PNAS RESEARCH ARTICLE | SOCIAL SCIENCES APPLIED MATHEMATICS

Triadic influence as a proxy for compatibility in social relationships

Miguel Ruiz-García[a,b,c,1,2], Juan Ozaita[c,1], María Pereda[b,d], Antonio Alfonso[e], Pablo Brañas-Garza[e], José A. Cuesta[b,c,f], and Angel Sánchez[b,c,f]

Edited by Susan Fiske, Princeton University, Princeton, NJ; received September 2, 2022; accepted February 14, 2023

Networks of social interactions are the substrate upon which civilizations are built. Often, we create new bonds with people that we like or feel that our relationships are damaged through the intervention of third parties. Despite their importance and the huge impact that these processes have in our lives, quantitative scientific understanding of them is still in its infancy, mainly due to the difficulty of collecting large datasets of social networks including individual attributes. In this work, we present a thorough study of real social networks of 13 schools, with more than 3,000 students and 60,000 declared positive and negative relationships, including tests for personal traits of all the students. We introduce a metric—the "triadic influence"—that measures the influence of nearest neighbors in the relationships of their contacts. We use neural networks to predict the sign of the relationships in these social networks, extracting the probability that two students are friends or enemies depending on their personal attributes or the triadic influence. We alternatively use a high-dimensional embedding of the network structure to also predict the relationships. Remarkably, using the triadic influence (a simple one-dimensional metric) achieves the best accuracy, and adding the personal traits of the students does not improve the results, suggesting that the triadic influence acts as a proxy for the social compatibility of students. We postulate that the probabilities extracted from the neural networks—functions of the triadic influence and the personalities of the students—control the evolution of real social networks, opening an avenue for the quantitative study of these systems.

social networks | triadic influence | relationship prediction | machine learning

Significance

Relationships are complicated. Individual features and the influence of other people can determine the fate of friendships. However, how rigorously can these effects be quantified? We have collected the relationships and personality traits of more than 3,000 students in 13 schools. We are able to identify the effect that personality and influence have in the construction of social networks and recover the probability of being friends or enemies depending on these variables. We postulate that the time evolution of social relations is dominated by the probabilities defined in this work.

Positive relationships help individuals thrive in society, whereas negative ones can jeopardize our chances of success and happiness. Social relationships arise from interactions between individuals and have been studied on different time scales and contexts (1, 2). As a result, social networks are formed, with individuals as nodes and interactions as links (3), and they can be studied and characterized using a complex network approach (4) in order to assess the many implications of social structure in our lives (5). A great deal of research has been carried out on social networks by aggregating the interactions that occur over a certain period of time to define links, starting from the pioneering work of Moreno (6). However, such an approach does not capture the dynamics of relationships, which is necessary to advance our understanding of the field (7). Large efforts have been devoted to this question in recent years, mainly using empirical data with different degrees of time resolution, such as, e.g., letter exchanges (8), mobile phone communications (9, 10), spatial mobility (11), or face-to-face interactions (12–14). (See also ref. 15 for a review.) All these analyses have led to many interesting insights into the evolution of relationships, but the issue of the mechanisms that explain how/why these relationships are created and evolve remains elusive.

Several models have been proposed to explain different aspects of the empirical observations. The first attempts were devoted to reproduce some of the structural properties observed in social networks, such as the small world phenomena (16) or the rich-get-richer effect (17, 18). Starnini et al. (19) proposed a simple model based on random walks and individual attractiveness to describe face-to-face interactions. For social networks, Jin et al. (20) studied networks with exponential decay of tie strengths to represent friendships. Other approaches have resorted to exponential random graph models (21) or stochastic actor-oriented models (22). Finally, regression models that incorporate a selection of individual traits have also been considered for online social networks (23). Still, none of these approaches sheds light on friendship formation in real life, taking into account the characteristics of the individuals and how some relationships can influence others.

In this paper, we contribute toward the understanding of friendship formation by adopting a different point of view, namely that of link prediction in networks (24).

Author contributions: M.R.-G., J.O., M.P., J.A.C., and A.S. designed research; M.R.-G. and J.O. performed research; M.R.-G. and J.O. contributed new reagents/analytic tools; M.R.-G. and J.O. analyzed data; J.A.C., A.S., A.A. and P.B.-G. collected data; and M.R.-G., J.O., M.P., A.A., P.B.-G., J.A.C., and A.S. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

[1]M.R.-G. and J.O. contributed equally to this work.

[2]To whom correspondence may be addressed. Email: miguel.ruiz.garcia@ucm.es.

This article contains supporting information online at http://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2215041120/-/DCSupplemental.

Published March 22, 2023.

The problem of link prediction, as originally formulated, is about temporal networks: Given the graph of connections between certain entities or nodes during some interval, the task is to predict the set of links in a later interval. Notwithstanding this definition, the same idea applies to many different situations, such as recommendation systems (25), bioinformatics (26), scientific collaboration networks (27), criminal networks (28), or even estimating the reliability of network data (29), to name a few. In the case of online social networks, link prediction has been considered, for example, by Song et al. (30) or Hao (31). (See ref. 32 for a review.) Much less has been explored regarding real-world social networks, in particular, friendship networks (33), due to the difficulty of collecting data on reasonably complete social networks that include personal attributes in real settings. For this reason, the discussion has been devoted in many cases to ego networks (i.e., data on disconnected individuals who mentioned their friends) and to the meaning of friendship (34).

In this work, we study social networks collected in 13 complete high schools in Spain, containing more than 3,000 individuals and 60,000 declared relationships between them. All students completed tests including information about their self-declared gender, cognitive results, and other variables that measured their *selfishness/prosociality*. Performing link prediction on this data, we are able to extract the probability that two students will be friends/enemies depending on their personalities. We also studied how this probability is affected by other relationships, defining a metric that we have termed triadic influence. Although we analyze static networks, our results suggest that the probabilities that we extract determine the mechanisms that control the initial formation of relationships and the evolution of the whole social network.

Results

Data collection was carried out in 13 schools in different areas of Spain, with a total of 3,395 students. They were asked to choose with whom they were related within their school by picking names from a school list. Then, they had to rate the relationship as very bad, bad, good, or very good, which we codified as -2, -1, $+1$, and $+2$, respectively. We recovered 60,566 declared relationships; *SI Appendix* for more details. In addition, we also collected data on the students' gender (self-reported), cognitive skills (measured by the cognitive reflection test, CRT), and their prosociality (*Methods* for details on these individual features). With this information, we build a directed weighted network, with each link representing a relationship that goes from the nominator to the nominee—two nodes can be connected by links in both directions—weighted by the reported rating. Additionally, each node represents one student and has his/her individual attributes (gender, CRT, and prosociality). Fig. 1 presents a sketch of the kind of social network that we will study. We have included several figures studying the structure of these social networks in *SI Appendix*, Figs. S1–S4.

In this work, we study the correlations between the personal features of both students and the type of relationship between them as well as the influence of other students on that relationship. We have used artificial neural networks to perform link prediction within our dataset from two complementary viewpoints: The first one focuses on the local structure, using the personality traits of both students and the influence of the nearest neighbors as described in the next section; the second one uses only the structural information of the network—the undirected and unweighted graph—to predict relationships. In what follows, we discuss these two approaches separately.

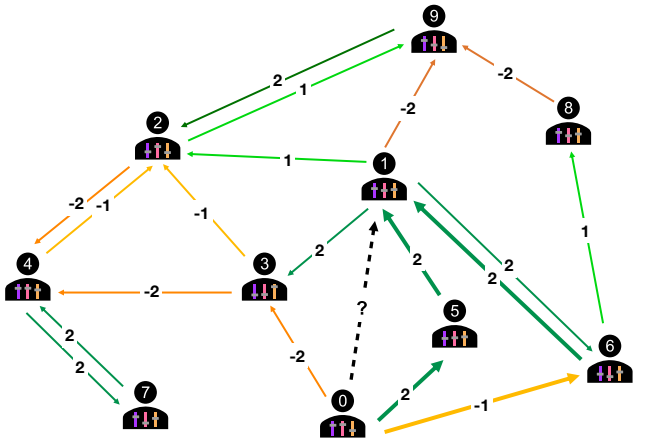

Fig. 1. Diagram of a social network that includes personality traits and computation of the triadic influence. To predict the relationship from node 0 to node 1, we can use the individual features of both students (represented by the sliders within their body) and/or the triadic influence I_{01}. The directions of these relationships are marked by arrows going from the nominator to the nominee, whereas the weight/intensity is represented with colors and edge labels (dark green, close friend; green, friend; yellow, dislike; orange, enemy). Thick arrows highlight the relationships that enter the calculation of I_{01}. To compute I_{01}, we select all directed paths of length 2 from node 0 to node 1 (0 $\rightarrow$ node $\rightarrow$ 1). In this example, they are 0-5-1 and 0-6-1. The path 0-3-1 is not a directed path (the direction of the edges is 0 $\rightarrow$ 3 $\leftarrow$ 1) and therefore is not included in the calculation of I_{01}. Thus, $I_{01} = w_{05}w_{51} + w_{06}w_{61} = 2 \cdot 2 + (-1) \cdot 2 = 2$.

Predicting with the Personality Traits and the Influence of the Nearest Neighbors. Fig. 1 shows a sketch of the social network with all the information available to perform link prediction. It shows the students (nodes) with their traits (sliders) and relationships of different types between them. In this section, we use only local properties of the network to predict the relationship between two students, namely the individual features of both students (e.g., nodes 0 and 1 in Fig. 1) and the directed weighted paths of length 2 between them. Specifically, we define a variable that we term triadic influence as $I_{ij} \equiv \left(W^2\right)_{ij} = \sum_k w_{ik}w_{kj}$, where w_{ik} is the weight of the link that goes from node i to node k (Fig. 1 for an example). The triadic influence condenses into one scalar the influence of third parties; e.g., if node i declares node k as a friend and k does the same with j, it adds a positive number to I_{ij} (your friend's friends are likely to be your friends), whereas a path containing links of the opposite sign will lead to a negative contribution (your enemy's friends or your friend's enemies are likely to be your enemies). I_{ij} adds up the contribution from all directed paths of length 2 between i and j. Interestingly, there is a connection between the concept of triadic influence and social balance theory that gives further insight into its meaning. Social balance theory (35–37) is an attempt to explain the dynamics of signed networks by classifying local motifs into stable or unstable. A key role in the theory is played by triangles: Triangles with an odd number of negative links (e.g., two persons who are enemies while sharing a common friend) are unstable, eventually evolving into a more balanced configuration by changing one link's sign or removing one link. In this context, the triadic influence adds up in one scalar the contribution of all the triads that are closed by that specific link, taking into account that our social network is weighted and directed. If I_{ij} is positive, it indicates that more triads will be socially balanced if the link ij is positive, and the opposite for a negative value of the triadic influence.

For simplicity, we will train a neural network to correctly classify all relationships in the network into two classes: friends

and enemies (*Methods* for more details). We used different combinations of the triadic influence and the individual characteristics of the students as input for the deep neural network (NN) and trained it to output the correct value for each relationship in the training dataset (*Methods* for a full description of the neural network and the training process). With our procedure, we obtain the probability that two students relate through a relationship belonging to one of the two classes (friends or enemies) as a function of the corresponding inputs. To avoid using a misleading metric of performance, since our classes are unbalanced—there are more declared friends than enemies—we assess the performance of our method using the balanced accuracy on the test dataset (38). To compute it, after training the NN, we feed it with all relations in the test dataset and assign the label "friend" or "enemy" to the class with the highest probability. The balanced accuracy is then computed as

$$\text{bAcc} = \frac{1}{2}\left(\frac{N_+^C}{N_+^T} + \frac{N_-^C}{N_-^T}\right),$$

where N_α^C is the number of samples belonging to class α (+ friend or − enemy) that were correctly classified from the total number of samples belonging to that class (N_α^T). This is more informative than other performance metrics because if either the NN classified everything in the same class or guessed at random, we would obtain bAcc $= 1/2$ regardless of the number of samples in each class, whereas if all relations were correctly predicted, then bAcc $= 1$ (*Methods*).

