

Statistical Model Selection Methods Applied to Biological Networks

Michael P.H. Stumpf^{1,*}, Piers J. Ingram¹, Ian Nouvel¹, and Carsten Wiuf²

¹ Centre for Bioinformatics, Department of Biological Sciences, Wolfson Building, Imperial College London, London SW7 2AZ, UK

² Bioinformatics Research Center, University of Aarhus, 8000 Aarhus C, Denmark
m.stumpf@imperial.ac.uk

Abstract. Many biological networks have been labelled scale-free as their degree distribution can be approximately described by a powerlaw distribution. While the degree distribution does not summarize all aspects of a network it has often been suggested that its functional form contains important clues as to underlying evolutionary processes that have shaped the network. Generally determining the appropriate functional form for the degree distribution has been fitted in an ad-hoc fashion.

Here we apply formal statistical model selection methods to determine which functional form best describes degree distributions of protein interaction and metabolic networks. We interpret the degree distribution as belonging to a class of probability models and determine which of these models provides the best description for the empirical data using maximum likelihood inference, composite likelihood methods, the Akaike information criterion and goodness-of-fit tests. The whole data is used in order to determine the parameter that best explains the data under a given model (*e.g.* scale-free or random graph). As we will show, present protein interaction and metabolic network data from different organisms suggests that simple scale-free models do not provide an adequate description of real network data.

1 Introduction

Network structures which connect interacting particles such as proteins have long been recognised to be linked to the underlying dynamic or evolutionary processes [13, 3]. In particular the technological advances seen in molecular biology and genetics increasingly provide us with vast amounts of data about genomic, proteomic and metabolomic network structures [15, 22, 19]. Understanding the way in which the different constituents of such networks, — genes and their protein products in the case of genome regulatory networks, enzymes and metabolites in the case of metabolic networks (MN), and proteins in the case of protein interaction networks (PIN) — interact can yield important insights into basic biological mechanisms [16, 24, 1]. For example the extent of

* Corresponding author.

phenotypic plasticity allowed for by a network, or levels of similarity between PINs in different organisms presumably depend on topological (in a loose sense of the word) properties of networks.

Our analysis here focuses on the degree distribution of a network, *i.e.* the probability of a node to have k connections to other nodes in the network. While it is well known that this does not offer an exhaustive description of network data, it has nevertheless remained an important characteristic/summary statistic of network data. Here we use $\Pr(k)$ to denote a theoretical model for the degree distribution, or $\Pr(k; \theta)$ if the model depends on an (unknown) parameter θ (potentially vector-valued), and $\hat{\Pr}(k)$ to denote the empirical degree distribution.

Many studies of biological network data have suggested that the underlying networks show scale-free behaviour [8] and that their degree distributions follow a power-law, *i.e.*

$$\Pr(k; \gamma) = k^{-\gamma} / \zeta(\gamma) \quad (1)$$

where $\zeta(x)$ is Riemann's zeta-functions which is defined for $x > 1$ and diverges as $x \rightarrow 1 \downarrow$; for finite networks, however, it is not necessary that the value of γ is restricted to values greater than 1.

These powerlaws are in marked contrast to the degree distribution of the Erdős-Rényi random graphs [7] which is Poisson, $\Pr(k; \lambda) = e^{-\lambda} \lambda^k / k!$. The study of random graphs is a rich field of research and many important properties can be evaluated analytically. Such Poisson random networks (PRN) are characterized by most nodes having comparable degree; the vast majority of nodes will have a connectivity close to the average connectivity.

The term "scale-free" means that the ratio $\Pr(\alpha k) / \Pr(k)$ depends on α alone but not on the connectivity k . The attraction of scale-free models stem from the fact that some simple and intuitive statistical models of network evolution cause powerlaw degree distribution. Scale-free networks are not, however, the only type of network that produces fat-tailed degree distributions.

Here we will be concerned with developing a statistically sound approach for inferring the functional form for the degree distribution of a real network. We will show that relatively basic statistical concepts, like maximum likelihood estimation and model selection can be straightforwardly applied to PIN and MN data. In particular we will demonstrate how we can determine which probability models best describe the degree distribution of a network. We then apply this approach in the analysis of real PIN data from five model organisms and MN data. In each case we can show that the explanatory power of a standard scale-free network is vastly inferior compared to models that take the finite size of the system into account.

