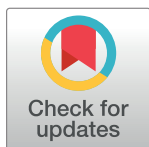


COLLECTION REVIEW

Predicting affinity ties in a surname network

Marcelo Mendoza^{1,2*}, Naim Bro¹**1** Millennium Institute of Foundational Research on Data, Santiago, Chile, **2** Department of Informatics, Universidad Técnica Federico Santa María, Santiago, Chile* marcelo.mendoza@usm.cl

Abstract

From administrative registers of last names in Santiago, Chile, we create a surname affinity network that encodes socioeconomic data. This network is a multi-relational graph with nodes representing surnames and edges representing the prevalence of interactions between surnames by socioeconomic decile. We model the prediction of links as a knowledge base completion problem, and find that sharing neighbors is highly predictive of the formation of new links. Importantly, We distinguish between grounded neighbors and neighbors in the embedding space, and find that the latter is more predictive of tie formation. The paper discusses the implications of this finding in explaining the high levels of elite endogamy in Santiago.

OPEN ACCESS

Citation: Mendoza M, Bro N (2021) Predicting affinity ties in a surname network. PLoS ONE 16(9): e0256603. <https://doi.org/10.1371/journal.pone.0256603>

Editor: Hocine Cherifi, University of Burgundy, FRANCE

Published: September 2, 2021

Copyright: © 2021 Mendoza, Bro. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: All files (code and data) are available from: Mendoza, Marcelo, Bro, Naim (2021): Predicting affinity ties in a surname network in the github repository: <https://github.com/marcelomendoza/SA19k>.

Funding: Marcelo Mendoza and Naim Bro acknowledge funding support from the Millennium Institute for Foundational Research on Data. Marcelo Mendoza was funded by the National Agency of Research and Development (ANID) grants Programa de Investigación Asociativa (PIA) AFB180002 and Fondo Nacional de Desarrollo Científico y Tecnológico (FONDECYT) 1200211.

Competing interests: The authors have declared that no competing interests exist.

Introduction

In countries where citizens inherit both parents' last names, surnames are a source to uncover the population structure of society. In previous research [1], we built a network from surname pairs that revealed strong patterns of socioeconomic segregation in Santiago, Chile. In this study, we model the formation of affinity ties between surnames, and find that the formation of new links is highly predictable.

The method predicts links as a knowledge base completion problem. We evaluate several state-of-the-art knowledge base completion models and find that TuckER, which is based on a three-way tensor decomposition, obtains the best results.

The findings bear implications for the literature that examines the mechanisms behind group segregation. Traditional explanations focus on homophily, the propensity of people to connect to others of their own kind—the “birds of a feather” mechanism [2]. Advances in network statistical modelling in the last decade have disentangled homophily from other, network-specific mechanisms that help segregate groups [3]. People not only tend to prefer connecting with their in-groups, they are also prone to close triangles. This is the notion of triadic closure or “friend-of-a-friend” (FOAF) mechanism: A and B are likelier to become friends if they have C as a common friend [4–7]. FOAF independently exacerbates the separation of groups, because the friend-of-a-friend is often an in-group.

Our experiments show that triadic closure operates not only on the empirical network. We compute the embeddings of all nodes, and find that those that share a common neighbor in the *embedding space* are very likely to connect. Moreover, proximity in the embedding space accounts for more than half of the triadic closures, and it is a stronger predictor of link

formation than proximity in the empirical network. Sharing a similar location in the embedding space can be thought of as being structurally similar.

The substantive implication is that class divisions in Chile do not occur only because people prefer connecting and marrying others of their own socioeconomic status, but because they connect to others occupying similar locations in the network of interactions.

Along with the publication of this study, we release the surname affinity network SA19k (Surname Affinity 19k entities), a knowledge base for the prediction of surname affinity ties in Santiago, Chile. This is a multi-relational graph, where nodes are surnames and links are affinity relations divided by decile in the socioeconomic distribution.

Materials and methods

Data

We use the data introduced in Bro & Mendoza [1]. This dataset is based on the Chilean electoral registry of 2012, which contains the full name and the address of all individuals eligible to vote for political authorities in Chile. We only use the data of residents in Santiago, Chile, totaling 4,652,933 individuals. We also use the Territorial Well-being Index of 2012 [8], which indexes the mean socioeconomic status of every census administrative unit down to the block level, totaling 39901 blocks.

Individuals' socioeconomic status was assigned based on the mean socioeconomic level of the blocks where they live. Socioeconomic status was transformed into a 0–100 range using the formula $z_i = \frac{x_i - \min(x)}{\max(x) - \min(x)}$, where $x = (x_1, \dots, x_n)$ and z_i is the i -th element of the normalized vector. For every person in the list, the final dataset contains the paternal and maternal surnames, an assigned socioeconomic status, and the block identification.

SA19k: A paternal-maternal surname knowledge base

We build a knowledge base that comprises paternal-maternal surname ties and income distributions based on the data described above. To create the knowledge base, we build an undirected network based on paternal-maternal surname affinities, where each node represents a surname. To create the edges, we count the number of individuals who share a pair of surnames, irrespective of the paternal-maternal or maternal-paternal order. Each pair of nodes was connected with ten relations, representing the deciles of the income distribution. The edge's weight represents the number of individuals in each decile of the income distribution. This network has 76649 surnames and 1750802 edges.

Following Mateos *et al.* [9], we remove surname pairs if their weight denoted by n_{ss} is less than a threshold $\frac{k \cdot n_{s1} \cdot n_{s2}}{N}$, where $k > 1$ is the security parameter of the method, n_{s1} is the number of occurrences of the surname s_1 , n_{s2} is the number of occurrences of the surname s_2 , and N is the number of individuals in the sample. In a uni-relational graph, $\frac{k \cdot n_{s1} \cdot n_{s2}}{N}$ is k times the expected number of co-occurrences of both surnames if the surnames are linked at random. A high value of k generates a sparse network where ties are indicative of affinity.

Mateos *et al.* [9] report that k must be tuned to balance reliability and sparseness. We test with $k = 10, 20$, and 30 , finding that $k = 20$ allows us to obtain a network that does not underrepresent any of the income distribution deciles. In addition, we removed rare surnames with less than twenty occurrences. Next, we removed surnames that fall in the periphery of the network using a k -core analysis of the network. The nodes that fall below a threshold $k = 2$ were removed, preserving the triangles that form the atomic units of relationships that characterize

the network. The resulting network contains 19041 nodes and 187563 edges, with an average degree of 19.7. We named this network **SA19k** (Surname Affinity **19k** entities).

Fig 1 shows the distribution of node degrees and socioeconomic distributions by decile. SA19k's degree distribution exhibits that a few nodes have many connections while most nodes have low connectivity. In specific, most nodes have around 5–7 neighbors (**Fig 1a**) and a small number of nodes have more. The income decile with the highest presence in SA19k is the one with the highest income (d10), with almost 29% of the network nodes (**Fig 1b**). The lowest decile has a minor presence, and each one of the deciles 2 to 9 have around 8% of the network nodes. This shows that surnames at the higher end of the socioeconomic distribution (d9-d10) are more prone to clustering.

Fig 2 visualizes SA19k, with colors red and blue corresponding to the main communities identified using Louvain's modularity maximization based community detection algorithm [10]. Nodes in the red community are more present in deciles 1–7, and nodes in the blue community belong to deciles 8–10. The network has a high modularity score of 0.802, which increases along with socioeconomic status.

Surname affinity prediction model

We study SA19k using predictive modeling. SA19k can be seen as a knowledge graph for surname affinity conditioned on income distribution. A predictive model based on SA19k can reveal unexpected connections between surnames. Our approach to predictive modeling is based on knowledge graph completion, where a model learns representations of nodes and edges to predict missing links. Link prediction reveals triplet-based facts, where each triplet represents a source and target surname, and each edge corresponds to a decile in the income distribution.

