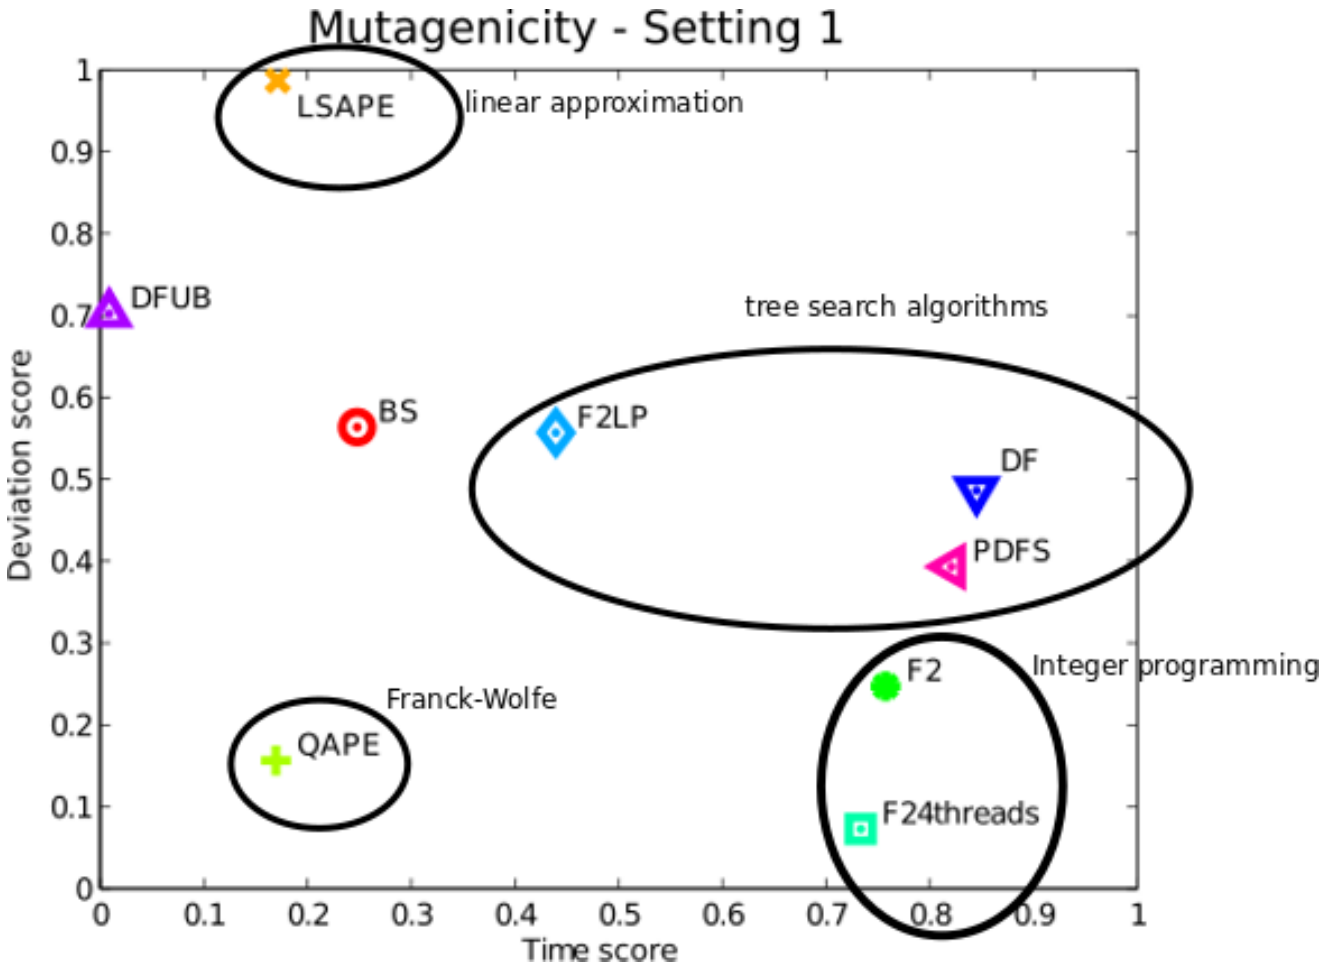

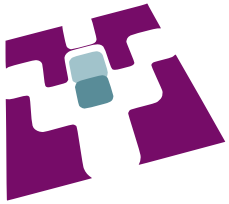

Graph Edit distance: Conclusion

- ☺ Allows to dissociate graph representation and graph metric.
- ☺ Constitutes a fine and intuitive metric between graphs.
- ☹ It is $\mathcal{NP}$-hard to compute,
- ☹ It is not conditionally definite negative (Krein space).
  - Most of machine learning machinery should be adapted [Loosli et al., 2016].
  - Weak properties: E.g. the median is usually not unique.