Fig. 2 collects the accuracies achieved using the NN to predict the relationships between students with different combinations of predictors. We first study relationships $i \to j$ with at least one directed path of length 2 from i to j; see Fig. 1 and *Methods* for more details. The results are shown in the four upper bars of Fig. 2. (*SI Appendix* for the distribution of relationships per number of directed paths of length 2, *SI Appendix*, Fig. S2.) We train the classifier using four sets of predictors: (1) triadic influence and personal information (gender, CRT, and prosociality) of the pair of nodes, (2) triadic influence, (3) personal information, and (4) only students' prosociality. Just as a clarification, in case (1), we use as input for the NN the triadic influence (a scalar) and the individual traits of both students (a 6-dimensional array) to predict the correct label of that relation (friend or enemy). *Methods* for a detailed explanation of how the value of the considered features—gender, CRT, and prosociality—are gathered and computed.

The highest balanced accuracy, 86%, is achieved using the triadic influence as input, either in combination with personal information of both students (1) or alone (2). It is remarkable that such a high accuracy for the prediction of the nature of a relationship (friend/enemy) can be obtained with just a scalar (the triadic influence) and that a 6-dimensional array containing information about both students' characteristics does not improve on that. This suggests that the triadic influence is encoding information about the prosociality of i and j as well as their gender and CRT. We postulate that I_{ij} will probably also encode (at least partially) any other relevant information for the determination of the sign of a relationship, such as political views, hobbies, sexual orientation, etc., …because our friends (and enemies) reflect on us our own idiosyncrasy ("known by the company we keep"). This suggests that I_{ij} can act as a proxy for personal compatibility when individual traits are not available.

On the other hand, using only the personal traits of both students (3) yields bAcc $= 60\%$. We studied the three attributes (gender, CRT, and prosociality) separately, and prosociality turned out to be the most predictive. Surprisingly, although gender homophily is important for the creation of links, it does not seem to be as relevant when predicting the sign of the relationship; *SI Appendix*, Figs. S6 and S7 for more details. In fact, using only students' prosociality to predict their relationship (4) already yields bAcc $= 57\%$, above the accuracy of a random guess (50%). Note that prosociality is calculated with students' answers to three simple questions (*Methods*). It is really remarkable that such a simple metric is already predictive of the nature of the social relationship between two individuals.

Finally, we study separately the relationships that do not have directed paths of length 2 connecting i to j (i.e. $(A^2)_{ij} = 0$, with A_{ij} being the adjacency matrix of the network); therefore, there is no triadic influence between i and j. These results are shown by the two *Bottom* bars of Fig. 2. Since this dataset is much smaller (2% of all relationships, i.e., 1, 211 out of a total

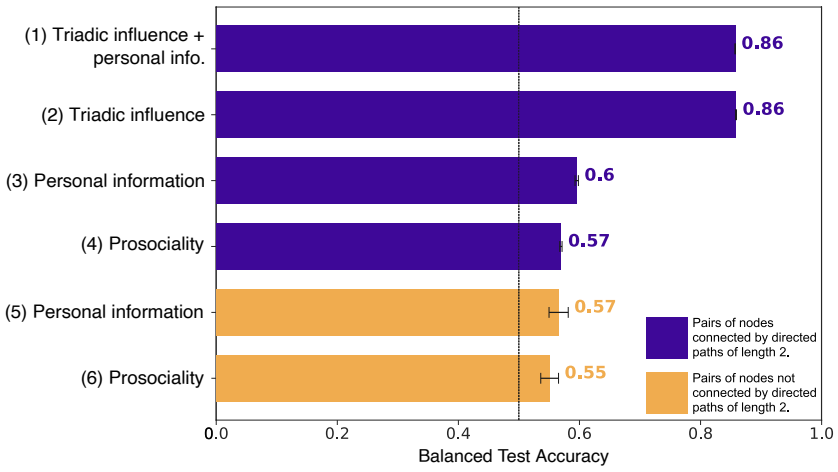

Fig. 2. Balanced test accuracy for different choices of information used to train the NN. Purple bars correspond to relationships where there is at least one directed path of length 2 from i to j ($(A^2)_{ij} > 0$, A_{ij} being the adjacency matrix of the network). We train the classifier using four sets of predictors: (1) triadic influence and personal information (gender, CRT, and prosociality), (2) triadic influence alone, (3) personal information alone, and (4) just students' prosociality. In all four cases, we trained 10 different NN with random initializations and show here the mean bAcc. Yellow bars correspond to the bAcc for relationships that have no directed paths of length 2. In this case, we use just two sets of predictors: (5) personal information and (6) students' prosociality. These cases use 10-fold cross-validation to estimate the performance of the prediction. Error bars represent the SE of the mean in all cases.

of 60, 566; *SI Appendix*, Fig. S2 for more details), we assess the performance of the classifier using 10-fold cross-validation to ensure that our results are robust. We study two sets of predictors: (5) the complete personal information of the students (gender, CRT, and prosociality) and (6) just the prosociality. The mean bAcc for the 10 realizations within 10-fold cross-validation is 57% for (5) and 55% for (6). Note that the mean bAcc seems to decrease compared to the case when $(A^2)_{ij} > 0$ (purple bars), although the significance of this difference is low given that error bars corresponding to cases (3) and (5) as well as (4) and (6) either overlap or are very close.

Interpreting the Probabilities Learned by the Neural Network. It is important to note that until now, we have chosen to assess the performance of our prediction using bAcc for the sake of simplicity. However, the NN learns more than this; in particular, it learns to predict the probability that a relationship belongs to each of the classes in the dataset. (*Methods* for a detailed explanation of how this is achieved through the minimization of the cross-entropy loss function.) The great advantage of using low-dimensional inputs is that we can interpret what the NN is learning. We can plot the probability that a sample belongs to a class (friend/enemy) as a function of the different predictors. In Fig. 3*A* we plot this probability as a function of the triadic influence. We use the 10 different NNs trained for Fig. 2 (2) and plot the average probability of being friends and enemies for a pair of students with a given triadic influence. The colored area around both curves represents the SD of the probabilities. The probability of being friends saturates to 1 when the triadic influence $I_{ij} \gg 1$ and drops to 0 if the triadic influence $I_{ij} \lesssim 0$ (the probability of being enemies is the complementary because both add up to 1). The probability curves of being friends and being enemies cross around $I_{ij} \approx 5$. Note that this is the only information used when computing the accuracy bAcc because we identify each relationship with the most probable one, as predicted by the NN. However, the probabilities learned by the neural network (which minimize the cross-entropy loss, *Methods*) contain much more information and could be used to generate ensembles of social networks or to simulate their evolution using stochastic Markov chains. It is worth mentioning that, although the probability curves change abruptly around $I_{ij} \approx 0$, this change slows down as the triadic influence increases, thus displaying an asymmetric behavior on both sides of the crossing point $I_{ij} \approx 5$. These probabilities are reminiscent of the asymmetric behavior presented by the distribution of

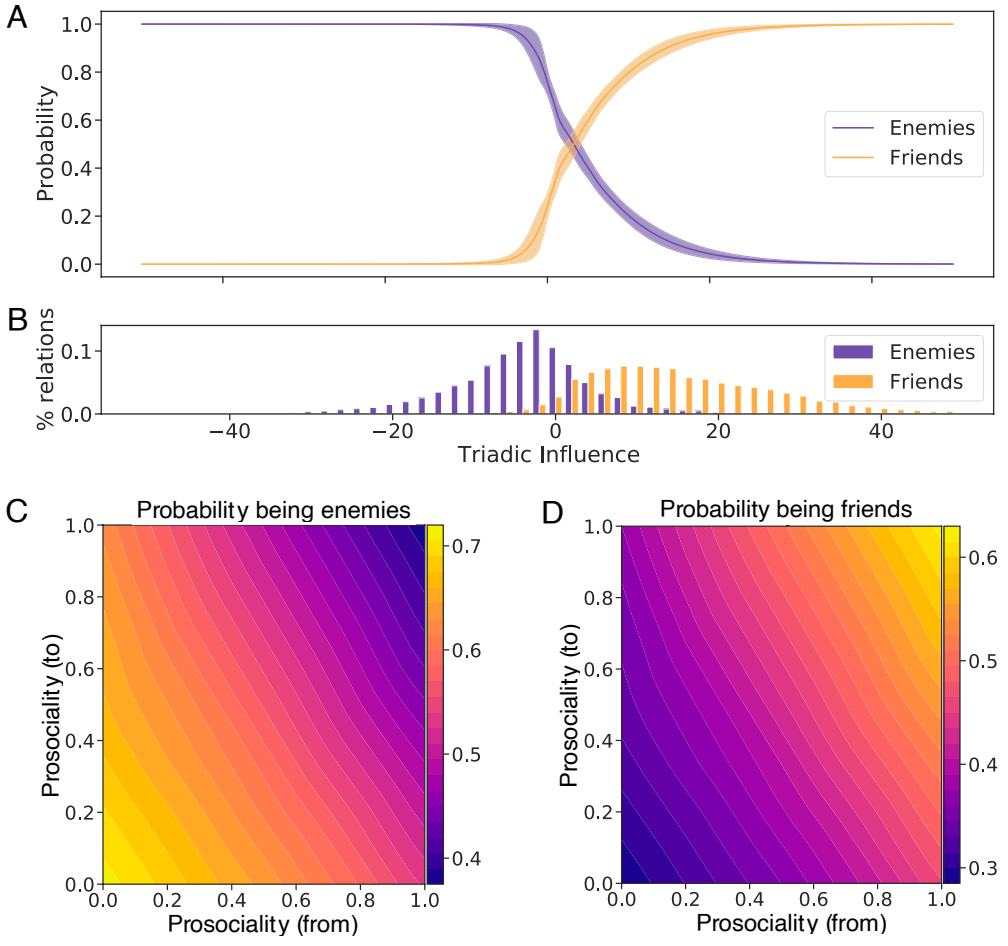

Fig. 3. Probabilities of being friends/enemies as a function of the triadic influence and prosociality. Panel (*A*) shows the probability learned by the NN as a function of the triadic influence. We performed 10 simulations that led to the accuracy shown in the (2) bar in Fig. 2. Continuous lines in panel (*A*) correspond to the mean, whereas the shaded area corresponds to one SE of the mean. Panel (*B*) shows the distribution of friends/enemies as a function of the triadic influence. Note that the probabilities in panel (*A*) display an asymmetry reminiscent of the distribution of the data. Panel (*C*) and (*D*) display the mean probabilities learned by the 10 NN used in Fig. 2 (4); they show the probability of having a friendly/enmity relationship as a function of the prosociality of both students, the nominator (from) and nominee (to). Both probabilities are normalized to 1.

friend/enemy relationships, shown in panel (*B*). A linear model can capture the transition at $I_{ij} \approx 5$, but it cannot capture the asymmetry in the probabilities (*SI Appendix*, Fig. S5).

Fig. 3 *C* and *D* display the probability of being enemies and friends, respectively, as a function of the prosociality of both students (nominator/nominee), averaged over the 10 simulations used for case (4) of Fig. 2. Similarly to the case of the triadic influence, even though bAcc is fully determined by the curve where the probability is 0.5, the profiles shown in these figures convey much more information. In particular, we can see that the probability that two students with 0 prosociality are enemies is above 70%, which is in line with what one would expect: Selfish people declare to have more enemies and are declared enemies more often than altruists (*SI Appendix*, where this can also be directly observed in the raw data, *SI Appendix*, Fig. S4). Alternatively, two highly prosocial students are friends with a probability higher than 60%. Note also that both colormaps are approximately symmetric with respect to the diagonal. This implies reciprocity: The probability that i declares j as a friend is approximately the same as the probability that j does the same with i.

