

---

# Correspondence Analysis of Government Expenditure Patterns

---

**Hsiang Hsu, Flavio P. Calmon, José Cândido Silveira Santos Filho\***

John A. Paulson School of Engineering and Applied Sciences  
Harvard University  
Cambridge, MA 02138

{hsianghsu, fcalmon, candido}@g.harvard.edu

**Andre P. Calmon**

Technology and Operations Management  
INSEAD  
Fontainebleau, France  
andre.calmon@insead.edu

**Salman Salamatian**

Research Laboratory of Electronics  
Massachusetts Institute of Technology  
Cambridge, MA 02139  
salmansa@mit.edu

## Abstract

We analyze expenditure patterns of discretionary funds by Brazilian congress members. This analysis is based on a large dataset containing over 7 million expenses made publicly available by the Brazilian government. This dataset has, up to now, remained widely untouched by machine learning methods. Our main contributions are two-fold: (i) we provide a novel dataset benchmark for machine learning-based efforts for government transparency to the broader research community, and (ii) introduce a neural network-based approach for analyzing and visualizing outlying expense patterns. Our hope is that the approach presented here can inspire new machine learning methodologies for government transparency applicable to other developing nations.

## 1 Introduction

Over the last decade, an increasing number of the World's governments and, in particular, the executive and legislative branches of these governments, have made data on their activities and expenditures publicly available (Bates, 2012). This government-led *open data* movement seeks to increase transparency, reduce corruption, make government activities more accessible to citizens and, ultimately, strengthen democratic institutions (Janssen et al., 2012).

The open data trend in the public sector has led to many data science and machine learning (ML) based initiatives that seek to quantify, model, and evaluate the performance of public administration. In particular, *Public Expenditure Analysis* (PEA) (Shah, 2005), which investigates how government budgets are spent, has become an active area of research in social and political science (Lopez et al., 2016; de Sousa et al., 2017; Garry and Rivas Valdivia, 2017; Odhiambo, 2018).

Within this context, the goal of this paper is to apply machine learning tools to perform PEA on data from a developing country whose executive and legislative branches have recently been marred by multiple budget misuse problems (Winter, 2017; Cagni, 2017). Specifically, we apply a neural network(NN)-based technique (Hsu et al., 2018) to examine, visualize, and interpret the expenditure

---

\*J. C. S. Santos Filho is also with the Department of Communications, School of Electrical and Computer Engineering, University of Campinas, Campinas, SP, Brazil (e-mail: candido@decom.fee.unicamp.br).

of discretionary funds by congress members of the Brazilian House of Congress (*Câmara dos Deputados*). This data has been made publicly available by the Brazilian government for about ten years, yet remains widely untouched by advanced ML techniques. We have translated the dataset to English and made it publicly available to the broader research community through the accompanying repository (Hsu and Calmon, 2018). Our hope is that this dataset serves as a benchmark for new methodological ML-based approaches for ensuring and evaluating government transparency. We note that, more often than not, open government data is analyzed using a “descriptive” approach (e.g., finding outlier expenses, computing aggregate expenditure per congress member), as opposed to using more systematic ML techniques. Our ultimate goal is to reverse this trend.

There are a few reasons why we focus on the discretionary expenses by the Brazilian Congress. First, the Brazilian government has a large open data initiative (The Brazilian Ministry of Planning, 2012). Up to now, this data has been analyzed mostly through descriptive analytics. For example, the *Operação Serenata de Amor* (Musskopf, 2016) has cleaned this data and made it available in a format that is easy to analyze, yet we are unaware of any efforts that use advanced ML techniques directly. Second, over the last few years many members of the Brazilian congress have been involved in high-profile budget misuse problems which have made global headlines (Winter, 2017), creating a natural test dataset for identifying budget misuse (some reports indicate that over 30% of Brazilian congress members as of 2018 are under investigation (Cagni, 2017; Sardinha, 2018)). This data can be used to validate methodological approaches that are adaptable to other countries.

From a methodological standpoint, we use a generalization of *Correspondence Analysis* (CA) to continuous variables and high-dimensional data to visualize and interpret expenditure patterns by congress members. This approach is more suitable for the investigated dataset than traditional methods such as Principal Components Analysis (PCA) and Canonical Correlation Analysis (Hotelling, 1936). The potential use cases of the data and the method we present are four-fold:

1. Anomalous expenditure discovery and prediction in order to perform proactive reactions against budget misuse.
2. Clustering of congress members in terms of their discretionary expenditure pattern.
3. Interpretation and visualization of the expenditures, creation of algorithmic watchdogs for misuse.
4. Inspiration for new methodological approaches for government transparency transferable to other civic projects that aim at similar goals.