Library: `https://github.com/Ryurin/Python_GedLib`

Graph Kernels

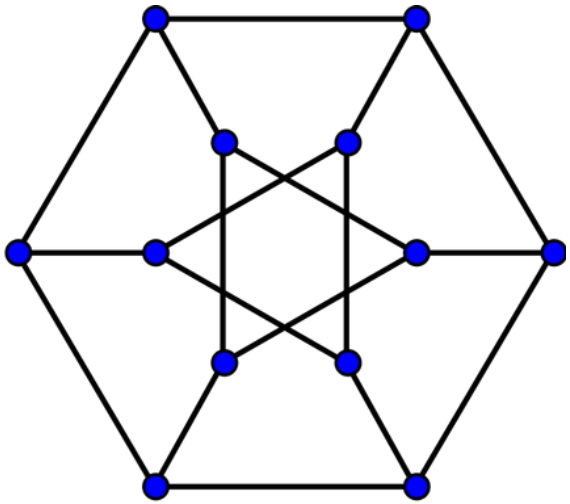

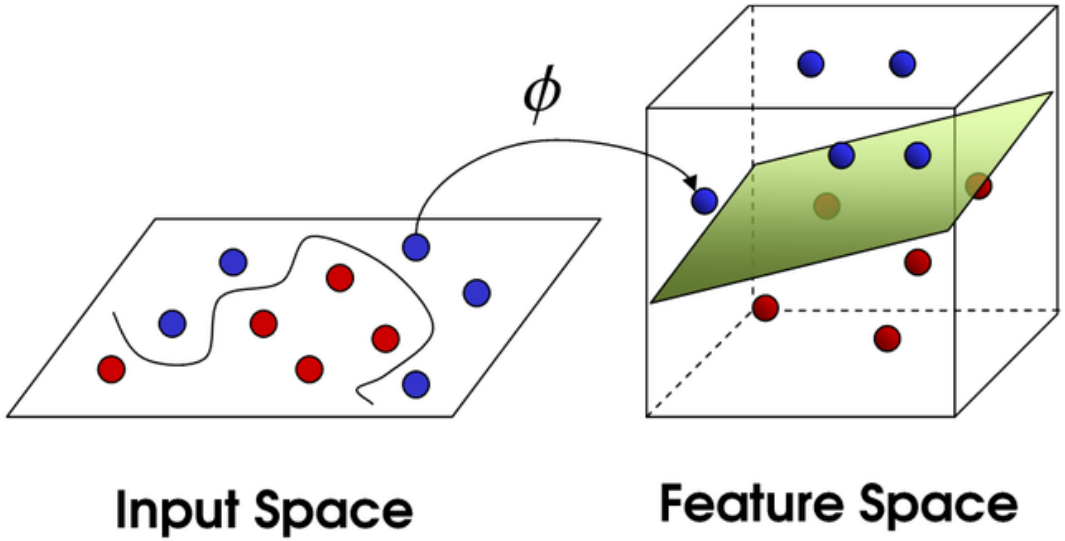

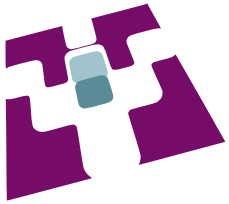

Kernels : Definition

- A kernel k is a **symmetric** similarity measure on a set χ

$$\forall (x, y) \in \chi^2,\ k(x, y) = k(y, x)$$

- k is said to be **definite positive** (d.p.) iff k is symmetric and iff:

$$\left.\begin{array}{rcl} \forall (x_1, \ldots, x_n) & \in & \chi^n \\ \forall (c_1, \ldots, c_n) & \in & \mathbb{R}^n \end{array}\right\} \sum_{i=1}^{n}\sum_{j=1}^{n} c_i c_j k(x_i, x_j) \geq 0$$

- $K = (k(x_i, x_j))_{(i,j)\in\{1,\ldots,n\}}$ is the Gramm matrix of k. k is d.p. iff:

$$\forall c \in \mathbb{R}^n - \{0\},\ c^t K c \geq 0$$

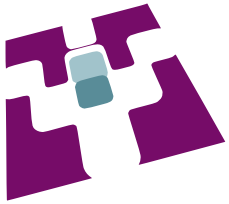

Kernels and scalar products

[Aronszahn, 1950] :

A kernel k is d.p. on a space χ
if and only if
it exists

- one Hilbert space $\mathcal{H}$ and
- a function $\varphi : \chi \to \mathcal{H}$

such that:

$$k(x, y) =< \varphi(x), \varphi(y) >$$

Open the way for rich interactions between graphs and usual machine learning methods: SVM, kPCA, MKL,...

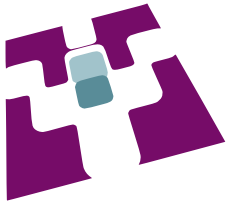

Graph Kernel: methods

$$K(G, G') = \sum_{x \in \mathcal{B}(G)} \sum_{y \in \mathcal{B}(G')} k(x, y)$$

where $\mathcal{B}(G)$ is a bag of patterns deduced from G [Haussler, 1999].

- Linear Patterns:
  - Random Walk Kernel [Kashima et al., 2003, Gärtner et al., 2003
  - n order path kernel: [Ralaivola et al., 2005, Dupé and Brun, 20
  - Shortest Path [Hermansson et al., 2015]
- Non linear patterns:
  - Tree Pattern kernel [Mahé and Vert, 2009, Shervashidze and Borgwardt, 2009]
  - Graphlet Kernel [Shervashidze et al., 2009]
  - Treelet kernels [Gaüzère et al., 2014b]

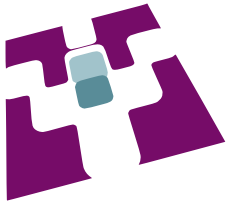

Some results

Method	RMSE(^oC)	Time(s)	
		Learning	Prediction
Gaussian edit distance	10.27	1.35	0.05
Random Walks	18.72	19.10	0.57
Path Kernel	12.24	7.83	0.18
Tree Pattern Kernel	11.02	4.98	0.03
Treelet Kernel (TK)	8.10	**0.07**	**0.01**
TK + MKL	**5.24**	70	**0.01**

Boiling point prediction on acyclic molecule dataset using 90% of the dataset as train set and remaining 10% as test set.

