

Visualization Methods for Longitudinal Social Networks and Stochastic Actor-Oriented Modeling[★]

Ulrik Brandes, Natalie Indlekofer, and Martin Mader

Department of Computer & Information Science, University of Konstanz

Abstract

As a consequence of the rising interest in longitudinal social networks and their analysis, there is also an increasing demand for tools to visualize them. We argue that similar adaptations of state-of-the-art graph-drawing methods can be used to visualize both, longitudinal networks and predictions of stochastic actor-oriented models (SAOM), the most prominent approach for analyzing such networks. The proposed methods are illustrated on a longitudinal network of acquaintanceship among university freshmen.

Key words: longitudinal social networks, dynamic network visualization, stochastic actor-oriented models, model fit

1 Introduction

With the ever-increasing availability of time-varying data and the diffusion of advanced modeling methods, research on longitudinal network analysis is widening tremendously. This is especially true for the interest in dependencies

[★] Research partially supported by DFG under grants Br 2158/3-2 (as part of ESF project 05-ECRP-026) and GRK 1042. Krists Boitmanis (Institute of Mathematics and Computer Science, University of Latvia) took part in early phases of this research. We are indebted to Britta Renner and Manja Vollmann (Department of Psychology, University of Konstanz) for providing us with the network data used for illustration. Editor Tom Snijders and the anonymous reviewers provided very helpful comments that led to significant improvement of the paper.

Email addresses: Ulrik.Brandes@uni-konstanz.de (Ulrik Brandes),
Natalie.Indlekofer@uni-konstanz.de (Natalie Indlekofer),
Martin.Mader@uni-konstanz.de (Martin Mader).

among tie dynamics and actor attributes, and more concretely the co-evolution of networks and behavior. Since the temporal dimension constitutes an additional, qualitatively different level of complexity, the demands on visualization tools are even higher than they are anyway in static network analysis (Bender-deMoll and McFarland, 2006).

Social network visualization is a field of growing interest in itself (Klov Dahl, 1981; Freeman, 2000; Brandes, Kenis, and Raab, 2006), and partly so because very different approaches are suitable for specific use cases (Brandes, Freeman, and Wagner, in preparation). For the present case, we assume to have longitudinal network data given in the form of panel data, i.e., as a time-ordered sequence of interrelated network observations that possibly differ in actor composition, structure, and attributes. In social sciences, this is the most common form of longitudinal network data today, and often due to data collection in waves or aggregation of dyadic events over time intervals. The latter is frequently done to allow for the application of the same methods that are common for static networks, and various forms of aggregation are described in Bender-deMoll and McFarland (2006).

We here define our problem area as that of visualizing a given sequence of snapshots of an evolving social network (rather than, say, an unordered collection of networks, an event stream, or a process taking place on a network). The task is further limited to producing a corresponding sequence of diagrams which may or may not serve as the basis of an animation (rather than, say, a merged view of the entire evolution). The characterizing trade-off in this situation is between the individual quality of each snapshot and the persistence of features over the sequence (Brandes and Wagner, 1997). In other words, each diagram should be a good representation of the corresponding cross-sectional network, and at the same time, a mental map of the structure should be preserved as much as possible to relate the individual frames with less cognitive effort (Misue, Eades, Lai, and Sugiyama, 1995).

The motivation behind this task is to facilitate visual exploration of longitudinal network data in a generic way. By using a specific methodology, however, analysts take a specific perspective that is generally in need of targeted visualization designs. As a concrete example, we here focus on the most prominent approach to longitudinal social network analysis, stochastic actor-oriented modeling (Snijders 2005; Snijders, van de Bunt, and Steglich 2010b), and show that with little adaptation, the same visualization techniques can be applied to reveal such a model’s predictions and interrelate them with the actual observation. Our approach is likely to generalize to other models as well.

The remainder of this article is organized into three main parts. Since the crucial technical challenge in network visualization is to find a suitable layout

for the underlying graph structure, we start by providing background on layout algorithms for static graphs in Section 2, and outline a method known as *stress minimization* that is central to our approaches. In Section 3, we review the dynamic graph drawing problem, and propose specific instantiations of stress minimization designed for visual exploration of dynamic graphs. In the third part, we introduce two targeted visualization approaches for stochastic actor-oriented models (SAOMs) in Section 5, after recalling the formal basics of SAOMs in Section 4. The first of these approaches shall help assess congruence of simulations and observations w.r.t. their underlying graph structure, whereas the second one is to point to inhomogeneities across actors, if any. We conclude with a brief discussion that includes directions for future work.

1.1 Running Example

We use a longitudinal network of acquaintanceship among university students as a running example. The data is courtesy of Britta Renner and Manja Vollmann (Department of Psychology, University of Konstanz) and was collected in 15 waves between October 2008 and February 2009.

Students provided, among many other things, their current perceived level of acquaintanceship with each other on a scale from 1 (lowest) to 7 (highest). We dichotomized each observation using 5 as a threshold. Of the 78 freshmen majoring in Psychology, only nine did not participate in an initial screening, never answered any questionnaire, or never made a nomination resp. were never nominated at a level above the threshold.

The example networks thus consist of acquaintanceship nominations among 69 students (18 male, 51 female) that form a connected component when aggregated over all waves.

The data constitutes a realistic scenario in which our methods may be applied, but is used here solely for illustrative purposes. No attempt at justifying models or drawing conclusions will be made.

2 Graph Drawing Methods for Static General Graphs

Social network visualization can draw on two major streams of research, information visualization of networks (Herman, Melançon, and Marshall, 2000) and graph drawing (Di Battista, Eades, Tamassia, and Tollis, 1999; Kaufmann and Wagner, 2001). Roughly speaking, the focus in information visualization is on visualization design, navigation, and interactivity, whereas properties and

95 construction of geometric representations are more central to graph drawing.

96 We here restrict our scope to the most common graphical representation for so-
97 cial networks, node-link diagrams (referred to as *sociograms* in [Moreno 1953](#)),
98 in which actor-representing vertices are depicted as points (or, more precisely,
99 graphical elements described by a single position), and tie-representing edges
100 are depicted as lines linking their endpoints. We will not, in general, make
101 the distinction between actors, nodes, vertices, and points, and between ties,
102 links, edges, and lines.

103 The central task in creating node-link diagrams is to determine positions for
104 its elements, referred to as the diagram’s *layout* in the following. This is be-
105 cause positional differences are the most accurately perceived graphical at-
106 tributes ([Cleveland and McGill, 1984](#)), and layout with complex dependencies
107 is the most challenging problem algorithmically. If the layout is of low quality,
108 even the best graphical design (in terms of using other graphical attributes
109 such as shape, color, size, etc.) or interaction mechanisms can only attenuate
110 the problems of poor legibility and interpretation artifacts.

111 While graph structure is represented completely in plain node-link diagrams,
112 the other attributes of a network can be incorporated by varying graphical
113 attributes as mentioned above. Clearly, these choices are more dependent on
114 the data and context, and in general easier to implement.

115 2.1 Graph Layout

116 In addition to distinct vertex positions to avoid ambiguity, the following ob-
117 jectives are commonly considered relevant for application-independent lay-
118 out ([Bertin, 1983](#); [Purchase, Cohen, and James, 1997](#)).

- 119 • Edges should be of more or less the same length.
- 120 • Vertices should be distributed well over the drawing area.
- 121 • The number of meaningless edge crossings should be kept small.
- 122 • Symmetries in graph structure should be visible in geometric symmetries.

123 For specific applications and purposes, there may be many more criteria to
124 observe. For most of them, optimization is computationally intractable even
125 in isolation, at least for general graphs. Since, in addition, the various criteria
126 are frequently contradictory, general-purpose graph-drawing algorithms are
127 usually heuristic in nature.

128 Even though social networks exhibit some general tendencies such as sparse-
129 ness and local clustering, they do not constitute a formally boundable class of
130 graphs that allows for specific optimization algorithms. Due to their general

applicability, conceptual simplicity, wide availability, and ability to produce satisfactory results in general, the most popular class of methods used for social network layout are *force-directed* or *energy-based* methods (Brandes, 2001), colloquially known as *spring embedders* (Eades, 1984).

The most widely available, and often only, layout algorithm in common software tools for social network analysis is the spring embedder variant of Fruchterman and Reingold (1991). It is a force-directed method in which a graph is likened to a physical system of repelling objects (the vertices) and springs of a given length (the edges) binding adjacent vertices together. Vertices are iteratively repositioned based on the forces exerted on them, so that the system moves toward a force equilibrium. The approach is easy to implement and yields acceptable results for small graphs, and it can be tuned for specific purposes by introducing additional or alternative forces.

There is, however, clear experimental evidence (Brandes and Pich, 2009) that this and related force-directed methods do not scale well to larger graphs, both in terms of quality and efficiency. It is almost ironic that a current variant of the earliest computer-implemented method for drawing social networks (Kruskal and Seery, 1980, already applied in the late 60s), turns out to be far superior.

This favorable approach, known as *stress minimization*, is an instance of a family of dimension-reduction methods referred to as *multidimensional scaling* (see, e.g., Cox and Cox 2001). It is based on an objective function called *stress* (Kruskal and Wish, 1978) and was re-popularized in graph drawing by Gansner, Koren, and North (2004). Details are given next, but it should be noted that the same objective function was also used in the spring embedder of Kamada and Kawai (1988), although with an inferior minimization method.

2.2 Stress Minimization

We next describe the workhorse of our approach, stress minimization, in detail. Let $G = (V, E)$ be an undirected graph defined by a set V of n vertices, and a set E of m edges. Given a matrix D of vertex *dissimilarities* δ_{ij} , $i, j \in V$, the purpose of stress minimization is to determine positions $p_i = \langle x_i, y_i \rangle \in \mathbb{R}^2$ for every vertex $i \in V$ such that the Euclidean distances in the plane resemble the given dissimilarities as closely as possible, i.e.,

$$\delta_{ij} \approx \|p_i - p_j\| ,$$

where $\|\cdot\|$ denotes the Euclidean norm. For any given layout $P = (p_1, \dots, p_n)$ this is quantified using a parameterized *stress function* $\text{stress}(P)$,

$$\text{stress}(P) = \sum_{i < j} \omega_{ij} (\delta_{ij} - \|p_i - p_j\|)^2, \quad (1)$$

where $W = (\omega_{ij})_{i,j \in V}$ is a weight matrix whose entries determine the contribution of each pair $i, j \in V$. Since stress is defined as the weighted sum of squared distance-representation errors, the objective is to find a layout of minimum stress.

In graph drawing, graph-theoretic distances (i.e., lengths of shortest paths) are a plausible choice for dissimilarities (Kamada and Kawai, 1988; Gansner, Koren, and North, 2004). The stress term of each dyad then corresponds to the squared error of representing a shortest path as a straight line with unit length edges. Because these distances are clearly not realizable for any non-trivial graph, weights $\omega_{ij} = \delta_{ij}^{-2}$ discount representation errors for distant pairs, thus emphasizing local accuracy.

Since no closed form is known to compute a layout with minimum stress directly, the approach appears to share some of the drawbacks of other force-directed methods. In particular, iterative stress reduction in general only yields a local minimum which may be far from an optimal layout. The experiments of Brandes and Pich (2009) suggest, however, that low-stress layouts can be obtained routinely and efficiently using a two-step process: In the first step, an initial layout is determined using *classical scaling* (Torgerson, 1952), the initial, spectral-decomposition variant of multidimensional scaling that has an essentially unique solution in which large distances are represented well, and that can be approximated very quickly (Brandes and Pich, 2007). In the second step, the representation of small distances is improved by iteratively and monotonically reducing the stress using an optimization technique called majorization (de Leeuw, 1977) that can be implemented by a simple, localized process moving one vertex at a time and considering only a sparse version of the stress function. See Gansner, Koren, and North (2004) for details.

The kind of results obtained is illustrated in Figure 1, using two observations from our running example.