We model SA19k as a knowledge graph as follows. Let \mathcal{V} be a set of nodes representing entities and let \mathcal{R} be a set of relations. We define $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{R} \times \mathcal{V}$ as the set of edges connecting pairs of entities in $\mathcal{V} \times \mathcal{V}$ through a specific relation $r \in \mathcal{R}$. A knowledge graph $\mathcal{G}(\mathcal{V}, \mathcal{E}, \mathcal{R})$ is a set of observed triplets $\langle h, r, t \rangle$, such that $h, t \in \mathcal{V} \times \mathcal{V}$ and $r \in \mathcal{R}$. The nodes h, t represent head and tail entities connected through r . We distinguish between head and tail entities when r is a directed relationship that starts from h and goes to t . $\mathcal{G}(\mathcal{V}, \mathcal{E}, \mathcal{R})$ is an undirected knowledge graph if every $r \in \mathcal{R}$ has the same meaning in both directions. In the case of SA19k, we model relationships in a non-directional way, without distinguishing between heads and tails and therefore, using maternal and paternal surnames interchangeably.

Knowledge graph completion is the task of using \mathcal{G} to predict missing facts. Given $r \in \mathcal{R}$, tail prediction addressed the $\langle h, r, ? \rangle$ task, i.e., for a known head entity h and a given relation r , we infer from \mathcal{G} the most likely tail node to which h will connect through r . Analogously, head prediction is a $\langle ?, r, t \rangle$, i.e., predicts the most likely head entity h to connect t through r .

Machine learning-based approaches for knowledge graph completion use \mathcal{G} to capture latent features of entities and relations. These features embody the semantic of the original knowledge graph, vectorizing nodes and relationships. For this purpose, the method learns latent representations, known as embeddings, defining a scoring function ϕ to estimate the probability of $\langle h, r, t \rangle$ of being observed in the empirical network. The scoring function operates on the embeddings of h, r , and t , represented in vector notation by $\mathbf{h}, \mathbf{r}, \mathbf{t}$. Then, the method calculates $\phi(\mathbf{h}, \mathbf{r}, \mathbf{t})$, fitting the embeddings to maximize the score of the observed facts, and also minimizing the score of unobserved facts in \mathcal{G} .

During the prediction phase, for a given head entity h and a given relation r , missing links are inferred determining the entity that gets the highest score according to the scoring

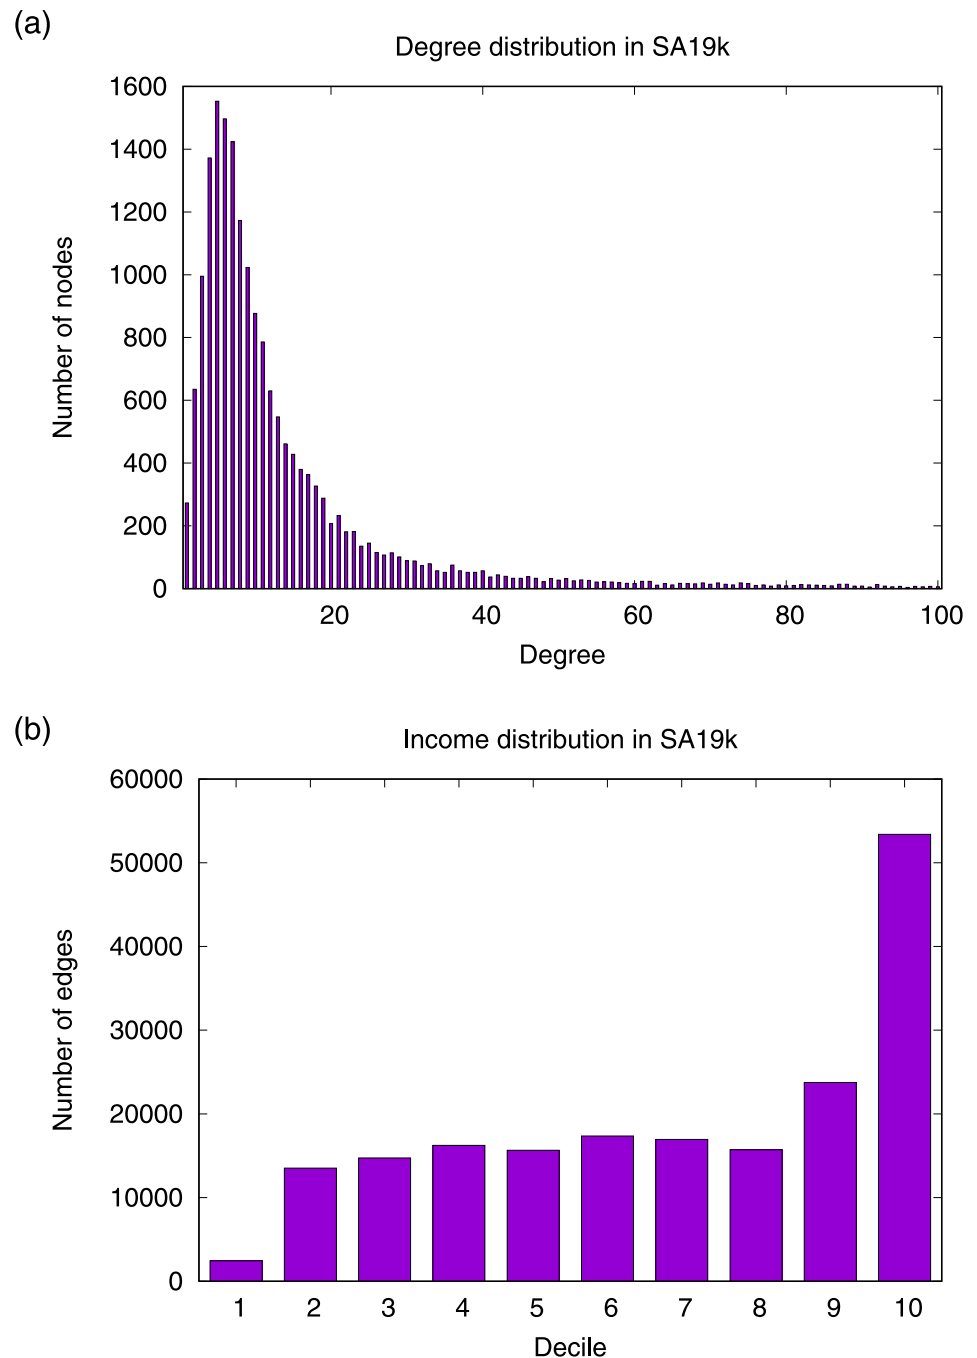


Fig 1. Node degree and income decile distribution in SA19k. a) Even though many nodes have around 5–7 neighbors, the degree distribution shows the presence of hubs; b) Most of the affinity relationships occur in the highest income decile, totaling the 29% of the edges of SA19k. On the other hand, affinities in the lowest decile are almost nonexistent. The other income deciles are quite similar, with around 8% of the edges per decile.