2 Statistical Tools for the Analysis of Network Data

Here we are only concerned with methods aimed at studying the degree distribution of a network. In particular we want to quantify the extent to which a given functional form can describe the degree distribution of a real network. Given a probability model (*e.g.* power-law distribution or Poisson distribution) we want

to determine the parameters which describe the degree distribution best; after that we want to be able to distinguish which model from a set of trial model provides the best description. Here we briefly introduce the basic statistical concepts employed later. These can be found in much greater detail in most modern statistics texts such as [12]. Tools for the analysis of other aspects of network data, *e.g.* cluster coefficients, path length or spectral properties of the adjacency matrix will also need to be developed in order to understand topological and functional properties of networks.

There is a well established statistical literature that allows us to assess to what extent data (*e.g.* the degree distribution of a network) is described by a specific probability model (*e.g.* Poisson, exponential or powerlaw distributions). Thus far, determining the best model appears to have been done largely by eye [13] and it is interesting to apply a more rigorous approach, although in some published cases maximum likelihood estimates were used to determine the value of γ for the scale-free distribution.

2.1 Maximum Likelihood Inference

Since we only specify the marginal probability distribution, *i.e.* the degree distribution, we take a composite likelihood approach to inference, and treat the degrees of nodes as independent observations. This is only correct in the limit of an infinite sized network and finite sized sample ($n \ll N$, where N denotes network size and n the sample size). Composite likelihood methods are becoming increasingly popular in cases for which the full likelihood is difficult to specify and/or the full likelihood is intractable to calculate numerically. In our case

Table 1. Network models and their degree distributions, $\Pr(k; \theta)$. Wherever it appears, C denotes the normalizing constant such that $\sum_k \Pr(k; \theta) = 1$.

| Network type | Degree distribution $\Pr(k; \theta)$ | | Model |
|--|--|--|------------|
| Poisson | $\exp(-\lambda) \frac{\lambda^k}{k!}$ | for all $k \geq 0$ | M1 |
| Exponential | $C \exp(-k/k)$ | for all $k \geq 0$ | M2 |
| Gamma | $\frac{k^{\gamma-1} e^{-k}}{\Gamma(\gamma)}$ | for all $k \geq 0$ | M3 |
| Scale-free | 0 | for $k = 0$ | M4 |
| | $k^{-\gamma} / \zeta(\gamma)$ | for $k > 0$ | |
| Truncated scale free network | 0 | for $k < L$ and $k > M$ | M4a |
| | $k^{-\gamma} / \sum_{i=L}^M k^{-\gamma}$ | for $L \leq k \leq M$ | |
| Scale-free network with exponential cut-off | 0 | for $k < k_0$ and $k > k_{\text{cut}}$ | M4b |
| | $(k + k_0)^{-\gamma} \exp(-k/k_{\text{cut}})$ | for $k_0 \leq k \leq k_{\text{cut}}$ | |
| Lognormal | $C \frac{e^{-\ln((k-\theta)/m)^2 / (2\sigma^2)}}{(k-\theta)\sigma\sqrt{2\pi}}$ | for all $k \geq 0$ | M5 |
| Stretched exponential | 0 | for $k < 0$ | M6 |
| | $C \exp(-\alpha k / \bar{k}) k^{-\gamma}$ | for $k > 0$ | |

the full likelihood is difficult to specify. Reference [11] provides an overview of composite likelihood methods.

For a given functional form or model $\Pr(k; \theta)$ of the degree distribution we can use maximum likelihood estimation applied to the composite likelihood in order to estimate the parameter which best characterizes the distribution of the data. The composite likelihood of the model given the observed data $K = \{k_1, k_2, \dots, k_n\}$ is defined by

$$L(\theta) = \prod_{i=1}^n \Pr(k_i; \theta), \quad (2)$$

and taking the logarithm yields the log-likelihood

$$\text{lk}(M) = \text{lk}(\theta) = \sum_{i=1}^n \log(\Pr(k_i; \theta)). \quad (3)$$

The maximum likelihood estimate (MLE), $\hat{\theta}$, of θ is the value of θ for which Eqns. (2) and (3) become maximal. For this value the observed data is more probable to occur than for any other parameters.

Here the maximum likelihood framework is applied to the whole of the data. This means that in fitting a curve —such as a powerlaw $k^{-\hat{\gamma}}/\zeta(\hat{\gamma})$, where $\hat{\gamma}$ denotes the MLE of the exponent γ — data for all k is considered. If a powerlaw-dependence were to exist only over a limited range of connectivities then the global MLE curve may differ from such a localized power-law (or equivalently any other distribution).

