A Tale of Two Communities: Assessing Homophily in Node-Link Diagrams

Wouter Meulemans[1(✉)] and André Schulz[2(✉)]

[1] giCentre, City University London, London, UK
wouter.meulemans@city.ac.uk
[2] LG Theoretische Informatik, FernUniversität in Hagen, Hagen, Germany
andre.schulz@fernuni-hagen.de

Abstract. Homophily is a concept in social network analysis that states that in a network a link is more probable, if the two individuals have a common characteristic. We study the question if an observer can assess homophily by looking at the node-link diagram of the network. We design an experiment that investigates three different layout algorithms and asks the users to estimate the degree of homophily in the displayed network. One of the layout algorithms is a classical force-directed method, the other two are designed to improve node distinction based on the common characteristic. We study how each of the three layout algorithms helps to get a fair estimate, and whether there is a tendency to over or underestimate the degree of homophily. The stimuli in our experiments use different network sizes and different proportions of the cluster sizes.

1 Introduction

Networks do not exist without a surrounding context. An object in a network is typically equipped with a set of characteristics (e.g., age, race, or gender in a social network). These characteristics have an influence on the network structure; often, nodes of a network are partitioned into clusters, based on (some) characteristics. Detecting, measuring and understanding network structures and dependencies is an important task in network analysis. In social networks one of these effects is *homophily*. Simply speaking, homophily is a principle that asserts that individuals are more likely to have relationships with similar individuals [10]. There are two main mechanisms behind homophily: (i) individuals with the same characteristics might have a stronger tendency to form relationships (this principle is known as *selection*), (ii) individuals change their behavior to align with their friends (this is known as *socialization*) [3, Sect. 4.2].

Whether homophily is present in a certain network (and to what extent) can be detected by comparing the number of links between nodes of the same cluster (same-cluster) with the number of links between nodes of different clusters (cross-cluster); see Sect. 1.1 for details. This yields a simple formula for the degree of homophily in a network. In this work we study the following questions.

Question 1. Can an observer detect the degree of homophily in a node-link diagram of a network? Is there a tendency for overestimation or underestimation?

There exist many layout algorithms for node-link diagrams. We expect that the drawing style has a big impact on answering Question 1. Hence, a natural subsequent question is:

Question 2. Which node-link diagram layout is best suitable for detecting homophily? Are there general design principles to improve homophily detection?

We deliberately set the scope to node-link diagrams, since they are probably the most popular style for visualizing networks. Tasks like path tracing can be performed well on node-link diagrams and many users are familiar with their methodology. Other methods for displaying networks (e.g., matrix views, hive plots [9], NodeTrix [6]) or summarizing them (e.g., histograms) may be more effective for enabling homophily assessment, but are out of scope for our study.

For a fair evaluation of the layout methods we include one additional task in our experiments. In particular, we ask for participants to answer shortest-path questions, to detect whether a better homophily assessment diminishes the versatility of the layout and makes other (path-tracing) tasks harder.

Homophily can also occur when there are more than two clusters. More clusters make it harder to detect homophily, simply because there is more information. Moreover, the notion of homophily can be extended to more clusters in slightly different ways. Due to these considerations we decided to restrict our investigations to the most basic case with only two clusters.

Our studies are not necessarily tied to *social* networks. Clusters exist in all kind of networks and it is a natural question to ask whether there exists a bias for cross-cluster or same-cluster links.

1.1   Homophily

Homophily is a natural phenomenon, but it is not always present in (social) networks. In fact, opposite effects might occur, that is, individuals favor to form bonds with individuals that have different characteristics. To understand networks within their surrounding context, we need methods to detect and to measure the effect of homophily with respect to a certain characteristic. We follow the presentation of Easley and Kleinberg [3] to derive such a framework.