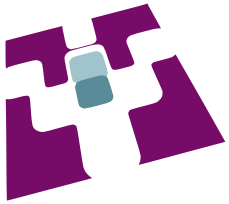

Graph kernels:

- Pros:
  - Graph kernels provide an implicit embedding of graphs,
  - Its open the way to the application of many statistical tools to graphs,
- Cons:
  - Graph kernels are usually based on a notion of bag which only provides a rough similarity measure.
  - The graph feature extraction process has been moved to the design of a similarity measure (the kernel). Such a measure remains largely "handcrafted".
- Libraries:
  - `https://github.com/jajupmochi/py-graph`
  - `http://chemcpp.sourceforge.net/html/index.html`

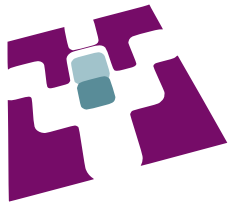

Graph Neural Network

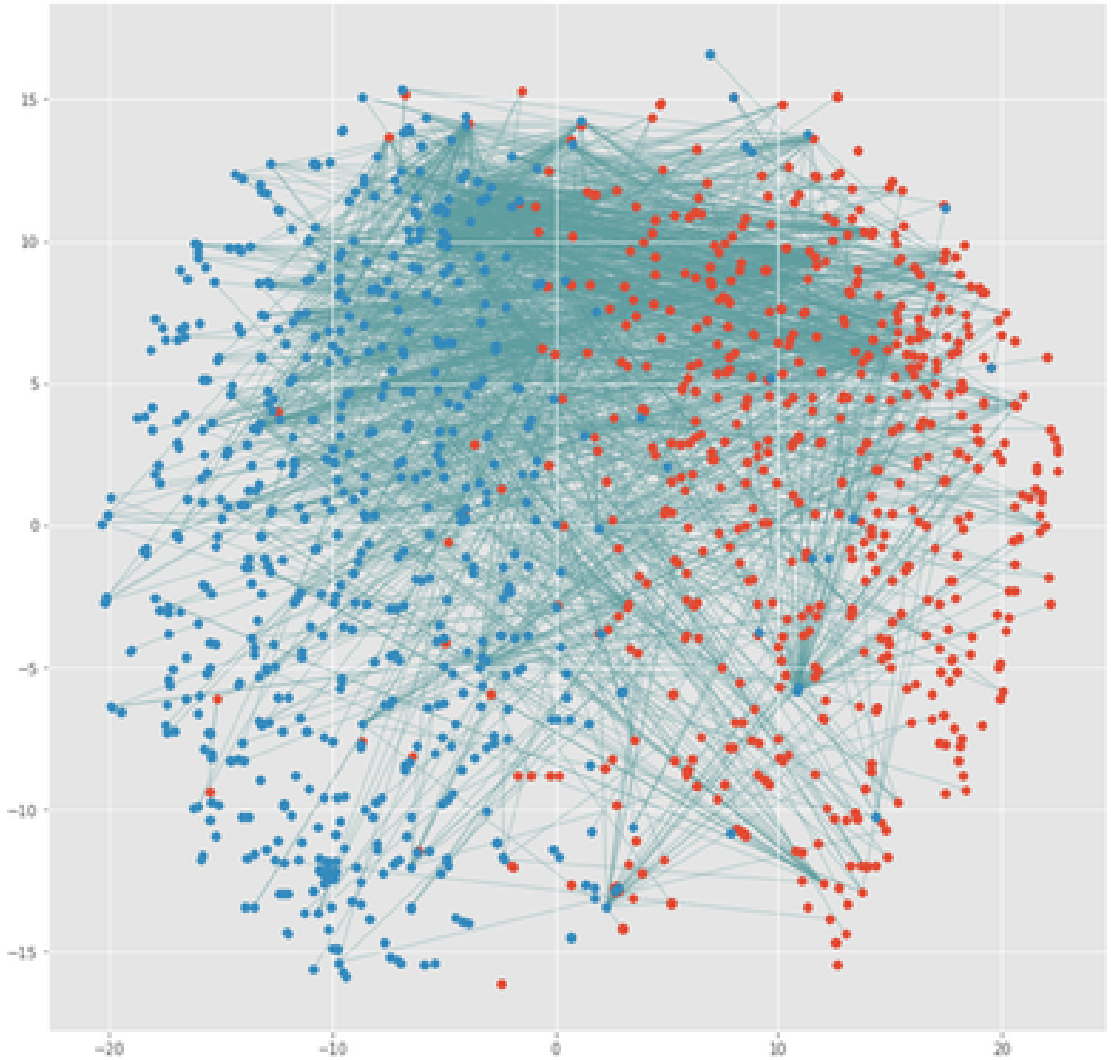

Graph Neural Networks: Three main steps

1. Agregation,
2. Decimation,
3. Pooling

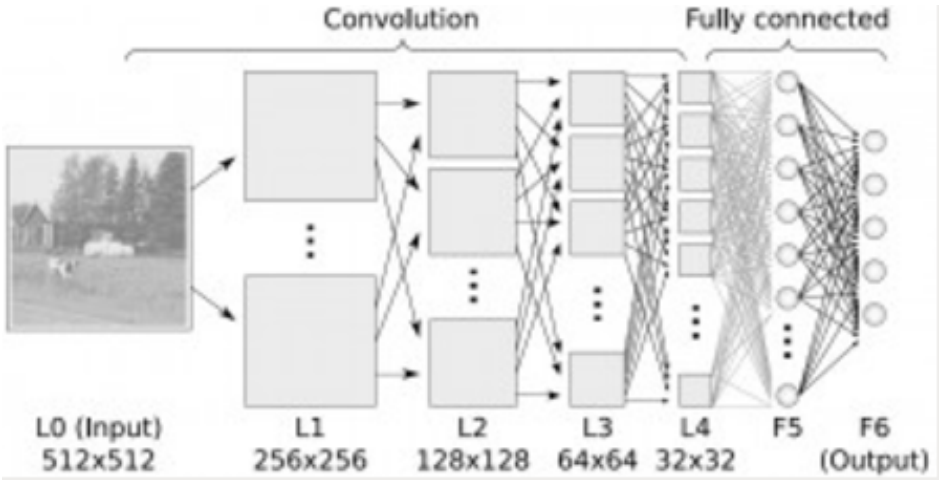

