Exact p -values for global network alignments via combinatorial analysis of shared GO terms

REFANGO: Rigorous Evaluation of Functional Alignments of Networks using Gene Ontology

Wayne B. Hayes

Received: date / Accepted: date

Abstract Network alignment aims to uncover topologically similar regions in the protein-protein interaction (PPI) networks of two or more species under the assumption that topologically similar regions perform similar functions. Although there exist a plethora of both network alignment algorithms and measures of topological similarity, currently no “gold standard” exists for evaluating how well either is able to uncover functionally similar regions. Here we propose a formal, mathematically and statistically rigorous method for evaluating the statistical significance of shared GO terms in a global, 1-to-1 alignment between two PPI networks. We use combinatorics to precisely count the number of possible network alignments in which k proteins share a particular GO term. When divided by the number of all possible network alignments, this provides an explicit, exact p -value for a network alignment with respect to a particular GO term.

Keywords Network alignment · Gene Ontology · GO terms

1 Introduction

Network alignment aims to uncover similar network connection patterns between two or more networks under the assumption that common network topology (which may be easily observable) correlates with common function (which is more difficult to observe).

In this paper we focus on the simplest case to analyze: a pairwise global network alignment (PGNA). Pairwise, because exactly two networks are being aligned; and global, because we insist that every node in the smaller network

---

W. Hayes
Dept. of Computer Science, UC Irvine
Tel.: +1-949-824-1753
Fax: +1-949-824-4056
E-mail: whayes@uci.edu

(ie., the one with fewer nodes) is mapped to exactly one node in the larger network. Extending the analysis to the non-1-to-1 case—as well as to local and multiple network alignments—appears nontrivial and is left for future work.

Network alignment algorithms abound and their number is increasing rapidly. There is an urgent need for a *rigorous* way to evaluate these methods for their ability to identify functionally similar proteins across species. Unfortunately, there are almost as many ways of evaluating network alignments as there are alignment algorithms. The research area is severely hampered by the lack of any “gold standard” for evaluating the functional similarity that is uncovered by any given network alignment. One of the most common methods for evaluating biological significance involves using the Gene Ontology’s GO term hierarchy. While other biologically relevant information exists, such as KEGG pathways (any others?), in this paper we focus on leveraging GO terms to produce a mathematically and statistically rigorous “gold standard” measure of the statistical significance of an alignment in which exactly k aligned protein pairs all share a common GO term g .

2 Method: GO-term p -values by exhaustive enumeration of alignments

2.1 Network alignment and functional similarity

Given two networks G_1, G_2 , let the node sets V_1, V_2 represent n_1 and n_2 proteins respectively, and the edge sets E_1, E_2 represent protein-protein interactions (PPIs). Assuming (without loss of generality) that $n_1 \leq n_2$, a pairwise global network alignment (PGNA) is a 1-to-1 mapping $f : V_1 \rightarrow V_2$ in which every node in V_1 is mapped to exactly one node in V_2 .

Once an alignment is specified, we need to measure the extent to which functionally similar proteins have been mapped to each other. Many existing methods evaluate their alignments using GO terms from the Gene Ontology [1], and most often evaluate the functional similarity of each pair of aligned proteins independent of all the others, and then average across the pairs. While the score of each pair may be meaningful, taking an average assumes that each pair is independent of all the others. This is not true because the pairings themselves are inter-dependent via the alignment itself, which is built globally. For example, in a 1-to-1 alignment, each node in each network can appear at most once across the entire alignment, a property which destroys the independence assumption needed for a meaningful average across aligned node pairs.

There are other considerations as well: most GO terms have inter-dependencies with most other GO terms via the GO hierarchy [2]; most proteins have more than one GO annotation, and it is difficult to create a measure that correctly evaluates similarity between two proteins with different sets of GO terms that only partially overlap; and finally, since most GO terms annotate many proteins, it is nontrivial to compute the significance of aligning a set of protein

pairs while accounting for both the frequency and inter-relationships between GO terms that may appear in multiple pairs across the set of aligned pairs.

Existing solutions to evaluate the relationships between GO terms, as well as between the gene products that they annotate, are currently all *ad hoc*[2]. Specifically in the context of network alignment, to our knowledge all of the existing methods are also *ad hoc*, and lack any strong mathematical or statistical rigor. (The details of these other methods are irrelevant and entirely unrelated to our work, and so are relegated to the Appendix, though we offer a discussion of their pitfalls in Section §4.) If we are to make progress in the network alignment community, we need a way to compare network alignments for the extent to which they uncover functional similarity among its aligned proteins. Given the plethora of *ad hoc* methods currently being used, there is an urgent need for a “gold standard” measure of the significance of functional similarity uncovered by a network alignment. Here we propose such a gold standard, which is both mathematically and statistically rigorous.

2.2 Computing the total number of possible alignments

In the following exposition, we must discuss in great detail the combinatoric structure of a given alignment. To aid visualization, we use what I call the “Pegs and Holes” analogy: given networks G_1, G_2 with n_1, n_2 nodes, we imagine G_2 ’s nodes as n_2 identical “holes” drilled into a large board, and G_1 ’s nodes as n_1 identical “pegs” that can each fit into any hole. To enforce the global 1-to-1 property, there are two cases:

1. $n_1 \leq n_2$, so every peg is placed into some hole, leaving $n_2 - n_1$ empty holes. There are $\binom{n_2}{n_1}$ ways to choose which holes to use, and $n_1!$ ways to place the pegs.
2. $n_1 > n_2$, so every hole is filled with some peg, leaving $n_1 - n_2$ pegs unplaced. There are $\binom{n_1}{n_2}$ ways to choose which pegs to place, and $n_2!$ ways to place them.

