GEORGE MASON UNIVERSITY AI STRATEGIES TEAM
& THE STIMSON CENTER
PRESENT

2023 GLOBAL ARTIFICIAL INTELLIGENCE INFRASTRUCTURES REPORT

Authors:

J.P. Singh
Amarda Shehu
Caroline Wesson
Manpriya Dua

With a Foreword from David A. Bray

2023 Global Artificial Intelligence Infrastructures Report*

J.P. Singh[1], Amarda Shehu[2], Caroline Wesson[1] and Manpriya Dua[2]

[1]Schar School of Policy and Government, George Mason University
[2]Department of Computer Science, George Mason University

Foreword from David Bray[3]
[3]Distinguished Fellow, The Stimson Center

Suggested citation: Singh, J.P., Amarda Shehu, Caroline Wesson, and Manpriya Dua. *The 2023 Global Artificial Intelligence Infrastructures Report.* With a Foreword from David Bray. AI Strategies Team and the Institute for Digital Innovation, George Mason University, and the Stimson Center, Washington DC. August 2023.

*This research is supported by a $1.389 million grant from the Minerva Research Initiative. Please send queries on this report to aipolicy@gmu.edu

Contents

List of Tables

List of Figures

1 The AI Strategies Team at George Mason

The AI Strategies team at George Mason provides evidence for how cultural values and institutional priorities shape artificial intelligence (AI) infrastructures in national and global contexts, in order to better understand the effects of comparative AI contexts for security.

AI Strategies is funded by a three-year, $1.39 million grant that was awarded to George Mason University to study the economic and cultural determinants for global artificial intelligence (AI) infrastructures—and describe their implications for national and international security. The team began work on the project on April 15th 2022. This is the first yearly AI infrastructure report.

The grant was awarded by the Department of Defense's reputed Minerva Research Initiative, a joint program of the Office of Basic Research and the Office of Policy that supports social science research focused on expanding basic understanding of security.

Team Members

Currently the team is comprised of three social scientists, two computer scientists, a computer science PhD student, and three political science, public policy, and philosophy doctoral and Master's students. Given this diverse team, the research reflects unique analytical creativity: the team members have worked together for over a year on this research, and learned to build off one another's strengths to understand the landscape of national AI infrastructures and how to apply NLP methodologies to empirically base their comparisons and contextualize the subject matter and country expertise. Read our research team's biographies in Section 8.

2 Foreword: Ensuring We Build the Right Foundation to Evaluate Trust in AI and Societies

Dear Readers,

As a field, Artificial Intelligence (AI) has been around since the mid-20th century. In 1955, U.S. computer scientist John McCarthy coined the term. Later in 1959, McCarthy collaborated with Marvin Minsky to establish the Artificial Intelligence Project, nowadays MIT's CSAIL (Computer Science and Artificial Intelligence Laboratory). In parallel to McCarthy and Minsky, U.S. political scientist Herbert Simon completed a PhD in 1943 exploring decision-making in administrative organizations and pursued research that later influenced the fields of computer science, economics, and cognitive psychology. In 1957, Simon partnered with Allen Newell to develop a General Problem Solver separating information about a problem from the strategy required to solve it.

All four individuals – McCarthy, Minsky, Simon, and Newell – would go on to receive the ACM A.M. Turing Award during their respective careers. In the almost six and a half decades that followed, AI research developed several flavors of systemic approaches to include: Logical Reasoning and Problem-Solving Algorithms, Expert Systems, Statistical Inferences and Reasoning, Decision Support Systems, Cognitive Simulations, Natural Language Processing, Machine Learning, Neural Networks, and more.

Though AI has many subcategories and has had many flavors of approaches since the 1950s, within the last few years, a subset of Neural Networks built on the transformer architecture have revolutionized natural language processing and given rise to what are now known as Large Language Models (LLMs). Just in the last year, LLMs such as ChatGPT and variants, have activated significant public interest, excitement, and anxiety with regards to the future of AI. While the full extent of the public, business, community, and individual value of LLMs remains to be seen, the ability of these models to provide responses to effective engineered prompts regarding the generation of predictive text, synthesized images, as well as the full gamut of multimedia audio and even video outputs has captured the public zeitgeist.

I. A valuable compass reading as to where different nations have decided to steer approaches to AI

Pundits globally have indicated both excitement and concerns about whether machines may be able to perform work previously thought only performable by humans as well as whether they may be able to produce content and interactions that appear human. It is precisely at this moment that this 2023 Global Artificial Intelligence Infrastructures Report by J.P. Singh, Amarda Shehu, and their doctoral students is so prescient. By bridging together multiple fields, including the best of computer science, economics, political science, and public policy, in a collaborative manner akin to the best work of Herbert Simon – Singh and

Shehu have produced a valuable compass reading pointing to where different nations have decided to steer their approaches to AI for the future ahead. Their report presents both rigorous and much needed insights that demystify some of the current fervor around the future AI and societies, namely:

- First, the report shares convincing evidence that humanity's AI-associated future will not be set by just the United States and China alone – there exist different AI strategies being pursued by multiple nations beyond just these two large nations, with different objectives and proposed paths outlined in these national AI policies.
- Second, while there is no singular grand strategy across the fifty-four national AI plans analyzed in this report – Singh and Shehu find similar choice elements in the national strategies analyzed. The researchers dub these similar choice elements a collective 'AI Wardrobe,' a term coined by Caroline Wesson, one of the doctoral students in the team, to relate the various choices each country can make in assembling a tailored AI national strategy outfit.
- Third – country clusters are apparent among the different national strategies that were analyzed for this report – to include the European Union, East Asia, Spain leading a Latin America cluster, and the United Kingdom leading a British influence cluster. Whether or not these clusters will result in closer AI-related business interactions, nation-to-nation civil relations, and geopolitical ties amongst the countries more closely aligned with regards to their national AI strategies represents a crucial area to watch both now and in future.

II. Building the necessary foundation an interdisciplinary mix of fields to tackle Trust, AI, and Societies

Juxtaposed against the global zeitgeist regarding AI, this important 2023 report exists amid a deeper milieu of important questions regarding trust within and across nations. In October 2017, the Pew Research Center found that less than forty-five percent of residents living in the United States under the age of twenty-five years old thought capitalism was in their opinion a good force in society. Contemporary studies at the time also found declining levels of trust among a similar age demographic in the essentialness of living in a democracy – not just for the United States but also for Sweden, Australia, the Netherlands, New Zealand, and the United Kingdom. Together these global trends blend to create a central question – namely:

Can nations invoke strategies that result in Trust in AI and societies? – and a corollary: Can nations encourage Trust in AI and societies, while facing growing distrust in their economic and political systems?

Readers should note that trust can be defined as the willingness to be vulnerable to the actions of an actor not directly controlled by you – a definition that works for both human and AI actors. Multiple studies have established that the antecedents of trust include the perceptions of benevolence, competence, and integrity of the actor to an individual. If perceptions of these three antecedents are positive, then trust is more likely. If perceptions are negative or absent, then trust is less likely.

The clustering of similar choice elements in this report, specifically the set of elements that comprise the report's described "AI Wardrobe", represent an important tool for leaders in the public and private sectors to assess if a national AI strategy has the requisite elements to address challenging questions of improving Trust in AI and societies.

Cumulatively, this question of Trust in AI and societies represents an essential one for nations' AI strategies with regards to their expressed objectives. In terms of expressed objectives, though LLMs and their outputs have captured the current public consciousness of 2023, there are so many more outcomes for which AI can be employed by nations, communities, and networked groups of people working to shared outcomes beyond just generative content. Readers are invited, after seeing the analysis and results in the report, to consider more expansive objectives for AI and societies, to include exploring how:

- Can AI improve human understanding of decisions we need to make now?
- Can AI help improve understanding the impact of our decisions (or lack thereof) on possible local and global futures?
- Can AI help improve human collaborations across sectors and geographies, potentially tipping and cueing humans that there are other humans with similar projects underway?
- Can AI help improve identification and reconciliation of misaligned goals and incentives – be they community, regional, or global – for important peacekeeping activities?
- Can AI help improve public safety, international security, and global preparedness for disruptions both natural and human-caused in the world?
- Can AI help improve the operations and resilience of networked, digital technologies for both organizational and public benefit – especially in an era of increasing internet devices?
- Can AI help improve the "essential fabrics" of open societies to include freedom of speech, freedom to think differently, and the need for an educated public to help inform pluralistic discussions all amid a digital tsunami of data?
- Can AI help improve education, focus, and entrepreneurial activities to tackle big, thorny, "hairball" issues like climate change, immediate & long-term food security, natural resources, and future sustainability for a planet of 9+ billion people?

These important questions represent a few of the important, shared outcomes to be explored and achieved through AI strategies that bring together human communities. While this report does not answer them all, it does indicate the different objectives being pursued by different nations with regards to their AI strategy as well as their performative declarations meant for the broader international community. Furthermore, this report both embodies and demonstrates the importance of interdisciplinary teams for AI research and AI education. Working across multiple disciplines is essential for both research and education especially as policymakers, business leaders, and students alike learn to explore and advance the necessary AI technical, commercial, civil, and ethical concepts required for a more positive future ahead.