Agregation: The problem

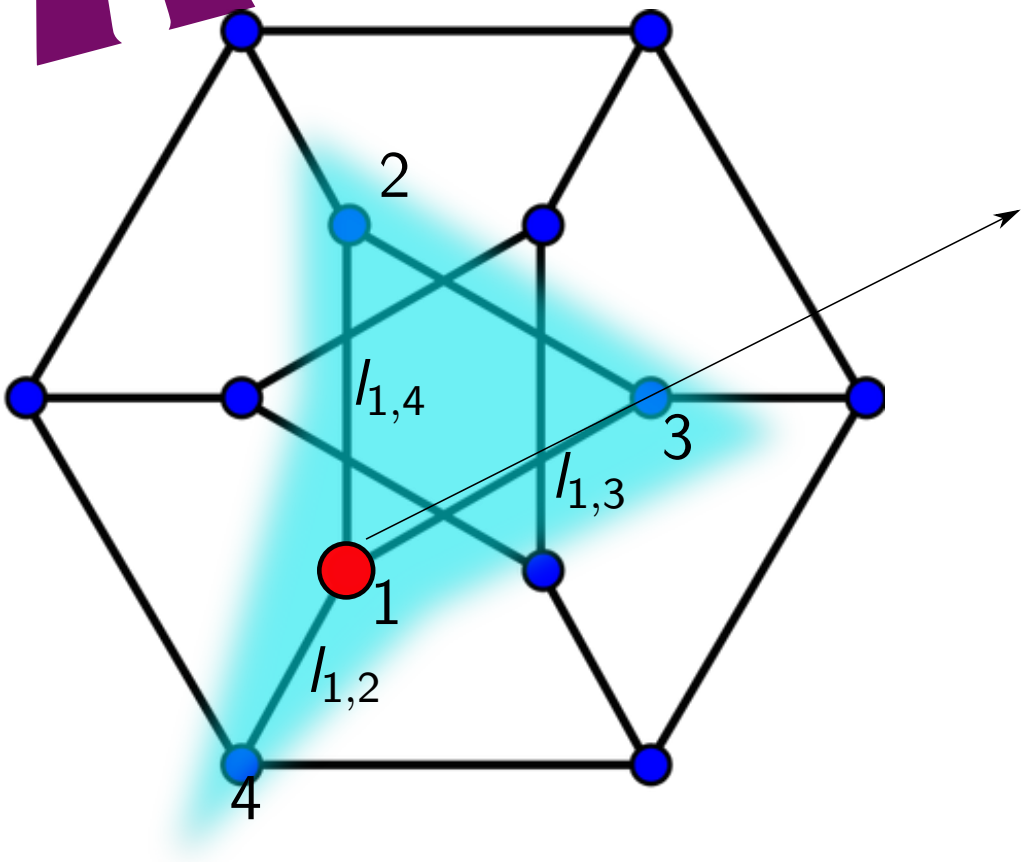

$$h_1 = f_w \quad (l_1, \{l_{1,2}, l_{1,3}, l_{1,4}\}, \{h_2, h_3, h_4\}, \{l_2, l_3, l_4\})$$

$$\begin{cases} h_v &= f_w(l_v, l_{CON(v)}, h_{\mathcal{N}(v)}, l_{\mathcal{N}(v)}) \\ o_v &= g_w(h_v, l_v) \end{cases}$$

with $CON(v) = \{(v, v') \; v' \in \mathcal{N}(v)\}$

Iteration

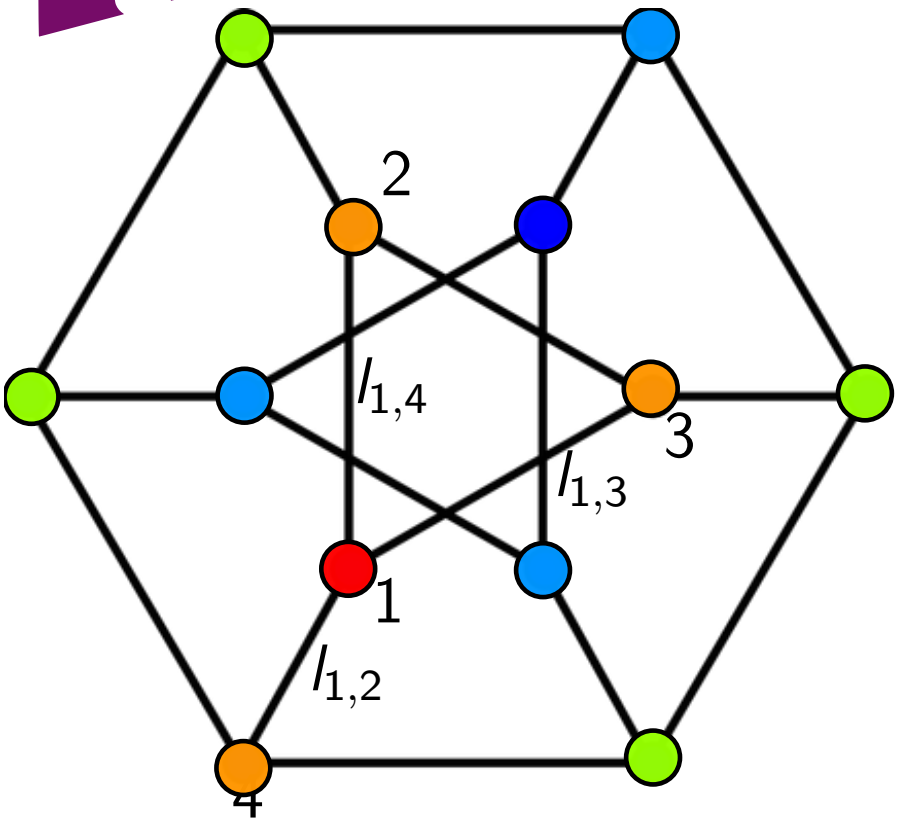

$t = 0$

$t = 1$

$t = 2$

$t = 3$

$t = 4$

$$\begin{cases} h_v^t &= f_w(l_v, l_{CON(v)}, h_{\mathcal{N}(v)}^{t-1}, l_{\mathcal{N}(v)}) \\ o_v &= g_w(h_v^T, l_v) \end{cases}$$

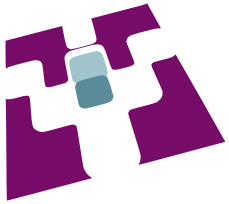

Graph Convolution

The problem

- Using images we learn $w_0 \ldots, w_8$:

w_5	w_1	w_6
w_3	w_0	w_4
w_7	w_2	w_8

  w_1 denotes the weigh of the pixel above the central pixel.

