CARMA 2026

8th Int. Conf. on Advanced Research Methods and Analytics
1-3 July 2026, Valencia, Spain
DOI: https://doi.org/10.4995/CARMA2026.2026.21976

Do Semantic Networks Encode Latent Traits? A Bayesian Semi-Supervised IRT Approach

Alex Cucco[1], Lara Fontanella[1], Sara Fontanella [2], Luigi Ippoliti [3], Pasquale Valentini [3]

[1]Department of Socio-Economic, Managerial, and Statistical Studies, G. d'Annunzio University Chieti-Pescara, Italy ,[2]National Hearth and Lung Institute, Imperial College London, London, United Kingdom, [3]Department of Economics, G. d'Annunzio University Chieti-Pescara, Italy.

How to cite: Cucco, A.; Fontanella, L.; Fontanella, S.; Ippoliti, L.; Valentini, P. (2026). Do Semantic Networks Encode Latent Traits? A Bayesian Semi-Supervised IRT Approach. In: 8th International Conference on Advanced Research Methods and Analytics (CARMA 2026). Valencia, 1-3 July 2026. https://doi.org/10.4995/CARMA2026.2026.21976

Abstract

This study investigates whether semantic networks encode psychometrically derived latent trait information. Data comprised responses to a composite questionnaire assessing stereotypes toward migrants from both university students and Facebook users. The questionnaire combined the Stereotype Attitude Scale (6-point Likert items) with an open-ended question eliciting respondents' general opinions about migrants. Text responses were embedded in a vector space using multilingual transformers, and semantic networks were constructed via cosine similarity. IRT-derived traits show positive autocorrelation over the semantic network, suggesting that network structure relates to trait information. Building on this finding, we developed a semi-supervised Bayesian framework incorporating graph-structured priors to estimate traits for individuals lacking scale observations. The framework enables trait inference from text while propagating uncertainty. Preliminary results suggest that computational linguistic methods can complement traditional psychometric approaches for measuring latent constructs from natural language data.

Keywords: *Bayesian IRT; Semantic Network; Gaussian Process.*

1. Introduction

Text data represent a naturally occurring and heterogeneous source of information. Open-ended survey responses, social media posts, clinical narratives, and naturalistic discourse capture how individuals freely express their opinions, personal positions, and attitudes in their own linguistic terms, reflecting the complexity and heterogeneity of individual conceptualizations. In contrast, traditional psychometric approaches operationalize psychological constructs through standardized scales with predefined items. While such instruments provide methodological

Article in Press

This work is licensed under a Creative Commons License CC BY-NC-SA 4.0
Editorial Universitat Politècnica de València

rigor, they inherently constrain the expressive range of respondents and may not fully capture the heterogeneity of individual perspectives.

Textual information is often readily available at scale, whereas administering comprehensive psychometric questionnaires is frequently costly, time-consuming, or logistically infeasible. This raises an important methodological question: can we extract latent trait information from naturally occurring text when direct psychometric measurement is unavailable? Semantic networks, graph-based formalizations of conceptual associations derived from textual data, offer a promising conceptualization for addressing this question. The theoretical foundation for the present investigation rests on the premise that latent traits shape cognitive organization and communicative patterns. Under this premise, individuals with similar trait levels should produce semantically related language, resulting in network proximities; conversely, the topological structure of semantic networks should reflect the underlying trait structure of the population. Despite the plausibility of this mechanism, empirical validation of whether semantic networks encode trait-relevant information remains crucial and context-dependent. To investigate this question empirically, we examine attitudes toward non-EU migration. We collected both structured psychometric assessments and open-ended text responses from participants across university students and Facebook users, enabling empirical evaluation of whether semantic network structure derived from text responses relates to standardized trait estimates. We employ Exploratory Graph Analysis (EGA) (Golino & Epskamp, 2017) to assess dimensionality, evaluate item stability and metric invariance. Item Response Theory (IRT) was also applied to further evaluate metric invariance across groups, ensuring comparable trait estimates. Subsequently, we construct semantic networks from participants' text responses and examine whether IRT-derived trait estimates exhibit autocorrelation within the network topology. Significant autocorrelation would constitute evidence that semantic network structure encodes trait-relevant information. Building on empirical evidence of this relationship, we develop a semi-supervised Bayesian hierarchical model that leverages observed associations between semantic network geometry and latent traits. The framework enables trait estimation for individuals lacking direct scale observations by propagating information through the semantic network structure while maintaining uncertainty quantification and propagation. This approach reconceptualizes semantic networks as a predictive instrument for trait inference, enabling estimation from text data.