<https://doi.org/10.1371/journal.pone.0256603.g001>

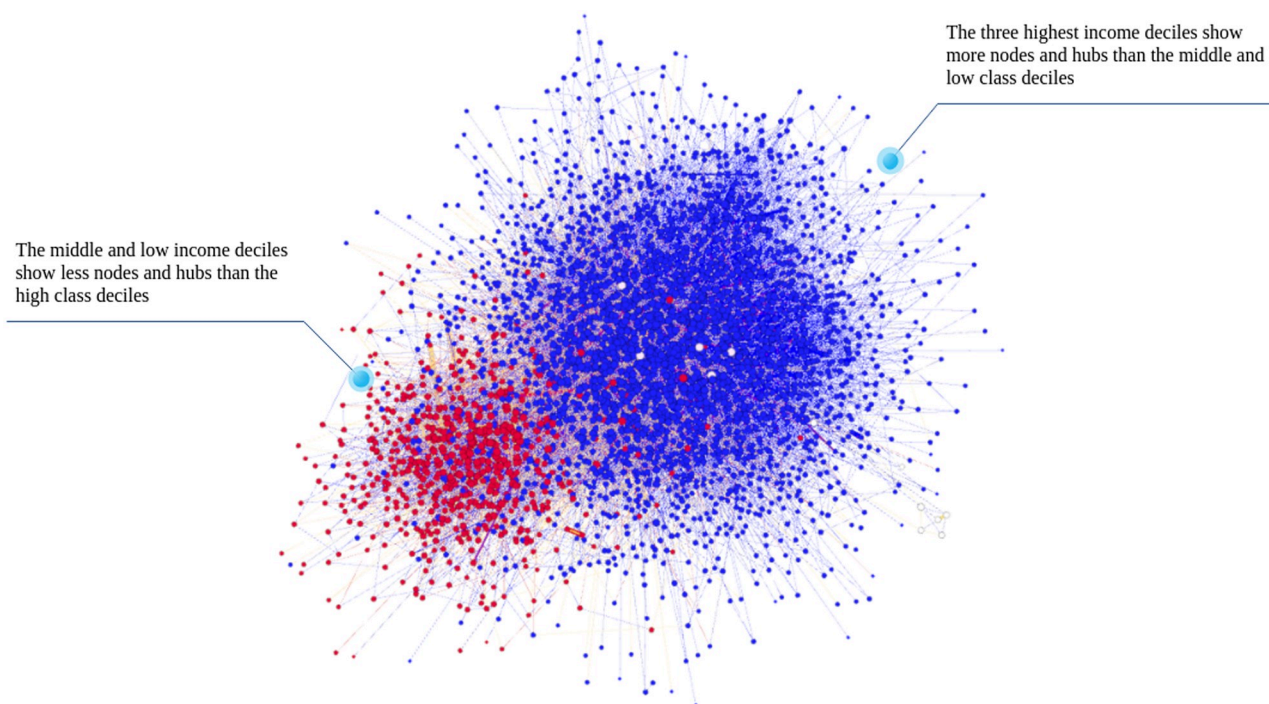


Fig 2. A visualization of SA19k. SA19k is highly modular, with modularity score = 0.802. Low-middle, and high-income deciles are grouped into two modularity-based partitions. The blank nodes do not belong to any of them. Low-middle income distribution deciles (in red color) show less nodes and hubs than the rest of the network. High income deciles (in blue color) show more nodes and hubs than the rest of the network. The size of each node represents its degree.

<https://doi.org/10.1371/journal.pone.0256603.g002>

function:

$$\arg \max_{e \in \mathcal{V}} \phi(\mathbf{h}, \mathbf{r}, \mathbf{e}).$$

Similarly, the link prediction task can be defined for head entities, looking for the entity $e \in \mathcal{V}$ that maximizes the scoring function $\phi(\mathbf{e}, \mathbf{r}, \mathbf{t})$ for a fixed r and t .

Machine learning-based approaches for knowledge graph completion

To cover the broadest possible range of methods and architectures in the evaluation, we identified representative methods of different model families, taking care that these methods achieve state-of-the-art performances in knowledge graph completion and have open-source implementations that favor the reproducibility of the reported results. We identified three families of models:

- **Tensor decomposition models:** These models are based on tensor decomposition techniques of the KG adjacency matrix, which is modeled as a three-dimensional tensor. The tensor representation consists of a set of stacked relationship adjacency matrices. These tensors are decomposed into a combination of low-dimensional vectors used as embeddings for the entities and relations of the KGs. Since tensor decompositions, in general, do not overfit the empirical KG, they are expected to have better generalization capabilities than other methods. We identified three methods of this family that show competitive results in benchmark data [11] and that have implementations that allow reproducing their results. These methods

are DistMult [12] and ComplEx [13], both belonging to bilinear type models since they represent the relationship matrix as a bi-dimensional tensor. Another prominent method of this family is the TuckER [14] method, a non-bilinear method based on Tucker's tensor decomposition [15].

- **Geometric models:** Geometric models interpret relationships as geometric transformations in the latent space of entity embeddings. These methods adjust the parameters of the node embeddings and the geometric transformation function that represents the relations. Given a head embedding, the transformation in the latent space maps it to the tail embedding of the observed fact. The parameters of the embeddings and relationships are adjusted to minimize the distance in the latent space between the head and tail embeddings connected by empirical facts. Geometric models need to verify constraints to ensure that the transformations in latent space are affine. We identified four geometric models representing this family: TransE [16] and TransH [17], based on translational type transformations; CrossE [18], based on translations with additional embeddings for KG relationships; and RotatE [19], which uses rotations (translations in polar coordinates) to represent relationships between entities in latent space.
- **Deep learning models:** These models use layered neural network architectures to learn the embeddings of the KG entities and relationships. The embeddings are learned in conjunction with the weights and biases of each layer of the network, which makes these models gain in expressiveness but assume risks of overfitting. We identified three methods from this family of models that show competitive results in benchmark data [11]. ConvE [20] uses a convolutional network architecture to encode the head entity and relationship embeddings, which produce an outcome in the network output layer that encompasses the KG's set of tail embeddings. The network is trained to maximize the scores of the observed facts, minimizing a categorical cross-entropy loss function. ConvTransE [21] reuses this architecture so that the transformation in the latent space between head and tail embeddings corresponds to translations. Finally, SACN [21] extends ConvTransE by adding a decoder that reconstructs the observed facts, replacing the loss function with a logit softmax.

Model selection

For model selection, we partitioned SA19k into training, validation, and testing folds. We did this by sampling triplets at random. Both the validation and testing partitions have 5000 triplets each, leaving the training partition with 177563 triplets. The proportions between training/validation/testing triplets are typical in knowledge base completion, and very similar to those used to study predictability in the Wordnet database [20]. In [S1 Table](#), we present the results of the validation; with TuckER producing the best statistics. The results in [S1 Table](#) show that TuckER obtains the highest precision in the top positions of each list of ranked entities, both in top-1 and top-3 lists. While its top-10 results are competitive but outperformed by RotatE, at MRR, TuckER also outperforms its competitors.

Tucker decomposition for link prediction

The best performing method in SA19K, TuckER [14], is a link prediction method based on Tucker's tensor decomposition [15]. The representation of the knowledge graph used by this method corresponds to a tensor (a multidimensional array) produced by the tensor products of N vector spaces. The order of a tensor corresponds to its number of dimensions, 3 in the

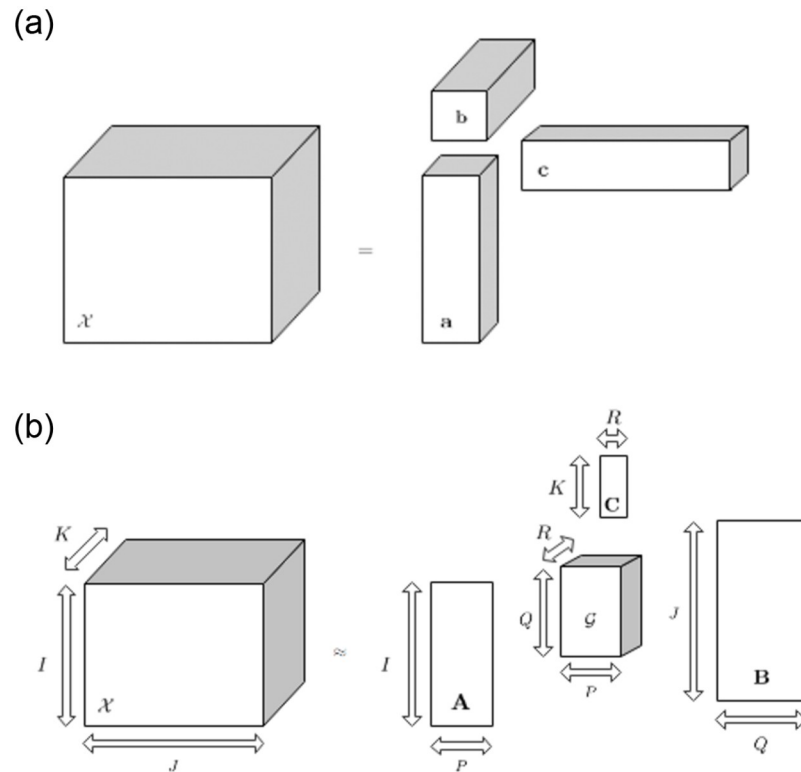


Fig 3. a) Three-way tensors are produced by the tensor products of 3 vector spaces; **b)** The Tucker decomposition factorizes a three-way tensor in a core tensor and three matrices.

<https://doi.org/10.1371/journal.pone.0256603.g003>

case of the usual Tucker decomposition. We will denote the input $\langle i, j, k \rangle$ of a three-way tensor \mathcal{X} by x_{ijk} .

