# Not just social networks: How people infer relations from mutual connections

**Abstract**

People can infer relationships from incomplete information about social networks. We examined whether these inferences depend on domain-specific knowledge about social relationships or instead depend on domain-general statistical reasoning. In five preregistered experiments, participants (total $N = 1424$) saw two target entities and their connections to others in social, semi-social, and non-social networks. In Experiments 1 and 2, participants made similar judgments across social and non-social networks: with greater proportion of mutual connections and number of connections, the two entities were judged as more likely to be connected to each other. These findings support the domain-general account. The next experiments provided further support for this account, while also investigating the question of whether people use mutual connections to infer the broader structure of networks. In Experiments 3 & 4, participants were asked whether entities connected to both targets were connected to each other, and judgments were hardly affected by network information. In Experiment 5, participants judged connections were more likely when entities were connected to both targets rather than when they were connected to only one. Overall, the findings support the domain-general account of network inferences, and further suggest that participants' inferences primarily concerned target entities and not the broader structure of the network.

*Keywords:* inferring relations, proportional information, network inferences, domain general

**Not just social networks: How people infer relations from mutual connections**

Websites like Facebook and LinkedIn send users recommendations for who to add to their networks. These notifications often tell users how many friends or connections they share with others (Li et al., 2017). People use this network information about mutual connections (i.e., friends in common) to infer social relationships (Sehl et al., 2023). They also misremember two individuals as friends with each other when they have a mutual friend (Freeman, 1992; see also DeSoto, 1960; Markovits & Dumas, 1999).

In this paper, we ask whether people's use of mutual connections to infer relationships depends on domain-specific knowledge about social relationships or instead reflects domain-general reasoning. Some prior work has examined whether people's processing of social networks is domain-specific or domain-general. For instance, recent work suggests that brain structures supporting the construction and retrieval of social networks also support other relational structures, like spatial representation (Parkinson & Wheatley, 2013; Peer et al., 2021; Schafer & Schiller, 2018; for a review, see Smith et al., 2020). This suggests that people encode, represent, and recall social networks using *generalized* cognitive mapping abilities. However, other work instead suggests domain-specificity. For instance, there is no correlation in people's performance in learning social and non-social networks (Tompson et al., 2019; for a review, see Momennejad, 2022). People also recall social networks more accurately (Janicik & Larrick 2005; Simpson et al. 2011), suggesting a benefit for processing the social domain. None of this work, though, has investigated how people infer missing connections in networks.

On a domain-specific account, when people use mutual connections to infer relationships, they do so by drawing on information specific to the social domain. In line with this account, people often infer relationship by drawing on information specific to this domain. For example,

young children infer friendships by attending to who individuals share secrets with (Liberman & Shaw, 2018) and favor when distributing resources (Liberman & Shaw, 2017). People also infer that individuals are more likely to be friends if they are similar on attributes such as language (Liberman et al., 2017), race (Roberts et al., 2017; Shutts et al., 2013), skills and experiences (Afshordi, 2019), and perceived personality traits (Alt et al., 2022; for reviews, see Basyouni & Parkinson, 2022; Liberman & Shaw, 2019).

Alternatively, on a domain-general account, people treat mutual connections as a purely statistical indicator of friendships. This possibility is broadly consistent with findings showing that statistical inferences underlie many other social judgments, including inferences about preferences (e.g., Kushnir et al., 2010; Heck et al., 2021), friendships (Son et al., 2021), group membership (e.g., Gershman et al., 2017; Vélez et al., 2019), and emotions (e.g., Doan et al., 2020; Asaba et al., 2019). For instance, people are more likely to judge individuals are friends if they share a rare preference over a common one (Vélez & Gweon, 2020), suggesting that friendships can be inferred from the statistics of the social environment.

To test between the domain-specific and domain-general accounts, we capitalize on findings from a recent study of how people infer relationships in social networks (Sehl et al., 2023). In this study, people inferred the likelihood that two target individuals were friends based on their other friends in the network. They judged targets were friends when they had a high proportion of mutual friends—that is, when targets' friends were mostly mutual rather than separate. People also judged targets were friends if they each had large social networks (i.e., many friends) rather than small ones.[1]

[1] Although these two kinds of information—proportion of mutual friends and overall number of friends—seem distinct, both might be processed by a single computational mechanism (Son et al., 2023).

These findings could reflect either domain-specific or domain-general reasoning about relationships. On the domain-specific account, people might infer that target individuals with many friends in common frequently encounter one another and have therefore had ample opportunity to get to know one another. People may likewise expect the individuals are friends if both are sociable or popular.[2] By contrast, on the domain-general account, people might reason statistically. For example, they might estimate how many mutual friends two target individuals should have by chance alone. If they have more mutual friends than expected, people may see this as a suspicious coincidence (Griffiths & Tenenbaum, 2007). They might conclude there is a reason for the high degree of social overlap, and explain this by positing that the target individuals are friends with one another.

To test between the accounts, we examine people's inferences about *non-social* networks, like websites that interlink and neurons that transmit chemicals to one other. If people's inferences about connections in non-social networks resemble their inferences about social networks, this implies that people rely on domain-general reasoning. We also investigate whether people use statistical information to infer the broader structure of the network. For example, if two targets have many mutual friends, people might infer that the targets and their mutual friends are all members of a clique, since this could explain why the targets have many more mutual friends than would be expected by chance alone (for work on inferring cliques, see Brashears & Quintane, 2015; Kilduff et al., 2008). Reasoning about cliques could be efficient—instead of tracking a host of individual relationships, it is more efficient to categorize individuals into social groups (Basyouni & Parkinson, 2022). Alternatively, people's inferences could be shallower. For instance, upon learning that two individuals have many mutual friends, people might directly

[2] Broadly consistent with this, people consider the number of friends Facebook users have to make a wide range of social judgments, such as their popularity and personality traits (Scott, 2014; see also Tong et al., 2008; Utz, 2010).

infer that they are likely to be friends, without considering or working out implications for other members of the network.[3]

## General Method

In five experiments, participants saw the networks of two target entities. Some networks were social—for example, two target people and other people they were friends with. Other networks were less social or completely non-social. For example, one kind of network involved connections between interlinked websites, and another kind involved connections between neurons in a worm brain. Our first experiments directly investigated the question about domain-specificity versus domain-generality. In these experiments, participants judged whether the two targets in the network were connected to one another in different domains. Our later experiments investigated whether participants infer broader aspects of network structure in varying domains, and in these studies participants instead judged whether the other members of the network were connected.