Suppose we study a social network in which the individuals are either female or male. We want to decide whether there exists homophily with respect to gender. Assume that a fraction p of the population is male and a fraction $q = 1-p$ is female. In a network without homophily we expect that a random link is male-male with a probability of p^2, female-female with a probability of q^2, and cross-gender with a probability of $2pq$. As a consequence there is evidence for homophily, if the fraction of cross-gender links is considerably less than $2pq$ (and "heterophily" if the fraction is considerably more). We use this "homophily test" to derive a measure for the *degree of homophily* of a network as follows:

Definition 1. *Given a network with two clusters (one with a fraction p of the nodes and the other with a fraction $q = 1 - p$). We say that the* degree of homophily *is 0 if there are only cross-cluster links, it is 1 if there are no cross-cluster links, and it is 1/2 if there is no homophily, that is we have 2pq cross-cluster links. For all other situations we linearly interpolate between these values.*

1.2 Network Visualization

There exists a vast literature on various methods for network visualization and analysis with clusters, hierarchies or other auxiliary data, e.g. [2,7,12]. Though an extensive review is out of scope, we briefly review a number of methods and describe those we used for our user study.

Force-Directed. The *force-directed* layout (see for example [4]) has been a popular network visualization method since its inception. The core idea of this method is to mimick a physical system in which nodes repel each other and links behave like springs, pulling their ends together. We include this in our study, providing a good baseline visualization method to which to compare. We use the implementation provided by the javascript library D3.js[1].

Polarized. It is straightforward to modify the classical force-directed method to pull the nodes of the two clusters apart. We modified the D3.js layout algorithm by adding a force that moves the nodes left or right, depending on the cluster. We refer to this as the *polarized* layout. This tends to pull the clusters apart, though a clear separation is not guaranteed.

Bipartite. As an extreme form of separating the clusters (*bipartite* layout) we place all vertices of one cluster equidistantly on a vertical line on the left, the other vertices on a vertical line on the right. Cross-cluster links are drawn as straight-line segments forming a 2-layer bipartite drawing of this subnetwork. Same-cluster links are drawn as semi-circles as in an arc diagram. After obtaining an initial ordering of the vertices from a barycentric layout, we use the method of Baur and Brandes [1] to reduce the number of crossings, within one round applying sifts to all nodes of one cluster before the other. We remark that this layout style can also be found as an unchosen design alternative in the interactive visualization system described by Ghani et al. [5]. Their reason for not using this design is that they aim to support many clusters: the same-cluster links drawn as arcs would lead to severe clutter.

Other Methods. Many other methods exist for visualizing (clustered) networks. An example of a method we also considered is by Jusufi et al. [8]. Though potentially useful for assessing homophily, the resulting layouts appear suitable mostly for high-level overview tasks; path tracing is likely to be difficult due to the bundling of the edges. Hence, we did not include this method in our study.

---

[1] http://www.d3js.org/

2   Experimental Design

2.1   Hypotheses

With this user study, we wish to investigate the following three hypotheses:

H1 Homophily assessment is easiest with Bipartite layouts, followed by Polarized
layouts and hardest with Force-Directed layouts.

H2 It is harder to assess homophily in networks with differently sized clusters.

H3 Finding a shortest path between two nodes is easiest with Force-Directed
layouts, followed by Polarized layouts and hardest with Bipartite layouts.

Underlying our main hypothesis (H1), driving the design of this experiment,
is the idea that visually separating the node clusters makes it easier to assess
homophily. By pulling the nodes apart, we separate the same-cluster and cross-
cluster links, thus potentially making it easier to assess the ratio between them.
Cluster separation is stronger for the Bipartite layout than for the Polarized
layout; it is not taken into account for Force-Directed layouts at all.

Whether there is homophily, depends not only on the ratio between same-
and cross-cluster links, but also on the relative size of the two clusters. We
hypothesize that it is easier to assess homophily when the two clusters are of
equal size: we may then simply assess whether there are more or less same-cluster
links in comparison to cross-cluster links.

Cluster separation may have a negative effect on tasks that are not influenced
by the clusters, such as path-tracing tasks. We instantiate this by considering
the task of finding the shortest path between two highlighted nodes. The cluster
separation may pull neighboring nodes apart, causing longer links in the visu-
alization. Such links become harder to follow. Moreover, switch-backs between
the two clusters may be counterintuitive to the idea of a "shortest" path.

