

What Drives the International Development Agenda? An NLP Analysis of the United Nations General Debate 1970-2016

Alexander Baturó
School of Law and Government
Dublin City University
Email: alex.baturó@dcu.ie

Niheer Dasandi
School of Government and Society
University of Birmingham
Email: n.dasandi@bham.ac.uk

Abstract—Surprisingly little is known about agenda setting for international development in the United Nations (UN), despite it having a significant influence on the process and outcomes of development efforts. This paper addresses this shortcoming using a novel approach that applies natural language processing techniques to countries’ annual statements in the UN General Debate. Every year UN member states deliver statements during the General Debate on their governments’ perspective on major issues in world politics. These speeches provide invaluable information on state preferences on a wide range of issues, including international development, but have largely been overlooked in the study of global politics. This paper identifies the main international development topics that states raise in these speeches between 1970 and 2016, and examine the country-specific drivers of international development rhetoric.

I. INTRODUCTION

Decisions made in international organisations are fundamental to international development efforts and initiatives. It is in these global governance arenas that the rules of the global economic system, which have a huge impact on development outcomes are agreed on; decisions are made about large-scale funding for development issues, such as health and infrastructure; and key development goals and targets are agreed on, as can be seen with the Millennium Development Goals (MDGs). More generally, international organisations have a profound influence on the ideas that shape international development efforts [1].

Yet surprisingly little is known about the agenda-setting process for international development in global governance institutions. This is perhaps best demonstrated by the lack of information on how the different goals and targets of the MDGs were decided, which led to much criticism and concern about the global governance of development [2]. More generally, we know little about the types of development issues that different countries prioritise, or whether country-specific factors such as wealth or democracy make countries more likely to push for specific development issues to be put on the global political agenda.

The lack of knowledge about the agenda setting process in the global governance of development is in large part due to the absence of obvious data sources on states’ preferences

about international development issues. To address this gap we employ a novel approach based on the application of natural language processing (NLP) to countries’ speeches in the UN. Every September, the heads of state and other high-level country representatives gather in New York at the start of a new session of the United Nations General Assembly (UNGA) and address the Assembly in the General Debate. The General Debate (GD) provides the governments of the almost two hundred UN member states with an opportunity to present their views on key issues in international politics – including international development. As such, the statements made during GD are an invaluable and, largely untapped, source of information on governments’ policy preferences on international development over time.

An important feature of these annual country statements is that they are not institutionally connected to decision-making in the UN. This means that governments face few external constraints when delivering these speeches, enabling them to raise the issues that they consider the most important. Therefore, the General Debate acts “as a barometer of international opinion on important issues, even those not on the agenda for that particular session” [3]. In fact, the GD is usually the first item for each new session of the UNGA, and as such it provides a forum for governments to identify like-minded members, and to put on the record the issues they feel the UNGA should address. Therefore, the GD can be viewed as a key forum for governments to put different policy issues on the international agenda.

We use a new dataset of GD statements from 1970 to 2016, *the UN General Debate Corpus* (UNGDC), to examine the international development agenda in the UN [4].¹ Our application of NLP to these statements focuses in particular on structural topic models (STMs) [5]. The paper makes two contributions using this approach: (1) It sheds light on the main international development issues that governments prioritise in the UN; and (2) It identifies the key country-specific factors associated with governments discussing development issues in

¹UNGDC is publicly available at the Harvard Dataverse at <http://dx.doi.org/10.7910/DVN/0TJX8Y>




Fig. 1. *Optimal model search.* Semantic coherence and exclusivity results for a model search from 3 to 50 topics. Models above the regression line provide a better trade off. Largest positive residual is a 16-topic model.

their GD statements.

THE UN GENERAL DEBATE AND INTERNATIONAL DEVELOPMENT

In the analysis we consider the nature of international development issues raised in the UN General Debates, and the effect of structural covariates on the level of developmental rhetoric in the GD statements. To do this, we first implement a structural topic model [5]. This enables us to identify the key international development topics discussed in the GD. We model topic prevalence in the context of the structural covariates. In addition, we control for region fixed effects and time trend. The aim is to allow the observed metadata to affect the frequency with which a topic is discussed in General Debate speeches. This allows us to test the degree of association between covariates (and region/time effects) and the average proportion of a document discussing a topic.

A. Estimation of topic models

We assess the optimal number of topics that need to be specified for the STM analysis. We follow the recommendations of the original STM paper and focus on exclusivity and semantic coherence measures. [6] propose semantic coherence measure, which is closely related to point-wise mutual information measure posited by [7] to evaluate topic quality. [6] show that semantic coherence corresponds to expert judgments and more general human judgments in Amazon’s Mechanical Turk experiments.