We will not elaborate on algorithmic issues any further, but would like to emphasize that stress minimization is reasonably easy to implement, and yields much better results than other force-directed methods (Brandes and Pich, 2009), including multi-level approaches (Walshaw, 2001; Gajer, Goodrich, and Kobourov, 2004; Hachul and Jünger, 2004). At the same time the method is very flexible, because there are several degrees of freedom that allow for sophisticated layout modeling. We may alter the given network’s structure resp. stress terms, admissible positions, dissimilarities, and weights to incorporate hard and weak constraints (see, e.g., Dwyer, Marriott, and Wybrow 2008; Gansner and Hu 2008; Brandes and Pich 2011).

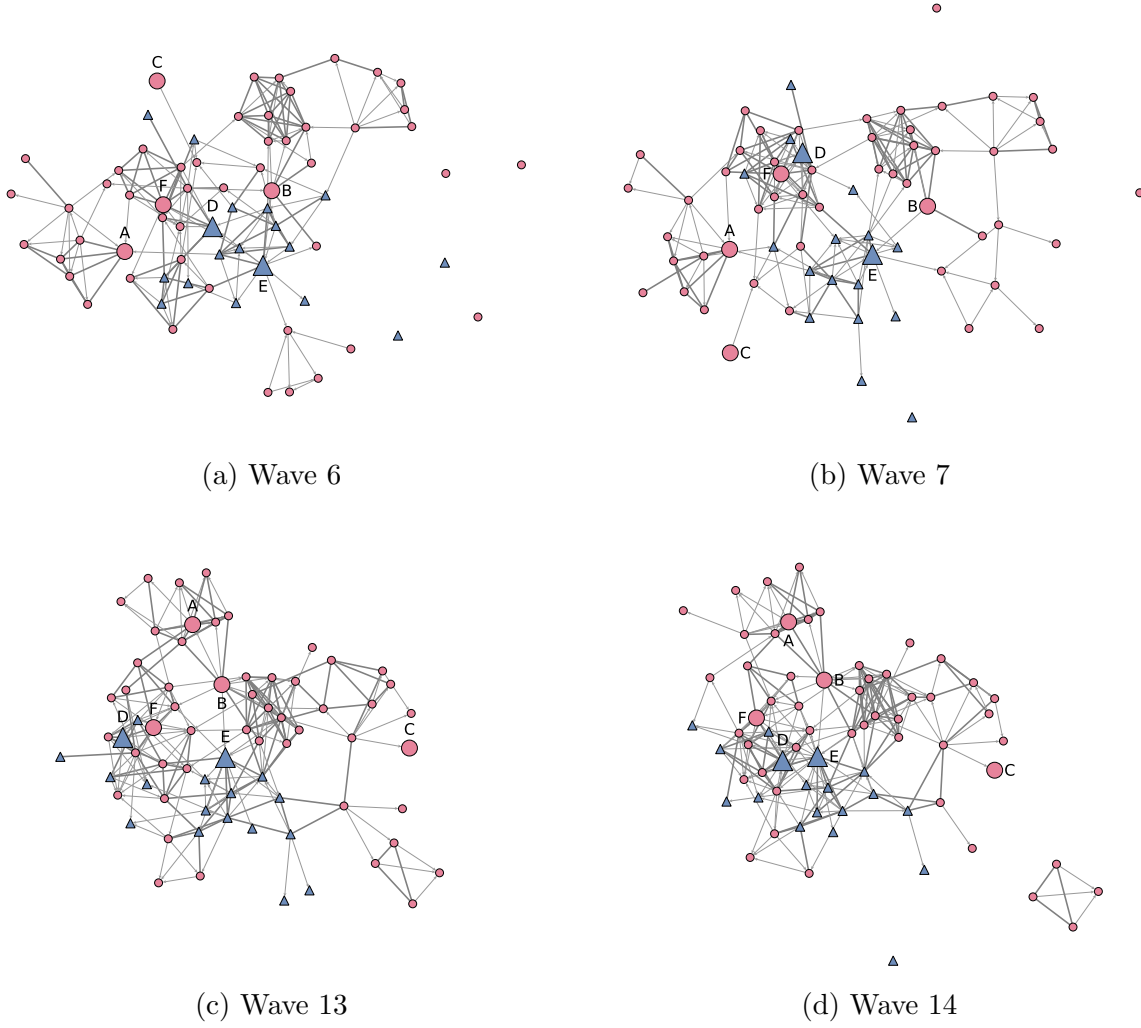


Fig. 1. Four observations of an acquaintanceship network laid out using stress minimization. Red circles represent female students, whereas blue triangles represent males. If an acquaintanceship tie is reciprocated, it is represented by a thicker line segment, otherwise an arrow indicates the direction of nomination. Some vertices are indexed for later reference.

Indeed, this will be our approach for both, dynamic layout and visual model assessment, as described below.

3 Dynamic Network Visualization

While producing layouts of single networks can already be challenging, determining coherent layouts for several observations of a longitudinal network is even more difficult (Branke, 2001). This is because, in addition to the criterion of faithfully displaying structural properties, the sequence of layouts should




Fig. 2. Knowledge of future changes can inform the choice between otherwise equally good layouts.

convey the evolution of the network.

The latter is facilitated by maintaining an observer’s *mental map* (Misue, Eades, Lai, and Sugiyama, 1995; Saffrey and Purchase, 2008; Purchase and Samra, 2008). The goal is to ease comparison between a state and its predecessor by largely retaining those parts of a drawing in which little structural change occurred, so that the time a viewer spends on familiarizing with the drawing is not wasted. Since layout algorithms for the static, single-graph case are based solely on the current graph’s structure, they have to be modified to strike the right balance between structural changes and positional stability.

Two scenarios need to be distinguished, because they differ in the amount of information available when a layout is determined. Either, the entire sequence of graphs is known in advance, i.e., before any layout is required (generally referred to as *offline scenario*, graphs are given together), or future graphs are unknown at the time that an intermediate state is to be laid out (*online scenario*, graphs are presented to the algorithm one at a time). In an online scenario, it is more difficult to maintain dynamic stability, because future changes are not known at the time that vertices have to be placed.

Let us illustrate this phenomenon by the small example in Figure 2 consisting of a 4-star in the first observation, and an additional tie between the two gray vertices in second observation. In an online scenario, there is no preference for any of the two layouts of the first observation, whereas in an offline scenario we can actually choose the one that will lead to better quality and less movement.

We will take advantage of the fact that the sequence of graphs that constitute a longitudinal social network is usually known in advance, and therefore assume to be in an offline scenario for the remainder of this paper.

The simplest (and most common) approaches to take stability into account are based on variants of the spring embedders of Fruchterman and Reingold

(1991) or Kamada and Kawai (1988), in which the iterative computation for each graph in the sequence is initialized with the preceding layout (Bender-deMoll and McFarland, 2006; Huang, Eades, and Wang, 1998; Geipel, 2007; Groh, Hanstein, and Wörndl, 2009). An implicit assumption is that consecutive graphs are similar in general, so that the initial layout is not too far from a locally optimal one. The method is therefore easy to implement, more efficient than computing a layout from scratch, and applicable in both on- and offline scenarios. Note also that both algorithms repeatedly move one vertex at a time to improve the current layout, thereby creating trajectories from the initial to the final position. These may be useful for animation. In practice, however, linear or sinusoidal coordinate interpolation offer better control and are more commonly used to determine such trajectories. More sophisticated transition methods are described, for instance, in Friedrich and Eades 2002; Friedrich and Houle 2002; Nesbitt and Friedrich 2002.

In spite of the convenient properties enumerated so far, the approach is rather problematic, because it does not address stability in a controlled way and may hence result in excessive and unnecessary movement of vertices. Moreover, by failing to make use of the existing knowledge about the future in offline scenarios, the approach is biased towards earlier configurations, and therefore prone to suffer from poor local minima precisely as already illustrated in Figure 2. Hence, layout quality tends to degrade over the course of the sequence.

Among the first to address stability directly were Böhringer and Paulisch (1990), and North (1996) provides a generic problem statement. The trade-off between readability and stability is formalized in Brandes and Wagner (1997) and a similar principle for offline scenarios is proposed in Diehl and Görg (2002). Current proposals generally take one of the following three approaches.

Aggregation. All graphs in the sequence are aggregated into a single graph that has one vertex for each actor. The position of each individual vertex instance in the sequence is determined from a layout of the aggregated graph. This approach is used, e.g., in Brandes and Corman (2003); Dwyer and Gallagher (2004); Moody, McFarland, and Bender-deMoll (2005).

Linking. All graphs in the sequence are combined into a single graph that has one vertex for each occurrence of an actor, and an edge is created between vertices representing the same actor in consecutive graphs. A layout of this graph directly yields positions for all vertex instances in the sequence. This approach is used, e.g., in Dwyer and Gallagher (2004); Erten, Kobourov, Le, and Navabi (2004); Dwyer, Hong, Koschützki, Schreiber, and Xu (2006).

Anchoring. Using auxiliary edges, vertices are connected to immobile copies fixed to a desired location which may be, for instance, the previous position in an online scenario, or a reference position in an offline scenario. This approach is used, e.g., in Lyons, Meijer, and Rappaport (1998); Brandes and Wagner (1997); Frishman and Tal (2008).

280 Since the method advocated here is a combination of aggregation and an-
 281 choring, we describe in more detail how these can be instantiated in a stress-
 282 minimization framework.

283 3.1 Aggregation

284 Maximum stability is obtained when a vertex maintains its position through-
 285 out the entire sequence of diagrams. This is called the *flip book* approach in
 286 Moody, McFarland, and Bender-deMoll (2005). We argue that suitable posi-
 287 tions can be obtained by applying stress minimization to an aggregation of
 288 the input sequence.

289 Given a sequence $G^{(1)} = (V, E^{(1)}), \dots, G^{(T)} = (V, E^{(T)})$ of T graphs with
 290 corresponding dissimilarities $D^{(t)}$, $1 \leq t \leq T$, we are to determine layouts
 291 $P^{(1)}, \dots, P^{(T)}$ such that a vertex is placed at the same position throughout.
 292 In other words, we are looking for one layout P for the vertices in V and let
 293 $P^{(t)} = P$ at all times $t = 1, \dots, T$.

294 We combine all shortest path information into a single stress function using
 295 the mean shortest path distance per dyad, and use the same position variables
 296 p_i for each instance $i^{(t)}$ of the same actor i ,

$$\text{stress}(P) = \sum_{i < j} \omega_{ij} \left(\bar{\delta}_{ij} - \|p_i - p_j\| \right)^2, \quad (2)$$

297 where $\bar{D} = (\bar{\delta}_{ij})_{i,j \in V}$, $\bar{\delta}_{ij} := \frac{1}{T} \sum_{t=1}^T \delta_{ij}^{(t)}$, contains the mean shortest-path dis-
 298 tances, and

$$\omega_{ij} = \frac{1}{\bar{\delta}_{ij}^2} \cdot \frac{1}{1 + \text{VAR}(\delta_{ij})},$$

299 where $\text{VAR}(\delta_{ij}) := \frac{1}{T} \sum_{t=1}^T \left(\delta_{ij}^{(t)} - \bar{\delta}_{ij} \right)^2$ is the variance of distances within a
 300 dyad across all observations. Thus, representation accuracy of dyads that are
 301 connected via short paths most of the time is emphasized. By additionally
 302 scaling with the variance, priority is given to structures that are relatively
 303 stable throughout the sequence.

304 Since, technically, the aggregate stress function consists of the same type of
 305 terms as before, we can also use the same algorithms for its minimization:
 306 Layout computation is initialized by classical scaling of mean distances in the
 307 aggregated graph, i.e., a graph with all vertices and an edge between two




Fig. 3. Global layout obtained by aggregation of all stress terms. An edge is drawn if two vertices are adjacent in any graph of the input sequence.

vertices if this edge exists in any of the observations. Subsequently, $\text{stress}(P)$ is reduced via majorization.