Preregistrations, materials, data, and analysis code for all experiments are available online on the Open Science Framework: https://osf.io/ezrxq/. All procedures and statistical analyses were preregistered prior to data collection. All reported analyses were preregistered. In all experiments, participants were recruited from CloudResearch (Litman et al., 2017). They were located throughout the United States and completed over 100 HITs with a 95% approval rating.

Participants saw network diagrams where two target entities had lines connecting them to other entities in the network. Participants first read instructions describing what the entities and

[3] This question may speak to the domain-generality or domain-specificity issue. People are likely to be familiar with social cliques, but may be less accustomed to cliques in other kinds of networks (Fortunato, 2010; Girvan & Newman, 2002). So, if their inferences depend on knowledge specific to the social domain, they might be especially likely to base inferences about whether networks include cliques on this social knowledge.

lines represented. All between-subjects conditions were randomly assigned. For each, they judged the likelihood that two entities were connected to each other on a 7-point scale from 1 (extremely unlikely) to 7 (extremely likely).

Participants completed two post-test multiple-choice attention checks. In Experiments 1-2, participants were shown a sample diagram and reported how many entities were connected to each of the two targets. In Experiments 3-5, participants were asked to identify the judgment they made and what they saw in the diagrams. Participants were excluded if they failed either attention check or neglected to answer on one or more test trials.

In all experiments, the main analyses used linear generalized estimating equations models (GEEs). We used GEEs because they provide outputs similar to Type III ANOVAs while also accounting for within-subject correlations introduced by repeated measures (see Frank et al., 2023). Table 1 shows the factors of the main analyses in Experiments 1, 2, and 4, and their results. We used pairwise comparisons using a Bonferroni correction to follow-up on interactions, and conducted single-sample tests using intercept-only models.

In Experiments 1-3, we also next ran secondary analyses to test if responses in each domain depended on proportion or number. Table 2 shows the factors of these secondary analyses for Experiments 1-3 and their results.

## Experiment 1

### Method

**Participants.** We tested 282 participants ($M_{age}$ = 40, range = 20-75, 98 female, 182 male, 2 other). Data were excluded from an additional 56 participants.

**Materials and Procedure.** Participants saw networks for one of three domains (manipulated between-subjects): friendships between people, links between websites, and flights

between airports. There were four types of trials, yielded by crossing two within-subjects factors. The first factor was the number of connections—targets had 6 or 12 connections. The second factor was the proportion of mutual connections to the total number of connections—targets had had a high proportion (5:6 or ~83%) or a low proportion of mutual connections (1:2 or 50%), see Figure 1. Participants completed two trials from each type. They completed 8 test trials, and trial order was randomized across participants. For each trial, they judged the likelihood that the two target entities were connected to each other on a 7-point scale from 1 (extremely unlikely) to 7 (extremely likely). Test questions were specific to the domain; for example, in the friendship domain, participants judged, "How likely is it that the two highlighted people are friends?".

**Results**

Targets were more likely to be seen as connected when they each had 12 connections rather than 6, $p < .001$, and when they had a 5:6 proportion rather than 1:2, $p < .001$; see Figure 2. Judgments did not differ across domains, $p = .451$.

These main effects were qualified by an interaction between domain and proportion, $p < .001$. It resulted because the effect of proportion was greatest for people ($M_{diff} = 1.21$), then for websites ($M_{diff} = 0.88$), and then for airports ($M_{diff} = 0.69$, all $p$s $< .001$). There was also an interaction between domain and number, $p = .005$. This effect was also greatest for people ($M_{diff} = 0.65$), then for websites ($M_{diff} = 0.34$), and then for airports ($M_{diff} = 0.68$, all $ps < .001$).

We next ran the secondary analysis to test if responses in each domain depended on proportion or number. To this end, we ran 2x3 GEEs comparing judgments in the 6/12 and 5/6 trials across the domains. This comparison is informative because there are more mutual connections in 6/12 trials than 5/6 trials, however, there is a greater proportion of mutual connections in 5/6 trials than 6/12 trials. The analysis revealed that participants judged targets

were more likely to be connected in 5/6 trials than 6/12 trials, $p < .001$, providing evidence that judgments were based on proportion. There were no differences in judgments across domains, $p = .530$, but there was a domain by trial type interaction, $p = .003$. This interaction resulted because judgments were higher in the 5/6 trials than the 6/12 trials for people and websites, $p$s $< .001$, however there was no difference between these trials for airports, $p = .965$.

**Discussion**

Participants largely showed the same pattern across domains, suggesting that their use of proportion and number is not specific to the social domain. This said, participants showed a somewhat different pattern with airports. They may have reasoned that it would be redundant to have a direct flight between airports that fly to many common destinations—why have a direct flight between San Francisco and Milan, when they both fly to Chicago, Toronto, Venice, London, Munich, Frankfurt, and Zurich? We next examined judgments about other domains.

## Experiment 2

**Method**

**Participants.** We tested 244 participants ($M_{age} = 40$, range = 20-74, 102 female, 141 male, 1 other). Data were excluded from an additional 87 participants.