Exclusivity scores for each topic follows [8]. Highly frequent words in a given topic that do not appear very often in other topics are viewed as making that topic exclusive. Cohesive and exclusive topics are more semantically useful. Following [9] we generate a set of candidate models ranging between 3 and 50 topics. We then plot the exclusivity and semantic coherence (numbers closer to 0 indicate higher coherence), with a linear regression overlaid (Figure 1). Models above the regression line have a “better” exclusivity-semantic coherence trade off. We select the 16-topic model, which has the largest positive residual in the regression fit, and provides

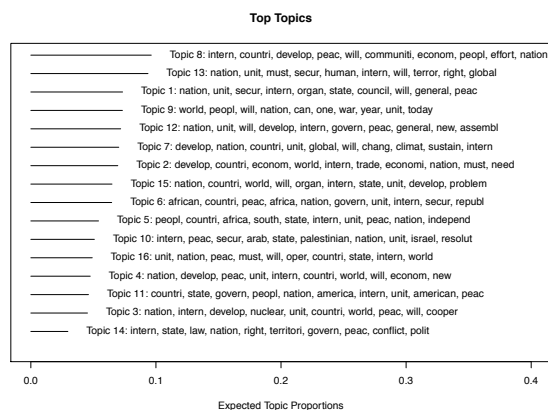


Fig. 2. *Topic quality.* 20 highest probability words for the 16-topic model.

higher exclusivity at the same level of semantic coherence. The topic quality is usually evaluated by highest probability words, which is presented in Figure 2.

B. Topics in the UN General Debate

Figure 2 provides a list of the main topics (and the highest probability words associated with these topics) that emerge from the STM of UN General Debate statements. In addition to the highest probability words, we use several other measures of key words (not presented here) to interpret the dimensions. This includes the FREX metric (which combines exclusivity and word frequency), the *lift* (which gives weight to words that appear less frequently in other topics), and the *score* (which divides the log frequency of the word in the topic by the log frequency of the word in other topics). We provide a brief description of each of the 16 topics here.

Topic 1 - Security and cooperation in Europe.

The first topic is related to issues of security and cooperation, with a focus on Central and Eastern Europe.

Topic 2 - Economic development and the global system.

This topic is related to economic development, particularly around the global economic system. The focus on ‘trade’, ‘growth’, ‘econom-’, ‘product’, ‘growth’, ‘financ-’, and etc. suggests that Topic 2 represent a more traditional view of international development in that the emphasis is specifically on economic processes and relations.

Topic 3 - Nuclear disarmament.

This topic picks up the issue of nuclear weapons, which has been a major issue in the UN since its founding.

Topic 4 - Post-conflict development.

This topic relates to post-conflict development. The countries that feature in the key words (e.g. Rwanda, Liberia, Bosnia) have experienced devastating civil wars, and the emphasis on words such as ‘develop’, ‘peace’, ‘hope’, and ‘democrac-’ suggest that this topic relates to how these countries recover and move forward.

Topic 5 - African independence / decolonisation.




Fig. 4. Comparing Topics 2 and 7 quality. 50 highest probability words contrasted between Topics 2 and 7.

development, with key words including ‘climat-’, ‘sustain’, ‘environ-’, ‘educ-’, ‘health’, ‘women’, ‘work’, ‘mdgs’, ‘peac-’, ‘govern-’, and ‘right’. Therefore, Topic 7 captures many of the issues that feature in the recent Sustainable Development Goals (SDGs) agenda [10].

Figure 4 calculates the difference in probability of a word for the two topics, normalized by the maximum difference in probability of any word between the two topics. The figure demonstrates that while there is a much high probability of words, such as ‘econom-’, ‘trade’, and even ‘develop-’ being used to discuss Topic 2; words such as ‘climat-’, ‘govern-’, ‘sustain’, ‘goal’, and ‘support’ are used in association with Topic 7. This provides further support for the Topic 2 representing a more economic view of international development, while Topic 7 relating to a broader sustainable development agenda.

We also assess the relationship between topics in the STM framework, which allows correlations between topics to be examined. This is shown in the network of topics in Figure 5. The figure shows that Topic 2 and Topic 7 are closely related, which we would expect as they both deal with international development (and share key words on development, such as ‘develop-’, ‘povert-’, etc.). It is also worth noting that while Topic 2 is more closely correlated with the Latin America topic (Topic 11), Topic 7 is more directly correlated with the Africa topic (Topic 6).