III. For Trust, AI, and Societies, what if the Turing Test is the wrong test for AI?

This report represents a vanguard assembly of an interdisciplinary mix of fields to include the best of computer science, economics, political science, and public policy. Ultimately for AI to succeed in benefiting nations, communities, and networked groups of people, we must understand human nature more. We humans are products of natural selection pressures. Darwinian evolution is akin to a "blind watchmaker" – and as a result evolution has not prepared us to encounter the true alienness of AI. It is risky for humans to think AI is aligned to the same things we want and value, especially when the alignment problem of an AI to specific outcomes remains an unsolved challenge for several neural network approaches. In addition, we humans anthropomorphize lots of things including animals, weather, inanimate objects, as well as machines and now AI – even if those things do not act, think, or behave at all like us humans. Furthermore, training an AI depends heavily on the datasets employed, meaning both extant human datasets as well as our human choices regarding AI may amplify some of the more socially beneficial or detrimental elements of human nature. These elements include the considerable number of known human biases that each of us possess, to include confirmation bias, sunk cost bias, "in vs. out group" biases (aka, xenophobia), and many more biases though, fortunately, these biases can be mitigated some by education and experiences. By both providing a valuable compass reading as to where different nations have decided to steer their approaches to AI for the future ahead, and building the necessary foundation for bringing together an interdisciplinary mix of fields to study national AI strategies – Singh, Shehu, and their students enable readers to ask what I professionally consider to be the crucial question of the 2020s, specifically: what if the Turing Test is the wrong test for AI?

It is important to remember the original Turing test – designed by computer science pioneer Alan Turing himself – involved Computer A and Person B, with B attempting to convince an interrogator Person C that they were human, and that A was not. Meanwhile Computer A was trying to convince Person C that they were human. In reading the findings and conclusions of this 2023 report, I invite readers to consider what if this test of a computer "fooling us" is the wrong test for the type of AI that our society needs, especially if we are to improve extant levels of trust among humans and machines collectively?

After all, consider the current state of 2023 LLMs where benevolence of the machine is indeterminate, competence is questionable as existing LLMs are not fact-checking and can provide misinformation with apparent confidence and eloquence, and integrity is absent as the LLMs can with some variability change their stance if user prompts ask them to do as such. These crucial questions regarding the antecedents of trust associated with AI should not fall upon these digital innovations alone. First, these are systems designed and trained by humans. Second, ostensibly the 2023 iteration of generative AI models will improve in the future ahead. Third, and perhaps most importantly, readers who care about the national AI strategies present in 2023 around the world also should carefully consider the other "obscured boxes" present in human societies, such as decision making in organizations, community associations, governments, oversight boards, and professional settings.

All of which brings us back, in conclusion to the earlier corollary to the central question of Trust in AI and societies, namely: Can nations encourage Trust in AI and societies, while facing growing distrust in their economic and political systems? It could be that for the near future, both members of the public and representative leaders both in the public and private sectors need to take actions that remedy the perceptions of benevolence, competence, and integrity – namely Trust – both in AI and societies (sans AI) simultaneously. As mapping positive, deliberative paths forward to improve the state of Trust in AI and Societies is important, this important 2023 report delivers a prescient view of the current expressed state of fifty-four different national AI strategies to help us understand the present and consider the next steps necessary for the future ahead.

David Bray

David Bray

Dr. David A. Bray is a Distinguished Fellow with the Stimson Center and the Business Executives for National Security (BENS). He also is a CEO for different "under the radar" tech and data ventures – and has served in a variety of leadership roles in turbulent environments, including bioterrorism preparedness and response, Executive Director for a bipartisan National Commission on R&D, non-partisan leadership twice global CIO 100 award-winner, work with the U.S. Navy and Marines, and advisor to the U.S. Special Operation Command on the challenges of countering disinformation online. He has received the Joint Civilian Service Commendation Award and the National Intelligence Exceptional Achievement Medal. David served as Executive Director for the People-Centered Internet coalition Chaired by Internet co-originator Vint Cerf and is a Senior Fellow with the Institute for Human-Machine Cognition. Business Insider named him one of the top "24 Americans Who Are Changing the World" and he was named a Young Global Leader by the World Economic Forum. He previously gave the AI World Society Distinguished Lecture to the United Nations on UN Charter Day.

3 Summary & Recommendations

In 2016, the United States published its National Artificial Intelligence Research and Development Strategic Plan, usually understood in policy communities as the first statement of its AI infrastructure strategy (Select Committee on Artificial Intelligence, 2016). Since then over 60 countries have announced their national or sectoral AI policies.

This report employs computer science techniques to analyze the published national AI plans of 54 countries. In other words, we employ AI to analyze AI strategies. The report includes an analysis of 213 documents on AI strategies. Apart from national plans, the set includes reports and publications from various government departments, ministries, nation commissions, bodies appointed to forward recommendations for specific issues and sectors.

Our computer science methodology, specifically Latent Dirichlet Analysis (LDA) (Blei, Ng and Jordan, 2003), is calibrated to recognize embedded or latent *topics* that each document contains. It does so through providing probabilities of words that are most likely to occur together in each document. All documents are analyzed together for a pre-specified number of topics, ascertained through rigorous methodological criteria. The choice of the number of topics reflects fulfillment of various methodological LDA criteria for model stability (consistency) and topic stability (coherence). A document may feature a dominant topic, or a document may contain two or more topics. Further, we employ a technique known ensemble-LDA (e-LDA) to provide stable results assessed over multiple model specifications.

Collectively we present the most detailed and comprehensive empirical analysis undertaken of national AI infrastructures to date. This analysis provides comparisons and contrasts across 54 national strategies and a granular look at what these strategies contain. We note the priorities that are contained in documents, but our analysis also points out the policy depth for particular countries. Policy depth refers to the extent to which countries have covered the entire gamut of issues that comprise an infrastructure, and the institutional and financial resources they have committed to these issues. For example, AI policies from leading powers such as United States and China contain depth for basic research capabilities in science and mathematics, while the European Union policies contain the most depth for data governance and ethics. For example, one of the strategic objectives stated in the Chinese AI strategy states: "by 2025, China will achieve major breakthroughs in basic theories for AI, such that some technologies and applications achieve a world-leading level and AI becomes the main driving force for China's industrial upgrading and economic transformation" (State Council, 2017).

We make three major claims:

- There is no grand strategy or conclusion that applies to all AI infrastructures. Countries and clusters of countries feature different objectives and how to achieve them.
- Countries are pursuing a variable mix of similar elements in their national strategies. We propose and utilize the concept of 'AI Wardrobes' to show the various elements available for putting together an AI infrastructure and the variable ways in which countries are putting together these wardrobes.
- Clusters of countries pursuing similar strategies are identifiable. Our machine learning algorithms are able to point out some obvious clusters from the European Union, Latin America, and East Asia. But there are also surprises. United Kingdom leads a British influence cluster. Spain is prominent in the Latin American cluster.

Our three major claims are made at three different levels:

- We analyze 54 plans that are taken to be national. These are often 'performative'. They are as much about national priorities as they are declarations meant for the international community. But they reveal the broad trajectory and differences among national strategies.

- We analyze 213 documents including the national plans that national governments, commissions and departments have published on their AI infrastructures. Unlike, the performativity and differences among national plans, the intra-national plan reveal fewer national differences but a few countries have more policy depth than others. We notice countries that are at the early stages of policies regarding their AI infrastructures, versus those that have detailed regulatory and sub-sector policies.

- We also analyze the 213 documents, regardless of country labels, and here we see the broad topics that stand out in country plans. These include transportation, education, data ethics, and regulation. Looking at the documents we can then understand the countries that dominate these topics and also some broad differences among them.

Based on our analysis we present three policy recommendations:

- Comparative analyses like ours provide countries sign posts and guidelines for their ambitions. There is no one size fits all for designing national AI infrastructures. Different countries have different capabilities and priorities.

- Regulating AI will depend on country preparedness and political systems. Grand pronouncements such as fears about sacrificing our human rights or privacy to machine-led systems in our media about AI need a reality check. Several countries, generally with democratic systems, are putting together or struggling to put together systems of accountability, while others barely feature any such concerns. This provides room to think about governance issues, rather than ceding this authority to machines (or corporations) prematurely.

- AI policies have many good stories to tell about service provision. These include AI applications for health, education and research, and transportation.

4 Introduction

National governments regularly publish strategies for economic development, healthcare, transportation, and other areas of importance for improving governance and the lives of citizens within the country. As artificial intelligence (AI) has become an increasingly important and revolutionary technology countries have started to draft and publish national AI strategies to articulate their direction and vision for the technology and its application to different sectors of society. These strategies also assist governments with goal setting, resource allocation, organizational coordination, and international cooperation in the development of AI infrastructures and technologies.

Our *2023 AI Infrastructures Report* identifies which values, priorities, goals, and policy mechanisms have influenced the development of national AI infrastructures, as represented through national AI strategies. We utilize the term "AI infrastructure" to refer to machine learning (ML), natural language processing (NLP), and associated technologies that bring together data and algorithms which touch all aspects of human life including ordering goods, reviewing employment and housing applicants, diagnosing diseases, regulating traffic patterns, and instituting drones and robots as elements of warfare. National AI strategies give us insight into the complex narratives of technological diffusion, security concerns, and the implications AI will have for prosperity.

First, this policy report will provide an overview of the landscape of these policies which includes both national policies, intra-national policies, and in some cases international policies and agreements. Second, the report briefly describes the research methodology. Our research employs a natural language processing methodology, Latent Dirichlet Allocation (LDA), which produces sets of topics that capture the key words with probability distributions that occur together for each document in a designated set of documents. This methodology goes further than the typically used text-mining techniques by employing AI based techniques that consider probability and context of words rather than just word frequency. Next, the majority of this policy report will provide an overview of our results on the 54 national AI strategies at three levels: the national level accounts for the top-level policy document, the intra-national level analyzes several policy documents in each country, and a third level that analyzes all documents collectively to see which topics stand out globally and the main country documents that address these topics. Although we include the European Union in our analysis of documents, we use the shorthand 'countries' for our analysis.

4.1 National Landscapes

In 2017, the first set of countries published national artificial intelligence (AI) infrastructure plans, the first published came from Canada and China (State Council, 2017; CIFAR, 2017). Meanwhile, the United States government published a key document in 2016, the National Artificial Intelligence Research and Development Plan, which later helped to inform the 2019 and 2023 updates of this document and several other additional plans authored by the White House, Department of Defense, and other departments and offices (Select Committee on Artificial Intelligence, 2016; "Select Committee on Artificial Intelligence", 2023). Other countries which were first movers in publishing their AI strategies include the European Union (2018), Germany (2018), and India (2018) (Niti Aayog, 2018; German Federal Government, December 2020; European Commission, 2018). Currently there are over 50 countries which have published national AI strategies. These countries are diverse geographically, by income, and technological capabilities, but there are only two countries from Sub-Saharan Africa – Mauritius and Uganda. These national AI strategies articulate the previous, current, and future efforts of the authoring nation to shape and support the development of AI.

