Edge direction and the structure of networks

Jacob G. Foster[1], David V. Foster, Peter Grassberger, and Maya Paczuski

Complexity Science Group, Department of Physics and Astronomy, University of Calgary, 2500 University Drive NW, Calgary, AB, Canada T2N 1N4

Edited by H. Eugene Stanley, Boston University, Boston, MA, and approved April 6, 2010 (received for review November 2, 2009)

Directed networks are ubiquitous and are necessary to represent complex systems with asymmetric interactions—from food webs to the World Wide Web. Despite the importance of edge direction for detecting local and community structure, it has been disregarded in studying a basic type of global diversity in networks: the tendency of nodes with similar numbers of edges to connect. This tendency, called assortativity, affects crucial structural and dynamic properties of real-world networks, such as error tolerance or epidemic spreading. Here we demonstrate that edge direction has profound effects on assortativity. We define a set of four directed assortativity measures and assign statistical significance by comparison to randomized networks. We apply these measures to three network classes—online/social networks, food webs, and word-adjacency networks. Our measures (*i*) reveal patterns common to each class, (*ii*) separate networks that have been previously classified together, and (*iii*) expose limitations of several existing theoretical models. We reject the standard classification of directed networks as purely assortative or disassortative. Many display a class-specific mixture, likely reflecting functional or historical constraints, contingencies, and forces guiding the system's evolution.

Complex networks reveal essential features of the structure, function, and dynamics of many complex systems (1–4). While networks from diverse fields share various properties (3, 5–7) and universal patterns (1, 3), they also display enormous structural, functional, and dynamical diversity. A basic measure of diversity is assortativity by degree (hereafter assortativity): the tendency of nodes to link to other nodes with a similar number of edges (4, 8, 9). Despite its importance, no disciplined approach to assortativity in directed networks has been proposed. Here we present such an approach and show that measures of directed assortativity provide a number of insights into the structure of directed networks and key factors governing their evolution.

Assortativity is a standard tool in analyzing network structure (4) and has a simple interpretation. In assortative networks with symmetric interactions (i.e., undirected networks), high degree nodes, or nodes with many edges, tend to connect to other high degree nodes. Hence, assortative networks remain connected despite node removal or failure (9), but are hard to immunize against the spread of epidemics (10). In disassortative networks, conversely, high degree nodes tend to connect to low degree nodes (8, 9); these networks limit the effects of node failure because important nodes (with many edges) are isolated from each other (11). Assortativity has a convenient global measure: the Pearson correlation (r) between the degrees of nodes sharing an edge (8, 9). It ranges from -1 to 1, with ($r > 0$) in assortative networks and ($r < 0$) in disassortative ones. Earlier work proposed a simple classification of networks on the basis of assortativity, in which social networks are assortative and biological and technological networks are disassortative (4, 8, 9). Recent work suggests that this classification does not hold for undirected networks: Many online social networks are disassortative (12). We go further, demonstrating that the simple assortative/disassortative dichotomy misses fundamental features of networks where edge direction plays a crucial role. In fact, we show that many networks are neither purely assortative nor disassortative, but display a mixture of both tendencies. These patterns provide a classification scheme for networks with asymmetric interactions.

In directed networks, an edge from source to target (A → B) represents an asymmetric interaction; for example, that Web site A contains a hyperlink to Web site B, or organism A is eaten by organism B. Edge direction is essential to evaluate and explain local structure in such networks. For instance, motif analysis (13, 14) identifies local connection patterns that appear more frequently in the real-world network than in ensembles of randomized networks. In this context, edge direction distinguishes functional units like feed-forward and feedback loops. Taking edge direction into account also overturns the simple picture of the World Wide Web (WWW) as having a short average distance between all Web pages (15) in favor of a richer picture of link flow into and out of a dense inner core (16). More recently, attempts to identify communities in directed networks have demonstrated that ignoring edge direction misses key organizational features of community structure in networks (17–19). Hence it is striking that assortativity in directed networks has been studied only by ignoring edge direction entirely (8) or by measuring a subset of the four possible degree-degree correlations (9, 20). All four degree-degree correlations were addressed in the specific contexts of earthquake recurrences (21) and the WWW (22) using the average neighbor degree, e.g., $\langle k'^{\mathrm{out}}\rangle_{\mathrm{nn}}(k^{\mathrm{in}})$, as a measure rather than the Pearson correlation. However, it is easier to interpret and assign statistical significance to the Pearson correlation. Moreover, the average neighbor degree cannot be easily used to quantify the diversity of a given network or to compare networks of various sizes, unlike the Pearson correlation. Incorporating edge direction into familiar assortativity measures based on the Pearson correlation is an essential step to better characterize, understand, and model directed networks. Indeed, since they scale as $\mathcal{O}(E)$, where E is the number of edges in the network, our directed assortativity measures can be evaluated for large networks that are beyond the reach of current motif analysis or community detection algorithms.

Here we analyze online and social networks, food webs, and word-adjacency networks. Classes of directed networks show common patterns across the four directed assortativity measures: $r(out,in)$; $r(in,out)$; $r(out,out)$; and $r(in,in)$. The first element in the parentheses labels the degree of the source node of the directed edge, and the second labels the degree of the target node. Thus $r(in,out)$ quantifies the tendency of nodes with high in-degree to connect to nodes with high out-degree, and so on; see Fig. 1.

We compare the real-world network with an ensemble of randomized networks. This comparison allows us to assign statistical significance to each measure.* We use that significance to define an Assortativity Significance Profile for each network. This profile allows us to distinguish between networks grouped together

---

Author contributions: J.G.F., D.V.F., P.G., and M.P. designed research; J.G.F., D.V.F., P.G., and M.P. performed research; J.G.F., D.V.F., P.G., and M.P. analyzed data; and J.G.F., D.V.F., P.G., and M.P. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

*To our knowledge, statistical significance has been assigned to assortativity measures in only one publication (23).

