

RESEARCH ARTICLE

The latent cognitive structures of social networks

Izabel Aguiar¹  and Johan Ugander²

¹Stanford University, Institute for Computational and Mathematical Engineering, Stanford, CA, USA and ²Stanford University, Department of Management Science and Engineering, Institute for Computational and Mathematical Engineering, Stanford, CA, USA

Corresponding author: Izabel Aguiar; Email: izabel.p.aguiar@gmail.com

Abstract

When people are asked to recall their social networks, theoretical and empirical work tells us that they rely on shortcuts, or heuristics. Cognitive social structures (CSSs) are multilayer social networks where each layer corresponds to an individual's perception of the network. With multiple perceptions of the same network, CSSs contain rich information about how these heuristics manifest, motivating the question, *Can we identify people who share the same heuristics?* In this work, we propose a method for identifying *cognitive structure* across multiple network perceptions, analogous to how community detection aims to identify *social structure* in a network. To simultaneously model the joint latent social and cognitive structure, we study CSSs as three-dimensional tensors, employing low-rank nonnegative Tucker decompositions (NNTuck) to approximate the CSS—a procedure closely related to estimating a multilayer stochastic block model (SBM) from such data. We propose the resulting latent cognitive space as an operationalization of the sociological theory of *social cognition* by identifying individuals who share *relational schema*. In addition to modeling cognitively *independent*, *dependent*, and *redundant* networks, we propose a specific model instance and related statistical test for testing when there is *social-cognitive agreement* in a network: when the social and cognitive structures are equivalent. We use our approach to analyze four different CSSs and give insights into the latent cognitive structures of those networks.

Keywords: Cognitive social structures; social cognition; multilayer networks; stochastic blockmodels; tensor decompositions

1. Introduction

The study of social networks often concerns itself with the notion of a *true* network. For a given relationship type, the true network is typically defined by what is captured when asking each individual about their own relationships (Wasserman and Faust, 1994). A common extension of this approach is to uncover multiple true networks by using multiple name generators (e.g., Campbell and Lee, 1991; Banerjee et al., 2013). Yet other approaches use interaction data to contrast physical contact networks to participants' faulty recollections (e.g., Killworth and Bernard, 1976; Bernard et al., 1979; Freeman et al., 1987). The study of cognitive social structures (CSSs) (Krackhardt, 1987) (sometimes referred to simply as "Krackhardt data"), where *perceptions* of the social network are collected from each individual within it, makes a radical departure from these approaches. Krackhardt proposed the idea that individuals' perceptions provide valuable information about the network, and through them, multiple notions of the network exist. Much of the subsequent body of work on CSSs has since focused on measuring how these perceptions differ from some sense of a common or underlying network. In this work we stray from the conception of a true network and instead intend to learn what perceived networks tell us about the joint social and cognitive structures that CSSs capture and what we can learn from considering these varied perceptions as a whole.

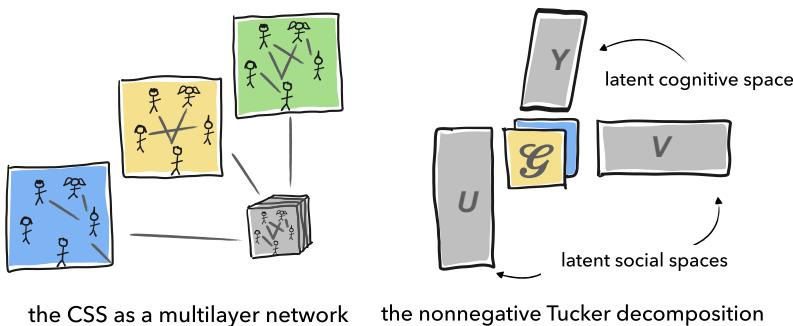


Figure 1. In this work, we analyze cognitive social structures (CSSs) as multilayer networks represented by an $N \times N \times N$ adjacency tensor (left). The *frontal slices* of the adjacency tensor are visualized in blue, yellow, and green. The N frontal slices of the CSS are adjacency matrices representing the perception each person has of their network. We use the nonnegative Tucker decomposition (NNTuck) to model the CSS with a multilayer stochastic block model, which decomposes the adjacency tensor into latent social spaces and a latent cognitive space (right).

CSSs capture information present, but usually ignored, in our social networks: our (varied) perceptions of the social structure. Newcomb (1961) first introduced the concept of incorporating this perspective in social network analysis by asking participants about ties involving both themselves and others, and Krackhardt (1987) expanded this perspective in his work that laid the foundations for the formal study of CSSs. For the purposes of this work we will consider the CSS as a *multilayer network*, wherein each layer is defined by a different individual's perception of the social network (see De Domenico et al., 2013, for a review of multilayer networks).

The present work makes the following contributions. First, our work proposes examining the CSS of a population to study the latent structure of the perceptions of individuals within the network. We approach this task using nonnegative Tucker decompositions. The nonnegative Tucker decomposition (NNTuck), which decomposes a tensor into three nonnegative *factor matrices* and one nonnegative *core tensor*, is a multilayer network extension of the degree-corrected, mixed-membership stochastic block model (SBM) proposed in Ball et al. (2011). Analogous to how the factor matrices in the single-layer SBM identify *node communities*, the additional third factor matrix in the NNTuck identifies *layer communities*. As such, the third factor matrix of the NNTuck allows for the adjacency tensor to be low rank in the layer dimension and identifies interdependencies between the layers of a network. In the case of the CSS multilayer network, the NNTuck identifies interdependencies between the *cognitions* of the network. For a visualization of this modeling contribution, see Figure 1.

Second, we introduce the NNTuck as a useful operationalization of the sociological theories of *relational schema* (Baldwin, 1992) and *social cognition* (Howard, 1994), theories that each discuss the ways in which we carry and perceive social structures in our minds, and the shortcuts we use to do so. We propose that the NNTuck can be used as a tool for identifying sets of individuals with commonly shared relational schema, doing so in an unsupervised manner without deciding *a priori* *which* schema to examine (as do, e.g., Janicik and Larrick, 2005; Kilduff et al., 2008). The interpretation of the relational schema at play can instead be interpreted after the latent cognitive spaces have been identified.

Third, we propose statistical tests for assessing different types of cognitive structure in a social network. Notably, we propose a definition and test for determining if the latent cognitive structure is mirrored by (or mirrors) the social structure, a property we call *social-cognitive agreement* (SCA). Finally, and most broadly, our work examines the CSS as a three-dimensional tensor, connecting the study of CSSs to useful formalisms from multilinear algebra for studying such rich data objects. We use the tools developed here to uncover insights into the cognitive and social

spaces of empirical networks, showing how the NNTuck can be used to explore how social and cognitive structure can significantly alter across time.

The structure of the work is as follows. In Section 2, we discuss related work. We introduce the nonnegative Tucker decomposition (NNTuck) of multilayer networks (Aguilar, et al., 2022) in Section 3, where we also discuss the sociological concepts that we draw upon and propose the interpretation of the NNTuck in the context of analyzing CSSs. In Section 3.3, we define and discuss the notion of *social-cognitive agreement* useful for analyzing CSS datasets, and in Section 3.5, we present for statistical tests for assessing latent cognitive structure in social networks. In Section 4, we investigate the use of the NNTuck to find latent structure in four CSS datasets from Krackhardt (1987) and Hunter (2019). We conclude with a discussion and promising open questions for future work.

2. Related work

This work is by no means the first which suggests that there is a latent, or shared, cognitive structure present in the perceptions of social networks. Even preceding Krackhardt's original CSS work, Carley (1986) showed how the shared "frames" developed in the social networks of the members of a college dormitory impacted shared cognition. In Krackhardt (1987), the author motivates his proposal of CSSs by discussing schemas, invoking research that finds people rely on patterns for their recollections (Freeman et al., 1987), and suggesting that Heider's (1958) balance theory is a possible cognitive model underlying peoples' perceptions. Later, DiMaggio's (1997) work on *Culture and Cognition* cites Krackhardt's (1987) work, stating that "networks are crucial environments for the activation of schemata, logics, and frames."

Although there was a strong emphasis on the theory of shared schema leading up to and within Krackhardt (1987), the bulk of the CSS work that has followed has focused on quantifying the "accuracy" of the perceptions of members in a social network. For example, Krackhardt (1990) found that those with more accurate perceptions of the network are perceived as being more powerful, and Brands and Kilduff (2014) relate the misperceptions of women's roles in a network to the characteristics that their peers attribute to them. For a comprehensive review of CSS work through 2013, see Brands (2013).

Our work connects four separate literatures: (i) modeling latent structure in a CSS, (ii) studying relational schema in network perceptions, (iii) multilayer network extensions of the SBM, and (iv) identifying layer interdependence in a multilayer network. As such, we discuss related work in the literature of each of these four categories.

2.1. Latent space models of CSSs

From a modeling perspective, our work is similar to work which has aimed to identify latent structure in the perceptions of the network by considering the CSS data object in its entirety. In Kumbasar et al. (1994), the authors conduct a correspondence analysis of a $N^2 \times N$ "stacked" matrix representation of a CSS to find and compare the latent self-perceived and group-perceived social spaces of each person. In Sosa and Rodríguez (2021), the authors propose a three-dimensional extension of the latent space model for social networks developed by Hoff et al. (2002). Their model identifies the self-perceived latent position of each individual, and the posterior probabilities of a Bayesian model are used to assess "whether the perception of an individual about its own position in social space agrees with the judgments of other actors." The authors propose that this information can be used to create a weighted (single-layer) network, where the weight of an edge between persons i and j corresponds to how much cognitive agreement i and j share. Relatedly, in Sewell (2019), the author proposes a latent space model which describes the deviations between individuals' perceptions through bias and variance, explaining individuals' tendencies to report more or less dense networks, and their confidence in the network perception,

respectively. More recently, De Bacco et al. (2023) proposed a latent model to uncover an unobserved underlying network from multiply reported network data (such as a CSS) by identifying latent propensities for each reporter to over or under report edges between others, as well as each reporter's tendency to report reciprocated ties between people. These related papers which have focused on identifying latent spaces in a CSS set the precedent that we hope to continue here. In contrast to this previous work, however, the present work centers a latent space model which also describes a generative process of the varying perceptions, in order to identify who in the network shares relational schema.

2.2. Relational schema in network perceptions

From a conceptual perspective, our work is most similar to that of Janicik and Larrick (2005), Carnabuci et al. (2018), Menon and Smith (2014), Kilduff et al. (2008), Emery et al. (2011), Brashears (2013), and Brashears and Quintane (2015), which all aim to identify specific relational schema individuals use when recalling their social networks. In contrast to our proposed method, these prior works all identify *specific* relational schema to study the presence of in the social network, whereas we aim to identify *people* who share a general relational schema that can be interpreted and identified *a posteriori*. Janicik and Larrick (2005) specify two relational schema that influence social network perceptions: balance schema (Heider, 1958) and linear-ordered schema (De Soto and Kuethe, 1959). Their work goes on to study how people use these schema to fill in incomplete parts of a social network. Carnabuci et al. (2018) take an experimental approach to infer the relational schema that participants rely upon by measuring how learning rates change in conditions which assume different relational schema in the network. Similarly, Menon and Smith (2014) study how psychological priming impacts the density and characteristics of people's perceived networks. Kilduff et al. (2008) study multiple CSS datasets to understand how schemas like small-world assumptions (Milgram, 1967) and perceived popularity manifest themselves in CSSs. In Emery et al. (2011), the goal is to identify how specific relational schemas impact the emergence of leadership in a social network. In Brashears (2013), the author conducts experiments to understand how compression heuristics aid in the recall of large social networks, finding that people have better recall when structural assumptions of relational schema like kinship and triadic closure are accurate. In Brashears and Quintane (2015), the same experimental data is reanalyzed and the recall of triads and groups is compared to what is expected from exponential random graphs. We aim to enrich the study of these varied questions about cognitive structure. In contrast to this previous work, the operationalization we introduce identifies individuals with shared relational schema without necessitating *which* relational schema to identify ahead of time and enables us to contrast social structure with cognitive structure.

2.3. Multilayer extensions of the SBM

The model we propose for analyzing CSSs builds upon a diverse set of work on SBMs for multilayer networks. Most similarly to the NNTuck work of Aguiar et al. (2022) is that of Schein et al. (2016) and De Bacco et al. (2017). In De Bacco et al. (2017), a multilayer SBM is proposed by estimating a separate SBM for each layer of the network, keeping the social space fixed across layers. In Schein et al. (2016), the authors propose a Bayesian Poisson Tucker Decomposition (BPTD) as a generalization of the degree-corrected, mixed-membership SBM (dc-mm-SBM) (Ball et al., 2011) to study multilayer networks. The model proposed in De Bacco et al. (2017) is a specific instance of the NNTuck wherein the third factor matrix in the decomposition is constrained to be the identity matrix (see Definition 1 in Section 3 for more details). Whereas Schein et al. (2016) estimate the SBM with an MCMC algorithmic approach, the NNTuck multiplicative updates procedure seeks to maximize the log-likelihood to obtain a point estimate of the same model (see Section 3.1 and Algorithm 1 for more details). See Aguiar et al. (2022) for more discussion on how the NNTuck relates to and generalizes other multilayer SBMs.

2.4. Layer interdependence in multilayer networks

Viewing the CSS more generally, as a multilayer network, identifying structure in the cognitive space (which we propose to do with the NNTuck) is similar to identifying *layer interdependence*, for which there have been a multitude of proposed methods. Beginning in the CSS literature, Krackhardt (1987) suggested differentiating layer similarity by comparing individual layers to a *consensus structure*. Battiston et al. (2014) and Kao and Porter (2018) introduce and use various similarity measures to identify layer communities. De Domenico et al. (2015) and De Domenico and Biamonte (2016) cluster similar layers using information-theoretic tools. In Stanley et al. (2016), layers are categorized into groups that were drawn from the same SBM. In De Bacco et al. (2017), the authors build multiple models using different subsets of the layers and layer interdependence is determined via each models' performance on a link prediction task. In contrast to this previous work, the NNTuck identifies layer interdependence by identifying *layer communities* which have the same underlying generative model. As such, the NNTuck allows for the adjacency tensor to be low rank in the layers as well as the nodes. Because identifying layer interdependence in a CSS with the NNTuck relates the varying perceptions to one another via a shared generative process (as opposed to this other work which uses similarity measures or information-theoretic tools), it is thus a natural operationalization of existing sociological theories.

3. A factor model of CSSs

In this section we propose the use of the nonnegative Tucker decomposition (NNTuck) (Aguiar et al., 2022) for analyzing CSSs. We introduce the NNTuck as a multilayer SBM, discuss the interpretation of the NNTuck in the and relate it to relevant sociological theories.

