

Analyzing the Differences in Human Reasoning via Joint Nonnegative Matrix Factorization

Daniel Brand* (daniel.brand@cognition.uni-freiburg.de)

Nicolas Riesterer* (riestern@cs.uni-freiburg.de)

Hannah Dames (damesh@cs.uni-freiburg.de)

Marco Ragni (ragni@cs.uni-freiburg.de)

Cognitive Computation Lab, University of Freiburg
Georges-Koehler-Allee 52, 79110 Freiburg, Germany

Abstract

Joint Nonnegative Matrix Factorization (JNMF) is a method for factor analysis that is capable of simultaneously decomposing two datasets into related latent state representations. Enabling factor analysis for contrasting applications, i.e., to find common and distinct structural patterns in data, JNMF has great potential for use in the field of cognitive science. Applied to experimental data, JNMF allows for the extraction of common and distinct patterns of behavior thereby extending the outcomes of traditional correlation-based contrasting methods. In this article, we introduce JNMF to the field of cognitive science and demonstrate its potential on the exemplary domain of syllogistic reasoning by comparing reasoning patterns for different personality factors. Results are interpreted with respect to the theoretical state of the art in syllogistic reasoning research.

Keywords: syllogistic reasoning; personality; nonnegative matrix factorization; data mining; cognitive modeling

Introduction

At the core of scientific progress lies a substantial amount of exploration. By letting oneself be guided by an insight-driven fascination about a field of science, explorative analyses of data lead to potentially unexpected results that could drive future evaluations and progress.

A standard method for explorative research is factor analysis (e.g., Murphy, 2012). Aiming at representing a dataset containing observed variables using a (usually lower) number of latent unobserved variables, so-called *factors*, factor analysis effectively performs a latent state decomposition. As a result, the actual observed variables in the data are modeled as combinations of the latent factors. One such method is *Nonnegative Matrix Factorization* (NMF), which performs a decomposition of data into a nonnegative latent state (Liu et al., 2006). In text mining, for instance, NMF is regularly used to automatically analyze extensive corpora of documents in order to extract representative topics for the texts (Pauca et al., 2004).

Often, the goal of scientific investigation is the contrasting of datasets. However, most standard methods for factor analysis are only able to provide factorizations for a single

dataset making them severely limited with respect to their applicability to contrasting tasks. If applied to two datasets, the most important factors are obtained. However, the main differences between both datasets are not necessarily the main factors. Therefore, there is no guarantee that the factors enabling contrasting of the datasets are found.

To make contrasting problems accessible via factor analysis, in the recent years, *Joint Nonnegative Matrix Factorization* (JNMF), an extension to NMF, has been proposed. JNMF simultaneously performs the factorization for two datasets while constraining the process ensuring that both resulting latent factorizations stay related to each other (Kim et al., 2015). Thereby, contrasting applications can be realized in a standardized framework.

In cognitive science, contrasting groups of individual can help us to identify interindividual differences in cognitive processes. For example, in the field of syllogistic reasoning, one goal is to discover the inferential strategies applied by different reasoners (Khemlani & Johnson-Laird, 2016). Here, approaches for contrasting could be a useful tool to disentangle the behavioral strategies different groups of individuals base their reasoning on. To identify factors that may relate to different reasoning behaviors, a recent study has investigated the connection between personality factors and reasoning performance (Dames et al., in preparation). However, currently the response patterns underlying these differences between groups have not been explored.

We see a large potential for JNMF application in the field of cognitive science. The method does not only perform a superficial evaluation of homogeneity between datasets, but allows for the extraction and subsequent analysis of in-depth patterns of behavior. To demonstrate the general potential of JNMF, in this article, we apply it to the exemplar domain of syllogistic reasoning in an attempt to uncover differences in reasoning patterns between personality factor groups.

Theoretical Background

Syllogistic Reasoning and Personality

In the following analyses, we will apply JNMF to data from the domain of human syllogistic reasoning (e.g., Khemlani

*Both authors contributed equally to this manuscript.

& Johnson-Laird, 2012), a form of deductive reasoning concerned with quantified categorical problems of the form

All A are B
All B are C

What, if anything, follows?