[1] To whom correspondence should be addressed. E-mail: jgfoster@uchicago.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.0912671107/-/DCSupplemental.

by other measures; indeed, we find that online and social networks, which have similar motif structure (14), have substantially *different* assortativity profiles. The class-specific profiles point to forces or constraints that may guide the structure, function and growth of that class (14, 24, 25). We also uncover limitations of several theoretical network models. For example, neither of two plausible models of word-adjacency networks [one proposed by Milo et al. (14), the other in this paper] can reproduce the directed assortativity profile we observe in the real-world networks. A standard model of the WWW (26) is similarly unsuccessful. On the other hand, the food web models (27) examined here reproduce the pattern of assortativity seen in different food webs. Hence our measures provide useful benchmarks to test models of network formation.

Table S1 provides descriptions and sources for all networks examined in this paper; Table S2 collects the full results including error estimates.

Results and Discussion

Since nodes in directed networks have both an in-degree and an out-degree, we introduce a set of four directed assortativity measures. Fig. 1 illustrates this set, with examples typical of assortative or disassortative networks. Let $\alpha, \beta \in \{in, out\}$ index the degree type, and j_i^α and k_i^β be the α- and β-degree of the source node and target node for edge i. Then we define the set of assortativity measures using the Pearson correlation:

$$r(\alpha,\beta) = \frac{E^{-1}\sum_i [(j_i^\alpha - \overline{j^\alpha})(k_i^\beta - \overline{k^\beta})]}{\sigma^\alpha \sigma^\beta} \qquad [1]$$

where E is the number of edges in the network, $\overline{j^\alpha} = E^{-1}\sum_i j_i^\alpha$, and $\sigma^\alpha = \sqrt{E^{-1}\sum_i (j_i^\alpha - \overline{j^\alpha})^2}$; $\overline{k^\beta}$ and σ^β are similarly defined. In each correlation the edges point from the node with the α-indexed degree to the node with the β-indexed degree (*Materials and Methods*). We assign errors by jackknife resampling (9) and plot 2σ-error bars in the figures.

To estimate statistical significance, we compare the degree-degree correlations for each real-world network to a null model. We use as our null model the ensemble of randomized networks with the same in- and out-degree sequence [number of nodes $n(k^{\text{in}}, k^{\text{out}})$ with in-degree k^{in} and out-degree k^{out}; hereafter degree sequence] as the original network (11–14, 24, 25) (*Materials and Methods*). The comparison distinguishes features accounted for by the degree sequence from those that might reflect other forces or constraints. Our method assigns each correlation $r(\alpha,\beta)$ a statistical significance through its Z score:

$$Z(\alpha,\beta) = \frac{r_{\text{rw}}(\alpha,\beta) - \langle r_{\text{rand}}(\alpha,\beta)\rangle}{\sigma[r_{\text{rand}}(\alpha,\beta)]}. \qquad [2]$$

This quantifies the difference between the assortativity measure of the real-world network $r_{\text{rw}}(\alpha,\beta)$ and its average value in the randomized ensemble $\langle r_{\text{rand}}(\alpha,\beta)\rangle$ in units of the standard deviation, $\sigma[r_{\text{rand}}(\alpha,\beta)]$. Larger networks typically have larger Z scores (see Table S2). To compare networks of various sizes, the Z scores are normalized (14) by defining an Assortativity Significance Profile (ASP), where $\text{ASP}(\alpha,\beta) = Z(\alpha,\beta)/[\sum_{\alpha,\beta} Z(\alpha,\beta)^2]^{1/2}$. This quantity is directly related to the Z score, and for a given network the normalization does not change the relative size of the significance measures. To separate less significant correlations, we indicate $|Z(\alpha,\beta)| < 2$ in all figures by an appropriately colored asterisk. A positive $Z(\alpha,\beta)$ or $\text{ASP}(\alpha,\beta)$ ("Z assortative") indicates that the real-world network is more assortative in that measure than expected based on the degree sequence. A negative $Z(\alpha,\beta)$ or $\text{ASP}(\alpha,\beta)$ ("Z disassortative") means that the original network is less assortative than expected.

Online and Social Networks. We first consider online and social networks. In an online network, edges represent hyperlinks. In the social networks considered here, edges represent positive sentiment. Online networks are built collaboratively and share motif patterns with social networks, leading them to be grouped in the same "superfamily" (14). Fig. 2*A* shows the ASP of the World Wide Web sample and two social networks studied in ref. 14. Each network differs significantly in its ASP, showing that the ASP discriminates between networks with similar motif structure. Fig. 2*B* shows the ASP of the WWW, Wikipedia (28), and a network of political blogs (29). All three networks are (*out*, *in*) Z disassortative, indicating that the small disassortative effects measured previously (9, 30) represent substantial deviations from expected behavior. This may reflect different growth mechanisms and/or functional constraints. The WWW and Wikipedia are also (*in*, *out*) Z assortative. This property indicates that pages with high in-degree [corresponding to "authorities" (31)] link to pages with high out-degree [useful pages (31)] more frequently than expected based on the degree sequence. Pages can be both authorities and useful, and in the WWW these "multihubs" are highly interconnected; this effect creates the (*in*, *out*) correlation, along with a tendency for low in-degree nodes to connect to low out-degree nodes. All three online networks show no assortative or disassortative tendency in the (*out*, *out*) or (*in*, *in*) measures, consistent with previous work on the average neighbor in-degree in Wikipedia (32). The effects of Z-assortative or -disassortative behavior can be huge, e.g., an increase of 268% in the number of connections from the top 5% of in-degree nodes (hereafter in-hubs) to the top 5% of out-degree nodes (hereafter out-hubs) in the real-world Leadership network, compared to the randomized ensemble. The smallest change is a 1.7% decrease (blogs, in-hub to out-hub). The (*in*, *out*) effect for the WWW is substantial: an 82.3% increase in connections from in-hubs to out-hubs.