We introduce some definitions needed to define the Tucker decomposition properly. An n -way tensor $\mathcal{X}^{I_1 \times \dots \times I_n}$ is rank-one if it can be expressed by the outer product of N vectors: $\mathcal{X} = a^{(1)} \circ a^{(2)} \circ \dots \circ a^{(n)}$, where $a^{(i)}$ represents the i -th vector of a sequence. Note that each entry of \mathcal{X} can be written as $x_{i_1 i_2 \dots i_n} = a_{i_1}^{(1)} \cdot a_{i_2}^{(2)} \cdot \dots \cdot a_{i_n}^{(n)}$. Three-way tensors are formed by $\mathcal{X} = a \circ b \circ c$, as shown in Fig 3a.

An essential operator of the Tucker decomposition is the N -mode product of a tensor with a matrix. Although N -mode products can be generalized to higher-order tensors, we focus only on the N -mode product between a tensor and a matrix since the Tucker decomposition uses it. The reader can check higher-order N -mode products and other types of tensor decomposition in [22]. Let \mathcal{X} be a tensor in $\mathbb{R}^{I_1 \times I_2 \times \dots \times I_n}$ and \mathcal{U} a matrix in $\mathbb{R}^{J \times I_N}$. The N -mode product between \mathcal{X} and \mathcal{U} along its N mode is denoted by $\mathcal{X} \times_N \mathcal{U}$ and has size $I_1 \times \dots \times I_{N-1} \times J \times I_{N+1} \times \dots \times I_n$. An entry in $\mathcal{X} \times_N \mathcal{U}$ is defined by:

$$(\mathcal{X} \times_N \mathcal{U})_{i_1 \dots i_{N-1} j i_{N+1} \dots i_n} = \sum_{i_N=1}^{I_N} x_{i_1 i_2 \dots i_n} \cdot u_{ji_N}.$$

The Tucker decomposition factorizes a tensor \mathcal{X} in a core tensor \mathcal{G} transformed by the N -mode tensor product along each of tensor's modes (dimensions). In the three-way case, where

\mathcal{X} in $\mathbb{R}^{I \times J \times K}$:

$$\mathcal{X} \approx \mathcal{G} \times_1 \mathbf{A} \times_2 \mathbf{B} \times_3 \mathbf{C} = \sum_{p=1}^P \sum_{q=1}^Q \sum_{r=1}^R g_{pqr} a_p \circ b_q \circ c_r,$$

where $\mathcal{G} \in \mathbb{R}^{P \times Q \times R}$ is called the core tensor of the decomposition, and $\mathbf{A} \in \mathbb{R}^{I \times P}$, $\mathbf{B} \in \mathbb{R}^{J \times Q}$, and $\mathbf{C} \in \mathbb{R}^{K \times R}$ can be considered the main components of each mode of \mathcal{X} . \mathcal{G} is dimensionally smaller than \mathcal{X} ($P < I$, $Q < J$, and $R < K$) and can be viewed as a compressed version of \mathcal{X} , as we show in Fig 3b. An element-wise Tucker decomposition entry is given by:

$$x_{ijk} \approx \sum_{p=1}^P \sum_{q=1}^Q \sum_{r=1}^R g_{pqr} \cdot a_{ip} \cdot b_{jq} \cdot c_{kr}, \quad i \in 1 \dots I, j \in 1 \dots J, k \in 1 \dots K.$$

Balazevic *et al.* [14] proposes a model for link prediction based on the Tucker tensor decomposition of the binary tensor representation of the knowledge graph. TuckER assumes that head and tail embeddings belong to the same vector space. Accordingly, TuckER defines an entity embedding matrix $\mathbf{E} \in \mathbb{R}^{n_e \times d_e}$, with $\mathbf{A} = \mathbf{C} = \mathbf{E}$. Relations are modeled using a relation embedding matrix $\mathbf{R} \in \mathbb{R}^{n_r \times d_r}$, which corresponds to the factor matrix \mathbf{B} of the Tucker decomposition. Then, the tensor representation \mathcal{X} of the knowledge graph is decomposed as:

$$\mathcal{X} \approx \mathcal{G} \times_1 \mathbf{E} \times_2 \mathbf{R} \times_3 \mathbf{E}.$$

The scoring function for triplet facts is defined for each triplet $\langle e_h, w_r, e_t \rangle$ using the element-wise tensor product per node:

$$\phi(\mathbf{e}_h, \mathbf{w}_r, \mathbf{e}_t) = \mathcal{G} \times_1 \mathbf{e}_h \times_2 \mathbf{w}_r \times_3 \mathbf{e}_t,$$

where \mathbf{e}_h and \mathbf{e}_t are rows in \mathbf{E} and \mathbf{w}_r is a row in \mathbf{R} . TuckER applies a logistic sigmoid to the scoring function obtaining the predicted probability p of a triplet being true.

For training, TuckER adds a reciprocal relation $\langle e_t, w_r^{-1}, e_h \rangle$ for every triplet fact $\langle e_h, w_r, e_t \rangle$. The training step uses a 1: N scoring, i.e., for a given triplet $\langle e_h, w_r, e_t \rangle$, TuckER scores $\langle e_h, w_r \rangle, \forall e_t \in \mathcal{E}$ and $\langle e_t, w_r^{-1} \rangle, \forall e_h \in \mathcal{E}$. Model fitting minimizes the Bernoulli log-likelihood loss function, that for a given triplet is defined by:

$$\mathcal{L} = -\frac{1}{n_e} \sum_{i=1}^{n_e} (y^{(i)} \log(p^{(i)}) + (1 - y^{(i)}) \log(1 - p^{(i)})),$$

where y is the binary label vector of observed/unobserved facts in the knowledge graph.

Results

Hyper-parameters tuning

The best fit reported in S1 Table was trained using Adam [23], a method for first-order gradient-based optimization of stochastic objective functions, based on adaptive estimates of lower-order moments. This method is suitable for problems with sparse gradients. We used batch learning, with batch size = 128, learning rate = 0.005, and decay rate = 1.0. We choose all hyper-parameters by search based on validation set performance. The dimensionality of the relational space was tested with values in {10, 20, 30}. The best result was obtained for $d_r = 10$, which matches the number of relationships in SA19K. The dimensionality of the entity space was tested with values in {100, 200, 500, 1000}. The best result was obtained for $d_e = 200$. TuckER considers three dropout factors. The first operates on the input, which corresponds to

the tensor representation of the entities. A second dropout factor applies to the tensor product between the tensor representations of the relationships and the Tucker core tensor. Finally, a third dropout factor is applied when calculating the tensor product between the previous tensor and the tensor representation of the entities. This final product corresponds to the tensor representations of the pair $\langle s, r \rangle$, which is used for the 1:N ranking-based training strategy. These parameters were evaluated using grid search in $\{0.2, 0.3, 0.4, 0.5\}$. The best configuration corresponds to factors $\{0.5, 0.2, 0.2\}$. Balazevic *et al.* [14] comment that low dropout factors are required for datasets with a high number of training triplets per relation, avoiding the risk of over-fitting. Our high input dropout value suggests that the number of training triplets is low regarding the number of entities available in SA19K. However, the proportion of examples per relationship is high due to the low number of relations in SA19K, which explains the lowest values obtained for the second and third dropout factors.

Performance per income decile

Fig 4 shows the relation matrices learned by TuckER for each income decile. Although TuckER does not force the learning of symmetric relationships, the learned matrices are approximately symmetric, suggesting a good fit of the model to SA19k. Heatmaps with more red entries indicate that TuckER encodes more information in that relationship than in the rest of the deciles. Fig 4 shows that TuckER encodes many information in relationships D1 and D5, which are the ones with the most burned entries in the tensor. On the other hand, D2-D4 show fewer burned entries in the tensor, suggesting that these relationships are less informative than the others. SA19K shows that low and middle-income deciles have less diverse relationships between surnames, with a low incidence of affinity ties in this network segment. The effort made by TuckER in D1 is related to the low incidence of ties in this layer (see Fig 1(b)). In D5, this income decile articulates the low-income deciles (D1-D5) with the middle and high-income deciles (D6-D10), suggesting a transition in SA19K. The relation matrices D6-D9 obtained by TuckER are quite similar to each other, showing homogeneity in this income range in terms of surname affinity. In D10, TuckER shows more burned entries in the tensor diagonal, which coincides with the high incidence of affinity ties in this segment of the network (see Fig 1(b)).