2.2   Method

Tasks. We used two Tasks: Bias and Path. The Bias task is targeted at Hypothe-
ses H1 and H2, asking participants to assess the homophily of a network. How-
ever, we avoided the use of the term homophily and used an informal descrip-
tion of "Bias" to avoid different behavior between people that knew homophily
beforehand and those that did not. In this paper, we shall use Bias to refer to
a participant's assessment and homophily for calculated values. For answering
Bias trials, participants were given a slider that internally allowed specifying a
value between 0 % (only cross-cluster links) and 100 % (only same-cluster links),
though no numbers were shown. The aim of the task was for participants to
estimate bias, without precisely counting nodes and links; they were instructed
accordingly. However, no time limit was given.

The Path task targets Hypothesis H3, asking participants to find the length
of a shortest path between two nodes in a given network. To provide their answer,
participants were given 5 radio buttons (for 2 to 6 steps). They were instructed
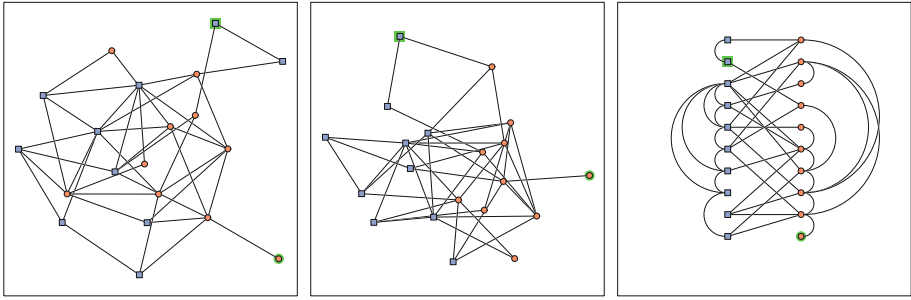
Fig. 1. Three stimuli used, all with Size 1, Balance B and Homophily 50 %. From left to right: force-directed (FD), polarized (P) and bipartite (B). The highlighted nodes were not highlighted for the Bias task.

to balance between answering correctly and answering quickly, ideally within 20 s. But it was mentioned explicitly that the time limit was not enforced.

The two tasks give rise to two sections in the study. To counter learning effects, the order of the sections was determined randomly per participant. Each section was preceded by a page explaining the task and an example question.

Stimuli. We have the following four independent variables for the stimuli:

– **Size.** Four different sizes: (1) 20 nodes, 40 links; (2) 20 nodes, 50 links; (3) 28 nodes, 60 links; (4) 40 nodes, 70 links;[2]
– **Balance.** Two ways to split into two clusters: (B) Balanced, an even split; (U) Unbalanced, one cluster contains 75 % of the nodes.
– **Homophily.** Five levels of degrees of homophily: 25 %; 37.5 %, 50 %; 62.5 %; 75 %;
– **Layout.** Three different layout algorithms (see Sect. 1.2 and Fig. 1): (FD) Force-Directed; (P) Polarized; (B) Bipartite.

To construct a network, we developed a simple random generator that takes as input the number of nodes for each cluster, the total number of links and the desired homophily. The desired homophily gives the fraction of links that should be cross-cluster; the remaining links were divided between the two clusters based on relative cluster sizes. The actual links added to the network were taken randomly (without replacement) from all possible links.

In each network, we also marked two arbitrary nodes (always one in each cluster) for the shortest-path task, controlling for the length of the shortest path. As we did not wish to introduce varying levels of difficultly for the shortest-path task, we need the stimuli to be of comparable difficulty, without always having the same number of steps as answer. We are mainly interested in the effects of Layout on task difficulty; using the same pair for each layout of the same network structure results in perfect balance. Nonetheless, we attempted to balance the

---

[2] Though small for social networks, we purposefully restricted to these sizes to ensure short trials and that path-tracing tasks would not become too difficult.

lengths across different levels of homophily and cluster balance to account for their possible effects on difficulty.