A syllogism consists of two premises which make quantified statements (using the quantifiers “All”, “Some”, “No” and “Some ... not”) for three terms (A, B, and C). The goal in syllogistic reasoning is to infer the quantified relation between the end terms, i.e., the terms appearing exclusively in one of the premises. In total, there are 64 distinct tasks with nine possible responses, eight of which are obtained by combining the end terms (A, C) in either direction with the four possible quantifiers. The ninth response option, “no valid conclusion” (NVC), denotes that no logically valid conclusion can be inferred from the premises.

Actively researched for over 100 years (Störring, 1908), with at least twelve competing cognitive theories (see Khemlani & Johnson-Laird, 2012), the domain is heavily invested in discussing the different observable patterns of human reasoning. The predictive performance of the existing cognitive theories aimed at explaining syllogistic reasoning, however, is worrisome (see Brand et al., 2019) and cannot entirely be attributed to random noise in the data (Riesterer et al., 2019). Such findings indicate that some factors so far neglected by the theories of reasoning may exist and could help to explain variance in the data. Hence, in a recent study we investigated whether in addition to individual differences in cognitive abilities also personality factors may be associated with the variance found in peoples’ reasoning performance (Dames et al., in preparation).

We employed the widely adopted personality model, the Big Five Factor model (McCrae & John, 1992). It distinguishes five domains of personality: extraversion, agreeableness, conscientiousness, neuroticism, and openness to experience. Although we were able to generally conclude that some personality traits are of significance for individuals’ reasoning performance, the statistical analysis we conducted is not able to capture differences in participants’ response patterns. In this study, participants’ reasoning performance was, for instance, negatively related with an individual’s conscientiousness. Other personality traits however (such as openness) led to inconclusive results due to great interindividual differences. As conscientiousness and openness were able to explain a substantial amount of variance in the data, we will exemplarily focus on these personality traits in the following analyses.

Conscientiousness Conscientiousness refers to the tendency to be planful, organized, goal-directed, and to follow rules (McCrae & John, 1992). Conscientiousness is associated with self-discipline, dutifulness, task orientation, and rule abiding. Although conscientiousness has consistently been demonstrated a good predictor for academic performance (Poropat, 2009), it seems to be negatively related

to intelligence (e.g., Carretta & Ree, 2018; Moutafi et al., 2004). In order to explain the observed relation to reasoning performance, we proposed—in line with studies on intelligence—that individuals may develop conscientious characteristics (e.g., plan more, being motivated, etc.) to compensate for a low reasoning ability in order to perform well (see also Rammstedt et al., 2016).

Openness to experience Openness to experience measures the tendency to be original, inventive, and open to new ideas and experiences (McCrae & John, 1992). Individuals scoring high on openness actively pursue novel experiences and engage in cognitively stimulating activities (e.g., reading newspapers or solving puzzles). It is assumed that such activities positively affect cognitive ability (see Ziegler et al., 2012). As a result, various studies consistently found a positive relationship between cognitive ability or intelligence and openness (e.g., Carretta & Ree, 2018; Moutafi et al., 2006; Rammstedt et al., 2016).

To explore how individuals who score either low or high on such trait characteristics respond differently, we adapted the *Batch-Processing Approach*, which is a recently proposed method for JNMF (Kim et al., 2015), to reasoning data. In the following paragraphs, we will briefly introduce the general concept of NMF as well as the specific algorithm we used.

Nonnegative Matrix Factorization (NMF)

Nonnegative matrix factorization is, similar to the well-known *Principal Component Analysis* (PCA; e.g., Murphy, 2012), a method to reduce the dimensionality of input data. NMF was successfully applied to a variety of problems, often outperforming other methods in terms of the interpretability of the latent representations. (Liu et al., 2006). For instance, in the field of text mining and topic modeling, NMF was successfully used to automatically extract patterns representing the relevant topics of the given document sets (Pauca et al., 2004). For data which is naturally represented by nonnegative values, it is advantageous to enforce nonnegativity in the latent state to preserve the meaning of values. This is often the case for behavioral data, especially when involving multiple-choice decisions. As it is unknown if the selected option was just slightly more appealing to the participant or if there was an active decision against certain options, a natural representation for the data would be to assign a positive value to the selected option and a “neutral” value of zero for the remaining options. For a dimensionality reduction, even though it would be mathematically correct, a latent representation including negative values can be undesirable, as it can easily lead to a misinterpretation in terms of biases against specific options (which is not justified given the data). Due to the nonnegative constraints, the latent representation obtained from an NMF preserves the original valence in these cases.