We require that the aggregated graph is connected, otherwise the data is split into connected components, and layouts are calculated separately for each of them. Note that it may still be the case that some individual networks of a sequence contain multiple disconnected components, resulting in infinite distances and undefined weights. Being in an offline scenario, other observations can be used to fill in gaps as follows. An infinite distance in a dyad is replaced by interpolating between the two finite distances observed previously and next for this dyad, and by adding a small constant, say 1, to emphasize temporary disconnectedness. Let t_{ij}^- be the most recent and t_{ij}^+ be the next observation in which actors i and j have finite distance $\delta_{ij}^{(t_{ij}^-)}$ and $\delta_{ij}^{(t_{ij}^+)}$. Then the interpolated distance for dyad $\{i, j\}$ is

$$\delta_{ij}^{(t)} = (1 - \beta_{ij}^{(t)}) \delta_{ij}^{(t_{ij}^-)} + \beta_{ij}^{(t)} \delta_{ij}^{(t_{ij}^+)} + 1 ,$$

where $\beta_{ij}^{(t)} = (t - t_{ij}^-)/(t_{ij}^+ - t_{ij}^-)$. In the special cases that there is no last or next finite distance, we do not interpolate, but use the one existing finite distances distance plus the same constant. Although unlikely, it may happen that a dyad has infinite distances at all times, even if the aggregated graph is connected. In this case, a sufficiently large distance Δ with a small weight Ω is used. We suggest $\Delta\sqrt{n}$ and $\Omega = 1/n$ with n the number of vertices. This is in analogy to the height and width of an equally spaced grid with n points.

The resulting aggregate layout for our running example is given in Figure 3, and Figure 4 contains samples from the corresponding sequence of diagrams.

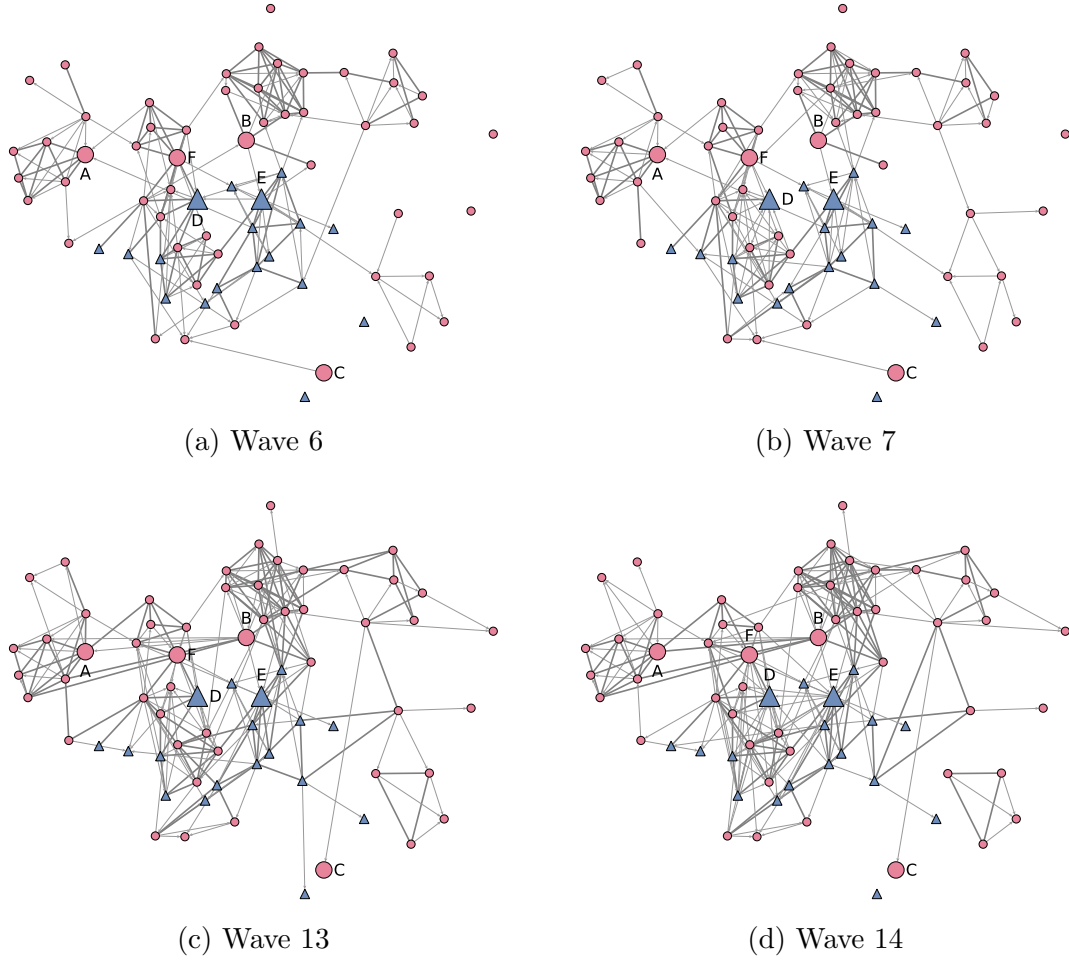


Fig. 4. Four observations of the acquaintanceship network using aggregation layout. While stability is perfect, there are inherent drawbacks as exemplified by labeled vertices (see main text).

The relatively stable major groups, i.e., the female groups around Actor A, and above Actor B, the male group around Actor E, and the mixed group involving Actor D and Actor F, are discernible. Due to perfect stability, occurrences of the same actor are easily indentified in each of the subfigures of Figure 4. However, a comparison with the layouts obtained by regular stress minimization shown in Figure 1 reveals that this extreme stability comes at the price of less desirable individual layouts:

- Subconfigurations may appear to be placed oddly at times. An example is the group around Actor A in Wave 13 and 14, where it would be expected to lie closer to the group to which it connects through Actor B. Its position leads to unnecessary long links representing the strong ties to Actor B that have been established by then. These also create much clutter in the group belonging to Actor F. See Figure 1d for comparison.
- The almost bipartite structure of the group around Actor A that could be

344 observed in Figure 1d is hardly recognizable.

- 345 • Actor C starts out being connected to an actor at the very bottom, but
- 346 then links up with an actor further up. Since this new neighbor is far away
- 347 from the first neighbor, any reasonable constant position can only result in
- 348 either an awkward or a long link connecting Actor C.
- 349 • Actor D and Actor E exhibit a very volatile connection, resulting in cor-
- 350 responding changes of geometric distances between them in Figure 1. Al-
- 351 though this volatility can be observed in the aggregate layout, it might be
- 352 easily missed if the actors were located in denser regions.

353 These qualitative examples are corroborated by quantitative measurement; for
 354 instances, stress of the layout of Wave 6 and Wave 7 as measured by Equa-
 355 tion 1 is almost twice as high as in the corresponding layouts of Figure 1, and
 356 more than thrice as high for Wave 13 and Wave 14. Because of the complex de-
 357 pendencies between structural characteristics, change, and relative positions,
 358 a comprehensive quantitative evaluation is beyond the scope of this work,
 359 though.

360 The above examples nevertheless illustrate that fixed positions are helpful in
 361 building a mental map of the overall configuration, but that it is desirable to
 362 allow for at least small deviations to represent better the specific configurations
 363 at individual time points.

364 3.2 *Anchoring*

365 The main idea of the Bayesian approach to online dynamic graph drawing
 366 ([Brandes and Wagner, 1997](#)) is an explicit modeling of the trade-off between
 367 layout quality as measured by an objective function, and layout stability with
 368 respect to the previous drawing as measured by a difference metric ([Bridge-](#)
 369 [man and Tamassia, 2000](#)). A conceptually similar, though computationally
 370 more demanding approach using instability thresholds for offline scenarios is
 371 proposed in [Diehl and Görg \(2002\)](#). Our approach here is a direct translation
 372 of the former, and a computationally more efficient variant of the latter.

373 Concretely, we use the aggregation approach of Section 3.1 to obtain a refer-
 374 ence layout that serves as a baseline for representing stable overall structures,
 375 and thereby facilitates the formation of a persistent mental map. For each in-
 376 dividual graph in the sequence we do, however, allow deviations from reference
 377 positions if they lead to improved representation of momentary structures. A
 378 stress function quantifying this compromise is

$$\text{stress}_\alpha(P^{(t)}) = (1 - \alpha) \cdot \underbrace{\sum_{i < j} \omega_{ij}^{(t)} \left(\delta_{ij}^{(t)} - \|p_i^{(t)} - p_j^{(t)}\| \right)^2}_{\text{individual layout quality}} + \alpha \cdot \underbrace{\sum_i \phi_i^{(t)} \|p_i^{(t)} - p_i\|^2}_{\text{stability}},$$

where $P = (p_i)_{i \in V}$ denotes the reference layout determined as described in the previous section, and weights $\phi_i^{(t)}$ allow for inter-vertex variation in deviation tolerance.

The stability term thus corresponds to a pointwise penalty for deviations from the reference layout, and the parameter $0 \leq \alpha \leq 1$ provides explicit control of the trade-off between quality (original stress) and stability. Note that minimizing stress_α for $\alpha = 0$ corresponds to regular stress minimization without control for stability, and $\alpha = 1$ corresponds to the aggregation method, since no deviation from the reference layout is tolerated.

Besides modeling stability explicitly as a layout objective, two additional measures are taken. First, we initialize individual layout computation with the reference layout for the first observation, and with the layout of the preceding observation for subsequent ones. Assuming that consecutive observations are structurally similar, stress minimization is expected to reach a similar local minimum as for the preceding observation. Hence, ambiguities will be resolved in favor of the preceding layout if deviations from the reference layout must occur due to structural changes. For the same reason we may also expect that only few iterations are needed to compute a layout. Second, we postprocess the layout thus obtained using *Procrustes rotation* (Sibson, 1978), i.e., by applying an affine transformation that minimizes the sum of squared deviations from reference positions without changing relative distances.

For now, we use constant stability weights $\phi_i^{(t)} := 1$ for all i and t . More sophisticated choices, however, may be useful to compensate for cases with highly varying degrees or localized structural change. Another potential use of stability weights is as normalizing factors, such that the quality and stability part in stress_α are on equal scales, and thus, parameter α can be interpreted more easily in between its extreme values. More technical and user-oriented experimentation is needed, though, to quantify dependencies on these parameters.

All fifteen networks of our running example are shown in Figure 5, now laid out using moderate anchoring with reference positions taken from the aggregate layout of Figure 3. Although more effort is needed to trace vertices from one layout to the next, major substructures remain stable and the individual quality of layouts has improved notably.