As a result, we have $4 \times 2 \times 5 = 40$ network structures. To each, we applied three Layouts; the resulting drawings were fitted to an SVG canvas, to allow for arbitrary resizing. The two node clusters were drawn using red circles and blue squares, the shade of the color chosen according to ColorBrewer[3]. Links were drawn in black with a small halo to increase separability between crossing links.

Mixed Design. With two tasks and 120 stimuli per task, we would need to give our participants 240 trials for a within-subjects design. This is far beyond what is reasonable for an online study, assuming 20 to 30 s per trial.

A between-subjects design is also not suitable, as we expect performance to be highly dependent on the participant's experience. Moreover, even if different people assess bias differently, it is possible that they are individually consistent: they perceive the homophily levels correctly, but assess the bias strength differently (see Fig. 2). To account for this Layout, Balance and Homophily are unsuitable as between-subjects factors.

Since network size is not directly of interest for our hypotheses and likely to be an obvious overall factor in increasing difficulty, we decided to use this as a between-subjects measure.

We now have 60 trials per participant. We aimed for a time investment of 20 to 25 min. A

Fig. 2. Bias versus Homophily for two fictive participants. The monotonicity of each line indicates per-participant consistency of bias assessments.

pilot study showed that with 60 trials, the actual completion time was around 30 min. Also, one of the pilot participants commented about the monotony of the questions, mentioning that less effort was put into the later trials for each task. We therefore decided to reduce the number of levels in Homophily-Balance interaction. We maintained the five levels of Homophily for the Balanced networks, but reduced it to three levels for Unbalanced networks. Maintaining five levels of Homophily for the Balanced networks provides a good baseline for investigating our main hypothesis on cluster separation and allows us to investigate for individual consistency. This reduced the number of trials to 48 (24 per Task). A second pilot study showed a completion time between 20 and 25 min as desired.

Again, to counter learning effects, the order of the 24 trials for each Task was randomized for each participant. Before each trial the participant was given a pause screen to reduce memory effects and at the same time allow them to pace themselves and reduce the possible impact of interruptions.

Apparatus. We developed our online user study, using a PHP webserver and a MySQL database. As is typical for online studies, we cannot control many aspects of the experimental environment (browser, OS, device, screen size,
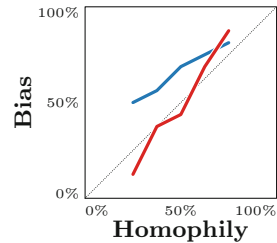
---

[3] http://www.colorbrewer2.org/.

interruptions, etc.). We requested participants to use a laptop or desktop, instead of a tablet or phone. They were also asked to avoid or minimize interruptions and indicate at the end of the study any that did occur. To ensure that the browser is appropriate for the user study, we gave them a simple test (setting a slider to a number indicated in a figure) before they could start the actual questions. Some background and preference information was asked after completing the two tasks, though this remained optional for what may be perceived as sensitive information (age, gender, country of residence).

We could not control the screen size, resolution or distance of the participant to his screen. Hence, participants were provided some simple controls to scale the webpage to be comfortably readable and fitting on their screen. Moreover, they were asked to use full-screen mode to reduce distractions.

Recruiting Participants. We recruited volunteers to participate using a mix of mailing lists, social networks and social media. Because we did not know how many people would participate in our study and we have a mixed design, we decided on the following procedure. The four levels of Size were to be filled up to 35 participants, in the following order: 2, 3, 1, 4. Any participants in excess of $4 \cdot 35 = 140$ would be divided equally over Size. If we would fall short of 140 participants, the participants would work on the same or a similar network size.

3   Results

The data set for analysis as well as all stimuli have been made available online[4].

Participants. In total, 105 people volunteered and completed the online questionnaire, which was open for participation for four weeks. We kept close watch at the number of participants and at the end of the second week, we were just in excess of 70 participants and thus decided to disable the last Size group (4).

After two weeks and continuously thereafter, we also inspected all comments left by participants. We excluded from analysis any participants for whom comments or timing indicated a serious interruption, distraction or technical difficulty during a trial, i.e., not during a pause screen or in between sections. Participants were explicitly asked to indicate their effort needed to distinguish nodes from different clusters and finding highlighted nodes; those who indicated having a hard time with this were excluded from analysis. In doing this exclusion while the study was open for participation, our online system assigned new volunteers to fill up the three remaining Size groups evenly. This resulted in three Size groups with 30 participants each.