Formally, for a given dataset consisting of n datapoints (e.g., participants in an experiment) with m features (e.g., responses) represented as a matrix $X \in \mathbb{R}_+^{m \times n}$ where \mathbb{R}_+ denotes the set of positive real numbers, NMF is defined as

a search for two matrices $W \in \mathbb{R}_+^{m \times k}$ and $H \in \mathbb{R}_+^{n \times k}$ with the goal of approximating X while imposing a decomposition into a latent state of dimensionality k :

$$X \approx WH^T. \quad (1)$$

By choosing $k \ll \min(n, m)$, the dimensionality reduction is imposed, resulting in latent representations of the input data. A column in the resulting matrix W is a vector defined in terms of the m features, reflecting an iconic pattern of the data (e.g., typical behavioral patterns of participants). In addition, a column in H describes the relationship between the patterns in W and the original datapoints. Currently, several algorithms for computing the NMF exist which optimize the method for specific use cases (Berry et al., 2007; Hoyer, 2004; Lee & Seung, 2001).

Joint Nonnegative Matrix Factorization (JNMF)

There exist extensions for NMF, called Joint Nonnegative Matrix Factorization, that allow to contrast datasets by performing simultaneous, joint decompositions (JNMF; e.g., Kim et al., 2015). For this paper, we adopted an approach originally developed for contrasting topics between two sets of documents, which relies on a simultaneous application of NMF to both datasets while using two regularizing terms to tie the W -matrices of both factorizations together (Kim et al., 2015).

Formally, the matrices are divided into two parts $W_c \in \mathbb{R}_+^{m \times k_c}$ and $W_d \in \mathbb{R}_+^{m \times k_d}$ with $k_c + k_d = k$, where columns of W_c represent latent patterns that represent commonalities, while the columns of W_d represent differences between both datasets.

In order to ensure that columns of both W_c -matrices are as common to each other as possible, their difference is minimized:

$$\|W_{1,c} - W_{2,c}\|_F^2 \quad (2)$$

where $\|\cdot\|_F$ denotes the Frobenius-Norm defined as the square root of the sum of the squares of the matrix. Similarly, the overall conformance of both W_d matrices is minimized:

$$\|W_{1,d}^T W_{2,d}\|_{1,1} \quad (3)$$

where $\|\cdot\|_{1,1}$ denotes the absolute sum of all entries of the matrix. When combined with the minimization of the reconstruction error, the following optimization problem for JNMF is obtained:

$$\min_{W_1, H_1, W_2, H_2 \geq 0} \frac{1}{n_1} \|X_1 - W_1 H_1^T\|_F^2 + \frac{1}{n_2} \|X_2 - W_2 H_2^T\|_F^2 + \alpha \|W_{1,c} - W_{2,c}\|_F^2 + \beta \|W_{1,d}^T W_{2,d}\|_{1,1} \quad (4)$$

where $X_1 \in \mathbb{R}_+^{m \times n_1}$ and $X_2 \in \mathbb{R}_+^{m \times n_2}$ are matrices representing the datasets (i.e., *responses* \times *participants* for syllogistic reasoning data) and α and β are regularization-factors for the commonality and the difference terms, respectively. For syllogistic reasoning, the resulting matrices W_1 and W_2 consist

of k latent response-vectors, which can be interpreted as response patterns, while H_1 and H_2 assign the patterns to the respective participants.

Method

The overall goal of our application of JNMF to the domain of syllogistic reasoning is the explorative contrasting of reasoning patterns. The common and distinct submatrices W_c and W_d generated by JNMF allow us to investigate which patterns of inference remain consistent and which emerge as distinct across personality factors. In this scenario, the submatrix W_d is of special interest since it allows for a contrasting of reasoning behavior.

Syllogistic Dataset

For our analyses, we relied on existing data collected in a previous study (Dames et al., 2020). In this study, participants took part in three experimental sessions each one week apart and completed in addition to the syllogistic reasoning task various cognitive ability tests and personality questionnaires. The dataset we used is based on $n = 106$ participants and contains the data from the syllogistic reasoning test as well as the Big Five personality instrument (assessed with the German Big-Five-Inventory-SOEP; BFI-S; Gerlitz & Schupp, 2005). In the syllogistic reasoning tasks, participants had to generate a conclusion for all 64 possible syllogism types (production task design; content concerned names of professions, sports, and hobbies).