Many countries have also published intra-national strategy documents, which are strategy documents from directorates, departments, and commissions. These documents are typically specific to the authoring organization's area of focus such as education, start-ups and business-expansion, cybersecurity, data protection, health, and transportation. Some examples of such documents include Australia's "Artificial Intelligence and Emerging Technologies in School" strategy published by the Australian Department of Education and Training (2019) or China's "Regulations for the Promotion of the Development of the Artificial Intelligence Industry in Shanghai Municipality" which was published by The Standing Committee of the 15th Shanghai Municipal People's Congress (2022). There are varying levels of overlap between these strategies and greater national AI strategies. Ultimately though, these strategies also have important implications for the development of AI infrastructures within a country, and specifically for the infrastructures directly related to the targeted domain area.

4.2 Stages in the Development of AI Strategies

Given the recent advent of national AI strategies, some countries have more developed articulated strategies than others. Some national strategies are focused on essential infrastructures needed to develop and deploy AI systems, whereas other strategies articulate elements of how specific ministries and offices will create the next step in AI infrastructures, technologies, and regulatory environments. In many cases, AI infrastructures in one country are influenced by similar steps in other countries. Our results, presented later in this report, have identified this empirically.

4.3 Existing Narratives in AI Infrastructures

There are many overarching narratives that surround the development of AI, national AI plans, and the different ways in which AI can be employed. These narratives have made their rounds through the media, politicians' speeches, political platforms, debates, and in some cases, policy. This research seeks to understand what narratives exist in AI policies around the world. Our methodology for doing this is detailed later in this report. It is important though to address some of the dominant narratives in this space, as they directly relate to the narratives which emerge from this research.

Many national AI plans tout the uniqueness of AI as a revolutionary technology which we have yet to see before. This statement is said in a matter-of-fact manner, without much concern for the historical progression of technology. We have had revolutionary technologies before – the printing press, electricity, the internet, and others – what makes AI so different? This question is not typically answered, but instead the argument is made by pointing out the many different aspects of life which will be changed by the roll out and continuous innovation of AI. To push back against this narrative is not to reject the massive impacts AI will have on life as we know it, but it is important to realize that other technologies have been introduced into society that have completely changed the way in which we live. We have thus adapted to these technologies and learned how to use existing institutions and governance structures to manage how these technologies impact life. It will ultimately be up to a variety of societal and governance institutions to determine if AI is a job replacer or augmenter, a tool for social good, or a tool for a dystopian future. Many of these national AI plans seek to deal with shaping the future of how AI impacts everyday life – but it is important to note that what narratives they employ determine their ultimate policy trajectory.

Another major narrative of AI involves the potential for the technology to be utilized by governments as a tool for oppression. A concern often articulated in various media is that in authoritarian states, governments will utilize AI to surveil, police, censor, suppress, and manipulate their populations. There is empirical proof that this has occurred in a few countries already. What is important to note though is that these tools will be available to all governments, authoritarian or democratic. Empirically what we demonstrate later in our analysis of the national AI plans is that pluralist states are more likely to open the process for developing systems to govern and support the development and roll out of AI with consultation from civil society. This is less true of non-pluralist or authoritarian governments. We do not refute the latter narrative, but more research needs to be done on the mechanisms which lead to the use of AI tools by governments for these purposes. Regardless, there is overlap between how pluralist and non-pluralist countries seek to develop AI and implement it for the purposes of economic development, security, and other broader uses.

Our machine learning based analysis of the national AI plans around the world better understands what narratives are influencing national strategies to develop and deploy AI. Utilizing our methodology we identify the ways in which specific narratives might lead to different mechanisms to shaping AI infrastructures. It is important to note that many stories about AI in popular discourse take a negative tone. This report takes an evidence-based approach to veer away from both negativity or overly optimistic AI scenarios. Many of the national AI plans emphasize the importance of inclusion, diversity, and transparency as the technology continues to develop and be deployed.

4.4 AI Wardrobes

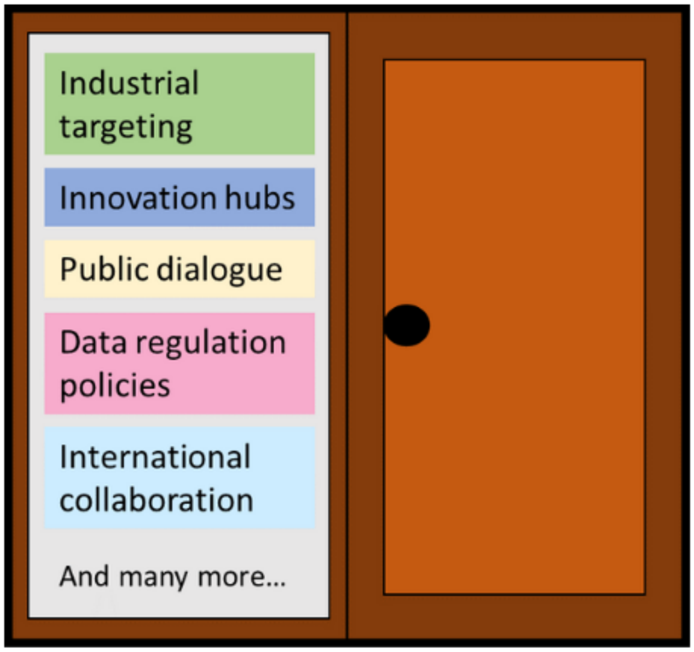

Figure 1: The AI Wardrobe

To best understand the way in which countries build their national AI policies, we have created the concept of AI wardrobes. Using this concept, we argue that the universe of national AI strategies can be conceived as a global wardrobe and a country's national AI plan is an outfit. How countries present their national AI infrastructure narrative or story is an outfit which they selected from the shared universe of options within the AI wardrobe. Therefore, within the wardrobe for national AI strategies there are concepts such as "trustworthy AI," "inclusion," "data privacy," "R&D strategies," and "international collaborations". These concepts are different articles of clothing and just as in reality, clothing can come in different patterns, shapes, and with different embellishments. Countries can select which policy elements they want to utilize. In some cases, countries might simply pick a "shirt" such as emphasizing "workforce training" and wear it. In others, a country might adapt it to their needs, stylizing it with a pin or a button to represent their own distinct flair and style. China places a great deal of emphasis on "talent" in its education strategies. Some clothing pieces in the wardrobe also might be "hand-me-downs" or elements which countries directly copy or duplicate from others. Many EU countries adopted the language on inclusion and diversity from the EU's own plan. Meanwhile, other pieces of clothing might be newly produced, or original ideas. India's start-up policies or Germany's emphasis on *mittelstand* or SME policies are examples (Niti Aayog, 2018; German Federal Government, December 2020). Ultimately, countries get to decide what elements of AI policy to put into their national AI strategy, and they get to determine how to put these pieces together to create their own outfit in their own distinct style.

It is our goal with this research to ground the narratives surrounding the development and deployment of AI in empirical analysis. To do this we have developed a methodology which utilizes computer science techniques, specifically LDA and eLDA. This is detailed in the next section of this report.

5 Methods and Data-Set

There is existing research which compares AI infrastructures, but much of it relies on more traditional methodologies, which include close text readings (Bareis and Katzenbach, 2022) and keyword frequency analyses in combination with other qualitative methods (Robinson, 2020; Wilson, 2022; Foffano, Scantamburlo and Cortés, 2022). Some uses of natural language processing (NLP) methods have been used (Hine and Floridi, 2022), but the corpus of input texts is only a subset of national AI plans rather than the universe of national AI plans, which our research covers. Additionally, much of the work in this space is focused on two country comparisons with a heavy emphasis on the U.S. and China.

In March 2023, a report entitled *Building Trust in AI: A Landscape Analysis of Government AI Programs* by Susan Aaronson was published by the Centre for International Governance Innovation (Aaronson, 2023). This report leverages the OECD AI policy website and database (discussed in this section) and asks important questions on government efforts to develop AI capabilities and trustworthy AI, evaluate their own efforts in developing AI, and the development of best practices on AI and trustworthy AI at the OECD. This research is comprehensive and provides an overview of different policy approaches to AI around the globe, but it has a different purpose than the research of our team. The CIGI research is more focused on broad policy mechanisms, evaluation, and the development of best practices for policy rather than understanding the important values and narratives which motivate the development and implementation of AI policies at a micro level. Our research uncovers the motivations and narratives which exist in national AI policies and compares these narratives across countries, which yield data that can help us better understand the global diffusion of values in AI infrastructures. We do this using a computer science methodology, LDA (Blei, Ng and Jordan, 2003).

Our LDA methodology employed in this research goes beyond existing methodologies in analyzing national AI policies in two important ways. *Firstly*, we have analyzed the universe of existing national AI policies. *Secondly*, we are utilizing an ensemble LDA methodology (eLDA) which builds on the existing and pre-trained LDA model. Latent Dirichlet Allocation or LDA is an unsupervised NLP algorithm, which we can run through on a set of documents, and the algorithm outputs the main topics which exist in the document set or in specific documents within the set. As a conceptual aid, one can think of a topic as being the "main idea" or "key theme." LDA identifies the main topics in a set of documents by analyzing patterns among the words within the documents specified. This is done through a process known as Gibbs sampling, assigning probabilities to both words and topics until a converge is established; the algorithm effectively "tells" us which words are associated with which topics, and which topics are associated with which documents, each weighted according to its "prominence." As previously mentioned, LDA is an unsupervised model, meaning that we as researchers do not tell the algorithm what topics we want it to search for; the model determines this through analyzing the probabilities of words in specific contexts throughout each document and ultimately the set of documents being utilized. It is important to note that LDA also does not simply consider the frequency of words within a document but considers the context of the word through analyzing the words used around each word and thus determining specific meanings of words. Ultimately the LDA algorithm is applied to our set of documents, and after several iterations it determines a set of topics which are likely to have generated the final collection of words in each policy document. Each document in our dataset is modeled by a mixture of topics produced by the LDA algorithm. Subject matter experts can then analyze the set of words which inform a topic to determine what a specific topic means. See Figure 2 for an example.