Predicting with the Structural Information of the Social Network Alone. In the previous sections, we use local information—individual features and triadic influence—to predict relationships. Complementary to this, in this section, we will attempt to make the same predictions using only the structure of the network—excluding weights, link directions, and individual features—hoping to shed light on the role played by the structure of the network for the creation of different relationships. We will merge labels $\{+1, +2\}$ into a unique "friends" label and labels $\{-1, -2\}$ into a unique "enemies" label so that predictions can be binary. In order to do that, we will create a node embedding by assigning to each node a d-dimensional array of features—which will replace the array of individual features used in the previous section. A 128-dimensional embedding is created with Node2Vec (39), an algorithm that explores the neighborhood of each node using biased random walks (*Methods* for more details and *SI Appendix*, Figs. S8 and S9). The embeddings of all nodes are then used as inputs to train different models in order to predict the relationships in the network. We show here the case where we train a neural network, although we have also used a random forest (*SI Appendix*, Fig. S10) obtaining similar results.

We create the embeddings for all nodes once and keep them throughout. We then train a neural network to predict the relationship between pairs of students (friends/enemies) using both their embeddings as input. This is akin to using the individual features of both students in the previous section; only this time embeddings encode information about the environment surrounding each node. We have trained and tested the neural network using two alternative treatments: In treatment I, we have chosen at random 20% of the relationships from all high schools as the test dataset and trained the neural network using the rest of the relationships; in treatment II, we have created a test dataset with all the relationships inside one specific age level from one high school and trained the model using all the other relationships. For treatment I, we trained the neural network 390 times, every time changing the train and test datasets as well as the initialization of the neural network (the embeddings do not change). For treatment II, there are 39 different age levels within the 13 high schools that we study, and we trained 10 different neural networks for each age level—390 simulations in total.

The results corresponding to treatments I and II are summarized in Fig. 4. In this figure, we show the accuracy as a histogram after carrying out treatments I and II for the 390 simulations—purple and orange bars, respectively. For treatment I, where we train and test on random relationships, the average accuracy is ~75%, and the accuracy is always above 60% (purple bars). However, when we test on a complete age level that was excluded from the training dataset the performance degrades (treatment

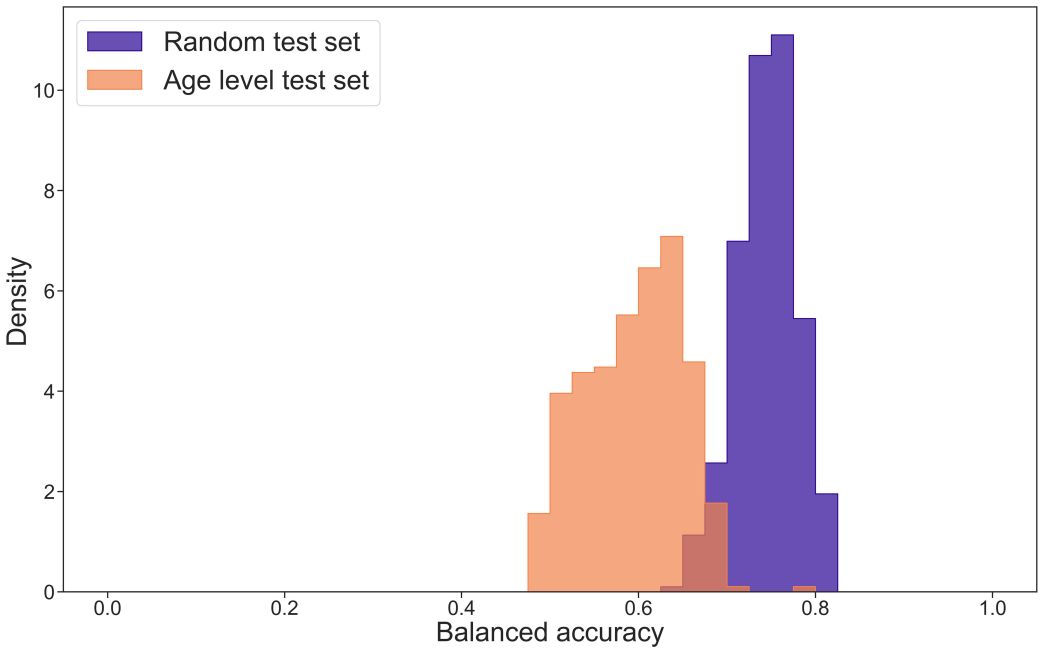

Fig. 4. Distribution of balanced accuracy for the 13 high schools. Each histogram is composed of a sample of $N = 390$ points, which are different simulations for the same treatment. The histograms are normalized so that the area under the curve is 1. The purple (dark) histogram represents treatment I where we use a random pick of edges as the test set. The orange (light) histogram represents treatment II, where we pick a specific age level from a high school as the test set. The same figure for a random forest model is included in (*SI Appendix*, Fig. S10).

II), the mean accuracy is now ∼60% (orange bars), and there are some instances where the model is not doing better than a dummy model (bAcc ∼ 50%).

The fact that the model has predictive power using only structural information shows that there is a structural difference between the environments of friendly and adversarial relationships. Besides, since the predictive power of the model decreases when testing on an isolated age level, this suggests that the structure of most age levels contains specific information that is not present in the rest of the data. We have used two dimensionality-reduction techniques to plot the embeddings corresponding to the relationships (*SI Appendix*, S11–S14). We observe that the relationships form clusters corresponding to the different age levels contained in each school. This proves that the relationships belonging to different age levels occupy different regions of input space. Therefore, when we validate taking relationships at random, we are testing the model in regions of input space that have been used during training (interpolation), whereas when we test on relationships in a complete age level, we are testing outside the regions explored during training (extrapolation), explaining the decrease in accuracy observed in Fig. 4.

Discussion

In this paper, we have applied techniques for link prediction to gain insight into the mechanisms behind the formation and evolution of social networks. This has been possible due to the large amount of data that we have collected, comprising individual features of more than 3, 000 students as well as their corresponding network of personal relationships—over 60, 000 connections. The picture of the network dynamics that emerges from our work is as follows. Some initial relationships appear between pairs of students, promoted by their prosocial stance. As a matter of fact, we have shown that the prosociality of both students by itself is capable of predicting isolated relationships significantly better than a purely random guess. This is actually a very strong claim because many of those initial relationships are now hidden among many other relationships that emerged afterward, and the isolated ones that we can find now are probably very sensitive to noise or trolling (e.g., students that label randomly other peers as friends/enemies). We hypothesize that isolated relationships continue to emerge until directed paths of length 2 dominate the dynamics of network formation. As discussed in previous sections, paths of length 2 are equivalent to intermediate students who can get two of their contacts in touch with each other. This mediation, quantified by the triadic influence, is an extremely good predictor of relationships, with accuracies as high as 86%. Interestingly, when we focus on relationships that are not isolated (there are directed paths of length 2 connecting both students), prosociality is still a good predictor of them. This suggests that some of these relationships might have originated as isolated relationships and that prosociality is still important even when the relationship is not isolated. Complementary to this, we have observed that the accuracy achieved by the triadic influence does not improve if we also provide personal information about the students. This implies that the triadic influence somehow subsumes the information on the students' characteristics, rendering it irrelevant to predict relationships. It is still an open question whether information obtained from more elaborated personality tests could improve on the predictions achieved by the triadic influence alone.

On the other hand, we have used state-of-the-art algorithms to create an embedding for each student that contains information about their surrounding, considering only the undirected and unweighted networks. We have shown that this structural information can be used to predict the type of relationship between two students. The embedding of each node is created using a random walk exploration of its surrounding, the depth of which is a parameter that we can vary (*Methods*). Depending on the typical length of the exploring random walks, this method can gather different structural information. The maximum length of the random walks used in this study is $L = 4$ (*SI Appendix*, Fig. S6). Therefore, the Node2Vec algorithm is exploring the local structure of each student. This aligns with the results achieved using the triadic influence, suggesting that the closest contacts in the network—the local environment—are the ones that influence the creation/transformation of relationships the most. Although predictions using the triadic influence achieve higher accuracies, it is remarkable that this method can predict the sign of a relationship using only structural information (without using the weights or directions of the edges).

Interestingly, ref. 40 suggests that individuals with similar genotypes may not be actively selected into friendships. Instead, they may be placed into these contexts by institutional mechanisms outside their control. Our conclusions could be interpreted similarly; the triadic influence may act as a social force that encourages students that are compatible (incompatible) to have positive (negative) relationships, akin to the popular knowledge "to be judged by the company you keep." In this case, prosociality would be still a good predictor of the relationship even though it was the social context—the triadic influence in our case—which promoted the relationship. This raises an important point that we want to stress: Predictability does not imply causality. Another situation that highlights the difficulty of disentangling cause and effect is that at the time we collected the data, many relationships that nucleated in isolation due to prosociality alone were now surrounded by multiple directed paths of length 2, and we have shown that the triadic influence is a very good predictor of the label of these relationships, even if their existence predated the paths entering the computation of the triadic influence. Therefore, while our results suggest a nucleation mechanism based on individual traits followed by growth and evolution of the network dominated by the triadic influence, they do not prove that this is indeed the case. In order to assess to what extent this idea describes what is actually happening in real networks, a possibility would be to use the probabilities that we have learned through our link prediction techniques to simulate growing/evolving networks and then compare these simulations with real data. In particular, it will be extremely interesting to collect data for the same network at different times to test the plausibility of different mechanisms of network evolution based on the probabilities learned here. If our proposal remains a good candidate to explain how networks form and evolve, then specific questions of interest arise, such as when the paths of length 2 begin to dominate over the primitive relationships existing in a network or how a local change in the sign of a relationship can lead to a cascade of changes with global effects on the social network.

Finally, it is worth mentioning that our results come from data from a large number of surveys but from a very specific population, namely, teenagers in secondary schools in Spain. Thus, the generality of our results should be validated by gathering similar data from other collectives and performing similar analyses.

Materials and Methods

Data Collection. Surveys were conducted in 13 Spanish high schools (mandatory education, 11 to 15 y of age). The study was approved by the Ethics Committees of Universidad Carlos III de Madrid and Universidad Loyola Andalucía, and the surveys were subsequently carried out in accordance with the approved guidelines. Consent was obtained from the schools which adopted this as a research project of their own and in turn got informed consent from the participants' parents. Students participated always voluntarily and signed informed consent prior to beginning the survey. The surveys were delivered through a computer interface and included direct questions about their relationships as well as some others aimed at identifying personal attributes. To elicit relationships, students could choose from a list containing all the other students in their own school. The number of classes participating in the study in each school depended on the availability of time and the decisions of the school direction. The data corresponding to one of the schools, also included in this work, were presented in full detail in ref. 41. For each student, we collected the following:

- **General data:** School ID, age level, class, and a student ID assigned by the software for the purpose of this study.
- **List of relationships:** All the relationships declared by the student (very good, good, bad, and very bad) were collected with the student IDs of the nominees and the corresponding labels (+2, +1, −1, −2).
- **Individual traits:**
  - Gender, which included 1,789 males, 1,720 females, and 4 nonbinary people.
  - Cognitive reflection test (CRT), computed using the answer to 3 questions about logic (42, 43), and yielding values 0, 1, 2, and 3.
  - Prosociality, evaluated through the answer to the three following questions about sharing (q_i ranks the level of selfishness of each answer):
    - What do you prefer? A) 10€ for you and 10€ for your partner ($q_1 = 0$) B) 10€ for you and 0€ for your partner ($q_1 = 1$).
    - What do you prefer? A) 10€ for you and 10€ for your partner. ($q_2 = 1$) B) 10€ for you and 20€ for your partner ($q_2 = 0$).
    - What do you prefer? A) 10€ for you and 10€ for your partner ($q_3 = 0$) B) 20€ for you and 0€ for your partner ($q_3 = 1$).

