



HAL
open science

Mapping Forward-Looking Mitigation Studies at Country Level

Claire Lepault, Franck Lecocq

► **To cite this version:**

Claire Lepault, Franck Lecocq. Mapping Forward-Looking Mitigation Studies at Country Level. 2020.
hal-03078474v1

HAL Id: hal-03078474

<https://hal.science/hal-03078474v1>

Preprint submitted on 16 Dec 2020 (v1), last revised 14 Jun 2021 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Mapping Forward-Looking Mitigation Studies at Country Level

Claire Lepault et Franck Lecocq

Corresponding author: claire.lepault@enpc.fr



Centre international de recherche sur l'environnement et le développement

Unité mixte de recherche CNRS - Ecole des Ponts ParisTech - Cirad - EHESS - AgroParisTech

Site web: www.centre-cired.fr Twitter: @cired8568

Jardin Tropical de Paris – 45bis, avenue de la Belle Gabrielle 94736 Nogent-sur-Marne Cedex

Mapping Forward-Looking Mitigation Studies at Country Level

Claire Lepault¹ and Franck Lecocq¹

¹CIREN, AgroParisTech, Université Paris Saclay, CNRS, ENPC, CIRAD, EHESS, 45 bis, avenue de la Belle Gabrielle, 94736 Nogent-sur-Marne Cedex, France

Abstract

We survey the rapidly expanding literature on country-level mitigation pathways using systematic mapping techniques. We build a database of 4691 relevant papers from the Web of Science and Scopus. We analyze their abstracts and metadata using language processing techniques and an innovative topic modeling approach based on two machine-learning models. We find that the country distribution of papers matches countries share of GHG emissions, with very few papers about currently low-emitting countries. Most papers have either 2030 or 2050 as time horizon, the former being more frequent in recent publications. Topic-wise, all areas of mitigation are covered, but major differences across countries reflect both specific circumstances and gaps in the literature, for example around finance.

Keywords: Mitigation, Forward-looking, National, Pathways, Topic Modeling

1 Introduction

The Paris Agreement signed in 2015 emphasizes nationally determined contributions (NDCs) as the building block of global action against climate change, today and over time as countries are expected to ramp up their ambition over time in subsequent NDCs. An increasing number of countries have also communicated long-term low greenhouse gas emission development strategies under Article 4 of the Agreement and/or adopted mid-century mitigation goals.

Though there exist principles for mitigation that are general enough to apply everywhere (e.g., decarbonize electricity, electrify end-uses, promote energy efficiency, enhance carbon sinks), building effective mitigation strategies at country level requires to take into account local economic, social, technological, institutional and cultural circumstances, all the more so that the mitigation objectives are ambitious. To inform such process, country-specific analysis is required.

The academic literature on mitigation pathways at country level, however, remains poorly known. Large number of countries, large number of research teams with diverse backgrounds (energy, macroeconomics, environment, etc.) and lack of institutions that would bring them together all conspire to make it less easy to grasp than the literature on mitigation pathways at global level, which originates from a limited number of research teams worldwide and benefits from well-developed institutions such as the mitigation scenario databases hosted by IIASA or the Integrated Assessment Modeling Consortium. To our knowledge, while global mitigation literature has been extensively surveyed, in particular in the IPCC 4th and 5th Assessment Reports, there exists no survey of the literature at country level despite its relevance for policy-making under the Paris Agreement.

In this paper, we bridge this gap by providing a comprehensive overview of the literature on mitigation pathways at country level. Specifically, we ask: How comprehensive is the geographical coverage of this

literature? Up to what time horizons does it consider mitigation strategies? And what are the main aspects of mitigation it addresses?

We contribute to the growing set of papers that mobilize big data and machine learning to analyse large and fast-growing areas of the academic literature on climate change (1; 2; 3; 4; 5; 6; 7). We first harvest forward-looking mitigation studies at country level from the Web of Science and Scopus databases (see Methods), resulting in a dataset of 4,691 publications. Second, we use language processing methods to extract such information as the case-study location¹, time horizon of the analysis or modeling tool. Third, we analyse the abstracts to identify the main topics addressed by the studies in the database. We improve the topic modeling method proposed by Lamb et al. (1) by reducing subjectivity bias and by fine tuning the parameters of the method to maximize the explanatory power of the topics.

Besides the main results presented in the following, the dataset of papers and topics produced in this research (see Supplementary Material) is of interest of its own as it provides a useful tool for researchers, policymakers and stakeholders to ‘zoom in’ on particular topics and/or countries of interest to inform policy processes and/or identify research gaps. We have striven here to provide the method and results in a clear, transparent and fully reproducible way, with view to making the results easier to communicate (8; 9; 10).

2 Methods

Database construction

To find papers describing mitigation pathway(s) at the national level, we search the academic databases Web of Science and Scopus for references that meet the following three conditions: (i) include the name of a country in the title,² (ii) include "mitigation" or a synonym in the title, abstract, or keywords,³ and (iii) include a year in the period [2025-2100] in the title, abstract or keywords.