- Using graphs:

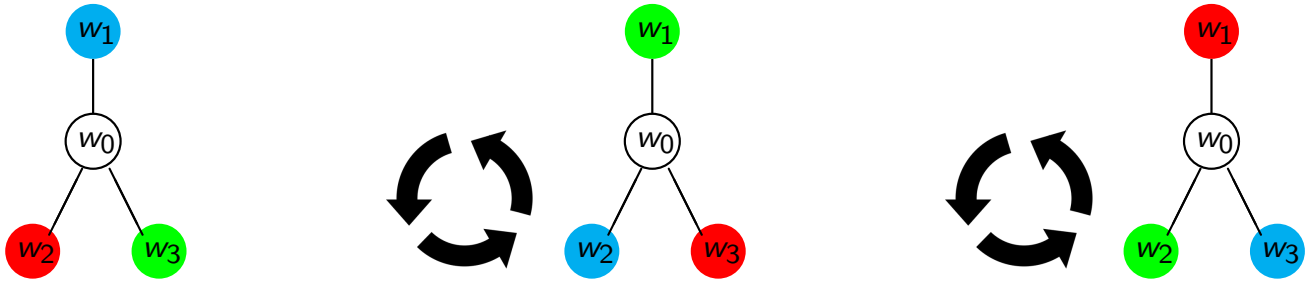

Without embedding nothing distinguishes the cyan,red and green neighbors.

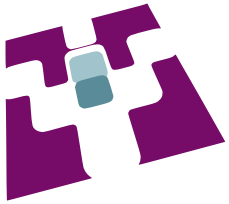

How to become permutation invariant

$$h_v^t = f_w(l_v, l_{CON(v)}, h_{\mathcal{N}(v)}^{t-1}, l_{\mathcal{N}(v)})$$

$$h_v^t = \sum_{v' \in \mathcal{N}(v)} f(l_v, l_{v,v'}, l_{v'}, h_{v'}^{(t-1)})$$

where f may be:

- An affine function [Scarselli et al., 2009],

$$f(l_v, l_{v,v'}, l_{v'}, h_{v'}^{(t-1)}) = A^{(l_v, l_{v,v'}, l_{v'})} h_{v'}^{(t-1)} + b^{(l_v, l_{v,v'}, l_{v'})}$$

- A MLP [Massa et al., 2006]

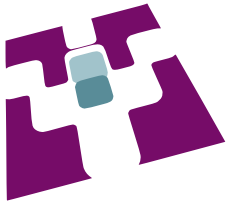

More complex agregation functions

- A long Short-term Memory [Hochreiter and Schmidhuber, 1997, Peng et al., 2017, Zayats and Ostendorf, 2018]
- A Gated Reccurent Unit [Li et al., 2016]

$$
\begin{aligned}
h_v^{(1)} &= [x_v^T, 0] && (1)\\
a_v^{(t)} &= A_v^T [h_1^{(t-1)\ T}, \ldots, h_{|V|}^{(t-1)\ T}]^T + b && (2)\\
z_v^t &= \sigma(W^z a_v^{(t)} + U^z h_v^{(t-1)}) && (3)\\
r_v^t &= \sigma(W^r a_v^{(t)} + U^r h_v^{(t-1)}) && (4)\\
\tilde{h}_v^{(t)} &= \tanh\left(W a_v^{(t)} + U\left(r_v^t \odot h_v^{(t-1)}\right)\right) && (5)\\
h_v^t &= (1 - z_v^t) \odot h_v^{(t-1)} + z_v^t \odot \tilde{h}_v^t && (6)
\end{aligned}
$$

z_v^t: update gate, r_v^t: reset gate, A_v: weight by edges types.

- Learned weight by edge type:
  $a_v^{(t)} = \sum_{w \in \mathcal{N}(v)} A_{l_{v,w}} h_w^{(t-1)}$ [Gilmer et al., 2017]

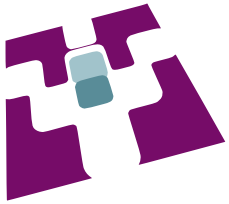

Graph attention Networks

[Velickovic et al., 2018]

- Not all neighbors have a same importance for update:

$$\alpha_{v,v'} = softmax_{v'}(e_{v,v'}) = \frac{exp(e_{v,v'})}{\sum_{v'' \in \mathcal{N}_i} exp(e_{v,v''})}$$

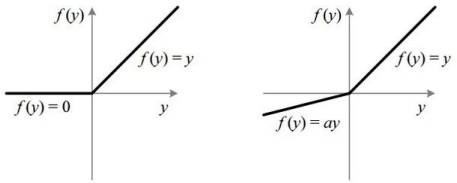

- With : $e_{v,v'} = LeakyReLU(a^T[Wh_v||Wh_{v'}])$
  a, W : weight vector and matrix.
- Update rule:

$$h'_v = \sigma(\sum_{v' \in \mathcal{N}_v} \alpha_{v,v'} W h_{v'})$$

- With K features:

$$h'_v = ||_{k=1}^{K} \sigma(\sum_{v' \in \mathcal{N}_v} \alpha_{v,v'}^k W^k h_{v'})$$

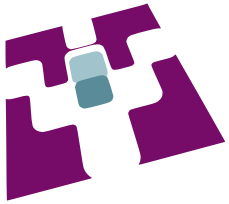

Graph Convolution

Input image

Convolution Kernel

$$\begin{bmatrix} -1 & -1 & -1 \\ -1 & 8 & -1 \\ -1 & -1 & -1 \end{bmatrix}$$

Feature map

The spectral approach [Defferrard et al., 2016]

- Graph Laplacian:

$$L = D - A \text{ with } D_{ii} = \sum_{j=1}^{n} A_{ij}$$

A adjacency matrix of a graph G.