3.1. The nonnegative Tucker decomposition

The degree-corrected, mixed-membership SBM (dc-mm-SBM) (Ball et al., 2011) is a generative model of a single-layer network that assumes each of the N nodes in the network have soft (“mixed”) membership in each of K different communities, and that an affinity matrix captures the Poisson rates at which nodes in each community form an edge with one another. More specifically, for a network of N nodes represented by an adjacency matrix $A \in \mathbb{Z}_{0+}^{N \times N}$, the dc-mm-SBM assumes that each node i has a nonnegative outgoing membership vector, $\mathbf{u}_i \in \mathbb{R}_+^K$, and a non-negative incoming membership vector, $\mathbf{v}_i \in \mathbb{R}_+^K$. Each membership vector represents the node’s soft assignment to $K \leq N$ groups when forming outgoing and incoming edges, respectively. The $K \times K$ nonnegative affinity matrix \mathbf{G} describes the rates of connections between nodes in each community. Then, for $N \times K$ matrices \mathbf{U} and \mathbf{V} with rows defined by \mathbf{u}_i and \mathbf{v}_i , the NNTuck assumes that $A \sim \text{Poisson}(\mathbf{U}\mathbf{G}\mathbf{V}^\top)$. While the model allows for adjacency matrices A consisting of nonnegative count data, it is common practice to employ the dc-mm-SBM for model-based analyses of binary network data. Note that constraining all the elements of \mathbf{u}_i , \mathbf{v}_i , and \mathbf{G} to be nonnegative allows for these parameters to be interpretable as membership weights and affinities.

The nonnegative Tucker decomposition (NNTuck) as a generative model is a multilayer network extension of the dc-mm-SBM. dc-mm-SBM. It builds on related work on multilayer SBMs (e.g, Schein et al., 2016; De Bacco et al., 2017; Tarrés-Deulofeu et al., 2019) by generalizing the SBM to a multilayer setting and using tensor decompositions to jointly model the layers, much like how matrix decompositions underlie traditional (single-layer) SBMs.

Consider a multilayer network with N nodes and L layers represented by adjacency tensor $\mathbf{A} \in \mathbb{Z}_{0+}^{N \times N \times L}$. The NNTuck multilayer SBM again assumes that each node i has nonnegative membership vectors $\mathbf{u}_i \in \mathbb{R}_+^K$ and $\mathbf{v}_i \in \mathbb{R}_+^K$, representing the node’s soft assignment to $K \leq N$ groups when forming outgoing and incoming edges, respectively (note that for an undirected network, $\mathbf{u}_i = \mathbf{v}_i$). Furthermore, the NNTuck assumes that each layer ℓ has a nonnegative vector $\mathbf{y}_\ell \in \mathbb{R}_+^C$ describing the layer’s soft membership to each of $C \leq L$ communities. Just as matrices

U and V in the dc-mm-SBM describe latent community structure in the *nodes* of single-layer networks, the factor matrix Y in the NNTuck describes latent structure in the *layers* of a multilayer network. Finally, the NNTuck assumes a nonnegative core tensor $\mathcal{G} \in \mathbb{R}_+^{K \times K \times C}$ whose frontal slices are C different affinity matrices. Let $\mathbf{u}_i, \mathbf{v}_i, \mathbf{y}_\ell$ be the rows of nonnegative matrices U, V , and Y , respectively. Then the NNTuck multilayer SBM assumes that

$$\mathcal{A} \sim \text{Poisson}(\mathcal{G} \times_1 U \times_2 V \times_3 Y). \quad (1)$$

When estimating this model from multilayer network data, maximizing the log-likelihood of observing \mathcal{A} under the model given by Eq. (1) is equivalent to minimizing the KL-divergence between \mathcal{A} and $\hat{\mathcal{A}} = \mathcal{G} \times_1 U \times_2 V \times_3 Y$. This equivalence is essential to why the NNTuck estimated by minimizing the KL-divergence is synonymous with the multilayer SBM.

3.2. Interpretation of the NNTuck of a CSS

Through its multiple factor matrices, the NNTuck multilayer SBM models both the latent social structure—sometimes referred to as the Blau space (Blau, 1977; McPherson, 1983; McPherson and Ranger-Moore, 1991)—and the latent layer structure of a multilayer network. When viewing a CSS as a multilayer network, each layer corresponds to a different individual’s perception of the social network. Therefore, the latent layer structure captured by the NNTuck of a CSS is precisely the structure identified among the varied perceptions of the network, what we hereafter refer to as the *latent cognitive structure* of the CSS. In the setting of CSSs, then, the NNTuck identifies three latent spaces in a network: the outgoing and incoming social groups in the U and V factor matrices, respectively, and the Y factor matrix, which identifies C *cognitive groups* and each individual’s membership to each of these groups.

Relationship to social cognition and relational schema. Aiming to interpret this factor model through sociological theories, we will not review all relevant or adjacent sociological literature surrounding cognition, social representations, schema, or culture’s various impact on the three. While tempting, it is far beyond the scope of the present work. Instead, we focus on connecting the NNTuck to Baldwin’s (1992) theory of *relational schema* and Howard’s (1994) theory of *social cognition*.

A succinct summary of Baldwin’s theory of *relational schema* (emphasis our own) is that,

“[*relational schema*] describe **expectations about the nature of relationships between people**, defining what aspects of social interactions individuals will pay attention to, as well as **the attributes of others that are meaningful** within those interactions” (Brands, 2013).

Relational schema describe the generalizations or patterns we rely on when making assumptions about relationships or when making new social connections. In data collection for CSSs, each person is asked to report on their perceptions of the relationships between every pair of people in their network. It has been empirically shown that people rely on compression heuristics when recalling large networks (Brashears, 2013), and it is reasonable to assume that the perceptions reported in CSSs are the product of some such compression heuristics. In Section 4.3, we use the NNTuck to understand the cognitive latent space of a longitudinal CSS collected repeatedly over several weeks, allowing us to investigate the heuristics involved (and not involved) in that early formation of the CSS for that population.

About *social cognition*, Howard writes,

“Social cognition articulates explicitly how social structures are carried in individuals’ mental systems . . . The content of the group dimensions on which people are prone to categorize reveals the connections between cognitive and social structures.”

CSS datasets provide empirical access to “how social structures are carried in individuals’ mental systems,” and we propose that the NNTuck, through the Y factor matrix, can identify specific

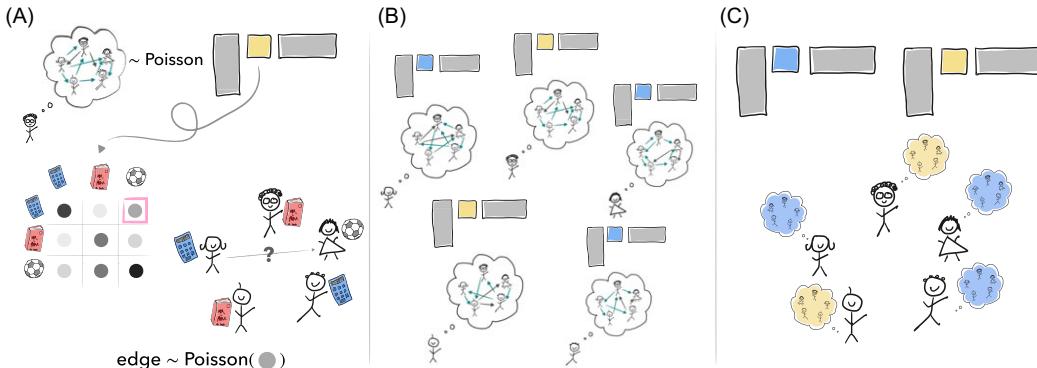


Figure 2. A visualization and example showing the connection between the stochastic block model (SBM), the nonnegative Tucker decomposition (NNTuck), and relational schema. (A) We assume that each person generates their perception of the network according to a stochastic block model (SBM). A person's perception of the existence of an edge is drawn according to the rate specified by the affinity matrix in each person's SBM. Previous empirical and theoretical work on how people store and recall large social networks suggests that a coarse model of relationship, like the SBM, well represents this cognitive process. (B) We *could* estimate a separate affinity matrix to describe each person's network perception. (C) However, the NNTuck allows us to identify when people share the same generative process for their perceptions, interpretable as sharing the same relational schema.

ways in which any differences in these perceptions can be attributed to differences in such mental systems.

Because the NNTuck is an extension of the SBM to multilayer networks, it serves as a *generative model* of multilayer networks, where the role of the Y factor matrix is to identify layers which can be modeled using the same generative SBM. In the context of the CSS, the Y factor matrix identifies individuals in the network who rely on the same generative process when recalling their perception of the network. As such, we propose that individuals in the same cognitive group can be thought of as having a shared set of relational schema. In this sense, too, the NNTuck provides an explicit operationalization of Howard's theory of social cognition, because it jointly models the connection between the cognitive and social structure in the network. In Section 4, we interpret possible relational schema underlying the shared cognitive spaces in three empirical datasets, and in Section 4.3, we directly tie our analysis using the NNTuck to Howard's theories. For a visualization of the connection between the NNTuck and relational schema, see Figure 2.

To formalize our interpretation of the NNTuck in terms of established social theories, we now turn to possible different assumptions of the structure and dimension of the "layer" factor matrix $Y \in \mathbb{R}^{L \times C}$ in the NNTuck of a CSS. The following definitions apply specifications for multilayer network models from Aguiar, et al. (2022) to the CSS context. They form the basis of our primary exploration of empirical CSS data in Section 4.

Definition 1 (Cognitively independent NNTuck). A *cognitively independent NNTuck* is a nonnegative Tucker decomposition where $C = L$ and Y has the constraint $Y = I$, meaning each individual in the network relies on a distinct relational schema.

Definition 2 (Cognitively dependent NNTuck). A *cognitively dependent NNTuck* is a nonnegative Tucker decomposition where Y has the constraint $C < L$, meaning each individual's social cognition can be described by a mixture of C relational schema.

Definition 3 (Cognitively redundant NNTuck). A *cognitively redundant NNTuck* is a nonnegative Tucker decomposition where $C = 1$ and we constrain Y to be the ones vector, $Y = [1, \dots, 1]^\top$, meaning each individual in the network relies on the exact same relational schema in perceiving their network.

3.3. Social-cognitive agreement

In addition to the above model specifications, we next propose a new specification that is particularly well suited to the CSS context. Considering that the NNTuck of a CSS can tell us both about the latent social and cognitive structures, a natural question to ask is how the two spaces are related to one another. Do our social surroundings directly influence how we conceptualize our social network, does our conceptualization influence how we socialize, and if so, can we empirically assess the ways in which they do?

The hypothesis that our cognitions are influenced by our peers was proposed in Carley (1986), where a main goal of the work was to “relate cognitive structure to social structure at an empirical level.” As Carley writes, “the social and cognitive processes cannot be decoupled.” Indeed, through a combination of quantitative and qualitative data collection, Carley identifies specific tightly knit groups in a community who share similar cognitive mappings of a concept. A similar idea was later echoed in Freeman (1992), where the author writes, “the individuals involved in any particular local community would be expected ultimately to produce very similar mental images of group structure in that community.”

These related hypotheses have a natural analog to a specification of the NNTuck. Namely, the social structure is the same as the cognitive structure if $U = Y$. We refer to this structural assumption as one of *social-cognitive agreement*.

Definition 4 (Social-cognitive agreement NNTuck). *An NNTuck with social-cognitive agreement (SCA) is a nonnegative Tucker decomposition where $C < L$ and we constrain $Y = U$. Individuals in the same social group share the same relational schema.*

We note that constraining Y and U to be the same does not imply that everybody in the same social/cognitive group will have the exact same perceptions. Rather, the model is constrained so that the *structure* we observe in the social space is also reflected in the structure of the cognitive space. The implication of this constraint is that the perceptions of people in the same social/cognitive space are modeled with the same generative model, which will result in *similar* (but not the same) perceived networks.

3.4. Model optimization details

Algorithmically, our primary method for maximizing the log-likelihood of observing a CSS dataset under a given NNTuck model specification is to employ the multiplicative updates algorithm of Kim and Choi (2007), an effective extension of the multiplicative updates algorithm for nonnegative matrix factorization (NMF) given by Lee and Seung (2001). For more details, see Algorithm 1 in Appendix B or Aguiar et al. (2022). The Kim and Choi (2007) algorithm, like the NMF algorithm it builds on, is guaranteed to find a local maxima of the nonconvex log-likelihood, but not necessarily the global maxima. That said, extensive empirical evaluations have shown it to perform very well when employed using best practices for nonconvex optimization (e.g., choosing the best of many random restarts of the optimization routine).

For the cognitively independent, dependent, and redundant NNTuck model assumptions (Definitions 1 to 3), this algorithm transfers easily to the relevant specifications. Briefly, when constraining $Y = I$ or $Y = \mathbf{1}$ as in the cognitively independent and redundant NNTucks, respectively, the algorithm maintains monotonic convergence to a local minimum by simply initializing these constraints and never updating Y after initialization. Similarly, when constraining $U = V$, as in the case of an undirected network, by initializing both $U = V$ and symmetry in the frontal slices of \mathcal{G} , monotonic convergence to a local minimum is maintained. The symmetric structure in the frontal slices of adjacency tensor \mathcal{A} (as is the case in an undirected network) and core tensor \mathcal{G} (which is an interpretable assumption for an undirected network) result in equivalence between the multiplicative update derived for U and the multiplicative update derived

for V . Thus, by simply initializing a symmetric \mathcal{G} and $U = V$, Algorithm 1 maintains this symmetry and equivalence.

In the case of estimating a SCA NNTuck (Definition 4), however, constraining $Y = U$ is a non-trivial constraint on the multiplicative updates algorithm. The structure in both the data \mathcal{A} and core tensor \mathcal{G} that would be necessary in order to ensure equivalence in the multiplicative updates of Y and U (as we noted for the symmetric network specification above), are contextually unreasonable to expect or assume. Specifically, equivalence in the multiplicative updates for Y and U requires that the 1-unfolding and the 3-unfolding¹ of \mathcal{A} are equivalent ($A_{(1)} = A_{(3)}$) and that the 1-unfolding and the 3-unfolding of \mathcal{G} are equivalent ($G_{(1)} = G_{(3)}$). In the context of the CSS, this structural assumption on \mathcal{A} means that, for every pair of people i and j , person i 's perception of j 's outgoing relationships is the same as person j 's perception of person i 's outgoing relationships. In order for this assumption to hold, each person would need to report their *own* outgoing relationships as their perceptions for everyone else's outgoing relationships (that they assume everybody goes to the exact same people for advice/friendship).

Thus, in order to estimate a SCA NNTuck, seemingly minor changes to the Kim and Choi (2007) algorithm nontrivially change its behavior. In considering other algorithmic approaches, Cambre et al. (1999) (and more recently Jin et al. (2022)) develop and discuss a method for estimating a *symmetric* Tucker decomposition (with no nonnegativity constraint) wherein all three factors (U, V, Y in our vocabulary), are constrained to be equivalent. However, in this and other related work, the focus is on developing algorithms for decomposing a data tensor which already has an appropriately symmetric structure. Such data symmetry is something that, again, is not reasonable to assume in the case of CSSs.