The above two cases are symmetric and so, without loss of generality, we assume $n_1 \leq n_2$. Then, the total number of all possible alignments is

$$\binom{n_2}{n_1} n_1! = \frac{n_2!}{(n_2 - n_1)!} \equiv P(n_2, n_1). \quad (1)$$

The function $P(\cdot, \cdot)$ of Eq. (1) is more commonly known as *k-permutations-of-n*, or $P(n, k)$. However, $P(n, k)$ is usually defined to be zero if $n < k$, whereas we will often need to compute the number of alignments when we don’t know which of the two values is larger. Thus, in this paper, we will adopt a modified permutation function $\pi(n_1, n_2)$ as follows

$$\pi(n_1, n_2) = \begin{cases} P(n_1, n_2), & \text{if } n_1 \geq n_2, \\ P(n_2, n_1), & \text{if } n_2 > n_1. \end{cases} \quad (2)$$

2.3 Counting alignments with exactly k matches

Given a particular GO term g , assume g annotates λ_1 pegs and λ_2 holes. A peg and the hole it sits in are, more technically, a pair of aligned nodes. We say that such a pair “match” with respect to GO term g if they are both annotated with g . Let $\underline{\lambda} = \min(\lambda_1, \lambda_2)$, and $\bar{\lambda} = \max(\lambda_1, \lambda_2)$. Given a random 1-to-1 alignment, we are going to compute the probability p that exactly k pairs of aligned nodes share g . In our analogy, this means that exactly k pegs—no more, no less—that are annotated with g sit in holes that are also annotated with g . To do this, we will use a combinatorial argument to enumerate all possible PGNAs that can exist that have exactly k matches. Given that number, we simply divide by Eq. (1) to get the probability that a randomly chosen alignment has exactly k matches.

2.3.1 Special cases

The following are special cases:

1. if $k > \underline{\lambda}$, then $p = 0$.
2. if $\underline{\lambda} = 0$, then $p = 1$ if $k = 0$ and $p = 0$ otherwise.
3. if $\lambda_2 = n_2$, then $p = 1$ if $k = \lambda_1$, and $p = 0$ otherwise.
4. if $\lambda_1 > n_2 - \lambda_2$ and $k < \lambda_1 - (n_2 - \lambda_2)$, then $p = 0$, otherwise $p > 0$ is computed below.

The last case arises when $\lambda_1 > n_2 - \lambda_2$, which means that there are more annotated pegs than non-annotated holes, necessitating that *at least* $\lambda_1 - (n_2 - \lambda_2)$ annotated pegs must align with annotated holes. (Recall we are computing the probability of *exactly* k aligned pairs sharing g , so k too small in this case gives $p = 0$.)

Below we describe the general case in detail. In broad outline, there are three steps: (i) create the required k matches by placing k annotated pegs into k annotated holes; (ii) arrange to place the remaining annotated pegs away from the annotated holes in order to keep k constant; (iii) place any remaining pegs (all of which are non-annotated) in any still-empty holes (some of which may be annotated). In each case we either sum, or multiply, as appropriate, the number of ways to perform the described action. In the end we have counted all the possible ways to create an alignment that has exactly k matches.

2.3.2 Creating exactly k matches

Out of the λ_1 pegs annotated with g , pick $k \leq \underline{\lambda}$ of them; there are $\binom{\lambda_1}{k}$ ways to do this. We will place these k pegs into k holes that are also annotated with g ; there are $\binom{\lambda_2}{k}$ ways to pick the holes, and $k!$ ways to place the k pegs into the k holes. Thus, the total number of ways to match exactly k pairs of nodes that share g is

$$M_k(\lambda_1, \lambda_2) = \binom{\lambda_1}{k} \binom{\lambda_2}{k} k!. \quad (3)$$

From this point onward, in order to keep k constant, we are committed to creating no more matches.

2.3.3 Enumerating the ways to use the remaining annotated holes

To ensure that no more node pairs are matched, we need to ensure that none of the remaining $(\lambda_1 - k)$ annotated pegs are placed into any of the remaining $(\lambda_2 - k)$ annotated holes. Thus, each annotated hole must either remain empty, or take an non-annotated peg. There are $n_1 - \lambda_1$ available non-annotated pegs, regardless of the value of k . Pick μ of them. Since these μ pegs are all non-annotated, they can go into any unoccupied annotated hole without changing k . However, there are lower and upper bounds on what μ can be, as follows:

- μ can be at most $\bar{\mu} \equiv \min(n_1 - \lambda_1, \lambda_2 - k)$, since $n_1 - \lambda_1$ is the total number of non-annotated pegs, and $\lambda_2 - k$ is the number of available annotated holes in which to place (some of) them.
- note that we have $n_1 - k$ pegs (of both types) remaining to place, and exactly $n_2 - \lambda_2$ non-annotated holes, into which some (or all) of the pegs can be placed. By the pigeon hole principle, if $(n_1 - k) > (n_2 - \lambda_2)$, then some of the pegs—and they can only be non-annotated pegs—*must* go into annotated holes. Thus, μ —which refers only to non-annotated pegs—must be at least $\underline{\mu} \equiv (n_1 - k) - (n_2 - \lambda_2)$ if $(n_1 - k) > (n_2 - \lambda_2)$; otherwise $\underline{\mu} = 0$.

2.3.4 Distributing the remaining pegs

For any $\underline{\mu} \leq \mu \leq \bar{\mu}$, we need to count how many alignments can be built when μ non-annotated pegs are placed into the $\lambda_2 - k$ available annotated holes, as well as what happens to all the remaining pegs. The process is as follows.