For each triplet in the test partition, TuckER evaluates the pair $\langle s, r \rangle$ with the rest of the network entities, using a 1:N ranking-based strategy. Accordingly, for the 19041 triplets evaluated, TuckER orders the triplets in descending order according to the softmax layer's logit. Then, the ordered list of triplets is evaluated by counting hits@n. A grounded triplet is a hit@n if it is among the first n results with the model's highest logit. It should be noted that the problem is very hard since the objective is to rank grounded triplets in the top positions of a long list of possible ties. According to the degree distribution in SA19K, on average, each entity has 20 affinity ties. Since the length of each entity list is 19041, the chances of generating a top@n list of maximum precision at random are $\prod_{i=1}^{20} \frac{21-i}{19041-i} \approx 0.00105^{20}$, a extremely small number. TuckER, on average, gets hits@1 = 0.214, hits@3 = 0.406, and hits@10 = 0.607, as shown in S1 Table. These results indicate that surname affinity ties are strongly predictable.

Fig 5a shows hits @1, @3, and @10 for each SA19k income decile in the test partition. While decile 1 is the one with the highest predictability, decile 10 is the least predictable. Between deciles 2 and 10, all hits @10 reach around 0.6. These results are related to the diversity of surnames in each socioeconomic segment. In [1], we showed that the diversity of surnames, quantified according to the alpha coefficient [24], is much greater in high socioeconomic segments than in low segments. Affinity tie prediction results indicate that the higher the surname diversity, the more difficult it is to predict the affinity ties. Regarding the

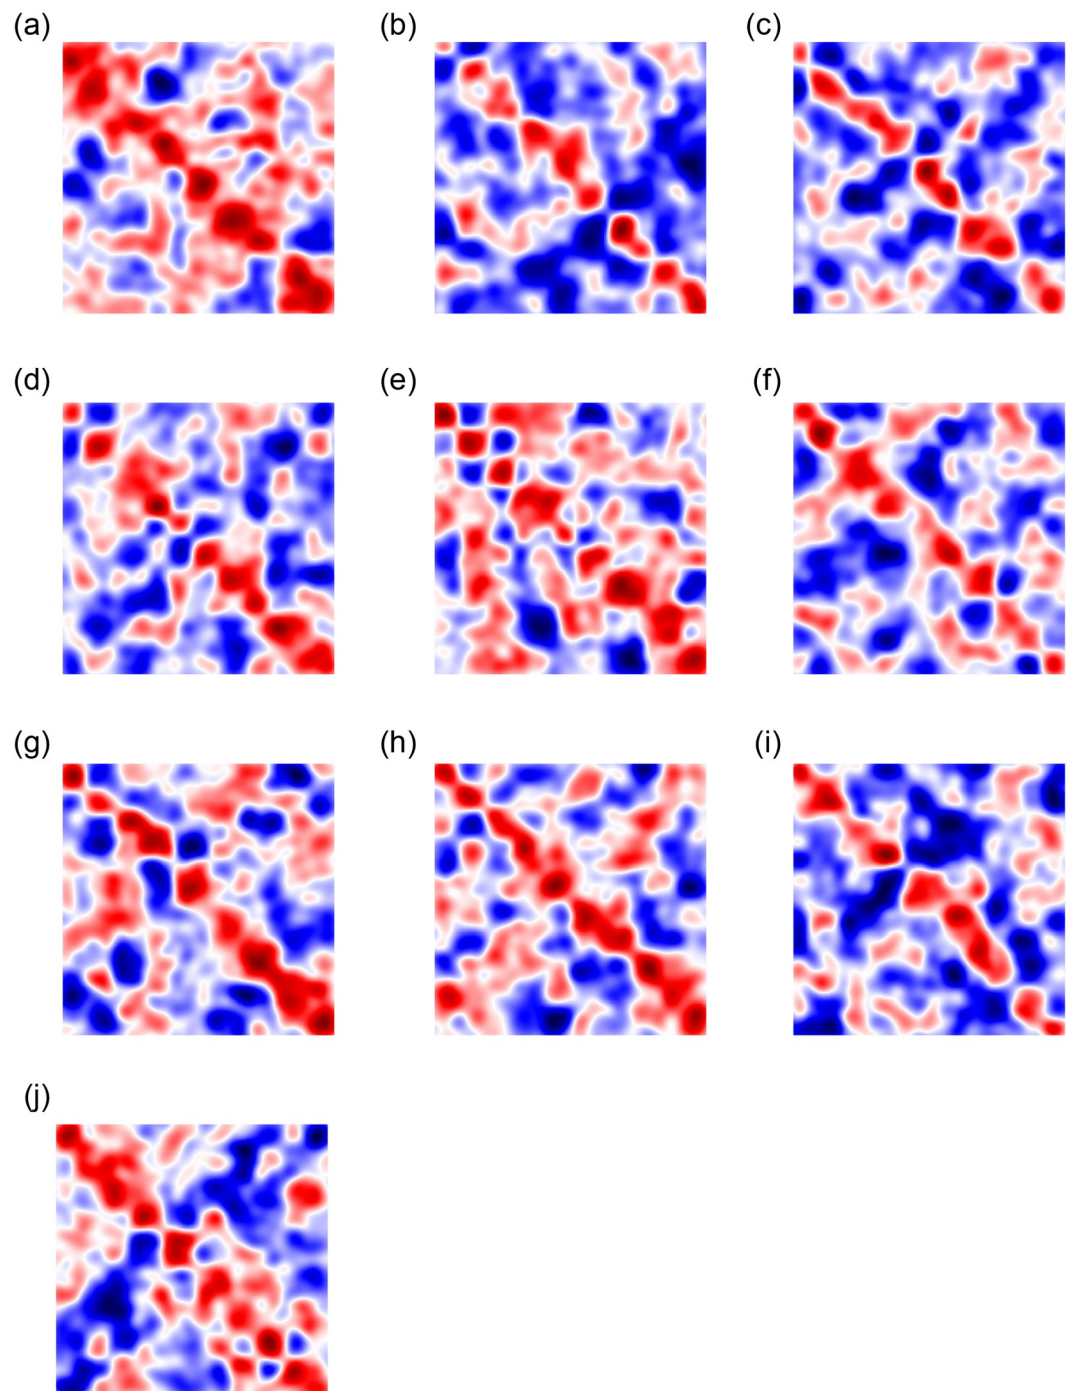


Fig 4. Heatmaps of the relation matrices learned with Tucker for each income decile. Red and blue colors indicate high and low-value entries in the relation matrices, respectively. These matrices are computed using k -mode tensor products between \mathbf{R} and the second mode of the core tensor \mathbf{W} . Although the model does not force the learning of symmetric relationships, all the learned relationships are approximately symmetric, indicating a good fit of the model to SA19k.