    The selfishness score is $s = q_1 + q_2 + q_3$, and the prosociality index is obtained as $p = 1 - (s/3)$. This task is based on (44) (see ref. 45 for details).

Predicting Relationships Using Local Information. Our social networks are directed graphs representing the relationships between all the students within each of the high schools of our study. We kept only the students who answered all the tests about their individual features (described above), a total of 3,395 students and 60566 relationships. Relationships are gathered in the weighted adjacency matrix W, with elements $w_{ij} \in \{-2, -1, 0, 1, 2\}$ corresponding to the value of the relationship that student i declares to have with student j ($w_{ij} = 0$ if there is no declared relationship). Note that $w_{ii} = 0$ and that W is not symmetric (relations are not necessarily reciprocal). Additionally, the individual traits described above (self-declared gender, CRT, and prosociality) are stored in the nodes n_i of the graph. A key quantity used in this work is the triadic influence $I_{ij} \equiv (W^2)_{ij} = \sum_k w_{ik} w_{kj}$. It quantifies the aggregated contribution of the directed paths of length 2 that go from i to j. Note that triadic influence considers only directed paths from i to j and that $I_{ij} \neq I_{ji}$ in general.

In order to use a neural network to predict the declared relationships between students, we would like to avoid having highly unbalanced classes, and therefore, we define a task with only two classes: friends (we consider here only +2 relationships) or enemies (we merge here relationships −2 and −1). We have also considered a more unbalanced case, with the friend class corresponding to relationships with labels +1 and +2, and the results were qualitatively analogous. In any case, when we compute the triadic influence I_{ij}, we keep all the labels in the network $\{-2, -1, 1, 2\}$ (Fig. 1 for an example). In this section, we use a deep neural network with one hidden layer, ReLu activation (see e.g. ref. 46), and 100 hidden units. The input dimension depends on the data we want to use to predict the relationship. Our neural network is a nonlinear function of the inputs and the internal parameters (numbers that change their value during training), which outputs a vector of dimension two. Let us call these outputs $f(\mathcal{I}, \mathcal{W})_i$, where $\mathcal{I}$ stands for the inputs corresponding to one specific relationship (triadic influence, gender of both students …), $\mathcal{W}$ are the internal parameters of the network, and $i = 0, 1$ indicates one of the two classes in our dataset (friends/enemies). Then, these outputs are put into a SoftMax function (see e.g. ref. 46) such that

$$q(\mathcal{I}, \mathcal{W})_i \equiv \frac{e^{f(\mathcal{I},\mathcal{W})_i}}{e^{f(\mathcal{I},\mathcal{W})_0} + e^{f(\mathcal{I},\mathcal{W})_1}},$$

where $q(\mathcal{I}, \mathcal{W})_i$ can be interpreted as the probability that a specific sample, characterized by inputs $\mathcal{I}$, belongs to class $i = 0, 1$. Training the neural network amounts to minimizing a loss function such that $q(\mathcal{I}, \mathcal{W})_i$ resembles the actual probability distribution $p(\mathcal{I})_i = \delta_{i,\ell(\mathcal{I})}$ for each sample–$\ell(\mathcal{I})$ being the label of that input data and $\delta_{i,j} = 1$ if $i = j$ and 0 otherwise. We use the cross-entropy loss function

$$\mathcal{L} = -\sum_{k,i} p(\mathcal{I}_k)_i \log(q(\mathcal{I}_k, \mathcal{W})_i) = -\sum_k \log(q(\mathcal{I}_k, \mathcal{W})_{\ell(\mathcal{I}_k)}),$$

where the index k runs over all samples in the dataset. Note that if $q(\mathcal{I}_k, \mathcal{W})_{\ell(l_k)} = 1$ for all k, the network would predict with 100% certainty the correct label for all samples. In this situation, $\mathcal{L} = 0$, indicating that for the set of parameters $\mathcal{W}$, the function $\mathcal{L}$ reaches an absolute minimum. Hence, training the neural network amounts to minimizing $\mathcal{L}$ with respect to the parameters $\mathcal{W}$. We have used stochastic gradient descent with an initial learning rate of 0.1 and a decaying factor of 0.99. We use a minibatch of size 20, and unless otherwise stated, we minimize for 200 steps and compute the accuracy in the final step. We observe that 200 minimization steps are enough to find a minimum of the loss function, which does not decrease further by using more steps or larger minibatches. Since we do not use all the data during training, we simply oversample the class with the smallest number of samples so that each minibatch has the same number of samples from each class. In the case of the prediction of isolated relationships (two *Bottom* bars of Fig. 2), the dataset is greatly reduced. To ensure that our results are robust, we use a 10-fold cross-validation approach and report the mean value and an error bar representing the SD from the mean. In this case, we train for 1,000 minimization steps using a dynamical loss function with oscillations of amplitude 10 and a period of 5 minimization steps. A dynamical loss function weights the contribution of each class to the loss function with proportionality factors that oscillate during minimization. This process changes the topography of the loss function landscape (47) and helps the model find deeper and wider minima of the loss function. (See ref. 48 for further details.)

Predicting Relationships Using Global Information. The steps followed in the process of creating the embeddings and predicting the class of a relationship are as follows:

- Passing the graph as an object to Node2Vec (39) yields a 128-dimensional vector for each node (an embedding). Node2Vec is defined by the two hyperparameters (p, q), which describe the space explored by the random walks. We use $(p = 1, q = 4)$ after doing a hyperparameter optimization. The characterization of the typical random walk in this process can be found in *SI Appendix*, Figs. S8 and S9.
- We merge the embeddings of each pair of nodes that are connected in the graph to create the embedding of each edge (relationship), leading to other vectors of 128 components, **e**.
- The structural representation for each edge, **e**, is the input that we use to predict the label, friends/enemies of the relationships in the training dataset. We oversample the training data (test data are left untouched) using the

SMOTE technique (49). This method produces new samples by interpolating close existing points in the 128-dimensional space.
- We apply two different machine learning procedures: a random forest and an artificial neural network.

The artificial neural network was implemented in the standard library TensorFlow (50) with one input layer of 128 neurons and 3 hidden layers–the sizes of the network layers are 128, 64, 32, and 8–and we use the ReLu activation function. The final output included a sigmoid function. To select the size of the input layer–the embedding dimension–its size was increased until the accuracy reached a plateau. The number of hidden layers has been chosen in a similar way obtaining the best results in a cross-validation procedure. The number of neurons in each hidden layer was changed sequentially to optimize the final accuracy. We also used a random forest model following previous designs in the literature (51) which provided similar results (*SI Appendix*, Fig. S10).

Data, Materials, and Software Availability. The data and codes to reproduce the results contained in this manuscript are available at https://github.com/miguel-rg/triadic-influence (52) and https://zenodo.org/record/7647000#.Y-5eDtLMJH4 (53).

ACKNOWLEDGMENTS. This work has been partly supported by grant PGC2018-098186-B-I00 (BASIC) funded by MCIN/AEI/10.13039/501100011033 and by "ERDF A way of making Europe". M.R.-G. acknowledges support from the Spanish Ministry of Science and Innovation and NextGenerationEU through the Ramón y Cajal program (RYC2021-032055-I) and from the CONEX-Plus program funded by Universidad Carlos III de Madrid and the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No. 801538. P.B.-G. acknowledges support from MCIN (PID2021-126892NB-I00), Agencia Andaluza de Cooperación Internacional para el Desarrollo (AACID-0I008/2020), Universidad de Granada (B.SEJ.280.UGR20), and Junta de Andalucía (PY18-FR-007).

Author affiliations: [a]Departamento de Estructura de la Materia, Física Térmica y Electrónica, Universidad Complutense Madrid, Madrid 28040, Spain; [b]Grupo Interdisciplinar de Sistemas Complejos (GISC), Madrid 28911, Spain; [c]Departamento de Matemáticas, Universidad Carlos III de Madrid, Leganés 28911, Spain; [d]Grupo de Investigación Ingeniería de Organización y Logística (IOL), Departamento Ingeniería de Organización, Administración de empresas y Estadística, Escuela Técnica Superior de Ingenieros Industriales, Universidad Politécnica de Madrid, Madrid 28006, Spain; [e]LoyolaBehLAB, Department of Economics and Fundación ETEA, Universidad Loyola Andalucía, Córdoba 14004, Spain; and [f]Instituto de Biocomputación y Física de Sistemas Complejos (BIFI), Universidad de Zaragoza, Zaragoza 50018, Spain