- Matrix L is real symmetric semi definite positive:

$$L = U \Lambda U^T$$

U orthogonal, Λ real(positive) diagonal matrix.

- A classical result from signal processing:

$$x * y = \mathcal{F}^{-1}(\hat{x}.\hat{y})$$

*: convolution operation, $\mathcal{F}^{-1}$ inverse Fourrier transform, $\hat{x}$ fourrier transform of x, '.' term by term multiplication.

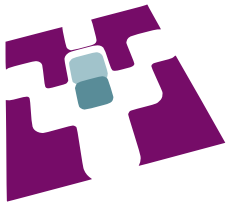

Graph Convolution

The spectral approach

- If x is a signal on G, $\hat{x} = U^T x$ can be considered as its "Fourrier" transform. We have:

$$U\hat{x} = UU^T x = x$$

  U is thus the inverse Fourrier transform.

- By analogy:

$$z * x = U(\hat{z} \odot \hat{x}) = U\left(U^T z \odot U^T x\right) = U\left(diag(U^T z)U^T x\right)$$

  $\odot$: Hadamard product.

- Let $g_\theta(\Lambda)$ be a diagonal matrix. The filtering of x by g_θ is:

$$y = U\left(g_\theta(\Lambda)U^T x\right) = \left(Ug_\theta(\Lambda)U^T\right)x$$

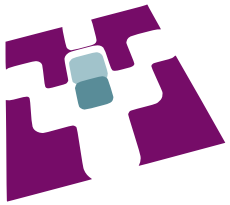

Graph convolution

The spectral approach

- If:

$$g_\theta(\Lambda) = \sum_{i=0}^{K-1} \theta_i \Lambda^i$$

Then:

$$y = \left(U g_\theta(\Lambda) U^T\right) x = U \left(\sum_{i=0}^{K-1} \theta_i \Lambda^i\right) U^T x = \left(\sum_{i=0}^{K-1} \theta_i L^i\right) x$$

- One parameter per ring:
  - Lx : one step (direct) neighborhood,
  - $L^2 x$: two step neighborhood (idem for $L^3, L^4, \dots$)
- Problem: Computing L^i for $i \in \{0, \dots, K-1\}$ is problematic for large matrices (SVD computation)

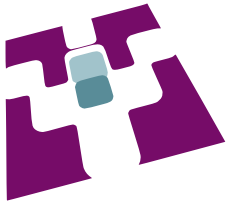

Graph convolution

The spectral approach

- Let us consider Chebyshev polynomial $T_k(x) = 2xT_{k-1}(x) - T_{k-2}(x)$, with $T_0 = 1$ and $T_1(x) = x$.

$$g_\theta(\Lambda) = \sum_{i=0}^{K-1} \theta_i \Lambda^i \rightarrow g_\theta(\Lambda) = \sum_{i=0}^{K-1} \theta_i T_i(\tilde{\Lambda})$$

  $\tilde{\Lambda}$ normalized version of Λ.

- we have:

$$\tilde{x}_k = 2\tilde{L}\tilde{x}_{k-1} - \tilde{x}_{k-2} \text{ with } \tilde{x}_0 = x \text{ and } \tilde{x}_1 = \tilde{L}x$$

  $\mathcal{O}(K|\mathcal{E}|)$ operations to get $\tilde{x}_k$.

- If $K = 2$ it simplifies to [Kipf and Welling, 2017]: $y = \theta L' x$ where L' is a regularized version of the normalized Laplacian.

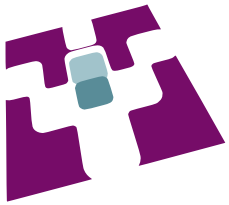

Graph convolution

Non spectral approaches

- [Simonovsky and Komodakis, 2017]

$$y_i = \frac{1}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} F_\theta(L(j,i))x_j + b$$

F: Parametric function of θ which associates one weigh to each edge label $L(j,i)$.

- [Verma et al., 2017]:

$$y_i = \frac{1}{|\mathcal{N}(i)|} \sum_{m=1}^{M} \sum_{j \in \mathcal{N}(i)} q_{\theta_m}(x_j, x_i) W_m x_j + b$$

$q_{\theta_m}(.,.)$ m^{th} learned soft-assignment function. W_m weight matrix.

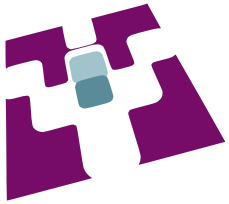

Graph Propagation

Recurent networks

[Hochreiter and Schmidhuber, 1997]

[Massa et al., 2006]

[Scarselli et al., 2009]

[Li et al., 2016]

[Gilmer et al., 2017]

[Peng et al., 2017]

[Zayats and Ostendorf, 2018]

Convolution

[Bruna et al., 2014]

[Defferrard et al., 2016]

[Kipf and Welling, 2017]

[Simonovsky and Komodakis, 2017]

[Verma et al., 2017]

Agregation

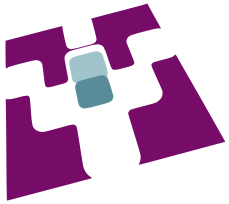

What's next ?

- Graph Downsampling, Graph pooling, Graph final decision: Some solutions but still the jungle.