Data Preprocessing To make the data accessible to JNMF, we converted the categorical responses to their onehot-encoded representations, i.e., nine-dimensional vectors identifying a response by setting the value of the corresponding dimension to 1 while keeping the remaining dimensions at 0. For example, “All A are C” is encoded as (1, 0, 0, 0, 0, 0, 0, 0, 0) and “All C are A” as (0, 1, 0, 0, 0, 0, 0, 0, 0), inflating the data by a factor of nine. As a result, each individual is represented as a vector of dimensionality 576 , which leads to a total data matrix of dimensionality 576×106 . Since we are interested in contrasting individuals with different personalities, we split the data into two groups of individuals, scoring low and high based on the median values, for openness and conscientiousness. From this we obtain two input matrices X_{low} and X_{high} . Applying JNMF to X_{high} and X_{low} yields two pairs of matrices $W_{high} \in \mathbb{R}_+^{576 \times k}$, $H_{high} \in \mathbb{R}_+^{n_{high} \times k}$ and $W_{low} \in \mathbb{R}_+^{576 \times k}$, $H_{low} \in \mathbb{R}_+^{n_{low} \times k}$, where n_{high} and n_{low} denote the number of participants in the respective groups. The W matrices contain the identified response patterns (common and distinct), while the H matrices contain the association between the pattern and the participants.

Applying JNMF

Since we were aiming for a first intuition for the differences between the groups of reasoners and expected them to be small, we set $k_c = k_d = 1$. This results in the algorithm computing one general common pattern and one main distinct

Table 1: Results from applying JNMF to the syllogistic data while considering different personality factor groups (high vs. low). α and β show the hyperparameter values found by the grid search. Errors denote the results for the optimization targets of JNMF. Finally, importances based on the H matrices for the common and distinct patterns as well as the specific importances of the high and low patterns are shown.

Property	α	β	Errors			Importances			
			Common	Distinct	Reconstruction	Common	Distinct	High	Low
Conscientiousness	300	250	0.0199	0.0051	0.0011	0.6775	0.3225	0.3086	0.3334
Openness	500	250	0.0122	0.0038	0.0011	0.7415	0.2585	0.2718	0.2475

pattern for each group, which is well suited for obtaining a coarse overview of the behavioral properties of the data. For an in-depth analysis of different patterns that occur within a group, higher values for k_c or k_d are favorable but introduce the risk of mistakenly capturing noise if the patterns are too weak.

A core difficulty of using the batch-processing algorithm for JNMF (Kim et al., 2015) is the choice of hyperparameters α and β , which are used to regularize the matrix factorization procedure. To select optimal hyperparameter values, we applied a grid search optimizing for the reconstruction error, the deviance of the commonality between the common submatrices $W_{1,c}$ and $W_{2,c}$, as well as the degree of distinctiveness between $W_{1,d}$ and $W_{2,d}$. However, as for most optimization algorithms with random initialization, the algorithm might still find different local minima across multiple runs. To minimize the impact of this problem, we repeatedly ran the algorithm for each parameter configuration.

In the following section, we present an analysis of the matrices, and therefore reasoning patterns, resulting from the application of the JNMF using the best values obtained from the grid search. The respective errors and specific parameters are depicted in Table 1.

Results

Figure 1 illustrates the patterns from the W matrices obtained from applying JNMF to the dataset split by conscientiousness. The patterns for high and low correspond to columns in the respective W_d matrices, while the mean of the respective common columns in $W_{c,1}$ and $W_{c,2}$ is shown. For better clarity, the vectors with a dimensionality of 576 were re-transformed to matrices with dimensions of 64×9 , depicting the weights for the 9 responses for each of the 64 syllogisms in a heatmap. Due to space limitations, we only present the pattern for conscientiousness¹. However, openness showed similar results but with the high and low patterns reversed.

Figure 1 suggests that the differences between the high and common pattern are greater than between low and common. The main reason for this are differences in participants' likelihoods to respond NVC: the low pattern is largely focused around NVC whereas this response option is almost absent in the high pattern. This is an important observation since NVC

response behavior can be discussed as one of the major differences between deliberative and heuristic reasoning (Ragni et al., 2019). Following this interpretation, the patterns would suggest that conscientious reasoners are less likely to engage in deliberative reasoning.

The low errors (see Table 1) illustrate the success of the algorithm's applications with maximums of 0.02 for the common and 0.005 for the distinct patterns as well as a reconstruction error of 0.001. This allows us to investigate the patterns with low risk of interpreting random noise induced by the algorithm's initialization.