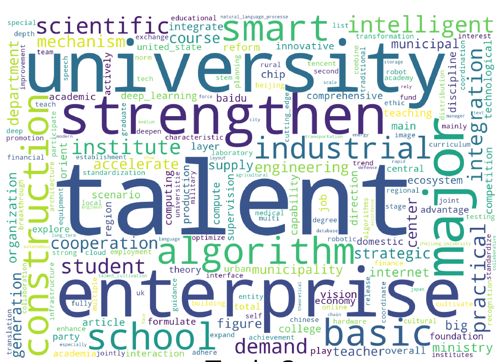

Topic 2

Figure 2: A sample word cloud of a topic based on education and talent. The font size of each word follows the probability assigned by the algorithm to that word in this particular topic (as determined by analysis of the documents in the corpus).

Our methodology goes even a step further, as we are utilizing ensemble LDA (eLDA). With traditional LDA, there may be differences in the resulting topics each time one runs the algorithm. This happens because the way in which the model calculates the probabilities which determine the topics. eLDA stabilizes LDA results by building instead an ensemble of models and then distilling the topics in common among the various ensembles, therefore stabilizing the methodology. The shared, stable topics are then utilized to offer a stable meta-model. The ensemble approach eliminates the impact of model sensitivity on topics, thus improving the methodology's accuracy and reliability. In our methodology, we build several ensembles and compare the resulting meta-models by analysis of the coherence of the topics; we utilize various state-of-the-art topic coherence metrics to essentially measure measure how interpretable topics are to humans. The meta-model with the highest topic coherence is utilized to obtain results for interpretation. Ultimately, the output of this methodology is a set of topics which represent our universe of national AI policies. These topics make the narratives, goals, and mechanisms (and even motivations) a country is utilizing to shape the development of their AI infrastructures more explicit.

Our methodology emphasizes a narrative based approach to understanding AI infrastructures in comparative contexts. This approach gives a far more comprehensive and detailed analysis of the "AI wardrobe" and the different "outfits" which countries have selected from it. Existing comparative studies of national AI strategies have tried captured the broad narratives which exist in these strategies through broad interpretations or focus on a few key words. We build on these works in a meaningful way by bringing an empirically based approach, one that accounts for far more granularity, to analyzing these documents.

5.1 Our Dataset: Introducing the National AI Policies

Our dataset is made up of 54 countries who have published national AI policies. These countries span different regions and are diverse in levels of national income, technological development, and type of political system. We have identified 54 national strategy documents for each of these countries and 167 intra-national documents. Therefore, a total of 221 text documents have been included in our eLDA model which inform the resulting topics output from the model. These documents thus cover several different wardrobe aspects, or illustratively we could say, these documents are several different blouses, pants, skirts, and jackets in a variety of colors and patterns.

Practically, given our methodology our team needed to collect each of the existing national AI strategies and as many intra-national policies available. Fortunately, there are existing resources which have compiled many of these documents. The most complete resource is OECD's AI Policy Observatory. This database features a "country dashboard" for 70 countries. Each country dashboard is a page which hosts all of the national and intra-national AI policies, as well as other policies relevant to AI, for a specific country that are known by OECD. In some cases, countries may have intra-national documents but not a national level policy.

Our team started the document collection process by using the OECD database. Throughout this process though we identified some problems with the database, making it impossible to simply download all of the resources and input them into our eLDA model. To arrive at our final dataset our team visited each of the country dashboards and carefully investigated each document hosted on each dashboard to determine if the resource was (1) an official government policy document, (2) relevant for AI policy, and (3) not a stub. Some documents were published in non-English languages, and we utilized Google translate and other NLP tools for the purpose of translating these documents to run them through our eLDA model. Our team also found documents outside of the OECD database. These countries include China, Qatar, and Russia. In some cases, our team also located documents not hosted on the OECD website through web searches. These searches were necessary due to many resources on the OECD database being reported as web-pages with dead URLs. We learned through our data collection process that it is fairly common for AI documents that are hosted on web-pages to be removed or hosted elsewhere as government administrations change, websites are overhauled, or policies are changed/updated.

For our analysis we also needed to designate which policies were the national plans and which were intra-national plans. We have built our methodology to have a few different levels of results. One level of analysis is just national plans, another analyzes both national plans and intra-national plans. We designated national plans utilizing our own system rather than relying on the existing classifications from the OECD website. Our system for categorizing these documents relies on a government's designation of their national plan. In some cases, there are multiple documents designated as a country's national plan; this is because the national government identifies multiple documents as such.

6 The Empirical Findings

We analyzed a total of 54 national plans and 167 intra-national plans, making a total of 221 documents in our final model. Our national plans reveal broad national policy priorities, while intra-national plans show levels of policy depth. We also performed document specific analyses which show shared policy issues (as identified by topics) across countries. We have organized our results below, emphasizing the topics resulting from the eLDA model, and which country documents are best summarized by which topics. For ease of understanding we have provided word clouds, which illustrate the key words that make up each topic – thus allowing us to summarize the meaning captured in each topic.

6.1 Comparing National AI Infrastructures: Policy Priorities Revealed

From our 54 national plan documents the following results were produced. It is important to note that in some cases countries had two or three official documents that made up their "national plan" due to updated versions of the policies being published. This is the case for the national plans for the United States, Canada, China, Estonia, Finland, France, Germany, India, and Japan. The number of topics selected for our final eLDA model was 15, meaning that the model sorted the dominating themes into 15 different topics, and then calculated the probability of each national plan featuring that topic in its contents. The choice of 15 fulfills our methodological criteria for model stability (consistency) and topic stability (coherence).

The eLDA model produces a helpful heat-map (Figure 3) containing these results. This heat-map features topics on the horizontal axis (labeled 1-8) and countries (as represented by their national AI plans) on the vertical axis. The shading in the heat-map indicates topics which have a higher probability of presence in a country's national AI plan and are shaded darker. You will also notice on this heat-map that there are dendograms on the left and top. The dendogram on the top of the heat-map shows how topics relate to one another while the dendogram on the left-hand side shows how different countries relate to one another. These relationships are determined by the Hellinger distance (Nikulin, 2001) and agglomerative hierarchical clustering.

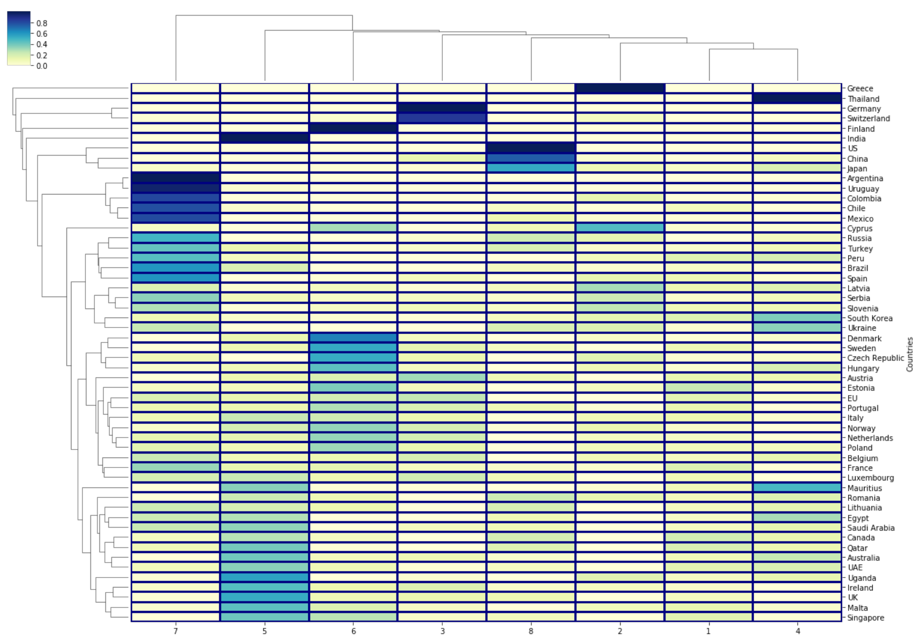

Figure 3: National AI Infrastructures: Heat-map relating country - topic distributions.

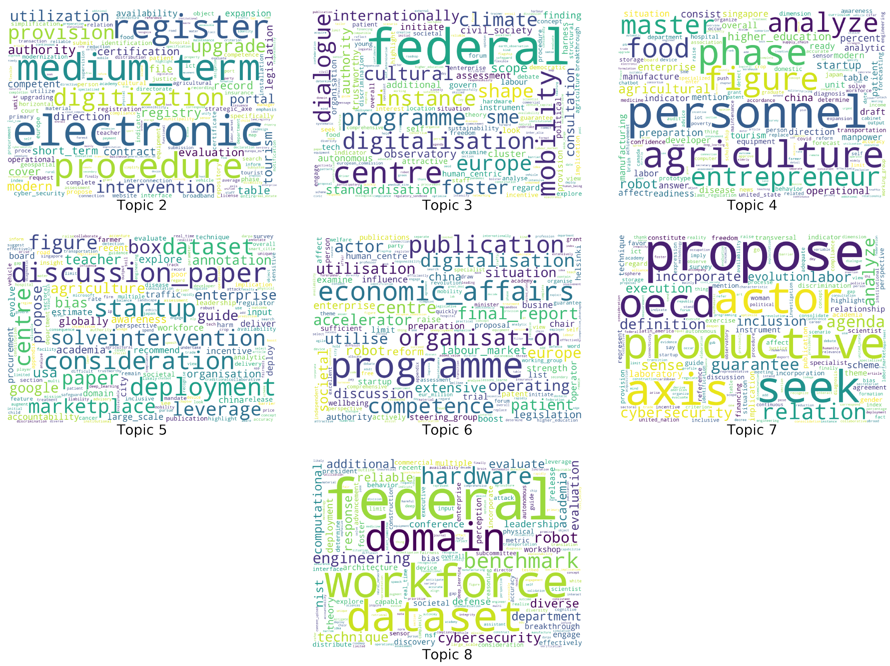

Figure 4: National AI Infrastructures Topic Word Clouds
Topic 1 does not hold much significance in the corpus and is omitted.