58 participants are male, 28 are female and 4 did not specify a gender. The men/women ratio was even for Size 1, but the other Sizes have a ratio of approximately 3 : 1. Our participants are skewed towards male mathematicians and computer scientists. The average age of our participants is 35.9. Participants with Size 1 where older on average (39.1) and younger with Size 3 (33.9). In terms of country of residence, a majority of the participants live in Europe (61)
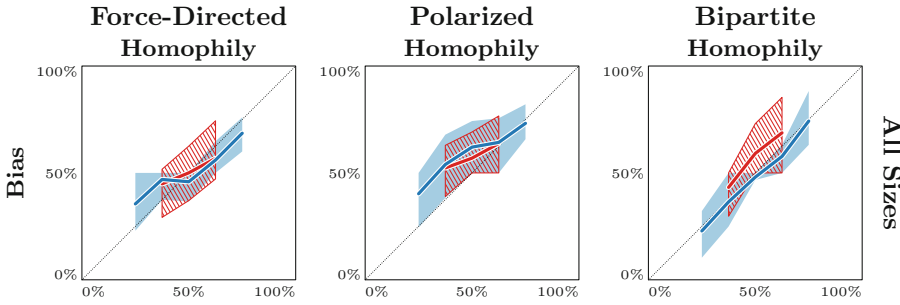

Fig. 3. Bias-Homophily charts. Thick lines indicate the average Bias; shaded blue (Balanced) and hashured red areas (Unbalanced) indicate the 25- and 75-percentile.

with a strong emphasis on Germany (27) and the Netherlands (22). Nine participants live in North America, three in Asia and one in South America; 16 participants did not provide a country of residence.

Hypothesis H1. The results of the bias estimation are summarized in Fig. 3. Each chart shows how the response Bias is correlated to the calculated degree of homophily. It suggests that the Bipartite layout leads to a stronger perception of Bias, with greater agreement (less variability) between the participants. For the Balanced case, the line is close to the diagonal, indicating that the average answer lies close to the calculated homophily. Notably, the Polarized layout and Unbalanced networks have greater variability; Polarized and Unbalanced-Bipartite lead to overestimating homophily (same-cluster links). The results for Balanced-Bipartite are centered on the diagonal, whereas the results for Balanced-Force-Directed are above the diagonal for Homophily below 50 % and below the diagonal for Homophily above 50 %. This suggests that the distinction between different levels of Homophily is clearer for the Bipartite layout.

We cannot simply classify answers as "correct" or "incorrect" as the participants were asked to give an estimation of Bias, without providing them a precise formula of how to determine it. Hence, we score each answer of Bias b for a stimulus with degree of homophily h with a deviation $b - h$. A positive value indicates an overestimation of cross-cluster links, whereas a negative value indicates an overestimation of same-cluster links. The deviation and response times for this task are summarized in Fig. 4. We performed RM-ANOVA on the deviation, to see if there are significant[5] differences of bias estimation for different layouts. The analysis showed a significant effect of Layout on deviation ($p < 0.001$). It also indicated an interaction effect of Layout and Size ($p < 0.001$). A posthoc Tukey HSD test with Bonferroni adjustment showed a significant difference between Layout FD and P ($p < 0.001$) and between Layout P and B ($p < 0.001$). However, no significant difference between Layout FD and B was found.