Additionally, Table 1 contains importance values for the obtained patterns that were derived from the values of the H matrices, which represent individual scaling factors for the patterns in W . The importances were calculated from the sum of the respective columns in the H matrices normalized by the total sum of the H matrices. By computing the proportion of values H contains for the common and distinct patterns, we can estimate the importance the algorithm attributes to them. The importances for the high and low patterns in Table 1 extend this by detailing the importances of the patterns for their respective personality groups. In general, groups seem to be very common (above 67% importance for all personality factors) with differences playing only a minor role. Moreover, there exist differences between the low and high pattern importances, suggesting that one group is slightly more consistent than the other (e.g., openness with 0.27 and 0.24 for high and low, respectively).

Furthermore, we can investigate the importance on the level of single participants. Figure 2 shows the distribution of the normalized H values for conscientiousness. The distribution for the common pattern suggests, that most participants are represented by the common pattern. The left-skewed distribution for the high and low patterns indicates that the pattern is not important for most participants in the respective group. However, for some participants the H values are high, which means that the pattern describes a subgroup of participants apparent in the respective group. This is to be expected, as the influence of conscientiousness, although reaching significance on a correlational level (Dames et al., in preparation), should not be strong enough to produce patterns that are consistent across all participants.

¹All results and materials are available on GitHub: github.com/Shadownox/cogsci-jnmf

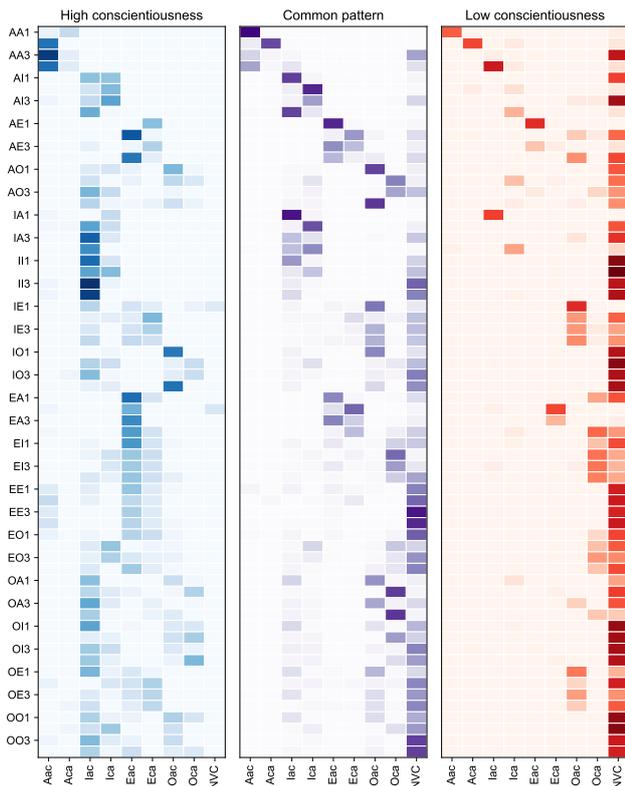


Figure 1: Heatmap of the patterns in the W matrices resulting from applying JNMF to datasets with high and low conscientious reasoners. For the 64 syllogisms (rows) and nine responses (columns), cell colors reflect the weights of the respective response to the pattern. Syllogisms and tasks are labeled using abbreviations A, I, E, and O for the quantifiers “All”, “Some”, “No”, and “Some ... not”, respectively, numbers reflecting the syllogistic figure, i.e., the order of terms in the premises (for an overview, see Khemlani & Johnson-Laird, 2012), as well as “ac” and “ca” representing the order of terms in the conclusion.

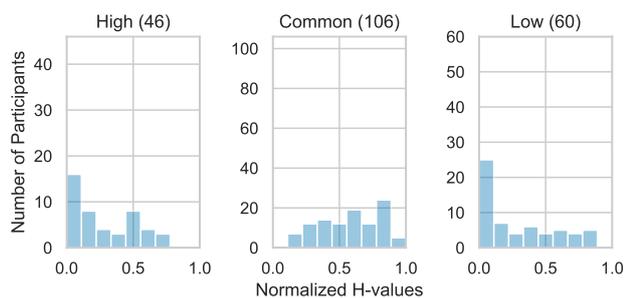


Figure 2: Illustration of the weights JNMF assigns between each pattern and reasoner for conscientiousness. Bars denote the frequencies of pattern values (bin widths of 0.1) in terms of the number of individuals they are assigned to. Numbers in the condition titles (high, common, low) denote the total number of individuals contained in this group.