<https://doi.org/10.1371/journal.pone.0256603.g004>

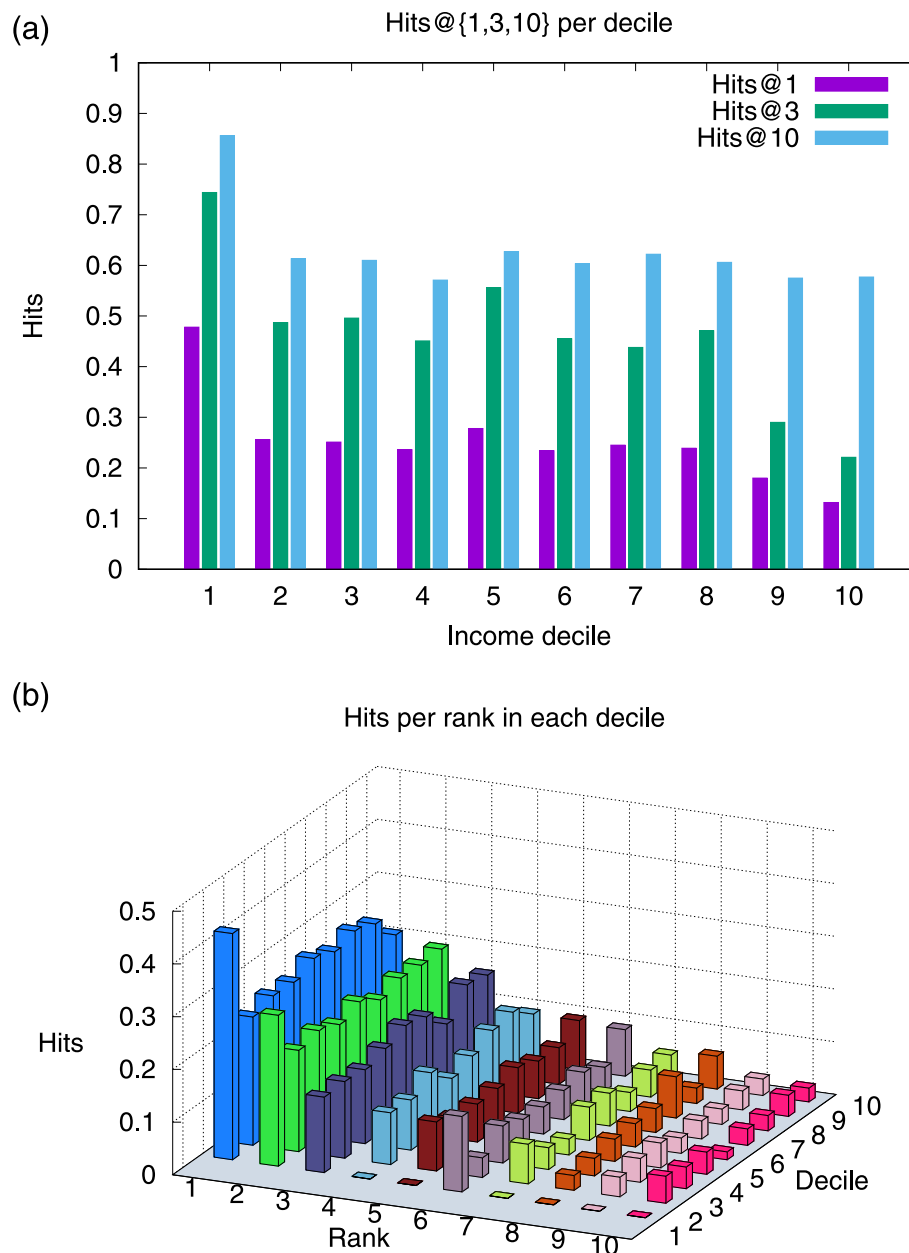


Fig 5. TuckER manages to rank high many grounding triplets, which allows the model to perform well on hits @1, @3, and @10; a) While the results in hits@1 show a better performance in the lower deciles, in hits@10, the results are more even; b) When disaggregating the results in absolute hits per rank position, we see that TuckER does not have significant differences between deciles, placing the vast majority of hits in the first five positions of each list.

<https://doi.org/10.1371/journal.pone.0256603.g005>

precision of the method, the greatest variations are found in hits@1, where the affinity of surnames becomes less predictable as the income decile increases. Fig 5b shows the results disaggregated at the rank-level, i.e., the absolute number of hits by ranking position. We can see that the hits decrease as the position rank increases, which indicates that the TuckER logits help to rank in the top positions the triplets with the highest probability of generating affinity ties. Most of the hits are in the top five ranking positions.

Explaining surname affinity predictions

We evaluate the ability of our model to forecast the emergence of new triangles. We consider three sources of information. The first source consists of the original network, which we call grounded triplets. If two unconnected nodes have many neighbors in common, a quantifiable measure known as the shared nearest neighbors fraction (SNN), the model expects them to connect. SNN measures the fraction of neighbors in common over the union of the neighborhoods of two entities. SNN is a very relevant measure in clustering analysis [25], since it reveals clues for high-dimensional data clustering of varying density. In social networks, SNN relates to the friend of a friend mechanism (FOAF), which explains that two strangers have a high probability of engaging if they have friends in common. To account for this effect, we measure the SNN between the source and target entity of each new grounded triplet in the same income decile. The second source of information is the SNN fraction in nearest income deciles. If two unknown nodes share many neighbors in neighboring income deciles, the model expects them to connect. We estimate this effect by counting the fraction of SNN in the same, and the two nearest deciles. The third source of information is the neighborhood in the embedding space. We measure this effect by counting the fraction of SNN in the embedding space, according to the k -mode tensor product of the entities of the triplet along with the relation matrix that represents each income decile. To measure the SNN in the embedding space, we compute a k NN query using the Euclidean distance with $k = 50$, for each entity embedding of a grounded predicted triplet. If the neighbors of two disconnected nodes intersect in the embedding space (i.e., both entities show a high fraction of SNN), the model expects them to connect.

Fig 6a shows that the SNN fraction varies by source of information. The lowest fraction of SNN in predicted triplets is obtained from grounded triplets of the same decile. This pattern is more robust in the lowest income deciles, losing relevance as the socioeconomic level increases. After adding near income deciles (we use the two closest deciles in the evaluation), the fraction of SNN increases in all deciles. This pattern shows that this second source of information is more informative than the first. Finally, after adding friends using neighbors in the embedding space, the fraction of SNNs increases. This finding shows that the node embedding space allows capturing ungrounded shared neighbors, representing valuable information for all income deciles. The embedding space is the most informative dimension for the affinity prediction task. In specific, this space is very relevant in income deciles 1, 2, and 7. In fact, in all income deciles, the grounded predicted triplets share more neighbors in this space than in the empirical network. This finding suggests that the entity embedding space can encode high-order information that is different to the one encoded from the neighborhoods of the empirical network.

Fig 7 shows examples of grounded predictions. All of them are characterized by connecting two micro-clusters adding an interaction between them. In some cases (Fig 7a and 7d), we observe an absorption pattern, where the interaction occurs between micro-clusters of very different sizes. In other cases, the interaction occurs between clusters of similar sizes (Fig 7b and 7c). In all these cases, the new interaction (marked with a red dotted line) creates a new path between both clusters, decreasing the resulting graph's diameter.

Most of the predicted triplets (57% on average across deciles) come from micro-clusters without direct connections like those shown in Fig 7. TuckER manages to learn proximity relationships between nodes obtained by working with the tensor factorization of the entire network. This principle allows identifying long-term interaction patterns between unconnected nodes, which cannot be detected by locality-based methods such as those based on FOAF patterns. Fig 6b shows the fraction of grounded predictions per income decile that corresponds to high SNN fractions in the same decile and/or in close income deciles, which on average

