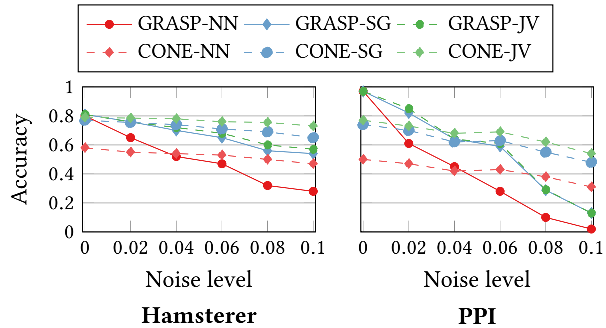

Figure 6: Accuracy with different assignment algorithms.

6.2 Impact of degree distribution

The graphs used in the experiments approximately follow a power law distribution. We estimate the power law exponents using the method in [18]. Table 2 lists the results.

Graph	*Power Law Exponent (γ)*
Arenas	1.56
Facebook	1.32
Hamsterster, PPI	1.45
Voles	1.64
MultiMagna	1.46
HighSchool	1.36

Table 2: Estimated power law exponents.

We study the impact of the power-law exponent on alignment accuracy. Figure 7 shows the performance for three synthetic power-law graphs vs. noise. GRASP achieves nearly 100% accuracy with low exponent, which explains the good performance in the previous experiments. GRASP falls short on graphs with highly skewed distribution, while CONE exhibits a more consistent performance.

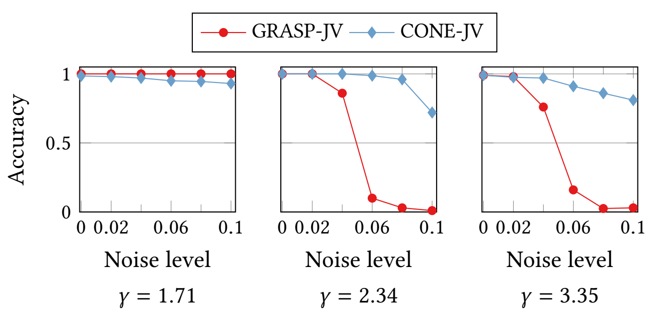

Figure 7: Accuracy of GRASP-PPR and CONE with JV assignment on generated power law graphs.

6.3 Real World Networks

So far we have analyzed behavior on synthetic noise. Here we consider networks in which noisy modifications occur naturally:

- **Voles** [5]. A time-evolving wildlife contact network. The full graph is matched to snapshots at time steps where 80%, 85%, 90%, and 99% of edges have been added.
- **HighSchool** [8]. A time-evolving graph of contact patterns among high school students. The full graph is matched to snapshots with 80%, 85%, 90% and 99% of edges added.
- **MultiMagna** [28]. Different versions of the same protein-protein-interaction network.

Figure 8 shows results on these real world networks. CONE-Align consistently outperforms GRASP-PRR with ICP and Voting.

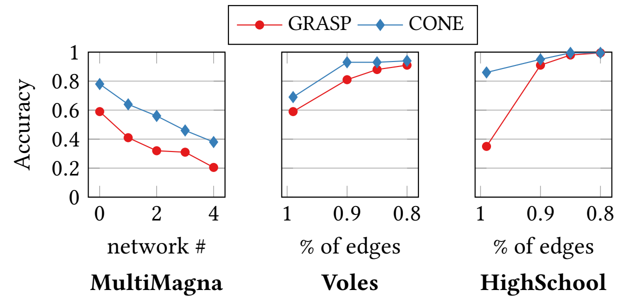

Figure 8: Accuracy on real data sets.