Connection to Cognitive Theories

In an attempt to ground the obtained results on the theoretical work from the field of syllogistic reasoning research, we compute the match between the identified patterns and the predictions of the most prominent cognitive theories (as reported in Khemlani & Johnson-Laird, 2012). For this, the predicted conclusions of the cognitive theories were represented as 64×9 matrices, i.e., in the same way as the patterns. In cases where theories predicted multiple conclusions, a value of $\frac{1}{|\text{conclusions}|}$ was used for the corresponding responses. As a measure for the congruency we used the cosine similarity between the pattern and the theory-vector.

Figure 3 shows the results for conscientiousness and openness depicting the congruency between the obtained patterns with the conclusions predicted by syllogistic theories to illustrate their behavioral differences.

For conscientiousness, the radar plot confirms the observations detailed above, showing that reasoners in the low group seem more aligned with the inference mechanisms postulated by models inspired by logic principles such as PSYCOP (Rips, 1994), which assumes formal rule-based inferences. The pattern for highly conscientious reasoners, on the other hand, has its highest congruency with the Matching heuristic (Wetherick & Gilhooly, 1995) representing a classical fast-and-frugal system 1 heuristic (Evans, 2003) operating only on structural properties of syllogisms.

Those participants also show a tendency towards congruency with models which do not include NVC in their predictions (e.g., Atmosphere, Matching), which is in line with the observation from Figure 1.

Overall, the low pattern aligns better with the theories of syllogistic reasoning indicating a more consistent reasoning behavior. The high pattern, on the other hand, generally scores lower, which could be an indicant for more noise in its pattern. This might limit its ability to be aligned with a theory’s prediction. Results for openness show the opposite with the high pattern being more consistent with logical theories and vice versa.

Validation of Pattern Predictions

In a final analysis, we want to make sure that the obtained patterns are indeed able to provide suitable descriptions for the distinct behavior of the personality factor groups they were extracted from. To this end, we subject patterns as models to an evaluation of their predictive performance. We expect them to be able to better capture their own group than the other. Note, that by treating patterns as models we do not expect them to be highly accurate on an absolute scale since they only reflect their group’s distinct behavior without considering the much more important common parts of the group’s reasoning behavior (cf. Table 1). As such, we focus on an interpretation of the changes in predictive accuracies between the groups.

Figure 4 shows that both personality factors agree in that the patterns obtained from JNMF truly capture the distinct be-

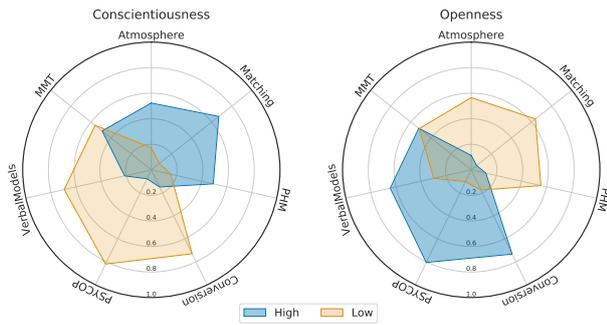


Figure 3: Radar plots depicting the match between JNMF patterns (for reasoners with high and low personality factor values) and cognitive theories for human syllogistic reasoning (taken from Khemlani & Johnson-Laird, 2012). Distances to the center reflect congruencies with the respective theories.

behavior for their respective groups (accuracies decrease when applied to the opposite group). Also, as expected, the common pattern performs better on an absolute scale than the distinct models confirming its generally higher importance (cf. Table 1).

For conscientiousness, as expected, the performance of the low pattern is more similar to common than to high pattern's performance (follows the observations from Figure 1). Moreover, the high group is represented worse by the pattern extracted from it. This could hint at inconsistent response behavior for example due to noise instead of actual systematic reasoning strategies that could be postulated based on the observations from Figure 3.

In comparison to conscientiousness, openness exhibits inverse results. Overall, the differences in performances between the data groups are much smaller than for conscientiousness, which suggest an overall lower significance of the patterns.