**Materials and Procedure.** The design and procedure were identical to Experiment 1. The only difference is that participants made judgments about diagrams in three other domains (manipulated between-subjects; see Figure 1). Two domains were non-social: electrical parts that are wired together, and neurons in a worm brain that transmit chemicals to each other. The third domain was people who worked on a project together. This network is social, but arguably less so than friendships—collaborations could reflect arbitrary assignments determined by an organization.

**Results**

Targets were more likely to be seen as connected when they each had 12 connections rather than 6, $p < .001$, and when they had 5:6 proportion rather than 1:2, $p < .001$; see Figure 2. Judgments also differed by domain, $p < .001$. Ratings were higher for electrical parts than for neurons and people, $p$s $\leq .028$, but did not differ between neurons and people, $p = .084$. No interactions were significant.

The secondary analysis revealed that ratings were higher in 5/6 trials than 6/12 trials, $p < .001$, providing support that participants' judgments were based on proportion. Judgments also differed by domain, $p < .001$. Ratings were higher for electrical parts than neurons and people, $p$s $\leq .013$, but there was no difference in ratings for neurons and people, $p = .053$. There was no significant interaction, $p = .383$.

**Discussion**

Participants showed the same response pattern in all three domains. Together with findings from Experiment 1, these results support the domain-general account: Participants used information about mutual connections and number of connections similarly for social and non-social networks alike. Findings only differed somewhat when participants considered flights between airports.

In the next experiments, we turned to our second question—do people use information about targets' mutual connections and number of connections to make broader inferences about their social network, such as whether the network contains cliques? To examine this, we investigated another kind of judgment: Are targets' mutual connections connected with one another, as would be expected if they are all members of the same clique or group? We

compared these with judgments of whether the targets are connected (as investigated in the previous experiments) to see if both kinds of judgments depend on the same information.

## Experiment 3

### Method

**Participants.** We tested 233 participants ($M_{age}$ = 40, range = 20-75, 102 female, 130 male, 1 other). Data were excluded from an additional 3 participants.

**Materials and Procedure.** Participants completed two trials for each cell in a 2x2 within-subjects design, manipulating the number of friends (6, 12) and proportion of mutual friends (1:2, 5:6), see Figure 3. Participants judged, "How likely is it that the two highlighted people are friends?". For half of participants, the two highlighted people were the targets (i.e., identical judgments as Experiment 1). For the other half of participants, the highlighted people were two of the targets' mutual friends (i.e., participants made judgments about non-targets). In these trials, the two highlighted mutual friends were always separated by one other person in the row.

In this experiment, we also switched to a slightly different style of diagram to depict the networks (i.e., one using curved lines rather than straight ones). As the results from the targets judgment show, we still observed the same pattern of results obtained with the original diagram style.

### Results

There were main effects of number (12>6 friends, $F(1) = 7.48$, $p = .006$), and proportion (5:6>1:2, $F(1) = 253.08$, $p < .001$, but not judgment type, $F(1) = 0.66$, $p = .417$, see Figure 3. The main effects were qualified by three significant two-way interactions, all $p$s < .001, but there was no significant three-way interaction. Given multiple interactions involving whether

judgments were about targets versus mutual friends, we followed-up on these results by separately examining them in two separate 2x2 GEEs.

For judgments about targets, there were main effects of number of friends (12>6, $F(1) = 30.12$, $p < .001$) and proportion (5:6>1:2, $F(1) = 335.95$, $p < .001$), with no interaction, $F(1) = 2.45$, $p = .117$. By contrast, for judgments about mutual friends, there was a main effect of proportion (5:6>1:2, $F(1) = 9.87$, $p = .002$), no main effect of number of friends, $F(1) = 2.52$, $p = .113$, and an interaction, $F(1) = 9.21$, $p = .002$. This interaction resulted because ratings were higher in 5:6 proportion trials when targets had 6 friends each rather than 12, $p < .001$. However, in 1:2 proportion trials, there was no difference in number of friends, $p = .672$.

The secondary analysis revealed that responses were greater in 5/6 trials than 6/12 trials, $p < .001$. There was no difference across judgment type, $p = .282$, however there was an interaction, $p < .001$. The interaction resulted because although both judgments were higher in the 5/6 than 6/12 trials, this effect was greater for judgments about targets, $M_{diff} = 0.93$, $p < .001$, than those about mutual friends, $M_{diff} = 0.30$, $p < .001$.

**Discussion**

Judgments that targets are friends with one another were *not* paralleled by judgments about their mutual friends, as would be expected if participants inferred cliques. Participants again judged targets were more likely to be friends when they had a higher proportion of mutual friends and a greater number of friends overall. With judgments about mutual friends though, neither pattern emerged. Given this, the findings suggest participants did not use social network information to make inferences about the broader structure of the network.

We next looked at other domains to further investigate inferences about whether mutual connections are connected to each other. The next experiment did not examine judgments about targets, as this was already investigated across domains in Experiments 1 and 2.

## Experiment 4

### Method

**Participants.** We tested 328 participants ($M_{age}$ = 42, range = 22-74, 162 female, 164 male, 5 other). Data were excluded from an additional 3 participants.

**Materials and Procedure.** Participants saw networks from one of three domains (manipulated between-subjects): people who worked on a project together, websites that link to each other, or neurons in a worm brain that transmit chemicals to each other. They completed two trials for each cell in a 2x2 within-subjects design, manipulating the number of connections (6, 12) and proportion of mutual connections (1:2, 5:6). However, this time, participants judged how likely it was that two mutual entities were connected to each other.

### Results

There were main effects of number (12>6 friends, $p$ = .001) and proportion (5:6>1:2, $p$ < .001); see Figure 4. There was also a main effect of domain, $p$ <.001, as ratings were lower for neurons than people and websites, $p$s < .001, but ratings for people and websites did not differ, $p$ = .991. There were no interactions.