As such, we move forward with altering Algorithm 1 to accommodate the SCA constraint, resigning ourselves to a estimation algorithm without guaranteed monotonic convergence to a local optima. That said, this was only ever a weak guarantee relative to guarantees of reaching a global optima, as enjoyed by many optimization procedures (e.g., for convex log-likelihoods).

We alter the multiplicative updates from Kim and Choi (2007) in the following ways. First, we initialize nonnegative factors Y , V , and nonnegative core tensor \mathcal{G} . Then, we set $U = Y$. For each multiplicative iteration, we proceed by first updating V , then updating U , and then updating Y , all according to the same multiplicative updates as in Algorithm 1. As our key modification, after updating Y and before updating \mathcal{G} , we set both U and Y to be equal to their average. This step ensures that the model remains “feasible” within the constrained model class, essentially updating both factor matrices to reflect the gradients in both roles (as U and Y factors). We then update \mathcal{G} as usual. As a final minor modification, given the lack of monotonicity guarantee, Algorithm 2 also features a modified termination criteria to ensure the return of the minimum solution found during a given solution trajectory. We describe this approach in full in Algorithm 2 in Appendix B, and provide a python implementation in our code repository (see our Data Availability Statement). In Appendix B, we also provide further details on our decision to choose this modification over others.

This heuristic optimization procedure for returning a maximum likelihood estimate of an NNTuck with SCA is provided as an initial feasible approach to the question of modeling SCA in CSSs. It is an interesting open question if the multiplicative updates algorithm of Kim and Choi (2007), or some other algorithm, can be modified to produce a maximum likelihood estimate with local or global optimality guarantees. For example, Arora et al. (2012) showed that under “approximate separability” (Donoho and Stodden, 2003) assumptions on a data matrix, the NMF has a guaranteed solution and an associated convex optimization problem. Their proposed algorithm overcame prior limitations of heuristics for approximating an NMF by minimizing the Frobenius norm error of the approximation. Establishing identifiability and convergence conditions for the nonnegative Tucker decomposition has been addressed in, e.g., Xu (2015) and Sun and Huang (2023). In Xu (2015), the author proposes an algorithm for estimating the nonnegative Tucker decomposition with global convergence guarantees, but does so for minimizing only

the Frobenius loss. In Sun and Huang (2023), the authors propose a regularization that makes the nonnegative Tucker decomposition identifiable, but the associated loss based on KL-divergence is nonconvex. Extending this work to estimating the nonnegative Tucker decomposition (and the nonnegative Tucker decomposition with SCA) with global convergence guarantees under loss given by the KL-divergence is promising future work.

3.5. Statistical tests for cognitive structure

The above structural assumptions are all articulated in terms of constraints on the parameters of the given models, where the constrained models all lie within a subspace of the more general (cognitively independent) model. As such, these assumptions are highly testable from data using likelihood ratio tests (LRTs). We adapt three statistical tests from Aguiar et al. (2022) to have vocabulary specific to the interpretation of the CSS and introduce an additional statistical test for SCA, all for studying the cognitive structure of a social network through its CSS.

Definition 5 (Cognitive independence). *For a multilayer network let model I be the cognitively independent NNTuck and let model II be the cognitively dependent NNTuck with $C < L$. A CSS has cognitive independence at level α if the likelihood ratio test (LRT) with $(L - C)K^2 - LC$ degrees of freedom is significant at level α .*

Definition 6 (Cognitive dependence). *A CSS has cognitive dependence at level α if the LRT described above is not significant at level α for a pre-specified C .*

Definition 7 (Cognitive redundancy). *A CSS has cognitive redundancy at level α if the LRT comparing the cognitively redundant NNTuck to the $C = 2$ cognitively dependent NNTuck with $K^2 + 2L$ degrees of freedom is not significant at level α .*

Definition 8 (Social-cognitive agreement). *A CSS has social-cognitive agreement at level α if the LRT comparing a cognitively dependent NNTuck with K' and C' , not necessarily equal, to a social-cognitive agreement NNTuck with the constraint $K = C$ and $\mathbf{U} = \mathbf{Y}$ with $2N(K' - K) + K'^2C' + C'N - K^3$ degrees of freedom is not significant at level α . The cognitively dependent NNTuck must be such that both $C' \geq K$ and $K' \geq K$ is true.*

The use of the standard LRT, as in the analogous tests for layer interdependence in Aguiar et al. (2022), relies on Wilks' Theorem (Wilks, 1938) and its assumptions. Wilks' Theorem provides the standard theory for the asymptotic validity of LRTs. In the context of statistically comparing different NNTuck models, we do not satisfy the necessary assumptions or regularity conditions in order for Wilks' Theorem to hold: the nonnegative Tucker decomposition is neither identifiable² (Kolda and Bader, 2009; Chen et al., 2020), nor can we be certain that our estimation of the NNTuck has reached a global maximum of the log-likelihood. Because the NNTuck is not identifiable, the concept of which models are nested within one another is subtle, as one can rewrite the NNTuck by absorbing the parameters of the factor matrices into the core tensor, or vice versa. As such, the notion of nested models and the difference in parameters (and subsequently, determining the degrees of freedom of the χ^2 distribution for the standard LRT) is nuanced, especially in the context of a three-mode tensor decomposition.

As a means of addressing the first shortcoming, we utilize the split-LRT from Wasserman et al. (2020), which does not require identifiability or nested models. Although the split-LRT does not require the same strict regularity conditions as does the standard LRT, it does lead to lower powered tests. Notably, although the power deficit of the split-LRT is somewhat inherent in the test, it may be exacerbated in our particular setting due to our sample size and the presence of many nuisance parameters in our model³ (Tse and Davison, 2022). In Tse and Davison (2022), Strieder and Drton (2022), and Spector et al. (2023), the authors propose ways to improve the power when testing statistical models with nuisance parameters. For instance, both Tse and Davison (2022)

and Strieder and Drton (2022) discuss the importance of an optimal splitting of the data when conducting the split-LRT. In the context of the CSS, wherein we are splitting the data of a 3-mode tensor, we have many choices for how to do so, and our naïve approach of splitting the data at random may not be optimal. Determining the optimal split of a data tensor and implementing other suggestions for increasing the power of the split-LRT in the context of the NNTuck is a topic for future work.

To address the second shortcoming, of reaching a global maximum, we use the NNTuck corresponding to the highest log-likelihood over multiple initializations of Algorithm 1. However, when estimating a SCA NNTuck we use Algorithm 2, which, as discussed in the previous section, is not monotonic in KL-divergence nor does it come with any convergence guarantees, even to a local minima. As such, we caution against strong interpretation of the results of the split-LRT when testing for SCA and again highlight addressing this algorithmic limitation as subject of future work. For more discussion on the standard and split-likelihood ratio tests, see Appendix C.

3.6. Modeling limitations

We conclude this section by discussing some limitations of using the NNTuck for studying CSS datasets. First, the NNTuck models the CSS using a multilayer stochastic block model. As such, the latent cognitive space identifies sets of perceptions of the network can be well modeled by the same (degree-corrected, mixed-membership) SBM. It is in this sense that we say the NNTuck identifies people who share relational schema, because their perceptions can be modeled by same generative process. A significant focus of the active literature on SBMs works to enrich their modeling structure to more realistically model empirical networks, noting that the dc-mm-SBM is still quite coarse (e.g., (Peixoto, 2014)). However, as we discuss above, both theoretical social theory as well as experimental and empirical studies on how people store and recall large social networks suggests that a coarse model of relationship well represents this cognitive process (see also, e.g., Mullainathan et al., 2008). In the context of CSS datasets, wherein each person is asked to report their perception of the social network, we believe that the coarse reasoning expressed in the SBM models this situation well.

Although the NNTuck does not tell us *what* the relational schema are, we believe that the identified spaces can be contextually interpreted (as we do in the empirical applications below). This modeling approach allows for flexibility in identifying people who share relational schema and interpreting that schema *a posteriori*, as opposed to deciding *a priori* which schema to study in a given dataset, which other work on relational schema necessitates (e.g., Kilduff et al., 2008). However, there may be relational schema that may not be well described using a stochastic block model, and latent space models which support these schema, or other generative models of multilayer network formation, would be better suited for modeling such cases. For example, a relational schema based on the assumption of triadic closure would likely be better modeled by a multilayer extension of the recent SBM developments presented in Peixoto (2022).

With these limitations in mind, we turn to four empirical applications of the NNTuck.

4. Empirical cognitive structure

We now use the models and tests introduced in previous sections to analyze four different CSS datasets, with the aim of exploring and interpreting the social and cognitive structures in each. In each of the below datasets, we interpret each person's responses as a layer of a multilayer network, representing the CSS as an $N \times N \times N$ adjacency tensor \mathcal{A} , where each layer is represented as a frontal slice of the tensor.

Krackhardt advice/friendship. In the original CSS work by Krackhardt (1987), $N = 21$ managers in a high tech firm were asked to report on their perceptions of the *advice network* and the

friendship network of the firm, forming two CSS data sets. We also have information about each employee’s affiliation to one of four departments (where the president of the company doesn’t belong to any department), their relative hierarchy to one another (president, vice president, supervisor), their tenure at the firm, and their age. In what we call the “Krackhardt Advice CSS” each person was asked who approached whom for work-related advice, and in the “Krackhardt Friendship CSS” each person was asked who was friends with whom.

Hunter friendship. Here, we analyze longitudinal CSSs from Hunter (2019) which tracks the friendship network perceptions of college juniors from around the country in a summer leadership class over the course of six weeks. In each week, each student was asked to report on their perceptions of the friendships between students in the class. We also have data about each student’s self-reported gender (in the case of this class, only “male” and “female” were reported), race, academic major, and undergraduate institution. We also know which of 10 dorm rooms and 4 study groups each student was assigned to. Although a CSS was collected at each of six weeks,⁴ we focus on the CSS data corresponding to the first and sixth week. Going forward, we refer to each dataset as the “Hunter Friendship, Week One CSS” and “Hunter Friendship, Week Six CSS”, respectively.

We analyze each CSS dataset above through a three-step procedure of (1) model selection, (2) statistical testing, and (3) visualization and interpretation.

For (1) model selection, our main goal is to assess which dimensions of the latent spaces (K and C) should be used for the analysis. To do so, we set up a cross-validation link prediction task. By splitting each dataset into a train and test set, we are able to assess how well each model performs in predicting *unseen* data, which can help account for overfitting that may happen with models that have more parameters. The construction of the cross-validation approach is such that for each link prediction task we construct five different *masking tensors* which each split the data into different *train* and *test* sets. For masking tensor \mathcal{M} , $\mathcal{M}_{ij\ell} = 1$ indicates that the link between i and j in layer ℓ is in the train set. Conversely, $\mathcal{M}_{ij\ell} = 0$ indicates that the link between nodes i and j in layer ℓ is in the test set (is *missing* or *unknown*), and will be held out in the estimation of the NNTuck. In the *tubular* link prediction task, which we use here, masking is done tube-wise (in the tensorial sense), meaning edges are always observed or unknown across all layers. Missing link (i, j) in layer k implies that link (i, j) is missing in all layers ($\mathcal{M}_{ijk} = 0 \Rightarrow \mathcal{M}_{ij\ell} = 0, \forall \ell$). For b -fold cross-validation, tubes (i, j, \cdot) in the tensor are missing with uniform and independent probability $1/b$. We select the NNTuck with the highest training set log-likelihood from 20 runs of the multiplicative updates algorithm with different random initializations. Then, test-AUC is averaged across the five different maskings. This process is repeated for varying model assumptions and dimensions (K, C). Parameter choices for K and C may then be determined by considering the pair with the highest test-AUC or by considering the cost of an increase in test-AUC relative to the complexity of the model (e.g., see our discussion of parameter choice in the Krackhardt Advice CSS).

Next, for step (2), statistical testing, we use the appropriate dimensions K and C determined from the above parameter sweep to perform statistical tests for cognitive redundancy, independence, and SCA using the corresponding LRTs.

Finally, as step (3), we estimate an NNTuck of the dataset using either Algorithm 1 or Algorithm 2 (the latter for models that feature SCA) with the model specification decided by the outcome of the LRT. We then visualize and interpret the social and cognitive spaces identified by these models.

4.1. Krackhardt advice CSS

For this CSS with $N = 21$ managers, the test-AUC under different model specifications is shown in Figure 3 (left). We choose to examine the factor matrices corresponding to the cognitively

Table 1. The determinations from the split-LRT for the Krackhardt and Hunter CSS datasets. See Appendix C for a discussion on the likelihood ratio test (LRT), for motivation for why we use the split-LRT as opposed to the regular LRT, and to compare these determinations to those of the standard LRT. Of note is that the split-LRT suggests that both the Krackhardt Friendship CSS and the last week of the Hunter Friendship CSS are well modeled with social-cognitive agreement (SCA). The other two CSS datasets are better explained when we allow for the social space to differ from the cognitive space

Dataset	Test	Split-LRT determination
Krackhardt Advice	H_0 : Redundant	Reject H_0
	H_1 : $C = 2$	
	H_0 : Dependent $K = C = 3$	Fail to reject H_0
	H_1 : Independent	
Krackhardt Friendship	H_0 : SCA $K = C = 3$	Reject H_0
	H_1 : Dependent $K = C = 3$	
	H_0 : Redundant	Fail to reject H_0
	H_1 : $C = 2$	
Hunter Friendship, Week One	H_0 : Dependent $K = 3, C = 5$	
	H_1 : Independent	Fail to reject H_0
	H_0 : SCA $K = C = 3$	Fail to reject H_0
	H_1 : Dependent $K = 3, C = 5$	
Hunter Friendship Week Six	H_0 : Redundant	Reject H_0
	H_1 : $C = 2$	
	H_0 : Dependent $K = 2, C = 3$	Fail to reject H_0
	H_1 : Independent	
	H_0 : SCA $K = C = 2$	Reject H_0
	H_1 : Dependent $K = 2, C = 3$	
	H_0 : Redundant	Reject H_0
	H_1 : Dependent $C = 2$	
	H_0 : Dependent $K = 2, C = 3$	Fail to reject H_0
	H_1 : Independent	
	H_0 : SCA $K = C = 2$	Fail to reject H_0
	H_1 : Dependent $K = 2, C = 3$	

dependent NNTuck (see Definition 2) with $K = 3$ and $C = 3$. We choose these parameters over others with a higher test-AUC due to the observation that higher values of either K or C result in a minimal increase in test-AUC.⁵ We perform a LRT to compare the goodness-of-fit of the nested, cognitively dependent NNTuck with $K = C = 3$, to the full cognitively independent NNTuck. In doing so, we fail to reject the null hypothesis that the data was generated from the smaller cognitively dependent model (see Table 1).