Due to differences in coverage of peer-reviewed journals, results of the searches from WoS and Scopus differ significantly, with 944 references that appear only in Scopus and 574 only in WoS. The two selections are then merged into one database, and duplicates are eliminated. The search expressions and the resulting database of 4691 publications, obtained November 14, 2020, can be found in the supplementary materials to this article.

Limiting the search for country names to the title of the reference is based on the observation that papers focusing on national mitigation pathways typically have the name of the country in the title. Conversely, attempts using search equations with country names in the abstract led to harvesting too many irrelevant papers. Finally, adding a year is critical to restricting the search to papers talking about future pathways. Without this condition, the vast literature on current mitigation policies, for instance, would also be embarked in the search.

Additional treatments

This initial database is post-treated using the Pandas library (11) of the Python software. We search country names, demonyms and acronyms in the title to associate each publication to a set of country(ies). When the title of a publication contains two different country names, two entries are created. The first database build in this way contains 4884 rows and is used to analyse the geographical coverage of the database. Actually 153 papers refer in the title to at least two countries. By taking into account each publication only once, the last database represents 4691 publications.

¹that is not provided in the metadata from WoS et Scopus

²The list of countries from the United Nations Statistics Division is available at:

<https://unstats.un.org/unsd/methodology/m49/overview/>

Demonyms and acronyms have been added afterwards by the authors as well as the terms ‘European Union’ and ‘EU’.

³Mitigation synonyms are "low carbon"; "decarboni*ation"; ("carbon" OR "CO2" OR "GHG" OR "greenhouse gas") NEAR/3 "reduc*", the last expression means one of the words in parentheses must be separated by a maximum of three words from the term "reduc*" to select the publication.

Two additional parameters are added to the database. First, we search the title, keywords and abstracts of each publication for horizon year in [2025; 2100]. Then, we search for model names, using the list from a comparative review of scenario modeling tools for national pathways to the Sustainable Development Goals (12). This review (12) lists 80 models, which we searched for under 91 different names. We also checked for the presence of the term "computable general equilibrium" and the associated acronym "CGE" not associated with any of the previous models.

Topic modeling

Overview

The references identified above are analyzed through topic modeling, using the Non-Negative Matrix Factorization classification method (13). The method has two steps. First, we create a corpus of terms and weigh their frequency. The resulting abstract \times term matrix is the Term Frequency-Inverse Document Frequency (TF-IDF) matrix. Second, we search for topics (i.e., weighted sets of terms) such that the product of the matrices describing the weight of each topic in each abstract (abstract \times topic) and the matrix describing the weight of each term in each topic (topic \times term) approximates the TF-IDF matrix as closely as possible. Since the algorithm that generates topics has exogenous parameters, we systemically explore a wide range of parameters and identify the dominant set of topics in this space.

Corpus identification

We identify the corpus from the abstracts in the database. To do so, abstracts are pre-treated: All characters are put in lower case, punctuation signs, connectors and commonly used words are deleted, and words are grouped according to common radicals. Since the country scope and time horizon of each paper is already identified through the search equation, country names and time horizons are deleted. Terms related to mitigation listed in the search equation are also deleted, since by construction of the database each abstract contains at least one of them. Finally, we exclude terms that are either too rare (i.e., that appear in less than 1% of the abstracts) or too frequent (i.e., that appear in more than 95% of the abstracts). The final corpus contains 1300 terms.

TFI-IDF matrix construction

We measure the weight of each term using the *TFIDF* index, defined for each abstract a and each term t as follows (14):

$$TFIDF_{at} = \frac{tfidf(a, t)}{\sqrt{\sum_{i=1}^T tfidf(a, i)^2}} \quad (1)$$

Where

$$tfidf(a, t) = tf(a, t) \times \left[\log \left(\frac{A}{df(t)} \right) + 1 \right] \quad (2)$$

With $tf(a, t)$ the number of occurrences of term t in abstract a and $df(t)$ the number of abstracts containing term t .⁴ The TF-IDF index thus weighs a particular term in a particular abstract if it appears frequently in that abstract but not frequently in the rest of the corpus.

⁴we use the Python Sci-Py library to create the TF-IDF matrix.

Topic identification

We use the Non-negative Matrix Factorization (NMF) method to identify relevant clusters of words (hereafter topics). The algorithm searches for the set of K topics such that the product of the non-negative matrices abstract-topic $W_{A \times K}$ and topic-terms $H_{K \times T}$ best approximates $TFIDF_{A \times T}$. The W matrix can be interpreted as the weight of each topic in each abstract, while the H matrix represents the weight of each term in each topic.