At this level, we can report several major findings across clusters:

- There are clusters and there are distinct leaders that do not cluster with many states. The major clusters are: EU, Latin America, historic British influence. The major leaders are China, Germany, Japan, and the United States.
- Topics feature divergences but there are many convergences. At a macro level, we see correlation matrices among topics. That of Germany is related to the EU, for example. At a micro level, we see fine distinctions. Germany emphasizes standards, while the United States emphasizes benchmarks, but they are both about technical thresholds.
- Country plans can contain dominant but multiple topics. The U.S. plan is unique with the dominance of topic 8. However, the Chinese plan has dominant topic 8 but also topic 3, which is dominant for Germany and Switzerland.

We are able to identify seven distinct clusters in this heat-map. These clusters have largely broken down in by geographic organizations which make sense due to overarching geopolitical associations and historical relationships.

Table 1 presents the different clusters, their associated topic produced by the eLDA model, and the member countries. Note that some countries are only partially included in a cluster, this is because they may have high probabilities for more than one topic, which is possible when using the eLDA methodology and more accurately captures the reality that some plans may have been influenced by more than one pre-existing plan or competing interests/priorities domestically.

Cluster Name	**Topic # & Top 10 Words**	**Associated Countries**
Greece Cluster	2 – Electronic, medium term, procedure, register, digitization, provision, upgrade, intervention, utilization, tourism	Greece, Cyprus
German-Swiss Cluster	3 – Federal, centre, mobility, programme, digitalization, dialogue, Europe, instance, shape, SME	Germany, Switzerland
Economic growth & development	4 – Personnel, agriculture, phase, figure, entrepreneur, analyze, master, food, agricultural, robot	Thailand, South Korea, Ukraine, Mauritius
Commonwealth - British Influence	5 – discussion paper, deployment, startup, consideration, centre, intervention, dataset, paper, solve, figure	India, Singapore, Malta, United Kingdom, Ireland, Uganda, United Arab Emirates, Australia, Qatar, Canada, (Partially includes Saudi Arabia)
EU Cluster	6 – programme, economic affairs, publication, organization, competence, digitalization, centre, actor, final report, utilization	European Union (Partially includes Spain)
Latin America - Spain influence	7 – propose, productive, axis, seek, actor, OECD, relation, guarantee, agenda, analyze	Spain, Argentina, Uruguay, Colombia, Chile, Mexico, Peru, Brazil, (Partially includes Russia)
Science & Technology First Movers	8 – federal, workforce, dataset, domain, hardware, benchmark, engineering, robot, cybersecurity, technique	United States, China, Japan

Table 1: Clusters, Topics, Top 10 Words, and Countries
Note: Cluster 1 omitted because no major countries were included in it

The following analysis of topics is not chronological: leads with first movers, pace-setters (Germany), then EU, UK/Commonwealth, then Latin America/Spain.

Topic 8 *Top 10 words*: federal, workforce, dataset, domain, hardware, benchmark, engineering, robot, cybersecurity, technique

Topic 8 contains themes that have a high probability to be found in the national plans of the United States, China, and Japan. All of these countries are leaders already in technology and innovation – therefore it makes sense that they might have well articulated and formulated plans towards developing and deploying AI domestically. This topic emphasizes datasets, hardware, the workforce and benchmarks (rather than standards) and basic science capabilities. Thus indicating some level of understanding of the needs for the AI industry, and a desire to be a part in providing the proper environment for these elements of AI infrastructures to be developed and utilized. The emphasis also on basic science and engineering is important to note – as these lend themselves to strategies that prioritize research and the cutting edge. Additionally, we see the use of the term "cybersecurity" which has yet to show up in other topics – indicating that these countries have concerns about elements of security. This is in line with the existing geopolitical conflicts between the U.S. and China. Another primary feature of this topic are clear statements about aspirations to be global leaders in the AI industry.

Topic 3 *Top 10 words*: Federal, centre, mobility, programme, digitalization, dialogue, Europe, instance, shape, SME

Topic 3 is closely related with the strategies of Germany and Switzerland. These plans emphasize the importance of federal action in the AI space and recognize the importance of dialogue and consultation. Switzerland's strategy is a broader plan, "Digital Switzerland Strategy" which puts an emphasis on AI and its importance for the country to be a leader in the technology. Meanwhile Germany's AI strategy is solely focused on the involvement in the development and applications of AI. Both the German and Swiss plans heavily discuss the importance for dialogue and coordination within the European region and other international organizations.

Topic 6 *Top 10 words*: programme, economic affairs, publication, organization, competence, digitalization, centre, actor, final report, utilization

Topic 6 is closely related to many of the countries within the European Union. The top words that define this cluster are closely related to specific EU programs and issue spaces of specific interest to the countries in the union. These words suggest a large amount of strategic action proposed in these national plans. It is likely that these plans suggest to development of new programs and organizations, or reorientation of existing ones, to support the development of AI. There is also an element here of assessment, in the references to publications and final reports. This reflects upon the processes which typically occur within the EU to coordinate member actions in specific areas. Specific countries included here are: Denmark, Hungary, and Estonia – among others.

Topic 5 *Top 10 words*: discussion paper, deployment, startup, consideration, centre, intervention, dataset, paper, solve, figure
This topic, like topic 4, is focused heavily on economic growth and development. Yet there is a difference in approach. This topic is more focused on two things, first strategizing which we see through words such as “discussion paper,” “deployment,” “consideration,” “centre,” “intervention,” and “paper.” These plans have a high probability of emphasizing different tasks which need to be done by a variety of actors to support the development of AI. It appears there is a whole of government type approach. Aside from that, there is a focus on startups and datasets, which are important for developing a robust domestic AI industry. Many of the countries which have a high probability of their plans containing this topic have ties to British colonialism or influences during history. These states are full of ambition and seek to manage the development of AI through a variety of different policy levers.

Topic 7 *Top 10 words*: propose, productive, axis, seek, actor, OECD, relation, guarantee, agenda, analyze
Topic 7 is has a high probability of appearing in the national strategies of many Latin American countries (including Argentina, Uruguay, Colombia, Chile and Mexico) but also countries such as Russia, Turkey, and Spain. Some of these countries clustering together may inherently make sense due to historical and cultural legacies. Meanwhile others are more surprising. The topic seems to emphasize broad processes and objectives (the term axis is a machine translation of Spanish word *ejes* or objective). Topic 7 speaks to country plans which are formulating different government tasks to undertake to develop AI (“propose,” “actor,” “guarantee,” and “agenda”). These plans also seem to be interested in productivity and working to catalyze productivity with AI and effective plans to harness the technology for economic benefits. Finally, the reference to OECD indicates that these countries might either be interested in collaborations, or desire to learn from other leaders in the OECD (which has done work on analyzing AI policies around the globe – as referenced earlier in this paper).

Topic 4 *Top 10 words*: Personnel, agriculture, phase, figure, entrepreneur, analyze, master, food, agricultural, robot
Topic 4 is characterized as a topic which encapsulates plans focused primarily on economic growth and development. There is a heavy emphasis on the agriculture industry in this topic, but also on entrepreneurship and personnel. The top ten words in this topic are indicative of strategies which have a strong likelihood of utilizing AI for the purpose of catalyzing economic development. These countries include Thailand, Ukraine, South Korea, and Mauritius. Each of these plans makes multiple mentions to the importance of applying AI to the

agriculture sector, with three of the four plans having designated sections to the topic. Additionally, all emphasize the importance of entrepreneurship and innovation as key elements for developing the technology and using AI to spark further company and device creation. Ultimately these plans appear to that a holistic approach, not limited to each country's strength but touching on a variety of important industries and aspects to develop and deploy AI into for economic development.

Topic 2 *Top 10 words*: Electronic, medium term, procedure, register, digitization, provision, upgrade, intervention, utilization, tourism
In this level of analysis topic two was most closely associated with the national AI strategies of Greece and to a slightly smaller probability, Cyprus. This means that there is a high probability that the word distributions associated with topic two will appear in Greece and Cyprus' national AI strategies. These plans appear to be highly associated with applying AI and other technologies towards medium term goals, likely with particular interest in specific industries of importance, such as tourism. There appears to be recognition of a need to upgrade and digitize to see forward additional AI aspirations.

6.2 Comparing Intra-national AI Strategies: Policy Depth

For the analysis of our intra-national policy documents we included both the 54 national AI policy documents and the 167 intra-national documents. In other words, the results developed from running the eLDA model on our total dataset of 221 text policy documents. As mentioned earlier, intra-national policy documents are documents relevant to AI from various government agencies, departments, commissions, or institutions. We have provided a breakdown of the top ten countries which have published intra-national AI documents in Table 2. Examples of these documents include autonomous vehicle policies, AI education policies, and policies for data protections as they relate to AI.

Country/Region	**Number of Documents**
European Union	28
United Kingdom	17
China	16
Japan	16
United States	15
Colombia	15
India	6
Germany	5
France	5

Table 2: Top 10 Countries for Number of AI Documents

These plans have important implications for how AI is developed, used in various industries, but also for how a variety of spheres of society are prepared for an AI enabled future. The analysis produced from this model shows the different policy themes which exist amongst the featured countries and their whole of government approaches to AI policy.