We inspect the latent spaces identified by this NNTuck in Figure 4. We see that the groupings of the social space of the advice network can be mostly attributed according to departmental affiliation within the firm: the three social groups mostly correspond to the left, bottom, and right, departmental groupings (visualized here in separate clusters), whereas the president (person 6) doesn't strongly belong to any of the three groups. The exception to this is the departmental

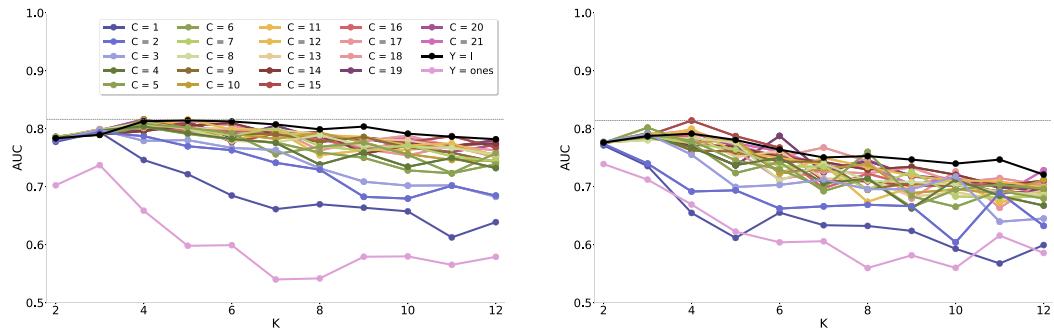


Figure 3. The test-AUC averaged across a tubular fivefold cross-validation task for the Krackhardt advice (left) and friendship (right) CSS datasets. The pink and black lines correspond to the NNTuck model assumptions of cognitive redundancy and independence, respectively. Each other colored line corresponds to a different value of C in assuming cognitive dependence in the CSS, and the x-axis corresponds to different choices of the social latent space parameter K . Based on this cross-validation task, we choose to examine the social and cognitive factor matrices of the advice and friendship CSS datasets corresponding to the cognitively dependent NNTuck with $K = C = 3$, and $K = 3, C = 5$, respectively.

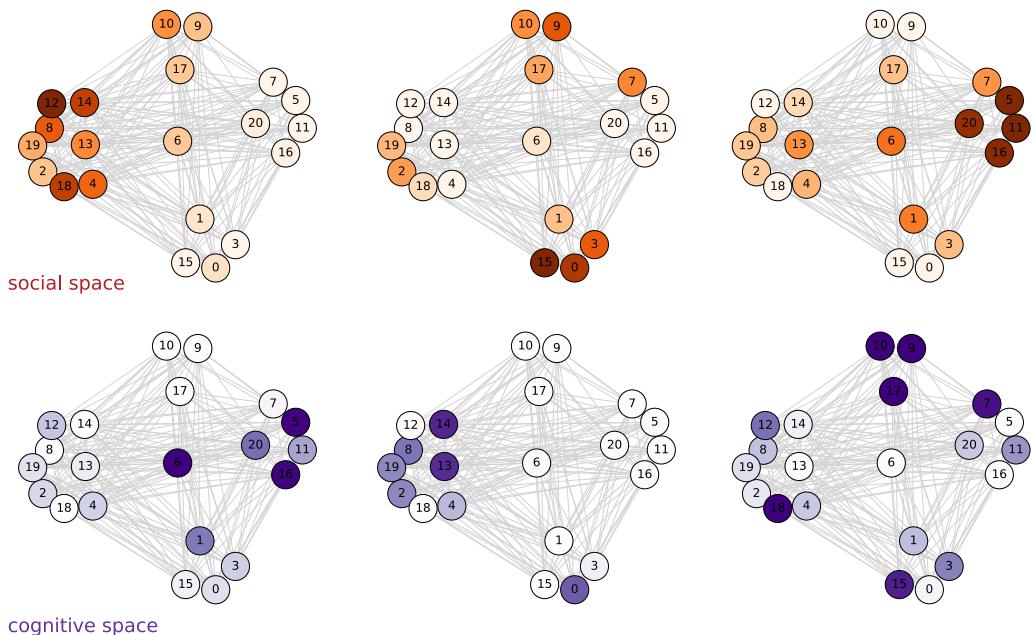


Figure 4. The latent social and cognitive spaces in the high tech firm from Krackhardt (1987), identified by estimating a cognitively dependent NNTuck of the advice CSS with $K = C = 3$. The plotted network is of the network's *consensus structure*, with an edge shown if at least 50% of the network perceived its existence. Each node's position is determined by the departmental affiliation and hierarchy structure of the firm, where the person in the middle is the president, persons 1, 3, 17, and 20 are vice presidents, and the rest are supervisors. Each node is colored according to its proportional membership to each group, where a darker color denotes more proportional membership. We see that persons 5 and 16 belong mostly to the same cognitive space as the president, persons 0 and 13 belong mostly to the same cognitive space as person 14, and everyone else belongs to the third cognitive space.

grouping that we view in the upper center of the network visualization, containing nodes 9, 10, and 17, which doesn't clearly belong to any of the three social groups. The cognitive space, however, seems to group employees according to different attributes.

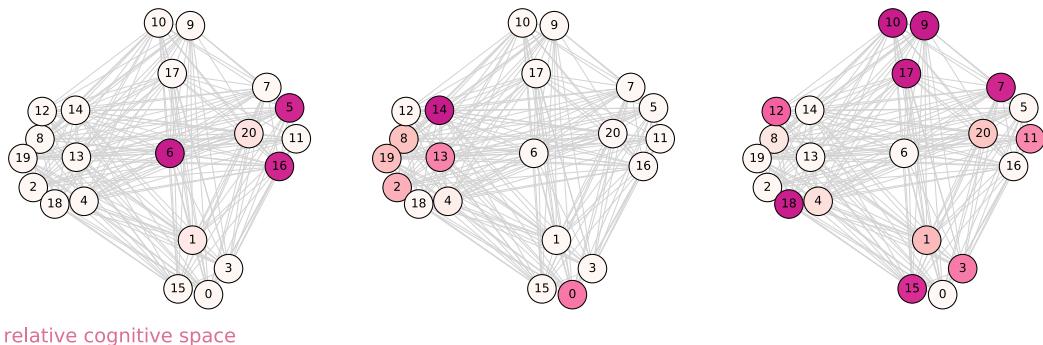


Figure 5. The latent cognitive space of the Krackhardt (1987) advice CSS, rewritten relative to the relational schema of the president of the company, the supervisor we refer to as person 14, and person 10. Each node is colored according to its proportional membership to each cognitive group, where dark pink denotes more membership. Note that, because this plot shows the cognitive membership of each node *relative* to persons 6, 14, and 10, person 6 (the president) has his entire membership in the first cognitive group, and persons 14 and 10 have their entire membership in the second and third cognitive groups.

As is, these cognitive groupings aren't entirely interpretable. To further interpret the identified cognitive space, we can rewrite the Y factor matrix and the core tensor in the basis of $C = 3$ individuals in the network. We choose the three individuals according to the heuristic proposed in Aguiar et al. (2022), where the aim is to choose C layers such that the corresponding rows of Y are (nearly) linearly independent. Doing so, we identify person 6, person 14, and person 10 as the reference perspectives to consider. We transform Y such that the $C = 3$ frontal slices of \mathcal{G} correspond exactly to the affinity matrix from which persons 6, 14, and 10 generate their perceptions, respectively. Thus, the 6th, 14th, and 10th rows of the transformed Y^* matrix will be $[1, 0, 0]$, $[0, 1, 0]$, and $[0, 0, 1]$, respectively, and all other rows will represent each individual's relational schema *relative* to these three people. We inspect this transformed Y^* matrix, which identifies the *relative cognitive space*, by plotting each individual's proportional cognitive membership in Figure 5.

Whereas we know that person 6 is the president of the firm, all we nominally know about person 14 is that he is a supervisor of a department, and that person 10 is a long-standing employee. Digging further into the original analysis of this dataset, however, we find that person 14 is identified as having a particularly interesting perception of his advice network (the friendship network is not discussed). Quoting from Krackhardt (1987) (emphasis our own),

“. . . his own perception is that he is very active in the network: advice is sought from him by 12 people; he actively seeks advice from 20 others; and he is on the crossroads of this network as evidence by his 81.15 betweenness score. This self-evaluation is not shared by his coworkers . . . Only three of the 12 indegrees he claims to have are confirmed . . . Also, only nine of his 20 outdegree nominations are confirmed . . . And finally, his betweenness in the LAS just about disappears (betweenness = 0.70). The Consensus Structure reveals that people generally think that no one approaches him for advice, that he goes to only five people for advice, and that he is not in between any other pair of people.”

Thus, the relational schema that the 21 different individuals in this firm rely upon when recalling the advice network can be concisely described by the relational schema of just three distinct people: the president of the firm, this supervisor with an overly optimistic view of his surrounding advice network, and someone who seems to represent the remaining employees. We continue to analyze this network and these interpretations in the next section, as we explore and contrast the social and cognitive spaces of the friendship network.

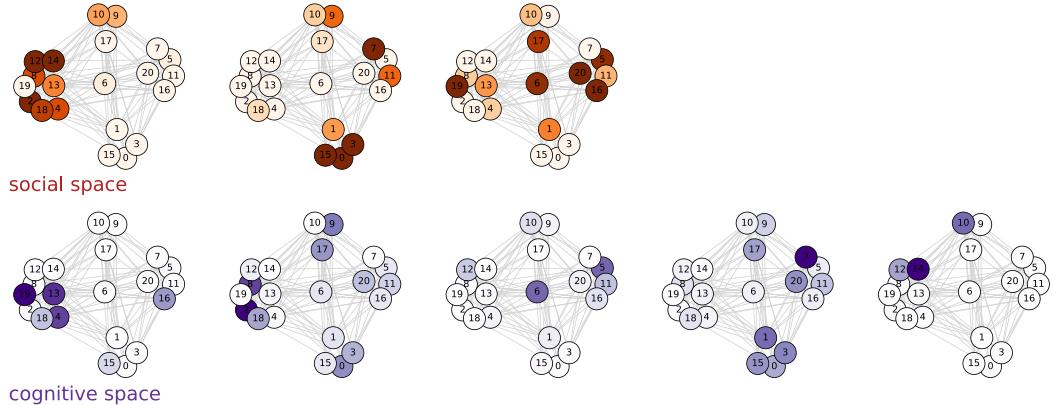


Figure 6. The latent social and cognitive spaces in the friendship CSS of the high tech firm from Krackhardt (1987), identified by estimating a cognitively dependent NNTuck.

4.2. Krackhardt friendship CSS

We turn now to analyze the same firm as above, with the CSS now representing perceptions of the friendship relationships. In modeling this CSS, the test-AUC for different model specifications suggests using $K = 3$, $C = 5$ as preferable to $K = 3$, $C = 3$, what we used to analyze the advice network. The former has a much lower test-AUC (see Figure 3, right). We see that for these values of $K = 3$ and $C = 5$, the standard LRT fails to reject the null hypothesis that the CSS is explained by the simpler cognitively dependent NNTuck.

We inspect the social and cognitive spaces with this model specification in Figure 6. We observe that this decomposition identifies a similar three-dimensional social space to that identified in the advice CSS above. Again, the three social dimensions mostly identify three different department affiliations within the firm. Differently from the advice network, however, we see that the president (person 6) almost entirely belongs to the third social space, whereas in the decomposition of the advice CSS, his social membership was mostly spread across all three groups. Additionally, we see that person 7's social membership does not correspond to his departmental affiliation, but to that of another department. Whereas we lack ethnographic data that might support these friendship alliances, it is curious to wonder how and why these two people belong to different social groups when considering relationships defined by friendship as opposed to advice.

In comparing the identified cognitive spaces in the advice and friendship CSSs for this firm, we see that, again, person 6 and person 5 belong to the same cognitive space. We also observe again that person 14 is identified as belonging to a different cognitive space than most of the others. Although the friendship CSS is not discussed in Krackhardt (1987), it is reasonable to assume that person 14 might have been just as overly optimistic in reporting his perception of the friendship network as he was in the advice network. Why is his relational schema, which sets him apart from many others in the firm, so notably different from that of his colleagues? While there are any number of reasons that could describe this (internal firm politics, his socioeconomic status, his upbringing), it's also possible that he is singled out by the NNTuck because of the noise that he contributed to the CSS by exaggerating connections in his perceived network. In this sense, the NNTuck has the potential to identify a cleaner model of the data, if we were to remove the noisy perceptions.

We see that using the NNTuck to understand the advice and friendship CSSs gives us a much cleaner and richer framework for analyzing these social networks. The original data analysis of this firm was able to identify person 14 as someone who had a noticeably abnormal perception of his social network by comparing his perception of the advice network to an aggregation of his coworkers' perceptions. The NNTuck, however, and the sociological theories that it builds upon

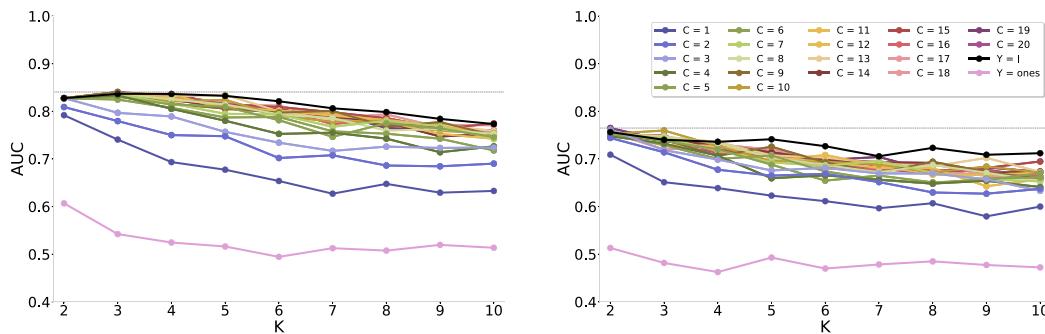


Figure 7. The test-AUC averaged across a tubular fivefold cross-validation task for the Hunter friendship week one (left) and week six (right) CSS datasets. We choose to examine the social and cognitive factor matrices of both CSS datasets corresponding to the cognitively dependent NNTuck with $K = 2$ and $C = 3$.

and operationalizes, gives us the vocabulary to say that he has a different *relational schema* than most of his coworkers. Likewise, for the advice network, we are able to identify two other groups with shared relational schema, as well as visualize the employees who share relational schema across the groups (for instance, person 8 partially shares the relational schema of both person 14 and person 10). Furthermore, we are able to visualize how both the social and cognitive spaces change when considering the advice network in contrast to the friendship network.