All codes for downloading, translating and parsing the dataset is available at (Hsu and Calmon, 2018); the dataset itself is made available by Brazilian government in (Musskopf, 2016). In the rest of this paper, we first describe the main ML tool used, namely CA using neural networks (Section 2), and then describe the dataset and numerical results (Section 3).

## 2 Neural Network-based CA for Visualizing and Interpreting Expenditure

CA is an exploratory multivariate statistical technique that converts data into a graphical display with orthogonal factors. In a similar vein to PCA and its kernel variants (Hoffmann, 2007), CA is a technique that maps the data onto a low-dimensional representation. By construction, this new representation captures possibly non-linear relationships between the underlying variables, and can be used to interpret the dependence between two random variables  $X$  and  $Y$  from observed samples. CA has the ability to produce interpretable, low-dimensional visualizations (often two-dimensional) that capture complex relationships in data with entangled and intricate dependencies. This has led to its successful deployment in fields ranging from genealogy and epidemiology to social and environmental sciences (Tekaiia, 2016; Sourial et al., 2010; Carrington et al., 2005; ter Braak and Schaffers, 2004; Ormoli et al., 2015; Ferrari et al., 2016).

CA considers two random variables  $X$  and  $Y$  with  $|\mathcal{X}| < \infty$ ,  $|\mathcal{Y}| < \infty$ , and their joint distribution  $P_{X,Y}$  (cf. Greenacre (1984) for a detailed overview). Given samples  $\{x_k, y_k\}_{k=1}^n$  drawn independently from  $P_{X,Y}$ , a two-way contingency table  $\mathbf{P}_{X,Y}$  is defined as a matrix with  $|\mathcal{X}|$  rows and  $|\mathcal{Y}|$  columns of normalized co-occurrence counts, i.e.  $[\mathbf{P}_{X,Y}]_{i,j} = (\# \text{ of observations } (x_i, y_j) = (i, j))/n$ . Moreover, the marginals are defined as  $\mathbf{p}_X \triangleq \mathbf{P}_{X,Y} \mathbf{1}_{|\mathcal{Y}|}$  and  $\mathbf{p}_Y \triangleq \mathbf{P}_{X,Y}^T \mathbf{1}_{|\mathcal{X}|}$ . Consider a matrix  $\mathbf{Q} \triangleq \mathbf{D}_X^{-1/2} (\mathbf{P}_{X,Y} - \mathbf{p}_X \mathbf{p}_Y^T) \mathbf{D}_Y^{-1/2}$ , where  $\mathbf{D}_X \triangleq \text{diag}(\mathbf{p}_X)$  and  $\mathbf{D}_Y \triangleq \text{diag}(\mathbf{p}_Y)$ , and let

the singular value decomposition of  $\mathbf{Q}$  be  $\mathbf{Q} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$ . Let  $d = \min\{|\mathcal{X}|, |\mathcal{Y}|\} - 1$ , and  $\{\sigma_i\}_{i=1}^d$  be the singular values, then we have the following definitions (Greenacre, 1984):

- Orthogonal factors of  $X$ :  $\mathbf{L} \triangleq \mathbf{D}_X^{-1/2}\mathbf{U}$ .
- Orthogonal factors of  $Y$ :  $\mathbf{R} \triangleq \mathbf{D}_Y^{-1/2}\mathbf{V}$ .
- Factor scores:  $\lambda_i = \sigma_i^2, 1 \leq i \leq d$ .
- Factor score ratios:  $\frac{\lambda_i}{\sum_{i=1}^d \lambda_i}, 1 \leq i \leq d$ .

The first and second columns of  $\mathbf{L}$  and  $\mathbf{R}$  can be plotted on a two-dimensional plane (with each row corresponding to a point) producing the *factoring plane*. The remaining planes can be produced by plotting the other columns of  $\mathbf{L}$  and  $\mathbf{R}$ . The factor score ratio quantifies the correlations captured by each orthogonal factor, and is often shown along the axes in factoring planes.

**Deep Neural Networks for Correspondence Analysis.** The contingency table-based approach for CA has three fundamental limitations. First, it is restricted to data drawn from *discrete* distributions with finite support, since contingency tables for continuous variables will be highly dependent on a chosen quantization which, in turn, may jeopardize information in the data. Second, even when the underlying distribution of the data is discrete, reliably estimating the contingency table (i.e., approximating  $P_{X,Y}$ ) may be infeasible due to limited number of samples. This inevitably hinges CA on the more statistically challenging problem of estimating  $P_{X,Y}$ . Third, building contingency tables is not feasible for *high-dimensional* data. This limitation can be circumvented by using a novel neural network-based approach for CA introduced in (Hsu et al., 2018).