Article in Press

2. Material and Methods

2.1. Participants and Design

Data were collected through a cross-sectional survey (N = 1,472) employing a multi-channel recruitment strategy. Two distinct samples were included: university students (N = 851) and Facebook users (N = 621), the latter recruited through targeted advertisements disseminated

across Italy. The student subsample represents a younger, highly educated, and relatively homogeneous population, whereas the Facebook subsample is more heterogeneous with respect to age, and educational background. Neither subsample is representative of the Italian population, and the non-probabilistic sampling design imply inherent limitations, including potential social desirability bias among students and self-selection bias among Facebook users. Nonetheless, the dual-channel recruitment strategy allows for preliminary comparative analyses across distinct social and demographic contexts.

2.2. Instruments

Attitudes were assessed through a composite questionnaire combining structured scales and an open-ended question. Stereotypes toward non-EU migrants were measured using the Stereotype Attitude Scale developed through a comprehensive review of prejudice literature and analysis of social media discourse, with items deliberately balanced between positive and negative stereotypes to capture contemporary attitudes toward immigrant groups. The scale was developed building on previous literature (Anderson 2018, Panno et al. 2023, Varela et al. 2013). Items were rated on a 6-point Likert scale (1 = completely disagree; 6 = completely agree). An open-ended question – *What is your opinion on migrants?* - allowed participants to express their attitudes, capturing heterogenous viewpoints beyond standardized measures. Complete data for both the open question and the Stereotype Attitude Scale were available for 918 observations.

2.3. Exploratory Analysis

We conducted Exploratory Graph Analysis (EGA) using the EGAnet (Golino & Christensen 2026) package in R to identify the dimensional structure of the Stereotype Attitude Scale. Item stability (Christensen & Golino 2021) was evaluated through bootstrap resampling with items flagged for removal if stability coefficients fell below 0.80. Items violating metric invariance criteria were then identified for exclusion. To further support measurement equivalence across groups, we also employed Item Response Theory (IRT) via the *mirt* (Chalmers 2012), package in R. Nested models were compared using Bayesian Information Criterion (BIC): configural (all parameters free), metric (equal factor loadings), scalar (equal loadings and intercepts), and fully constrained (all parameters equal).

2.4. Semantic Network Analysis

To capture the semantic structure underlying discourse on immigration emerging from the responses to the open question, we leveraged on a Semantic Network Analysis (SNA) (Pronello et al. 2024). Raw textual responses were embedded in a 768-dimensional vector space using the multilingual transformer model *sentence-transformers/stsb-xlm-r-multilingual* (Reimers & Gurevych 2019). Semantic similarity between response pairs was quantified using cosine

Article in Press

similarity, computed via the *text2vec* package in R. An undirected semantic network $G = \{V, E\}$ was constructed, where V is a set of nodes representing individual responses and E a set of edges connecting response pairs through cosine similarity.

2.5. Bayesian Framework

Linking Semantic Structure to Latent Traits

We hypothesized that latent stereotype attitude is smoothly distributed across the network, as densely clustered networks indicate internally consistent belief systems. To validate this hypothesis, we computed Moran's I (Arthur 2025) on the semantic network using trait estimates obtained from IRT estimation on the entire sample through a Bayesian approach where observations are considered independent. Moran's I for network quantifies autocorrelation of trait values across the network, with positive values indicating that semantically similar responses (proximal in the network) correspond to similar stereotype attitudes. We then developed a semi-supervised Bayesian model to address a realistic data collection scenario: open-ended text responses are readily available from all participants, but administering the questionnaire is feasible only for a subset. We leverage the semantic network to infer latent trait for individuals lacking scale observations while propagating estimation uncertainty through the network structure. We employ Item Response Theory (IRT) for ordinal responses to model the latent trait. Given $i = 1, \dots, N$ observations and $j = 1, \dots, J$ items, for ordinal item responses with K categories, and assuming a unidimensional latent model, the cumulative probability of the observed response Y is specified as:

$$\Pr\left(Y_{ij} = k \mid \theta_i, \lambda_j, \boldsymbol{\gamma}_j\right) = \Phi\left(\lambda_j\theta_i - \gamma_{jk-1}\right) - \Phi\left(\lambda_j\theta_i - \gamma_{jk}\right) \quad (1)$$

where $\theta_i \in \mathbb{R}$ denotes the latent trait for individual i, λ_j represents the discrimination (loading) parameter for item j, and $\boldsymbol{\gamma}_j$ are item-specific thresholds. In a Bayesian framework, we employ data augmentation: each observed ordinal response Y_{ij} is generated from an underlying continuous latent variable Z_{ij}:

$$Z_{ij} = \lambda_j\theta_i + \epsilon_{ij}, \epsilon_{ij} \sim \mathcal{N}(0,1) \quad (2)$$

and the observed ordinal response is obtained via threshold mechanism:

$$Y_{ij} = k \; if \gamma_{j,k-1} < Z_{ij} \leq \gamma_{j,k} \quad (3)$$

The latent trait scores θ are hypothesized to exhibit smoothness over the semantic network structure. Specifically, we model the latent scores by specifying a Gaussian Process (GP) prior with a covariance function defined through the Matern-graph kernel (Borovitskiy et al. 2021):

Article in Press

$$\left(\frac{2\nu}{\kappa^2}\boldsymbol{I} + \mathbf{L}\right)^{-\nu} \tag{4}$$

where $\boldsymbol{L}$ is the Laplacian operator, ν controls the smoothness of the process and κ the length scale. In the case of a process defined over a graph, the Laplacian discrete operator can be derived in the classical way

$$\mathbf{L} = \mathbf{D} - \mathbf{A} \tag{5}$$

where **A** is the adjacency matrix and **D** is a diagonal matrix storing the degree of each node. Alternatively, **L** can be the normalized Laplacian. The identifiability of the model is guaranteed by fixing the latent score mean to zero and variance to one.

3. Results

3.1. Measurement Model and Metric Invariance

EGA identified an initial three-dimensional structure. After bootstrap stability evaluation, two items with stability coefficients below 0.80 were removed. Subsequently, items violating metric invariance were excluded, yielding a final unidimensional model including 12 items that respect metric invariance (Figure 1).

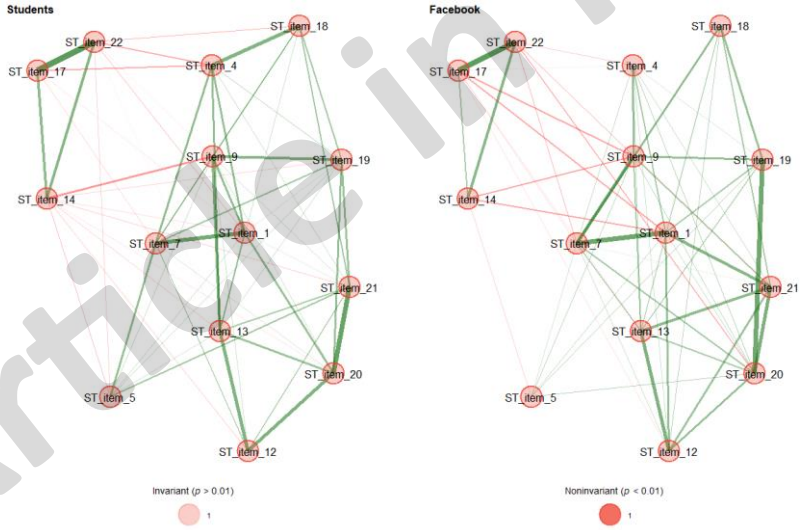


Figure 1. Exploratory Graph Analysis: comparison among students and Facebook users.