Models of online network growth should reproduce the qualitative features of each online $r(\alpha,\beta)$, $Z(\alpha,\beta)$, and $\text{ASP}(\alpha,\beta)$. We tested a directed preferential attachment model for the WWW (*Materials and Methods*) (26). This model fails to generate any of the ASP characteristics of the WWW (Fig. 2*C*). As shown in Fig. 2*D*, $r(in, out)$ is small in the growth model, whereas $r(in, out) = 0.2567$ is large for the WWW. This difference arises because the growth model fails to generate many connections between multihubs or between low in- and low out-degree nodes.

Thus $r(\alpha,\beta)$ and $\text{ASP}(\alpha,\beta)$ for the three online networks cannot be attributed to the degree sequence or simple models of network growth. The (*out*, *in*) Z disassortativity may reflect that hyperlinking and (more generally) information have a hierarchical structure, e.g., the existence of distinct "high-level" topics—much as disassortativity in protein interaction networks captures

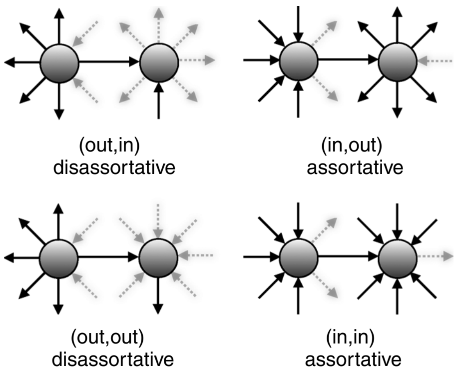

Fig. 1. The four degree-degree correlations in directed networks. The fuzzy edges indicate that nodes can have any number of edges of this type, as they do not enter into the specific correlation. For each correlation we show an example typical of assortative or disassortative networks.

