Expectation Complete Graph Representations Using Graph Homomorphisms

Maximilian Thiessen* [1] Pascal Welke* [2] Thomas Gärtner [1]

[1]TU Wien [2]University of Bonn

TL;DR

Through the power of random features we devise efficiently computable and expectation complete graph embeddings.

Expressiveness

Graph representation methods are compared to each other in terms of expressiveness. That is, their (theoretical) ability to compute different representations for pairs of non-isomorphic graphs.

For example, MPNNs are at most as expressive as the 1-WL isomorphism test.

High expressiveness is necessary for learning: If your method cannot distinguish two graphs, it cannot learn a function that behaves differently on these graphs.

Completeness

$\mathcal{G}$ the set of all graphs, V a vector space (e.g., $\mathbb{R}^d$)

A graph embedding $\varphi : \mathcal{G} \to V$ is permutation-invariant if for all isomorphic graphs

$$G \simeq H : \varphi(G) = \varphi(H)$$

A permutation-invariant graph embedding φ is complete if for all non-isomorphic graphs

$$G \not\simeq H : \varphi(G) \neq \varphi(H)$$

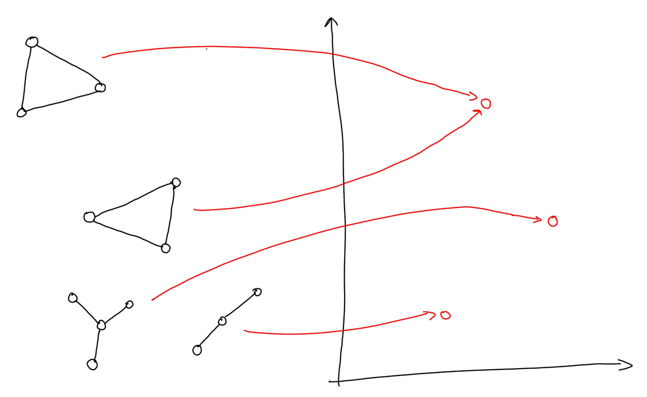

Originated from complete graph kernels [Gärtner et al., COLT 2003]

Problem

Why do we care about complete graph embeddings?

Allow us to learn/approximate any permutation-invariant function!

Unfortunately computing any such embedding is at least as hard as deciding graph isomorphism

- not known to be NP-hard and not known to be computable in polynomial-time

Typical solution: drop completeness for efficiency

- most practical graph kernels, GNNs, Weisfeiler Leman test, k-WL test, ...

Our solution: keep completeness in expectation!

Complete in Expectation

Let $\varphi_X : \mathcal{G} \to V$ depend on a random variable X drawn from a distribution $\mathcal{D}$ over a set $\mathcal{X}$

We call φ_X complete in expectation if the expectation

$$\mathbb{E}_{X\sim\mathcal{D}}[\varphi_X(\cdot)] = \sum_{t\in\mathcal{X}} \Pr(X = t)\varphi_t(\cdot)$$

is a complete graph embedding

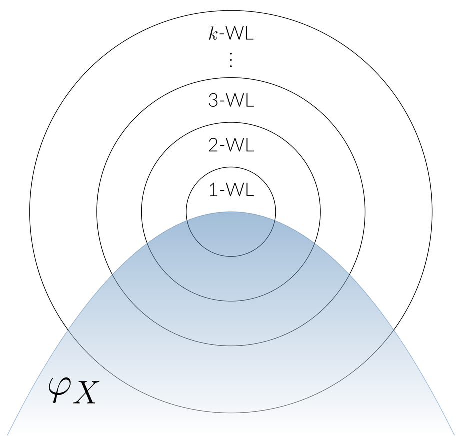

What is the benefit?

Sampling $X_1, X_2, X_3, \dots$ will eventually make the joint embedding $(\varphi_{X_1}(G), \varphi_{X_2}(G), \varphi_{X_3}(G), \dots)$ arbitrarily expressive

Our Approach: Sampling from the Lovász Vector

Let $\mathcal{G}_n$ be the set of all graphs with at most n vertices.

- the parameter n is typically the size of the largest graph in the sample.