---

[5] Significance is reported using a p-value: it indicates the probability that our observation is incorrect, i.e., occurring by chance rather than the different conditions. For an accessible introduction to HCI experimental design and analysis, we refer to [11].
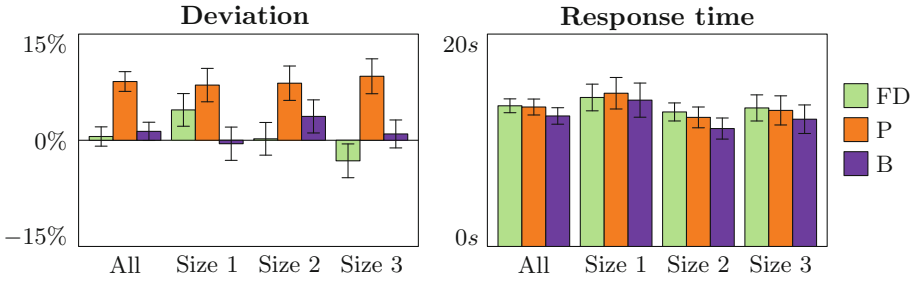
Fig. 4. Average deviation (left) and response time (right) for the Bias task, per Layout and Size. Error bars indicate 95 %-confidence intervals.
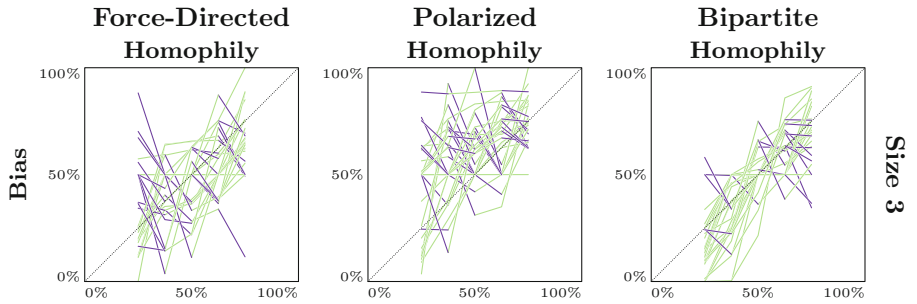

Fig. 5. Bias-Homophily charts for each Layout for Size 3, for the Balanced cases. Each line represents a participant. Monotonicity defects are colored dark purple.

A significant effect of Layout on response time ($p < 0.001$) was found, with the post-hoc test indicating a faster response time for the Bipartite layout.

We may thus partially accept Hypothesis H1: the Bipartite layout indeed outperforms the Polarized layout. However, because the Force-Directed layout outperforms the Polarized layout, it may be that our underlying argument for this hypothesis, i.e., cluster separation, is not the main effect in this difference.

The above focuses on the difference between Bias and Homophily and response time as an indicator of bias assessment. However, this does not readily mean that it is easier for a single participant to consistently assess bias. Let us now briefly turn towards an informal investigation of individual consistency (see also Sect. 2.2). If we chart each participant as a line in a Bias-Homophily plot, we would ideally see only monotonically increasing lines. The reality is of course different: Fig. 5 shows such a plot for Size 3. We observe that there are a lot less defects (decreasing parts of a line) from monotonicity for Layout B than for Layout FD and P. Over all Sizes, there are 90 defects for Layout FD, 113 for P and 60 for B; 24.4 % of the participants had no defect for Layout FD, 10 % for Layout P and 44.4 % for Layout B. This suggests a better homophily perception for the Bipartite layout, but the percentage of people without defects remains rather low. This may in part be explained by the lack of repetitions and

training tasks in our study. Unfortunately, this was unavoidable to keep a low time investment of the volunteers.

Hypothesis H2. To investigate the effects of Balance, we again refer to Fig. 3. We observe an increased variability for the Unbalanced cases. For Bipartite layout, we also observe a skew towards overestimating same-cluster links.

For the analysis, we filtered the 25 % and 75 % answers from the data set, as these levels were not used in the Unbalanced case. RM-ANOVA with the resulting data revealed a significant effect of Balance on deviation ($p < 0.05$) and response time ($p < 0.001$). We accept Hypothesis H2.