Since the number of topics is selected small relative to the total number of abstracts (typically less than 5%), there is no algorithm of polynomial complexity that converges to a unique solution.⁵ However, one can iteratively converge to local solutions by solving the optimization problem 3, in which $\|\cdot\|_{Fro}$ and $\|\cdot\|_1$ are the Frobenius and L_1 norms respectively, and where $\alpha \geq 0$ and $0 \leq l_1 \leq 1$ are coefficients.

$$\min_{W, H \geq 0} \frac{1}{2} \|X - WH\|_{Fro}^2 + \alpha [l_1 (\|W\|_1 + \|H\|_1) + (1 - l_1) (\|W\|_{Fro}^2 + \|H\|_{Fro}^2)] \quad (3)$$

The first term of 3 ensures convergence of the WH product towards $TF-IDF$, while the second imposes additional constraint on the structure of W and H . The L_1 regularisation (second term) favours the presence of null coefficients in the matrixes, thereby limiting the number of topics each abstract is related to, and limiting the number of terms each topic contains. The minimization of the L_2 regularisation (third term), on the other hand, tends to favor coefficients that are close to one another.

We initialize the NMF algorithm using the Non-Negative Double Singular Value Decomposition method (15; 16). To ensure our results are reproducible, we set the random seed of the algorithm to 1511.

Parameter optimization in the NMF method

The set of topics identified with the NMF method is contingent on the choice of the triplet K (number of topics), α (intensity of regularization relative to the optimization criteria) and l_1 (regularization parameter). We thus build a performance measure for each set of topic, and then select the triplet (K, α, l_1) that produces the highest ranking set. To our knowledge, this is the first time that the algorithm below is used to select the best triplet (it has previously been used to select K (17)).

The performance measure, named "coherence", of a given topic is assessed by looking at how similar the semantic environments of each of the terms that compose the topics are. The performance of a set of topic is the average of the coherence measure of each individual topic.

Following O'Callaghan et al. (17), we use the Word2vec word embedding algorithm (18; 19) to produce a vectoral representation of the semantic environment of each term within the corpus of abstracts. Word2vec is a two-layer neuronal network which maps words into vectors that account for the textual environment of the word. Words that share similar contexts are characterized by similar multi-dimensional vectors.

We use the Skip-Gram method to train the neural network. This approach seeks to predict the semantic context of a term. The error is computed based on the corrected prediction of the words surrounding this term. The coherence of each topic coherence is then the mean of the pairwise cosine similarities between the terms that characterize the topic. Precisely, for a topic k , the coherence index $TCW2V_k$ is computed as follows:

$$TCW2V_k = \frac{1}{\binom{N}{2}} \sum_{j=2}^N \sum_{i=1}^{j-1} similarity(wv_{tj}, wv_{ti}) \quad (4)$$

Where N is the number of terms that we choose to characterize each topic,⁶ wv_{kj} is the vector associated to term j characterizing topic k , wv_{ki} is the vector associated to term i characterizing topic k and $similarity(A, B)$ is the cosine similarity of vectors A and B , defined as:

⁵NMF methods are still preferred to principal component analysis methods (for which such algorithm exists) for textual analysis because in the former, several topics can apply to a single abstract.

⁶For each topic, we retain the five terms with the highest ranking score in the topic - terms matrix H .

$$\text{similarity}(A, B) = \frac{\sum_{i=1}^D A_i B_i}{\sqrt{\sum_{i=1}^D A_i^2} \sqrt{\sum_{i=1}^D B_i^2}} \quad (5)$$

We assign each set K with the mean of the scores of each individual topic within K .

We then systematically search the set of triplets (K, α, l_1) in the range $[2, 40] \times ([0.00, 0.31] \cup [0.40, 1.00]) \times [0.0, 1.0]$ (with increments of 1, (.01; 0.1) and 0.1 respectively) for the highest-scoring set of topics.

The topic coherence score maximizing set is (39, 0.1, 0.9).

Relationship across topics

To visualize how topics relate to each other (Figure A6), we use *LDAvis* (20), a system initially developed to explore topic-term relationships in a fitted Latent Dirichlet Allocation (LDA) model. The intertopic distance is based on the Jensen Shannon divergence calculated from the H matrix coefficients characterizing the topic-terms relationships. Principal Components Analysis (PCA) then projects the set of intertopic distances onto two dimensions. In online supplementary material, the whole visualization is available and represents the individual terms that are the most useful for interpreting each topic. In particular, it enables to look at the corpus-wide frequency of a given term as well as the topic-specific frequency of the term.

Relationship between abstracts and topics

To define if a paper is related to a topic, we normalize the W matrix such as the sum of the coefficients of each row is equal to one. By this way, an abstract is decomposed according to the 39 topic-dimensions, which enables to compare the coefficients from one topic-abstract combination to another. We thus apply the following transformation to the $W_{A \times K}$ matrix:

$$W_{ak}^* = \frac{W_{ak} \times \sum_{i=1}^T H_{ki}}{\sum_{k=1}^K (WH)_{ak}} \quad (6)$$

First, we realize an association of topics to publications based on normalized coefficients of the W matrix. We choose the threshold $s = 0.02$ (as per Lamb et al. (1)) based on the coefficients of the matrix $W_{A \times K}^*$, according to the specification :

$$\begin{aligned} \text{abstract } a \text{ is related to topic } k & \text{ if } W_{ak}^* > s \\ \text{abstract } a \text{ is not related to topic } k & \text{ if } W_{ak}^* \leq s \end{aligned} \quad (7)$$