In Figure 5 we have provided a heat-map which shows what country's combined national plans and intra-national plans with topics created from the eLDA model. This diagram shows the different countries included in the analysis on the vertical axis and the different topics on the horizontal axis. The darker blue the spaces are colored, the higher the probability is that that specific topic is within the set of documents associated with that country. In Figure 6 we have a breakdown of the contents of each of the topics generated from our eLDA model.

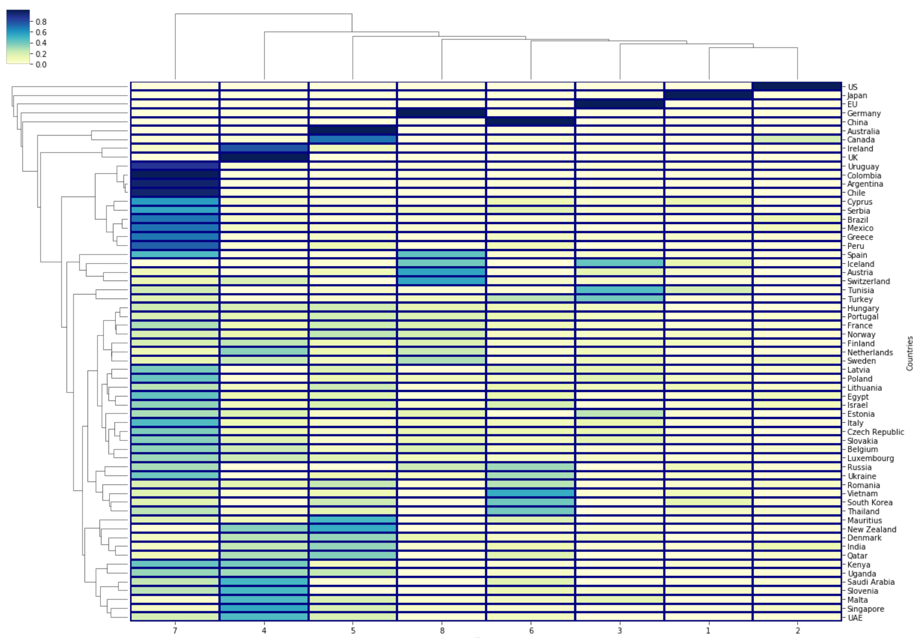

Figure 5: Intra-National AI Infrastructures: Heat-map relating country - topic distributions.

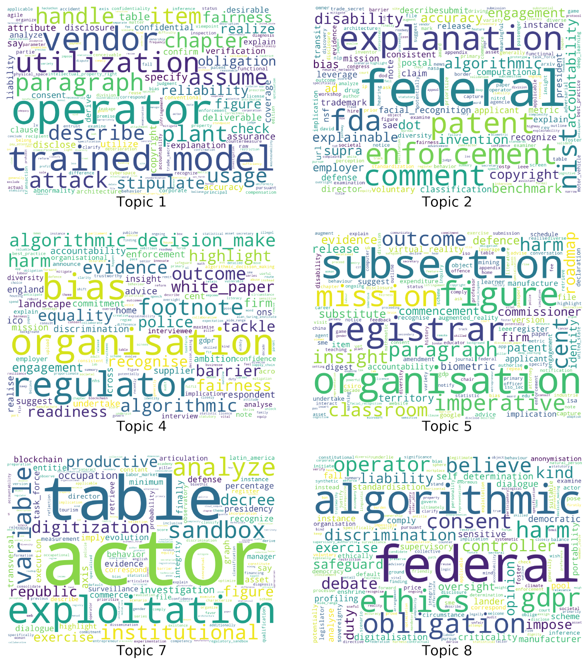

Figure 6: Intra-National AI Infrastructures Topic Word Clouds

Here are a few broad themes at the intra-national level:

- The eLDA methodology resulted in eight different topics. Before this analysis was completed, the expectation was that an LDA with all international documents included would result in more focused, sharper topics. The results were, in fact, different – with a large amount of coherence around three dominant topics (topics 7, 4, and 5). These topics capture the necessary fundamentals a country needs to develop an AI infrastructure. The countries which are associated with topics are different from the four countries and one region we have identified as global leaders in AI research and/or ethics: Germany, China, the European Union, Japan, and the United States.
  Their associated topics are listed in Table 3.

Country/Region	Topic # & Top 10 Key Words
Japan	1 – operator, trained model, vendor, utilization, paragraph, handle, attack, item, assume, plant
United States	2 – federal, explanation, enforcement, patent, comment, FDA, NIST, algorithmic, disability, accuracy
European Union	3 – liability, member state, damage cause, fault, parliament, compensation, stress, GDPR, obligation, trustworthy
China	6 – trustworthy, attack, figure, scenario, theory, municipal, municipality, chip, college, reform
Germany	8 – federal, algorithmic, ethics, obligation, GDPR, harm, consent, believe, operator, discrimination

Table 3: Topics for Leading AI Countries and EU

- Figure 7 illustrates different topics and their probabilities to have related word distributions appear in national plans of the U.S., Russia, Germany, EU, China, Canada, and India. We can conclude from this figure that some countries contain multiple topics while others contain only one. We argue that this is likely because some countries have elements of AI policies which are originally derived while others borrow elements from other pre-existing policies, thus illustrating some level of policy diffusion. Another reason for this is AI policies which are in their earlier stages of development in terms of articulating their goals and mechanisms for developing and utilizing AI systems. Figure 8 illustrates the number of countries which have some levels of all the topics in their word distributions, and this is modeled by numbers of countries with each topic. Of these topics, 7 & 4 are largely associated with fundamentals of any AI policy, and 5 & 8 are associated with basics of data governance and regulatory issues (see topic breakdown in Figure 6).

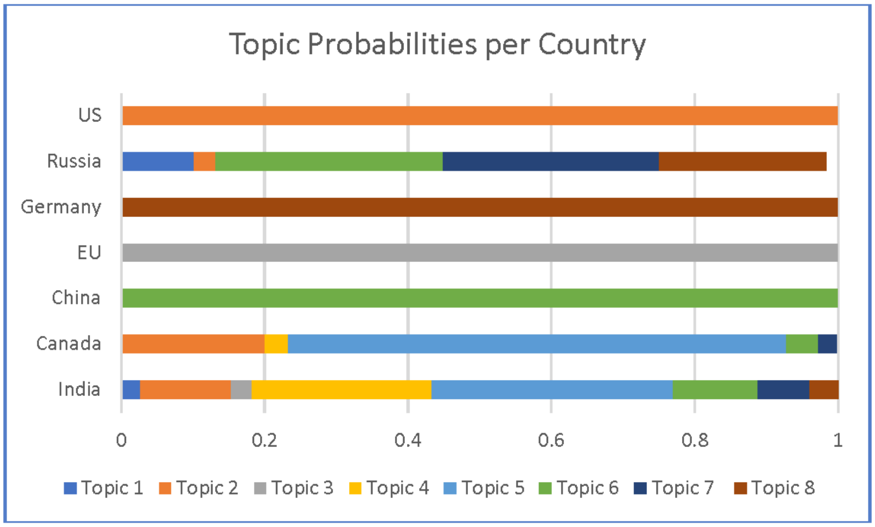

Figure 7: Intra-national Policy Documents: Topic Probabilities Per Country

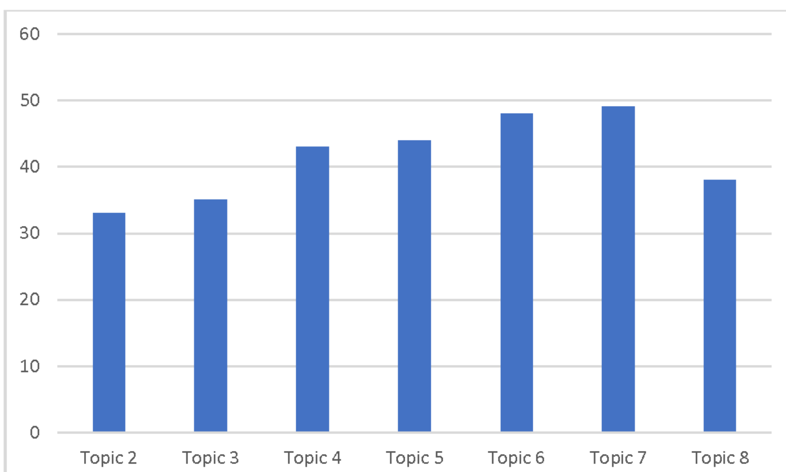

Figure 8: Country Counts Per Topic

Topic 1 *Top 10 words*: operator, trained model, vendor, utilization, paragraph, handle, attack, item, assume, plant

Topic 1 is closely associated with documents from Japan. Many words that are related to this topic are highly technical including "trained model" and "operator," indicating that some intra-national documents by the Japanese government are likely sufficiently detailed. Additionally, we see words such as "attack" indicating concerns about security, which relates to Japan's national plan being associated with the topic containing a reference to "cybersecurity." The Japanese government has published several intra-national documents including "Act of Protection of Personal Information," "AI Utilization Guidelines," "Machine Learning Quality Guidelines," and "Social Principles of Human-Centric AI." These documents are highly specific, and some quite technical. This shows a depth to the government's approach to AI policy. There are multiple sets of guidelines published, which also lend themselves to regulatory perspectives and a high degree of government guidance to development of technology.