Here, we summarize the neural network-based approach for CA in (Hsu et al., 2018). Consider two neural networks F-Net and G-Net, which encode  $X$  and  $Y$  to  $\mathbb{R}^d$  respectively. We denote the outputs from the F and G-Net of  $X$  and  $Y$ , respectively, as

$$\tilde{\mathbf{f}}(X) \triangleq [\tilde{f}_1(X), \dots, \tilde{f}_d(X)]^\top \in \mathbb{R}^{d \times 1}, \text{ and } \tilde{\mathbf{g}}(Y) \triangleq [\tilde{g}_1(Y), \dots, \tilde{g}_d(Y)]^\top \in \mathbb{R}^{d \times 1}. \quad (1)$$

The solution of the optimization problem

$$\min_{\mathbf{A} \in \mathbb{R}^{d \times d}, \tilde{\mathbf{f}}, \tilde{\mathbf{g}}} \mathbb{E} \left[ \|\mathbf{A}\tilde{\mathbf{f}}(X) - \tilde{\mathbf{g}}(Y)\|_2^2 \right], \quad \text{subject to} \quad \mathbb{E} \left[ \mathbf{A}\tilde{\mathbf{f}}(X)(\mathbf{A}\tilde{\mathbf{f}}(X))^\top \right] = \mathbf{I}_d \quad (2)$$

recovers the orthogonal factors of  $X$  and  $Y$  (Hsu et al., 2018). Using theoretical results from orthogonal Procrustes problem (Gower and Dijksterhuis, 2004), we can further simplify the objective function (2) into an unconstrained version:

$$\min_{\tilde{\mathbf{f}}, \tilde{\mathbf{g}}} -2\|\mathbf{C}_f^{-\frac{1}{2}}\mathbf{C}_{fg}\|_d + \mathbb{E}[\|\tilde{\mathbf{g}}(Y)\|_2^2], \quad (3)$$

where  $\mathbf{C}_f = \mathbb{E}[\tilde{\mathbf{f}}(X)\tilde{\mathbf{f}}(X)^\top]$ ,  $\mathbf{C}_{fg} = \mathbb{E}[\tilde{\mathbf{f}}(X)\tilde{\mathbf{g}}(Y)^\top]$ , and  $\|\mathbf{Z}\|_d$  is the  $d$ -th Ky-Fan norm, defined as the sum of the singular values of  $\mathbf{Z}$  (Horn et al., 1990). Denoting by  $\mathbf{A}$  and  $\mathbf{B}$  the whitening matrices for  $\tilde{\mathbf{f}}(\mathbf{X})$  and  $\tilde{\mathbf{g}}(\mathbf{Y})$ , the orthogonal factors of  $X$  and  $Y$  are given by  $\mathbf{f}(X) = [f_0(X), \dots, f_d(X)]^\top = \mathbf{A}\tilde{\mathbf{f}}(X)$  and  $\mathbf{g}(Y) = [g_0(Y), \dots, g_d(Y)]^\top = \mathbf{B}\tilde{\mathbf{g}}(Y)$ . The factor score  $\lambda_i$  is given by  $\mathbb{E}[f_i(X)g_i(Y)]$ ,  $1 \leq i \leq d$ . The loss (3) is unconstrained over the space of all finite variance functions of  $X$  and  $Y$ , and therefore is trainable via back-propagation using the common loss function (3). For more information about optimization details, see (Hsu et al., 2018).

### 3 The Data: Brazilian House of Congress Discretionary Spending

**Description and Pre-processing of the Dataset.** We investigate data on discretionary funding reimbursements from the Brazilian House of Congress. This data was made openly and freely available (in Portuguese) by The Brazilian Ministry of Planning (2012). Each Brazilian congress member receives a certain amount of discretionary funding for supporting parliamentary activity (*Cota para o Exercício da Atividade Parlamentar – CEAP*) (The Brazilian House of Congress, 2018). This fund is used to reimburse travel, food, phone bills, postal services, cabinet costs, etc. The limit that each congress member can spend depends on their state of origin, with a maximum monthly cap of around BRL\$45k (about USD\$13k) (The Brazilian House of Congress, 2018). Brazilian Congress