II. EXPLAINING THE RHETORIC

We next look at the relationship between topic proportions and structural factors. The data for these structural covariates is taken from the World Bank’s World Development Indicators (WDI) unless otherwise stated. Confidence intervals produced by the method of composition in STM allow us to pick up statistical uncertainty in the linear regression model.

Figure 6 demonstrates the effect of wealth (GDP per capita) on the the extent to which states discuss the two international development topics in their GD statements. The figure shows that the relationship between wealth and the topic proportions




Fig. 5. Network of topics. Correlation of topics.




Fig. 6. Effect of wealth. Main effect and 95% confidence interval.

linked to international development differs across Topic 2 and Topic 7. Discussion of Topic 2 (economic development) remains far more constant across different levels of wealth than Topic 7. The poorest states tend to discuss both topics more than other developing nations. However, this effect is larger for Topic 7. There is a decline in the proportion of both topics as countries become wealthier until around \$30,000 when there is an increase in discussion of Topic 7. There is a further pronounced increase in the extent countries discuss Topic 7 at around \$60,000 per capita. However, there is a decline in expected topic proportions for both Topic 2 and Topic 7 for the very wealthiest countries.

Figure 7 shows the expected topic proportions for Topic 2 and Topic 7 associated with different population sizes. The figure shows a slight surge in the discussion of both development topics for countries with the very smallest populations. This reflects the significant amount of discussion of development issues, particularly sustainable development (Topic 7) by the small island developing states (SIDs). The discussion of Topic 2 remains relatively constant across different population sizes, with a slight increase in the expected topic proportion for the countries with the very largest populations. However, with




Fig. 7. *Effect of population*. Main effect and 95% confidence interval.




Fig. 8. *Effect of ODA*. Main effect and 95% confidence interval.

Topic 7 there is an increase in expected topic proportion until countries have a population of around 300 million, after which there is a decline in discussion of Topic 7. For countries with populations larger than 500 million there is no effect of population on discussion of Topic 7. It is only with the very largest populations that we see a positive effect on discussion of Topic 7.

We would also expect the extent to which states discuss international development in their GD statements to be impacted by the amount of aid or official development assistance (ODA) they receive. Figure 8 plots the expected topic proportion according to the amount of ODA countries receive. Broadly-speaking the discussion of development topics remains largely constant across different levels of ODA received. There is, however, a slight increase in the expected topic proportions of Topic 7 according to the amount of ODA received. It is also worth noting the spikes in discussion of Topic 2 and Topic 7 for countries that receive negative levels of ODA. These are countries that are effectively repaying more in loans to lenders than they are receiving in ODA. These countries appear to raise development issues far more in their GD statements, which is perhaps not altogether surprising.




Fig. 9. *Effect of democracy*. Main effect and 95% confidence interval.

We also consider the effects of democracy on the expected topic proportions of both development topics using the Polity IV measure of democracy [11]. Figure 9 shows the extent to which states discuss the international development topics according to their level of democracy. Discussion of Topic 2 is fairly constant across different levels of democracy (although there are some slight fluctuations). However, the extent to which states discuss Topic 7 (sustainable development) varies considerably across different levels of democracy. Somewhat surprisingly the most autocratic states tend to discuss Topic 7 more than the slightly less autocratic states. This may be because highly autocratic governments choose to discuss development and environmental issues to avoid a focus on democracy and human rights. There is then an increase in the expected topic proportion for Topic 7 as levels of democracy increase reaching a peak at around 5 on the Polity scale, after this there is a gradual decline in discussion of Topic 7. This would suggest that democratizing or semi-democratic countries (which are more likely to be developing countries with democratic institutions) discuss sustainable development more than established democracies (that are more likely to be developed countries).

We also plot the results of the analysis as the difference in topic proportions for two different values of the effect of conflict. Our measure of whether a country is experiencing a civil conflict comes from the UCDP/PRIO Armed Conflict Dataset [12]. Point estimates and 95% confidence intervals are plotted in Figure 10. The figure shows that conflict affects only Topic 7 and not Topic 2. Countries experiencing conflict are less likely to discuss Topic 7 (sustainable development) than countries not experiencing conflict. The most likely explanation is that these countries are more likely to devote a greater proportion of their annual statements to discussing issues around conflict and security than development. The fact that there is no effect of conflict on Topic 2 is interesting in this regard.