2.2 Model Selection and Akaike Weights

We are interested in determining which model describes the data best. For non-nested models (as are considered here, *e.g.* scale-free versus Poisson) we cannot use the standard likelihood ratio test but have to employ a different information criterion to distinguish between models: here we use the Akaike-information criterion (AIC) to choose between different models [2, 10]. The AIC for a model $\Pr(k; \theta)$ is defined by

$$\text{AIC} = 2(-\text{lk}(\hat{\theta}) + d), \quad (4)$$

where $\hat{\theta}$ is the maximum likelihood estimate of θ and d is the number of parameters required to define the model, *i.e.* the dimension of θ . Note that the model is penalized by d . The model with the minimum AIC is chosen as the best model and the AIC therefore formally biases against overly complicated models. A more complicated model is only accepted as better if it contains more information about the data than a simpler model. (It is possible to formally derive the AIC from Kohn-Sham information theory.) Other information criteria exist, *e.g.* the Bayesian information criterion (BIC) offers a further method for penalizing more complex models (*i.e.* those with more parameters) unless they have significantly higher explanatory power (see [10] for details about the AIC and

model selection in statistical inference). In order to compare different models we define the relative differences

$$\Delta_j^{\text{AIC}} = \text{AIC}_j - \min_j(\text{AIC}), \quad (5)$$

where j refers to the j th model, $j = 1, 2, \dots, J$, and \min_j is minimum over all j . This in turn allows us to calculate the relative likelihoods (adjusted for the dimension of θ) of the different models, given by

$$\exp(-\Delta_j^{\text{AIC}}/2). \quad (6)$$

Normalizing these relative likelihoods yields the so-called Akaike weights w_j ,

$$w_j = \frac{\exp(-\Delta_j^{\text{AIC}}/2)}{\sum_{j=1}^J \exp(-\Delta_j^{\text{AIC}}/2)}. \quad (7)$$

The Akaike weight w_j can be interpreted as the probability that model j (out of the J alternative models) is the best model given the observed data and the range of models to choose from. The relative support for one model over another is thus given by the ratio of their respective Akaike weights. If a new model is added to the J existing models then the analysis has to be repeated. The Akaike weight formalism is very flexible and has been applied in a range of context including the assessment of confidence in phylogenetic inference [20]. In the next section we will apply this formalism to PIN data from five species and estimate the level of support for each of the models in table 1.

2.3 Goodness-of-Fit

In addition to the AIC or similar information criteria we can also assess a model's performance at describing the degree distribution using a range of other statistical measures. The Kolmogorov-Smirnoff (KS)[12] and Anderson-Darling (AD) [4, 5] goodness-of-fit statistics allow us to quantify the extent to which a theoretical or estimated model of the degree distribution describes the observed data. The former is a common and easily implemented statistic, but the latter puts more weight on the tails of distributions and also allows for a secular dependence of the variance of the observed data on the argument (here the connectivity k). The KS statistic is defined as

$$D = \max |\hat{P}(k) - P(k)|, \quad (8)$$

where $\hat{P}(k)$ and $P(k)$ are the empirical and theoretical cumulative distribution functions, respectively, for a node's degree *i.e.* $P(k) = \sum_{i=1}^k \text{Pr}(i)$ and $\hat{P}(k) = \sum_{i=1}^k \hat{\text{Pr}}(i)$. If $P(k)$ depends on θ , $P(k)$ is substituted by $P(k; \hat{\theta}) = \sum_{i=1}^k \text{Pr}(i; \hat{\theta})$, the estimated cumulative distribution. This statistic is most sensitive to differences between the theoretical (or estimated) and observed distributions around the median of the data, *i.e.* the point where $P(k) \approx 0.5$. Given that we will

also be considering a number of fat-tailed distributions this is somewhat unsatisfactory and we will therefore also use the AD statistic (Anderson and Darling discussed a number of statistics [4, 5]) which is defined as

$$D^* = \max \frac{|\hat{P}(k) - P(k)|}{\sqrt{P(k)(1 - P(k))}} \quad (9)$$

(again $P(k)$ might be substituted for $P(k; \hat{\theta})$).