To check how relevant that classification is, we build another classification based on the title of the paper. Specifically, we associate a paper with abstract a to a topic k if $W_{ak}^* > 0.02$ and if at least one of the 5 terms characterizing topic k is in the title of the paper with abstract a . Figure A5 presents the number of papers per topics in each classification. As the Figure illustrates, these distributions are similar, the one on the bottom being scaled down from the one on top. Since the presence of a word characterizing a topic in the title of a paper is a strong indication that the paper is indeed related to that particular topic, the comparison between the two distributions is a good indication that our initial classification is relevant.

3 Results

Papers are distributed in proportion to countries GHG emissions

Figure 1 maps the countries in which we found forward-looking mitigation literature. Overall, 136 countries (plus the European Union) appear in the database. However, the geographic distribution of papers is particularly skewed towards China, which represents (24.3%) of all papers in the database. Distant second are the US (9.0%), followed by the UK (6.0%), the EU as a region (5.3%) (excluding papers related to individual EU Member States), and India (4.9%). Regionally, nearly half of the papers (47.8%) focus on Asia, nearly a third on Europe (30.0% including 24.7% for individual European countries, and 5.3% for the EU as a whole), a sixth on the Americas (18.1%, including 7.1% on Latin America and 11.0% on North America). Other continents account for less than 5% of the publications. Africa, in particular, is very poorly represented (4.8%), with all but 7 countries with less than 10 papers, and nearly half with no paper at all.




Figure 1: Country distribution of studies in the database. Papers related to the EU as a whole are not represented in this map (Source: Authors).

The share of each country in the database appears well correlated with the share of each country in global GHG emissions (Figure 2). This is not surprising since the larger the problem, the more likely it is to attract the attention of the (domestic and foreign) research community, either *suo motu* or at the request of governments or of other interested parties. A prominent exception to the proportionality rule amongst large emitters is the UK, which is much more represented in the literature than it is in global GHG emissions. This might translate the strength of the UK research community on mitigation, and/or the fact that with the adoption of the Climate Change Act in 2008, the UK has a longer history of national climate policies than most high-emitting countries. Another exception is Russia, which on the contrary has significantly fewer papers than its share of emissions would suggest. Likewise, this might reflect a research community that has invested less on mitigation and/or translate public policies with other priorities. Amongst countries with middle- or low-emissions, OECD countries, particularly in Europe (e.g., Finland, Switzerland, Sweden) have more papers in the database than their shares of GHG emissions would suggest. Developing countries, on the other hand, tend to be closer to the line or below.




Figure 2: Country share of papers in the database (y-axis) against country share in World GHG emissions (x-axis). Oblique line the $x=y$ line. Sources: Authors, GHG emissions data (including emissions from land use, land-use change and forestry) from the CAIT database for the year 2016.

The tail of the distribution is also relevant for policy making. Of a total of 197 Parties to the UN Framework Convention on Climate Change, 143 have less than 10 papers in the database, 127 less than 5, and 65 have none. In their survey of urban climate mitigation case studies, Lamb et al. (1) similarly found a very uneven distribution of papers by country. While it is difficult to determine a threshold below which the number of forward-looking publications on mitigation would be "insufficient" to inform policies, 10 papers or less (to be compared with 39 major topics in the database, see below) leaves little chance that even the different sectoral aspects of mitigation be adequately covered. Policymakers and stakeholders in Africa, in particular, have for the most part scant scientific literature to rely on, despite rapidly increasing emissions. Informing strategies to limit growth in GHG emissions (and ultimately start reducing them) while continuing to other development goals needs a major shift in the focus of research towards the continent.

Paris Agreement has spurred increased attention to 2030 time horizon

We set the time horizon of each paper in the database as the highest number in the [2025, 2100] range found in the abstract or title (see Methods). The distribution of time horizons (Figure 3.b) presents two very clear peaks in 2030 and 2050 respectively, each accounting for 34% of all papers. Horizons beyond 2050 represent 14% of all papers, a major difference with the literature on mitigation at the global level, in which the 2100 time horizon is the norm. This translates a difference in research questions. Forward-looking mitigation studies at the global level are typically conducted to assess mitigation scenarios against long-term temperature goals, whereas forward-looking studies at the national level have typically the objective to assess more detailed policy packages. For that purpose, 2050 is already a long time horizon.




Figure 3: Papers in the database by time horizon and publication year (left); and by region and time horizon (right). On panel a, the first bar accounts for all papers published up to 2002. On panel b, cross-hatched bar plots indicate aggregation of papers with different time horizons.