6.4 Efficiency

We now compare the efficiency of best-of-breed versions, namely GRASP-PPR with ICP and voting and CONE with NetMF embeddings, across Assign variants. Figure 9 shows the runtime for each steps. For GRASP, we report precomputation including Embed and Align, and the time for voting. For CONE-Align, we report Embed, Align and Assign, separately. While CONE-Align is faster than GRASP-PPR on small graphs, its Align is slower than BA and ICP; therefore, GRASP-PPR outperforms CONE-Align on large enough graphs, especially when using nearest neighbor matching.

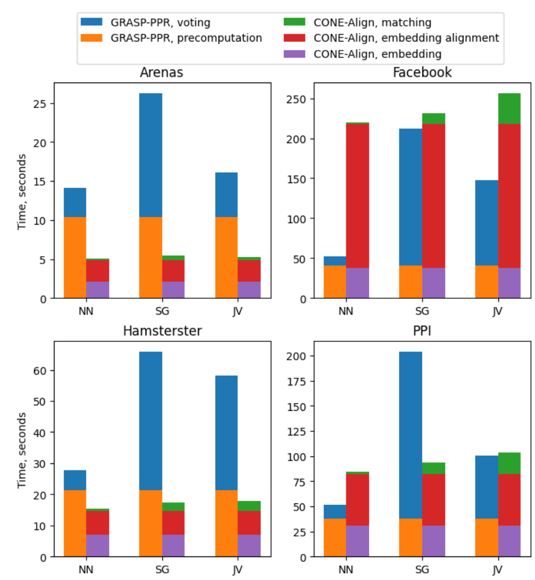

Figure 9: Time with GRASP-PPR (ICP, voting) and CONE with three matching methods on four datasets.

7 CONCLUSION

We revisited state-of-the-art graph alignment algorithms and identified an overarching modular framework that allows for components exchange. By virtue of this framework, these algorithms are amenable to several enhancements. We offered and analyzed alternatives that improve performance vs. that of default variants, as our experiments with synthetic and real noise corroborate. Our results pave the way to further advances in graph alignment.

REFERENCES

[1] Ahmet E Aladağ and Cesim Erten. 2013. SPINAL: scalable protein interaction network alignment. *Bioinformatics* 29, 7 (2013), 917–924.

[2] Jon Louis Bentley. 1975. Multidimensional binary search trees used for associative searching. *CACM* 18, 9 (1975), 509–517.

[3] Paul J Besl and Neil D McKay. 1992. Method for registration of 3-D shapes. In *Sensor fusion IV: control paradigms and data structures*, Vol. 1611. International Society for Optics and Photonics, 586–606.

[4] Xiyuan Chen, Mark Heimann, Fatemeh Vahedian, and Danai Koutra. 2020. CONE-Align: Consistent Network Alignment with Proximity-Preserving Node Embedding. In *CIKM*. 1985–1988.

[5] Stephen Davis, Babak Abbasi, Shrupa Shah, Sandra Telfer, and Mike Begon. 2015. Spatial analyses of wildlife contact networks. *Journal of the Royal Society, Interface* 12, 102 (2015).

[6] Katerina Doka, Mingqiang Xue, Dimitrios Tsoumakos, and Panagiotis Karras. 2015. k-Anonymization by freeform generalization. In *AsiaCCS*. 519–530.

[7] Soheil Feizi, Gerald Quon, Mariana Mendoza, Muriel Medard, Manolis Kellis, and Ali Jadbabaie. 2019. Spectral alignment of graphs. *IEEE Trans. Netw. Sci.* (2019).

[8] Julie Fournet and Alain Barrat. 2014. Contact patterns among high school students. *PloS one* (2014).

[9] Mark Heimann, Haoming Shen, Tara Safavi, and Danai Koutra. 2018. REGAL: Representation Learning-based Graph Alignment. In *CIKM*. 117–126.

[10] Judith Hermanns, Anton Tsitsulin, Marina Munkhoeva, Alexander Bronstein, Davide Mottin, and Panagiotis Karras. 2021. GRASP: Graph Alignment through Spectral Signatures. In *APWeb-WAIM*.