### Discussion

Number and proportion influenced participants' judgments of whether mutual connections are connected with each other. The effects were very small, though. For instance, the *greatest* mean difference was only 0.21 (i.e., the effect of proportion for coworkers), whereas in judgments about targets from previous experiments, the *smallest* mean difference for proportion

was 0.61. This pattern suggests that participants' use of these factors to infer the broader network structure was limited at best.

Our final experiment further investigates if people make broader inferences about the network (i.e., inferences about network members besides the targets). Specifically, we examine if people base these judgments on another feature of networks: Whether entities in the network are mutually connected to both targets or only connected to one target.

## Experiment 5

### Method

**Participants.** We tested 337 participants ($M_{age}$ = 42, range = 20-78, 170 female, 161 male, 6 other/unspecified). Data were excluded from an additional 8 participants.

**Materials and Procedure.** Participants saw networks for one of three domains (manipulated between-subjects): people who were friends with each other, websites that link together, or neurons in a worm brain that transmit chemicals to each other, see Figure 5. There were four types of trials, yielded by crossing two within-subjects factors: the number of connections the target entities had (6, 12) and judgment type (mutual-connection or single-connection). For both judgment types, participants rated how likely it was that the two highlighted entities were connected. For mutual-connection judgments, the two highlighted entities were both mutual connections of the targets. For single-connection judgments, the two highlighted entities were connected to only one of the targets (i.e., connected to one target and not the other). Participants completed 8 test trials, two trials from each type.

### Results

A 2(number) x 2(judgment) GEE was conducted for each domains (people, websites, neurons). Each domain yielded similar results: participants' judgments were greater for mutual

connection than single connection judgments, *p*s < .001, and there was no effect of number, *p*s ≥ .295; see Figure 5. There were differences in the presence of interactions. For websites and neurons, there was no interaction between judgment and number, *p*s ≥ .295, but this interaction was significant in the people domain, $p = .003$. It resulted because mutual connection judgments were greater when people had 12 friends rather than 6, $M_{diff} = 0.15$, $p = .018$, but there was no difference in single connection judgments across number of connections, $p = .134$.

**Discussion**

Participants used entities' connections to infer the broader structure of the network. Participants were more likely to infer connections when both entities were mutually connected to targets compared with when entities were only connected to one target. This suggests that the entity's missing connection disrupted participants' detection of a clique or grouping. Furthermore, findings again did not substantially differ across domains—the only exception was a relatively small interaction in the friends domain which was not mirrored with websites or neurons. Overall, this suggests that these inferences arise from domain-general reasoning about networks.

## General Discussion

Overall, the findings support the domain-general account. We consistently found similar patterns of results across domains, and judgments about social domains (e.g., friendships) did not substantially differ from those for non-social domains (e.g., electrical parts and neurons). This is not to deny that people have domain-specific knowledge about social relationships—they clearly do (e.g., Afshordi & Liberman, 2021). Our results, however, show that people's inferences about networks do not depend on this domain-specific knowledge.

One concern for the domain-general account is that we did observe some differences across domains. One difference is that people were more sensitive to proportion than number in all domains except airports. Another difference is that ratings were overall higher in some domains than others. For example, ratings about neurons were overall lower than corresponding ratings in other domains. Also, in the final experiment, it was only in the friends domain where number of connections impacted judgments about mutual connections. All these differences may seem to speak against the domain-general account because they suggest domain-specific contributions to reasoning about networks. Even so, the many parallels we observed across domains (e.g., sensitivity to proportion over number), suggest a strong contribution from domain-general processes. Moreover, the few domain-specific effects we observed mostly do not support the idea that inferences about *social* networks are shaped by domain-specific knowledge. For example, participants treating proportion and number equally in the airports domain suggests an effect of domain-specific information about airports rather than social relationships. Furthermore, while some non-social domains had lower ratings overall, the main patterns remained consistent *within* each domain. For example, within each domain, participants were more likely to infer targets were connected when they had a high rather than low proportion of mutual connections.

Our findings also address the second question of whether people use mutual connections and other statistical information to infer the broader structure of the network. Overall, the results suggest that people's inferences are shallow—they only made nuanced inferences about targets' connection to one another, and not about relations between others in the network. That is, if participants had inferred cliques, then the factors that impacted their judgments about targets' connectedness (i.e., proportion and number) should have similarly affected judgments about

mutual connections. Instead, in Experiments 3 and 4, judgments were hardly affected by proportion and number. This conclusion of shallow inferences broadly fits with the proposal that people are limited in their ability to infer missing relationships in social networks, especially when inferring relationships in the broader network based on a subset of known relationships (Basyouni & Parkinson, 2022). Nonetheless, the final experiment revealed that people had some ability to draw distinctions about whether entities in the broader network are connected. This experiment examined people's judgments about mutually-connected entities, and also about entities only connected to one target each. Across all domains, participants judged that entities were more likely to be connected when both were connected to the targets, rather than when they were only connected to one target each.

Although we generally found similar inferences across domains, it remains possible that people expect them to systematically differ in some regards. For example, people might have expectations for how many connections and mutual connections different entities typically have. People may see social networks like friendships as having many connections, but see networks of neurons and websites to have fewer mutual connections. Furthermore, people's expectations about some domains could be uniform, where all entities are assumed to have the same number of connections (e.g., neurons), while others may be variable and depend on the entity in question (e.g., popular people versus lonely people). Notably, though, these potential differences across domains are entirely consistent with the use of domain-general reasoning.

## Declarations

### Funding

This work was supported by grants to SD and OF from the Natural Sciences and Engineering Research Council of Canada.

### Conflicts of interest/Competing interests

The authors have no competing interests to declare that are relevant to the content of this article.

### Ethics approval

These experiments were approved by the Office of Research Ethics at the University of Waterloo. The research in this paper was conducted in compliance with APA ethical guidelines.