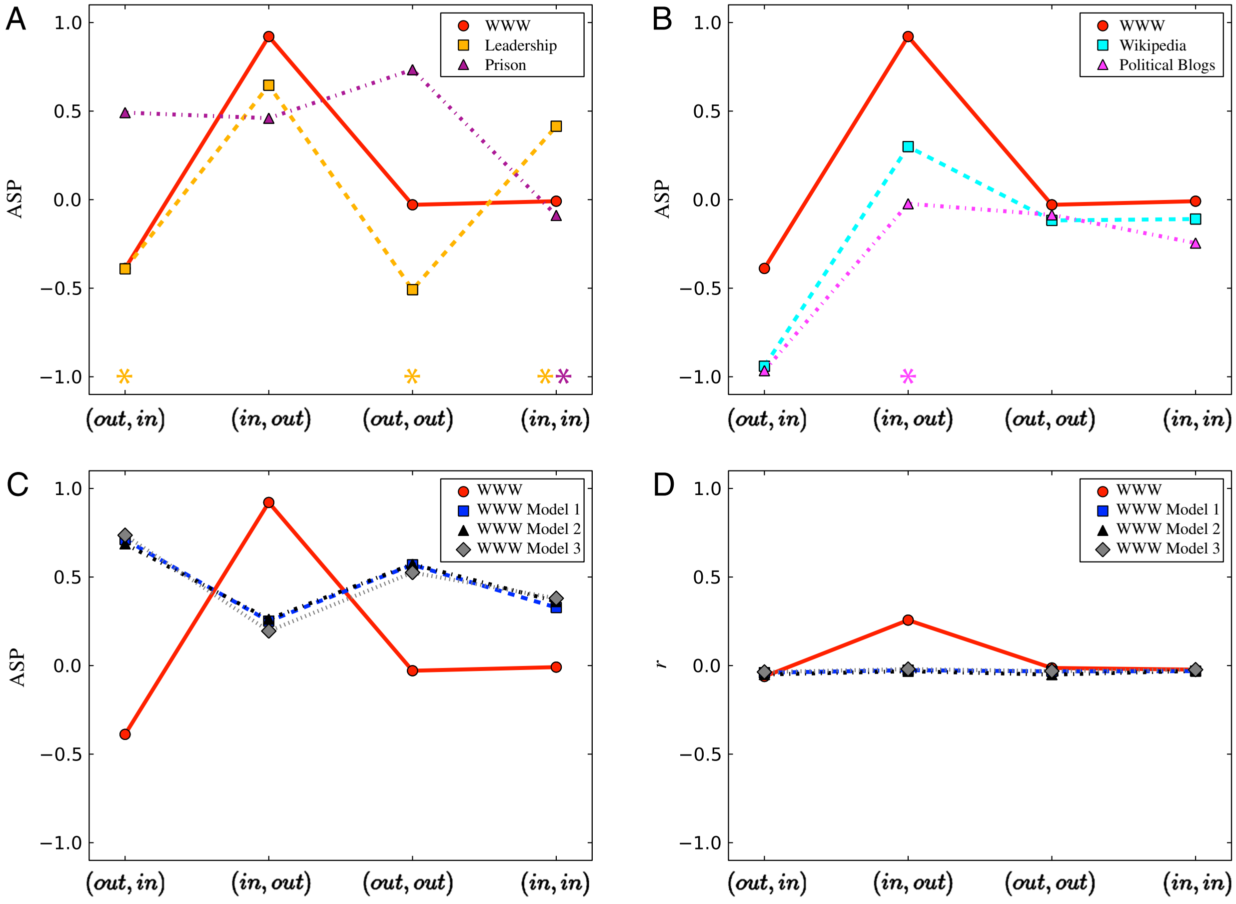

Fig. 2. Online networks differ from social networks and growth models. (*A*) The ASP for a subset of the WWW [edges represent hyperlinks (14)] and two social networks [students in a leadership class and prisoners, edges represent positive sentiment (14)]. The three networks differ substantially, despite having similar motif patterns (14). In cases where $|Z| < 2$, the corresponding ASP is marked with an appropriately colored asterisk. Only Prison (*in*, *in*) has $|Z| < 1$. (*B*) The ASP for the WWW, a snapshot of Wikipedia (28), and a collection of political blogs (29). All three online networks are more (*out*, *in*) disassortative than expected from the degree sequence alone; the WWW and Wikipedia are significantly (*in*, *out*) assortative. The blog network has $Z(in, out) = -0.609$ and does not differ significantly from the ensemble in this measure. All other *Z* scores are significant. (*C* and *D*) Three realizations of the WWW growth model (26) fail to reproduce the $\text{ASP}(\alpha, \beta)$ or $r(\alpha, \beta)$ of the WWW. Errors in *r*, estimated via jackknife (9), are smaller than the symbols.

the existence of weakly connected modules (11). The large (*in*, *out*) assortativity and *Z* assortativity of the WWW are especially pertinent for how users navigate the Web. High in-degree nodes (authorities) may gain their status by aggregating links to useful pages (with high out-degree). This pairing of trusted authorities and useful pages would provide broad access to relevant information on the Web. We find that more than half of the authorities (in-hubs) are also useful (out-hubs): Hence they may become authorities by themselves being useful. We further find that these multihubs interconnect preferentially, whereas pages with low in-degree connect preferentially to pages with low out-degree. These results are consistent with the bow-tie structure revealed by a much more computationally costly analysis (16): a densely interconnected and highly navigable core, with less trusted or useful pages clumping into small clusters or chains.

Food webs. We now turn to food webs (33). Recall that a directed edge from species A to species B means that A is eaten by B. Food webs from diverse ecosystems display universal properties, e.g., a common form for the in- and out-degree distributions (34, 35). Previous work indicated that food webs are disassortative in the (*out*, *in*) measure (9). As shown in Fig. 3*A*, although $r(out, in)$ is disassortative for all food webs (36–40), we see a wide range of values from *Z* disassortative to *Z* assortative in the (*out*, *in*) ASP measure of Fig. 3*B*. Thus, once the degree sequence is taken into account, no common pattern remains in this measure.

In contrast, food webs are both disassortative and *Z* disassortative in the (*in*, *out*) measure. This means that organisms with many prey species are eaten by organisms with few predator species (and vice versa) more frequently than expected. This tendency captures the structuring of ecosystems into trophic levels (33) and is consistent with an overall “spindle” shape to the food web (fewer species in the upper and lower levels and a greater number in the middle) (41). The small number of species at lower trophic levels follow from the practice of aggregating the lowest units of the food web into one or a few nodes broadly labeled “plant,” “detritus,” etc. The consumers of these lowest units have low in-degrees and are in turn consumed by predators of low trophic level (with high out-degrees). The food webs are assortative and *Z* assortative in both the (*out*, *out*) and (*in*, *in*) measures (though in the case of Ythan only slightly); because species at the same trophic level should have similar in- and out-degrees, these results may indicate that species are eating species at the same or similar trophic level—a signature of omnivory (42)—more frequently than expected based on the degree sequence. The effects of *Z*-assortative or -disassortative behavior on linking between hubs range from a <1% increase (Little Rock, in-hub to in-hub) to a 135% increase (Coachella, in-hub to in-hub).

To identify the origin of these patterns, we built two theoretical models for each web (*Materials and Methods*). The *cascade* model assigns each species a random “niche” value and allows species to eat species of lower value with some probability (27). The *niche*