Figure 3.a presents papers by publication year and by time horizon. It shows a rapid expansion of the forward-looking mitigation literature at country level over the past two decades, with a clear inflexion point in 2014: The annual increment of publications is markedly higher in the 2014-2020 period than in the 2007-2013 period. Such inflexion does not appear when looking at the climate change literature as a whole. When doing so, on the contrary, Callaghan et al. (2) find annual increments more or less constant over the whole 2007-2020 period. The timing of the inflexion point (2014) suggests that the negotiation and adoption of the Paris Agreement, with its strong emphasis on national-level mitigation, may have spurred increased interest in national-level mitigation.

This hypothesis is further supported by the fact that 2014 also marks an inflexion in the distribution of time horizons across papers. As can be seen in Figure 3.a, the share of papers with a time horizon up to 2030 was decreasing before 2014. But from 2014 onward, this trend is reversed. Since 2030 is the time horizon of nearly all the Intended NDCs communicated in 2015 (and of most the NDCs to date), this finding suggests again that the signature of the Paris Agreement has spurred increased interest in mitigation at the national level. The adoption of the Sustainable Development Goals (SDGs) in 2015, also with a 2030 as a

time horizon, may also have played a role, though mitigation is only one of the 17 SDGs.⁷

Finally, Figure 3.b shows that the distribution of time horizons differs by region⁸. Europe and North America represent more 62% of the literature with a 2050 time horizon, against 25% of the literature with a 2030 time horizon. Conversely, Asia represents 68% of the literature with a 2030 time horizon, against 40% up to 2050. This suggests that research in Europe and North America is already focused on mid-century time horizons, consistent with the mid-century mitigation strategies that several European countries and the EU have adopted. Whereas the focus in Asia would be more on the conditions under which NDCs can be achieved by 2030. If this explanation is correct, then we should soon see an increase in the share of papers about 2050 and 2060 time horizons in Asia following the recent announcement of the long-term mitigation objective by China.

Studies offer a comprehensive but uneven coverage of major mitigation issues

To capture the content of the papers we have identified, we use an original topic modeling approach based on two machine learning models. Here a 'topic' is a set of words that best describe the content of the title and of the abstract of a paper. All country names, associated demonyms and years are removed before searching for topics since they are analysed separately. As a result, they cannot be part of the set of terms that characterize topics. We search for topics using a non-negative matrix factorization (NMF) algorithm that is able to learn parts of faces and semantic features of text (13). This machine learning model has already been used to map climate-related literature (1; 2). However, a key challenge when using the NMF algorithm resides in the selection of exogenous parameters, chief among which the number of topics. To reduce the risk of arbitrariness, we assess the similarity of the semantic environment of the words characterizing the topic with the actual abstracts using a two-layer neural network (18; 19) (see Methods).

Table 1 presents the resulting list of topics⁹, with the five most important words for each, ranked by number of papers attached. Nearly all the papers in the database (4687 out of 4691) are related to topic No.1, characterized by the words "policy - development - economic - countri - use". This is not surprising, since papers on mitigation scenarios at country level typically discuss policy implications, including in the abstract. More interesting is the fact that the corpus is then split nearly in half between papers related to topic No.2 (*Climate Change*) and papers related to topic No.3 (*Energy Efficiency*). The two ensembles are largely disjointed, as can be seen from the mapping of the strength of the pairwise combinations of topics (Figure A7). Papers associated with topic No.2 (*Climate Change*) tend to be also associated with topics such as *Drought, Flood, Water, Crop Yield, Forest, Land Use, Agriculture* or *Air Pollution*. Whereas papers associated with topic No.3 (*Energy Efficiency*) tend to be associated with topics such as *Hydrogen, Steel/Iron, Nuclear, Peak, Oil, CCS, Wind/Solar* or *Buildings*. The other topics can be organized in five groups: (i) methods (*Scenarios* and *Systems*), (ii) policies (e.g., *Costs* or *Targets/INDC*), (iii) sectors ; (iv) air pollution; and (v) climate change impacts (*Drought, Flood* and *Crop Yield*). The latter are not all primarily about mitigation, as the search equation also picks forward looking impact assessment or adaptation study at the national level that have in the abstract the word "mitigation" or a demonym.

⁷A breakdown of papers by publication year and by region (Figure A1) shows a rapid increase in the share of papers on Asian countries since 2016, mostly driven by China. This is consistent with the hypothesis that the Paris Agreement spurred increased interest in national level mitigation, given the prominent attention given to the Chinese NDC.

⁸When a paper is about two countries in different regions (e.g., China and US), it is attributed the region of the first country to appear in the list. There are only 72 such papers in the database so we consider the potential bias negligible.

⁹Topics are characterized by word stems rather than by full words. The "title" of the topic is our work.