### Consent to participate

Informed consent was obtained from all individual participants included in the study.

### Consent for publication

All participants agreed that their anonymized data may be further used for research purposes.

### Availability of data and materials

The data that support the findings of this study are openly available via the Open Science Framework and can be accessed at https://osf.io/ezrxq/.

### Code Availability

The code for data analysis are available in the Open Science Framework repository: https://osf.io/ezrxq/.

## References

Afshordi, N. (2019). Children's inferences about friendship and shared preferences based on reported information. *Child Development*, *90*(3), 719-727. https://doi.org/10.1111/cdev.13237

Afshordi, N., & Liberman, Z. (2021). Keeping friends in mind: Development of friendship concepts in early childhood. *Social Development*, *30*(2), 331-342. https://doi.org/10.1111/sode.12493

Alt, N. P., Parkinson, C., Kleinbaum, A. M., & Johnson, K. L. (2022). The face of social networks: Naive observers' accurate assessment of others' social network positions from faces. *Social Psychological and Personality Science*, *13*(1), 118-126. https://doi.org/10.1177/19485506211003723

Basyouni, R., & Parkinson, C. (2022). Mapping the social landscape: tracking patterns of interpersonal relationships. *Trends in Cognitive Sciences*, *26*(3), 204-221. https://doi.org/10.1016/j.tics.2021.12.006

Brashears, M. E., & Quintane, E. (2015). The microstructures of network recall: How social networks are encoded and represented in human memory. *Social Networks*, *41*, 113-126. https://doi.org/10.1016/j.socnet.2014.11.003

DeSoto, C. B. (1960). Learning a social structure. *The Journal of Abnormal and Social Psychology*, *60*(3), 417-421. https://doi.org/10.1037/h0047511

Fortunato, S. (2010). Community detection in graphs. *Physics Reports*, *486*(3-5), 75-174.

Frank, M., Braginsky, M., Cachia, J., Coles, N., Hardwicke, T. E., Hawkins, R. D., Mathur, M. B., & Williams, R. (2023). *Experimentology: An open science approach to experimental psychology methods.* https:// experimentology.io/

Freeman, L. C. (1992). Filling in the blanks: A theory of cognitive categories and the structure of social affiliation. *Social Psychology Quarterly*, 118-127. https://doi.org/10.2307/2786941

Gershman, S. J., Pouncy, H. T., & Gweon, H. (2017). Learning the structure of social influence. *Cognitive Science*, *41*, 545-575. https://doi.org/10.1111/cogs.12480

Girvan, M., & Newman, M. E. (2002). Community structure in social and biological networks. *Proceedings of the National Academy of Sciences*, *99*(12), 7821-7826. https://doi.org/10.1073/pnas.122653799

Janicik, G. A., & Larrick, R. P. (2005). Social Network Schemas and the Learning of Incomplete Networks.*Journal of Personality and Social Psychology, 88*(2), 348–364. https://doi.org/10.1037/0022-3514.88.2.348

Kilduff, M., Crossland, C., Tsai, W., & Krackhardt, D. (2008). Organizational network perceptions versus reality: A small world after all? *Organizational Behavior and Human Decision Processes*, *107*(1), 15-28. https://doi.org/10.1016/j.obhdp.2007.12.003

Li, Z., Fang, X., & Sheng, O. R. L. (2017). A survey of link recommendation for social networks: Methods, theoretical foundations, and future research directions. *ACM Transactions on Management Information Systems (TMIS)*, *9*(1), 1-26. https://doi.org/10.1145/3131782

Liberman, Z., & Shaw, A. (2017). Children use partial resource sharing as a cue to friendship. *Journal of Experimental Child Psychology*, *159*, 96-109. https://doi.org/10.1016/j.jecp.2017.02.002

Liberman, Z., & Shaw, A. (2018). Secret to friendship: Children make inferences about friendship based on secret sharing. *Developmental Psychology*, *54*(11), 2139. https://doi.org/10.1037/dev0000603

Liberman, Z., & Shaw, A. (2019). Children use similarity, propinquity, and loyalty to predict which people are friends. *Journal of Experimental Child Psychology*, *184*, 1-17. https://doi.org/10.1016/j.jecp.2019.03.002

Liberman, Z., Woodward, A. L., & Kinzler, K. D. (2017). Preverbal infants infer third-party social relationships based on language. *Cognitive Science*, *41*, 622-634.

Litman, L., Robinson, J., & Abberbock, T. (2017). TurkPrime.com: A versatile crowdsourcing data acquisition platform for the behavioral sciences. *Behavior Research Methods*, *49*(2), 433-442. https://doi.org/10.3758/s13428-016-0727-z

Markovits, H., & Dumas, C. (1999). Developmental patterns in the understanding of social and physical transitivity. *Journal of Experimental Child Psychology*, *73*(2), 95-114. https://doi.org/10.1006/jecp.1999.2496

Momennejad, I. (2022). Collective minds: Social network topology shapes collective cognition. *Philosophical Transactions of the Royal Society B*, *377*(1843), 20200315. https://doi.org/10.1098/rstb.2020.0315

Parkinson, C., & Wheatley, T. (2013). Old cortex, new contexts: Re-purposing spatial perception for social cognition. *Frontiers in human Neuroscience*, *7*, 645. https://doi.org/10.3389/fnhum.2013.00645

Peer, M., Hayman, M., Tamir, B., & Arzy, S. (2021). Brain coding of social network structure. *Journal of Neuroscience*, *41*(22), 4897-4909. https://doi.org/10.1523/jneurosci.2641-20.2021

Roberts, S. O., Williams, A. D., & Gelman, S. A. (2017). Children's and adults' predictions of black, white, and multiracial friendship patterns. *Journal of Cognition and Development*, *18*(2), 189-208. https://doi.org/10.1080/15248372.2016.1262374