[11] Roy Jonker and Anton Volgenant. 1987. A shortest augmenting path algorithm for dense and sparse linear assignment problems. *Computing* 38, 4 (1987), 325–340.

[12] Gunnar W. Klau. 2009. A new graph-based method for pairwise global network alignment. *BMC Bioinf.* 10, S-1 (2009).

[13] Danai Koutra, Hanghang Tong, and David Lubensky. 2013. BIG-ALIGN: Fast Bipartite Graph Alignment. In *ICDM*. 389–398.

[14] Jérôme Kunegis. 2013. Konect: the Koblenz network collection. In *WWW*. ACM, 1343–1350.

[15] Jure Leskovec and Andrej Krevl. 2014. SNAP Datasets: Stanford Large Network Dataset Collection. http://snap.stanford.edu/data.

[16] Chung-Shou Liao, Kanghao Lu, Michael Baym, Rohit Singh, and Bonnie Berger. 2009. IsoRankN: spectral methods for global alignment of multiple protein networks. *Bioinformatics* 25, 12 (2009), i253–i258.

[17] Huda Nassar, Nate Veldt, Shahin Mohammadi, Ananth Grama, and David F. Gleich. 2018. Low Rank Spectral Network Alignment. In *WWW*. 619–628.

[18] Mark EJ Newman. 2005. Power laws, Pareto distributions and Zipf's law. *Contemporary physics* 46, 5 (2005), 323–351.

[19] Sadegh Nobari, Panagiotis Karras, HweeHwa Pang, and Stéphane Bressan. 2014. L-opacity: Linkage-Aware Graph Anonymization. In *EDBT*. 583–594.

[20] Maks Ovsjanikov, Mirela Ben-Chen, Justin Solomon, Adrian Butscher, and Leonidas J. Guibas. 2012. Functional maps: a flexible representation of maps between shapes. *ACM Trans. Graph.* 31, 4 (2012), 30:1–30:11.

[21] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. 1999. *The PageRank citation ranking: Bringing order to the web.* Technical Report. Stanford InfoLab.

[22] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. 2014. Deepwalk: Online learning of social representations. In *KDD*. 701–710.

[23] Jiezhong Qiu, Yuxiao Dong, Hao Ma, Jian Li, Kuansan Wang, and Jie Tang. 2018. Network embedding as matrix factorization: Unifying DeepWalk, LINE, PTE, and node2vec. In *WSDM*. 459–467.

[24] Leonardo FR Ribeiro, Pedro HP Saverese, and Daniel R Figueiredo. 2017. struc2vec: Learning node representations from structural identity. In *KDD*. 385–394.

[25] Rohit Singh, Jinbo Xu, and Bonnie Berger. 2008. Global alignment of multiple protein interaction networks with application to functional orthology detection. *PNAS* 105, 35 (2008), 12763–12768.

[26] Anton Tsitsulin, Davide Mottin, Panagiotis Karras, and Emmanuel Müller. 2018. VERSE: Versatile graph embeddings from similarity measures. In *TheWebConf*. 539–548.

[27] Anton Tsitsulin, Marina Munkhoeva, Davide Mottin, Panagiotis Karras, Ivan V. Oseledets, and Emmanuel Müller. 2021. FREDE: Anytime Graph Embeddings. *Proc. VLDB Endow.* 14, 6 (2021), 1102–1110.

[28] Vipin Vijayan and Tijana Milenković. 2017. Multiple network alignment via multi-MAGNA++. *IEEE/ACM Transactions on Computational Biology and Bioinformatics* 15, 5 (2017), 1669–1682.

[29] Mingqiang Xue, Panagiotis Karras, Chedy Raïssi, Panos Kalnis, and Hung Keng Pung. 2012. Delineating social network data anonymization via random edge perturbation. In *CIKM*. 475–484.

[30] Si Zhang and Hanghang Tong. 2016. FINAL: Fast Attributed Network Alignment. In *KDD*. 1345–1354.