Topic 2 *Top 10 words*: federal, explanation, enforcement, patent, comment, FDA, NIST, algorithmic, disability, accuracy

This topic is closely associated with the United States. In fact, the U.S. features only this topic it its topic composition, therefore it is incredibly important for understanding the U.S. national and intra-national documents. We can note that some of the terms in this topic are specifically related to the U.S. since they are agencies in the American Government – FDA and NIST. Other words include "federal" which indicates that the federal government does have a role and plans related to AI. Additionally, there are several words related to governance – "enforcement" and "patent" specifically. There is also an element of AI policy being open for shaping by the public through the word "comment." Some intra-national documents published by the U.S. include documents from the U.S. Patent and Trade Office, the "NIST Principles for Explainable AI," the "Plan for Federal Engagement in Developing Technical Standards and Related Tools," and the "Artificial Intelligence and Algorithmic Fairness Initiative."

Topic 3 *Top 10 words*: liability, member state, damage cause, fault, parliament, compensation, stress, GDPR, obligation, trustworthy

Topic 3 is closely associated with the European Union, but not with specific EU countries. This is likely due to the nature of EU approaches versus the approaches of member states. The EU has published several intra-national (intra-regional in this case) documents. One of these includes the document referenced as "GDPR" which is the EU's data protection law. Documents from the EU seem to emphasize trustworthy and fair AI systems. This is evident through words such as "trustworthy," "obligation," "liability," and "compensation." It also appears that there might be some allusion to different concerns related to the roll out of AI. Some plans published by the EU include "Ethics Guidelines for Trustworthy AI," "Open Data Directive," the "Framework of ethical aspects of artificial intelligence, robotics, and related technologies," and the "Sectoral Considerations on the Policy and Investment Recommendations for Trustworthy Artificial Intelligence." These different documents show the massive range in publication topics but also the emphasis on ethics, dialogue, and trust.

Topic 4 *Top 10 words*: Organization, bias, regulator, footnote, algorithmic decision make, evidence, harm, algorithmic, outcome, barrier

Topic four contains several countries. This topic captures many elements that are critical for the regulation of AI. Terms that indicate efforts or plans towards regulating AI include "regulator" and "evidence." This topic also emphasizes some new elements such as "harm" and concern over "bias" therefore, there are likely elements of concern about AI in these documents, but also likely plans for regulation to mitigate the impacts of these potentially negative impacts of AI. It is important to note that most of the countries which are also in topics 7 and 5 also have overlap with this topic. This is due to the fact that topics 7 and 5 are also topics which relate to basic aspects of AI infrastructures.

Topic 5 *Top 10 words*: subsection, organization, registrar, figure, mission, imperative, cent, paragraph, outcome, classroom

Topic 5 is another topic which contains a variety of countries and overlaps with topics 4 and 7. This topic primarily emphasizes economic organization, featuring words such as mission, imperative, outcome, firm, manufacture, and GDP. This topic being present indicates plans that may focus on the uses of AI for economic development and growth. This is important in the view of many countries. There are also aspects of this topic which allude to education such as "classroom" which connects the importance of education to the development of AI – but also the consideration of the implications of AI for education. Two countries which heavily feature topic 5 include Australia and Canada.

Topic 6 *Top 10 words*: trustworthy, attack, figure, scenario, theory, municipal, municipality, chip, college, reform

Topic 6 is heavily featured in the intra-national sets of AI strategies published by China and South Korea. This topic emphasizes a variety of distinct themes in AI strategies. There is an emphasis on security but also on trustworthiness – thus focusing on both outside threats to AI systems and needs of those domestically to trust AI systems sufficiently. Another key term is "chip" which is fitting given both China and South Korea's involvement in semiconductor manufacturing, which both countries likely hope to grow and improve in future. There are also references to universities – emphasizing likely both education and research, which are key elements of strategies for workforce development.

Topic 7 *Top 10 words*: actor, table, exploitation, analyze, variable, sandbox, institutional, digitization, productive, decree

Topic 7 should be considered the most basic organizer of AI infrastructures. This term is one we are using to describe the basic elements needed for AI infrastructures to be created and support the development of the technology. *Of the 54 countries included in this analysis, 42 of them feature some elements of this topic.* Important themes within this topic are: actors, institutions, execution, economics (commerce, blockchain, productive), processes (dialogue, sandbox), and capabilities (digitization). These terms are all important pieces of designing the right environment for innovation in AI, but also for continuing to evaluate plans and effectiveness as well. Countries which have a large portion of this topic include most of the Latin American states, many of the European Union states (but not Germany, Ireland, or the EU plan), and developing countries. This topic overall captures the 3Cs of AI infrastructures: competencies, concerns, and capabilities.

Topic 8 *Top 10 words*: federal, algorithmic, ethics, obligation, GDPR, harm, consent, believe, operator, discrimination

Topic 8 is associated closely with Germany and shows a high degree of ethical and societal concerns. This is seen through the prevalence of words such as "ethics," "GDPR," "harm," "consent," and "discrimination." It is also interesting to see the word "believe" which may indicate an element of norms being important in these intra-national plans, and these norms may also relate to ethics and fairness in the AI infrastructure space. Some documents published by the German government that were included in this analysis include the "Interim Report One Year Strategy," the "Opinion of the Data Commission," and "Ethical Guidelines for Self-Driving Cars." The emphasis on data protection and safety in self-driving cars aligns with the themes of this topic as well.

6.3 Document Level Analysis: Analyzing the 5 Most Important Topics in Our Intra-national Documents

To round out our analysis of our dataset of AI policy documents we ran our eLDA model on our set of 167 intra-national policy documents. We opted to run this analysis at the document level rather than the country level. We did this in order to establish the most popular topics throughout these documents. Once we obtained these top topics, we were able to then connect the documents associated with each topic to their publishing country. For this analysis documents from 37 countries were included, as some of the countries which have national plans do not have intra-national plans. Additionally, we have added Kenya to this analysis, a country which has intra-national documents but not a national level AI plan. We have broken down and classified the five top topics in these documents and presented this information in Table 4. Here is our macro analysis of the intra-national documents:

Topic Number	Topic Name Descriptor
1	Data & Governance
2	Education & Training
3	Economy
4	Contracts & Liability
5	Transport

Table 4: Main Topics Across Documents

- Our results show that 27 of the 38 included countries have published documents that feature elements of all five of the top topics that eLDA identified. Meanwhile, there are 11 countries that do not address all 5 of these topics. These countries are: Egypt, Estonia, Greece, Iceland, Kenya, Malta, Norway, Singapore, Thailand, Tunisia, and Vietnam.
- Our eLDA model determined that topics 2 and 3 are the most similar, showing a connection between education, training, and the economy. This is likely due to the manner in which these topics are discussed – as education and training are seen as important for developing an AI-ready workforce. Topics 1 and 4 also have overlap, primarily due to the expectation that data and governance issues are related to contracts and property.

Now we present 5 word-clouds which show the universe of documents in our intra-national dataset (figures 9, 10, 11, 12 and 13).

In these word-clouds the font sizes indicate the documents which have a higher probability of possessing word distributions that constitute this topic. This means that these documents have a higher probability of containing this specific topic as a major theme.

Topic 1: Data and governance *Top 10 words*: component, dataset, algorithm, attack, external, section, internal, input, check, stage
The word-cloud presented for topic 1 has a large amount of diversity for countries publishing documents associated with data and governance. This is aligned with the results we found from our intra-national policy depth analysis findings. Data and governance are important basic foundations for AI infrastructures. Some documents which have high probabilities of containing word distributions associated with this topic include: Japan's Machine Learning Quality Guideline, China's Security Specification and Assessment Methods for ML Algorithms, EU's Independent High Level Expert Group on AI, the United State's NIST Principles for Explainable AI, and the EU's Robustness and Explainability of AI JRC Technical Report. Words associated with this topic include the following: algorithm, component, dataset, indicator, attack, check, fairness, reliability, and accuracy.

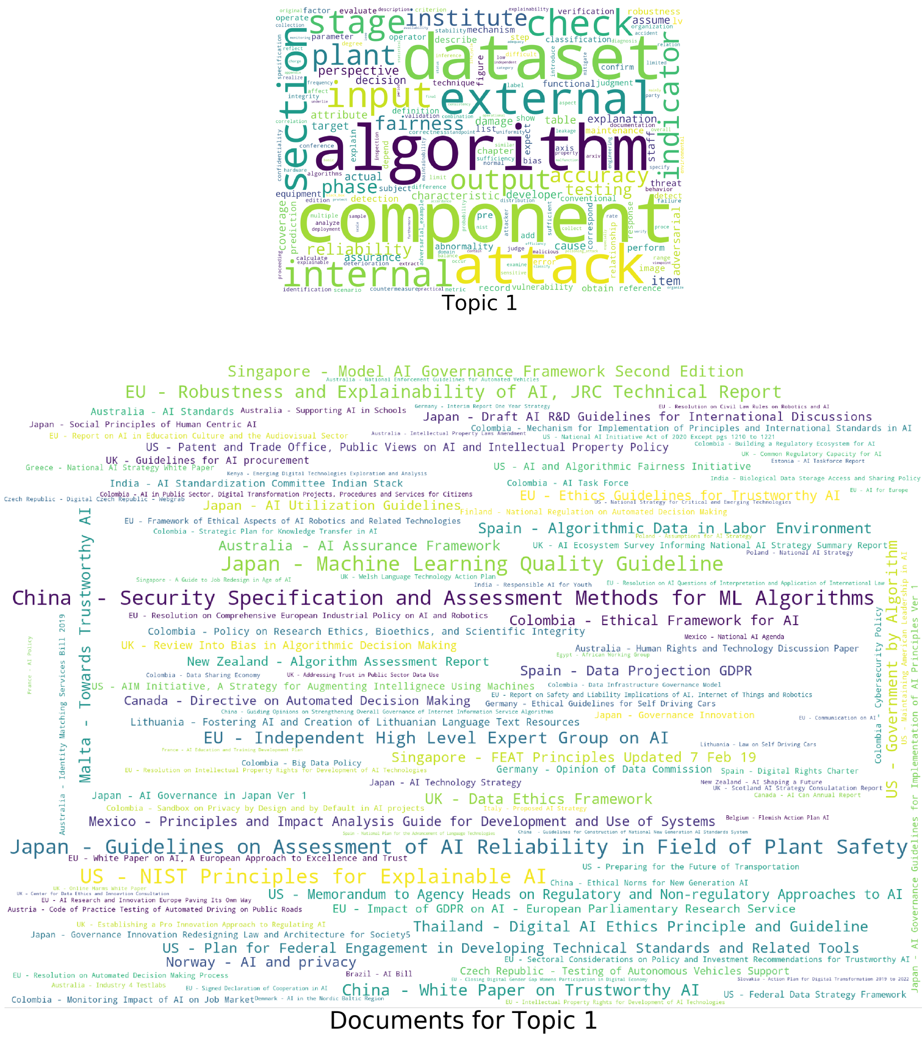

Figure 9: Top Panel: Datasets and Governance Topic Word Cloud (topic 1 out of 5 topics learned by the model over intra-national documents). Bottom Panel: Datasets and Governance Document Cloud relating the probabilities with which this topic emerges over countries (national and intra-national documents), utilizing larger font to visibly relate a document where the topic is most prominent (probabilistically).