Schafer, M., & Schiller, D. (2018). Navigating social space. *Neuron*, *100*(2), 476-489. https://doi.org/10.1016/j.neuron.2018.10.006

Shutts, K., Roben, C. K. P., & Spelke, E. S. (2013). Children's use of social categories in thinking about people and social relationships. *Journal of Cognition and Development*, *14*(1), 35-62. https://doi.org/10.1080/15248372.2011.638686

Simpson, B., Markovsky, B., & Steketee, M. (2011). Power and the perception of social networks. *Social Networks*, *33*(2), 166-171. https://doi.org/10.1016/j.socnet.2010.10.007

Scott, G. G. (2014). More than friends: Popularity on Facebook and its role in impression formation. *Journal of Computer-Mediated Communication*, *19*(3), 358-372. https://doi.org/10.1111/jcc4.12067

Sehl, C. G., Friedman, O., & Denison, S. (2023). The social network: How people infer relationships from mutual connections. *Journal of Experimental Psychology: General*, *152*(4), 925. https://psycnet.apa.org/doi/10.1037/xge0001330

Smith, E. B., Brands, R. A., Brashears, M. E., & Kleinbaum, A. M. (2020). Social networks and cognition. *Annual Review of Sociology*, *46*(1), 159-174. https://doi.org/10.1146/annurev-soc-121919-054736

Son, J. Y., Bhandari, A., & FeldmanHall, O. (2021). Cognitive maps of social features enable flexible inference in social networks. *Proceedings of the National Academy of Sciences*, *118*(39), e2021699118. https://doi.org/10.1073/pnas.2021699118

Son, J. Y., Bhandari, A., & FeldmanHall, O. (2023). Abstract cognitive maps of social network structure aid adaptive inference. *Proceedings of the National Academy of Sciences*, *120*(47), e2310801120. https://doi.org/10.1073/pnas.2310801120

Tompson, S. H., Kahn, A. E., Falk, E. B., Vettel, J. M., & Bassett, D. S. (2019). Individual differences in learning social and nonsocial network structures. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *45*(2), 253. http://dx.doi.org/10.1037/xlm0000580

Tong, S. T., Van Der Heide, B., Langwell, L., & Walther, J. B. (2008). Too much of a good thing? The relationship between number of friends and interpersonal impressions on Facebook. *Journal of Computer-Mediated Communication*, *13*(3), 531-549. https://doi.org/10.1111/j.1083-6101.2008.00409.x

Utz, S. (2010). Show me your friends and I will tell you what type of person you are: How one's profile, number of friends, and type of friends influence impression formation on social network sites. *Journal of Computer-Mediated Communication*, *15*(2), 314-335. https://doi.org/10.1111/j.1083-6101.2010.01522.x

Van Kreveld, D., & Zajonc, R. B. (1966). The learning of influence structures. *Journal of Personality*. https://doi.org/10.1037/e513662009-001

Vélez, N., Bridgers, S., & Gweon, H. (2019). The rare preference effect: Statistical information influences social affiliation judgments. *Cognition*, *192*, 103994. https://doi.org/10.1016/j.cognition.2019.06.006

Vélez, N., & Gweon, H. (2020). Preschoolers use minimal statistical information about social groups to infer the preferences and group membership of individuals. In B. Armstrong, S. Denison, M. Mack, & Y. Xu (Eds.), *Proceedings of the 42rd Annual Meeting of the Cognitive Science Society* (pp. 227-233). Cognitive Science Society.

**Table 1.** Results of the main 2x2x3 analyses in Experiments 1, 2, and 4.

| Main effects & interactions | *Experiment 1* | | *Experiment 2* | | *Experiment 4* | |
|---|---|---|---|---|---|---|
| | *F* | *p* | *F* | *p* | *F* | *p* |
| Number | 121.50 | < .001*** | 33.69 | < .001*** | 11.24 | < .001*** |
| Proportion | 326.95 | < .001*** | 128.30 | < .001*** | 41.15 | < .001*** |
| Domain | 0.80 | .451 | 7.79 | < .001*** | 4.77 | < .001*** |
| Number*domain | 5.31 | .005** | 1.02 | .362 | 0.16 | .850 |
| Proportion*domain | 9.63 | < .001*** | 1.16 | .313 | 0.05 | .952 |
| Number*proportion | 0.00 | .992 | 0.22 | .640 | 2.47 | .116 |
| Number*proportion*domain | 0.24 | .788 | 1.93 | .145 | 0.04 | .962 |