1. M. Jackson, *Social and Economic Networks* (Princeton University Press, Princeton, 2010).
2. D. Easley, J. Kleinberg, *Networks, Crowds, and Markets: Reasoning About a Highly Connected World* (Cambridge University Press, Cambridge, 2010).
3. S. Wasserman, K. Faust, *Social Network Analysis: Methods and Applications* (Cambridge University Press, Cambridge, 1994).
4. M. E. J. Newman, *Networks: An Introduction* (Oxford University Press, Oxford, 2010).
5. R. Dunbar, Structure and function in human and primate social networks: Implications for diffusion, network stability and health. *Proc. R. Soc. A* **476**, 20200446 (2020).
6. J. L. Moreno, *Who Shall Survive? Foundations of Sociometry, Group Psychotherapy, and Sociodram*. (Beacon House, 1934).
7. M. S. Granovetter, The strength of weak ties. *Am. J. Soc.* **78**, 1360–1380 (1973).
8. J. G. Oliveira, A. L. Barabasi, Darwin and Einstein correspondence patterns. *Nature* **437**, 1251 (2005).
9. J. P. Onnela *et al*., Structure and tie strengths in mobile communication networks. *Proc. Natl. Acad. Sci. U.S.A.* **104**, 7332–7336 (2007).
10. J. Ureña Carrion, J. Saramäki, M. Kivelä, Estimating tie strength in social networks using temporal communication data. *EPJ Data Sci.* **9**, 37 (2020).
11. D. Brockmann, L. Hufnagel, T. Geisel, The scaling laws of human travel. *Nature* **439**, 462–465 (2006).
12. C. Cattuto *et al*., Dynamics of person-to-person interactions from distributed RFID sensor networks. *PLoS One* **5**, e11596 (2010).
13. M. Leecaster *et al*., Estimates of social contact in a middle school based on self-report and wireless sensor data. *PLoS One* **11**, e0153690 (2016).
14. V. Gelardi, J. Godard, D. Paleressompoulle, N. Claidière, A. Barrat, Measuring social networks in primates: Wearable sensors versus direct observations. *Proc. R. Soc. A* **476**, 20190737 (2020).
15. P. Holme, J. Saramäki, Temporal networks. *Phys. Rep.* **519**, 97–125 (2012).
16. D. Watts, S. Strogatz, Collective dynamics of 'small-world' networks. *Nature* **393**, 440–442 (1998).
17. A. L. Barabási, R. Albert, Emergence of scaling in random networks. *Science* **286**, 509–512 (1999).
18. G. Bianconi, A. L. Barabási, R. Albert, Competition and multiscaling in evolving networks. *Europhys. Lett.* **54**, 436–442 (2001).
19. M. Starnini, A. Baronchelli, R. Pastor-Satorras, Modeling human dynamics of face-to-face interaction networks. *Phys. Rev. Lett.* **110**, 168701 (2013).
20. E. M. Jin, M. Girvan, M. E. J. Newman, Structure of growing social networks. *Phys. Rev. E* **64**, 046132 (2001).
21. S. Hanneke, W. Fu, E. P. Xing, Discrete temporal models of social networks. *Electron. J. Stat.* **4**, 585–605 (2010).
22. T. A. B. Snijders, G. G. Van de Bunt, C. E. G. Steglich, Introduction to stochastic actor-based models for network dynamics. *Soc. Netw.* **32**, 44–60 (2010).
23. J. Peter, P. M. Valkenburg, A. P. Schouten, Developing a model of adolescent friendship formation on the internet. *Cyberpsychol. Behav.* **8**, 423–430 (2005).
24. D. Liben-Nowell, J. Kleinberg, The link-prediction problem for social networks. *J. Am. Soc. Inf. Sci. Tec.* **58**, 1019–1031 (2007).
25. L. Lü *et al*., Recommender systems. *Phys. Rep.* **519**, 1–49 (2012).
26. E. M. Airoldi, D. M. Blei, S. E. Fienberg, E. P. Xing, T. Jaakkola, Mixed membership stochastic block models for relational data with application to protein-protein interactions. *J. Mach. Learn. Res.* **9**, 1823–1856 (2008).
27. M. E. J. Newman, Clustering and preferential attachment in growing networks. *Phys. Rev. E* **64**, 025102 (2001).
28. G. Berlusconi, F. Calderoni, N. Parolini, M. Verani, C. Piccardi, Link prediction in criminal networks: A tool for criminal intelligence analysis. *PLoS One* **11**, e0154244 (2016).
29. R. Guimerà, M. Sales-Pardo, Missing and spurious interactions and the reconstruction of complex networks. *Proc. Natl. Acad. Sci. U.S.A.* **106**, 22073–22078 (2009).
30. H. H. Song, T. W. Cho, V. Dave, Y. Zhang, L. Qiu, "Scalable proximity estimation and link prediction in online social networks" in *IMC 2009: Proceedings of the 9th ACM SIGCOMM conference on Internet measurement* (ACM, New York, NY, USA, 2009), pp. 322–335.
31. Z. Hao, Link prediction in online social networks based on the unsupervised marginalized denoising model. *IEEE Access* **7**, 54133–54143 (2019).
32. A. Kumar, S. S. Singh, K. Singh, B. Biswas, Link prediction techniques, applications, and performance: A survey. *Physica A* **553**, 124289 (2020).
33. I. Tamarit, "Ego-centred models of social networks: the social atom," PhD thesis, Universidad Carlos III de Madrid, Madrid, Spain (2019).
34. V. L. Buijs, G. Stulp, Friends, family, and family friends: Predicting friendships of dutch women. *Soc. Netw.* **70**, 25–35 (2022).
35. F. Heider, Attitudes and cognitive organization. *J. Psychol.* **21**, 107–112 (1946).
36. D. Cartwright, F. Harary, Structural balance: A generalization of Heider's theory. *Psychol. Rev.* **63**, 277 (1956).
37. F. Harary, On the measurement of structural balance. *Behav. Sci.* **4**, 316–323 (1959).
38. K. H. Brodersen, C. S. Ong, K. E. Stephan, J. M. Buhmann, "The balanced accuracy and its posterior distribution" in *2010 20th International Conference on Pattern Recognition* (IEEE, 2010), pp. 3121–3124.
39. A. Grover, J. Leskovec, "node2vec: Scalable feature learning for networks" in *KDD 2016: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (ACM, 2016), pp. 855–864.
40. J. D. Boardman, B. W. Domingue, J. M. Fletcher, How social and genetic factors predict friendship networks. *Proc. Natl. Acad. Sci. U.S.A.* **109**, 17377–17381 (2012).
41. D. Escribano, V. D. Martelli, F. J. Lapuente, J. A. Cuesta, A. Sánchez, Evolution of social relationships between first-year students at middle school: From cliques to circles. *Sci. Rep.* **11**, 11694 (2021).
42. P. Brañas-Garza, P. Kujal, B. Lenkei, Cognitive reflection test: Whom, how, when. *J. Behav. Exp. Econ.* **82**, 101455 (2019).
43. P. Brañas Garza, L. Ductor, J. Kovárík, The role of unobservable characteristics in friendship network formation (2022).
44. H. Fehr, B. Bernhard, B. Rockenbach, Egalitarianism in young children. *Nature* **454**, 1079–1083 (2008).
45. A. Alfonso-Costillo *et al*., The adventure of running experiments with teenagers. PsyArXiv [Preprint] (2022). https://psyarxiv.com/zphx9 (Accessed 7 March 2023).
46. I. Goodfellow, Y. Bengio, A. Courville, *Deep Learning* (MIT Press, Cambridge, MA, USA, 2016). http://www.deeplearningbook.org.
47. M. Ruiz-García, A. J. Liu, E. Katifori, Tuning and jamming reduced to their minima. *Phys. Rev. E* **100**, 052608 (2019).
48. M. Ruiz-García, G. Zhang, S. S. Schoenholz, A. J. Liu, "Tilting the playing field: Dynamical loss functions for machine learning" in *Proceedings of the 38th International Conference on Machine Learning* (2021), vol. 139, pp. 9157–9167 (2021).
49. N. V. Chawla, K. W. Bowyer, L. O. Hall, W. P. Kegelmeyer, Smote: Synthetic minority over-sampling technique. *J. Artif. Intell. Res.* **16**, 321–357 (2002).
50. M. Abadi *et al*., "Tensorflow: Large-scale machine learning on heterogeneous systems" in *OSDI 2016: Proceedings of the 12th USENIX Conference on Operating Systems Design and Implementation* (ACM, 2015), pp. 265–283.
51. A. Longa, G. Pellegrini, G. Santin, Pytorch geometric tutorial (2021). https://antoniolonga.github.io/Pytorch_geometric_tutorials/. Accessed 7 March 2023.
52. M. Ruiz-García *et al*., Triadic influence as a proxy for compatibility in social relationships. Github. https://github.com/miguel-rg/triadic-influence. Deposited 21 February 2023.
53. M. Ruiz-García *et al*., Triadic influence as a proxy for compatibility in social relationships. Zenodo. https://zenodo.org/record/764700#.ZAEny-zMJox. Deposited 16 February 2023.

Supporting Information for

Triadic influence as a proxy for compatibility in social relationships

Miguel Ruiz-García*, Juan Ozaita*, María Pereda, Antonio Alfonso, Pablo Brañas-Garza, José A. Cuesta and Angel Sánchez

To whom correspondence should be addressed. E-mail: miguel.ruiz.garcia@uc3m.es

This PDF file includes:

Supporting Information Text

1. Statistical analysis of the data

Our dataset compromises 13 schools with 3395 students and 60566 declared relationships, where we have already removed the students that did not answer all the questions about personality (around 3% of the students) and the relationships that included them. From the 13 schools considered, 3 of them are in the Region of Madrid and the rest are in Andalucía. In this Supplementary Materials we will name the schools in Andalucía as t11_1, ..., t11_10, whereas the schools in Madrid will be t1, t2 and t6.

Students can declare to have a very good (+2), good (+1), bad (−1) or very bad (−2) relationship with any other student in their school. The distribution of the four types of relationships changes among schools although some features are common, see Fig. S1. For all schools, the most numerous relationships are good (+1) relationships, accounting between 40 and 50% in most schools. On the contrary, the least abundant relationships are the very bad ones (−2). In addition to this, we see some differences between the schools in Madrid and Andalucía but we leave a detailed study of these features for a future work.

According to the main text, the most important information to predict a relation between students i and j is the triadic influence. This quantity accounts for the influence that third people have on other relationships, and uses the directed paths of length 2 that connect i to j. For this reason, it is interesting to know how many relationships have a particular number of paths of length 2 connecting the starting node (i) and the final node (j). Figure S2 encodes this information. We have plotted each school separately (color lines) and all the schools together (black line). From the plot, we can see that 2% of all relationships do not have any path of length 2 connecting the starting and final node. The violin plot helps to visualize each school separately. Although there is variation between schools, the peak of the joint distribution (black line) is around 5 or 6 paths of length 2.

Prosociality is computed for each student. Students answer 3 questions that lead to a scalar that takes values $0, 0.33, 0.66$ or 1 (see Methods). The proportion of students per prosociality value is similar between schools, see Fig. S3, and resembles the distribution of relationships (Fig. S1). The largest group of students in most schools is students with prosociality 0.66, while the smallest group corresponds to antisocial students (0). Again, there is an apparent difference between the schools in Andalucía and Madrid. The latter seem to have a larger proportion of very social students 1 compared to the schools in Andalucía. This seems to be in good agreement with the distribution of relationships where schools in Madrid showed more very good relationships (+2). We leave a detailed study of this phenomenon for future work.

Figure S4 shows the average number of friends/enemies that are nominated by (or that nominate to) students of different prosociality values. Clear general trends indicate that students of high prosociality nominate and are nominated as friends more often, whereas students with low prosociality belong to negative relationships in a larger proportion than highly prosocial students. These results are in good agreement with the probabilities learnt by the neural network in the main text, where students with low prosociality had a larger probability of being enemies whereas highly prosocial students have a larger probability of being friends.

In the case where we use a simple fully connected neural network, we aim to learn the probability that two students are friends or enemies based on descriptors that we have selected/defined *a priori*, such as the personal information of each student and a new metric proposed by us, the triadic influence. Our goal is to learn these probabilities because we expect them to be the foundations for future work that use them to simulate the evolution of social networks. However, even a linear model could use the triadic influence to predict the value at which the sign of a relationship would change, this highlights that the triadic influence is an extremely good predictor for the sign of relationships. Fig. S5 shows that a linear model can also capture the transition where the most probable relationship goes from enemy to friends. Panel (c) shows the percentage of friend and enemy relationships for each value of the triadic influence: there is a transition around 5 above which the percentage corresponding to friends is higher than enemies. Note that the linear model can capture the correct point at which the transition occurs, however, it cannot capture the asymmetry in the probability distribution that stems from the different distribution of friend/enemy relationships (one distribution is sharper whereas the other one is flatter).

Considering homophily in the formation and sign of the relationships. Homophily is present in the creation of links in our dataset. From Fig. S6, we see that female and male students form roughly twice as many relationships with other students that self-identify with the same gender than with students of the opposite gender. However, in our case we are more interested in the type (friend/enemy) of that relationship. In this case, homophily plays also a role. For example, from all the relationships that male students form with other male students, 86% are positive, whereas they only declare a 73% of the relationships with female students as positive. The case of female students is similar, although the effect of homophily is less strong. They declare a 71% of their relations with male students as positive versus a 78% for their relationships with other female students. However, although these effects exist, note that if we directly transformed these proportions to probabilities, the probability of being friends is always above 50%, regardless of the gender combination. One of the key ideas in our work is that the triadic influence acts as a proxy for social compatibility. This means that all the factors that can lead/promote a positive/negative relationship are (at least partially) contained in the triadic influence. This is true in the case of homophily, the distribution of triadic influence is very similar for all the gender combinations (male-male, female-female, male-female and female-male), see Fig. S7. In particular, there is no significant difference depending on the combination of genders, the transition occurs approximately in the same point ($I_{ij} \sim 5$) and because of that the triadic influence can be used as an universal predictor.

2. Predicting using only the structure of the network: additional information on the creation of the embedding and alternative results using a Random Forest

We include here some additional information to clarify how the Node2vec algorithm works. Node2Vec builds the embedding for each node using random walks. The configuration of hyperparameters that we have chosen ($p = 1$ and $q = 4$) leads to local exploration. To create the embedding of one node we use 420 random walks, each of them composed of 30 movement attempts. In order to get some intuition, we show in Fig.S8 the path followed by three random walks starting from different nodes, we show RWs that does not stay fixed in the initial node, although that is the most probable situation. The local exploration of the surrounding of the node by the random walks is confirmed by Fig.S9, which shows the distribution of lengths for the random walks with the hyperparameters used in our work (see Methods).