A clear limitation of this present analysis is the lack of data about the individuals in this organization: each person is a manager in a high tech firm in the pacific northwest; we don't know much about the internal dynamics of the company; each person is a male aged between 27 and 59. According to Howard (1994), "Gender, race, and age are social systems of differentiation that are especially prone to cognitive categorization." In this present dataset, where each individual has the same gender and race, we are unable to confirm or refute whether the latent cognitive space identifies these relational schema.

4.3. Hunter friendship CSS

We focus the analysis of this rich dataset by contrasting the identified social and cognitive spaces of week one and week six of the longitudinal study.

4.3.1. Week one

When considering the predictive power of various NNTuck models of the week one CSS, we see that the cognitively dependent NNTuck with $K = 2$ and $C = 3$ performs just as well as the cognitively independent NNTuck (see Figure 7, left). However, the LRT comparing these two models fails to reject the smaller, cognitively dependent, model (see Table 1). Considering this alongside the observation that the two models have nearly identical predictive power and similar log-likelihoods, we choose to examine the cognitively dependent model further.

We inspect the identified social and cognitive space, as well as the relative cognitive space, in Figure 8. Notably, we see that during week one of the course, the students' social space can largely be described by their self-identified gender. Similarly, the identified cognitive spaces are well aligned with gender identity, where we see the first group is mostly one gender and the second group is mostly another gender. Interestingly, the third identified cognitive group, which mostly includes students 3 and 6, has students of both gender identities. While we do not have enough data at hand to understand why these students have different relational schema from the others with their same gender identity, our analysis provides an entry point for further analysis.

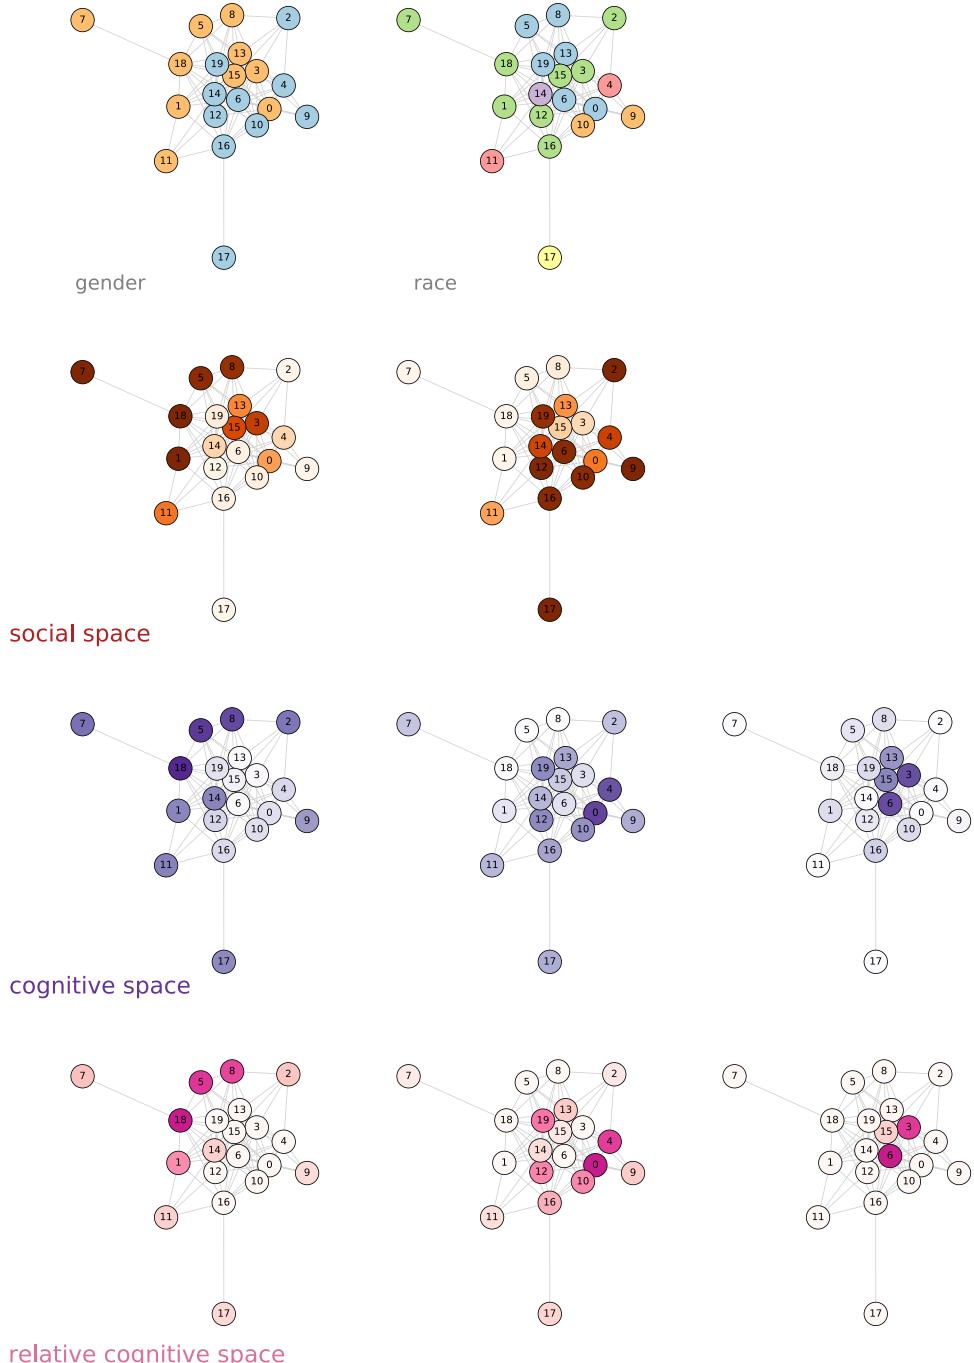


Figure 8. The latent social and cognitive spaces in week one of the Hunter Friendship CSS, identified by estimating a cognitively dependent NNTuck with $K = 2$ and $C = 3$. In the first row, we plot the gender and race identifiers for each of the 20 students and in the last row we plot the cognitive space relative to the relational schema of persons 0, 6, and 18. Note that both the social and cognitive spaces in week one align well with the self-identified gender of each student. The plotted network is of the network's *locally aggregated structure*, with an edge shown from node i to j if node i perceived its existence.

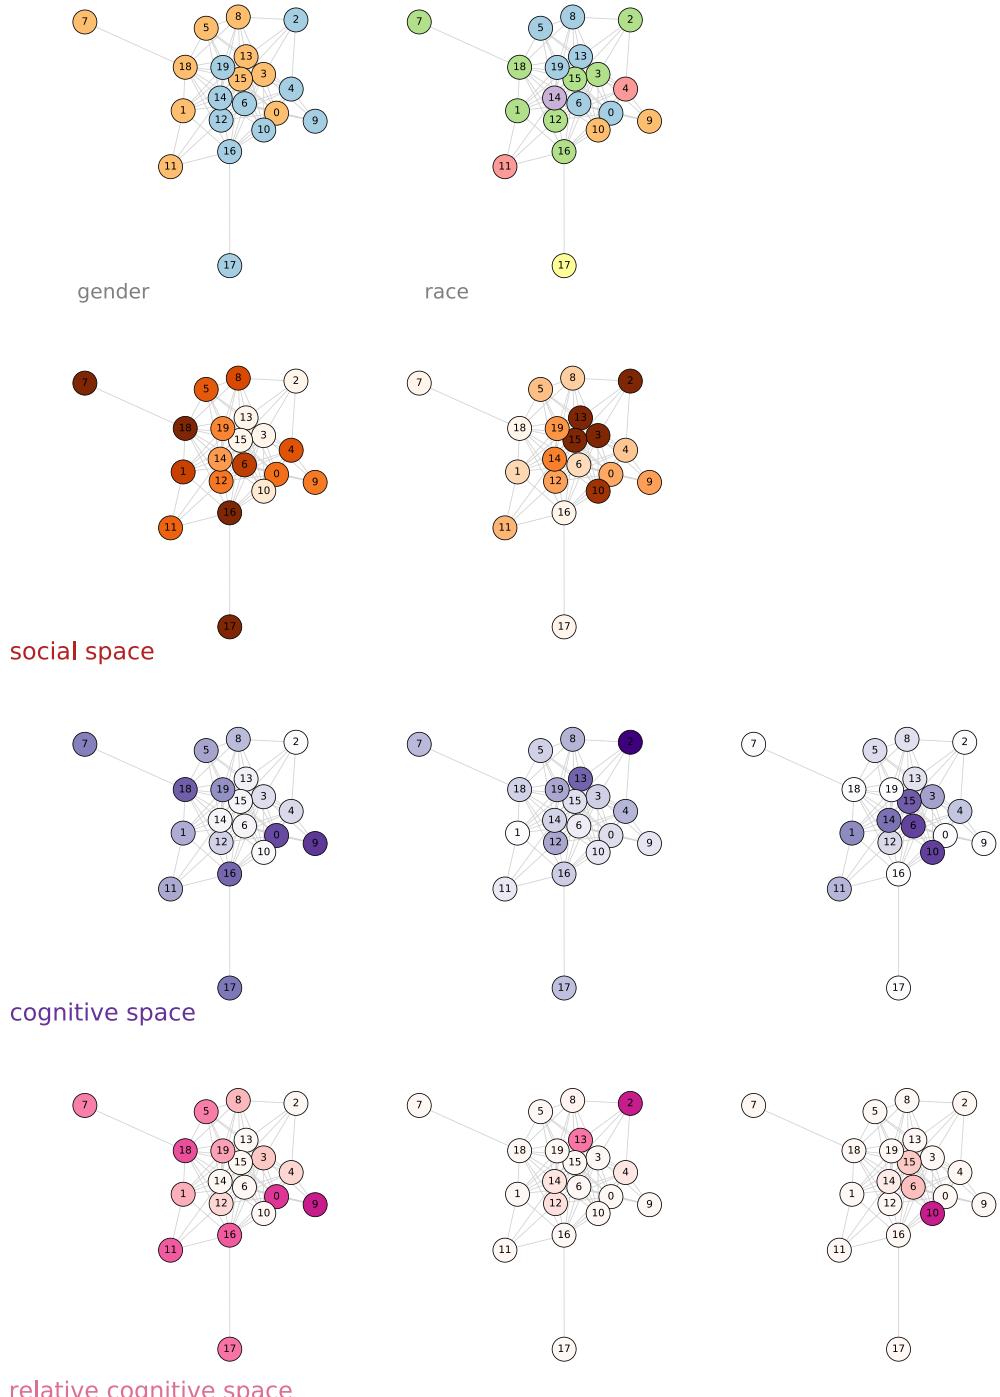


Figure 9. The latent social and cognitive spaces in the last week of the college leadership course friendship network from Hunter (2019), identified by estimating a cognitively dependent NNTuck with $K = 2$ and $C = 3$. In the last row, we also plot the cognitive space relative to students 9, 2, and 10. To visualize the differences in the cognitive spaces of these students, see Appendix D.

4.3.2. Week six

The test-AUC for different model specifications describing the last week of the friendship CSS lead us to inspect the NNTuck with $K = 2$ and $C = 3$ (see Figure 7, right). With this model choice we fail to reject the null hypothesis that the advice CSS was generated from a cognitively dependent NNTuck, and we explore this dependent NNTuck in further detail. Figure 9 shows the identified social and cognitive spaces.

The social spaces that are identified in the NNTuck estimation are not easily explained given the metadata we have (see Figure A1 in A for visualizations of available student attributes, including study groups, undergraduate institutions, and declared majors). Even without any external context about the students in each of the identified social and cognitive spaces in week six, the longitudinal nature of this dataset provides unique opportunities for other rich insights. Notably, we can learn a lot about both the students and the course by comparing these spaces to those identified in week one. We first observe that the identified social and cognitive spaces no longer correspond with the self-reported gender identity of the students. Secondly, we note that students 3 and 6, who shared relational schema in the first week, now have different relational schema, with person 3 mostly belonging to the first cognitive space and person 6 mostly belonging to the last cognitive space. While there are many possible interpretations for the social and cognitive spaces identified in the last week (see Appendix A for one interpretation), the most valuable observations would come from a contextual analysis of these findings done directly by the researchers (see, for example, Carley, 1986, wherein the author compares the findings of a community detection algorithm with her own ethnographic observations of the study participants.). To develop an understanding of the differences between the cognitive spaces identified in the last week of this CSS, we consider two different visualizations in Appendix D which display different ways to understand the differences in the generative processes identified by each cognitive space.

Overall, the above observations have the potential to say a lot about this classroom setting, the instructional materials, and the students enrolled in this course. The course began with a relatively clear divide along gender lines, both in the social environment and the students' relational schema. Again, we quote Howard (1994) in her theory of social cognition wherein she notes that gender is one of the "social systems of differentiation that are especially prone to cognitive categorization." Furthermore, Howard offers additional theory about *why* these initial relational schema occurred along gender lines, writing that "systems of classification must be relatively simple in order to provide a beginning place for interaction. This interactional requirement generalizes the principle of cognitive efficiency: interaction is easier if there are fewer cognitive distinctions to consider . . . the constant use, in interaction, of dichotomies such as gender may help to keep them simple."

Considering the last week of the course, however, after 6 weeks of instructional material, both the social and cognitive spaces became more complicated across identifiers. Whereas the students were prone to this easy heuristic of cognitive categorization—gender—in the first week of the course, we see richer and more complicated relational schema emerge by the end of the study. Again, in Howard's (1994) introduction of the theory of social cognition she writes, "if schemas are to be sustained and reproduced over time . . . they must be validated by the accumulation of resources that their enactment engenders. Schemas and resources constitute structure only when they imply and sustain each other over time." It could be argued that this departure from a gender-informed schema indicates the success of this particular course: that the students' initial schema was not validated and engendered.

5. Conclusion

From the analyses of the CSSs in this work, we conclude that individuals' *perceptions* of their social networks are often just as, if not more, important than some identified "true" network. These perceptions are rich with information, and arguably more valuable when considered all together.

By considering a CSS data object as a tensor, and estimating a factor model of that data tensor, we are able to simultaneously consider each individual's relational schema—how they categorize and compress their perceptions of their social world—and how it relates to others'.

This work highlights many exciting opportunities for future work. From a modeling perspective, it is interesting to connect how other more recent multilayer stochastic block models (for example, Josephs et al., 2023) relate to the NNTuck, and how their modeling interpretations in the context of CSS datasets overlap or differ. Mapping the seemingly disparate multilayer generative models is a fruitful direction for future work on both multilayer networks and their applications.