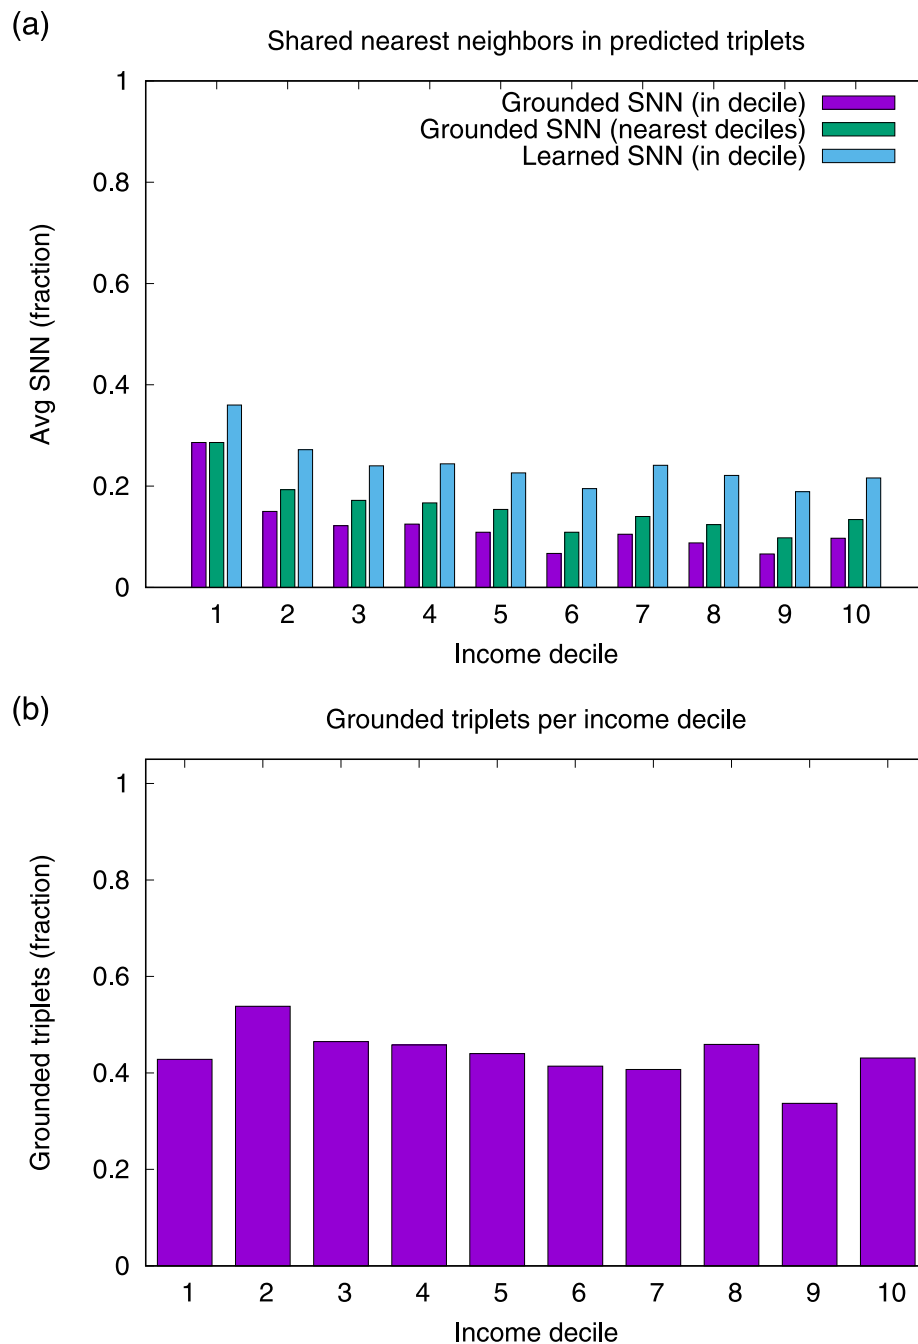


Fig 6. Shared nearest neighbors (SNN) and grounded triplets per income decile. a) The fraction of SNN (average) varies depending on which source of information we consider. The most informative source is proximity-based SNN, although there is significant support based on grounded SNN from the original network. b) Grounded triplets per income decile. The fraction of grounded triplets is always less than half of the predicted triplets, except in decile 2, where it reaches almost 60% of the predictions.

<https://doi.org/10.1371/journal.pone.0256603.g006>

represents the 43% of the successfully predicted triplets. Fig 6b shows that more than half of the predicted triplets are grounded on the entity embedding space. Accordingly, Tucker must make a significant effort to identify triplets based on proximity-based neighbors inferred from the entity embedding space, representing more than half of the correctly predicted triplets.

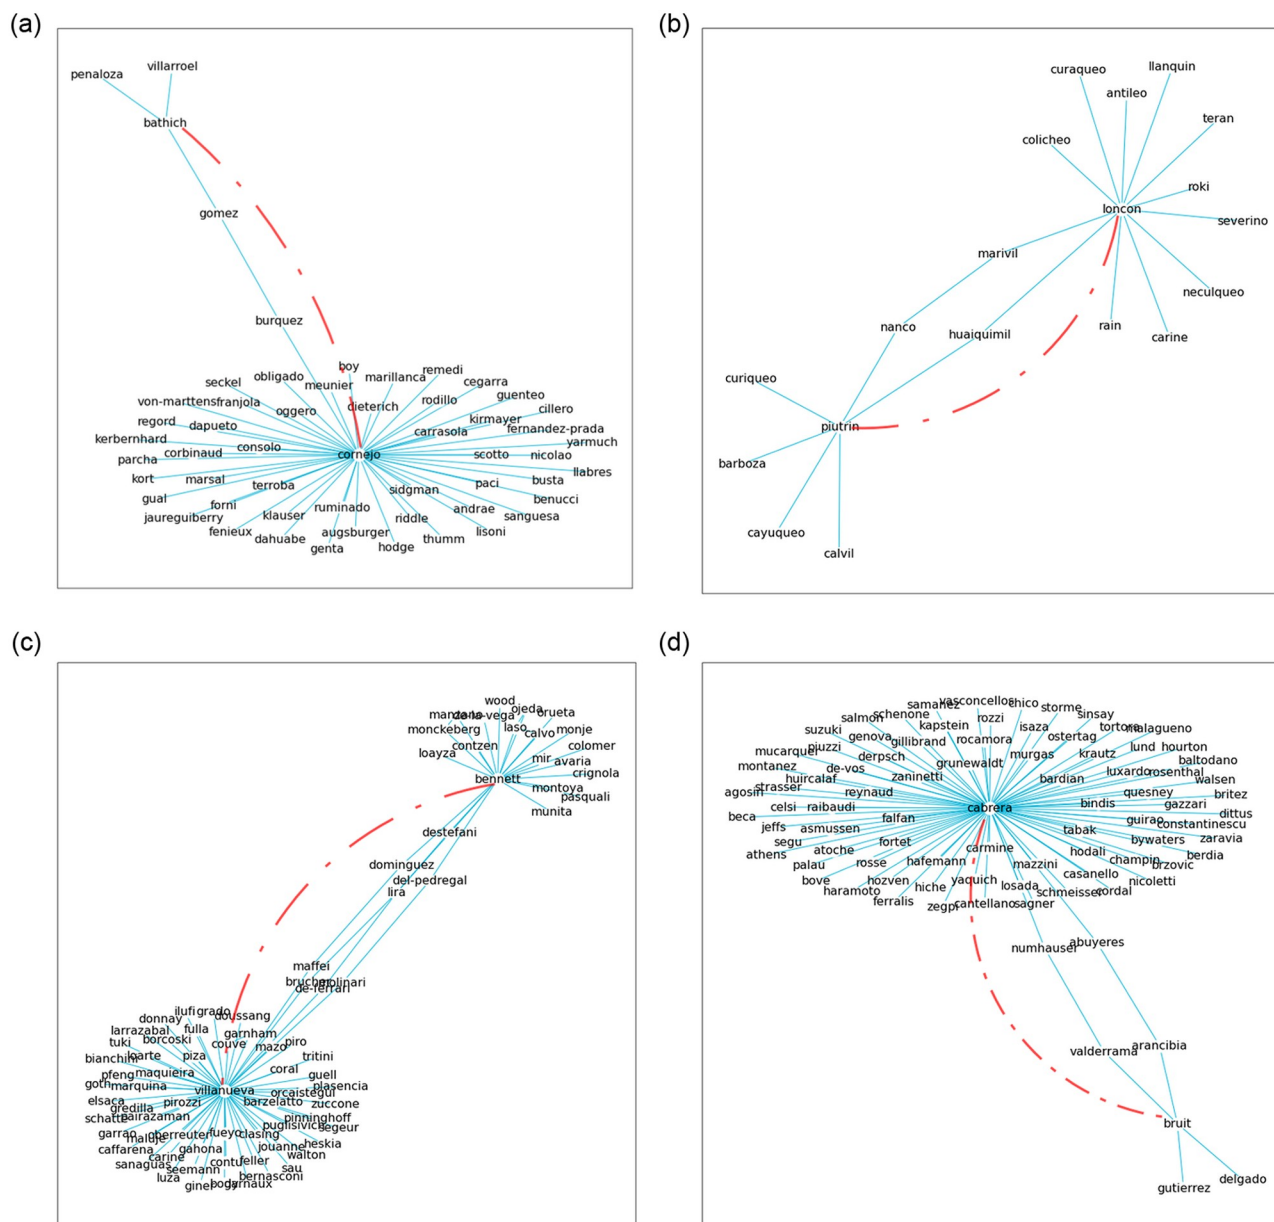


Fig 7. Grounded predictions generally show paths that connect two micro-clusters. Figures a and d show the connection of a larger cluster and a smaller one, exhibiting an absorption pattern. Figures b and c show the relationship of two clusters of similar sizes. In all the cases indicated, the new connection (marked with a red dotted line) creates a shorter path between both micro-clusters, decreasing the resulting graph's diameter.