General Discussion

In this article we proposed the use of matrix factorization for the explorative analysis of behavioral data from cognitive science and demonstrated its application to the domain of syllogistic reasoning. Matrix factorization is a valuable method for the development of intuition with successful applications to a wide variety of domains (see Liu et al., 2006). Joint nonnegative Matrix Factorization (JNMF; Kim et al., 2015), a recently introduced extension to standard Nonnegative Matrix Factorization (NMF), allows for the decomposition $X \approx WH^T$ for two datasets X_1 and X_2 simultaneously. As its output, the method produces two sets of W and H matrices representing the common and distinct patterns in the group. For behavioral data, this allows for the detailed contrasting of different groups of individuals.

Applied to syllogistic reasoning data and additional information about big five personality factors, JNMF allowed us to obtain patterns representing the common and distinct behaviors between groups of different personalities (high vs. low

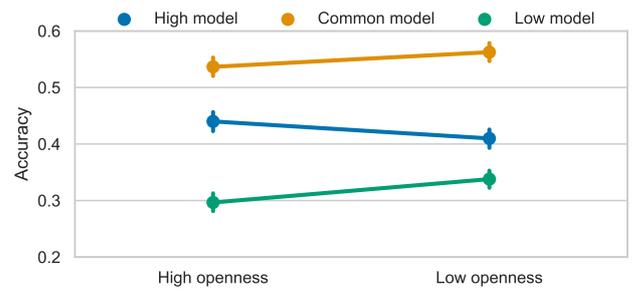
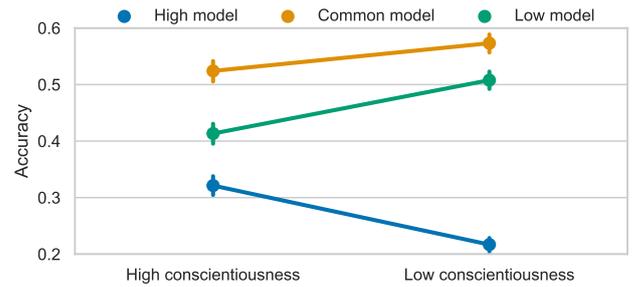


Figure 4: Evaluation of the obtained patterns in terms of their predictive performance (high, common, and low model). Patterns are expected to score higher on their respective groups of reasoners (high vs. low personality factors) and lower than common due to their limited importance (cf. Table 1).

values for conscientiousness and openness). Overall, JNMF managed to reliably represent the data in terms of its resulting W and H matrices. Our results were in line with our previous findings (conscientiousness is the most influential factor), but, crucially, allowed for a more in-depth analysis of the differences between groups by detailing the distinct reasoning patterns. By interpreting the scoring weights for the identified patterns contained in the H matrix, we were able to deduce their importances. Moreover, comparing the patterns to the conclusion predictions of cognitive theories, we were able to provide a method allowing for grounding results on the theoretical insight in the field. Finally, we presented a method for validating the distinctiveness of the patterns by interpreting them as predictive models, thereby making them accessible for statistical testing.

Although JNMF achieved promising results already, the proposed method still offers much potential for improvement. Domain-specific regularization of the factorization and sparsity constraints could help to reduce the complexity of the method while providing more precise patterns. Additionally, future work could focus on finding possibilities to disentangle the different behavioral strategies contained in the data by determining optimal values for k_c and k_d . In sum, JNMF proved to be a promising and reliable tool to identify different behavioral patterns. Importantly, this method adds to typically performed statistical predictor analyses (e.g., regressions) as the contrasting allows for a deeper understanding of the data.

Acknowledgements

This paper was supported by DFG grants RA 1934/2-1, RA 1934/3-1 and RA 1934/4-1 to MR.