Topic	Terms	Category	Sector	Relevant IPCC WG3 chapter	T0.02	T0.02Title
Policy-Dvlpt-Eco	polici_develop_econom_countri_use	context			4,687	1,324
Climate change	climat_chang_temperatur_futur_project	context			2,000	1,038
Energy efficiency	energi_effici_consumpt_save_demand	context			1,732	1,018
Scenario	scenario_model_bau_refer_three	method			1,428	553
Consumption	intens_structur_consumpt_factor_growth	other		5	936	221
Electricity	electr_generat_demand_grid_suppli	sectoral	energy	6	904	471
Power	power_generat_plant_sector_capac	sectoral	energy	6	797	416
Costs	cost_abat_option_benefit_margin	policy		12	761	229
Target/INDC	target_achiev_indc_meet_ndc	policy		13	706	200
Air Pollution	air_pollut_pm2_qualiti_health	other		17	674	252
Fuel	fuel_fossil_diesel_altern_biofuel	sectoral	energy	6	586	157
Transport	transport_sector_road_passeng_freight	sectoral	transport	10	567	298
System	system_model_transit_integr_pathway	method			543	339
Renewable Energies	renew_energi_sourc_share_res	sectoral	energy	6	540	359
Cement	industri_cement_sector_product_process	sectoral	industry	11	497	301
Technology	technolog_advanc_deploy_low_clean	technology		16	421	168
Land Use	land_use_soil_area_chang	sectoral	landuse	7	420	253
Permit market	price_et_market_trade_polici	policy		13	419	241
Vehicle	vehicl_fleet_car_hybrid_passeng	sectoral	transport	10	387	227
Coal	coal_fire_plant_natur_phase	sectoral	energy	6	360	126
Buildings	build_residenti_stock_sector_construct	sectoral	buildings	8	356	180
Agriculture	agricultur_food_product_livestock_farm	sectoral	landuse	7	351	188
Bioenergy	biomass_bioenergi_biofuel_residu_wood	sectoral	energy	7	345	160
Urban	urban_citi_area_popul_develop	policy		8	331	197
Wind/Solar	wind_solar_capac_instal_photovolta	sectoral		6	324	107
Forest	forest_sequestr_wood_sink_stock	sectoral	landuse	7	314	188
Crop yield	crop_yield_soil_wheat_fertil	sectoral	landuse	7	302	113
Water	water_resourc_basin_river_irrig	impacts		WG2	291	133
Flood Risk	flood_risk_coastal_sea_disast	impacts		WG2	284	127
Heat Pump	heat_pump_district_cool_boiler	sectoral	buildings	9	272	94
CCS	ccs_storag_captur_geolog_plant	technology		6	251	127
Waste	wast_landfil_solid_municip_treatment	sectoral	waste	11	222	73
Tax	tax_revenu_polici_equilibrium_model	policy		13	213	117
Oil	oil_product_crude_natur_export	sectoral	energy	6	200	72
Peak	peak_around_reach_earlier_non	policy		3	173	65
Nuclear	nuclear_plant_power_mix_new	sectoral	energy	6	160	87
Steel/Iron	steel_iron_product_materi_save	sectoral	industry	11	105	70
Hydrogen	hydrogen_cell_chain_produc_product	sectoral	energy	6	94	60
Drought	drought_precipit_sever_frequenc_index	impacts		WG2	89	39

Table 1: Description of Topics. For each paper in the database, the NMF algorithm provides the weight of each topic. We consider that a particular paper is related to a particular topic if the normalized weight of the topic in the paper is greater than 0.02 (column T0.02). To check the robustness of this method, we consider a second attribution process in which papers are linked to a topic if at least one of the five top words (the words most important in the topic) figures in the title of the paper (column T.02Title). We find that both methods produce similar distribution of papers across topics (the latter being more selective than the former) (Figure A5).


Using the outline of the IPCC Working Group III 6th Assessment Report as a rough mapping of the topics associated with mitigation (1, column 5), one can see that the forward-looking mitigation papers at the national level cover all IPCC WGIII AR6 sectoral chapters (6 to 11) as well as issues related to demand (5), policies (13) and innovation (16). The absence of international policies (Chapter 14) as a major topic is understandable since the search equation focuses on mitigation at the national level. The absence of a topic related to finance (Chapter 15) on the other hand confirms anecdotal evidence that few forward-looking national mitigation pathways have been analyzed along that lens so far. Finally, the lack of a standalone topic dedicated to SDGs (Chapter 17) may be related to the fact that if individual SDGs are discussed in the abstracts, it may be in a diffuse way that does not get picked up in a topic (except for *Air Pollution*). Among the sectors that are represented there is considerable imbalance: energy is the one with the largest number of related papers (27%) followed by LULUCF (9%), while the other sectors are much less represented. Though the attribution of topics to particular sectors may be debatable in some cases (for example, bioenergy could also be related to LULUCF), the general imbalance between sectors appears robust.

Topics and modeling tools are mainly identified in Asian studies

The distribution of countries for each topic reflects the overall distribution of countries in the database (see Figures A3 and A4). At one end of the spectrum, China has the largest number of papers for all topics except heat pump (preceded by the UK), nuclear (preceded by Japan, the UK and South Korea) and hydrogen (preceded by Germany, Japan and the UK). At the other, African countries appear only once in the top 5 for a topic (Ethiopia for *drought*). There are, however, differences across topics. Forward-looking mitigation studies about industrial sectors (*cement* or *steel-Iron*) have been conducted predominantly for China, while the distribution of papers is much more balanced across countries for topics such as *renewable energy* or *buildings*. The imbalance in research across countries in the *urban* topic is particularly surprising, since urban development issues are not confined to China. It is however, consistent with the finding of Lamb et al. (1) that urban case studies in China overwhelmingly dominate the literature.