This is confirmed when considering the issues we had with constant layout in

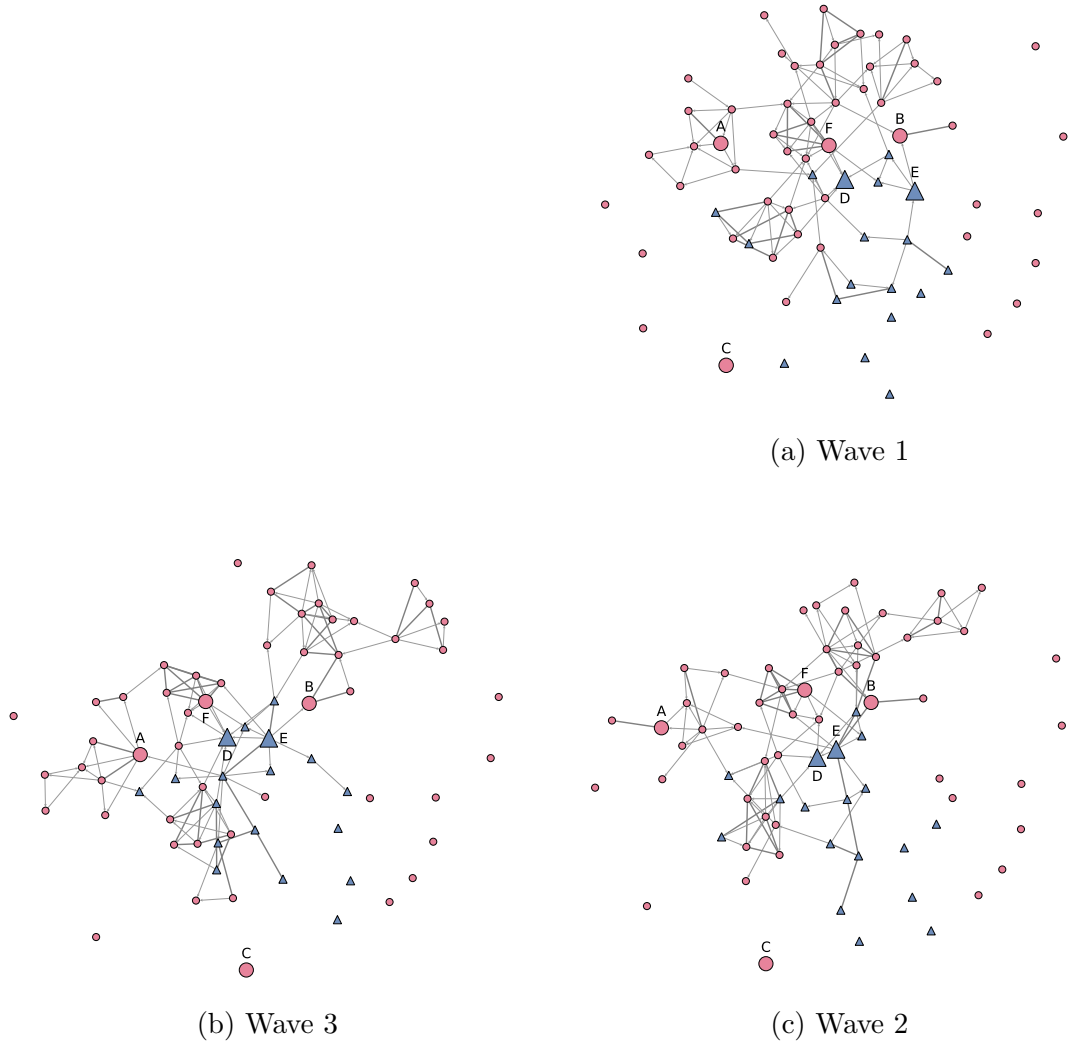




Fig. 5. Waves 1–3 of the acquaintanceship network drawn by anchoring with aggregate layout as a reference (stability at $\alpha = 0.15$). Waves 4–9 and 10–15 on next pages.

the previous section. The groups of Actors A and B move towards each other when they start being connected (Waves 11–15). The near-bipartite structure of the group involving Actor A in Wave 14 is apparent. Movement highlights the changing affiliations of Actors C and the volatile connection between Actor D and Actor E. While it is more difficult to locate them in individual layouts, deviations from their reference position are actually meaningful and therefore considered desirable.


Finally, we want to demonstrate the benefits of using offline information. Compare the above results of anchoring at reference positions from an aggregation layout (Figure 5) to an online approach based on initialization with the preced-




(d) Wave 4




(e) Wave 5




(f) Wave 7



(g) Wave 6




(h) Wave 8




(i) Wave 9


Fig. 5. Waves 4–9 (continued from previous page).




(j) Wave 10




(k) Wave 11




(ℓ) Wave 13



(m) Wave 12




(n) Wave 14




(o) Wave 15


Fig. 5. Waves 10–15 (continued from previous page).




(a) Wave 10



(b) Wave 11



(c) Wave 13



(d) Wave 12

Fig. 6. Waves 10–13 of our running example drawn by a common online approach. Substructures are similar, but the offline method (Figure 5) avoids the need for flipping subgroups such as the ones associated with Actors A and F from Wave 12 to 13.

ing layout and no control for stability (Figure 6). Note that the latter is similar to the online approach provided by the SoNIA¹ software tool (Bender-deMoll and McFarland, 2006), since the method employed there is conceptually the same, as is the basic objective function, stress, when using the Kamada-Kawai option (MultiCompKK).

Both approaches result in a similar configuration for Wave 10. In Wave 11, the group of Actor A establishes lasting connections to the groups of Actor B and momentary connections to the group of Actor E, forcing the groups to move towards each other. The online approach results in a positioning of the group of Actor A below the group of Actor F in Waves 11 and 12. However, with the

¹ <http://www.stanford.edu/group/sonia/>

dissolution of the momentary connections in Wave 13, those groups are forced to flip around each other, creating excessive movement between observations.

Anchoring with reference positions from an aggregation layout exploits offline information to avoid the need for such flips, and thus eliminates some movement that indicates structural changes inappropriately. Observe that there is much less difference in the configuration of substructures than suggested by the large movements in the online scenario.

4 Stochastic Actor-Oriented Models

Today, stochastic actor-oriented models (SAOM) (see, e.g., [Snijders 2005, 2001](#)) are the primary tool for the analysis of longitudinal social networks. Given panel data, i.e., sequences of network observations at discrete time points, as input, it is supposed that an unobserved evolution process took place inbetween consecutive states. The process model is composed of network-specific rules such as reciprocity, triadic closure, or homophily, and these rules are referred to as *network effects*. The aim of an analysis is to identify network effects that explain the observed evolution.

We briefly introduce the model in its simplest form, with only two consecutive network observations. For a detailed introduction and extended actor-oriented models see [Snijders, van de Bunt, and Steglich \(2010b\)](#); [Snijders, Steglich, and Schweinberger \(2007\)](#); [Steglich, Snijders, and Pearson \(2010\)](#).

Let $V = \{1, \dots, n\}$ denote the fixed set of actors, and \mathcal{A}_V the set of all adjacency matrices over V , i.e., for a matrix $A \in \mathcal{A}_V$ we have entry $a_{ij} = 1$ if there is a tie from i to j , and $a_{ij} = 0$ otherwise. Note that we are dealing with directed graphs from now on. The two observed networks are denoted by $A^{(pre)}, A^{(post)} \in \mathcal{A}_V$.

The evolution from $A^{(pre)}$ to $A^{(post)}$ is regarded as a stochastic sequence of elementary changes, called *micro-steps*, that are performed consecutively by randomly chosen actors. An actor is allowed to either create a single new outgoing tie, delete a single existing outgoing tie, or not change any tie at all. When an actor is designated to perform the next micro-step, probabilities of all feasible micro-steps are determined by the current network structure $A^{(cur)}$ and do not depend on previous states. These model assumptions are realized in a continuous-time Markov process on state space \mathcal{A}_V that is starting from the first observation, $A^{(pre)}$. The transition probability from the current state $A^{(cur)}$ to a possible next state A is positive only if the two networks differ by at most one tie a_{ij} for which i is the actor designated to perform the next micro-step. Since actors are assumed to strive for improving their position in the

network, the transition probability from $A^{(cur)}$ to A depends on the position enhancement achieved by changing to a_{ij} rather than any other feasible option.

The assessment of an actor i 's position in a network A is according to an objective function $f_i : \mathcal{A}_V \rightarrow \mathbb{R}$ that is assumed to be of the form

$$f_i(\theta; A) = \sum_{k=1}^K \theta_k \cdot s_{ki}(A) , \quad (3)$$

where each $s_k = (s_{k1}, \dots, s_{kn})^T$, $1 \leq k \leq K$, is a vector of statistics representing network effects, and $\theta = (\theta_1, \dots, \theta_K) \in \mathbb{R}^K$ are associated model parameters. The latter are fitted to the observed data, assuming that actors strive to maximize their objective function. Uncertainty stemming from the potential influence of chance and other unknown factors is usually represented by a random term added to the objective function, but is not important for the purpose of this paper.

Specifically, statistics $s_{ki}(A)$ count the number of occurrences of configurations in A that actor i is part of, and thus represent the k^{th} network effect. The selection of network effects is referred to as the *model specification*. Some common statistics are the number of

outgoing ties:	$s_{1i}(A) = \sum_{j=1}^n a_{ij}$
reciprocated ties:	$s_{2i}(A) = \sum_{j=1}^n a_{ij}a_{ji}$
transitive ties:	$s_{3i}(A) = \sum_{j=1}^n a_{ij} \cdot \max_{l \neq j, i} (a_{il}a_{lj})$
actors at distance two:	$s_{4i}(A) = \sum_{j=1}^n (1 - a_{ij}) \cdot \max_{l \neq j, i} (a_{il}a_{lj})$
transitive triplets:	$s_{5i}(A) = \sum_{j=1}^n \sum_{l=1}^n a_{ij}a_{il}a_{lj}$
3-cycles:	$s_{6i}(A) = \sum_{j=1}^n \sum_{l=1}^n a_{ij}a_{jl}a_{li}$

Such statistics indicate local rules of network evolution. Note that the entire formulation implies a rather strong assumption, namely that actors are homogenous in the sense that the same rules are working for them in the same way.

For illustration, consider three actors i, j_1, j_2 with i intending to connect with either j_1 or j_2 . Assume that the network positions of j_1 and j_2 are equivalent, except that there exists a tie from j_1 to i , but not from j_2 . In this setting, a positive reciprocity parameter (i.e., a positive weight θ_2 for the above s_{2i}) implies that i is more likely to connect with j_1 than with j_2 . A negative parameter would imply the opposite. In the following, we identify a model with its parameters θ .

497 Given a model θ and an initial network $A^{(pre)}$, the evolution predicted by θ can
 498 be simulated via a sequence of micro-steps conforming to the above rules. We
 499 denote such a sequence by $\mathcal{S}(\theta, A^{(pre)})$. One possible condition for terminating
 500 the simulation is that the current network state $A^{(cur)}$ differs from $A^{(pre)}$ in
 501 the same number of dyads as $A^{(post)}$. The probability that a network $A \in \mathcal{A}_V$
 502 is the final state of a simulation $\mathcal{S}(\theta, A^{(pre)})$ depends on θ , so that θ induces a
 503 probability distribution over all networks in \mathcal{A}_V conditional on $A^{(pre)}$.

504 Among several methods for fitting θ to observed data (Snijders 2001; Snijders,
 505 Koskinen, and Schweinberger 2010a; Koskinen and Snijders 2007), the one
 506 most commonly used is an instance of the *Method of Moments* (Snijders,
 507 2001).² Parameters are determined to match statistics

$$s_k(A^{(post)}) := \sum_{i=1}^n s_{ki}(A^{(post)}), \quad 1 \leq k \leq K, \quad (4)$$

508 of the second observation $A^{(post)}$ as closely as possible to their expected val-
 509 ues in predicted networks. Since the model is too complicated to determine
 510 expected values precisely, they are estimated from simulations. Therefore, the
 511 aim is finding parameters θ that satisfy the *moment equations*

$$E_\theta [S_k] = s_k(A^{(post)}), \quad 1 \leq k \leq K, \quad (5)$$

512 where S_k is a random variable for the value of s_k in a network resulting from
 513 a simulation $\mathcal{S}(\theta, A^{(pre)})$.

514 5 Diagnostic Visualization

515 Based on the view that a SAOM θ induces a probability distribution over \mathcal{A}_V ,
 516 we say that θ 's predictions fit the observed data if $A^{(post)}$ is highly probable
 517 under the distribution implied by θ , i.e., if it is likely that $A^{(post)}$ is the result
 518 of simulations $\mathcal{S}(\theta, A^{(pre)})$.

519 There are several potential causes for poor representation of data. The main is-
 520 sue is model mis-specification, i.e., the inclusion of effects that are not present
 521 in the data or exclusion of effects that actually drive the evolution process.
 522 Inappropriate specification is often evidenced by parameter estimates that ex-
 523 hibit large standard errors or, in extreme cases, by non-convergence of the

² An implementation is available in the R package `RSiena`, see <http://cran.r-project.org/web/packages/RSiena/>.

iterative estimation process. A problem that is even more difficult to cope with is inhomogeneity across actors, because it is hidden in the aggregated statistics (4). The fact that parameters θ do not depend on individual actors, but are estimated in order to fit the average behavior of all actors by satisfying the moment equation (5), basically suggests that all actors follow the same rules, except for differences implied by covariates. In many cases, this restrictive assumption is unjustified and leads to poor models. However, poor fit caused by inhomogeneity is often not apparent from estimates and their standard errors because parameters inferred from the aggregate quantities in (5) do not necessarily exhibit a high level of uncertainty, even if occurrences of statistics are distributed very inhomogeneously across actors.