We can use these statistics for two purposes: first, we can use them to compare different trial distributions as the "best" distribution should have the smallest value of D and D^* , respectively. Second, we can use these statistics to determine if the empirical degree distribution is consistent with a given theoretical (or estimated) distribution.

To evaluate the fit of a model, p -values can be calculated for the observed values of D and D^* using a parametric boot-strap procedure using the estimated degree distribution: for a network with N nodes we repeatedly sample N values at random from the maximum likelihood model, $\Pr(k, \hat{\theta})$ and calculate D^* and D^* , respectively, for each replicate. From L bootstrap-replicates we obtain the Null distribution of D and D^* under the estimated degree distribution. For sufficiently large L we can thus determine approximate p -values which allow us to test if the empirical degree distribution is commensurate with the estimated degree distribution.

3 Statistical Analysis of Biological Networks

Here we apply the analysis of the preceding sections to the study of PINs and metabolic networks. It is easy to find a straight-line fit to some degree interval for all of the datasets considered here. For a powerlaw to be meaningful it has to extend over at least two or three decades and with a maximum degree of $k_{\max} \lesssim 300$ this will be unachievable for the present data sets. We therefore use all the data and fit the model which yields the best overall a description of the degree distribution.

3.1 Analysis of PIN Data

In table 2 we show the maximum composite likelihoods for the degree distributions calculated from PIN data collected in five model organisms [23] (the protein interaction data was taken from the DIP data-base; <http://dip.doe-mbi.ucla.edu>). We find that the standard scale-free model (or its finite size versions) never provides the best fit to the data; in three networks (*C.elegans*, *S.cerevisiae* and *E.coli*) the lognormal distribution (M5) explains the data best. In the remaining two organisms the stretched exponential model provides the best fit to the data. The bold likelihoods correspond to the highest Akaike weights. Apart from the case of *H.pylori* (where $\max(w_j) = w_5 \approx 0.95$ for M5 and $w_6 \approx 0.05$) the value of the maximum Akaike weight is always > 0.9999 . For *C.elegans*,

however, the scale-free model and its finite size versions are better than the lognormal model, M5.

Table 2. Log-likelihoods for the degree distributions from table 1 in five model organisms. M1, M2, M3 and M4 have one free parameter, M4a, M4b and M5 have two free parameters, while M6 has three free parameters. The differences in the log-likelihoods are, however, so pronounced that the different numbers of parameters do not noticeably influence the AIC.

| Organism | M1 | M2 | M3 | M4 | M4a | M4b | M5 | M6 |
|-----------------------|--------|--------|--------|--------|--------|--------|---------------|---------------|
| <i>D.melanogaster</i> | -38273 | -20224 | -29965 | -18517 | -18257 | -18126 | -17835 | -17820 |
| <i>C.elegans</i> | -9017 | -5975 | -6071 | -4267 | -4258 | -4252 | -4328 | -4248 |
| <i>S.cerevisiae</i> | -24978 | -14042 | -20342 | -13454 | -13281 | -13176 | -12713 | -12759 |
| <i>H. pylori</i> | -2552 | -1776 | -2052 | -1595 | -1559 | -1546 | -1527 | -1529 |
| <i>E.coli</i> | -834 | -884 | -698 | -799 | -789 | -779 | -659 | -701 |

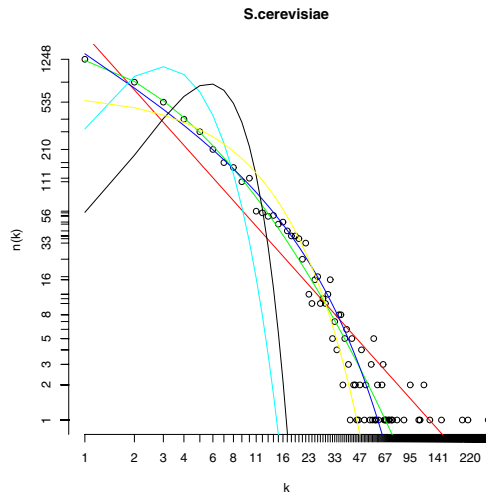


Fig. 1. Yeast protein interaction data (o) and best-fit probability distributions: Poisson (—), Exponential (—), Gamma (—), Power-law (—), Lognormal (—), Stretched exponential (—). The parameters of the distributions shown in this figure are the maximum likelihood estimates based on the real observed data.