Finally, to complement the results shown in the main text, we also include here the results for a second method using the embeddings as input. We have used a Random Forest (scikit-learn (1)). A Random Forest is an algorithm that divides the data into random ensembles of predictors and data. From these randomly chosen pieces of data, the algorithm builds and trains decision trees. The final decision is then made on the most popular answer within this population of decision trees. We used a maximum depth of 7 levels for our decision trees. Fig. S10 depicts qualitatively analogous results to the ones achieved with the neural network and shown in the main text.

3. Embedding visualization

In order to improve our understanding of the node embedding, we show in this section the results corresponding to two different methods for dimensionality reduction. We use Principal Component Analysis (PCA) and t-distributed Stochastic Neighbor Embedding (TSNE). In both cases, we reduce the dimension from 128 to 2, and color each point after applying these unsupervised techniques with other properties of the relationships, such as their sign.

We show in figures S11 and S12 two-dimensional representations of the embeddings. In these figures, each dot represents a relationship (a link) in the social network. Color represents the sign of each relationship (friends/enemies). Both methods show that the embedding is not linearly separable in two dimensions for the sign of the relationship.

As it can be seen in figures S11 and, especially, in S12, relationships tend to form clusters. We looked for the correlations of this clusters with other node variables, which we recall they are not using during the creation of the embedding. We made the same plots as before, but colored by age level (fig.S13 and fig.S14). The colors in this picture refer to the different age levels present in our dataset: the sixth (and last) grade of primary school, and the four different grades of secondary school (ESO in Spanish). Note that not all the schools had all levels. A dot of a color corresponding to an age level corresponds to a relationship where both individuals correspond to that age level. A green dot, for example, would correspond to a relationship where both students are in the second grade of secondary school. The label "Intergroup" classifies all the cases where individuals in the relationship belong to different age levels. For example, a student from the last grade of primary school who has a relationship with someone from the first grade of secondary school.

The conclusion from the pictures where color represents age level is much more clear, the clusters both in PCA and TSNE are related to the age levels that exist in every high school. This clusterization around age levels provides an intuition of why the embedding performs poorly when it has to guess the sign of a relationship for a particular age level that it was not included during training. This detriment of generalization may suggest that each age level contains particular information important for the analysis.

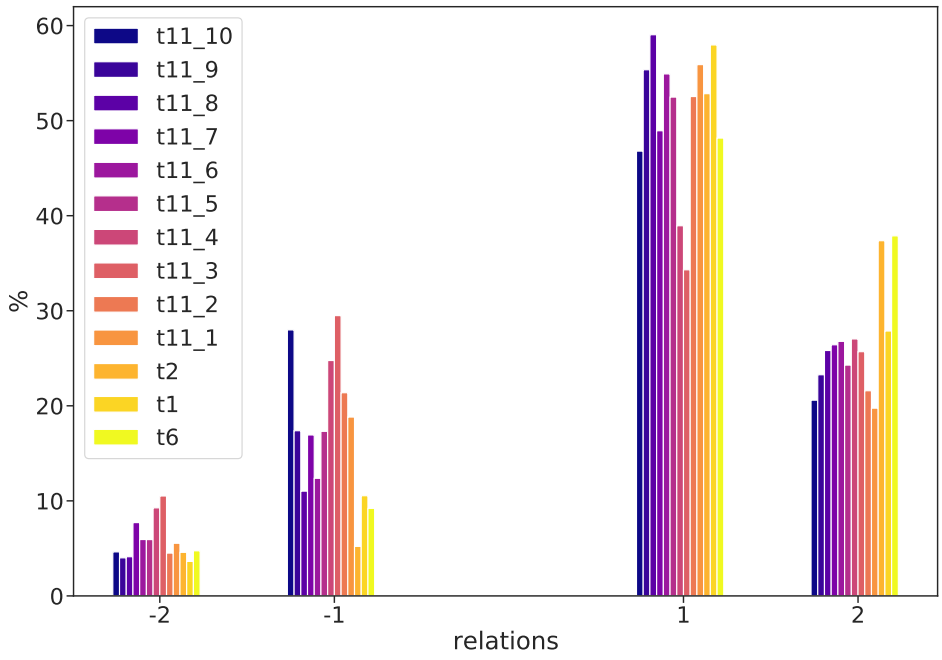

Fig. S1. Proportion of the different types of relationship at each school. Students declare relationships with values $-2, -1, 1, 2$, from nemesis to best friends. These relationships are directed, from student $i \to j$ with value W_{ij}. We have studied 13 schools with 3395 students in total, and 60566 declared relationships. Most common relationships are $+1$ followed by $+2$, -1 and -2. We observe similar percentages across schools.

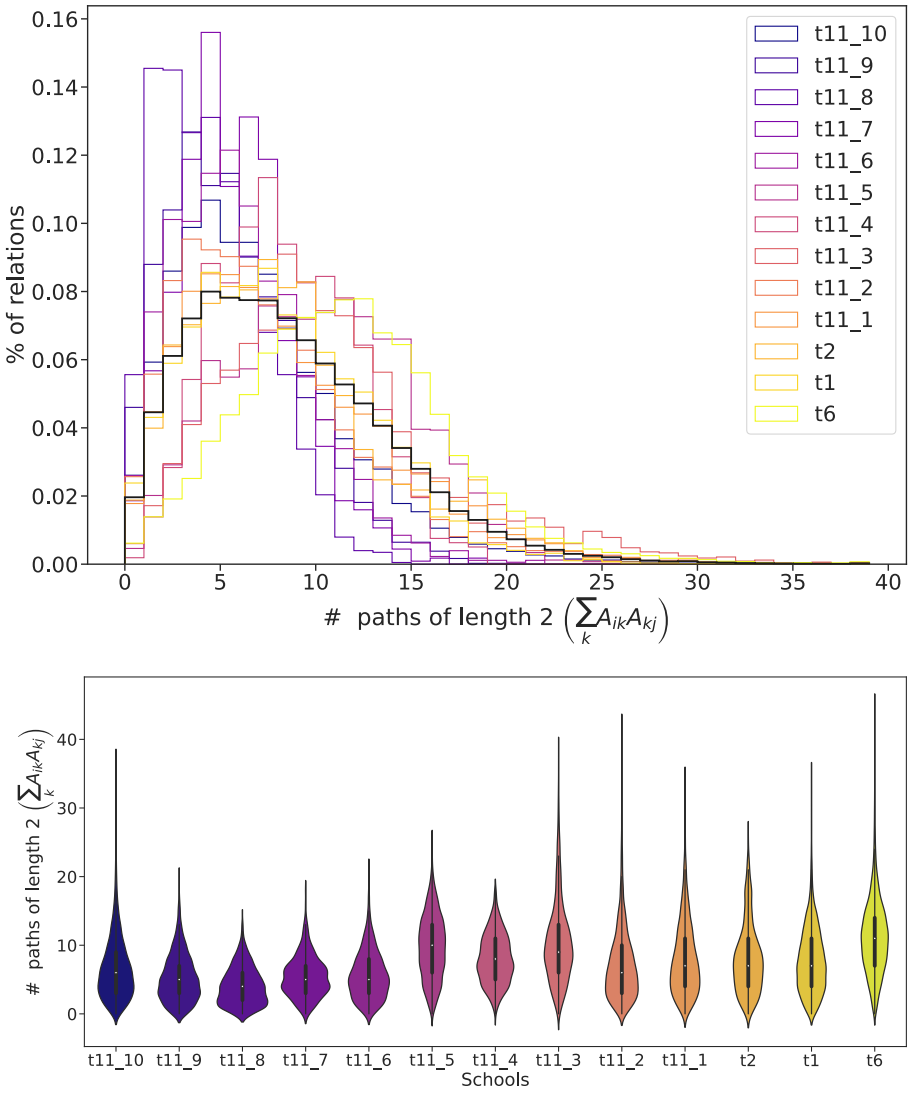

Fig. S2. Percentage of relationships according to the number of paths of length 2 connecting the same nodes. For each relationship from i to j we compute the number of paths of length 2 that go from i to j, after this we compute the number of relations with a specific number of length-2 paths. Upper panel displays a color line for each school (see legend) and a thicker black line for the distribution corresponding to the complete dataset (clustering all the schools together). Bottom panel show a violin plot for each school to ease visualization.

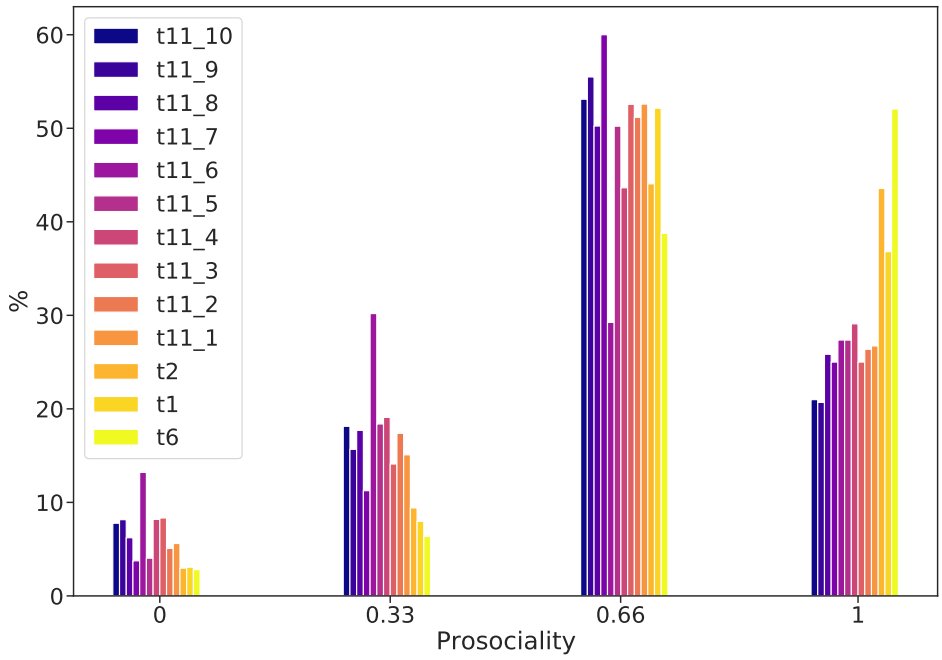

Fig. S3. Distribution of prosocial behavior across schools. Prosocial behavior takes values $0, 0.33, 0.67, 1$ and bars correspond to the percentage of students in each school that display a certain level of prosociality. Distributions are similar across schools although t1, t2 and t6 (schools in Madrid) seem to have a larger proportion of very prosocial students (1) when compared to the other schools from Andalucía.

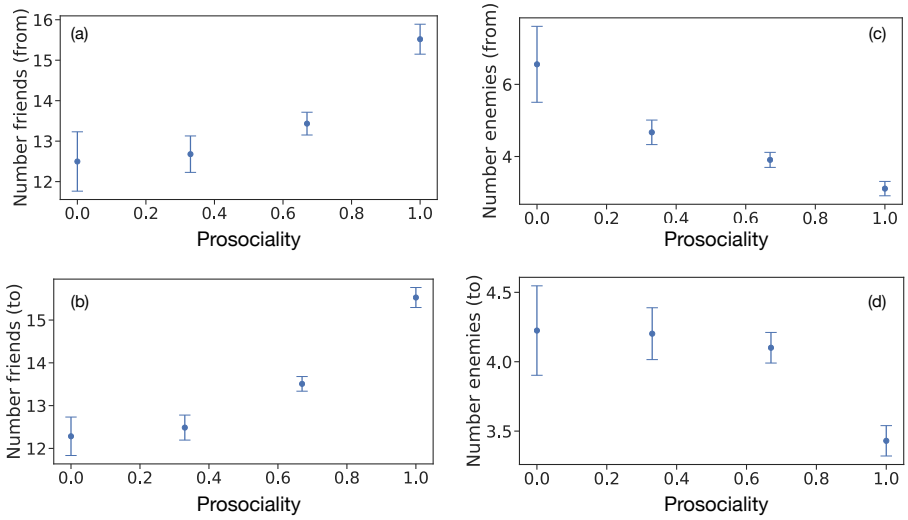

Fig. S4. Statistics for the number of friends and enemies depending on the prosociality of the students. We merge here -2 and -1 relationships into an enemy category and $+2$ and $+1$ relationships as friends. Panel (a) shows the average number of people nominated as friends by a student of a specific prosociality, whereas panel (b) shows the average number of people that nominate someone of a specific prosociality as a friend. Panels (c) and (d) are analogous to (a) and (b) but for enemies. Students with high (low) prosociality tend to nominate and be nominated as friends (enemies) more frequentely. Error bars correspond to the error of the mean.