In this section, we introduce methods that allow for visual exploration of model predictions and comparison with observed data. The first method reveals a global impression of model predictions that can be compared with the observed data in order to detect regions or structural patterns that are poorly represented. The second method is especially designed for assessing actor-inhomogeneity. Contrary to the first method, where the focus is on the global structure of connections, the second method considers local statistics in the neighborhood of individual actors. These graphical methods support model interpretation and can provide guidance in proper effect selection in addition to the significance tests (Schweinberger, 2010) already integrated into the Siena program (Ripley and Snijders, 2011).

5.1 Overall Fit

Our first goal is a graphical representation of the predictions made by a model. To assess how well such predictions and the data that was actually observed align, the representation shall be comparable to related visualizations of the observed data. What, however, is an appropriate graphical representation for a *distribution* of networks? And how could it be displayed whether a network is typical for a given distribution or not?

Since the actual distribution is computationally intractable, we start by sampling from it. Let $A^{(1)}, \dots, A^{(R)} \in \mathcal{A}_V$ be a set of network samples resulting from R simulations $\mathcal{S}(\theta, A^{(pre)})$. In order to reveal shared characteristics of these samples, such as clustering, we follow the idea of the stress minimization approach described in Section 3.1 where, too, information on structural features of several networks is aggregated into a single layout.

Even in the context of networks arising from uncoordinated myopic decisions of independently acting actors, the following arguments substantiate why an aggregate layout approach may be appropriate. Local actor decisions indirectly

depend on each other, since previous dynamics formed the structure that subsequent changes depend on. Therefore, actor-oriented models determine not only local dynamics, but also the global evolution of network characteristics. Purely local effects such as, for instance, reciprocation and transitive closure, thus reinforce existing global features such as segmentations into cohesive substructures and structural holes. Since the given observation $A^{(pre)}$ seeds an initial global structure, it is plausible that a model can not only predict occurrences of structural patterns but also their topological distribution.

In analogy to the **stress** in Equation (2), we therefore define a stress function **stress_{sim}** based on shortest-path distances between each pair of vertices in all simulated networks simultaneously. Observe that it does not matter that simulation runs do not have a meaningful ordering as the observations do. Because they are determined by the same principles, comparison of such drawings with a layout of $A^{(post)}$ based on the regular stress function yields insight into structural similarities or differences.

Let $\delta_{ij}^{(1)}, \dots, \delta_{ij}^{(R)}$ denote the lengths of undirected shortest paths between vertices i and j in the graphs of $A^{(1)}, \dots, A^{(R)}$. If i and j are in different connected components of $A^{(r)}$, we replace the conventional infinite distance by $\delta_{ij}^{(r)} = n$, so that all distance between distinct vertices are in the range $\{1, \dots, n\}$. Similarly to the aggregated stress function (2), we define a stress function based on samples $A^{(1)}, \dots, A^{(R)}$ by

$$\text{stress}_{sim}(P) = \sum_{i < j} \omega_{ij} \left(\bar{\delta}_{ij} - \|p_i - p_j\| \right)^2, \quad (6)$$

where $\bar{\delta}_{ij} := \frac{1}{R} \sum_{r=1}^R \delta_{ij}^{(r)}$ denotes the sample mean of simulated distances between vertices i and j . As in Section 2.2, local accuracy should be emphasized by reducing the impact of distant dyads $\{i, j\}$ identified by high sample means $\bar{\delta}_{ij}$. Our application also requires additional emphasis on the influence of highly confirmed distances indicated by small sample variances $s^2(\delta_{ij}) := \frac{1}{R-1} \sum_{r=1}^R \left(\delta_{ij}^{(r)} - \bar{\delta}_{ij} \right)^2$. We therefore propose to use weights

$$\omega_{ij} := \frac{1}{(1 + s^2(\delta_{ij})) \bar{\delta}_{ij}^2}.$$

where the newly introduced factor $\frac{1}{1+s^2(\delta_{ij})}$ increases the bias towards faithful representation of local structures, since dyads with higher average distances usually exhibit larger standard deviations.

Minimization of **stress_{sim}** yields a layout in which vertices that are likely (unlikely) to be connected by a short path are close together (far apart). Note

that (6) does not include any data of the second observation $A^{(post)}$, but only information resulting from model-based simulations starting from $A^{(pre)}$. However, if all sampled networks $A^{(r)}$ were identical to $A^{(post)}$, stress_{sim} would reduce to the basic stress function of $A^{(post)}$ for static stress minimization as introduced in Section 2.2. The layout of $A^{(post)}$ based on stress_{sim} would then be equal to the layout based on graph-theoretic distances. Therefore, in case of excellent and detailed model predictions, both methods would yield essentially the same layout of $A^{(post)}$. In particular, a drawing of $A^{(post)}$ based on stress_{sim} would exhibit short edges, and large distances between pairs of unconnected vertices. Long edges, on the other hand, appear where graph-theoretic distances are small in the observation $A^{(post)}$, but not in the samples $A^{(1)}, \dots, A^{(R)}$.

We stress that the assessment of model predictions should be in line with the inferential goal of SAOMs. This is not the prediction of individual ties, but the characterization of local processes by a limited number of network effects. Consider, for instance, the longitudinal network of Figure 7a together with a model $\bar{\theta}$ consisting of only two effects, **outgoing ties** and **reciprocated ties**. In this setting, it is equally probable that a simulation $\mathcal{S}(\bar{\theta}, A^{(pre)})$ ends up in $A^{(post)}$ or in network A depicted in Figure 7b. Therefore, simulations resulting in A do not necessarily invalidate the model even if A differs from $A^{(post)}$ in one third of all possible ties. Knowledge about poorly predicted ties may well be important for model evaluation, though, because it may inform the selection of effects.

To illustrate how our approach can be used for this kind of assessment, let us compare a layout of $A^{(post)}$ based on observed shortest-path distances in Figure 7c with a layout of expected shortest-path distances in Figure 7d. Comparing both layouts, we recognize that the reciprocated edges connecting the two gray vertices in Figure 7d are longer than other reciprocated edges in the same graph as well as corresponding edges in Figure 7c. Moreover, compared to their positions in Figure 7c, the white vertex and the black vertex are closer in Figure 7d. The fact that both pairs are equidistant in Figure 7d implies that, although only one pair is connected in $A^{(post)}$, both connections are equally probable under the model $\bar{\theta}$. As already stated, this does not necessarily translate into poor model-fit, but that the model fails to predict specific connections between these actors reliably. If context calls for more detailed explanations of observed relations, this may be seen as a hint at the inclusion of other effects, possibly based on actor covariates.

To facilitate comparison between layouts for $A^{(post)}$ according the observed (**stress**) and expected (stress_{sim}) distances, both computations are initialized with positions from a layout of $A^{(pre)}$ according to **stress**. While this ensures a certain degree of similarity, we refrain from adding stability via anchoring, because this would interfere with our reading of the simulation result. Instead,

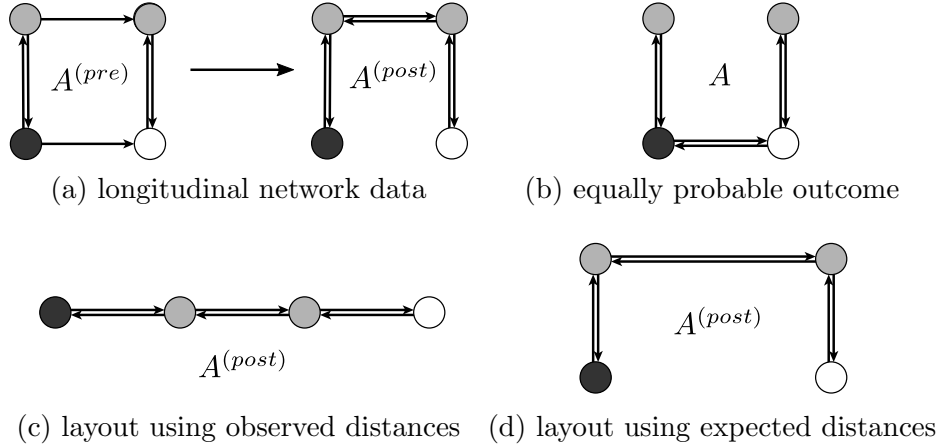


Fig. 7. How to read layouts based on model predictions (see text).

layouts are aligned using Procrustes rotation, which does not alter relative distances.

We apply the method to the sixth and thirteenth observation of our running example. Note however that in case $A^{(pre)}$ is disconnected, connections between different components are hardly predictable by a SAOM. As a consequence, the positioning of separate components can be quite arbitrary and therefore misleading. Since this is particularly true for singleton components, we omit the five isolated actors of Wave 6, but continue to refer to the reduced data as Waves 6 and 13.

The data was subjected to the three models listed in Table 1, with Wave 6 as the initial and Wave 13 as the target observation. The three models are increasingly complex. While Model 1 is elementary and dyadic, the additional effects of Model 2 generate triadic closure and thus the formation and reinforcement of clusters. Model 3 is a further refinement representing the formation of local hierarchies (if signs are as in our example).

The layout in Figure 8a serves as initialization for visualizing model predictions by minimizing the modified stress function stress_{sim} . The resulting layout of Wave 13 based on simulations of Model 2 is depicted in Figure 9a. To allow direct comparison, Wave 13 laid out according to stress is shown once more in Figure 9b. Layouts have been aligned by Procrustes rotation.

Informally speaking, the clustering in the layouts of Figure 9 is quite similar. The cohesive group of female students including Actor Y is easily identified in both. Other groups such as the males around Actors Q and Actor S, the females at the top left, or the mixed group next to Actor U can be matched as well. However, from their poor representation in Figure 9a we can conclude that several ties between groups as well as ties connecting marginal vertices like V or T with the center of the network are not predicted well by the model.







Network Effect	Model 1		Model 2		Model 3	
	θ	(s.e.)	θ	(s.e.)	θ	(s.e.)
 outgoing ties	-1.61	(0.10)	-1.95	(0.23)	-2.16	(0.23)
 reciprocated	2.64	(0.20)	1.96	(0.22)	2.60	(0.29)
 transitive ties			1.37	(0.24)	1.05	(0.26)
 actors at distance 2			-0.26	(0.06)	-0.22	(0.06)
 transitive triplets					0.36	(0.09)
 3-cycles					-0.65	(0.16)

Table 1

Analyzed models with parameters θ estimated using the **RSiena** package (Ripley and Snijders, 2011) for $A^{(pre)}$ = Wave 6 and $A^{(post)}$ = Wave 13 (ignoring isolates of Wave 6).

Another salient disparity is the location of Actor W who has a bridging position between two groups in Figure 9b, but is close to only one group in Figure 9a. This indicates that the model performs badly in predicting the creation of ties between W and the group at the top left. Note that there were no such connections in Wave 6. The position of Q can be explained in a similar way.

A weakness of this approach is that it is not always obvious whether a poor layout, indicated by long edges and close unconnected vertices, is caused by poor model predictions or by desired distances that are just difficult to realize in a two-dimensional drawing. Therefore, Figure 10 shows two complementary diagrams that are meant to improve interpretability of the layouts from Figure 9a.