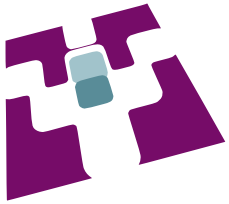

Conclusion

- Three graph metrics

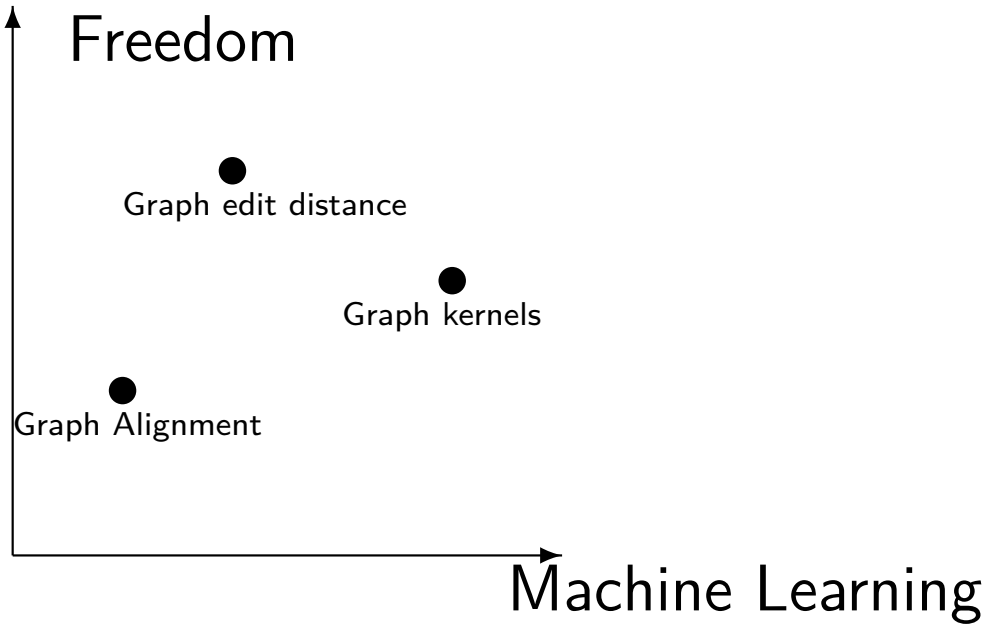

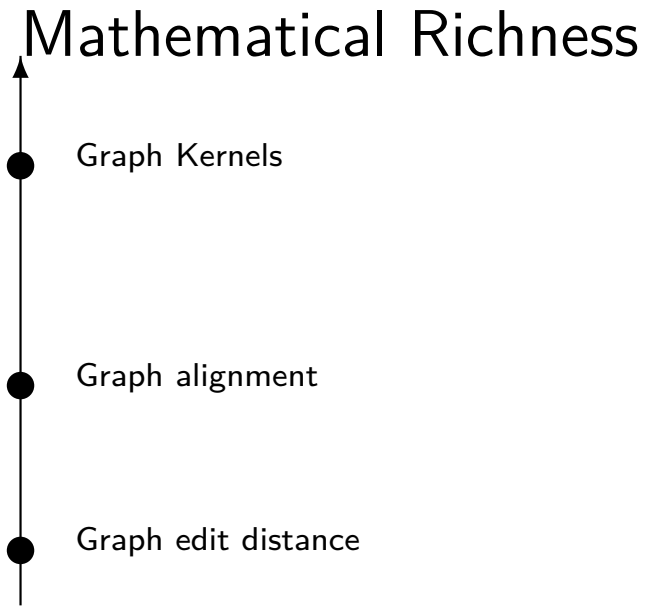

- Graph neural network:
  - Still in their infancy,
  - A great potential.

Bibliography

Abu-Aisheh, Z. (2016). *Anytime and distributed Approaches for Graph Matching*. PhD thesis, Université François Rabelais, Tours.

Ambauen, R., Fischer, S., and Bunke, H. (2003). Graph edit distance with node splitting and merging, and its application to diatom identification. In *Graph Based Representations in Pattern Recognition: 4th IAPR International Workshop, GbRPR 2003 York, UK, June 30 – July 2, 2003 Proceedings*, pages 95–106, Berlin, Heidelberg. Springer Berlin Heidelberg.

Aronszahn, N. (1950). Theory of reproducing kernels. *Transactions of the American Mathematical Society*, 68(3):337–404.

Blumenthal, D., Bougleux, S., Gamper, J., and Brun, L. (2018). Ring based approximation of graph edit distance. In *Proceedins of SSPR'2018*, Beijing. IAPR, Springer.

Boria, N., Bougleux, S., and Brun, L. (2018). Approximating ged using a stochastic generator and multistart ipfp. In Bai, X., Hancock, E. R., Ho, T. K., Wilson, R. C., Biggio, B., and Robles-Kelly, A., editors, *Proceedings of SSPR'2018*, pages 460–469. IAPR, Springer International Publishing.

Bougleux, S., Brun, L., Carletti, V., Foggia, P., Gaüzère, B., and Vento, M. (2017). Graph edit distance as a quadratic assignment problem. *Pattern Recognition Letters*, 87:38–46. Impact factor : 1.586.

Bruna, J., Zaremba, W., Szlam, A., and LeCun, Y. (2014). Spectral networks and locally connected networks on graphs. In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*.

Carletti, V., Gaüzère, B., Brun, L., and Vento, M. (2015). *Graph-Based Representations in Pattern Recognition: 10th IAPR-TC-15 International Workshop, GbRPR 2015, Beijing, China, May 13-15, 2015. Proceedings*, chapter Approximate Graph Edit Distance Computation Combining Bipartite Matching and Exact Neighborhood Substructure Distance, pages 188–197. Springer International Publishing, Cham.

Daller, E., Bougleux, S., Gaüzère, B., and Brun, L. (2018). Approximate graph edit distance by several local searches in parallel. In Fred, A., editor, *7th Internation Conference on Pattern Recognition Applications and Methods*.

Darwiche, M. (2018). *When Operations Research meets Structural Pattern Recognition: on the solution of Error-Tolerant Graph Matching Problems*. PhD thesis, University of Tours, tours, France.

Defferrard, M., Bresson, X., and Vandergheynst, P. (2016). Convolutional neural networks on graphs with fast localized spectral filtering. In Lee, D. D., Sugiyama, M., Luxburg, U. V., Guyon, I., and Garnett, R., editors, *Advances in Neural Information Processing Systems 29*, pages 3844–3852. Curran Associates, Inc.