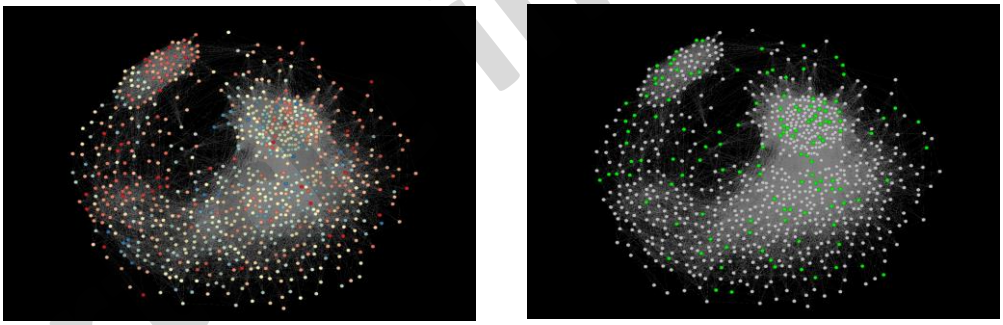
IRT model comparisons further support the metric invariance. The metric invariant model provided superior fit (BIC = 26,778.73) compared to the fully constrained model (BIC = 27,288.98), supporting equivalent factor loadings across groups. The configural model (BIC =

Article in Press

26,783.03) showed only marginally worse fit than the metric model, indicating that metric invariance constraints do not substantially compromise model fit while enhancing generalizability and interpretability across groups.

3.2. Semantic Network Structure

We constructed the network by connecting nodes with cosine similarity>0.6, producing a binary adjacency matrix. We focus on node with at least four connections resulting in 856 vertices. Although a weighted dense network could be used, estimating kernel parameters would be difficult; thus, the threshold-based approach provides a simpler, more interpretable alternative while still capturing the most relevant relationships. Latent posterior mean traits were first derived using a Bayesian model assuming independent observation applied to the entire sample. This was done to first evaluate the autocorrelation across the semantic network of the retrieved traits. Statistical significance of autocorrelation of the latent trait across the semantic network was assessed using two bootstrap tests through data permutation and edge-swapping. Null distributions from both randomization procedures confirmed that observed Moran's I values (I=0.32) fell in the extreme tales, validating the statistical reliability of the autocorrelation findings ($p < 0.001$), supporting the hypothesis that semantic network structure encodes trait-relevant information. Building on this results, we evaluated the performance of the semi-supervised Bayesian model with sematically-related observation by randomly selecting 15% of observations and treat them as missing (Figure 2b).



a) Semantic network colored by posterior mean latent trait values. Blue nodes indicate low stereotype scores (favorable attitudes toward non-EU migrants); red nodes indicate high scores (unfavorable attitudes)

b) Semantic network composition showing data availability for the Bayesian model. Gray nodes represent complete observations with both text and scale responses; green nodes represent the 15% of observations treated as missing (text only, no scale data)

Figure 2. Retrieved semantic Network

Article in Press

Table 1. Estimated Factor Loading with 95% CI from the Bayesian semi-supervised model

Item	Est	95% CI
Non-EU immigrants are invading us.	2.785	[2.525-
Non-EU immigrants are stealing jobs from Italians because they work	1.678	[1.5319-
Non-EU immigrants should be supported at home.	0.993	[0.887-
Non-EU immigrants come to Italy and seek medical care at our expense.	2.687	[2.443-
Immigrants' customs and traditions undermine the cultural identity of the	2.700	[2.435-
Non-EU immigrants who come to Italy are the worst.	2.852	[2.586-
The cultural diversity brought by non-EU immigrants has a positive impact	-1.013	[-1.122--
Immigrants are a valuable workforce that fosters Italian economic growth.	-1.111	[-1.225--
Non-EU immigrants are favored by the government, which welcomes them	1.399	[1.271-
Non-EU immigrants are a security threat.	2.819	[2.539-
Non-EU immigrants commit more serious crimes.	2.700	[2.431-
Non-EU immigrants are an asset to Italian society.	-1.319	[-1.449--

Table 1 reports the posterior mean of the retrieved factor loading for the 12 items leveraging the proposed semi-supevised model togheter with the 95% credible interval. The model's ability to retrieve latent scores was evaluated by comparing the posterior mean of imputed latent scores from the Bayesian semi-supervised approach with those obtained from the fully observed data, as shown in Figure 1. The high correlation of 0.62 between the two sets of retrieved scores across the missing observations demonstrates the model's potential. Notably, the non-perfect correlation may reflect the semi-supervised framework's bias correction mechanisms, implying that some divergence from fully observed estimates represents genuine adjustment for bias rather than pure prediction error. These results provide promising but preliminary evidence that the semi-supervised framework successfully leverages semantic network structure to estimate trait-relevant responses for unobserved individuals.