Above, we proposed the question of empirically studying and testing for SCA in a network, and future work can be dedicated to further studying its sociological (or cognitive) implications. As we briefly discussed, we find evidence that two of the CSS datasets we explore may be well explained by a model with SCA (see Table 1). *When* can we expect to see SCA? *What* does it mean when SCA is or isn't observed in a group? Do only certain types of social relationships reflect an equivalence in social and cognitive space? A domain-informed perspective and study of these questions (perhaps with an accompanying CSS dataset) which uses the tools proposed here (or other appropriate latent space models) to empirically study these specific questions would provide rich insight into the cognitive processes underlying social networks. Furthermore, although SCA is proposed to model and test for when the social and cognitive spaces are the same, there may be applications in which it is interesting to compare the differences between the identified social and the cognitive spaces when they're *not* constrained to be equivalent. In such a setting, one could consider quantifying the difference between the spaces using a cluster comparison measure such as Variation of Information (Meilă, 2003).

At a technical level, as discussed in the main text as well as in Appendix B, there are also several open questions regarding the optimization technique we propose for estimating the SCA NNTuck. As presented, it does not have monotonic convergence guarantees (and none of the multiplicative update algorithms for tensor factorization have any guarantees of global optimality). As such, we cannot confidently assume that the resulting estimated NNTuck satisfies the maximum assumption necessary in both the standard LRT (or split-LRT, as discussed in Appendix C). For future work aiming to identify and/or test SCA in a network, efforts to improve the suggested algorithm and test would be very helpful contributions.

As another direction for future work, we saw in the analysis of the Krackhardt CSS that the estimated cognitive space identified persons 6, 10, and 14 as having notably different relational schema from one another. In Section 4.2, we discuss how person 14 may have been singled out due to his notably dense and optimistic perceptions of the network. Although we focused this work on the introduction of the NNTuck as a tool for studying CSSs, interesting future work could further explore this question of identifying *noise* in CSS datasets, possibly comparing the cognitive space of the NNTuck to the latent methods proposed by Sewell (2019) and De Bacco et al. (2023).

In a separate direction, recent work on information diffusion has focused on how information and behaviors spread differently across different *layers* of multilayer social networks, aiming to identify which types of relationships are most important to identify for purposes of influence maximization (Kempe et al., 2003). This examination of information cascades in a multilayer perspective readily lends itself to interesting questions in the CSS space, namely, *how does information propagate across different network perspectives?* and *does information seeded in different cognitive spaces spread differently?* It's well established that seeding information in different parts of a social space leads to differences in diffusion (Krackhardt, 1996; Banerjee et al., 2013). To go a step further, we contend that perhaps differences in cognitive space—when departing from social space, see our earlier discussion of SCA—are even more significant. When people share a piece of information, they do not have an omniscient view of some ground truth network: they act and share according to their perceptions. Indeed, to quote Thomas and Thomas (1928), just as Krackhardt did in his original work, “if men define situations as real, they are real in their consequences.”

We believe factor-based analyses of CSS, using the NNTuck, can meaningfully help us better understand how cognition impacts information propagation.

Currently, CSS data is collected by asking each person *yes/no* questions about relationships in their surrounding network. Another interesting direction for future work is incorporating something akin to an *unknown* response into the space of possible answers for CSS surveys. Distinguishing between perceived relationships, educational guesses-at relationships, and pure speculation of relationships ought to have interesting impacts on how CSSs can be interpreted, including through properly adjusted NNTuck models, showing how and when relational schema are most important.

Finally, there is an obvious lack of CSS data from networks consisting of more than 30 people. This is, in most part, due to the immense burden (both on the survey administrator and participant) of asking each person in a network of size N to report on $N^2 - N$ distinct relationships. Collecting the entire CSS, however, may be unnecessary depending on the research objective. For example, if the end goal of collecting CSS data is to identify the generative NNTuck and the network's associated latent social and cognitive spaces, partial data may be more than sufficient. Extending the active matrix factorization for surveys work of Zhang et al. (2020) to tensors could potentially allow for efficient CSS data collection of much larger networks.

We conclude by urging that there is much to be learned from continuing to view the CSS as a tensor-object. We hope that the present work opens new research directions toward the aim of uncovering relational schema in social networks, and are eager for future work that continues in this pursuit.

Acknowledgements. We sincerely thank the two anonymous reviewers who provided generous and thoughtful feedback to make this manuscript stronger and more clear. We thank Keith Hunter for making his data available and for his generous availability for discussions. We thank Amir Goldberg and our MURI collaborators for their fruitful discussions and encouragement on this work.

Funding statement. IA acknowledges support from the NSF GRFP and the Knight-Hennessy Scholars Fellowship. JU acknowledges partial support from ARO (#76582-NS-MUR) and NSF (#2143176).

Competing interests. None.

Data availability. The friendship cognitive social structure from Krackhardt (1987) and the employees' attribute data is available in the `statnet` R package (Handcock et al., 2003). A transcription of the advice cognitive social structure from Krackhardt (1987) is available on the Github repository of the present work, as well as that of (Aguiar, et al., 2022). The advice and friendship longitudinal CSS datasets are generously available online (Hunter, 2019). Tools for this work, a notebook reproducing all figures and table values in this manuscript, as well as a tutorial notebook (all in python) can be found at <https://github.com/izabelaguiar/CSS>.

Notes

1 Tensor unfoldings are higher-order analogs of matrix vectorizations. See Kolda and Bader (2009) for explanations of tensor unfoldings and other tensor properties.

2 Recent work by Sun and Huang (2023) has proposed that imposing a “volume regularization” condition on the factor matrices of the nonnegative Tucker decomposition ensures identifiability of the decomposition. Implementing these conditions in contextually suitable ways for CSSs and, more generally, multilayer networks, is promising future work.

3 Although we are concerned with testing how well a particular Y fits the data, we also have to estimate parameters U , V , and G in each model. In this case, U , V , and G can be considered nuisance parameters.

4 The entire longitudinal CSS for each relationship type from Hunter (2019) could be meaningfully analyzed as a fourth-order tensor of size $20 \times 20 \times 20 \times 6$, studying temporal factors through tensor decomposition methods. That said, this data constitutes the only presently known instance of such fourth-order data, and we consider such an analysis to be beyond the scope of the present work.

5 Specifically, considering Figure 3 (left), we see that the NNTuck with $K = 4$ and $C = 20$ has only a slightly higher test-AUC (0.0178 higher) and comes at the cost of an additional 633 parameters.

References

Aguiar, I., Taylor, D., & Ugander, J. (2022). A factor model of multilayer network interdependence, arXiv preprint arXiv: 2206.01804.

Arora, S., Ge, R., Kannan, R., & Moitra, A. (2012). Computing a nonnegative matrix factorization—provably. *Proceedings of the Forty-fourth Annual ACM Symposium on Theory of Computing*, pp. 145–162.

Baldwin, M. W. (1992). Relational schemas and the processing of social information. *Psychological Bulletin*, 112(3), 461–484.

Ball, B., Karrer, B., & Newman, M. E. J. (2011). Efficient and principled method for detecting communities in networks. *Physical Review E*, 84(3), 036103.

Banerjee, A., Chandrasekhar, A. G., Duflo, E., & Jackson, M. O. (2013). The diffusion of microfinance. *Science*, 341, 1236498.

Battiston, F., Nicosia, V., & Latora, V. (2014). Structural measures for multiplex networks. *Physical Review E*, 89(3), 032804.

Bernard, H. R., Killworth, P. D., & Sailer, L. (1979). Informant accuracy in social network data IV: a comparison of clique-level structure in behavioral and cognitive network data. *Social Networks*, 2(3), 191–218.

Blau, P. M. (1977). A macrosociological theory of social structure. *American Journal of Sociology*, 83(1), 26–54.

Brands, R. A. (2013). Cognitive social structures in social network research: a review. *Journal of Organizational Behavior*, 34(S1), S82–S103.

Brands, R. A., & Kilduff, M. (2014). Just like a woman? Effects of gender-biased perceptions of friendship network brokerage on attributions and performance. *Organization Science*, 25(5), 1530–1548.

Brashears, M. E. (2013). Humans use compression heuristics to improve the recall of social networks. *Scientific Reports*, 3(1), 1513.

Brashears, M. E., & Quintane, E. (2015). The microstructures of network recall: how social networks are encoded and represented in human memory. *Social Networks*, 41, 113–126.

Cambre, J., De Lathauwer, L., & De Moor, B. (1999). Best rank (R, R, R) super-symmetric tensor approximation—a continuoustime approach. *Proceedings of the IEEE Signal Processing Workshop on Higher-order Statistics*, pp. 242–246.

Campbell, K. E., & Lee, B. A. (1991). Name generators in surveys of personal networks. *Social Networks*, 13(3), 203–221.

Carley, K. (1986). An approach for relating social structure to cognitive structure. *Journal of Mathematical Sociology*, 12(2), 137–189.

Carnabuci, G., Emery, C., & Brinberg, D. (2018). Emergent leadership structures in informal groups: a dynamic, cognitively informed network model. *Organization Science*, 29(1), 118–133.

Chen, Y., Moustaki, I., & Zhang, H. (2020). A note on likelihood ratio tests for models with latent variables. *Psychometrika*, 1012(4), 996–1012.

De Bacco, B., Contisciani, M., Cardoso-Silva, J., Safdari, H., Borges, G. L., Baptista, D., Sweet, T. Young, J., Koster, J., Ross, C. T., McElreath, R., Redhead, D., & Power, E. A. (2023). Latent network models to account for noisy, multiply reported social network data. *Journal of the Royal Statistical Society Series A: Statistics in Society*, 186(3), 355–375.

De Bacco, C., Power, E. A., Larremore, D. B., & Moore, C. (2017). Community detection, link prediction, and layer interdependence in multilayer networks. *Physical Review E*, 95, 042317.

De Domenico, M., & Biamonte, J. (2016). Spectral entropies as information-theoretic tools for complex network comparison. *Physical Review X*, 6(4), 041062.

De Domenico, M., Nicosia, V., Arenas, A., & Latora, V. (2015). Structural reducibility of multilayer networks. *Nature Communications*, 6, 6864.

De Domenico, M., Solé-Ribalta, A., Cozzo, E., Kivelä, M., Moreno, Y., Porter, M. A., Gómez, S., & Arenas, A. (2013). Mathematical formulation of multilayer networks. *Physical Review X*, 3, 041022.

De Soto, C. B. & Kuethe, J. L. (1959). Subjective probabilities of interpersonal relationships. *The Journal of Abnormal and Social Psychology*, 59, 290–294.

DiMaggio, P. (1997). Culture and cognition. *Annual Review of Sociology*, 23, 263–287.

Donoho, D., & Stodden, V. (2003). When does non-negative matrix factorization give a correct decomposition into parts? *Advances in Neural Information Processing Systems*, 16.

Emery, C., Carnabuci, G., & Brinberg, D. (2011). Relational schemas to investigate the process of leadership emergence. In *Academy of Management Proceedings*, Academy of Management Briarcliff Manor, NY, (pp. 1–6).

Freeman, L. C. (1992). Filling in the blanks: a theory of cognitive categories and the structure of social affiliation. *Social Psychology Quarterly*, 55(2), 118–127.

Freeman, L. C., Romney, A. K., & Freeman, S. C. (1987). Cognitive structure and informant accuracy. *American Anthropologist*, 89, 310–325.

Heider, F. (1958). *The psychology of interpersonal relations*. Hillsdale, NJ: Psychology Press.

Handcock, M. S., Hunter, D. R., Butts, C. T., Goodreau, S. M., & Morris, M. (2003). Statnet: an R package for the statistical modeling of social networks. <http://www.csde.washington.edu/statnet>.

Hoff, P. D., Raftery, A. E., & Handcock, M. S. (2002). Latent space approaches to social network analysis. *Journal of the American Statistical Association*, 97(460), 1090–1098.

Howard, J. A. (1994). A social cognitive conception of social structure. *Social Psychology Quarterly*, 57(3), 210–227.

Hunter, K. O. (2019). Krackhardt data over time: the evolution and emergent properties of network perceptions. Available at SSRN 3408547.

Janicik, G. A., & Lerrick, R. P. (2005). Social network schemas and the learning of incomplete networks. *Journal of Personality and Social Psychology*, 348(2), 348–364.

Jin, R., Kileel, J., Kolda, T. G., & Ward, R. (2022). Scalable symmetric Tucker tensor decomposition, arXiv preprint arXiv: 2204.10824.

Josephs, N., Amini, A. A., Paez, M., & Lin, L. (2023). Nested stochastic block model for simultaneously clustering networks and nodes, arXiv preprint arXiv: 2307.09210.

Kao, T.-C., & Porter, M. A. (2018). Layer communities in multiplex networks. *Journal of Statistical Physics*, 173(3-4), 1286–1302.

Kempe, D., Kleinberg, J., & Tardos, É. (2003). Maximizing the spread of influence through a social network. In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 137–146).

Kilduff, M., Crossland, C., Tsai, W., & Krackhardt, D. (2008). Organizational network perceptions versus reality: a small world after all? *Organizational Behavior and Human Decision Processes*, 107, 15–28.

Killworth, P., & Bernard, H. (1976). Informant accuracy in social network data. *Human Organization*, 35(3), 269–286.

Kim, Y.-D., & Choi, S. (2007). Nonnegative tucker decomposition. In *IEEE CVPR*

Kolda, T. G., Bader, B. W. (2009). Tensor decompositions and applications. *SIAM Review*, 51(3), 455–500.

Krackhardt, D. (1987). Cognitive social structures. *Social Networks*, 9(2), 109–134.

Krackhardt, D. (1990). Assessing the political landscape: structure, cognition, and power in organizations. *Administrative Science Quarterly*, 35(2), 342–369.

Krackhardt, D. (1996). Structural leverage in marketing. In *Networks in marketing* (pp. 50–59).

Kumbasar, E., Romney, A. K., & Batchelder, W. H. (1994). Systematic biases in social perception. *American Journal of Sociology*, 100(2), 477–505.

Lee, D. D., & Seung, H. S. (2001). Algorithms for non-negative matrix factorization. In *Neurips*

McPherson, J. M., & Ranger-Moore, J. R. (1991). Evolution on a dancing landscape: organizations and networks in dynamic Blau space. *Social Forces*, 70(1), 19–42.

McPherson, M. (1983). An ecology of affiliation. *American Sociological Review*, 48(4), 519–532.

Meilă, M. (2003). Comparing clusterings by the variation of information. In *Learning Theory and Kernel Machines: 16th Annual Conference on Learning Theory and 7th Kernel Workshop, COLT/Kernel 2003, August 24–27, 2003, Washington, DC, USA*. Springer. pp. 173–187.

Menon, T., & Smith, E. B. (2014). Identities in flux: cognitive network activation in times of change. *Social Science Research*, 45, 117–130.

Milgram, S. (1967). The small world problem. *Psychology Today*, 2, 60–67.