Dupé, F. X. and Brun, L. (2009). Edition within a graph kernel framework for shape recognition. In *Graph Based Representation in Pattern Recognition 2009*, pages 11–21.

Frank, M. and Wolfe, P. (1956). An algorithm for quadratic programming. *Naval Research Logistics Quarterly*, 3(1-2):95–110.

Gärtner, T., Flach, P. A., and Wrobel, S. (2003). On graph kernels: Hardness results and efficient alternatives. In *Computational Learning Theory and Kernel Machines, 16th Annual Conference on Computational Learning Theory and 7th Kernel Workshop, COLT/Kernel 2003, Washington, DC, USA, August 24-27, 2003, Proceedings*, pages 129–143.

Gaüzère, B., Bougleux, S., Riesen, K., and Brun, L. (2014a). *Structural, Syntactic, and Statistical Pattern Recognition: Joint IAPR International Workshop, S+SSPR 2014, Joensuu, Finland, August 20-22, 2014. Proceedings*, chapter Approximate Graph Edit Distance Guided by Bipartite Matching of Bags of Walks, pages 73–82. Springer Berlin Heidelberg, Berlin, Heidelberg.

Gaüzère, B., Grenier, P.-A., Brun, L., and Villemin, D. (2014b). Treelet kernel incorporating cyclic, stereo and inter pattern information in Chemoinformatics. *Pattern Recognition*, page 30 p.

Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., and Dahl, G. E. (2017). Neural message passing for quantum chemistry. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, pages 1263–1272.

Haussler, D. (1999). Convolution kernels on discrete structures.

Hermansson, L., Johansson, F. D., and Watanabe, O. (2015). Generalized shortest path kernel on graphs. In Japkowicz, N. and Matwin, S., editors, *Discovery Science*, pages 78–85, Cham. Springer International Publishing.

Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8):1735–1780.

Jain, B. J. (2014). Margin perceptrons for graphs. In *22nd International Conference on Pattern Recognition, ICPR 2014, Stockholm, Sweden, August 24-28, 2014*, pages 3851–3856.

Jain, B. J. (2016). Statistical graph space analysis. *Pattern Recognition*, 60:802 – 812.

Jain, B. J. and Wysotzki, F. (2004). Central clustering of attributed graphs. *Machine Learning*, 56(1-3):169–207.

Kashima, H., Tsuda, K., and Inokuchi, A. (2003). Marginalized kernels between labeled graphs. In *Machine Learning, Proceedings of the Twentieth International Conference (ICML 2003), August 21-24, 2003, Washington, DC, USA*, pages 321–328.

Kipf, T. N. and Welling, M. (2017). Semi-supervised classification with graph convolutional networks. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*.

Lerouge, J., Abu-Aisheh, Z., Raveaux, R., Héroux, P., and Adam, S. (2017). New binary linear programming formulation to compute the graph edit distance. *Pattern Recognition*, 72:254 – 265.

Li, Y., Tarlow, D., Brockschmidt, M., and Zemel, R. S. (2016). Gated graph sequence neural networks. In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*.

Liu, Z. and Qiao, H. (2014). GNCCP - graduated nonconvexityand concavity procedure. *IEEE Trans. Pattern Anal. Mach. Intell.*, 36(6):1258–1267.

Loosli, G., Canu, S., and Ong, C. S. (2016). Learning SVM in kreĭn spaces. *IEEE Trans. Pattern Anal. Mach. Intell.*, 38(6):1204–1216.

Mahé, P. and Vert, J. (2009). Graph kernels based on tree patterns for molecules. *Machine Learning*, 75(1):3–35.

Massa, V. D., Monfardini, G., Sarti, L., Scarselli, F., Maggini, M., and Gori, M. (2006). A comparison between recursive neural networks and graph neural networks. In *Proceedings of the International Joint Conference on Neural Networks, IJCNN 2006, part of the IEEE World Congress on Computational Intelligence, WCCI 2006, Vancouver, BC, Canada, 16-21 July 2006*, pages 778–785.

Peng, N., Poon, H., Quirk, C., Toutanova, K., and Yih, W. (2017). Cross-sentence n-ary relation extraction with graph lstms. *TACL*, 5:101–115.

Ralaivola, L., Swamidass, S. J., Saigo, H., and Baldi, P. (2005). Graph kernels for chemical informatics. *Neural Networks*, 18(8):1093–1110.

Riesen, K. and Bunke, H. (2009). Approximate graph edit distance computation by means of bipartite graph matching. *Image and Vision Computing*, 27(7):950 – 959. 7th IAPR-TC15 Workshop on Graph-based Representations (GbR 2007).

Scarselli, F., Gori, M., Tsoi, A. C., Hagenbuchner, M., and Monfardini, G. (2009). The graph neural network model. *IEEE Trans. Neural Networks*, 20(1):61–80.

Shervashidze, N. and Borgwardt, K. M. (2009). Fast subtree kernels on graphs. In *Advances in Neural Information Processing Systems 22: 23rd Annual Conference on Neural Information Processing Systems 2009. Proceedings of a meeting held 7-10 December 2009, Vancouver, British Columbia, Canada.*, pages 1660–1668.

Shervashidze, N., Vishwanathan, S. V. N., Petri, T., Mehlhorn, K., and Borgwardt, K. M. (2009). Efficient graphlet kernels for large graph comparison. In *Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics, AISTATS 2009, Clearwater Beach, Florida, USA, April 16-18, 2009*, pages 488–495.

Simonovsky, M. and Komodakis, N. (2017). Dynamic edge-conditioned filters in convolutional neural networks on graphs. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 29–38.

Velickovic, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., and Bengio, Y. (2018). Graph attention networks. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*.

Verma, N., Boyer, E., and Verbeek, J. (2017). Dynamic filters in graph convolutional networks. *CoRR*, abs/1706.05206.

Zayats, V. and Ostendorf, M. (2018). Conversation modeling on reddit using a graph-structured LSTM. *TACL*, 6:121–132.