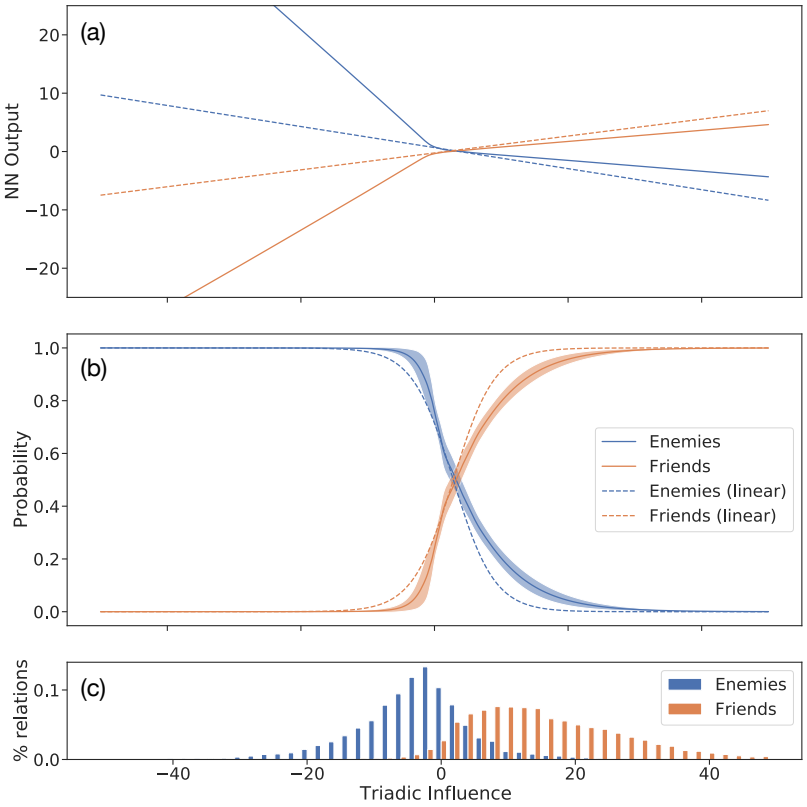

Fig. S5. Panels (a) and (b) show the output of the neural network and the probability of being friends/enemies predicted by these outputs. The continuous lines show the nonlinear model described in the manuscript whereas the dashed lines show a linear model (the same NN without the ReLu nonlinearities). Panel (c) shows the percentage of friend and enemy relationships for each value of the triadic influence.

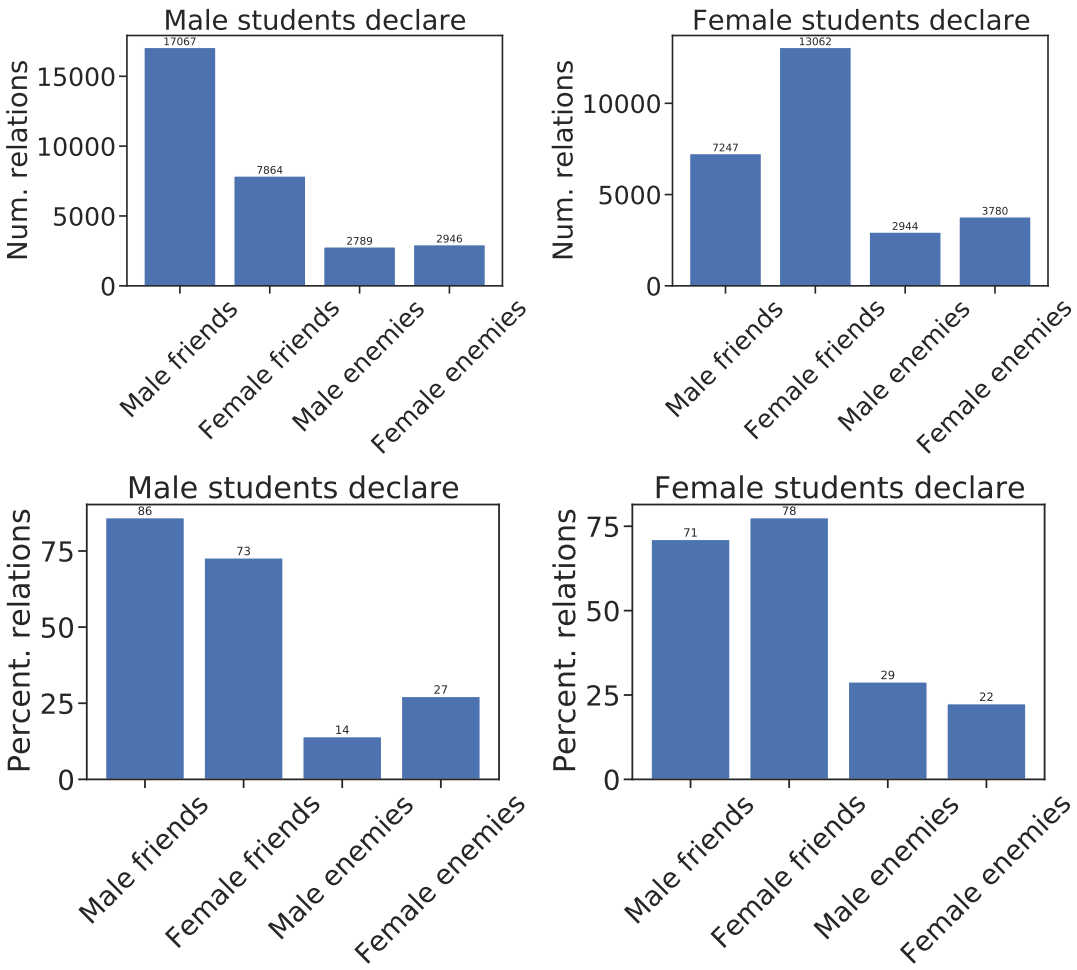

Fig. S6. Total number and percentage of relationships that connect students that self identify as male or female.

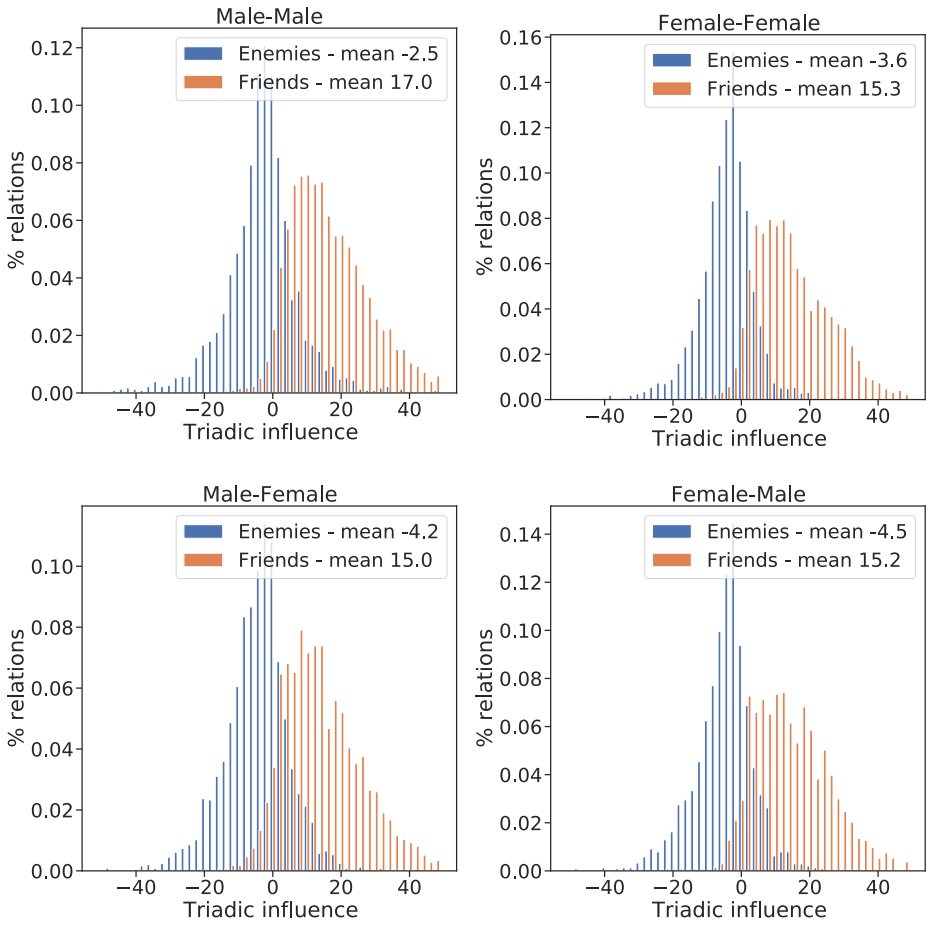

Fig. S7. Triadic influence distribution of friend/enemy relationships male-male, female-female, male-female and female-male relationships.

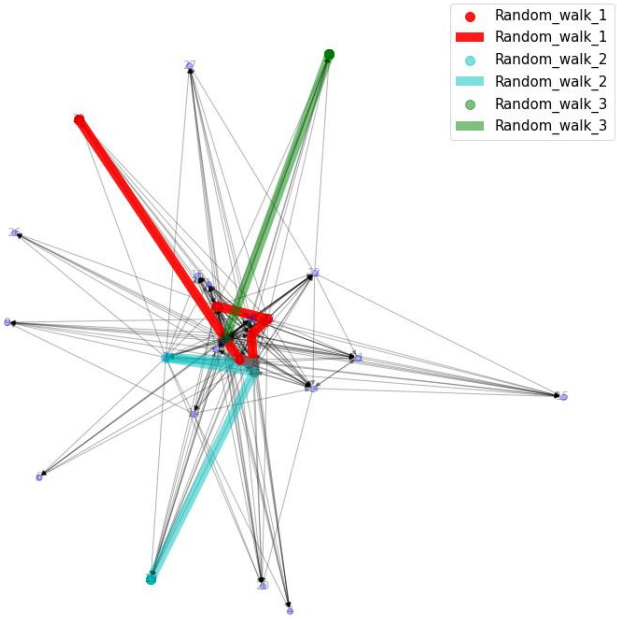

Fig. S8. Network representation of the students in one of the classes in the dataset, together with three examples of random walks used by *Node2Vec* to create the embedding of one node. The random walks are controlled by the hyperparameters (in our case $p = 1$ and $q = 4$), they measure the probability of exploring or staying in a certain node along the path. From the figure, we can see that Node2Vec is exploring the local environment of the node.