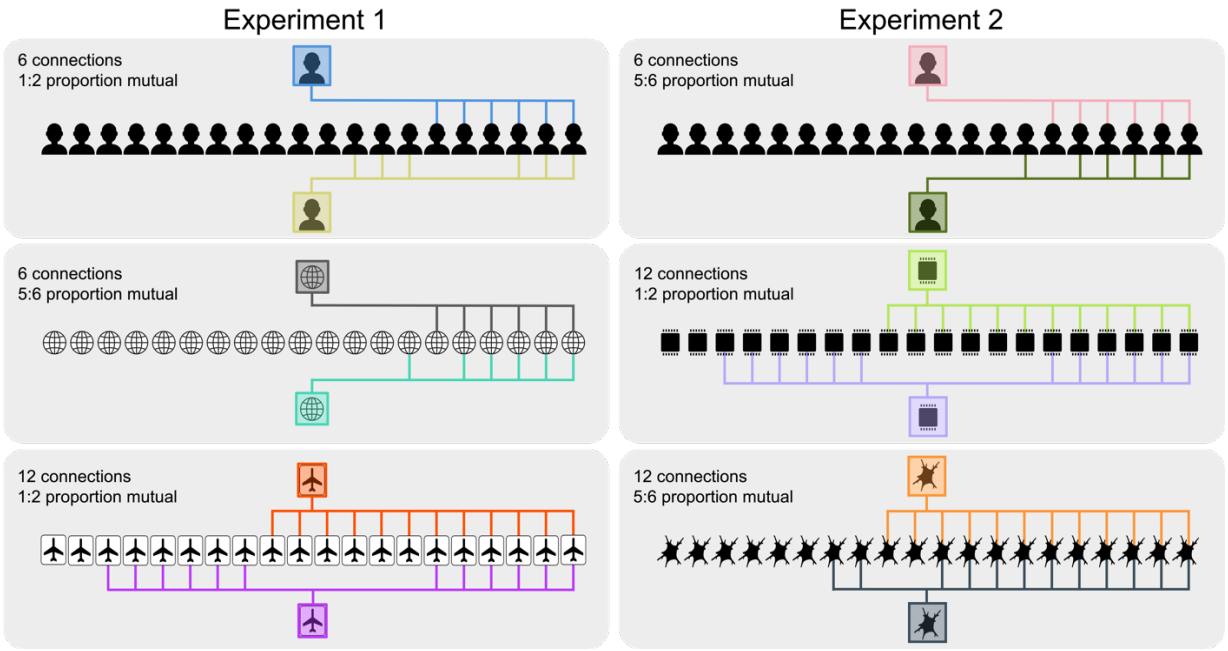
*Note.* All experiments featured a 2(number of connections) x 2(proportion of connections) x 3(domains) analysis. For all experiments, there were two levels for the number of connections (6, 12) and proportion of connections (1:2, 5:6). Domains differed across experiments: Experiment 1 had friendships, websites, airports; Experiment 2 had collaborators, electrical parts, neurons; Experiment 4 had collaborators, websites, and neurons. ** $p < .01$, *** $p < .001$.

**Table 2.** Results of the secondary 2x3 analyses in Experiments 1 and 2.

| Main effects & interactions | *Experiment 1* | | *Experiment 2* | | *Experiment 3* | | |
|---|---|---|---|---|---|---|---|
| | *F* | *p* | *F* | *p* | | *F* | *p* |
| Trial | 24.69 | .001*** | 15.27 | < .001*** | Trial | 92.96 | < .001*** |
| Domain | 0.64 | .530 | 9.19 | < .001*** | Judgment | 1.16 | .282 |
| Domain*trial | 5.67 | .003** | 0.96 | .383 | Trial*judgment | 24.56 | < .001*** |

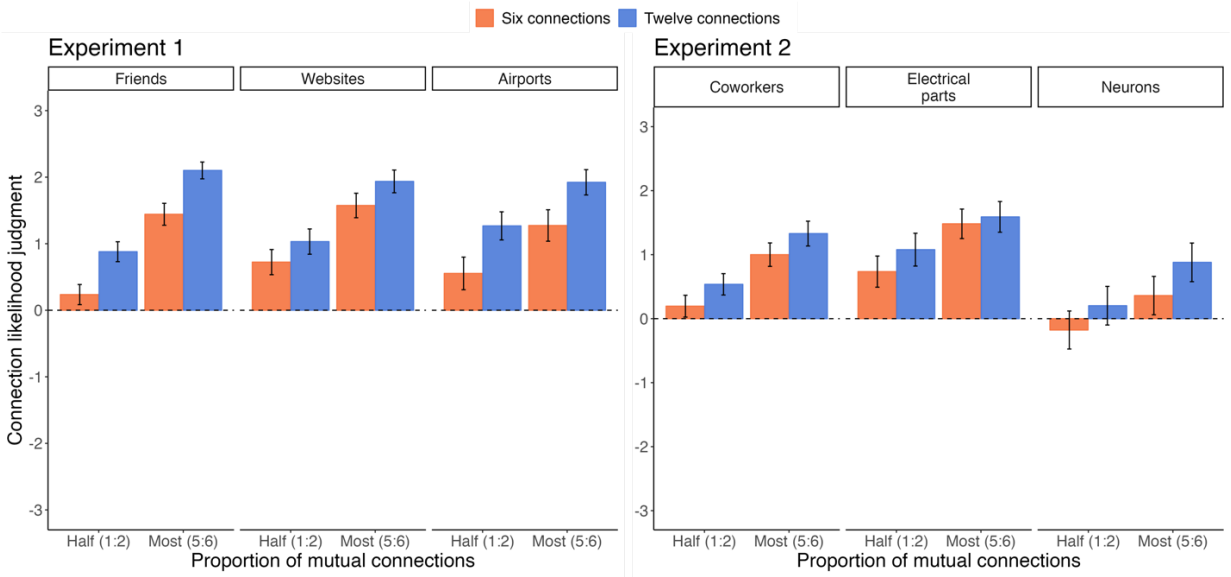
*Note.* Experiments 1 and 2 featured a 2(trial) x 3(domains) analysis. Domains differed across experiments: Experiment 1 had friendships, websites, airports; Experiment 2 had collaborators, electrical parts, neurons. Experiment 3 featured a 2(trial) x 2(judgment) analysis. For all experiments, there were two levels of trial (6/12 and 5/6). ** $p < .01$, *** $p < .001$.