The distribution of topics per country provides, in each country, a mapping of the issues identified as most important by the academic community when looking at mitigation. Each line in Figure 4 maps the distribution of topics in the forward-looking mitigation literature for a particular country. Since the colors of the cells refer to the share of each topic in the country's overall number of publications, lines are comparable. Patterns for countries emerge. For instance, literature on Indonesia is largely distributed with three major topics *Energy Efficiency*, *Target/INDC* and *Land Use* plus eight topics around power generation, forest and oil. Whereas the literature on Poland is more focused on *Power*, *Costs*, *Renewable Energies* and *Coal*. Reading the Table vertically similarly provides a view of the importance of topics per country, with again clusters emerging around e.g., *Oil*, *Nuclear* or *Land Use*.

Finally, we attempt to analyze the methods used in the papers to study mitigation at country level. This is not easy given the limited amount of information present in the metadata. We focus on models, checking metadata against a database of 80 scenario modeling tool for national pathways to the Sustainable Development Goals (12) (see Methods). We identify model names in only 13% of the abstracts (608). Compared with the country coverage of the overall dataset (Figure 1), Asia is even more represented (61% in this subset. For example, Thailand is three times more represented than in the general database (7% against 2.2%) (see Figure A8). At the other end of the spectrum, Africa is even less present (4%) with only 10 countries represented. Model-wise, General Equilibrium Models is the most common category of models in the corpus. The three individual models that dominate, on the other hand, LEAP, TIMES and MARKAL are all bottom-up. They are highly used for Asia and Europe (Figure A9). Unsurprisingly, these three models mainly used for energy-related questions (top 3 for topics *Energy Efficiency*, *Electricity*, *Power*, *Fuel*, *Transport*, *Vehicle*, *Wind/Solar*) (Figure A10). Finally, it is interesting to note that the Japanese AIM model is present in 65 publications (of which 59 in Asia), illustrating the importance of regional clusters. Although limited to the arguably small sample of papers that name their model (or model type) in the abstract, these findings nonetheless emphasize again the inequalities between countries.



2030 time horizon created by the (I)NDCs in the Paris Accord. Topic-wise, mitigation pathways have been developed with a broad range of focus encompassing all dimensions of mitigation, save for finance issues, that are lacking. Energy and to a lesser degree LULUCF activities are, however, very dominant relative to other sectors. Mapping topics at country level reveal current priorities but also gaps. From a methodological point of view, the paper builds upon and improve on Lamb et al. (1) by providing a systematic way to maximize the accuracy of the topic modeling.

This paper has several limitations. First, the term mitigation (or its demonyms) that we use in the search equation harvests too broad a set of papers, since papers about impacts and adaptation to climate change may still refer to mitigation in the abstract. Secondly, by construction, metadata (abstract, title and keywords) only embark a fraction of the content of the paper, potentially making classification of papers difficult. General sentences providing context, for example about climate change, may be easily recognizable as such in a full paper as they would only represent a tiny fraction of the overall word count, whereas in an abstract they may be confused with a substantive result of the paper. The ubiquitousness of the *climate change* topic is a demonstration of that risk. Third, despite instructions by Journals, abstracts are written in very different ways across papers. They not only differ in length but also in content, notably on the relative share between context and summary of the result. For the purpose of textual analysis, abstracts that are as close as possible to the method and key findings of the paper are preferable, though that may come at the expense of readability.

Finally, from a policy point of view, the large differences in coverage (both geographical and sectoral) observed above suggest that the academic community is not yet equipped to inform mitigation strategies in a large range of countries. Though these countries are not the ones with the most important emissions (at least today), they may also be the ones where opportunities to act early to avoid emissions growth in the future are most important. Obstacles are multiple, including lack of domestic research capacity, lack of data, lack of interest or incentive for foreign research teams to work on other national contexts. Yet providing countries with minimal forward-looking capacity appears both desirable from an equity point of view, and useful from a climate perspective, since ultimately better knowledge and anticipation is more likely to result into action.