Finally, we consider regional effects in Figure 11. We use the World Bank's classifications of regions: Latin America and the Caribbean (LCN), South Asia (SAS), Sub-Saharan Africa




Fig. 10. *Effect of conflict*. Point estimates and 95% confidence intervals.




Fig. 11. *Regional effects*. Point estimates and 95% confidence intervals.

(SSA), Europe and Central Asia (ECS), Middle East and North Africa (MEA), East Asia and the Pacific (EAS), North America (NAC). The figure shows that states in South Asia, and Latin America and the Caribbean are likely to discuss Topic 2 the most. States in South Asia and East Asia and the Pacific discuss Topic 7 the most. The figure shows that countries in North America are likely to speak about Topic 7 least.

The analysis of discussion of international development in annual UN General Debate statements therefore uncovers two principle development topics: economic development and sustainable development. We find that discussion of Topic 2 is not significantly impacted by country-specific factors, such as wealth, population, democracy, levels of ODA, and conflict (although there are regional effects). However, we find that the extent to which countries discuss sustainable development (Topic 7) in their annual GD statements varies considerably according to these different structural factors. The results suggest that broadly-speaking we do not observe linear trends in the relationship between these country-specific factors and discussion of Topic 7. Instead, we find that there are significant fluctuations in the relationship between factors

such as wealth, democracy, etc., and the extent to which these states discuss sustainable development in their GD statements. These relationships require further analysis and exploration.

III. CONCLUSION

Despite decisions taken in international organisations having a huge impact on development initiatives and outcomes, we know relatively little about the agenda-setting process around the global governance of development. Using a novel approach that applies NLP methods to a new dataset of speeches in the UN General Debate, this paper has uncovered the main development topics discussed by governments in the UN, and the structural factors that influence the degree to which governments discuss international development. In doing so, the paper has shed some light on state preferences regarding the international development agenda in the UN. The paper more broadly demonstrates how text analytic approaches can help us to better understand different aspects of global governance.

REFERENCES

- [1] D. Hudson and N. Dasandi, "The global governance of development: development financing, good governance and the domestication of poverty," *Handbook of the International Political Economy of Governance*. Cheltenham: Edward Elgar, pp. 238–258, 2014.
- [2] A. Saith, "From universal values to millennium development goals: Lost in translation," *Development and change*, vol. 37, no. 6, pp. 1167–1199, 2006.
- [3] C. Smith, *Politics and Process at the United Nations: The Global Dance*. Boulder, CO: Lynne Rienner, 2006.
- [4] A. Baturu, N. Dasandi, and S. J. Mikhaylov, "Understanding state preferences with text as data: Introducing the un general debate corpus," *Research & Politics*, vol. 4, no. 2, p. 2053168017712821, 2017.
- [5] M. E. Roberts, B. M. Stewart, D. Tingley, E. M. Airoldi *et al.*, "The structural topic model and applied social science," in *Advances in Neural Information Processing Systems Workshop on Topic Models: Computation, Application, and Evaluation*, 2013.
- [6] D. Mimno, H. M. Wallach, E. Talley, M. Leenders, and A. McCallum, "Optimizing semantic coherence in topic models," in *Proceedings of the conference on empirical methods in natural language processing*. Association for Computational Linguistics, 2011, pp. 262–272.
- [7] D. Newman, J. H. Lau, K. Grieser, and T. Baldwin, "Automatic evaluation of topic coherence," in *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics, 2010, pp. 100–108.
- [8] J. Bischof and E. M. Airoldi, "Summarizing topical content with word frequency and exclusivity," in *Proceedings of the 29th International Conference on Machine Learning (ICML-12)*, 2012, pp. 201–208.
- [9] M. Roberts, B. Stewart, and D. Tingley, "stm: R package for structural topic models 2014," *R package version 0.6*, vol. 21, 2016.
- [10] J. Waage, C. Yap, S. Bell, C. Levy, G. Mace, T. Pegram, E. Unterhalter, N. Dasandi, D. Hudson, R. Kock *et al.*, "Governing the un sustainable development goals: interactions, infrastructures, and institutions," *The Lancet Global Health*, vol. 3, no. 5, pp. e251–e252, 2015.
- [11] M. Marshall and K. Jagers, "Political regime characteristics and transitions, 1800-2003," *Polity IV Project*, 2003.
- [12] N. P. Gleditsch, P. Wallensteen, M. Eriksson, M. Sollenberg, and H. Strand, "Armed conflict 1946-2001: A new dataset," *Journal of peace research*, vol. 39, no. 5, pp. 615–637, 2002.