For the yeast PIN the best fit curves (obtained from the MLEs of the parameters of models M1-M6) are shown in figure 1, together with the real data. Visually, log-normal (green) and stretched exponential (blue) appear to describe the data almost equally well. Closer inspection, guided by the Akaike weights, however, shows that the fit of the lognormal to the data is in fact markedly better than the fit of the stretched exponential. But the failure of quickly decaying

distributions such as the Poisson distribution, characteristic for classical random graphs [7] to capture the behaviour of the PIN degree distribution is obvious.

Interestingly, common heuristic finite size corrections to the standard scale-free model improve the fit to the data (measured by the AIC). But compared to the lognormal and stretched exponential models they still fall short in describing the PIN data in the five organisms. Figure 2 shows only the three curves with

Table 3. The KS and AD statistics D and D^* for the scale-free, lognormal and stretched exponential model. The ordering of these models is in agreement with the AIC, but KS and AD statistics capture different aspects of the degree distribution. We see that the likelihood treatment, which takes in all the data, agrees better with the KS statistic. In the tails (especially for large connectivities k) the maximum likelihood fit sometimes —especially in the case of *E.coli*— can provide a very bad description of the observed data.

| Species | M4 | | M5 | | M6 | |
|-----------------------|------|-------|------|-------|------|-------|
| | D | D^* | D | D^* | D | D^* |
| <i>D.melanogaster</i> | 0.13 | 0.26 | 0.01 | 0.06 | 0.02 | 0.06 |
| <i>C.elegans</i> | 0.03 | 0.09 | 0.10 | 0.20 | 0.02 | 0.08 |
| <i>S.cerevisiae</i> | 0.17 | 0.33 | 0.01 | 0.04 | 0.03 | 5.99 |
| <i>H. pylori</i> | 0.13 | 0.26 | 0.01 | 0.05 | 0.02 | 0.12 |
| <i>E.coli</i> | 0.28 | 0.56 | 0.04 | 56.09 | 0.12 | 6072 |

the highest values of ω_j , which apart from *E.coli* are the log-normal, stretched exponential and power-law distributions; for *E.coli*, however, the Gamma distribution replaces the power-law distribution. These figures show that, apart from *C.elegans* the shape of the whole degree distribution is not power-law like, or scale-free like, in a strict sense. Again we find that log-normal and stretched exponential distributions are hard to distinguish based on visual assessment alone. Figures 1 and 2, together with the results of table 2, reinforce the well known point that it is hard to choose the best fitting function based on visual inspection. It is perhaps worth noting, that the PIN data is more complete for *S.cerevisiae* and *D.melanogaster* than for the other organisms.

The standard scale-free model is superior to the log-normal only for *C.elegans*. The order of models (measured by decreasing Akaike weights) is M6, M5, M4, M2, M3, M1 for *D.melanogaster*, M6, M4, M5, M2, M3, M1 for *C.elegans*, M5, M6, M4, M2, M3, M1 for *S.cerevisiae* and *H.pylori*, and M5, M3, M6, M4, M2, M1 for *E.coli*. Thus in the light of present data the PIN degree distribution of *E.coli* lends more support to a Gamma distribution than to a scale-free (or even stretched scale-free) model. There is of course, no mechanistic reason why the gamma distribution should be biologically plausible but this point demonstrates that present PIN data is more complicated than predicted by simple models. Therefore statistical model selection is needed to determine the extent to which simple models really provide insights into the intricate architecture of PINs. For