1. There are $\binom{n_1 - \lambda_1}{\mu}$ ways to choose μ non-annotated pegs, and $\pi(\lambda_2 - k, \mu)$ ways to align them with the open annotated holes. To simplify notation note that $n_1, n_2, \lambda_1, \lambda_2$ are all fixed; thus, let $\gamma_k(\mu) = \binom{n_1 - \lambda_1}{\mu} \pi(\lambda_2 - k, \mu)$.
2. Recall that there are still $\lambda_1 - k$ annotated pegs to be placed, and that they must be placed into non-annotated holes, so we must “reserve” $\lambda_1 - k$ non-annotated holes, which will be further accounted for below.
3. Once μ annotated holes are filled with non-annotated pegs, the rest of the annotated holes must remain empty; this leaves $n_1 - \lambda_1 - \mu$ non-annotated pegs to go into the $n_2 - \lambda_2$ non-annotated holes. Keeping in mind the “reservation” above, there are $n_2 - \lambda_2 - (\lambda_1 - k)$ available non-annotated holes. There are $\binom{n_2 - \lambda_2}{\lambda_1 - k}$ ways to choose which holes to use while reserving $\lambda_1 - k$ of them, and $\pi(n_1 - \lambda_1 - \mu, n_2 - \lambda_2 - (\lambda_1 - k))$ ways to place the pegs into the chosen holes; let $\delta_k(\mu) = \binom{n_2 - \lambda_2}{\lambda_1 - k} \pi(n_1 - \lambda_1 - \mu, n_2 - \lambda_2 - (\lambda_1 - k))$.
4. Finally, we place the remaining $\lambda_1 - k$ annotated pegs into the reserved holes of the same number; there are $(\lambda_1 - k)!$ ways to do this.

2.3.5 Summing the unmatched region of the alignment

Combining all of the above for fixed μ and then summing over all possible μ , the total number of ways that $n_1 - \lambda_1$ non-annotated pegs can be used to (partially or wholly) fill $\lambda_2 - k$ annotated holes, and then use all the remaining pegs and holes in a manner consistent with keeping k constant, is

$$U_k(\lambda_1, \lambda_2) \equiv (\lambda_1 - k)! \sum_{\mu=\mu}^{\bar{\mu}} \gamma_k(\mu) \delta_k(\mu). \quad (4)$$

2.3.6 Final tally for exactly k matches

Combining Eq.s (3) and (4), the total number of alignments in which exactly k aligned node pairs share GO term g is

$$C_k(\lambda_1, \lambda_2) \equiv M_k(\lambda_1, \lambda_2) U_k(\lambda_1, \lambda_2). \quad (5)$$

2.4 The probability of an alignment with exactly k matches

Eq. (5) counts all possible alignments in which exactly k aligned node pairs share GO term g . To get the probability p_k of the same event, we divide by Eq. (1):

$$p_k^g(n_1, n_2, \lambda_1^g, \lambda_2^g) = \frac{C_k^g(\lambda_1^g, \lambda_2^g)}{\pi(n_1, n_2)}, \quad (6)$$

where a superscript g has been added as appropriate to denote that this probability is specifically tied to GO term g .

Note this refers to *exactly* k matches. To measure the statistical significance of m matches, we sum Eq. (6) for k from m to λ^g .

2.5 Efficiently dealing with huge numbers

Though technically it is only an implementation detail, it is important to briefly discuss how to deal with the astronomically huge numbers involved in these calculations. Typical modern biological networks can have thousands to tens of thousands of nodes, and some GO terms annotate thousands of genes in each network. For example, in BioGRID 3.4.164 that we use below, the two biggest PPI networks in terms of number of nodes are *H. sapiens* and *A. thaliana*, which contain exactly 17,200 and 9,364 unique proteins, respectively, that are involved in physical interactions. Eq. (1) in this case is approximately 10^{38270} —an integer with over 38,000 digits in base-10, which is far above the values typically representable on modern hardware. Luckily, its logarithm is easy to represent in double precision floating point, and so all of the multiplications herein can be computed as the floating-point sum of logarithms. The sole complication is the summation in Eq. (4), which is a sum of *values*, not

logarithms. We use the following trick. Given two numbers a and b , assume we have at our disposal only their logarithms, $\alpha = \log(a)$ and $\beta = \log(b)$. Our goal is to estimate $\log(a + b)$. Without loss of generality, assume $a \leq b$. Then,

$$\log(a + b) = \beta + \log(1 + a/b) \quad (7)$$

$$= \beta + \log(1 + e^{\alpha-\beta}) \quad (8)$$

$$= \beta + f(e^{\alpha-\beta}), \quad (9)$$

where $f(x)$ is some function that can provide an accurate estimate of $\log(1+x)$ for any $|x| \leq 1$. One must be careful because if $|x|$ is below the machine epsilon ($\approx 10^{-16}$ in double precision), then $1+x$ evaluates to 1 because x is rounded away, and a direct evaluation of the expression $\log(1+x)$ gives zero. The solution is not hard: the built-in library function for log can evaluate $\log(1+x)$ with sufficient accuracy if $|x| > 10^{-6}$; for smaller values of $|x|$, we explicitly invoke the Taylor series, which is extremely accurate for small values of $|x|$. We have tested that this method gives values for $\log(a+b)$ that are accurate to almost machine precision for any $|x| \leq 1$.

3 Results

3.1 Numerical Validation

Staring at $C_k(\lambda_1, \lambda_2)$ in Eq. (5) and tracing back through the equations that define its components, it is not immediately obvious that the $C_k(\lambda_1, \lambda_2)$, when summed over all possible values of k , must add up to exactly $\pi(n_1, n_2)$ independent of the choice of λ_1, λ_2 . Yet if Eq. (5) is correct, then this must be the case since summing p_k in Eq. (6) across all k of must give exactly 1.