Figure 10a displays sample means $\bar{\delta}_{ij}$ of simulated graph-theoretic distances between all dyads. E.g., the position of R in Figure 9a close to the above female group suggests that connections between this group and R are likely. But in Figure 10a, the two more suspicious edges incident to R reveal that

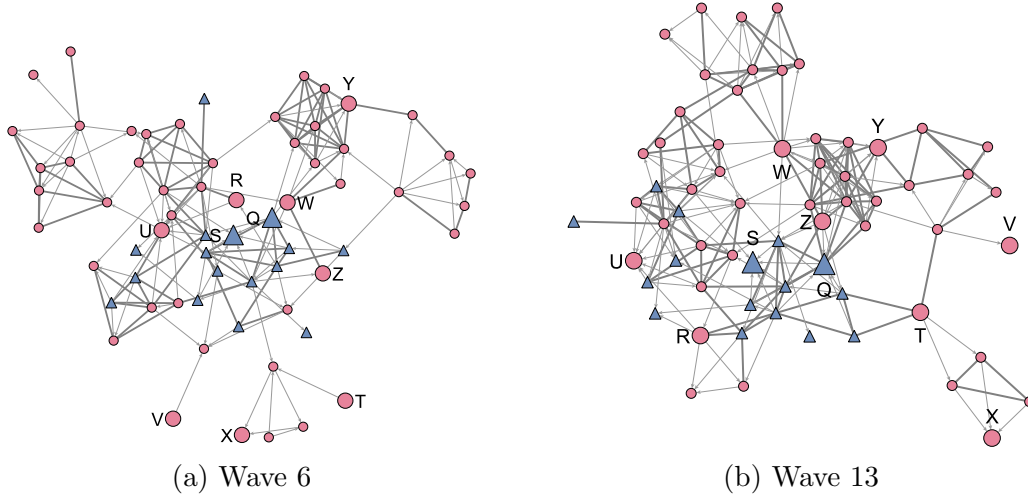


Fig. 8. Stress minimization layouts of Wave 6 and Wave 13 based on graph-theoretic distances. Vertices that are isolated in Wave 6 are omitted in both networks.


its position is rather caused by the realization of desired distances to vertices below it. Note that peripheral vertices such as V or T, exhibit large average graph-theoretic distances to all other vertices.

Additional information on model predictions is obtained by considering for each pair of vertices i and j the probability π_{ij} that edge a_{ij} exists in a simulated network A . This probability can be estimated by counting the number of simulated networks in which i and j are adjacent, i.e.,


$$\pi_{ij} := \frac{1}{R} \sum_{r=1}^R a_{ij}^{(r)}.$$

Even though our previous considerations implied that assessment of model predictions must not be based solely on individual dyads, the elementary comparison of entries $a_{ij}^{(post)}$ with the predicted probabilities π_{ij} can still be indicative of the plausibility of a model, where especially the coincidence of high values of π_{ij} with absent ties $a_{ij}^{(post)}$ points at poor predictions. Moreover, correspondence of tie probabilities with distances between vertices in the layout implied by stress_{sim} is another indicator for assumptions on model predictions.

Figure 10b displays tie probabilities π_{ij} implied by Model 2. In general, it indicates that most edges are very improbable, and that most of the highly probable edges are within groups that existed previously while edges between groups are rather unlikely with few exceptions of moderate probability. This is indeed what we would expect from a model in which existing groups are reinforced by positive triadic-closure and reciprocity parameters, and a negative density and distance-2 parameter.




(a) Layout based on simulations of Model 2




(b) Layout based on graph-theoretic distances

Fig. 9. Comparison between layouts of Wave 13 based on simulated and observed distances. The layout in (a) results from modifying Figure 8a by minimizing stress_{sim} based on Model 2. The layout in (b) is determined by using regular stress minimization based on graph-theoretic distances.



(a) Sample means of shortest paths in simulations of Model 2



(b) Tie probabilities implied by simulations of Model 2

Fig. 10. Sample means of shortest paths and tie probabilities implied by simulations of Model 2. Color intensities and widths of edges increase with their probability of existence π_{ij} in (b). In (a) they increase with the inverse of the expected shortest path $1/\bar{\delta}_{ij}$ between connected vertices.

700 The basic structures in Figure 10b are, as expected, very similar to those in
701 Figure 10a. But there are also differences such as the moderate probabilities of
702 connections between Actor X and two others, which are visible in Figure 10b
703 but not in Figure 10a. Apparently, X is isolated in many simulated networks so
704 that the average graph-theoretic distances to the two other actors, and so the
705 Euclidean distances in the layout, are large, even though direct connections
706 are fairly probable. Due to this shortcoming, our method may not be suitable
707 for extremely sparse networks or for longitudinal data with great differences
708 between consecutive observations. Our experiences so far suggest that a Jac-
709 card index (see [Snijders, van de Bunt, and Steglich \(2010b\)](#)) higher than .5 is
710 a reasonable threshold.

711 Finally, we use our method to compare the three models listed in Table 1.
712 Figure 11 shows for each model the layout determined in the same way as in
713 Figure 9a. Compared to Figures 11b and 11c, where the layouts reveal sev-
714 eral clusters of vertices, vertices in Figure 11a are homogeneously distributed.
715 This indicates that the effects included in Model 1 are not sufficient to pre-
716 dict the observed clustering that, on the other hand, is clearly identifiable in
717 the predictions of Model 2 and Model 3. The additional effects included in
718 Model 3 reinforce the basic structure already implied by Model 2, but achieve
719 no remarkable improvement. In conclusion, Model 2 appears to strike the best
720 balance between fit and parsimony.

721 5.2 Inhomogeneity

722 From analyzing the overall configuration predicted by a model, we now shift
723 focus to individual actors. The methods below facilitate model diagnostics
724 in general, and validation of actor homogeneity assumptions in particular.
725 Thereby, they may support improvement of a model specification.

726 Stochastic actor-oriented models are estimated in order to obtain simulated
727 networks that resemble second observation $A^{(post)}$ in terms of aggregated statis-
728 tics s_k defined in Equation (4). Considering, e.g., the **outgoing ties** effect,
729 simulated networks should contain the same number of ties as $A^{(post)}$. Whether
730 these ties are homogeneously distributed over the entire network or clustered
731 around some hubs, however, is not differentiated by the summarizing moment
732 equations (5). One way to deal with inhomogeneities is the inclusion of addi-
733 tional network effects such as popularity or similarity effects. As a prerequisite,
734 however, the analyst must first become aware of them.

735 In the context of exponential-family random graph models (ERGMs), [Hunter,](#)
736 [Goodreau, and Handcock \(2008\)](#) addresses the problem of potential inhom-
737 ogeneity by comparing distributions of statistics (e.g., the degree distribution)




Fig. 11. Layouts of Wave 13 based on simulated networks predicted by Model 1, Model 2, and Model 3.

of simulated and observed networks. In contrast to ERGMs, however, where simulations start from a random network and are made to converge to a stationary distribution, SAOMs start from a fixed initial network $A^{(pre)}$, from which a finite number of micro-steps leads to the resulting simulated network. Consequently, actors are distinguishable by their position and the structure of their environment in $A^{(pre)}$. This justifies, in fact, the assessment of model predictions on the local level of individuals, as an alternative to the global analysis of distributions of statistics. To identify actors or groups of actors deviating from model predictions, we turn to the fit of statistics in the neighborhoods of individual actors.

Let θ denote a model based on objective functions $f_i = \sum_{k=1}^K \theta_k s_{ki}$ and let $A^{(1)}, \dots, A^{(R)} \in \mathcal{A}_V$ be network samples resulting from R simulations $\mathcal{S}(\theta, A^{(pre)})$. As in Section 5.1, these sampled networks are regarded as representative for model predictions. For each actor, we determine sample means

$$\bar{s}_{ki} = \frac{1}{R} \sum_{r=1}^R s_{ki} \left(A^{(r)} \right)$$

of statistics associated with the effects included in the model. To assess model predictions for individual actors, we are interested in deviations of these sample means from observed values. Differences

$$\Delta_{ki} := s_{ki} \left(A^{(post)} \right) - \bar{s}_{ki}$$

indicate deficiencies of model predictions for actor i in terms of the network effect s_k . While differences are comparable across actors, they are not comparable across effects, because the latter operate on variable scales. Therefore, relative deviations $\tilde{\Delta}_{ki}$ are determined by dividing Δ_{ki} by estimated standard deviations of corresponding statistics in the entire network, i.e.,

$$\tilde{\Delta}_{ki} := \frac{\Delta_{ki}}{\sqrt{\text{VAR}_k}}$$

with

$$\text{VAR}_k = \frac{1}{R} \sum_{r=1}^R \left(\sum_{i=1}^n \left(s_{ki} \left(A^{(r)} \right) - \bar{s}_{ki} \right) \right)^2.$$

This normalization is chosen in analogy to the t -statistics used for checking the convergence of the parameter estimation algorithm (see [Ripley and Snijders 2011](#)), which equal the sums $\sum_{i=1}^n \tilde{\Delta}_{ki}$ of respective relative deviations.

As a heuristic summary measure for the quality of model predictions for actor i , we here use the sum of absolute values of relative deviations $|\tilde{\Delta}_{ki}|$ of all considered statistics,

$$\|\tilde{\Delta}_i\|_1 := \sum_{k=1}^K |\tilde{\Delta}_{ki}|,$$

although other definitions such as a generalized Mahalanobis distance incorporating covariances between effects are conceivable as well.

Figure 12 shows Wave 13 of the acquaintanceship network with vertex areas proportional to the associated values $\|\tilde{\Delta}_i\|_1$ resulting from Model 2 (see Table 1). Recall that these represent differences of four effects, outgoing ties (OT), reciprocated ties (RT), transitive ties (TT), and actors at distance 2 (AD2), between the actual statistics in Wave 13 and those in simulations starting from Wave 6.




Fig. 12. Wave 13 of the acquaintanceship network with vertex sizes indicating sums of absolute deviations of simulated and observed statistics associated with the effects included in Model 2. The larger a vertex, the larger its lack of fit.

This helps with comparing the quality of predictions across actors and detecting individuals or groups whose behavior is particularly deviant from model predictions. In Figure 12, Z is most salient. But also the bridging vertices Y, W, and T are subject to a mismatch between simulated and observed statistics. From a global viewpoint, vertices of the cohesive group around Z and, in attenuated form, vertices of the group next to U exhibit particularly high deviations. However, the causes for the observed inhomogeneities and therefore with possible model improvements are not apparent from this visualization since it does not point to the effects that are responsible for poor prediction. This motivates the extension of the above idea to separate visualizations for different effects with sizes of vertices proportional to their relative deviation $\tilde{\Delta}_{ki}$ and colors of vertices indicating the direction of deviation, i.e., whether effects are over- or underestimated by model predictions.

788 Applied to effects not yet included in the model, such diagrams may indicate
 789 whether they are sufficiently represented by already included effects or whether
 790 they should be included themselves. For effects considered in the model, it is
 791 expected that the sum of deviations over all actors satisfies

$$\sum_{i=0}^n \tilde{\Delta}_{ki} \approx 0, \quad (7)$$


792 since parameters θ are estimated to satisfy Equation (5). But deviations of
 793 single actors or subgroups can compensate each other mutually, such that (7)
 794 can be true even though individual deviations are large. Inhomogeneities thus
 795 require detailed inspection as provided in our visualizations.

796 Moreover, there may be actors whose behavior with respect to some effect is
 797 hardly determined by the model such that deviations from observed values are
 798 caused by uncertainty rather than poor prediction. For other actors the implied
 799 probability distribution may admit only slight variations of statistics such that
 800 deviations from observed values indeed contradict the model predictions. The
 801 degree of uncertainty is generally indicated by sample standard deviations of
 802 respective statistics. The ratio $\frac{\Delta_{ki}}{\text{SSD}_{ki}}$, where


$$\text{SSD}_{ki} = \sqrt{\frac{1}{R-1} \sum_{i=1}^R (s_{ki}(A^{(r)}) - \bar{s}_{ki})^2}$$

803 is the sample standard deviation of local statistics, indicates whether it is
 804 likely that observed statistics are sampled from the model-implied probabil-
 805 ity distribution. High values indicate contradictions with model predictions,
 806 whereas low values can either result from fitting predictions that yield low
 807 values of Δ_{ki} , or from indefinite predictions that yield high values of SSD_{ki} .