Mullainathan, S., Schwartzstein, J., & Shleifer, A. (2008). Coarse thinking and persuasion. *The Quarterly Journal of Economics*, 123(2), 577–619.

Newcomb, T. M. (1961). *The acquaintance process*. New York, NY: Holt, Rinehart & Winston.

Peixoto, T. P. (2014). Hierarchical block structures and high-resolution model selection in large networks. *Physical Review*, 4(1), 011047.

Peixoto, T. P. (2022). Disentangling homophily, community structure, and triadic closure in networks. *Physical Review X*, 12(1), 011004.

Schein, A., Zhou, M., Blei, D., & H., Wallach. (2016). Bayesian poisson Tucker decomposition for learning the structure of international relations. In *International Conference on Machine Learning* (pp. 2810–2819).

Sewell, D. K. (2019). Latent space models for network perception data. *Network Science*, 7, 160–179.

Sosa, J., & Rodríguez, A. (2021). A latent space model for cognitive social structures data. *Social Networks*, 65, 85–97.

Spector, A., Candes, E., & Lei, L. (2023). A discussion of Tse and Davidson, A note on universal inference. *Stat*, 12, e570.

Stanley, N., Shai, S., Taylor, D., & Mucha, P. J. (2016). Clustering network layers with the strata multilayer stochastic block model. *IEEE Transactions on Network Science and Engineering*, 3(2), 95–105.

Strieder, D., & Drton, M. (2022). On the choice of the splitting ratio for the split likelihood ratio test. *Electronic Journal of Statistics*, 16(2), 6631–6650.

Sun, Y., & Huang, K. (2023). Volume-regularized nonnegative Tucker decomposition with identifiability guarantees. In *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 1–5).

Tarrés-Deulofeu, M., Godoy-Lorite, A., Guimera, R., & Sales-Pardo, M. (2019). Tensorial and bipartite block models for link prediction in layered networks and temporal networks. *Physical Review E*, 99(3), 032307.

Thomas, W. I., & Thomas, D. S. (1928). *The child in America: behavior problems and programs*. New York, NY: Alfred A. Knopf. p. 1236498.

Tse, T., & Davison, A. C. (2022). A note on universal inference. *Stat*, 11(1), e501.

Wasserman, L., Ramdas, A., & Balakrishnan, S. (2020). Universal inference. In *Proceedings of the National Academy of Sciences* (pp. 16880–16890).

Wasserman, S., & Faust, K. (1994). *Social network analysis: methods and applications*. Cambridge, UK: Cambridge University Press.

Wilks, S. S. (1938). The large-sample distribution of the likelihood ratio for testing composite hypotheses. In *The Annals of Mathematical Statistics* (pp. 60–62).

Xu, Y. (2015). Alternating proximal gradient method for sparse nonnegative Tucker decomposition. *Mathematical Programming Computation*, 7(1), 39–70.

Zhang, C., Taylor, S. J., Cobb, C., & Sekhon, J. (2020). Active matrix factorization for surveys. *The Annals of Applied Statistics*, 14(3), 3362.

A. Hunter CSS metadata

We briefly examine and visualize (see Figure A1) the available metadata of the students and classroom from the longitudinal CSS study from Hunter (2019) which we analyze in Section 4.3. The “group” attribute visualized in Figure A1 describes the study groups the students were assigned to at the beginning of the study. Not only were students within a group expected to do in-class exercises together, but they also participated in a Prisoner’s Dilemma activity together in the middle of the course. At the end of the Prisoner’s Dilemma activity, each group was asked to describe the actions of the other groups in the class. These descriptions are reported in the “affect matrix” in Table A1. For more details, see Hunter (2019). Note that Group A and Group D had, and made, very different impressions of, and on, each group. Specifically, Group A described all groups’ actions with positive words like “honest” or “profitable,” whereas group D described all groups’ actions with negative words like “conflicted” and “sell out.” Similarly, all other groups described Group A’s actions with negative words like “filth” and “self-interested,” whereas all other groups described Group D’s actions with positive words like “loyal” and “altruistic.”

Although more context is needed to make any definitive conclusions, these differences lend more insight into the observations made in Section 4.3. Specifically, in the first week, students 3 and 6 had shared schema that distinctly differed from the other students’, whereas in the sixth week their schema differed (they also belong to separate social groups in both the first and sixth week). Interestingly, student 3 was in Group A and student 6 was in Group D. Again, interesting future work could examine the entire longitudinal CSS dataset as a fourth-order tensor to consider *when* these two students stopped sharing a relational schema, using the context and timing of the Prisoner’s Dilemma activity as additional information.

B. Algorithmic implementation details

Here, we provide more details of the algorithm used to estimate the NNTuck, as discussed in Section 3.4. We also explore the choice of the optimization algorithm discussed in Section 3.3 to maximize the log-likelihood of the social-cognitive agreement (SCA) NNTuck.

B.1 Kim and Choi’s (2007) multiplicative updates

For completeness, we first include the details of the algorithm we use to estimate the NNTuck, as given in Kim and Choi (2007), which we refer to throughout the text as Algorithm 1. Note that in practice we choose the NNTuck corresponding the highest log-likelihood (equivalent to the lowest KL-divergence) over 20 random initializations using Algorithm 1. See the red lines in

Table A1. The “affect matrix” from Hunter (2019). Each group was asked to describe the actions of the other groups following a Prisoner’s Dilemma activity during the course (e.g., group B described group D as “altruistic” and “long sighted.”)

	A	B	C	D
A	Opportunistic, Considerate	Honest, Trusting	Strategic, Profitable	Loyal, Communal
B	Self-Interested, Liars/Trifling	Argumentative, Hop-Scotch Rational	Reactive, Self-Interested	Altruistic, Long Sighted
C	Self-Interested, Bullied	Impressionable, Forgettable	Balanced, Adaptive	Racheted, Idealistic
D	Filth, Judas	Conflicted, Logical	Sell out, Capitalist	Socially Conscious, Resistant to capitalism

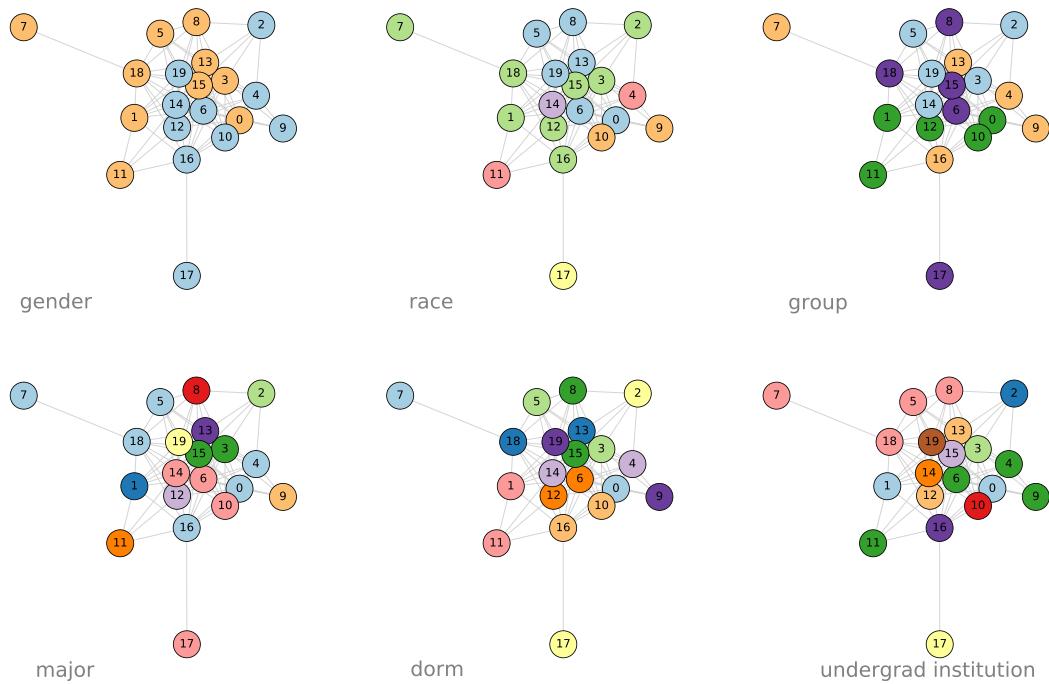


Figure A1. The metadata of the students in the leadership course studied in Hunter (2019).

Figure B1 to see the variation of KL-divergence of this algorithm over 20 random initializations. For further discussion of this algorithm, see Kim and Choi (2007) or Aguiar et al. (2022).

B.2 Comparison of SCA algorithms

Algorithm 2 is only one of many possibilities for how one could alter the algorithm from Kim and Choi (2007) to accommodate the structure of SCA. Here we compare our choice with two other options. For each description, we assume that Y and U are equivalent upon initialization.

Algorithm 1. Multiplicative Updates for minimizing KL-Divergence in the NNTuck (Kim and Choi, 2007)

Input: \mathcal{A}, K, C , Symmetric, Masked, \mathcal{M} , Independent, Redundant

Initialize $U, V \in \mathbb{R}_+^{N \times K}$, $Y \in \mathbb{R}_+^{L \times C}$, and $\mathcal{G} \in \mathbb{R}_+^{K \times K \times C}$ to have random, nonnegative entries.

if Symmetric: $V \leftarrow U$, $\mathbf{G}_\ell \leftarrow \mathbf{G}_\ell^\top \mathbf{G}_\ell$ for $\ell = 1, \dots, C$, and skip each V update step below.

if Independent: $Y \leftarrow I$ and skip each Y update step below.

if Redundant: $Y \leftarrow \text{ones}(C)$ and skip each Y update step below.

Initialize $\hat{\mathcal{A}} = \mathcal{G} \times_1 U \times_2 V \times_3 Y$.

if not Masked: $\mathcal{M} = 1^{N \times N \times L}$

while $\frac{KL(\mathcal{A} || \hat{\mathcal{A}}_t) - KL(\mathcal{A} || \hat{\mathcal{A}}_{t-1})}{KL(\mathcal{A} || \hat{\mathcal{A}}_t)} < \text{rel_tol}$:

$U \leftarrow U \circ \frac{[M_{(1)} \circ A_{(1)} / \hat{A}_{(1)}] [\mathcal{G} \times_2 V \times_3 Y]_{(1)}^\top}{M_{(1)} [\mathcal{G} \times_2 V \times_3 Y]_{(1)}^\top}$

$V \leftarrow V \circ \frac{[M_{(2)} \circ A_{(2)} / \hat{A}_{(2)}] [\mathcal{G} \times_1 U \times_3 Y]_{(2)}^\top}{M_{(2)} [\mathcal{G} \times_1 U \times_3 Y]_{(2)}^\top}$

$Y \leftarrow Y \circ \frac{[M_{(3)} \circ A_{(3)} / \hat{A}_{(3)}] [\mathcal{G} \times_1 U \times_2 V]_{(3)}^\top}{M_{(3)} [\mathcal{G} \times_1 U \times_2 V]_{(3)}^\top}$

$\mathcal{G} \leftarrow \mathcal{G} \circ \frac{[\mathcal{M} \circ \mathcal{A} / \hat{\mathcal{A}}] \times_1 U^\top \times_2 V^\top \times_3 Y^\top}{\mathcal{M} \times_1 U^\top \times_2 V^\top \times_3 Y^\top}$

$\hat{\mathcal{A}} \leftarrow \mathcal{G} \times_1 U \times_2 V \times_3 Y$

Return U, V, Y, \mathcal{G} .

Alternative 1: Update Y and U and then set both to their average (“averaged”). In this approach, for each multiplicative iteration, we proceed by first updating V , then updating U , and then updating Y , all according to the same multiplicative updates as in Algorithm 1. After updating Y and before updating \mathcal{G} , we set both U and Y to be equal to their average. Note that this is the algorithm we discuss in Section 3.3 and which is detailed in Algorithm 2.

Alternative 2: Enforce $Y = U$ after each update (“naive method”). In this approach, we update each factor, beginning with V , according to the update rules in Algorithm 1. After updating U , we set $Y = U$, and after updating Y , we set $U = Y$. We continue to update \mathcal{G} as usual. As in Algorithm 2, we observe empirically that the updates are frequently nonmonotonic, more often so than our chosen algorithm. As such, we keep track of and return the NNTuck corresponding to that with the minimal KL-divergence across iterations. This method has the downside that at each step, only either U or Y is being updated according to the correct gradient descent step, whereas the other factor is constrained to follow the gradient descent step of the other.

Alternative 3: Update Y and U using the average multiplicative update (“averaged updates”). Here, we again begin each iteration with updating V according to the update rules in Algorithm 1. We find the multiplicative update for both U and Y , which we call U_{update} and Y_{update} , respectively, but do not yet update Y or U . We then update both Y and U by multiplying each by the average of these updates, $\frac{1}{2}(U_{\text{update}} + Y_{\text{update}})$. \mathcal{G} is updated as usual. Again, we observe that these updates

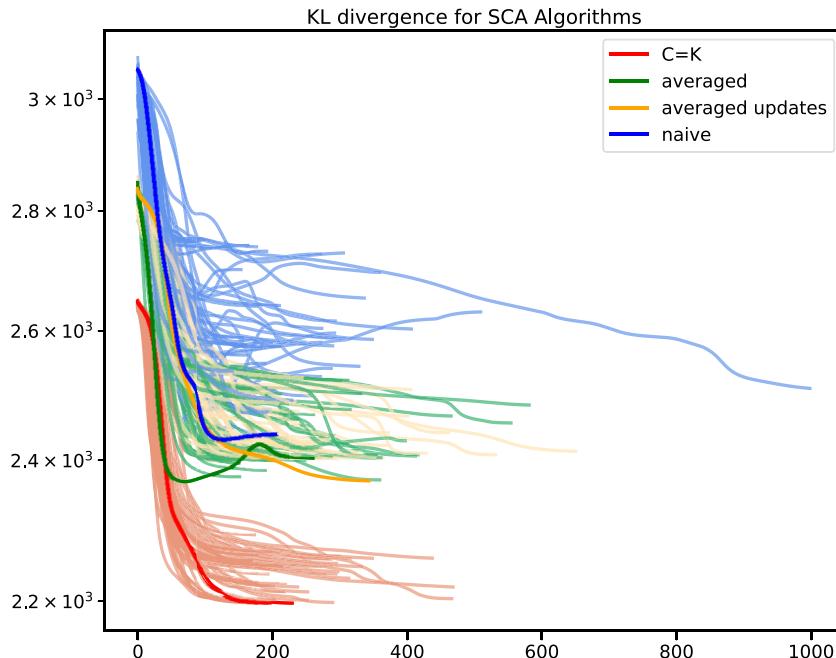


Figure B1. Comparing the KL-divergence across iterations for 50 random initializations in three different SCA algorithms (in blue, yellow, and green), to the convergence of a dependent NNTuck with $K = C$ (red) for Advice CSS with $K = C = 3$. The comparison of these SCA algorithms to the dependent NNTuck is not to compare overall difference in KL-divergence, but rather to show how the SCA algorithms are generally nonmonotonic.

are frequently nonmonotonic in KL-divergence, and we keep track of and return the NNTuck corresponding to that with the minimal KL-divergence across iterations.