In the calculation of p_k^g in Eq. (6), the values of k and g are fixed. For a fixed g , valid values of k range from zero to $\underline{\lambda}^g$. If our calculations are correct, then the sum across k of p_k^g should be exactly 1 for any fixed $g, n_1, n_2, \lambda_1, \lambda_2$. We tested this property in the following cases:

1. exhaustively for all $0 \leq \lambda_1 \leq n_1$ and $0 \leq \lambda_2 \leq n_2$ for all $0 \leq n_1 \leq n_2 \leq 100$;
2. as above but in steps of 10 in λ_i and n_i up to $n_2 = 1,000$;
3. as above but in powers of 2 in λ_i and n_i up to $n_2 = 32,768$;
4. several billion random quadruples of $(n_1, n_2, \lambda_1, \lambda_2)$ with n_2 chosen uniformly at random up to 100,000, n_1 chosen uniformly at random up to n_2 , and the λ 's chosen uniformly at random up to their n value.

We found in all cases that the difference from 1 of the sum over k of p_k^g was bounded by 10^{-9} . (Keep in mind that we had access only to the logarithms of the C_k ; that the actual sum across k had to be approximated term-by-term using Eq. (9); that the correct answer in log space is $\log(1) = 0$; and that all operations were performing in floating point, which incurs roundoff error.) Furthermore, in any particular case, the numerical (floating-point roundoff) error will be dominated by the sum over μ in Eq. (4), and so we would expect

Table 1 8 largest networks of BioGRID 3.4.164, sorted by node count.

nodes	common name	official name
17200	Human	<i>H. sapiens</i>
9364	Thale cress	<i>A. thaliana</i>
8728	Fruit fly	<i>D. melanogaster</i>
6777	Mouse	<i>M. musculus</i>
5984	Baker's yeast	<i>S. cerevisiae</i>
3194	Worm	<i>C. elegans</i>
2811	Fission yeast	<i>S. pombe</i>
2391	Rat	<i>R. norvegicus</i>

the error to be smaller (ie., sum closer to 1) when there are fewer terms in Eq. (4). The number of terms is well-approximated by $\min(n_1 - \lambda_1, n_2)$. Indeed, we find that if the sum was S , then the value $|S - 1|/\min(n_1 - \lambda_1, n_2)$ has mean $\approx 3 \times 10^{-14}$, standard deviation $\approx 3 \times 10^{-13}$, and was never observed to exceed 3×10^{-12} .

3.2 Validation against random alignments of real PPI networks

We downloaded the 8 largest protein-protein interaction networks from release 3.4.164 of BioGRID (cf. Table 1), and the GO database release of the same month. As many authors of network alignment papers do, we then split the GO database into two versions: one with all GO terms, and ones where sequence-based GO terms were disallowed. For each of the $\binom{8}{2} = 28$ pairs of networks and for both versions of the GO database, we generated 400 million random alignments, for a total of 22.4 billion random alignments. For each GO term g , we observed the integer frequency ϕ_k^g that g was shared by exactly k proteins when it annotated λ_1^g out of n_1 proteins in network G_1 and λ_2^g proteins out of n_2 in network G_2 . (Note that formally ϕ_k^g it has six parameters, $\phi_k^g(n_1, n_2, \lambda_1^g, \lambda_2^g)$, though we often abbreviate it to ϕ_k^g or even just ϕ_k or ϕ if context is clear.) It is a non-negative integer bounded by the number of random alignments, $N = 4 \times 10^8$, and dividing it by N gives an estimate of the probability that a randomly chosen alignment between G_1 and G_2 will contain exactly k aligned protein pairs that share g .

The estimated (ie., observed) probability ϕ_k^g/N can be compared to p_k^g of Eq. (6). Across the 22.4 billion random alignments, we observed 428,849 unique combinations of the six parameters $g, k, n_1, n_2, \lambda_1^g, \lambda_2^g$ that formally define ϕ_k^g . Figure 1 is a scatter plot of ϕ_k^g/N for all 428,849 of them, versus the theoretical value from Eq. (6). The agreement is excellent.

The scatter in Figure 1 increases towards the low end because events with probability near N^{-1} are rarely observed, and so the estimate of their probability contains significant sampling noise. In fact there is “width” to the scatter plot at all values of probability, but it is difficult to observe in Figure 1. To more clearly see the scatter, we compute the *ratio* of the observed to theoretical values of probability, which will have an expected value of 1 if Eq. (6) is an accurate and unbiased estimator of probability. Figure 2 plots the mean

Fig. 1 Scatter plot of the observed ϕ_k/N vs. theoretical p_k probability across 22.4 billion random alignments between pairs of networks from BioGRID 3.4.164. The vertical axis depicts the observed probability of an event, which is the observed frequency $\phi_k^g(n_1, n_2, \lambda_1, \lambda_2)$ divided by the number of samples $N = 4 \times 10^8$. The horizontal axis is the value given by Eq. (6) for the parameters of the observation. There are 428,849 observations plotted across all observed values of $n_1, n_2, \lambda_1^g, \lambda_2^g, k$.

and standard deviation (binned in powers of 2 of the number of samples) of $|1 - (\phi_k^g/N)/p_k^g|$ across all 428,849 observed frequencies, as a function of the number of samples that gave rise to the probability estimate. We can clearly see that the ratio approaches 1 asymptotically with the square root of the number of samples, consistent with sampling noise in ϕ .