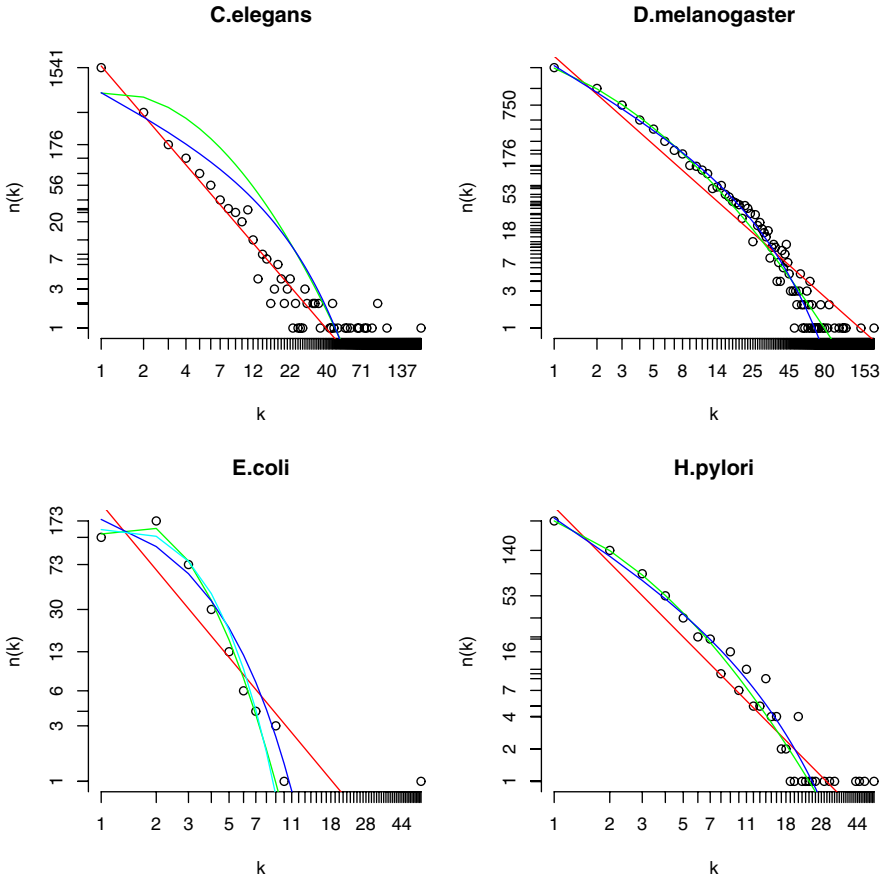


Fig. 2. Degree distributions of the protein interaction networks (o) of *C.elegans*, *D.melanogaster*, *E.coli* and *H.pylori*. The power-law (—), log-normal (—) and stretched exponential (—) models are shown for all figures; for *E.coli* the gamma distribution (—), which performs better (measured by the Akaike weights) than either scale-free and the stretched exponential distributions.

completeness we note that model selection based on BIC results in the same ordering of models as the AIC shown here.

In table 3 we give the values of D and D^* for the empirical degree distributions. The estimated cumulative distribution $P(k; \hat{\theta})$ was obtained from the maximum likelihood fits of the respective models. The results in table 3 show that the maximum likelihood framework (or the respective models) sometimes cannot adequately describe the tails of the distribution —at low and high values of the connectivity— in some cases, where $D^* \gg D$. The order of the different models suggested by D for the three models generally agrees with the ordering obtained from the AIC.

3.2 Analysis of MN Data

Metabolic networks aim to describe the biochemical machinery underlying cellular processes. Here we have used data from the KEGG database (www.genome.jp/kegg) with additional information from the BRENDA database (www.brenda.uni-koeln.de). The nodes are the enzymes and an edge is added between two enzymes if one uses the products of the other as educts.

Table 4. Log-likelihoods for the degree distributions from table 1 applied to metabolic networks. Human and yeast data were extracted from the global database.

| Data | M1 | M2 | M3 | M4 | M5 | M6 |
|---------------------|-------|-------|-------|-------|-------|--------------|
| KEGG | -8276 | -5247 | -7676 | -5946 | -5054 | -4178 |
| <i>H.sapiens</i> | -1619 | -1390 | -1565 | -1525 | -1312 | -1308 |
| <i>S.cerevisiae</i> | -2335 | -1485 | -2185 | -1621 | -1436 | -1427 |

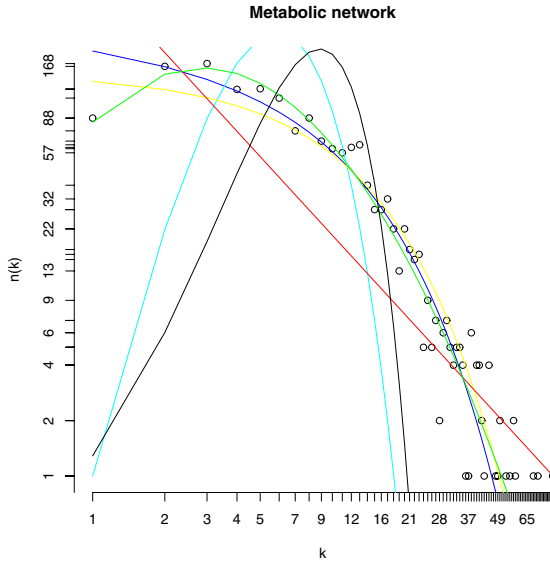


Fig. 3. Degree distributions of the metabolic network data (o) and best-fit probability distributions: Poisson (—), Exponential (—), Gamma (—), Power-law (—), Log-normal (—), Stretched exponential (—). The parameters of the distributions shown in this figure are the maximum likelihood estimates based on the real observed data. Ignoring low degrees ($k = 1, 2$) when fitting the scale-free model increases the exponent (*i.e.* it falls off more steeply) but compared to the other models no increased performance is observed.