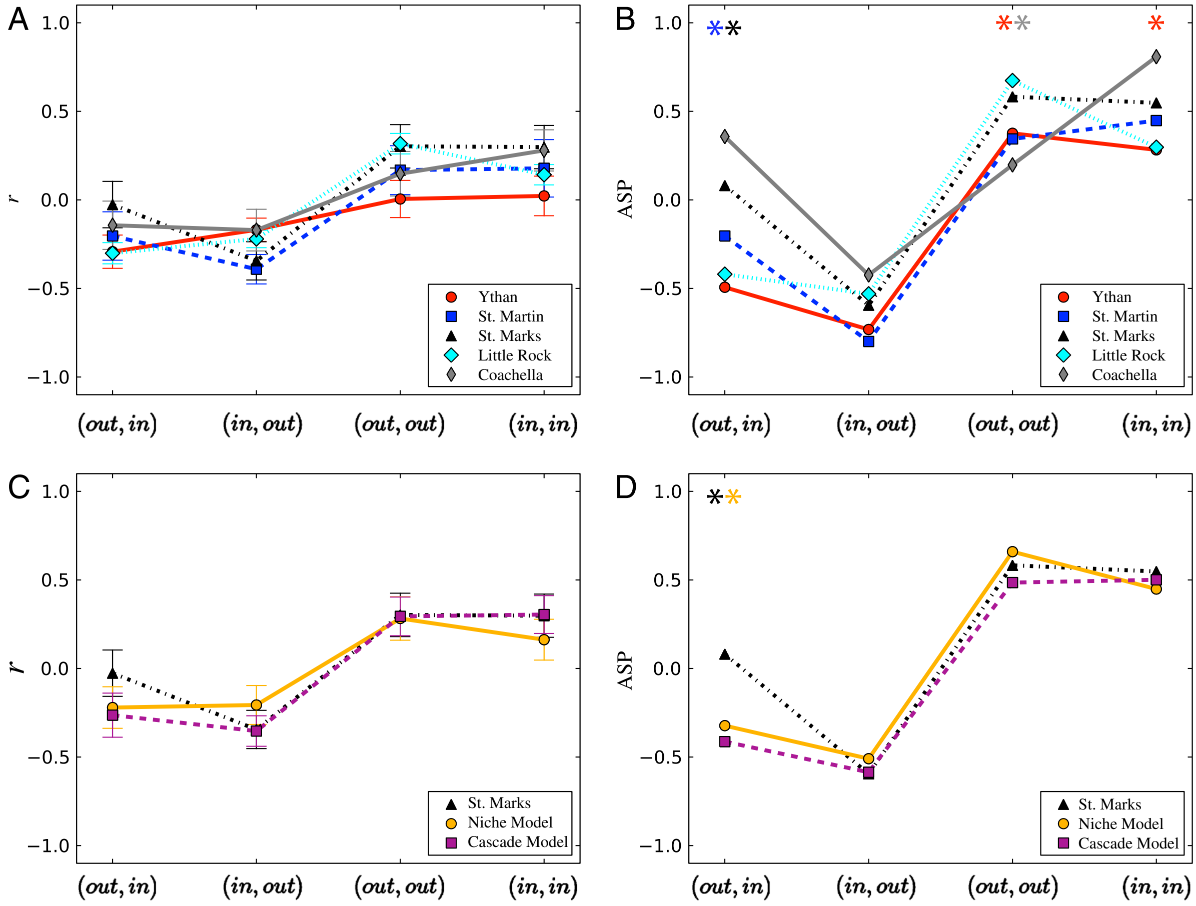

Fig. 3. Simple models largely explain directed assortativity patterns of food webs. A directed edge from A to B indicates that A is eaten by B. (*A*) $r(\alpha,\beta)$ for food webs from several diverse ecosystems (36–40). Errors are estimated by jackknife (11), and we plot $\pm 2\sigma$ error bars. Note the common pattern: disassortative in the first two and assortative in the second two measures. All networks save St. Marks [(*out*, *in*)] and Ythan [(*out*, *out*), (*in*, *in*)] obey this pattern including errors. (*B*) The ASP for these food webs. Controlling for the degree distribution highlights common *Z*-disassortative and *Z*-assortative behaviors in the all measures but (*out*, *in*). In cases where $|Z| < 2$, the corresponding ASP is marked with an appropriately colored asterisk. Only St. Marks (*out*, *in*) has $|Z| < 1$. (*c* and *d*) The cascade and niche models (27) reproduce most common behaviors robustly. Errors and significance levels indicated as above.

model relaxes this rigid hierarchy, permitting cannibalism and the eating of species with higher niche value (27). Fig. 3 *C* and *D* shows the $r(\alpha,\beta)$ and $\text{ASP}(\alpha,\beta)$ for the cascade and niche models of the St. Marks food web (38). The model webs shown are typical of the model and qualitatively reproduce the pattern observed in Fig. 3 *A* and *B*. The *ensemble* of niche model realizations for a given food web, however, displays large variance (see Table S3), favoring the cascade model as more predictable. These results suggest that ordering species along a single niche dimension largely explains the observed patterns in $r(\alpha,\beta)$ and $\text{ASP}(\alpha,\beta)$ for food webs. Neither model, however, typically generates the (*out*, *in*) *Z* assortativity seen in certain food webs.

Word-Adjacency Networks. Finally, we analyze word-adjacency networks, where edges point from each word to any word that immediately follows it in a selected text (14). For example, (for → example). These networks are strongly disassortative for $r(\alpha,\beta)$; see Fig. 4*A*. Fig. 4*B* shows that they are also strongly disassortative in their ASP. The effects on linking between high degree nodes are relatively small, ranging from a decrease of 3.8% (English book, out-hub to out-hub) to a decrease of 15.8% (Japanese book, out-hub to out-hub).

The in- and the out-degree of nodes in these networks are both increasing functions of word frequency (43); thus the correlation between a node's in-degree and out-degree is high ($r_{\text{auto}} > 0.86$). Very high frequency words generally have grammatical function but low "semantic content" (43). While the large r_{auto} guarantees similar values for all four measures, disassortativity could result from at least two possible mechanisms.

Milo et al. propose a bipartite model (*Materials and Methods*), with a few nodes of one type representing high-frequency "grammatical" words and many nodes of a second type representing low-frequency content words; grammatical words must be followed by content words, and vice versa (14). The Bipartite model reproduces the motif pattern of word-adjacency networks and is thus assigned to the same superfamily in this scheme (14). This model generates negative values across all $r(\alpha,\beta)$, as shown in Fig. 4*A*, but these values are too large compared to the real network. When compared to its rewired ensemble, however, the model reproduces the roughly equal, negative $\text{ASP}(\alpha,\beta)$ of the actual networks; see Fig. 4*B*. Thus our measures do not support the classification of the Bipartite model network with the real networks. Alternately, the observed disassortativity could result from a broad word-frequency distribution [Zipf's law (43)]. We scrambled the English text (44) to produce a text with identical word-frequency distribution but no grammatical structure (*Materials and Methods*). The Scrambled text model has $r(\alpha,\beta)$ very close to the empirical values (Fig. 4*A*), but it is *Z* assortative across all measures (Fig. 4*B*), unlike the real-world networks. Neither model yields the relative magnitude of $\text{ASP}(out, in)$ and $\text{ASP}(in, out)$, suggesting that this difference results from genuine linguistic structure.