3.3 Comparison with a simpler Poisson model

We introduce a Poisson-based model that correctly iterates across GO terms rather than protein pairs, though it is simplistic and only provides an approxi-

Fig. 2 Same data as Figure 1, except that, for each point, we have computed the distance D from 1 of the ratio of observed to predicted probability: $D = |1 - \frac{\phi_k^g/N}{p_k^g}|$. Each observed frequency ϕ_k^g (which we will henceforth abbreviate as ϕ) is converted to an observed probability ϕ/N , where N is the number of random alignments (4×10^8) per pair of networks. However, ϕ is also the number of samples used to create the observed probability estimate; higher ϕ gives a better estimate of the probability. We binned ϕ in powers of 2 (ie. the bin is $\lfloor \log_2(\phi) \rfloor$), and for each bin plotted the mean and standard deviation of D . We see that as the number of samples increases, the ratio approaches 1 as the square root of the number of samples, consistent with sampling noise.

mate p -value. Given that GO term g annotates λ_i out of n_i proteins in network G_i , then a randomly chosen protein u_i from network G_i has probability λ_i/n_i of being annotated by g . Thus, when a pair of proteins (u_1, u_2) is independently sampled from all possible pairs of proteins in $V_1 \times V_2$, the probability that they both share g is $\frac{\lambda_1}{n_1} \frac{\lambda_2}{n_2}$; note that, at this stage, this *not* an approximation—the probability is exact. Since multiple independent Poisson processes have a cumulative rate which is simply the sum of their individual rates, if we choose m

Fig. 3 Ratio of the Poisson-based p -value of Eq. (10) to the exact p -value of Eq. (6), across 562,000 random samples from *good* alignments between networks of BioGRID 3.4.164 [3]. As we can see, the p -value returned by the Poisson model gets progressively worse (underestimating statistical significance) as the λ values grow. Note that some points had ratios as high as 10^{70} , though we truncated the vertical axis at 10^5 . Each point has been perturbed in a random direction by a distance distributed as $N(0, 10^{-1})$ in log space, otherwise thousands of points land at the same integer co-ordinates, making it impossible to visualize the density of points across the plane.

such pairs of nodes, *each independently of all the others*, then the number of pairs that share g is modeled by a Poisson distribution with mean $\Lambda = m \frac{\lambda_1}{n_1} \frac{\lambda_2}{n_2}$, and the probability that k such pairs share g is

$$e^{-\Lambda} \frac{\Lambda^k}{k!}. \quad (10)$$

While this distribution correctly models the case where each protein pair is chosen independently and uniformly at random from all the others, it is only an approximation for the distribution of protein pairs that share g in a 1-to-1 network alignment, because the set of node pairs in an alignment are *not*

independent: each pair depends implicitly on all the others via the alignment itself, which is built globally and disallows any one node from being used more than once.

In relation to the GO term frequencies, Eq. (10) is a good approximation when $\lambda_1 \ll n_1$ and $\lambda_2 \ll n_2$, because then the probabilities are small and each pair that shares g has only a small influence on others. However, the approximation gets worse as either of λ_1 or λ_2 increases. To demonstrate this, we took an assortment of *good* alignments between the 3.4.164 BioGRID networks [3] which had some astronomically small p -values. Figure 3 plots the ratio of the Poisson-based p -value of Eq. (10) to the correct one of Eq. (6), as a function of $\lambda_1 + \lambda_2$. As we can see, the ratio is rarely less than 1 (ie., the Poisson-based p -value is almost always greater than or equal to the exact one), but can be huge if $\lambda_1 + \lambda_2$ is large—meaning the Poisson model grossly underestimates the statistical significance of matching a large number of pairs that share g .

4 Discussion

4.1 The fallacy of evaluating an alignment as an average across pairs of aligned nodes

There is a crucially important case that is often implicitly ignored in the literature by methods that evaluate GO-based significance of alignments by evaluating all aligned protein pairs, rather than evaluating all GO terms. This case is alluded to by phrases such as “consider the GO terms shared by a pair of aligned proteins...”. The problem is when there is a GO term g that exists in both networks (ie., λ_1^g and λ_2^g are both nonzero), but no pair of aligned proteins share it. Then the “consider...” phrase above implicitly misses the fact that g *could* have been shared by some aligned protein pairs, but was not. Unless taken care of explicitly, the alignment evaluation fails to penalize the alignment for failing to provide any matches for GO term g . In contrast, our method is correctly penalized for such cases: any GO term g for which $k = 0$ but both λ_1 and λ_2 are nonzero receives the appropriate penalty of a p -value with little statistical significance. Unfortunately, since many existing publications ignore this case, many published p -values claim far more statistical significance than actually exists.

A second major problem with existing *ad hoc* measures is that they do not scale even remotely monotonically with statistical significance. Take the Jaccard similarity, which is the most popular according to Table 2, though it has variously been called *GO Correctness* or *Consistency* (GOC), as well as *Functional Correctness/Consistency* (FC). Formally, given node $u \in V_1$ aligned to $v \in V_2$, let S_u, S_v be the set of GO terms annotating u, v , respectively. Then the Jaccard/GOC/FC between u and v is defined as

$$\text{FC}(u, v) \equiv \text{GOC}(u, v) \equiv \text{Jaccard}(u, v) \equiv \frac{|S_u \cap S_v|}{|S_u \cup S_v|}. \quad (11)$$

Given this similarity across all aligned pairs of proteins, the entire alignment is given an FC value equal to the mean across all aligned pairs.

It is easy to construct a scenario to demonstrate how badly the Jacard/GOC/FC measure can lead one astray. Consider the following simple system: network G has $n = 1000$ nodes. Each node is annotated with one, and only one, GO term. The first 900 nodes are annotated with GO term g_0 —ie., $\lambda^{g_0} = 900$. We will refer to these as the “common” nodes. The remaining 100 nodes are each individually annotated with their own unique GO term, with names $\{g_1, g_2, \dots, g_{99}, g_{100}\}$; thus, $\lambda^{g_i} = 1$ for all $i = 1, \dots, 100$. We will refer to these as the “specific” nodes. For simplicity, we will align G to itself, and assume that all 101 of the GO terms are *independent*, so that the p -value of the entire alignment is the product of the p -values across 101 GO terms.¹ Then, every pair of aligned nodes constitutes a *cluster*, and the only possible per-cluster FC scores are 0 and 1, so that the mean alignment-wide FC score is simply the fraction of node pairs that are *matched*, using the formal definition of “match” from Section §2.3.