Hypothesis H3. The performance for the Path task is summarized in Fig. 6. Not surprisingly, RM-ANOVA revealed a significant effect of Size on error rate ($p < 0.001$). To investigate the effects of Layout, we therefore split the data into three subsets, one of each Size. Subsequent analysis of these sets revealed a significant effect of Layout on error rate ($p < 0.001$). The post-hoc test showed that the difference in error rate is significant between the Force-Directed layout and the Bipartite layout for each level of Size ($p < 0.01$). The Force-Directed layout also has a lower error rate for Size 2 ($p < 0.05$) and Size 3 ($p < 0.001$) compared to the Polarized layout. The Bipartite layout has a higher error rate than the Polarized layout for Size 1 ($p < 0.001$) and Size 2 ($p < 0.05$); a hint of a lower error rate was found in Size 3 ($p < 0.1$).

Further investigating this higher error rate, we found that four out of eight stimuli for P in Size 3 had an error rate of 80 % or more (i.e., worse than expected with random answers), suggesting a misleading visualization. After manual inspection of these stimuli, we attribute this to ambiguity of links that pass close by or even through unrelated nodes (see also Sect. 4).

Size does not have a significant effect on response time, but Layout does ($p < 0.001$). The post-hoc test showed a significant difference between the Bipartite layout and the others (both $p < 0.001$) as well as difference between the Force-Directed and Polarized layout ($p < 0.05$).

Combining the results of error rate and response time, we may accept H3: Force-Directed outperforms Polarized, which in turn outperforms Bipartite.
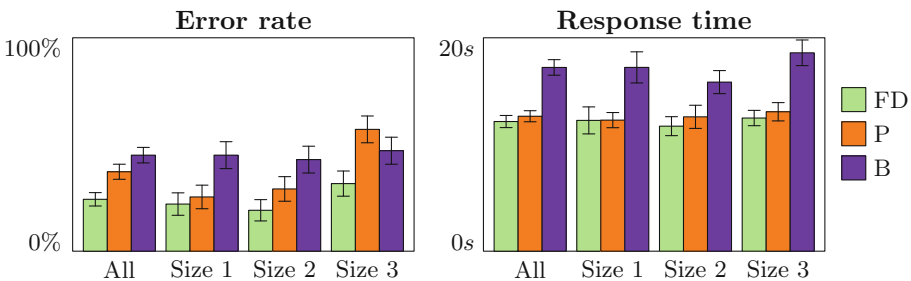

Fig. 6. Average error rate (left) and response time (right) for the Path task, per Layout and Size. Error bars indicate 95 %-confidence intervals.

Preferences. After completing the tasks, we asked the participants to indicate how hard they found it to perform the tasks with each type of network. However, the pilot study indicated that the distinction between the Force-Directed and Polarized layout was not clear while performing the tasks and thus hard to assess afterwards. As we did not want to introduce the three layouts beforehand, we opted to ask participants to rate (1–5) the Force-Directed layout and Bipartite layout only. The overall preference corresponds to the overall performance. For assessing bias, respondents clearly preferred the Bipartite layouts (mean $\mu = 3.5$, standard deviation $\sigma = 0.1$) over the Force-Directed and Polarized ones ($\mu = 2.0$, $\sigma = 0.1$). For finding shortest paths, this was reversed (Bipartite: $\mu = 2.4$, $\sigma = 0.1$; Force-Directed/Polarized: $\mu = 3.8$, $\sigma = 0.1$).

4   Discussion and Conclusion

Our results indicate that we may answer Question 1 positively: observers can indeed assess homophily, but this is affected by the layout and some layouts may in particular lead to an overestimation of homophily. We remark that individual consistency was not very strong, but this was leveled out by taking the average over all participants; see also the discussion below on training and repetition. To answer Question 2, the Bipartite layout performs best, followed by the classic force-directed method. The improved performance of the Bipartite layout, however, must be weighed against a loss in performance for other tasks.

Our results indicate that cluster separation by itself is not a general design principle to improve homophily perception. Future work may investigate such design principles in the context of homophily as well as further explore homophily with multiple clusters and defining a per-cluster homophily degree.

As with any user study, no experiment is flawless. We conclude our paper by discussing some aspects that may undermine our findings.

Bias Estimation. Explaining the Bias task was a difficult thing to do, without making the description overly long. In particular, we chose to go with a simple explanation of bias and not attempt to explain (degree of) homophily in detail. A participant's interpretation of Bias may thus inherently deviate from what is computed with our degree of homophily.