From figure 3.2 and table 4 it is apparent that the maximum likelihood scale-free model obtained from the whole network data does not provide an

adequate description of the MN data. This should, however, not be too surprising as the network is only relatively small with maximum degree $k_{\max} = 83$. The degree distribution appears to decay in an essentially exponential fashion but the stretched exponential has the required extra flexibility to describe the whole of the data better than the other models. Measured by goodness of fit statistics D and D^* , however, the log-normal model ($D = 0.02$ and $D^* = 0.06$) performs rather better than the stretched exponential ($D = 0.07$ and $D = 0.22$); the scale-free model again performs very badly ($D = 1.0$ and $D^* = 34.2$) for the metabolic network data.

4 Conclusions

We have shown that it is possible to use standard statistical methods in order to determine which probability model describes the degree distribution best. We find that the common practice of fitting a pure power-law to such experimental network data [13, 3] may obscure information contained in the degree distribution of biological networks. This is often done by identifying a range of connectivities from the log-log plots of the degree distribution which can then be fitted by a straight line. Not only is this wasteful in the sense that not all of the data is used but it may obfuscate real, especially finite-size, trends. The same will very likely hold true for other biological networks, too [17]. The approach used here, on the other hand, (i) uses all the data, and (ii) can be extended to assessing levels of confidence through combining a bootstrap procedure with the Akaike weights. What we have shown here, in summary, is that statistical methods can be usefully applied to protein interaction and metabolic network data.

Even though the degree distribution does not capture all (or even most) of the characteristics of real biological networks there is reason to reevaluate previous studies. We find that formally real biological networks provide very little support for the common notion that biological networks are scale-free. Other fat-tailed probability distributions provide qualitatively and quantitatively better descriptions of the degree distributions of biological networks. For protein interaction networks we found that the log-normal and the stretched exponential offer superior descriptions of the degree distribution than the powerlaw or its finite size versions. For metabolic versions our results confirms this. Even the exponential model outperformed the scale-free model in describing the empirical degree distribution (we note that randomly growing networks are characterized by an exponentially decreasing degree distribution). The best models are all fat-tailed—like the scale-free models—but are not formally scale-free. Unfortunately, there is as yet no known physical model for network growth processes that would give rise to log-normal or stretched exponential degree distributions.

There is thus a need to develop and study theoretical models of network growth that are better able to describe the structure of existing networks. This probably needs to be done in light of at least three constraints: (i) real networks are finite sized and thus, in the terms of statistical physics, mesoscopic systems; (ii) present network data are really only samples from much larger net-

works as not all proteins are included in present experimental setup (in our case *S.cerevisiae* has the highest fraction, 4773 out of approximately 5500-6000 proteins); the sampling properties have recently been studied and it was found that generally the degree distribution of a subnet will differ from that of the whole network. This is particularly true for scale-free networks. (iii) biological networks are under a number of functional and evolutionary constraints and proteins are more likely to interact with proteins in the same cellular compartment or those involved in the same biological process. This modularity—and the information already available *e.g.* in gene ontologies—needs to be considered. Finally there is an additional caveat: biological networks are not static but likely to change during development. More dynamic structures may be required to deal with this type of problem.

Quite generally we believe that we are now at a stage where simple models do not necessarily describe the data collected from complex processes to the extent that we would like them to. But as Burda, Diaz-Correia and Krzywicki point out [9], even if a mechanistic model is not correct in detail, a corresponding statistical ensemble may nevertheless offer important insights. We believe that the statistical models employed here will also be useful in helping to identify more realistic ensembles.

The maximum likelihood, goodness of fit and other tools and methods for the analysis of network data are implemented in the NetZ R-package which is available from the corresponding author on request.