Theorem. Let $\mathcal{D}$ be a distribution with full support on $\mathcal{G}_n$ and $G \in \mathcal{G}_n$. The graph embedding

$$\varphi_F(G) = \hom(F, G)e_F$$

with $F \sim \mathcal{D}$ is complete in expectation.

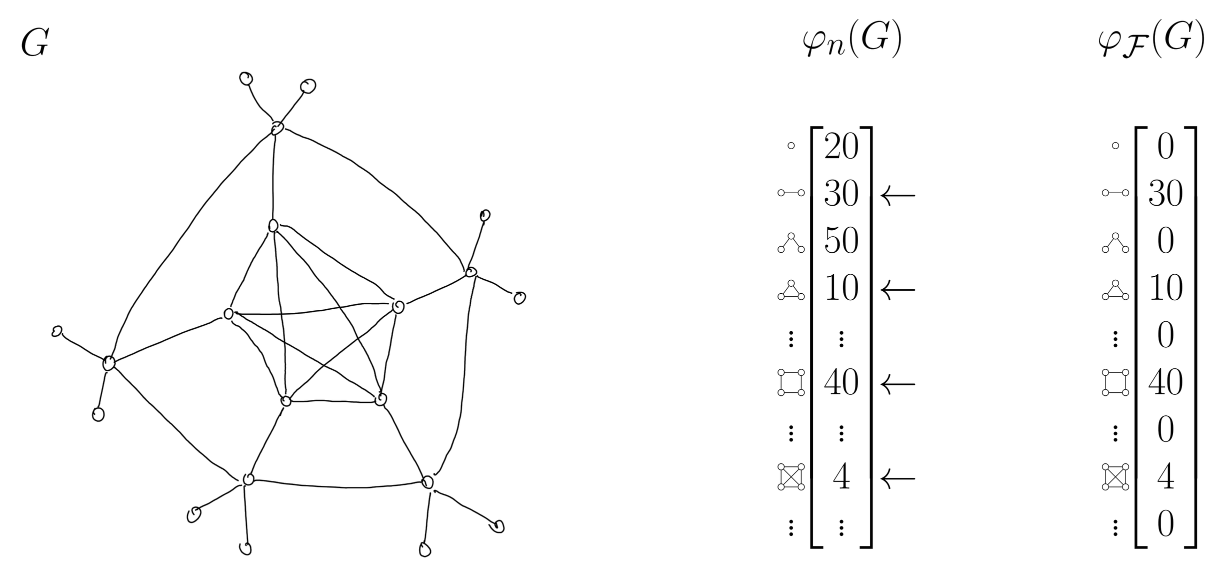

Proposed embedding: sample multiple pattern graphs F

- draw a finite sample $\mathcal{F}$ i.i.d from $\mathcal{D}$ and represent any graph $G \in \mathcal{G}_n$ by

$$\varphi_{\mathcal{F}}(G) = \sum_{F\in\mathcal{F}} \varphi_F(G)$$

- reduces the variance of the embedding
- currently $\ell = |\mathcal{F}|$ is a fixed hyperparameter (e.g., $\ell = 30$)

Efficient Sampling Scheme

Computing $\hom(F, G)$ is NP-hard in general.

If we take the treewidth of pattern F into account the runtime is [Díaz et al., 2002]:

$$\mathcal{O}\left(|V(F)||V(G)|^{\mathrm{tw}(F)+1}\right)$$

Idea: define distribution $\mathcal{D}$ on $\mathcal{G}_n$ s.t. runtime is polynomial in expectation!

Theorem. There exists a distribution $\mathcal{D}$ such that computing the expectation complete graph embedding $\varphi_F(G)$ takes polynomial time in $|V(G)|$ in expectation for all $G \in \mathcal{G}_n$.