808 Figure 13 and Figure 14 show Wave 13 with graphical vertex attributes rep-
 809 representing the quality of model predictions for individual actors regarding the
 810 four effects OT, RT, TT, and AD2. Visualizations in left columns illustrate re-
 811 sults obtained from Model 1. Right columns contain results of Model 2. Areas
 812 of vertices are proportional to absolute values of relative deviations $\tilde{\Delta}_{ki}$. Color
 813 indicates the sign of $\tilde{\Delta}_{ki}$, i.e., whether occurrences of effect-related configu-
 814 rations are under- (blue) or overestimated (red) by model predictions. Color
 815 intensity depicts the ratio $\frac{\Delta_{ki}}{\text{SSD}_{ki}}$ of deviation divided by standard deviation, so
 816 that intensely red- or blue-colored vertices hint at contradictions with model
 817 predictions. Light colors are either caused by small deviations Δ_{ki} or by large
 818 standard deviations SSD_{ki} , where the two causes can be distinguished by the
 819 size of vertices. Vertices are white if the absolute value of deviation is lower
 820 than the standard deviation. Note that a light-colored vertex does not neces-




(a) Model 1: outgoing ties (OT)
 $\sum_i \Delta_{ki} = 0.30, \sqrt{\text{VAR}_k} = 10.41$



(b) Model 2: outgoing ties (OT)
 $\sum_i \Delta_{ki} = 1.45, \sqrt{\text{VAR}_k} = 10.85$

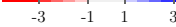


(c) Model 1: reciprocated ties (RT)
 $\sum_i \Delta_{ki} = 0.49, \sqrt{\text{VAR}_k} = 11.53$




(d) Model 2: reciprocated ties (RT)
 $\sum_i \Delta_{ki} = 0.60, \sqrt{\text{VAR}_k} = 12.41$

Fig. 13. Goodness-of-fit of Model 1 and Model 2 in terms of inhomogeneity analyzed for individual network effects.


The framed boxes in each diagram indicate the relation between vertex sizes and deviations Δ_{ki} . The color-scale  represents values $\frac{\Delta_{ki}}{\text{SSD}_{ki}}$.

sarily indicate a correct prediction but, depending on its size, rather indefinite predictions about the associated actor.


Figures 13b, 13d, 14b, and 14d, which represent Model 2, allow to distinguish effects responsible for remarkable deviations in Figure 12. While all effects contribute to the deviations of strikingly prominent Actors Z and T by underestimating their activity, W is only underestimated in terms of effects OT, TT, and AD2, but not in terms of RT. Likewise, effect AD2 does not contribute to the overestimation of Y.




(a) Model 1: transitive ties (TT)
 $\sum_i \Delta_{ki} = -79.80, \sqrt{\text{VAR}_k} = 12.85$



(b) Model 2: transitive ties (TT)
 $\sum_i \Delta_{ki} = 1.48, \sqrt{\text{VAR}_k} = 13.03$

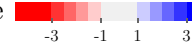


(c) Model 1: actors at distance 2 (AD2)
 $\sum_i \Delta_{ki} = 362.87, \sqrt{\text{VAR}_k} = 58.36$



(d) Model 2: actors at distance 2 (AD2)
 $\sum_i \Delta_{ki} = 1.40, \sqrt{\text{VAR}_k} = 54.61$

Fig. 14. Goodness-of-fit of Model 1 and Model 2 in terms of inhomogeneity analyzed for individual network effects.

The framed boxes in each diagram indicate the relation between vertex sizes and deviations Δ_{ki} . The color-scale  represents values $\frac{\Delta_{ki}}{\text{SSD}_{ki}}$.

829 The balanced coloring in the graphs representing Model 2 implies that (7)
 830 approximately holds for all effects.³ Regarding Model 1, the same applies to
 831 Figures 13a and 13c representing effects OT and RT. For these effects, which
 832 are included in both models, differences between the two models are hardly
 833 perceivable (see Figure 13). For TT and AD2, in contrast, the two models yield
 834 widely differing results (see Figure 14). Visualizations obtained from Model 1,

³ Actually, the graphs exhibit a few more red than blue vertices but most blue vertices are slightly larger. Hence, absolute values of positive deviations are larger so that the higher number of negative deviations is still balanced.

835 which includes neither of the two effects, are highly imbalanced. Compared
 836 to visualizations of Model 2, shifts to the extreme range of the color scale are
 837 noted. According to the algebraic signs of estimates in Model 2, the dominant
 838 blue in Figure 14a indicates that the observed data contains considerably more
 839 transitive ties than Model 1 predicts, whereas the dominant red in Figure 14c
 840 indicates an overestimation of the number of indirect neighbors. It is remark-
 841 able that many vertices indeed have an intense coloring but their moderate
 842 sizes indicate at the same time only small relative deviations $\tilde{\Delta}_{ki}$. Thus, model
 843 predictions for these actors must be quite definite which results in high values
 844 $\frac{\Delta_{ki}}{SSD_{ki}}$ due to minor standard deviations.

845 A closer look at the visualizations reveals that Model 2 outperforms Model 1
 846 not only in terms of TT and AD2 but also in terms of OT and RT. Especially,
 847 predictions for actors with considerably more incoming than outgoing ties in
 848 both observations such as U, S, or X are improved. This is in line with the
 849 parameter estimates of the two models. Actor S, for example, has six incoming
 850 but no outgoing ties in Wave 6 and Wave 13, which results in an objective
 851 function value of zero. In Model 1, this value is increased by 1.03 for S, if an
 852 incoming tie is reciprocated. Hence, reciprocation is likely, and this leads to an
 853 overestimation of OT and RT. In Model 2, any new tie that S builds decreases
 854 his objective function because of the negative parameter estimate for AD2.
 855 Since no newly created tie would be a transitive tie but most would increase
 856 the number of indirect neighbors, the observed behavior of S gets more likely
 857 under Model 2, as conveyed in the visualizations.

858 The apparent improvement of model predictions after inclusion of TT and
 859 AD2 implies that these effects are not represented indirectly in Model 1. This
 860 is further substantiated by the fact that for OT and RT model predictions are
 861 hardly affected by the addition of the two other effects. Obviously, correlations
 862 between the effects are only moderate, so that it seems advisable to include
 863 TT and AD2 in the model. A similar comparison between Model 2 and Model 3
 864 reveals that the additional inclusion of the effects **transitive triplets** and
 865 **3-cycles** yields only slightly improved predictions for a few actors.

866 6 Discussion

867 We have presented methods for explorative and diagnostic visualization of
 868 longitudinal social networks.

869 Explorative visualization of longitudinal social networks was treated as an of-
 870 fline dynamic graph-drawing problem, and proposed a corresponding instan-
 871 tiation of the generic stress-minimization framework for graph layout. Using
 872 anchoring at reference positions from a layout of the aggregate graph, inter-

873 mediate layouts produced by our method represent a compromise between
874 individual layout quality and persistence of an overall organization.

875 In summary, this method serves to visually explore the evolution of a dy-
876 namic network. Aggregation facilitates getting an overview of global structure,
877 whereas anchoring provides explicit control over the balance of readability and
878 stability, and therefore allows to put emphasis on the specifics of individual
879 networks in the sequence. It should be noted, though, that the approach is
880 based on the implicit assumption that there actually exists a relatively con-
881 stant global organization.

882 Directions for future work therefore include the development of variant ap-
883 proaches for other scenarios by taking advantage of the flexibility of stress
884 minimization. Also, more experimentation is needed to guide tuning of in-
885 volved parameters by considering the type of structural change present in the
886 longitudinal network at hand.

887 In the modeling-related part of the paper, we have shown that layouts repre-
888 senting predictions from stochastic actor-oriented models can be obtained by
889 aggregating over a set of simulated networks rather than a sequence of obser-
890 vations. By combining all information on dyadic predictions into one layout,
891 an impression of the predicted global structure is conveyed. Our approach
892 does assume that simulation results exhibit a certain degree of similarity in
893 terms of their global structure, although this is not directly ensured by the
894 model’s inferential goal of identifying local processes. We argued, however,
895 that conditioned on the initial structure in the first observation and because
896 of dependencies of dynamics on previous changes, local processes often do
897 yield globally similar patterns. This may no longer hold if observations differ
898 in, say, more than half of the dyads, but this would also be a problem for the
899 models themselves.

900 Our second diagnostic method focuses on the correspondence of predictions
901 and observations on the actor level by visualizing discrepancies between simu-
902 lated and observed local statistics. The aim is to either evaluate homogeneity
903 assumptions or to detect outliers or groups of actors with deviant behavior in
904 order to analyze their shared characteristics. The latter may indirectly suggest
905 model alterations, and one particular use-case is to experiment with additional
906 effects not yet included in a model. Note that, because of the Markov property
907 of SAOMs, both diagnostic methods extend to more than two observations.

908 The visualization methods presented here are to facilitate visual exploration of
909 data and models. Clearly, experimentation and further refinement will be nec-
910 essary to better assess their utility. The inherent flexibility of the basic layout
911 engine, stress minimization, may carry over to related visualization problems,
912 and we plan to look at exponential-family random graph models (ERGM; see,

913 e.g., [Snijders, Pattison, Robins, and Handcock 2006](#)) in particular.

914 Most of the methods described here are already available within visone,⁴ a
915 software tool for network analysis and visualization, and more will be added.

916 References

- 917 Bender-deMoll, S., McFarland, D., 2006. The art and science of dynamic net-
918 work visualization. *Journal of Social Structure* 7 (2).
- 919 Bertin, J., 1983. *Semiology of Graphics: Diagrams, Networks, Maps*. University
920 of Wisconsin Press.
- 921 Böhringer, K.-F., Paulisch, F. N., 1990. Using constraints to achieve stability
922 in automatic graph layout algorithms. In: *Proceedings of the ACM Human*
923 *Factors in Computing Systems Conference (CHI '90)*. ACM, The Association
924 for Computing Machinery, pp. 43–51.
- 925 Brandes, U., 2001. Drawing on physical analogies. In: Kaufmann, M., Wagner,
926 D. (Eds.), *Drawing Graphs: Methods and Models*. Vol. 2025 of *Lecture Notes*
927 *in Computer Science*. Springer-Verlag, pp. 71–86.
- 928 Brandes, U., Corman, S. R., 2003. Visual unrolling of network evolution and
929 the analysis of dynamic discourse. *Information Visualization* 2 (1), 40–50.
- 930 Brandes, U., Freeman, L. C., Wagner, D., in preparation. Social networks. In:
931 Tamassia, R. (Ed.), *Handbook of Graph Drawing and Visualization*. CRC
932 Press.
- 933 Brandes, U., Kenis, P., Raab, J., 2006. Explanation through network visu-
934 alization. *Methodology* 2 (1), 16–23, spanish translation in *REDES* 9(6),
935 2005.
- 936 Brandes, U., Pich, C., 2007. Eigensolver methods for progressive multidimen-
937 sional scaling of large data. In: Kaufmann, M., Wagner, D. (Eds.), *Proceed-*
938 *ings of the 14th International Symposium on Graph Drawing (GD'06)*. Vol.
939 4372 of *Lecture Notes in Computer Science*. Springer-Verlag, pp. 42–53.
- 940 Brandes, U., Pich, C., 2009. An experimental study on distance-based
941 graph drawing. In: *Proceedings of the 16th International Symposium on*
942 *Graph Drawing (GD'08)*. Vol. 5417 of *Lecture Notes in Computer Science*.
943 Springer-Verlag, pp. 218–229.
- 944 Brandes, U., Pich, C., 2011. More flexible radial layout. *Journal of Graph*
945 *Algorithms and Applications* 15 (1), 157–173.
- 946 Brandes, U., Wagner, D., 1997. A Bayesian paradigm for dynamic graph lay-
947 out. In: Di Battista, G. (Ed.), *Proceedings of the 5th International Sympo-*
948 *sium on Graph Drawing (GD'97)*. Vol. 1353 of *Lecture Notes in Computer*
949 *Science*. Springer-Verlag, pp. 236–247.
- 950 Branke, J., 2001. Dynamic graph drawing. In: [Kaufmann and Wagner \(2001\)](#),
951 pp. 228–246.