Figure 1: The first factoring plane of the expenditures of 595 congress members in Brazil from 2015 to 2018. The factor score ratios are shown with the axis; higher factor score ratio means more correlation captured by the orthogonal factor. Each trace (e.g. pink line) represents the expenditure pattern of a category for all congress members. Each grey dot represents a congress member without investigations, and each red dot represents one under investigations according to (Cagni, 2017), where the radius of red dots indicates the number of investigations (We did not independently verify the completeness/ correctness of the dataset, and recommend caution when using information about ongoing investigations to avoid potential errors (false positive) in analysis). Points and lines close to the center (the origin) indicate small correlation. Names, states, and parties of congress members are omitted.

has 513 seats distributed among 26 states and the Federal District. Brazil has several political parties, with over 30 parties being represented in Congress as of 2018. The term for a congress member is 4 years.

We have produced code in Python for automatically downloading, translating and parsing this data, as well as meta-data regarding the multiple features found in the dataset, available at (Hsu and Calmon, 2018). The dataset contains more than 7 million expenditure records from 2009 to 2018, including the category (e.g., fuel, food, office maintenance, airline tickets), values, date, and vendor that produced the receipt for the expenditure. Moreover, the states, parties, and names of the congress members are also included. In the analysis here, we present the records for the most recent term (i.e., 2015–2018), dropped missing data points, and eliminated categories that appear less than 500 times. The resulting dataset finally contains approximately 1.1 million expenditure records in 16 categories

of 595 congress members from 26 parties and 26 states and the Federal District (the number of congress members is greater than the number of seats since not all members finish their term). For the CA, we set  $X$  to be the categories and values of the expenditure and  $Y$  to be the congress member with their parties and states, and perform a 70%-30% training-validation split of the data.

**Neural Network and Training Configuration.** The F and G-Net are composed of two simple feed-forward neural networks with different structures. The F-Net has four layers with number of units 1000, 500, 300, 50 and G-Net has three layers with number of units 100, 50, 30. We adopt  $\tanh$  activation for hidden layers and the readout layer. We train for 20 epochs on the training set with a batch size of 256 using a gradient descent optimizer with a learning rate of 0.01. The result of the CA for expenditure analysis is shown in Fig. 1.

**Expenditure Pattern.** In Fig. 1 (see caption for instructions), we show the expenditure patterns of 16 categories and the congress member in a standard CA factor-plane plot (Greenacre, 1984). CA is performed using the NN-based approach described in the previous section. We summarize our observations below:

- Our generalized CA approach automatically clusters related expenses together since they have close patterns, e.g., aviation-related expenses “Airline Ticket Issue” and “Rental of aircrafts”, transportation-related expenses “River transport tickets” and “Rental of motor vehicles”, and daily expenses “Food”, “Fuel”, and “Security services”.
- There are certain categories of expenditures that are not correlated with specific congress members: “Food”, “Fuel and lubricants”, “River transport tickets”, “Rental of motor vehicles”, “Security services”, and “Taxi services and parking”.
- Categories that show high variations also have clear pattern. For instance, overlapping traces of “Publication subscription” and “Postal service”, and “Airline tickets”, “Consulting, research, and technical activities” and “Disclosure and advertisement of parliamentary activity” can be observed. Moreover, “Disclosure and advertisement of parliamentary activity” has a very large variation (pink line on the left-hand side of the graph). This may potentially indicate mishandling of expenses in these categories.
- Two categories exhibit outlying patterns: “Maintenance of an office” and “Lodging”. This might indicate that in different states, the expense on the two categories is dramatically different, or could be an indication of foul play. This can help direct further investigatory efforts.

**Charged Congress members.** We also collected information from publicly available sources on congress members that are currently under investigation (Cagni, 2017; Sardinha, 2018)<sup>2</sup>. We display in Fig. 1 those who are under multiple investigations. As we can see, the investigated congress members are concentrated near expenditure patterns that have large variation, i.e. outliers. This may indicate that congress members under multiple investigations also deviate from the mean use of discretionary funding, suggesting that discretionary funding may be predictive of misbehaviours — even though further investigation is required to confirm this statement. This analysis demonstrate how modern ML techniques can be applied to this large dataset to both visualize and interpret congress member behaviours.

## References

- Bates, J. (2012). “this is what modern deregulation looks like”: co-optation and contestation in the shaping of the uk’s open government data initiative. *The Journal of Community Informatics*, 8(2).
- Cagni, P. (2017). Os deputados sob investigação no supremo tribunal federal. Congresso em Foco, <https://congressoemfoco.uol.com.br/especial/noticias/os-deputados-sob-investigacao-no-supremo-tribunal-federal/>.
- Carrington, P. J., Scott, J., and Wasserman, S. (2005). *Models and methods in social network analysis*, volume 28. Cambridge university press.
- de Sousa, R. G., Paulo, E., and Marôco, J. (2017). Longitudinal factor analysis of public expenditure composition and human development in brazil after the 1988 constitution. *Social Indicators Research*, 134(3):1009–1026.