General recipe:

1. pick n as the maximum number of vertices in the training set
2. sample treewidth upper bound k
3. sample a maximal graph F' with treewidth k
4. take a random subgraph F of F'

E.g., $k \sim \mathrm{Poi}(\lambda)$ with $\lambda \leq \frac{1+d\log n}{n}$ guarantees runtime $\mathcal{O}\left(|V(G)|^{d+2}\right)$

Homomorphisms

Let F, G be graphs. A map $\varphi : V(F) \to V(G)$ is a graph homomorphism if φ preserves edges:

$$\{v, w\} \in E(F) \text{ implies } \{\varphi(v), \varphi(w)\} \in E(G) .$$

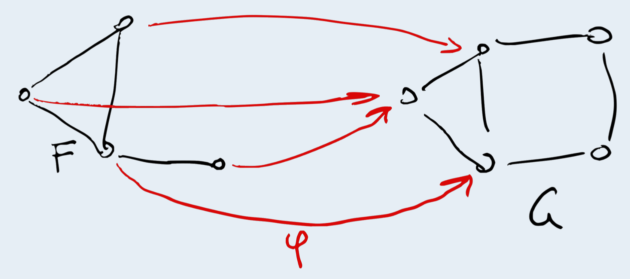

φ does not have to be injective (!)

$\hom(F, G)$: number of homomorphisms from F to G.

The Lovász Vector

Let $\varphi_n(G) = \hom(\mathcal{G}_n, G) = (\hom(F, G))_{F\in\mathcal{G}_n}$ denote the Lovász vector of G for $\mathcal{G}_n$.

Theorem [Lovász, 1968]. Two arbitrary graphs $G, H \in \mathcal{G}_n$ are isomorphic iff $\varphi_n(G) = \varphi_n(H)$.

That means that $\varphi_n(\cdot)$ is complete!

Properties of Homomorphism Counts

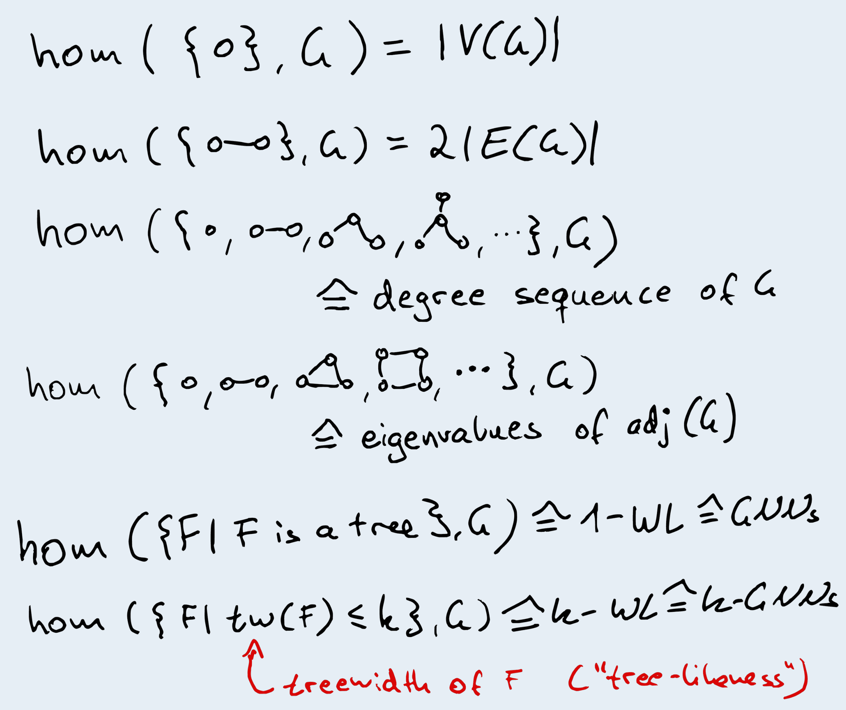

Counting subgraphs [Curticapean et al., STOC 2017]

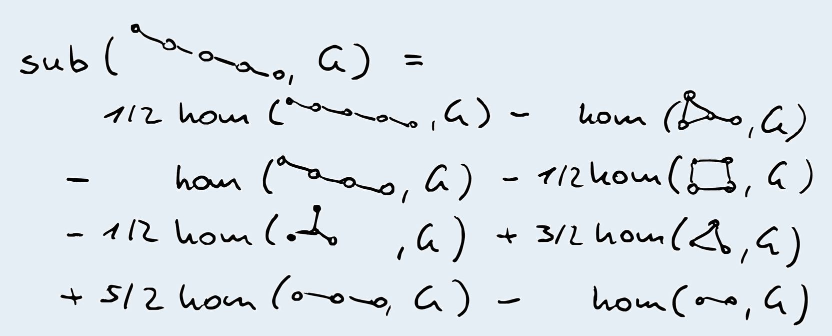

Universality [NT and Maehara, ICML 2020]: Any permutation-invariant function

$$f : \mathcal{G} \to \mathbb{R}^d$$

can be approximated arbitrarily well by a polynomial of

$$\{\hom(F, G) \mid F \in \mathcal{G}\}$$

Working on Arbitrary Graph Sizes

If we cannot restrict the size of graphs at inference time, we can define a kernel on $\mathcal{G}_\infty$ without restricting to $\mathcal{G}_n$ for some $n \in \mathbb{N}$. We define the countable-dimensional vector