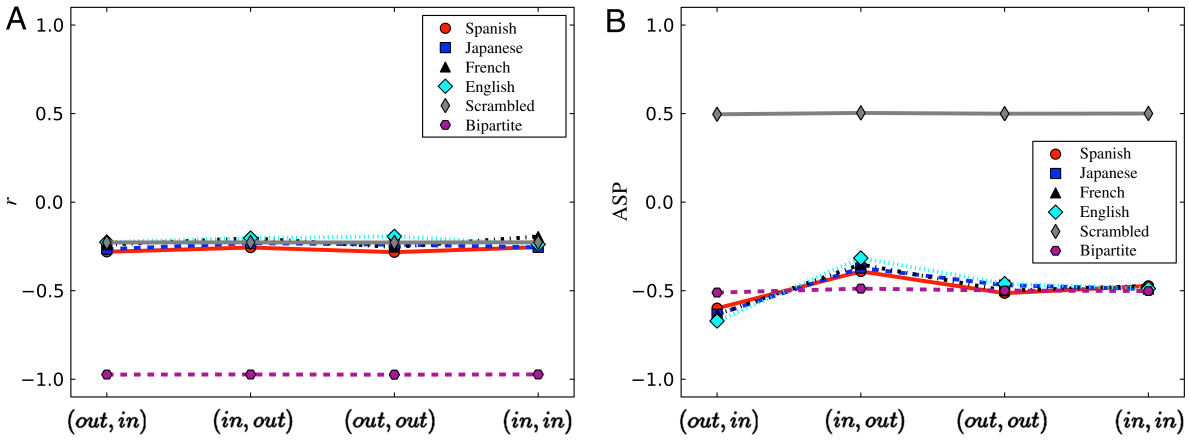

Fig. 4. Simple models cannot explain directed assortativity patterns of word-adjacency networks. A directed edge from word X to word Y indicates that X precedes Y at some point in the text under consideration. (*A*) $r(\alpha, \beta)$ for word-adjacency networks in four languages. The common pattern may result from grammatical structure (Bipartite model) or a broad word-frequency distribution (Scrambled text model). The Bipartite model (14) overestimates the $r(\alpha, \beta)$, as shown in *A*, while the Scrambled text model (45) produces realistic values. Errors in r as estimated by jackknife are smaller than the symbols. (*B*) The ASP for the same networks. The Bipartite model produces realistic values, while the Scrambled text model produces assortative values. The real-world networks are remarkably similar, despite ranging in size over an order of magnitude. All Z scores are highly significant.

Conclusions

Our results demonstrate the fundamental importance of edge direction and the advantages of assortativity—when properly extended—in the analysis of directed networks. Our most basic observation is that directed networks are structurally diverse: Many directed networks are not purely assortative or disassortative, but a mixture of the two. Our measures apply to any directed network, and we expect similar diverse but class-specific mixtures to arise in other directed networks. By comparison with randomized ensembles, we are able to detect statistically significant features such as (*in*, *out*) assortativity in the WWW.

Our measures display common patterns for classes of similar networks (see Figs. S1 and S2) and can be compared to a local analogue, the Triad Significance Profile (TSP). The TSP assigns each possible three-node subgraph (motif) a normalized Z score by comparing the number of appearances of the subgraph in a real-world network to the average number in a randomly rewired ensemble; classes of networks have similar TSPs (14). The measures $r(\alpha, \beta)$, $Z(\alpha, \beta)$, and $\mathrm{ASP}(\alpha, \beta)$ are more computationally tractable and scalable than motif analysis; they also discriminate between networks grouped together by TSP (online/social), while confirming the motif-based classification of word-adjacency networks (14), correctly grouping the online networks (although the political blogs only weakly), and classifying food webs together. As illustrated by all three classes, $r(\alpha, \beta)$ and $\mathrm{ASP}(\alpha, \beta)$ are best used together for exploring the structure of the real-world networks and testing theoretical models.

We tested models for all three network classes. The preferential attachment model of WWW growth (26) does not generate the observed (*in*, *out*) assortativity in the WWW. Neither the Bipartite (14) nor the Scrambled text model of word-adjacency networks generates realistic patterns in both $r(\alpha, \beta)$ and $\mathrm{ASP}(\alpha, \beta)$. We note that creating a mixture of assortative and disassortative behavior is nontrivial. While the WWW growth model fails to do so, both food web models (27) succeed. We suggest that they do so by remaining close to the basic features of the phenomenon. Our measures can be used to test models for any type of directed network and thus validate or falsify the prevailing theoretical understanding.

The straightforward interpretation of directed assortativity leads to a variety of questions: For example, do the overabundant connections between authorities and useful pages in the WWW reflect demands of network navigation, facilitating the spread of user flow—whereas the negative $r(in, out)$ in food webs reflects the opposite tendency to concentrate energy flows at higher trophic levels? Such questions suggest further applications of these concepts to build models better tailored to the reality of asymmetric interactions in complex networks.

Materials and Methods

Defining the Assortativity Measures. Newman (8, 9) defines r in terms of the excess degree, i.e., the degree of the node minus 1. The correlation coefficients are exactly the same if the degree itself is used (8). Identical Z-score results are obtained for any assortativity measure that is related to the Pearson coefficient $r(\alpha, \beta)$ by a linear transformation, e.g., the s metric of Alderson and Li (45); thus when statistical significance is properly measured, it is sufficient to use the Pearson coefficient.

Constructing the Null Model. We sample the ensemble of randomized networks with the same fixed degree sequence (FDS) (13, 24, 25) using a Monte Carlo rewiring algorithm. The algorithm starts with a directed network with a given in- and out-degree sequence $n(k^{\mathrm{in}}, k^{\mathrm{out}})$ and, by randomly swapping directed edges between nodes, samples from the FDS ensemble. If the starting network contains self-edges, we allow them in sampled networks; otherwise, we reject such rewiring steps. We always forbid multiple edges. To assure random sampling, we performed 10^5 edge swaps between samples for most ensembles, 10^6 for WWW and related models, and 10^7 for the Wikipedia network. Before sampling the FDS ensemble, we performed 10 times the number of intersample edge swaps on the starting network to ensure sampling of typical networks. We assume that errors in the ensemble averages are normally distributed and that after i samples the difference between the mean value of an observable up to that point $\langle A\rangle_i = i^{-1}\sum_{j=1}^{i} A_j$ and the final mean $\langle A\rangle$ is less than $bi^{-1/2}$ in absolute value, for some constant b. Plotting the difference as a function of $i^{-1/2}$ and choosing b to contain approximately 90% of the data points gives an estimate of the error in the final mean, reported in Table S2 as σ_r^{rand}.