In a random alignment of G to itself, each common node has a 90% chance of being aligned with another common node, so that the expected number of matched common nodes is 810; evaluating Eq. (6) we find that $p_{810}(1000, 1000, 900, 900) = 0.139$ —not statistically significant, as expected. On the other hand, each specific node has only a 0.1% chance of being aligned with its one and only match, so that in a random alignment we expect *none* (or very few) of the specific nodes to match. For this example, assume we do a bit better on the common nodes and match 820 of them, but match none of the specific nodes. The p -value of matching 820 common nodes is 0.0007. The 100 unmatched specific nodes each have $p_0(1000, 1000, 1, 1) = 0.999$, and $0.999^{100} = 0.90$. All told, this alignment has FC score of 0.82—making it look very good—and a p -value of about 0.0006.

Now consider a second alignment with the same FC score: we will correctly match just 10 of the specific nodes, and assume the other 90 are aligned with common nodes. This leaves precisely 810 common nodes to match each other, so the FC score is $(810+10)/1000 = 0.82$, as above. The p -value of 810 matched common nodes is above—0.139. However, each specific node has probability only 10^{-3} of matching in a random alignment, so the p -value of matching 10 of them is 10^{-30} .

Thus, both alignments have a mean FC of 0.82, yet—to the nearest order-of-magnitude—the first has a p -value of only $\approx 10^{-3}$, while the second has a

¹ The assumption of independence is not entirely unfounded; for example we could choose g_0 to be the *Cellular Component* (CC) GO term **GO:0005634**, which describes the location “nucleus”, and choose the remainder of GO terms to be *molecular functions* (MF) that tend to occur only outside the nucleus. In fact, in the Sept. 2018 release of the GO term database there are over 700 MF GO terms with the following properties: (a) they annotate exactly one protein (ie., each of over 700 GO terms g has $\lambda^g = 1$), and (b) for each such GO term, the one protein it annotates is *not* annotated with **GO:0005634**. The fact that over 700 such GO terms exist make our independence assumption plausible—at least in this artificial scenario.

p -value of $\approx 10^{-31}$. From a statistical significance standpoint, the second one is—quite literally—*astronomically* more significant.

The takeaway message is that any method that evaluates functional significance across pairs of aligned proteins, rather than across GO terms, is likely to lead to very misleading conclusions.

4.2 The problem with ignoring GO terms close to the root of the hierarchy

A common practice [2] involves arbitrarily ignoring GO terms in the top few levels of the GO hierarchy on the assumption that, when a GO term annotates so many proteins, a protein pair that matches it has little value. A known problem with this suggestion [2] is the definition of “top few levels”: even GO terms at the same level but different regions of the GO hierarchy can have vastly different values of λ , so that it is difficult to choose which GO terms to ignore.

From the network alignment perspective, ignoring these common GO terms has the opposite problem to that of §4.1 in that, rather than failing to *penalize* a bad alignment, this procedure fails to adequately *reward* alignments that are “good” in the following sense. Assume a GO term g annotates 10% of proteins in each network. This can be a substantial number of proteins (eg., over 1700 in human and almost 700 in mouse), and such a GO term is likely to be high in the hierarchy. However, if a network alignment matches a substantially larger fraction of this plethora of pairs than is expected at random, it is a sign that *large regions* of functional similarity are being correctly aligned to each other, even if individual proteins are not. In other words, perhaps similar pathways are being correctly mapped to each other even if the individual proteins in the pathway are incorrectly mapped. A network alignment that successfully matches such large regions should be rewarded for doing so, but if “common” GO terms are disregarded, this won’t happen.

4.3 Practical details regarding treatment of GO terms

Our analysis is easily adapted to evaluate network alignments based on any subset of GO terms. For example, one may wish to separately evaluate the three GO hierarchies of *Biological Process* (BP), *molecular Function* (MF), and *Cellular Component* (CC). Similarly, one should evaluate an alignment without the use sequence-based GO terms if sequence played any part in constructing the alignment.

Once we compute a rigorous p -value for each GO term g that appears in both networks, computing a GO-based p -value of the entire alignment requires combining the multitude of “per-GO-term” p -values into a “holistic” GO-based p -value for the entire alignment. Doing so rigorously is nontrivial, since GO terms are arranged in a hierarchy and are not independent of each other. While we provide here the first rigorous measure of the p -value of an alignment with

respect to an individual GO term, to our knowledge nobody has yet worked out how to rigorously account for the inter-dependence of GO terms. We leave for future work the job of providing such a holistic p -value.

Appendix: brief survey of existing GO-based measure of network alignments

Table 2 presents a list of alignment papers and the measures they use to evaluate functional similarity. Without exception, all of these methods evaluate each pair of aligned nodes individually, and then take some sort of average across pairs. (Some methods are not 1-to-1 and so the “pair” of aligned nodes we discuss must be generalized to a *cluster* of aligned nodes, but this generalization does not negate our point.) We are aware of no existing methods that consider the alignment from the perspective of one GO term’s performance globally across all clusters. Thus, all of these methods suffer the major drawbacks described in Section §4.1.

Below is a brief description of the methods.