**Figure 1.**



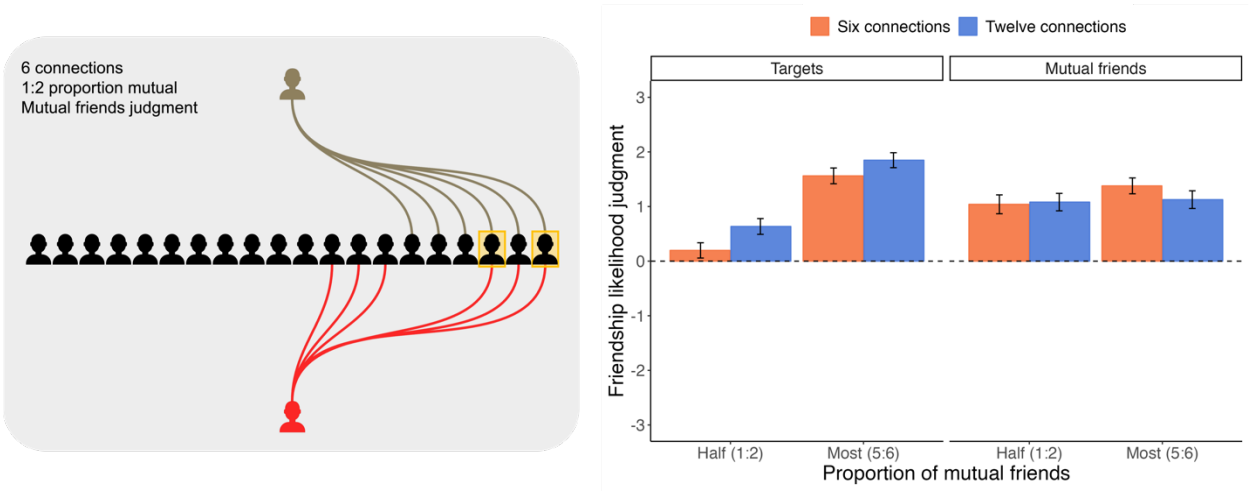

*Note.* Sample diagrams from Experiments 1 and 2. Networks in Experiment 1 featured (top-bottom): people who are friends, websites that are linked, and airports that have flights between them. Networks in Experiment 2 featured: people who worked on a project together, electrical parts that are wired together, or neurons in a worm brain that transmit chemicals to each other.

**Figure 2.**



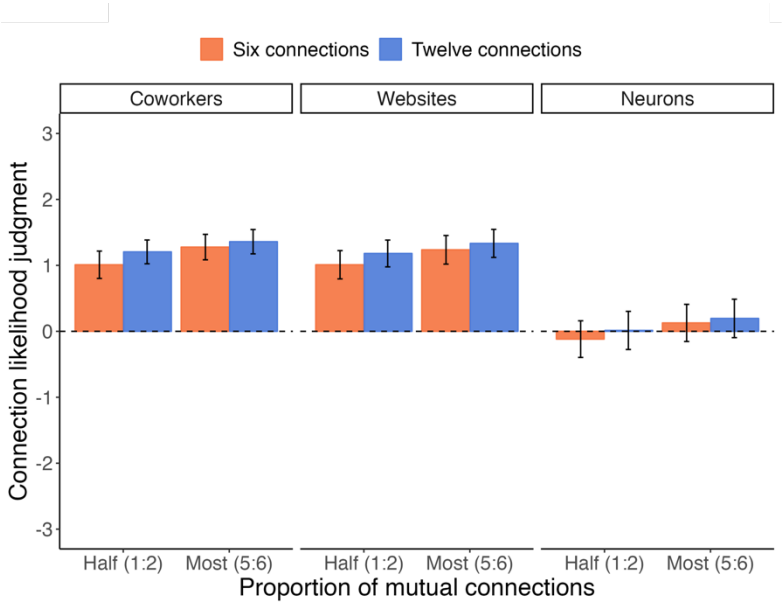

*Note*. Results from Experiment 1 (left) and 2 (right). Error bars in all figures represent 95% confidence intervals.

**Figure 3.**



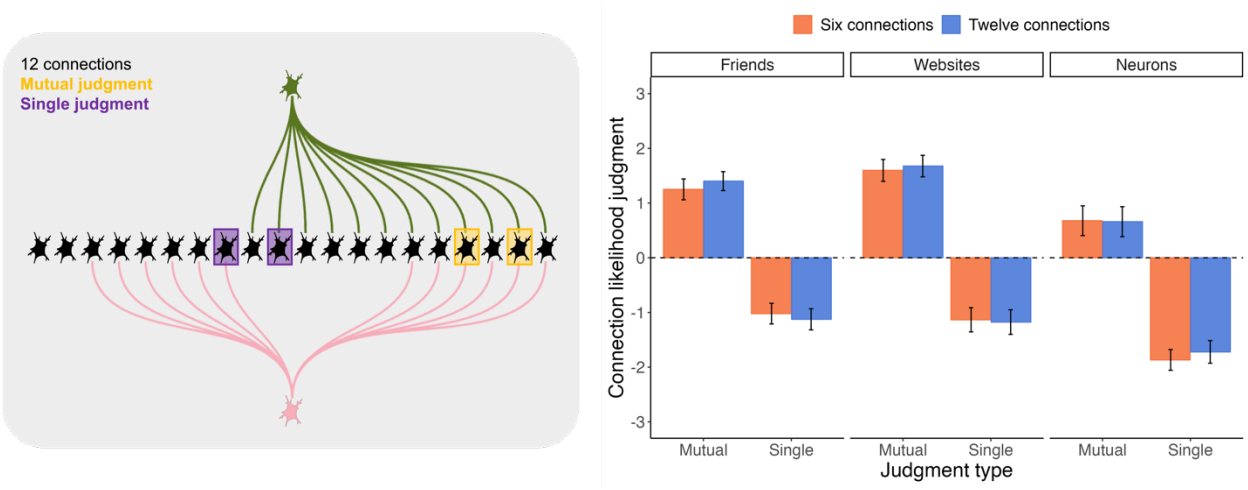

*Note.* Sample trial (left) and results (right) from Experiment 3. This sample trial shows a trial where participants were asked about two mutual friends.

**Figure 4.**




*Note.* Results from Experiment 4.

**Figure 5.**




*Note*. Sample trial and results from Experiment 5. Participants either saw targets *or* mutual friends highlighted. Highlights were always yellow but are purple in this figure to disambiguate judgment type.