World Wide Web Growth Model. The growth model for the World Wide Web is taken from ref. 26; we summarize it here in the original notation. This model constructs a directed network approximating the power-law in-degree and out-degree distributions of a target real-world network, $n(k^{\mathrm{in}}) \propto (k^{\mathrm{in}})^{-\nu_{\mathrm{in}}}$ and $n(k^{\mathrm{out}}) \propto (k^{\mathrm{out}})^{-\nu_{\mathrm{out}}}$. The model is parameterized by the number of nodes in the network, N; the average out-degree $\langle k^{\mathrm{out}}\rangle$; and the exponents ν_{in} and ν_{out}. At each step, with probability $p = 1/\langle k^{\mathrm{out}}\rangle$ a new node is born and attaches to an existing target node in the network, chosen with probability (depending on its in-degree i) $\propto A_i = i + \lambda$. Otherwise, with probability $q = 1 - p$, a directed edge is added between two existing nodes, with the source and target nodes selected with probability (depending on the out-degree of the source j and in-degree of the target i) $\propto C(j, i) = (i + \lambda)(j + \mu)$. Choosing λ, μ such that $\nu_{\mathrm{in}} = 2 + p\lambda$ and $\nu_{\mathrm{out}} = 1 + q^{-1} + \mu p q^{-1}$ generates the target exponents. We initialize the model with two unconnected nodes and run until the network has N nodes. We eliminate any multiple edges to yield a simple graph; this does not substantially alter the degree distributions or r values. For the WWW dataset $\nu_{\mathrm{in}} = 2.32$ and $\nu_{\mathrm{out}} = 2.66$. For the three model webs, the exponents are indistinguishable and are $\nu'_{\text{in}} = 2.2 \pm 0.2$ and $\nu'_{\text{out}} = 2.5 \pm 0.2$.

Cascade and Niche Models. The food webs models are taken from ref. 27; we summarize them here in the original notation. Both are parameterized by the number of species in the target food web, N, and the connectance $C = E/N^2$, where E is the number of edges. In the cascade model, every species is assigned a random niche value chosen uniformly from [0, 1]. With probability $P = 2CN/(N-1)$, a species will consume a species with a lower niche value. In the niche model, every species i is assigned a random niche value n_i as before; the species of smallest niche value is assigned to be the "basal species" (27). All other species consume every species falling within some range r_i. The center of the range c_i is chosen uniformly from $[0.5r_i, n_i]$. The range r_i is chosen such that the expected connectance is that of the real-world web by setting $r_i = n_i x_i$, where x_i is drawn from a beta distribution $f(x_i|1,\beta) = \beta(1-x_i)^{\beta-1}$, $0 < x_i < 1$ with expected value $E(x_i) = 1/(1+\beta) = 2C$. Both models yield the connectance of the real-world food web, on average. We do not check for disconnected or trophically identical species (species having identical in- and out-neighbors), as these are quite rare. For each food web, we generated 500 cascade model and niche model networks with E within 5% of the original food web. To identify typical networks (shown in the paper and Tables S1 and S2), we selected the model network with the smallest Euclidean distance to the ensemble average values of $r(\alpha,\beta)$. The standard deviations in each ensemble are shown in Table S3.

Bipartite and Scrambled Text Models for Word-Adjacency Networks. The Bipartite model (14) assumes that there are two categories of words: a few high frequency grammatical words and many low-frequency content words. Words of the first type alternate with words of the second type, resulting in a bipartite word-adjacency network. We build the model with $N_{\text{gram}} = 10$ and $N_{\text{cont}} = 1,000$. For all pairs of grammatical and content words we draw a random number x. If $x < p = .06$, we put an edge from the grammatical word to the content word; if $p < x < 2p$ we put an edge from the content word to the grammatical word; and if $2p < x < 2p + q$ for $q = .003$ we put an edge going each way. The values of p, q are taken from ref. 14. We constructed the Scrambled Text Model by randomly scrambling the order of the words in the underlying text for one of the word-adjacency networks [English, *On the Origin of Species* by Charles Darwin (45)]. The scrambling destroys any syntactic structure, although the high frequency of articles, prepositions, etc., remains. The assortativity across all ASP(α,β) of networks generated from the scrambled text is subtle. The high correlation between the in- and out-degrees of a node guarantees that all values will be similar. In the scrambled text, high frequency (high degree) words are more likely to follow one another. But since multiple links are disallowed, rewiring, on average, destroys links between high degree nodes, making the ensemble less assortative than the Scrambled Text word-adjacency network, and making all ASP(α,β) assortative.

ACKNOWLEDGMENTS. The authors warmly thank E. A. Cartmill for her enormously helpful and detailed comments on the manuscript, K. Brown, M. Cartmill, J. Davidsen, and Seung-Woo Son for their thoughtful reading of the manuscript, the reviewers for their useful comments, and Juyong Park for an insightful discussion. J.G.F. and P.G. acknowledge the support of iCORE. This work was funded in part by NSERC.