- **Jaccard / FC / GOC:** A common pairwise measure is the Jaccard similarity of Eq. (11), often called GOC (for GO “consistency” or “correctness”) or FC (functional consistency/correctness).
- **Common GO terms:** In a PGNA, choose an integer threshold h (usually 1–5), and count how many aligned pairs have at least h GO terms in common. No effort is made to account for the annotation frequency (λ in our notation) of any GO term.
- **Entropy:** Given a cluster of proteins S in which d GO terms $\{g_1, \dots, g_d\}$ appear at least once across all the proteins in S , the entropy is defined as $H(S) = -\sum_{i=1}^d p_i \log p_i$, where p_i is the fraction of all proteins in S that are annotated with GO term g_i . Entropy is always non-negative and lower values are better. The *normalized entropy* is $N(S) = H(S)/m$, where m is the number of unique GO terms in S . Alignments can then be scored using *Mean Entropy* (ME) or *Mean Normalized Entropy* (MNE), which is just the appropriate mean across all clusters S .
- **Resnik:** Based on Resnik’s measure of semantic similarity [4,5], it was originally designed only to evaluate the similarity between two GO terms by finding their *most informative common ancestor* in the GO hierarchy, and using an information-theory argument to compute their common information. Later it was extended to measure similarity between gene products, such as proteins, by taking some sort of mean or maximum between the GO terms of two proteins (see, eg., [6,7,2]).
- **Schlicker’s method [6]:** an extension of Resnik’s measure tailored more closely to genes and gene products.
- **Enrichment:** has been defined in various ways but usually measures whether the shared GO terms between a pair of aligned proteins (or more generally in a cluster) is “enriched” beyond what’s expected at random.
- **m -sim:** This measure is only used by MUNK [8], which is not technically a network alignment algorithm, though it is designed to find functionally similar genes or proteins between species. It is also the only method from Table 2 that takes into account the frequency of annotation of a GO term g (λ^g in our notation), by using only GO terms with λ below some threshold m .

Compliance with Ethical Standards

This work was unfunded, and the author declares no competing interests.

Algo	year	Jac	Com	MNE	Res	Sch	Enr	m-sim
GRAAL [9]	2010	.	✓	.	.	.	.	.
H-GRAAL [10]	2010	.	✓	.	.	.	.	.
MIGRAAL [11]	2011	.	✓	.	.	.	.	.
GHOST [12]	2012	.	.	.	✓	.	.	.
NETAL [13]	2013	.	✓	.	.	✓	.	.
SPINAL [14]	2013	✓	.	.	.	.	.	.
PIswap [15]	2013	✓	.	.	.	.	.	.
BEAMS [16]	2014	✓	.	.	.	.	.	.
NetCoffee [17]	2014	.	.	.	.	✓	.	.
MAGNA [18]	2014	.	✓	✓	.	.	.	.
HubAlign [19]	2014	✓	.	.	✓	✓	.	.
SiPAN [20]	2015	✓	.	.	.	.	.	.
FUSE [21]	2015	✓	.	✓	.	.	.	.
MeAlign [22]	2015	✓	.	.	.	.	.	.
OptNetAlign [23]	2015	✓	✓	.	.	.	.	.
LGRAAL [24]	2015	.	.	.	✓	.	.	.
WAVE [25]	2015	.	.	✓	.	.	✓	.
HGA [26]	2016	.	✓	✓	.	.	.	.
DirectedGr [27]	2016	.	.	.	.	.	✓	.
ModuleAlign [28]	2016	.	.	.	.	✓	.	.
ConvexAlign [29]	2016	✓	.	✓	.	✓	.	.
PROPER [30]	2016	✓	.	.	.	.	.	.
GMalign [31]	2017	✓	.	.	✓	.	.	.
INDEX [32]	2017	✓	✓	.	✓	.	.	.
Ualign [33]	2017	.	.	.	.	.	✓	.
SANA [34]	2017	.	.	.	✓	.	.	.
GLalign [35]	2018	.	.	.	✓	.	.	.
PrimAlign [36]	2018	✓	.	.	.	.	.	.
IBNAL [37]	2018	✓	.	.	.	.	.	.
MAPPIN [38]	2018	✓	.	✓	.	✓	.	.
multiMagna [39]	2018	.	✓	✓	.	.	.	.
MUNK [8]	2019	✓	.	.	.	.	.	✓

Table 2 Sample of published network alignment algorithm names, with their citation, year, and the method(s) they used to evaluate functional similarity. The rows are sorted by publication year; the columns are sorted by popularity of evaluation measure. **Header Legend:** Jac=Jaccard Similarity (called “GOC” and “FC” by some authors); Com=number of “common” GO terms in the cluster; MNE=Mean Normalized Entropy; Res=Resnik[4, 5]; Sch=Schlicker’s method[6]; Enr=Enrichment of GO terms in a cluster compared to average cluster; *m*-sim=similarity using only GO terms with frequency (λ in our notation) less than *m*.

References

1. The Gene Ontology Consortium, Nucleic Acids Research **36**(suppl 1), D440 (2008). DOI 10.1093/nar/gkm883. URL http://nar.oxfordjournals.org/content/36/suppl_1/D440.abstract
2. C. Pesquita, D. Faria, A.O. Falcao, P. Lord, F.M. Couto, PLoS comput biol **5**(7), e1000443 (2009)
3. W.B. Hayes, in *Proceedings of ISMB 2020—Intelligent Systems for Molecular Biology* (2020)
4. P. Resnik, in *Proceedings of the 14th International Joint Conference on Artificial Intelligence - Volume 1* (Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1995), IJCAI’95, pp. 448–453. URL <http://dl.acm.org/citation.cfm?id=1625855.1625914>
5. P. Resnik, et al., J. Artif. Intell. Res.(JAIR) **11**, 95 (1999)