Acknowledgements. We thank the Wellcome Trust for a research fellowship (MPHS) and a research studentship (PJI). CW is supported by the Danish Cancer Society. Financial support from the Royal Society and the Carlsberg Foundation (to MPHS and CW) is also gratefully acknowledged. We have furthermore benefitted from discussions with Eric de Silva, Bob May and Mike Sternberg.

References

- [1] I. Agraftoti, J. Swire, J. Abbott, D. Huntley, S. Butcher and M.P.H. Stumpf. Comparative analysis of the *Saccharomyces cerevisiae* and *Caenorhabditis elegans* protein interaction networks. *BMC Evolutionary Biology*, 5:23, 2005.
- [2] H. Akaike. Information measures and model selection. In *Proceedings of the 44th Session of the International Statistical Institute*, pages 277–291, 1983.
- [3] R. Albert and A. Barabási. Statistical mechanics of complex networks. *Rev. Mod. Phys.*, 74:47, 2002.
- [4] T.W. Anderson and D.A. Darling. Asymptotic theory of certain goodness-of-fit criteria based on stochastic processes. *Ann.Math.Stat.*, 23:193, 1952.
- [5] T.W. Anderson and D.A. Darling. A test of goodness of fit. *J.Am.Stat.Assoc.*, 49:765, 1954.
- [6] A. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286:509, 1999.
- [7] B. Bollobás. *Random Graphs*. Academic Press, London, 1998.
- [8] B. Bollobás and O. Riordan. Mathematical results on scale-free graphs. In S. Bornholdt and H. Schuster, editors, *Handbook of Graphs and Networks*, pages 1–34. Wiley-VCH, 2003.

- [9] Z. Burda and J. Diaz-Correia and A. Krzywicki. Statistical ensemble of scale-free random Graphs. *Phys. Rev. E*, 64,;046118,2001.
- [10] K. Burnham and D. Anderson. *Model selection and multimodel inference: A practical information-theoretic approach*. Springer, 2002.
- [11] D. R. Cox and Reid. A note on pseudolikelihood constructed from marginal densities. *Biometrika*, 91:729, 2004.
- [12] A. Davison. *Statistical models*. Cambridge University Press, Cambridge, 2003.
- [13] S. Dorogovtsev and J. Mendes. *Evolution of Networks*. Oxford University Press, Oxford, 2003.
- [14] S. Dorogovtsev, J. Mendes, and A. Samukhin. Multifractal properties of growing networks. *Europhys. Lett.*, 57:334;cond-mat/0106142, 2002.
- [15] T. Ito, K. Tashiro, S. Muta, R. Ozawa, T. Chiba, M. Nishizawa, K. Yamamoto, S. Kuhara, and Y. Sakaki. Towards a protein-protein interaction map of the budding yeast: A comprehensive system to examine two-hybrid interactions in all possible combinations between the yeast proteins. *PNAS*, 97:1143, 2000.
- [16] S. Maslov and K. Sneppen. Specificity and stability in topology or protein networks. *Science*, 296:910, 2002.
- [17] R. May and A. Lloyd. Infection dynamics on scale-free networks. *Phys. Rev. E*, 64:066112, 2001.
- [18] M. Newman, S. Strogatz, and D. Watts. Random graphs with arbitrary degree distribution and their application. *Phys. Rev. E*, 64:026118;cond-mat/0007235, 2001.
- [19] H. Qin, H. Lu, W. Wu, and W. Li. Evolution of the yeast interaction network. *PNAS*, 100:12820, 2003.
- [20] K. Strimmer and A. Rambaut. Inferring confidence sets of possibly misspecified gene trees. *Proc. Roy. Soc. Lond. B*, 269:127, 2002.
- [21] M.P.H. Stumpf, C. Wiuf and R.M. May Subnets of scale-free networks are not scale-free: the sampling properties of random networks. *PNAS*, 103:4221, 2005.
- [22] A. Wagner. The yeast protein interaction network evolves rapidly and contains few redundant duplicate genes. *Mol. Biol. Evol.*, 18:1283, 2001.
- [23] I. Xenarios, D. Rice, L. Salwinski, M. Baron, E. Marcotte, and D. Eisenberg. Dip: the database of interacting proteins. *Nucl. Acid. Res.*, 28:289, 2000.
- [24] S. Yook, Z. Oltvai, and A. Barabási. Functional and topological characterization of protein interaction networks. *Proteomics*, 4:928, 2004.