$$\overline{\varphi}_\infty(G) = \left(\hom_{|V(G)|}(F, G)\right)_{F\in\mathcal{G}_\infty}$$

where

$$\hom_{|V(G)|}(F, G) = \begin{cases} \hom(F, G) & \text{if } |V(F)| \leq |V(G)| , \\ 0 & \text{if } |V(F)| > |V(G)| . \end{cases}$$

That is, $\overline{\varphi}_\infty(G)$ is the projection of $\varphi_\infty(G)$ to the subspace that gives us the homomorphism counts for all graphs of *size at most of* G. Note that this is a well-defined map of graphs to a subspace of the ℓ^2 space, i.e., sequences $(x_i)_i$ over $\mathbb{R}$ with $\sum_i |x_i|^2 < \infty$.

Theorem. $\overline{\varphi}_\infty$ is complete.

Theorem. $\overline{\varphi}_X$ is complete in expectation.

The map $\overline{\varphi}_\infty$ even maps all graphs into an inner product space and allows to compute norms or distances, and to apply kernel methods.

Empirical Results

Our method with $\ell = 30$ sampled patterns and the $\overline{\varphi}_\infty$ embedding

Deterministic embeddings as baseline [NT and Maehara, ICML 2020]

- GHC-tree(6): all tree patterns up to size 6
- GHC-cycle(8): all cycle patterns up to size 8

Additionally:

- graph neural tangent kernel (GNTK) [Du et al., NeurIPS 2019]
- GIN [Xu et al., ICLR 2019]

Table 1. Cross-validation accuracies on benchmark datasets

method	MUTAG	IMDB-BIN	IMDB-MULTI	PAULUS25	CSL
GHC-tree(6)	89.28 ± 8.26	72.10 ± 2.62	48.60 ± 4.40	7.14 ± 0.00	10.00 ± 0.00
GHC-cycle(8)	87.81 ± 7.46	70.93 ± 4.54	47.41 ± 3.67	7.14 ± 0.00	100.00 ± 0.00
GNTK	89.46 ± 7.03	75.61 ± 3.98	51.91 ± 3.56	7.14 ± 0.00	10.00 ± 0.00
GIN	89.40 ± 5.60	70.70 ± 1.10	43.20 ± 2.00	7.14 ± 0.00	10.00 ± 0.00
WL-kernel	90.4 ± 5.7	73.12 ± 0.4	-	7.14 ± 0.00	10.00 ± 0.00
ours (SVM)	87.94 ± 0.01	70.37 ± 0.01	47.34 ± 0.01	100.00 ± 0.00	37.33 ± 0.1
ours (MLP)	88.55 ± 0.01	70.81 ± 0.01	48.29 ± 0.01	40.524 ± 0.01	13.27 ± 0.01

Relations to k-WL and k-GNNs

Theorem. Let $\mathcal{D}$ be a distribution with full support on the set of graphs with treewidth up to k. The resulting graph embedding $\varphi_F^{k\text{-WL}}(\cdot)$ with $F \sim \mathcal{D}$ has the same expressiveness as the k-WL test in expectation. Furthermore, there is a specific such distribution such that we can compute $\varphi_F^{k\text{-WL}}(G)$ in expected polynomial time $\mathcal{O}(|V(G)|^{k+1})$ for all $G \in \mathcal{G}_\infty$.

Future Work

Choose number of patterns ℓ and distribution $\mathcal{D}$ adaptively:

- stop sampling when expressive enough
- pick $\mathcal{D}$ based on the task or a given dataset
- frequent / interesting patterns

Going beyond expressiveness: similarity!

- if $G \approx H$ then $\varphi(G) \approx \varphi(H)$
- possible solution: cut distance (captures local and global properties)