⁴ <http://www.visone.info/>

- 952 Bridgeman, S. S., Tamassia, R., 2000. Difference metrics for interactive or-
 953 thogonal graph drawing algorithms. *Journal of Graph Algorithms and Ap-*
 954 *plications* 4 (3), 47–74.
- 955 Cleveland, W. S., McGill, R., September 1984. Graphical perception: Theory,
 956 experimentation, and application to the development of graphical methods.
 957 *Journal of the American Statistical Association* 79 (387), 531–554.
- 958 Cox, T. F., Cox, M. A. A., 2001. *Multidimensional Scaling*, 2nd Edition.
 959 CRC/Chapman and Hall.
- 960 de Leeuw, J., 1977. Applications of convex analysis to multidimensional scal-
 961 ing. In: Barra, J. R., Brodeau, F., Romier, G., van Cutsem, B. (Eds.), *Recent*
 962 *Developments in Statistics*. Amsterdam: North-Holland, pp. 133–145.
- 963 Di Battista, G., Eades, P., Tamassia, R., Tollis, I. G., 1999. *Graph Drawing:*
 964 *Algorithms for the Visualization of Graphs*. Prentice Hall.
- 965 Diehl, S., Görg, C., 2002. Graphs, they are changing. In: *Proceedings of the*
 966 *10th International Symposium on Graph Drawing (GD’02)*. Vol. 2528 of
 967 *Lecture Notes in Computer Science*. Springer-Verlag, pp. 23–30.
- 968 Dwyer, T., Gallagher, D. R., 2004. Visualising changes in fund manager hold-
 969 ings in two and a half-dimensions. *Information Visualization* 3 (4), 227–244.
- 970 Dwyer, T., Hong, S.-H., Koschützki, D., Schreiber, F., Xu, K., 2006. Visual
 971 analysis of network centralities. In: *Proceedings of the 2006 Asia-Pacific*
 972 *Symposium on Information Visualisation (APVis’06)*. pp. 189–197.
- 973 Dwyer, T., Marriott, K., Wybrow, M., 2008. Topology preserving constrained
 974 graph layout. In: *Proceedings of the 16th International Symposium on*
 975 *Graph Drawing (GD’08)*. Vol. 5417 of *Lecture Notes in Computer Science*.
 976 Springer-Verlag, pp. 230–241.
- 977 Eades, P., 1984. A heuristic for graph drawing. *Congressus Numerantium* 42,
 978 149–160.
- 979 Erten, C., Kobourov, S., Le, V., Navabi, A., 2004. Simultaneous Graph Draw-
 980 ing: Layout Algorithms and Visualization Schemes. In: *Proceedings of the*
 981 *11th International Symposium on Graph Drawing (GD’03)*. Vol. 2912 of
 982 *Lecture Notes in Computer Science*. Springer-Verlag, pp. 437–449.
- 983 Freeman, L. C., 2000. Visualizing social networks. *Journal of Social Structure*
 984 1 (1).
 985 URL <http://www.heinz.cmu.edu/project/INSNA/joss/vsn.html>
- 986 Friedrich, C., Eades, P., 2002. Graph drawing in motion. *Journal of Graph*
 987 *Algorithms and Applications* 6 (3), 353–370.
- 988 Friedrich, C., Houle, M. E., 2002. Graph drawing in motion II. In: *Proceedings*
 989 *of the 9 International Symposium on Graph Drawing (GD’01)*. Vol. 2265 of
 990 *Lecture Notes in Computer Science*. Springer-Verlag, pp. 122–125.
- 991 Frishman, Y., Tal, A., 2008. Online dynamic graph drawing. *IEEE Transac-*
 992 *tions on Visualization and Computer Graphics* 14 (4), 727–740.
- 993 Fruchterman, T. M., Reingold, E. M., 1991. Graph-drawing by force-directed
 994 placement. *Software—Practice and Experience* 21 (11), 1129–1164.
- 995 Gajer, P., Goodrich, M. T., Kobourov, S. G., 2004. A multi-dimensional ap-
 996 proach to force-directed layouts of large graphs. *Computational Geometry*

- 997 29 (1), 3–18.
- 998 Gansner, E., Hu, Y., 2008. Efficient node overlap removal using a proxim-
 999 ity stress model. In: Proceedings of the 16th International Symposium on
 1000 Graph Drawing (GD’08). Vol. 5417 of Lecture Notes in Computer Science.
 1001 Springer-Verlag, pp. 206–217.
- 1002 Gansner, E., Koren, Y., North, S., 2004. Graph drawing by stress majorization.
 1003 In: Proceedings of the 12th International Symposium on Graph Drawing
 1004 (GD’04). Vol. 3383 of Lecture Notes in Computer Science. Springer-Verlag,
 1005 pp. 239–250.
- 1006 Geipel, M. M., 2007. Self-organization applied to dynamic network layout.
 1007 International Journal of Modern Physics C 18 (10), 1537–1549.
- 1008 Groh, G., Hanstein, H., Wörndl, W., 2009. Interactively visualizing dynamic
 1009 social networks with DySoN. In: Proceedings of the IUI’09 Workshop on
 1010 Visual Interfaces to the Social and Semantic Web (VISSW’09). Vol. 443.
 1011 CEUR Workshop Proceedings.
- 1012 Hachul, S., Jünger, M., 2004. Drawing large graphs with a potential-field-based
 1013 multilevel algorithm. In: Proceedings of the 12th International Symposium
 1014 on Graph Drawing (GD’04). Vol. 3383 of Lecture Notes in Computer Sci-
 1015 ence. Springer-Verlag, pp. 285–295.
- 1016 Herman, I., Melançon, G., Marshall, M. S., 2000. Graph visualization and
 1017 navigation in information visualization: a survey. IEEE Transactions on
 1018 Visualization and Computer Graphics 6 (1), 24–43.
- 1019 Huang, M. L., Eades, P., Wang, J., 1998. On-line animated visualization of
 1020 huge graphs using a modified spring algorithm. Journal of Visual Languages
 1021 and Computing 9 (6), 623–645.
- 1022 Hunter, D. R., Goodreau, S. M., Handcock, M. S., 2008. Goodness of fit
 1023 of social network models. Journal of the American Statistical Association
 1024 103 (481), 248–258.
- 1025 Kamada, T., Kawai, S., 1988. A simple method for computing general positions
 1026 in displaying three-dimensional objects. Computer Vision, Graphics and
 1027 Image Processing 41, 43–56.
- 1028 Kaufmann, M., Wagner, D. (Eds.), 2001. Drawing Graphs: Methods and Mod-
 1029 els. Vol. 2025 of Lecture Notes in Computer Science. Springer.
- 1030 Klovdahl, A. S., 1981. A note on images of networks. Social Networks 3, 197–
 1031 214.
- 1032 Koskinen, J. H., Snijders, T. A. B., 2007. Bayesian inference for dynamic
 1033 social network data. Journal of Statistical Planning and Inference 137 (12),
 1034 3930–3938.
- 1035 Kruskal, J., Seery, J., 1980. Designing network diagrams. In: Proceedings of
 1036 the First General Conference on Social Graphics. pp. 22–50.
- 1037 Kruskal, J. B., Wish, M., 1978. Multidimensional Scaling. Vol. 07-011 of Sage
 1038 University Paper series on Quantitative Applications in the Social Sciences.
 1039 Sage Publications.
- 1040 Lyons, K. A., Meijer, H., Rappaport, D., 1998. Algorithms for cluster busting
 1041 in anchored graph drawing. Journal of Graph Algorithms and Applications

2 (1), 1–24.

Misue, K., Eades, P., Lai, W., Sugiyama, K., 1995. Layout adjustment and the mental map. *Journal on Visual Languages and Computing* 6 (2), 183–210.

Moody, J., McFarland, D., Bender-deMoll, S., 2005. Dynamic Network Visualization. *American Journal of Sociology* 110 (4), 1206–1241.

Moreno, J. L., 1953. *Who Shall Survive: Foundations of Sociometry, Group Psychotherapy, and Sociodrama*. Beacon House.

Nesbitt, K. V., Friedrich, C., 2002. Applying gestalt principles to animated visualizations of network data. In: *Proceedings of the 6th International Conference on Information Visualisation (IV’02)*. IEEE Press, pp. 337–343.

North, S. C., 1996. Incremental layout with DynaDag. In: Brandenburg, F. J. (Ed.), *Proceedings of the 3rd International Symposium on Graph Drawing (GD ’95)*. Vol. 1027 of *Lecture Notes in Computer Science*. Springer, pp. 409–418.

Purchase, H. C., Cohen, R. F., James, M., 1997. An experimental study of the basis for graph drawing algorithms. *ACM Journal of Experimental Algorithmics* 2 (4).

Purchase, H. C., Samra, A., 2008. Extremes are better: Investigating mental map preservation in dynamic graphs. In: *Proceedings of the 5th International Conference on Diagrammatic Representation and Inference (Diagrams 2008)*. Vol. 5223 of *Lecture Notes in Computer Science*. Springer-Verlag, pp. 60–73.

Ripley, R., Snijders, T. A., 2011. *Manual for siena version 4.0* (version february 6, 2011). Oxford: University of Oxford, Department of Statistics; Nuffield College. <http://www.stats.ox.ac.uk/siena/>.

Saffrey, P., Purchase, H., 2008. The “mental map” versus “static aesthetic” compromise in dynamic graphs: A user study. In: *Proceedings of the 9th Australasian User Interface Conference (AUIC 2008)*. pp. 85–93.

Schweinberger, M., 2010. Statistical modeling of network panel data: Goodness-of-fit. submitted for publication.

Sibson, R., 1978. Studies in the robustness of multidimensional scaling: Procrustes statistics. *Journal of the Royal Statistical Society. Series B (Methodological)* 40 (2), 234–238.

Snijders, T. A. B., 2001. The statistical evaluation of social network dynamics. *Sociological Methodology* 31, 361–395.

Snijders, T. A. B., 2005. Models for longitudinal network data. In: Carrington, P. J., Scott, J., Wasserman, S. (Eds.), *Models and Methods in Social Network Analysis*. New York: Cambridge University Press, pp. 215–247.

Snijders, T. A. B., Koskinen, J., Schweinberger, M., 2010a. Maximum likelihood estimation for social network dynamics. *Annals of Applied Statistics* 4 (2), 567–588.

Snijders, T. A. B., Pattison, P. E., Robins, G. L., Handcock, M. S., 2006. New specification for exponential random graph models. *Sociological Methodology* 36, 99–153.

Snijders, T. A. B., Steglich, C. E. G., Schweinberger, M., 2007. Modeling

1087 the co-evolution of networks and behavior. In: van Montfort, K., Oud, H.,
1088 Satorra, A. (Eds.), Longitudinal Models in Behavioral and Related Sciences.
1089 Lawrence Erlbaum, pp. 41–71.

1090 Snijders, T. A. B., van de Bunt, G. G., Steglich, C. E. G., 2010b. Introduction
1091 to stochastic actor-based models for network dynamics. *Social Networks*
1092 32 (1), 44–60.

1093 Steglich, C. E. G., Snijders, T. A. B., Pearson, M., 2010. Dynamic networks
1094 and behavior: Separating selection from influence. *Sociological Methodology*
1095 40, 329–393.

1096 Torgerson, W. S., 1952. Multidimensional scaling: I. Theory and Method. *Psy-*
1097 *chometrika* 17, 401–419.

1098 Walshaw, C., 2001. A multilevel algorithm for force-directed graph drawing.
1099 In: Marks, J. (Ed.), *Proceedings of the 8th International Symposium on*
1100 *Graph Drawing (GD 2000)*. Vol. 1984 of *Lecture Notes in Computer Science*.
1101 Springer, pp. 171–182.