4. Conclusion

This study demonstrated that semantic networks derived from text encode information about latent psychological traits. Results provided evidence that semantic network structure relates to trait-relevant information. Building on this finding, we developed a semi-supervised Bayesian framework incorporating graph-structured priors and graph Laplacian regularization to estimate traits for individuals lacking scale observations. This approach enables principled propagation of trait information across unobserved individuals while maintaining uncertainty quantification, reconceptualizing semantic networks as predictive instruments for trait inference. Our findings suggest that natural language structure carries information about underlying psychological constructs, supporting the premise that latent traits manifest in measurable network properties. Practically, the framework addresses a methodological challenge: enabling trait estimation from text data alone in settings where comprehensive psychometric assessment is infeasible. This capability has value for large-scale assessment and applied contexts where text data are

Article in Press

abundant but scale administration is costly or logistically constrained. Different network construction approaches may better capture trait-relevant semantic structure. Additionally, investigating multidimensional trait estimation, diverse text modalities (social media, clinical narratives, naturalistic discourse), and domain-specific validation would enhance the framework's applicability and robustness across different populations and contexts.

References

Anderson, J. R. (2018). The prejudice against asylum seekers scale: Presenting the psychometric properties of a new measure of classical and conditional attitudes. *The Journal of Social Psychology, 158(6), 694–710*. DOI: 10.1080/00224545.2017.1404958

Arthur, R. (2025). Correlation and autocorrelation of data on complex networks. *EPJ Data Science, 14(1), 1*. DOI: 10.1140/epjds/s13688-025-00525-1

Borovitskiy, V., Azangulov, I., Terenin, A., Mostowsky, P., Deisenroth, M., & Durrande, N. (2021). Matérn Gaussian processes on graphs. *In International Conference on Artificial Intelligence and Statistics* (pp. 2593–2601). PMLR.

Chalmers, R. P. (2012). mirt: A multidimensional item response theory package for the R environment. *Journal of Statistical Software*, 48(6), 1–29. DOI: 10.18637/jss.v048.i06

Christensen, A. P., & Golino, H. (2021). Estimating the Stability of Psychological Dimensions via Bootstrap Exploratory Graph Analysis: A Monte Carlo Simulation and Tutorial. *Psych*, *3*(3), 479-500. DOI: 10.3390/psych3030032

Golino, H. F., & Epskamp, S. (2017). Exploratory graph analysis: A new approach for estimating the number of dimensions in psychological research. PloS one, 12(6), e0174035. DOI: 10.1371/journal.pone.0174035

Golino, H., & Christensen, A. (2026). EGAnet: Exploratory graph analysis – A framework for estimating the number of dimensions in multivariate data using network psychometrics (Version 2.4.1). DOI: 10.32614/CRAN.package.EGAnet

Panno, A., Pellegrini, V., De Cristofaro, V., & Donati, M. A. (2023). A measure of positive and negative perception of migration: Development and psychometric properties of the Positive and Negative Perception of Immigrants Scale (PANPIS). *Analyses of Social Issues and Public Policy, 23(1), 73–105.* DOI: 10.1111/asap.12338

Pronello, N., Cucco, A., Del Gobbo, E., Fontanella, S., & Fontanella, L. (2024). Dynamics of online debates: insights from textual network analysis. *Annals of Operations Research, 1-21.* DOI: 10.1007/s10479-024-06315-8

Reimers, N., & Gurevych, I. (2019). Sentence-BERT: Sentence embeddings using Siamese BERT-Networks. *In Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing (pp. 3982–3992)*. Association for Computational Linguistics. DOI: 10.48550/arXiv.1908.10084

Varela, J. G., Gonzalez, Jr., E., Clark, J. W., Cramer, R. J., & Crosby, J. W. (2013). Development and preliminary validation of the Negative Attitude Toward Immigrants Scale. *Journal of Latina/o Psychology, 1(3), 155–170.* DOI: 10.1037/a0033707

Author in press