References

- [1] Lamb, W. F., Creutzig, F., Callaghan, M. W. & Minx, J. C. Learning about urban climate solutions from case studies. *Nature Climate Change* **9**, 279–287 (2019).
- [2] Callaghan, M. W., Minx, J. C. & Forster, P. M. A topography of climate change research. *Nature Climate Change* **10**, 118–123 (2020).
- [3] Aleixandre-Benavent, R., Aleixandre-Tudó, J. L., Castelló-Cogollos, L. & Aleixandre, J. L. Trends in scientific research on climate change in agriculture and forestry subject areas (2005–2014). *Journal of cleaner production* **147**, 406–418 (2017).
- [4] Belter, C. W. & Seidel, D. J. A bibliometric analysis of climate engineering research. *Wiley Interdisciplinary Reviews: Climate Change* **4**, 417–427 (2013).
- [5] Haunschild, R., Bornmann, L. & Marx, W. Climate change research in view of bibliometrics. *PLoS One* **11** (2016).
- [6] Li, W. & Zhao, Y. Bibliometric analysis of global environmental assessment research in a 20-year period. *Environmental Impact Assessment Review* **50**, 158–166 (2015).
- [7] Wang, B., Pan, S.-Y., Ke, R.-Y., Wang, K. & Wei, Y.-M. An overview of climate change vulnerability: a bibliometric analysis based on web of science database. *Natural Hazards* **74**, 1649–1666 (2014).
- [8] Donnelly, C. A. *et al.* Four principles to make evidence synthesis more useful for policy (2018).

- [9] Haddaway, N. R. & Macura, B. The role of reporting standards in producing robust literature reviews. *Nature Climate Change* **8**, 444–447 (2018).
- [10] Minx, J. C., Callaghan, M., Lamb, W. F., Garard, J. & Edenhofer, O. Learning about climate change solutions in the ipcc and beyond. *Environmental Science & Policy* **77**.
- [11] Wes McKinney. Data Structures for Statistical Computing in Python. In Stéfan van der Walt & Jarrod Millman (eds.) *Proceedings of the 9th Python in Science Conference*, 56 – 61 (2010).
- [12] Allen, C., Metternicht, G. & Wiedmann, T. National pathways to the sustainable development goals (sdgs): A comparative review of scenario modelling tools. *Environmental Science & Policy* **66**, 199–207 (2016).
- [13] Lee, D. D. & Seung, H. S. Learning the parts of objects by non-negative matrix factorization. *Nature* **401**, 788–791 (1999).
- [14] Salton, G. & Buckley, C. Term-weighting approaches in automatic text retrieval. *Information processing & management* **24**, 513–523 (1988).
- [15] Boutsidis, C. & Gallopoulos, E. Svd based initialization: A head start for nonnegative matrix factorization. *Pattern recognition* **41**, 1350–1362 (2008).
- [16] Belford, M., Mac Namee, B. & Greene, D. Stability of topic modeling via matrix factorization. *Expert Systems with Applications* **91**, 159–169 (2018).
- [17] O’callaghan, D., Greene, D., Carthy, J. & Cunningham, P. An analysis of the coherence of descriptors in topic modeling. *Expert Systems with Applications* **42**, 5645–5657 (2015).
- [18] Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S. & Dean, J. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, 3111–3119 (2013).
- [19] Mikolov, T., Chen, K., Corrado, G. & Dean, J. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781* (2013).
- [20] Sievert, C. & Shirley, K. Ldavis: A method for visualizing and interpreting topics. In *Proceedings of the workshop on interactive language learning, visualization, and interfaces*, 63–70 (2014).

Appendices

Supplementary Figures




Figure A1: Papers in the corpus by region and publication year. For ease of reading, all papers published up to 2002 are associated to year 2002.




Figure A2: Comparison of the regional distribution of papers according to horizon year and publication year;




Figure A3: Top 5 countries representing each topic. From left to right, and from top to bottom, topics are presented in the descending order of their representation in the whole database.




Figure A4: Top 5 countries representing each topic. From left to right, and from top to bottom, topics are presented in the descending order of their representation in the whole database.




Figure A5: Distribution of topics when papers : are are selected with $t = 0.02$ (a); are selected with $t = 0.02$ and contains in the title one of the five term characterizing the topic (b)




Figure A6: Intertopic Distance Map (via multidimensional scaling)




Figure A7: Heatmap of the papers common across topics (each cell corresponds to the number of common papers to both topics (row and column) on the total number of papers related to the topic-column)




Figure A8: Country distribution of studies mentioning at least one model in the abstract. As the map is at the country level, papers related to the European Union are not represented in this figure.




Figure A9: Histogram of the number of model-study associations




Figure A10: Histogram of the number of model-study association per topic. Topics, from left to right, are presented in the descending order of their representation in the whole database. For each topic, only the three most represented models are in distinct colors.

Supplementary Table

Number of topics	Number of papers	
	Threshold 0.02	Threshold 0.02 and title check
0	0	367
1	31	1104
2	238	1410
3	648	1050
4	956	539
5	972	177
6	832	37
7	517	5
8	283	2
9	136	0
10	56	0
11	18	0
12	2	0
13	2	0

Table A 1: Number of topics characterizing papers. For each paper in the database, the NMF algorithm provides the weight of each topic. We consider that a particular paper is related to a particular topic if the normalized weight of the topic in the paper is greater than 0.02 (column Threshold 0.02). To check the robustness of this method, we consider a second attribution process in which papers are linked to a topic if at least one of the five top words (the words most important in the topic) figures in the title of the paper (column Threshold 0.02 and title check).