<https://doi.org/10.1371/journal.pone.0256603.g007>

Discussion

The results show great predictability in the formation of ties between people holding certain surnames. Triadic closure is an empirical regularity well established in the social sciences [7]. If persons A and B share a common friend C, they are much likelier to become friends than in the alternative scenario where they do not share C as a friend. So far, most research has examined triadic closures either theoretically [4, 7], to statistically model networks [6], or in micro-

empirical contexts. This paper empirically investigates the formation of triads at the level of an entire population. It shows that surname-based groups are prone to form triads.

Chile is one of the most unequal countries in the world [26], with a pattern of social mobility where it is especially hard to enter the elites. [27]. Our previous research [1] showed that Santiago's upper class is highly endogamous, not unlike an ethnic minority in its tendency to form in-group marital unions. This pattern can be explained in two ways: from the preferences of wealthy individuals to connect with others similar to them, or from the structure of the network around them.

The results of this study reveal the importance of the second mechanism, a dimension of socioeconomic segregation that is independent from homogamy: the subjective preference that people may have to connect with others like themselves [28]. People also connect with and marry their socioeconomic in-groups because the latter inhabit the same locations in the network. In this sense, an upper-class person may not have a particular preference of marrying someone of her own status group, but she may end up doing so regardless because the network of social interactions constrains the marital options available to her.

Importantly, we make a distinction between empirical and more abstract forms of network proximity. Neighbors may come in the form of grounded triplets, but also as proximate points in the embedding space. We show that triadic closure occurs not only in the empirical network, but also in the embedding space. If nodes A and B are both similar to C, they have a good chance of forming a tie in the network. Link prediction tasks therefore would benefit from taking not only the empirical graph into account but also the embeddings of nodes.

Conclusions

This paper develops a model to predict the formation of new links in the network of surname affinity in Santiago, Chile (SA19k). It formulates the problem as a knowledge base completion task, and finds that the TuckER method produces the best results. The results show that the formation of affinity links between surnames in Santiago, Chile, is highly predictable. We find that proximity in the network explains a large proportion of new links. Importantly, the method distinguishes between proximity in the empirical network and proximity in the embedding space, and finds that the latter is more predictive of link formation than the former.

The results shed light on an unexplored dimension of socioeconomic segregation. People may or may not prefer connecting with and marrying others belonging to their same socioeconomic status. However, given that the network of social interactions places constraints on their options, they are still likely to connect with or marry in-groups. Extensions of this paper could incorporate proximity in the embedding space as a parameter in statistical models that aim to disentangle the effects of different types of network proximity in explaining the social segregation that we often observe in empirical networks.

Supporting information

S1 Table. Link prediction for SA19k. The results show the performance of the models in hits@{1,3,10}, and Mean Reciprocal Rank (MRR). (PDF)

Acknowledgments

Thanks to Josué Tapia and Andrés Cruz for helping in the data-building phase of this project. Other persons that helped are Daniel Alcatruz, Sebastián Huneeus, and Johans Peña.

References

1. Bro N., Mendoza M. Surname affinity in Santiago, Chile: A network-based approach that uncovers urban segregation, PLoS ONE, 16(1): e0244372, 2021. <https://doi.org/10.1371/journal.pone.0244372>
2. Goodreau J., Kitts J., Morris M. Birds of a Feather, or Friend of a Friend? Using Exponential Random Graph Models to Investigate Adolescent Social Networks. *Demography* 46:1, 2009. <https://doi.org/10.1353/dem.0.0045>
3. Wimmer A., Lewis K. Beyond and Below Racial Homophily: ERG Models of a Friendship Network Documented on Facebook. *American Journal of Sociology*, 116: 2 2010. <https://doi.org/10.1086/653658>
4. Burt R. Structural holes: the social structure of competition. Harvard University Press, 1995.
5. Simmel G. *Sociology: inquiries into the construction of social forms*. Brill, 2009.
6. Snijders T., van de Bunt G., Steglich C. *Introduction to stochastic actor-based models for network dynamics*. Social Networks, 2010.
7. Granovetter M. The strength of weak ties. *American Journal of Sociology*, 1973. <https://doi.org/10.1086/225469>
8. CIT (Centro de Inteligencia Territorial). Índice de Bienestar Territorial 2012. Santiago: Universidad Adolfo Ibáñez, 2012.
9. Mateos P., Longley P., O'Sullivan D. Ethnicity and Population Structure in Personal Naming Networks, PLoS ONE, 6(9): e21909399, 2011. <https://doi.org/10.1371/journal.pone.0022943>
10. Blondel V., Guillaume J., Lambiotte R., Lefebvre E. Fast unfolding of communities in large networks, *Journal of Statistical Mechanics: Theory and Experiment*, (10): P10008, 2008. <https://doi.org/10.1088/1742-5468/2008/10/P10008>
11. Rossi A., Barbosa D., Firmani D., Matinata A., Merialdo P. Knowledge Graph Embedding for Link Prediction: A Comparative Analysis. *ACM Transactions on Knowledge Discovery and Data*, 15(2): 14:1–14:49, 2021.
12. Yang B., Yih W., He X., Gao J., Deng L. Embedding Entities and Relations for Learning and Inference in Knowledge Bases, ICLR, 2015.
13. Trouillon T., Welbl J., Riedel S., Gaussier E., Bouchard G. Complex Embeddings for Simple Link Prediction, ICML, 2016.
14. Balazevic I., Allen C., Hospedales T. TuckER: Tensor Factorization for Knowledge Graph Completion, EMNLP/IJCNLP, 2019.
15. Tucker L. Some Mathematical Notes on Three-Mode Factor Analysis, *Psychometrika*, 31(3):279–311, 1966. <https://doi.org/10.1007/BF02289464>
16. Bordes A., Usunier N., Garcia-Duran A., Weston J., Yakhnenko O. Translating Embeddings for Modeling Multi-relational Data, NIPS, 2013.
17. Wang Z., Zhang J., Feng J., Chen Z. Knowledge Graph Embedding by Translating on Hyperplanes, AAAI, 2014.
18. Zhang W., Paudel B., Zhang W., Bernstein A., Chen H. Interaction Embeddings for Prediction and Explanation in Knowledge Graphs, WSDM, 2019.
19. Sun Z., Deng Z., Nie J., Tang J. RotatE: Knowledge Graph Embedding by Relational Rotation in Complex Space, ICLR, 2019.
20. Dettmers T., Minervini P., Stenetorp P., Riedel S. Convolutional 2D Knowledge Graph Embeddings, AAAI, 2018.
21. Shang C., Tang Y., Huang J., Bi J., He X., Zhou B. End-to-End Structure-Aware Convolutional Networks for Knowledge Base Completion AAAI, 2019.
22. Kolda T., Bader B. Tensor Decompositions and Applications. *SIAM Review* 51(3): 455–500, 2009. <https://doi.org/10.1137/07070111X>
23. Kingma D., Ba J. Adam: A Method for Stochastic Optimization. *International Conference on Learning Representations (ICLR)*, 2015.
24. Barraí I., Rodríguez-Laralde A., Mamolini E., Manni F., Scapoli C. Isonymy structure of USA population. *American Journal of Physical Anthropology*, 2001. 114:109–123. [https://doi.org/10.1002/1096-8644\(200102\)114:2%3C109::AID-AJPA1011%3E3.0.CO;2-I](https://doi.org/10.1002/1096-8644(200102)114:2%3C109::AID-AJPA1011%3E3.0.CO;2-I)
25. Zamora J., Mendoza M., Allende H. Hashing-based Clustering in High Dimensional Data. *Expert Systems with Applications*, 62:202–211, 2016. <https://doi.org/10.1016/j.eswa.2016.06.008>
26. OECD Broken Social Elevator?: How to Promote Social Mobility. OECD Publishing, 2018.

27. Torche F. Unequal But Fluid: Social Mobility in Chile in Comparative Perspective. *American Sociological Review*, 70:3, 2005. <https://doi.org/10.1177/000312240507000304>
28. McPherson M., Smith-Lovin L., Cook J. Birds of a Feather: Homophily in Social Networks. *Annual Review of Sociology*, 27:1, 2001. <https://doi.org/10.1146/annurev.soc.27.1.415>