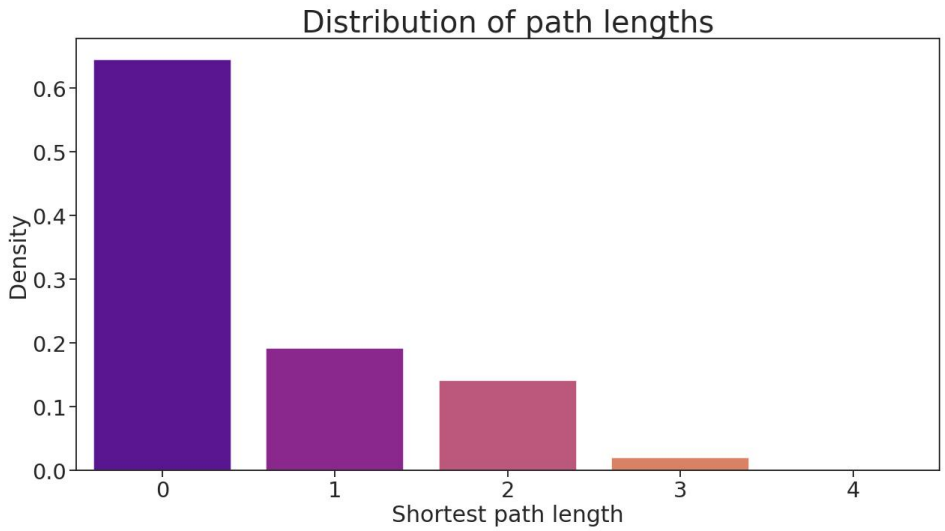

Fig. S9. Distribution of path lengths in the random walks used by Node2Vec. This length is the shortest path between the origin of the random walk and the farthest node in that walk. A walk of length 0 means staying at the same node. We can therefore understand the locality of the algorithm. The shortest path length has an average and standard deviation of 0.80 ± 1.07, which shows again that embeddings are mainly built with local structure.

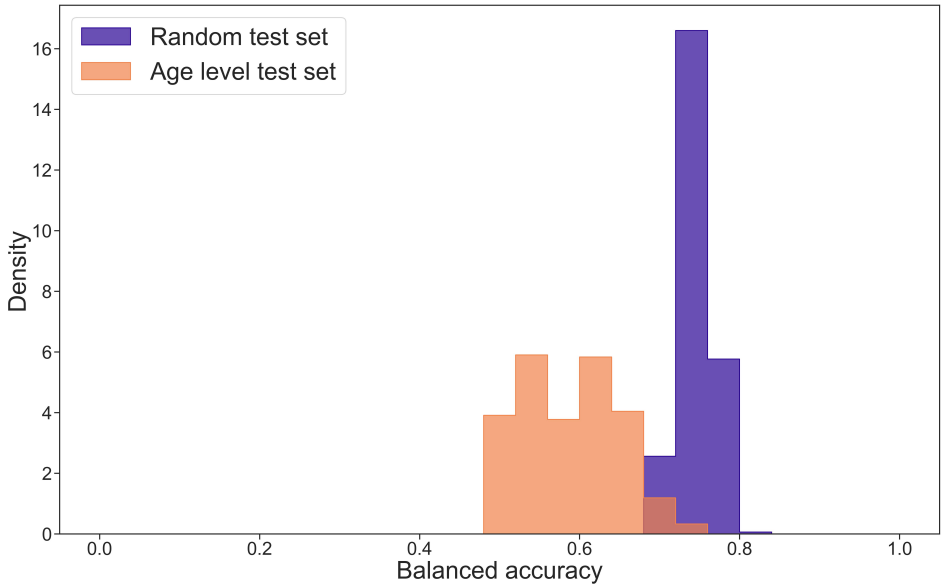

Fig. S10. Predictions achieved with a Random Forest. Distribution of balanced accuracy for the 13 high schools. Each histogram is composed of a sample of $N = 390$ points, that are different simulations for the same treatment, and then normalized such that the area of the histogram sums 1. The purple/dark histogram represents *treatment I* whereas the orange histogram represents *treatment II*, that is, using an specific age level from a high school as the test set.

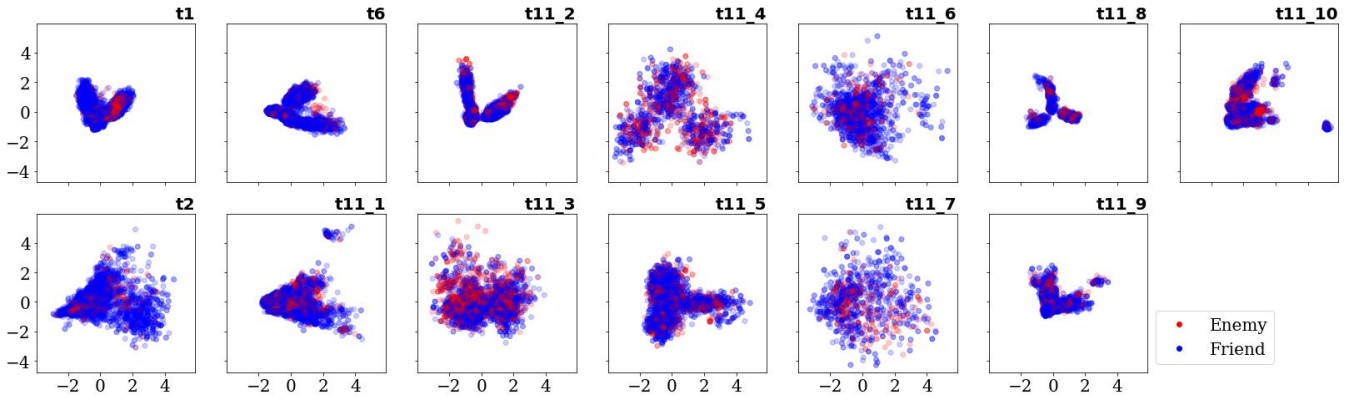

Fig. S11. Principal Component Analysis reduction to 2 dimensions for the embedding representation of each relationship, each panel represents a different high school. Color represents the type of relationship, blue represents a friendship and red an enmity. The data cannot be easily separated in this representation.

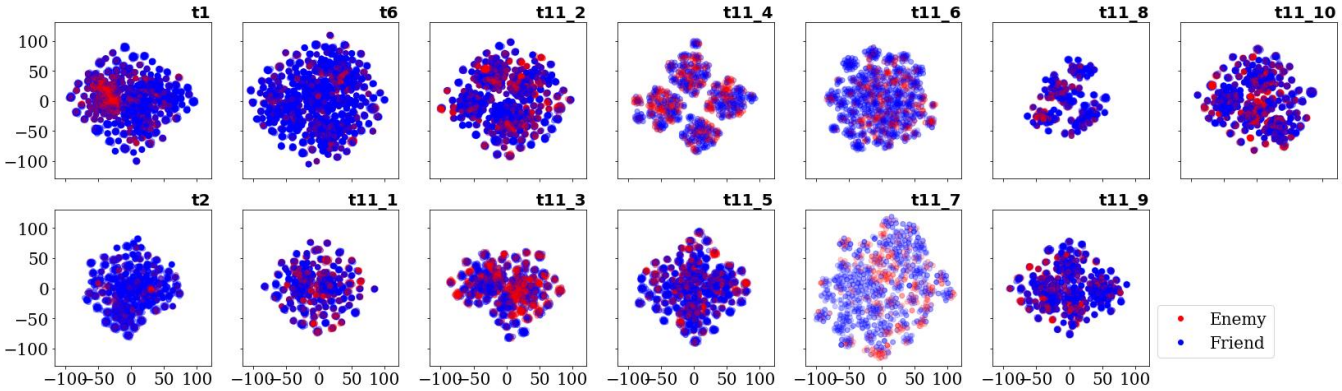

Fig. S12. t-distributed Stochastic Neighbor embedding representation in two dimensions for the relationships of the 13 high schools. As in the figure before, each panel corresponds to a high school and color represents the type of relationship. As in the case of PCA, there is no clear separation for the type of relationships.

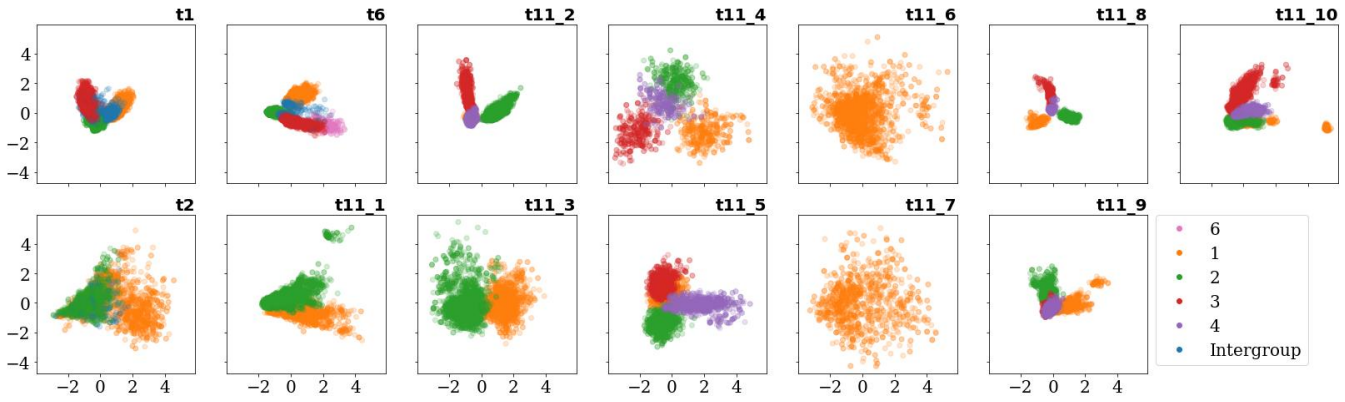

Fig. S13. Principal Component Analysis embedding representation in two dimensions for the relationships of the 13 high schools. Each panel represents a high school and the color represents the age level the students belongs to. Due to lack of data, there are high schools where only one age level is available. However, in high schools where there are more than one age level, the data clearly form clusters that map to the different age levels.

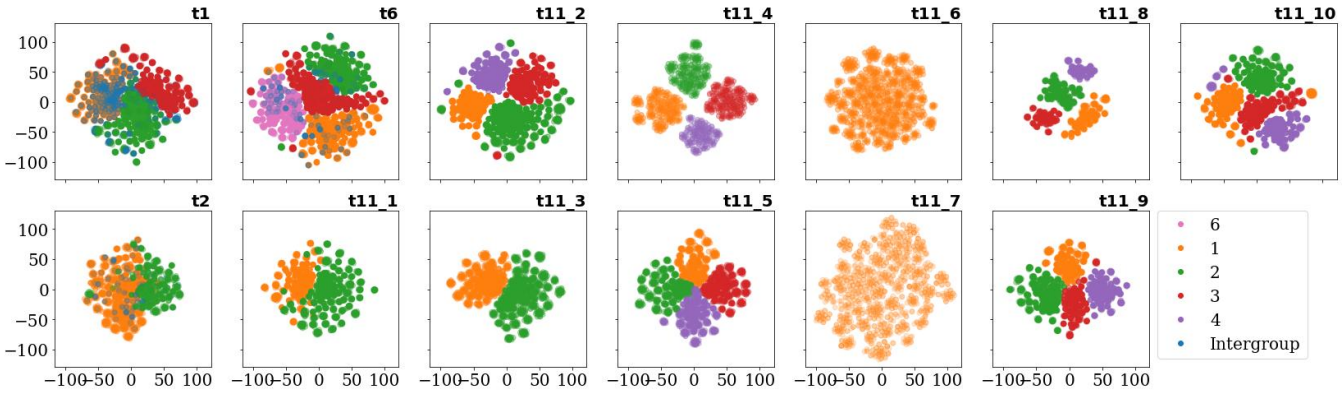

Fig. S14. t-distributed Stochastic Neighbor Embedding representation in two dimensions for the relationships within the 13 high schools, where color represents the age level the students belongs to. The analysis is similar to the one made for the previous figure, but in this case is even clearer. Relationships form clusters that correspond to the different age levels present in each high school.

References

1. F Pedregosa, et al., Scikit-learn: Machine learning in python. *J. Mach. Learn. Res.* **12**, 2825–2830 (2011).