1. Barabási A-L (2009) Scale-free networks: A decade and beyond. *Science* 235:412–413.
2. Newman MEJ (2003) The structure and function of complex networks. *SIAM Rev* 45:167–256.
3. Barabási A-L, Albert R (1999) Emergence of scaling in random networks. *Science* 286:509–512.
4. Barabási A-L, Oltvai Z (2004) Network biology: Understanding the cell's functional organization. *Nat Rev Genet* 5:101–113.
5. Watts DJ, Strogatz SH (1998) Collective dynamics of 'small-world' networks. *Nature* 393:440–442.
6. Albert R, Jeong H, Barabási A-L (2000) Error and attack tolerance of complex networks. *Nature* 406:378–382.
7. Ravasz E, Somera AL, Mongru DA, Oltvai ZN, Barabási A-L (2002) Hierarchical organization of modularity in metabolic networks. *Science* 297:1551–1555.
8. Newman MEJ (2002) Assortative mixing in networks. *Phys Rev Lett* 89:208701.
9. Newman MEJ (2003) Mixing patterns in networks. *Phys Rev E* 67:026126.
10. Eubank S, et al. (2004) Modeling disease outbreaks in realistic urban social networks. *Nature* 429:180–184.
11. Maslov S, Sneppen K (2002) Specificity and stability in topology of protein networks. *Science* 296:910–913.
12. Hu H-B, Wang X-F (2009) Disassortative mixing in online social networks. *EPL-Europhys Lett* 86:18003.
13. Milo R, et al. (2002) Network motifs: Simple building blocks of complex networks. *Science* 298:824–827.
14. Milo R, et al. (2004) Superfamilies of evolved and designed networks. *Science* 303: 1538–1542.
15. Albert R, Jeong H, Barabási A-L (1999) Diameter of the World-Wide Web. *Nature* 401:130–131.
16. Broder A, et al. (2000) Graph structure in the Web. *Comput Netw* 33:309–320.
17. Leicht EA, Newman MEJ (2008) Community structure in directed networks. *Phys Rev Lett* 100:118703.
18. Guimera R, Sales-Pardo M, Amaral LAN (2007) Module identification in bipartite and directed networks. *Phys Rev E* 76:036102.
19. Kim Y, Son S-W, Jeong H (2010) Link Rank: Finding communities in directed networks. *Phys Rev E* 81:016103.
20. Karrer B, Newman MEJ (2009) Random graph models for directed acyclic networks. *Phys Rev E* 80:046110.
21. Davidsen J, Grassberger P, Paczuski M (2008) Networks of recurrent events, a theory of records, and an application to finding causal signatures in seismicity. *Phys Rev E* 77: 066104.
22. Serrano MA, Maguitman AG, Boguñá M, Fortunato S, Vespignani A (2007) Decoding the structure of the WWW: A comparative analysis of Web crawls. *ACM Trans Web* 1(2):10.
23. Flack JC, Girvan M, de Waal FBM, Krakauer DC (2006) Policing stabilizes construction of social niches in primates. *Nature* 439:426–429.
24. Foster JG, Foster DV, Grassberger P, Paczuski M (2007) Link and subgraph likelihoods in random undirected networks with fixed and partially fixed degree sequences. *Phys Rev E* 76:036107.
25. Maslov S, Sneppen K, Zaliznyak A (2004) Detection of topological patterns in complex networks: Correlation profile of the internet. *Physica A* 333:529–540.
26. Krapivsky PL, Rodgers GJ, Redner S (2001) Degree distributions of growing networks. *Phys Rev Lett* 86:5401.
27. Williams RJ, Martinez ND (2000) Simple rules yield complex food webs. *Nature* 404:180–183.
28. Gleich D (2009) Available at http://www.cise.ufl.edu/research/sparse/matrices/Gleich/index.html.
29. Adamic LA, Glance N (2005) The political blogosphere and the 2004 U.S. election. *Proceedings of the 3rd International Workshop on Link Discovery* (ACM, Chicago), pp 36–44.
30. Zlatić V, Božičević M, Štefančić H, Domazet M (2006) Wikipedias: Collaborative web-based encyclopedias as complex networks. *Phys Rev E* 74:016115.
31. Kleinberg JM (1999) Authoritative sources in a hyperlinked environment. *J Assoc Comput Mach* 46:604–632.
32. Capocci A, et al. (2006) Preferential attachment in the growth of social networks: The internet encyclopedia Wikipedia. *Phys Rev E* 74:036116.
33. Williams R, Martinez ND (2004) Limits to trophic levels and omnivory in complex food webs: Theory and data. *Am Nat* 163:458–468.
34. Camacho J, Guimerà R, Amaral LAN (2002) Robust patterns in food web structure. *Phys Rev Lett* 88:228102.
35. Dunne JA, Williams RJ, Martinez ND (2002) Food-web structure and network theory: The role of connectance and size. *Proc Natl Acad Sci USA* 99:12917–12922.
36. Polis GA (1991) Complex trophic interactions in deserts: An empirical critique of food-web theory. *Am Nat* 183:123–155.
37. Martinez ND (1991) Artifacts or attributes? Effects of resolution on the Little Rock Lake food web. *Ecol Monogr* 61:367–392.
38. Christian RR, Luczkovich JJ (1999) Organizing and understanding a winter's seagrass foodweb network through effective trophic levels. *Ecol Model* 117:99–124.
39. Goldwasser L, Roughgarden J (1993) Construction and analysis of a large Caribbean food web. *Ecology* 74:1216–1233.
40. Huxham M, Beaney S, Raffaeli D (1996) Do parasites reduce the chances of triangulation in a real food web?. *Oikos* 76:284–300.
41. Bascompte J, Melián CJ (2005) Simple trophic modules for complex food webs. *Ecology* 86:2868–2873.
42. Stouffer DB, Camacho J, Jiang W, Amaral LAN (2007) Evidence for the existence of a robust pattern of prey selection in food webs. *Proc R Soc London Ser B* 274:1931–1940.
43. Ferrer i Cancho R, Solé RV (2001) The small world of human language. *Proc R Soc London Ser B* 268:2261–2265.
44. Darwin C (1859) *On the Origin of Species*, 6th Ed. Available at http://www.gutenberg.org/etext/2009.
45. Alderson D, Li L (2007) Diversity of graphs with highly variable connectivity. *Phys Rev E* 75:046102.