6. A. Schlicker, F.S. Domingues, J. Rahnenführer, T. Lengauer, BMC bioinformatics **7**(1), 302 (2006)
7. C. Pesquita, D. Faria, H. Bastos, A.E. Ferreira, A.O. Falcão, F.M. Couto, BMC bioinformatics **9**(5), S4 (2008)
8. J. Fan, A. Cannistra, I. Fried, T. Lim, T. Schaffner, M. Crovella, B. Hescott, M.D. Leiserson, Nucleic acids research **47**(9), e51 (2019)
9. O. Kuchaiev, T. Milenković, V. Memišević, W. Hayes, N. Pržulj, Journal of The Royal Society Interface **7**(50), 1341 (2010). DOI 10.1098/rsif.2010.0063
10. T. Milenković, W.L. Ng, W. Hayes, N. Pržulj, Cancer Informatics **9**, 121 (2010). DOI 10.4137/CIN.S4744. URL www.la-press.com/optimal-network-alignment-with-graphlet-degree-vectors-article-a2141
11. O. Kuchaiev, N. Pržulj, BIOINFORMATICS **27**, 1390 (2011). DOI bioinformatics/btr127. URL <http://dx.doi.org/10.1093/bioinformatics/btr127>
12. R. Patro, C. Kingsford, Bioinformatics **28**(23), 3105 (2012). DOI 10.1093/bioinformatics/bts592. URL <http://bioinformatics.oxfordjournals.org/content/28/23/3105.abstract>
13. B. Neyshabur, A. Khadem, S. Hashemifar, S.S. Arab, Bioinformatics **29**(13), 1654 (2013). DOI 10.1093/bioinformatics/btt202. URL <http://bioinformatics.oxfordjournals.org/content/29/13/1654.abstract>
14. A.E. Aladağ, C. Erten, Bioinformatics **29**(7), 917 (2013). DOI 10.1093/bioinformatics/btt071. URL <http://bioinformatics.oxfordjournals.org/content/29/7/917.abstract>
15. L. Chindelevitch, C.Y. Ma, C.S. Liao, B. Berger, Bioinformatics **29**(21), 2765 (2013). DOI 10.1093/bioinformatics/btt486. URL <http://bioinformatics.oxfordjournals.org/content/29/21/2765.abstract>
16. F. Alkan, C. Erten, Bioinformatics **30**(4), 531 (2014)
17. J. Hu, B. Kehr, K. Reinert, Bioinformatics **30**(4), 540 (2014). DOI 10.1093/bioinformatics/btt715. URL <http://dx.doi.org/10.1093/bioinformatics/btt715>
18. V. Saraph, T. Milenković, Bioinformatics **30**(20), 2931 (2014)
19. S. Hashemifar, J. Xu, Bioinformatics **30**(17), i438 (2014). DOI 10.1093/bioinformatics/btu450. URL <http://dx.doi.org/10.1093/bioinformatics/btu450>
20. F. Alkan, C. Erten, Bioinformatics **31**(14), 2356 (2015)
21. V. Gligorićević, N. Malod-Dognin, N. Pržulj, Bioinformatics p. btv731 (2015)
22. M. Gong, Z. Peng, L. Ma, J. Huang, IEEE/ACM transactions on computational biology and bioinformatics **13**(6), 1117 (2015)
23. C. Clark, J. Kalita, Bioinformatics **31**(12), 1988 (2015). DOI 10.1093/bioinformatics/btv063. URL <http://bioinformatics.oxfordjournals.org/content/31/12/1988.abstract>
24. N. Malod-Dognin, N. Pržulj, Bioinformatics (2015). DOI 10.1093/bioinformatics/btv130. URL <http://bioinformatics.oxfordjournals.org/content/early/2015/02/28/bioinformatics.btv130.abstract>
25. Y. Sun, J. Crawford, J. Tang, T. Milenković, in *Algorithms in Bioinformatics, Lecture Notes in Computer Science*, vol. 9289, ed. by M. Pop, H. Touzet (Springer Berlin Heidelberg, Germany, 2015), pp. 16–39. DOI 10.1007/978-3-662-48221-6_2. URL http://dx.doi.org/10.1007/978-3-662-48221-6_2
26. J. Xie, C. Xiang, J. Ma, J. Tan, T. Wen, J. Lei, Q. Nie, IEEE/ACM Transactions on Computational Biology and Bioinformatics (TCBB) **13**(3), 483 (2016)
27. A. Sarajlić, N. Malod-Dognin, Ö.N. Yaveroğlu, N. Pržulj, Scientific reports **6**, 35098 (2016)
28. S. Hashemifar, J. Ma, H. Naveed, S. Canzar, J. Xu, Bioinformatics **32**(17), i658 (2016)
29. S. Hashemifar, Q. Huang, J. Xu, Journal of Computational Biology **23**(11), 903 (2016)
30. E. Kazemi, H. Hassani, M. Grossgläuser, H.P. Modarres, BMC bioinformatics **17**(1), 527 (2016)
31. Y. Zhu, Y. Li, J. Liu, L. Qin, J.X. Yu, in *2017 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)* (IEEE, 2017), pp. 120–127
32. A. Mir, M. Naghibzadeh, N. Saadati, Biosystems **162**, 24 (2017)
33. N. Malod-Dognin, K. Ban, N. Pržulj, Scientific Reports **7**(1), 953 (2017)
34. N. Mamano, W.B. Hayes, Bioinformatics (Oxford, England) **33**, 2156 (2017)

35. M. Milano, P.H. Guzzi, M. Cannataro, IEEE/ACM transactions on computational biology and bioinformatics **16**(6), 1958 (2018)
36. K. Kalecky, Y.R. Cho, Bioinformatics **34**(13), i537 (2018)
37. A. Elmsallati, A. Msalati, J. Kalita, IEEE/ACM Transactions on Computational Biology and Bioinformatics (TCBB) **15**(1), 330 (2018)
38. W.E. Djeddi, S.B. Yahia, E.M. Nguifo, IEEE/ACM transactions on computational biology and bioinformatics **15**(6), 2060 (2018)
39. V. Vijayan, T. Milenković, IEEE/ACM Transactions on Computational Biology and Bioinformatics pp. 1–1 (2018). DOI 10.1109/TCBB.2017.2740381