References

- Berry, M. W., Browne, M., Langville, A. N., Pauca, V. P., & Plemmons, R. J. (2007). Algorithms and applications for approximate nonnegative matrix factorization. *Computational Statistics & Data Analysis*, 52(1), 155–173.
- Brand, D., Riesterer, N., & Ragni, M. (2019). On the matter of aggregate models for syllogistic reasoning: A transitive set-based account for predicting the population. In T. Stewart (Ed.), *Proceedings of the 17th International Conference on Cognitive Modeling* (pp. 5–10). Waterloo, Canada: University of Waterloo.
- Carretta, T. R., & Ree, M. J. (2018). The relations between cognitive ability and personality: Convergent results across measures. *International Journal of Selection and Assessment*, 26(2-4), 133–144.
- Dames, H., Klauer, K. C., & Ragni, M. (2020). *Individual differences in syllogistic reasoning performance: Why and when do individuals differ?* OSF. Retrieved from osf.io/x3wvf
- Dames, H., Klauer, K. C., & Ragni, M. (in preparation). The associations between personality, cognitive ability, and reasoning.
- Evans, J. S. (2003). In two minds: dual-process accounts of reasoning. *Trends in Cognitive Sciences*, 7(10), 454–459.
- Gerlitz, J., & Schupp, J. (2005). Zur erhebung der big-five basierten persönlichkeitsmerkmale im soep. *DIW Research Notes*, 4.
- Hoyer, P. O. (2004). Non-negative matrix factorization with sparseness constraints. *Journal of Machine Learning Research*, 5(Nov), 1457–1469.
- Khemlani, S., & Johnson-Laird, P. N. (2012). Theories of the syllogism: A meta-analysis. *Psychological Bulletin*, 138(3), 427–457.
- Khemlani, S., & Johnson-Laird, P. N. (2016). How people differ in syllogistic reasoning. In A. Papafragou, D. Grodner, D. Mirman, & J. C. Trueswell (Eds.), *Proceedings of the 38th Annual Conference of the Cognitive Science Society* (pp. 2165–2170). Austin, TX: Cognitive Science Society.
- Kim, H., Choo, J., Kim, J., Reddy, C. K., & Park, H. (2015). Simultaneous discovery of common and discriminative topics via joint nonnegative matrix factorization. In *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining - KDD '15*. ACM Press.
- Lee, D. D., & Seung, H. S. (2001). Algorithms for non-negative matrix factorization. In *Advances in neural information processing systems* (pp. 556–562).
- Liu, W., Zheng, N., & You, Q. (2006). Nonnegative matrix factorization and its applications in pattern recognition. *Chinese Science Bulletin*, 51(1), 7–18.
- McCrae, R. R., & John, O. P. (1992). An introduction to the five-factor model and its applications. *Journal of Personality*, 60(2), 175–215.
- Moutafi, J., Furnham, A., & Crump, J. (2006). What facets of openness and conscientiousness predict fluid intelligence score? *Learning and Individual Differences*, 16(1), 31–42.
- Moutafi, J., Furnham, A., & Paltiel, L. (2004). Why is conscientiousness negatively correlated with intelligence? *Personality and Individual Differences*, 37(5), 1013–1022.
- Murphy, K. P. (2012). *Machine learning*. MIT Press Ltd.
- Pauca, V. P., Shahnaz, F., Berry, M. W., & Plemmons, R. J. (2004). Text mining using non-negative matrix factorizations. In *Proceedings of the 2004 siam international conference on data mining* (pp. 452–456).
- Poropat, A. E. (2009). A meta-analysis of the five-factor model of personality and academic performance. *Psychological Bulletin*, 135(2), 322–338.
- Ragni, M., Dames, H., Brand, D., & Riesterer, N. (2019). When does a reasoner respond: Nothing follows? In A. K. Goel, C. M. Seifert, & C. Freksa (Eds.), *Proceedings of the 41st Annual Conference of the Cognitive Science Society* (pp. 2640–2646). Montreal, QB: Cognitive Science Society.
- Rammstedt, B., Danner, D., & Martin, S. (2016). The association between personality and cognitive ability: Going beyond simple effects. *Journal of Research in Personality*, 62, 39–44.
- Riesterer, N., Brand, D., & Ragni, M. (2019). Predictive modeling of individual human cognition: Upper bounds and a new perspective on performance. In T. Stewart (Ed.), *Proceedings of the 17th International Conference on Cognitive Modeling* (pp. 178–183). Waterloo, Canada: University of Waterloo.
- Rips, L. J. (1994). *The psychology of proof: Deductive reasoning in human thinking*. Cambridge, MA: MIT Press.
- Störring, G. (1908). *Experimentelle Untersuchungen über einfache Schlussprozesse*. Leipzig, Germany: W. Engelmann.
- Wetherick, N. E., & Gilhooly, K. J. (1995). ‘atmosphere’, matching, and logic in syllogistic reasoning. *Current Psychology*, 14(3), 169–178.
- Ziegler, M., Danay, E., Heene, M., Asendorpf, J., & Bühner, M. (2012). Openness, fluid intelligence, and crystallized intelligence: Toward an integrative model. *Journal of Research in Personality*, 46(2), 173–183.