That the Bias task was rather difficult as a result, was also evidenced by some of the participants' comments. In particular, a few participants commented that the Unbalanced condition was hard to assess, further supporting Hypothesis H2.

Visual Representation. By using both color and shape, we tried to make the distinction between the clusters very clear. This supports the Bias task, but may in fact be detrimental for the Path task: participants may have had a tendency to look for connections between same colored nodes before looking at different colored nodes or vice versa. We think that this effect is mitigated by always selecting two nodes of different clusters. Also, network visualizations may simply need to display the different clusters for a variety of possible reasons, while still supporting the task of following paths well.

Whereas the Force-Directed and Polarized layouts use only line segments, the Bipartite layout uses circular arcs for same-cluster links. This difference in graphic encoding likely has a strong effect on distinguishing same-cluster from cross-cluster links, but this is in a large part already done by the clear cluster separation of the layout. However, it may also affect how easily participants can estimate the number of such links and hence affect the bias estimation. Moreover, it is likely to influence how easy it is to follow links for the Path task. As a result, it may be difficult to ascribe the results to the effect of cluster separation alone. This in particular affects for Hypothesis H1, but also means that for the other hypothesis, the effects may be in large part due to the layout as a whole, rather than any single aspect of it.

Training and Repetition. Due to our short intended time investment, there was little time for training the participants. They were given only one example question, before the trials started and each condition was presented only once. As a result, responses to earlier trials may be less accurate (or in case of bias assessment, less consistent). This is countered to some degree by randomizing the order of the trials for each participant. However, as indicated, it undermines evaluating an individual's responses. This is particularly the case for consistency of bias assessment, which is therefore done only informally and treated as an indication. However, plotting the deviation of bias assessment over time did not reveal clear learning effects.

References

1. Baur, M., Brandes, U.: Crossing reduction in circular layouts. In: Hromkovič, J., Nagl, M., Westfechtel, B. (eds.) WG 2004. LNCS, vol. 3353, pp. 332–343. Springer, Heidelberg (2004)
2. Eades, P., Feng, Q., Lin, X., Nagamochi, H.: Straight-line drawing algorithms for hierarchical graphs and clustered graphs. Algorithmica **44**, 1–32 (2006)
3. Easley, D., Kleinberg, J.: Networks, Crowds and Markets: Reasoning About a Highly Connected World. Cambridge University Press, New York (2010)
4. Fruchterman, T.M.J., Reingold, E.M.: Graph drawing by force-directed placement. Softw. Pract. Exp. **21**(11), 1129–1164 (1991)
5. Ghani, S., Kwon, B.C., Lee, S., Yi, J.S., Elmqvist, N.: Visual analytics for multi-modal social network analysis: a design study with social scientists. IEEE TVCG **19**(12), 2032–2041 (2013)
6. Henry, N., Fekete, J.-D., McGuffin, M.J.: Node trix: a hybrid visualization of social networks. IEEE TVCG **13**(6), 1302–1309 (2007)
7. Holten, D.: Hierarchical edge bundles: visualization of adjacency relations in hierarchical data. IEEE TVCG **12**(5), 741–748 (2006)
8. Jusufi, I., Kerren, A., Liu, J., Zimmer, B.: Visual exploration of relationships between document clusters. In: Proceedings of International Conference on Information Visualization Theory and Applications, pp. 195–203 (2014)

9. Krzywinski, M., Birol, I., Jones, S.J.M., Marra, M.A.: Hive plots–rational approach to visualizing networks. Briefings Bioinf. **13**, 627–644 (2011)
10. McPherson, M., Smith-Lovin, L., Cook, J.M.: Birds of feather: homophily in social networks. Ann. Rev. Sociol. **27**, 415–444 (2001)
11. Purchase, H.C.: Experimental Human-Computer Interaction: A Practical Guide with Visual Examples. Cambridge University Press, New York (2012)
12. van den Elzen, S., van Wijk, J.J.: Multivariate network exploration and presentation: from detail to overview via selections and aggregations. IEEE TVCG **20**(12), 2310–2319 (2014)