<sup>2</sup>We did not independently verify the completeness/ accuracy of the dataset, and recommend caution when using information about ongoing investigations to avoid potential errors (false positive) in analysis.

- Ferrari, A., Vincent-Salomon, A., Pivot, X., Sertier, A.-S., Thomas, E., Tonon, L., Boyault, S., Mu-lugeta, E., Treilleux, I., Macgrogan, G., et al. (2016). A whole-genome sequence and transcriptome perspective on her2-positive breast cancers. *Nature communications*, 7:12222.
- Garry, S. and Rivas Valdivia, J. C. (2017). An analysis of the contribution of public expenditure to economic growth and fiscal multipliers in Mexico, Central America and the Dominican Republic, 1990-2015.
- Gower, J. C. and Dijksterhuis, G. B. (2004). *Procrustes problems*, volume 30. Oxford University Press on Demand.
- Greenacre, M. J. (1984). *Theory and applications of correspondence analysis*. London (UK) Academic Press.
- Hoffmann, H. (2007). Kernel PCA for novelty detection. *Pattern recognition*, 40(3):863–874.
- Horn, R. A., Horn, R. A., and Johnson, C. R. (1990). *Matrix analysis*. Cambridge University Press.
- Hotelling, H. (1936). Relations between two sets of variates. *Biometrika*, 28(3/4):321–377.
- Hsu, H. and Calmon, F. P. (2018). Camara Brazil. <https://github.com/HsiangHsu/Brazilian-Congress-Expenditure>.
- Hsu, H., Salamatian, S., and Calmon, F. P. (2018). Deep orthogonal representations: Fundamental properties and applications. *arXiv preprint arXiv:1806.08449*.
- Janssen, M., Charalabidis, Y., and Zuiderwijk, A. (2012). Benefits, adoption barriers and myths of open data and open government. *Information systems management*, 29(4):258–268.
- Lopez, G. H. N., Mori, E. S., Avila, L., Lozano, R., et al. (2016). A performance analysis of public expenditure on maternal health in Mexico.
- Musskopf, I. (2016). Operation serenata de amor. <https://serenata.ai/en/>.
- Odhiambo, N. M. (2018). Public expenditure and economic growth in Kenya: A multivariate dynamic causal linkage.
- Ormolì, L., Costa, C., Negri, S., Perenzin, M., and Vaccino, P. (2015). Diversity trends in bread wheat in Italy during the 20th century assessed by traditional and multivariate approaches. *Scientific reports*, 5:8574.
- Sardinha, E. (2018). Um em cada três deputados é acusado de crimes. Congresso em Foco, <https://congressoemfoco.uol.com.br/especial/noticias/um-em-cada-tres-deputados-e-acusado-de-crimes/>.
- Shah, A. (2005). *Public Expenditure Analysis*. The World Bank.
- Sourial, N., Wolfson, C., Zhu, B., Quail, J., Fletcher, J., Karunanathan, S., Bandeen-Roche, K., Béland, F., and Bergman, H. (2010). Correspondence analysis is a useful tool to uncover the relationships among categorical variables. *Journal of clinical epidemiology*, 63(6):638–646.
- Tekaia, F. (2016). Genome data exploration using correspondence analysis. *Bioinformatics and Biology insights*, 10:BBI-S39614.
- ter Braak, C. J. and Schaffers, A. P. (2004). Co-correspondence analysis: a new ordination method to relate two community compositions. *Ecology*, 85(3):834–846.
- The Brazilian House of Congress (2018). Cota para o Exercício da Atividade Parlamentar – CEAP. [http://www2.camara.leg.br/transparencia/aceso-a-informacao/copy\\_of\\_perguntas-frequentes/cota-para-o-exercicio-da-atividade-parlamentar](http://www2.camara.leg.br/transparencia/aceso-a-informacao/copy_of_perguntas-frequentes/cota-para-o-exercicio-da-atividade-parlamentar).
- The Brazilian Ministry of Planning (2012). Portal brasileiro de dados abertos. <http://dados.gov.br>.
- Winter, B. (2017). Brazil’s never-ending corruption crisis: Why radical transparency is the only fix. *Foreign Aff.*, 96:87.