Comparisons. We now estimate the SCA NNTuck of the Krackhardt Advice CSS with $K = C = 3$ using the proposed algorithms above. In Figure B1, we compare the above “naive” and “averaged updates” algorithms (in blue and yellow, respectively) with Algorithm 2 (in green). Each line shows the KL-divergence across iterations in each algorithm for 50 random initializations. The NNTuck corresponding to the highest log-likelihood (lowest KL-divergence) across the different initializations is highlighted in bold. We also include the KL-divergence of the unconstrained NNTuck optimized using Algorithm 1 (in red), which exhibits monotonic convergence to a local optima, whereas the other three algorithms do not (necessarily) have monotonic convergence. The performance of Algorithm 1 can be viewed as a lower bound of sorts (but for the fact that it also doesn’t have any global optimality guarantee), since the unconstrained model is strictly more general than the constrained SCA model that the three constrained algorithms are trying to estimate.

We see that Algorithm 2 (green) generally results in a SCA NNTuck with lower KL-divergence than the other two explored algorithms, both in terms of the best-of-50 estimate and the general ensemble of estimates under different initializations.

C. Supplemental discussion of statistical tests

We now expand on the discussion of assumptions underlying statistical tests initiated in Section 3.5. As further background, the likelihood ratio test (LRT) is a statistical test of two hypotheses about the underlying generative model of a given dataset. Broadly, the LRT compares

Algorithm 2. Multiplicative Updates for minimizing KL-Divergence in the social-cognitive agreement NNTuck

Input: \mathcal{A}, K, C , Symmetric, Masked, \mathcal{M} , Independent, Redundant

Initialize $Y, V \in \mathbb{R}_+^{N \times K}$ and $\mathcal{G} \in \mathbb{R}_+^{K \times K \times C}$ to have random, nonnegative entries.

Initialize $U = Y$

Initialize $\hat{\mathcal{A}} = \mathcal{G} \times_1 U \times_2 V \times_3 Y$.

Initialize $\text{min_nntuck} = (U, V, Y, \mathcal{G})$

Initialize $\text{min_KL} = KL(\mathcal{A} || \hat{\mathcal{A}})$

if not Masked: $\mathcal{M} = 1^{N \times N \times L}$

while $\frac{KL(\mathcal{A} || \hat{\mathcal{A}}_t) - KL(\mathcal{A} || \hat{\mathcal{A}}_{t-1})}{KL(\mathcal{A} || \hat{\mathcal{A}}_t)} < \text{rel_tol}$:

$$V \leftarrow V \circ \frac{[M_{(2)} \circ A_{(2)} / \hat{A}_{(2)}] [\mathcal{G} \times_1 U \times_3 Y]_{(2)}^\top}{M_{(2)} [\mathcal{G} \times_1 U \times_3 Y]_{(2)}^\top}$$

$$U \leftarrow U \circ \frac{[M_{(1)} \circ A_{(1)} / \hat{A}_{(1)}] [\mathcal{G} \times_2 V \times_3 Y]_{(1)}^\top}{M_{(1)} [\mathcal{G} \times_2 V \times_3 Y]_{(1)}^\top}$$

$$Y \leftarrow Y \circ \frac{[M_{(3)} \circ A_{(3)} / \hat{A}_{(3)}] [\mathcal{G} \times_1 U \times_2 V]_{(3)}^\top}{M_{(3)} [\mathcal{G} \times_1 U \times_2 V]_{(3)}^\top}$$

$$U, Y \leftarrow (U + Y) / 2$$

$$\mathcal{G} \leftarrow \mathcal{G} \circ \frac{[\mathcal{M} \circ \mathcal{A} / \hat{\mathcal{A}}] \times_1 U^\top \times_2 V^\top \times_3 Y^\top}{\mathcal{M} \times_1 U^\top \times_2 V^\top \times_3 Y^\top}$$

$$\hat{\mathcal{A}}_t \leftarrow \mathcal{G} \times_1 U \times_2 V \times_3 Y$$

if $KL(\mathcal{A} || \hat{\mathcal{A}}_t) < \text{min_KL}$:

$$\text{min_nntuck} = (U, V, Y, \mathcal{G})$$

$$\text{min_KL} = KL(\mathcal{A} || \hat{\mathcal{A}})$$

Return min_nntuck .

the goodness-of-fit in terms of the maximum log-likelihood of a *nested model* to that of a *full model*, where the nested model is a specific instance of the full model. The LRT compares the maximum log-likelihood of each model while considering the difference in the number of parameters between each model. The null hypothesis, then, is that both the nested and full model fit the data equally well, and thus the nested model (with less parameters) should be used. The alternate hypothesis is that the full model fits the data significantly better than the nested model, and thus the full model should be used. The standard LRT depends on Wilks' Theorem (Wilks, 1938), which broadly says that under the null hypothesis, as the sample size approaches infinity, the difference between the maximum likelihoods of each model will be χ^2 distributed with degrees of freedom given by the difference in the number of parameters.

As introduced in the main text, the setting in which we want to statistically compare different NNTuck models is not one which satisfies the necessary assumptions or regularity conditions

Table C1. The p-values and LRT determinations for the Krackhardt and Hunter CSS datasets using the standard LRT. See Section 4 to compare these determinations to those using the split-LRT. Of note is that the differences in the two tests are for the layer redundancy and SCA tests in the Krackhardt Friendship CSS, the layer dependency test in the Hunter Week One CSS, and the SCA test in the Hunter Week Six CSS. In all of these tests, the split-LRT failed to reject H_0 , and the standard LRT rejected H_0 . This observation may be due to the split-LRT being lower powered, which we discuss in Appendix C

Dataset	Test	LRT determination
Krackhardt Advice	H_0 : Redundant	Reject H_0
	H_1 : $C = 2$	p -value $<1e-16$
	H_0 : Dependent $K = C = 3$	Fail to reject H_0
	H_1 : Independent	p -value 0.887
	H_0 : SCA $K = C = 3$	Reject H_0
Krackhardt Friendship	H_1 : Dependent $K = C = 3$	p -value $<1e-16$
	H_0 : Redundant	Reject H_0
	H_1 : $C = 2$	p -value $<1e-16$
	H_0 : Dependent $K = 3, C = 5$	Fail to reject H_0
	H_1 : Independent	p -value 0.995
Hunter Friendship, Week 1	H_0 : SCA $K = C = 3$	Reject H_0
	H_1 : Dependent $K = 3, C = 5$	p -value 9.92e-13
	H_0 : Redundant	Reject H_0
	H_1 : $C = 2$	p -value $<1e-16$
	H_0 : Dependent $K = 2, C = 3$	Reject H_0
Hunter Friendship Week 6	H_1 : Independent	p -value 0.004
	H_0 : SCA $K = C = 2$	Reject H_0
	H_1 : Dependent $K = 2, C = 3$	p -value $<1e-16$
	H_0 : Redundant	Reject H_0
	H_1 : $C = 2$	p -value $<1e-16$
	H_0 : Dependent $K = 2, C = 3$	Fail to reject H_0
	H_1 : Independent	p -value 0.375
	H_0 : SCA $K = C = 2$	Reject H_0
	H_1 : Dependent $K = 2, C = 3$	p -value $<1e-16$

in order for Wilks' Theorem to hold: the nonnegative Tucker decomposition is neither identifiable (Kolda and Bader, 2009; Chen et al., 2020), nor can we be certain that our estimation of the NNTuck has reached a global maximum of the log-likelihood. Thus, we instead utilize split-LRT from Wasserman et al. (2020), which does not require identifiability of the model.

The second shortcoming, of reaching a global maximum log-likelihood, remains troublesome for the cognitively redundant, independent, and dependent LRTs (see Definitions 5 to 7), as well as the SCA LRT (Definition 8). Specifically, even the split-LRT requires that the model of the null hypothesis must correspond to the maximum log-likelihood. When estimating a redundant, dependent, or independent NNTuck, Algorithm 1 has guaranteed monotonic convergence to a *local* minima of the KL-divergence (equivalent to a local maxima of the log-likelihood). We attempt to overcome this local-only guarantee by choosing the NNTuck with the maximal

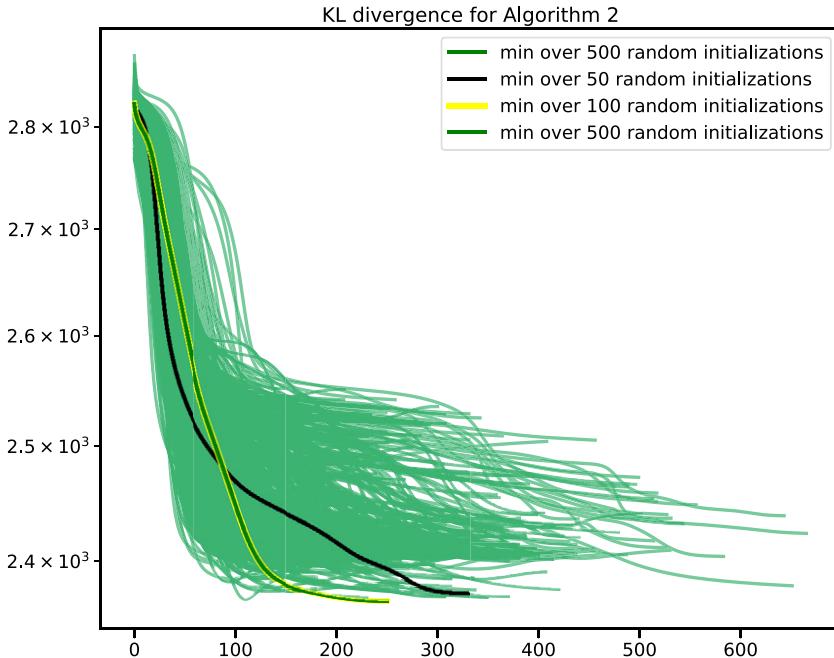


Figure C1. Comparing the KL-divergence of Algorithm 2 across iterations for 500 random initializations. We see that the KL-divergence of the NNTuck with the minimal KL-divergence over 50 random initializations is not much lower than the KL-divergence of the NNTuck with the minimal KL-divergence over 500 random initializations (which is the same as that over the first 100 random initializations). Thus, we estimate the SCA NNTuck using 50 random initializations of Algorithm 2.

log-likelihood over many random initializations. See the red lines in Figure B1 for a comparison of how the minimal KL-divergence varies over different random initializations.

However, when estimating a SCA NNTuck we use Algorithm 2, which is not even monotonic in KL-divergence nor does it come with any convergence guarantees, even to a local minima. Our approach is, as discussed in the main text, to keep track of the SCA NNTuck with the highest log-likelihood over many random initializations. See Figure C1 for a comparison of how the minimal KL-divergence varies over different random initializations when using Algorithm 2.

Although we hope that the heuristic of utilizing multiple random initializations is able to address the problem of only having local minima convergence guarantees (or no convergence guarantees at all, in the case of using Algorithm 2 for the SCA NNTuck), we are unable to provide any assurances or certificates. As such, we urge practitioners to approach the determinations of the SCA LRTs (both standard and split) with caution.

D. Visualizing differences in the cognitive spaces

In this section we highlight two possible visualizations for understanding the differences in the generative processes of each cognitive space in the last week of the Hunter Friendship CSS (see Section 4.3). The difference between each of the C cognitive spaces is the affinity matrix describing the SBM. As such, we first visualize the three different affinity matrices corresponding to the basis identified in considering the relative cognitive spaces. In Figure D1, (top) we visualize the affinity matrices corresponding to the perceptions of students 9, 2, and 10, respectively, each representing the basis of the three cognitive spaces. Interpreting the differences between these affinity matrices,

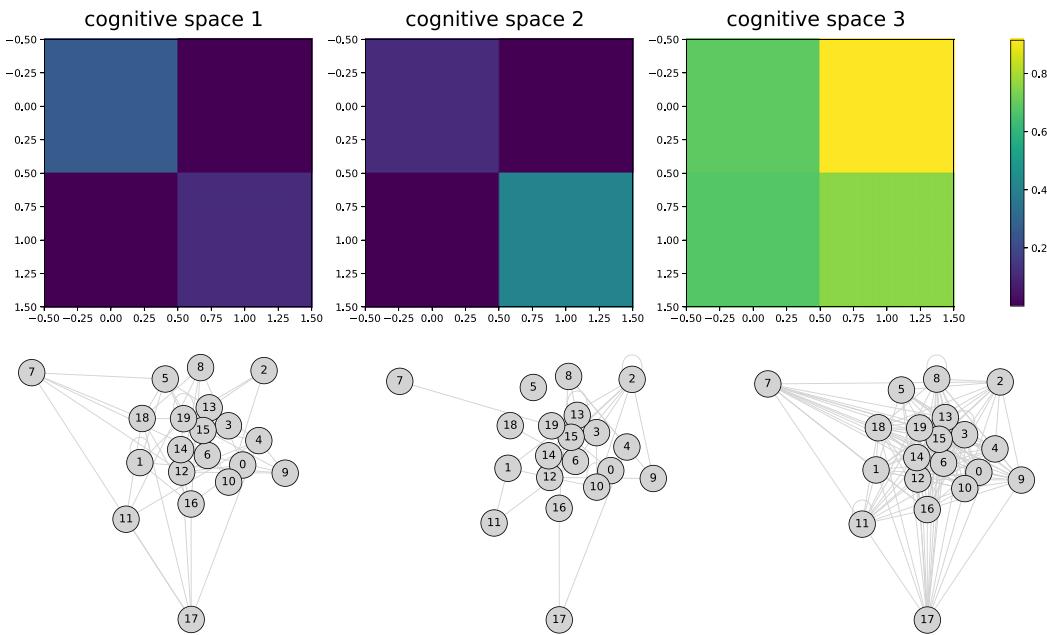


Figure D1. (Top) The affinity matrices corresponding to the three cognitive spaces identified in the last week of the Hunter Friendship CSS. (Bottom) Three synthetic networks showing a random realization of a network generated using each of the three affinity matrices.

we see that student 9's generative model assumes that the first social group has more relative in-group connections than the second social group, which is opposite of what student 2's model assumes. Student 10, conversely, assumes that the most connections occur between the first and second social group, and that all connections occur at a higher rate. In Figure D1, (bottom) we visualize one random initialization of a network generated from a stochastic block model with the same U and V matrices, but with affinity matrices corresponding to those in the top of the figure. Although it is difficult to interpret differences between random realizations of these models (because the randomness in an edge being present or not), we are clearly able to visualize the relatively more dense connections in Student 10's generative process.