Figure 10: Top Panel: Education and Training Topic Word Cloud.
Bottom Panel: Education and Training Document Cloud.

Figure 11: Top Panel: Economy Topic Word Cloud.
Bottom Panel: Economy Document Cloud.

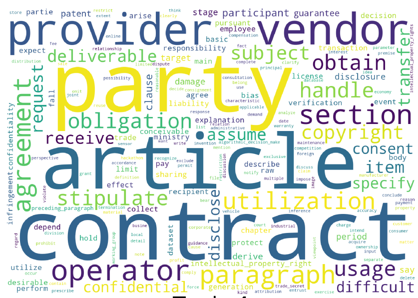

Topic 4

Documents for Topic 4

Figure 12: Top Panel: Contracts Topic Word Cloud.
Bottom Panel: Contracts Document Cloud.

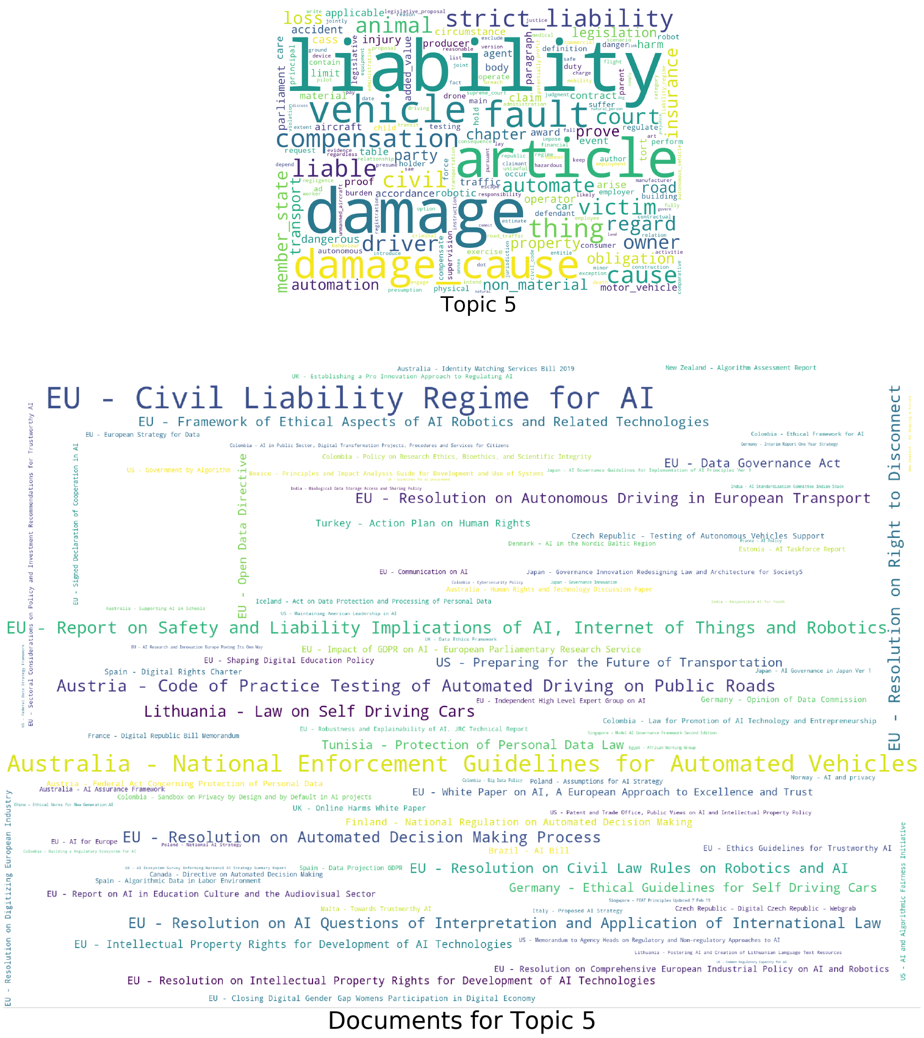

Figure 13: Top Panel: Transport Topic Word Cloud.
Bottom Panel: Transport Document Cloud.

Topic 2: Education and training *Top 10 words*: talent, university, enterprise, strengthen, major, algorithm, construction, basic, smart, school
In this word-cloud it is immediately apparent that documents from China have distributions of words that are associated with education and training. Key documents which are heavily associated with this topic are: China's AI Talent Training Report, China's AI Industry Talent Development Report (2019-2020), China's AI Innovation Action Plan for Institutions of Higher Education, China's Establishment of Ministry of Education AI Technology Innovation Expert Group, and China's Guidelines for Construction of National New Generation AI Standards System. In the associated topic word-cloud we can immediately notice words that tell us this topic is speaking to education and training – talent (specifically pertinent for China due to the 1000 Talent's Program), university, enterprise, strengthen, school, and industrial. In iterations our models prior to eLDA analysis, the Chinese approach to education encapsulated in the key word "talent" stood out as distinct from the other countries education strategies. We have reproduced word clouds here from those earlier iterations.

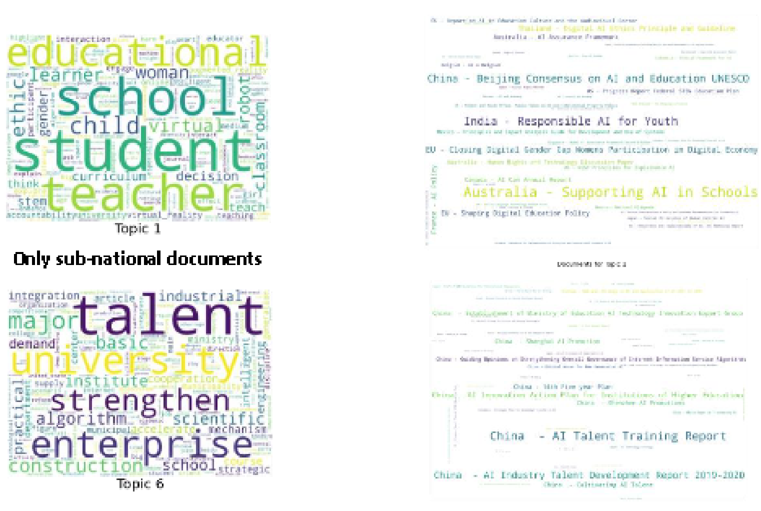

Figure 14: Education versus Talent Approaches

Topic 3: Economy *Top 10 words*: economy, infrastructure, job, growth, industrial, ministry, university, centre, fund, programme
In this topic there is a large range of countries publishing documents. This is noticeable due to the relatively similar size of many of the documents presented in the word-cloud. Some documents which have high probabilities of containing word distributions associated with the economy include: the United Kingdom's Council AI Road-map, the United Kingdom's Industrial Strategy Building A Britain Fit for Future, Kenya's Digital Economy Blueprint, Australia's Prosperity Through Innovation, Singapore's A Guide to Job Redesign in Age of AI, and Slovakia's Action Plan for Digital Transformation 2019 to 2022. The word-cloud associated with this topic thematically includes words such as: economy, growth, infrastructure, productivity, industrial, and job.

Topic 4: Contracts and liability *Top 10 words*: contract, party, article, provider, vendor, operator, section, paragraph, agreement, utilization
This topic is relatively diverse in terms of countries which are publishing documents containing word distributions associated with contracts and liability. It should be noted though that three of the largest documents associated with this topic are from Japan. They are Japan's Protection of Personal Information Act, Japan's Contact Guidelines on Utilization AI and Data Version 1, and Japan's Practical Guidebook on Providing Data for Employee Development in AI and Data Science. Other countries with documents with high probabilities of word distributions associated with this topic include: Iceland's Act on Data Protection and Processing of Personal Data, Australia's Intellectual Property Laws Amendment, and Colombia's Law for Promotion of AI Technology and Entrepreneurship. Some of the words associated with this topic are shown in the corresponding word-cloud. Words that are particularly important are: party, article, contract, vendor, operator, provider, obligation, usage, and agreement.

Topic 5: Transport *Top 10 words*: liability, damage, article, damage cause, vehicle, fault, compensation, thing, liable, strict liability
At first glance, one will notice that there are several European Union documents included in topic 5. This means that several EU documents contain word distributions that have a high probability of being related to the topic of transport. Some of the European Union documents include: EU's Civil Liability Regime for AI, EU's Report on Safety and Liability Implications of AI, Internet of Things and Robotics, and EU's Resolution on AI Questions of Interpretation and Application of International Law. Some other countries publishing documents related to this topic include Australia with their National Enforcement Guidelines for Automated Vehicles, Lithuania with their Law on Self Driving Cars, and Austria with their Code of Practice Testing of Automated Driving on Public Roads. Words associated with this topic include: liability, vehicle, fault, article, liable, damage, driver, damage cause, and compensation.

7 Conclusion

This policy report has provided a detailed analysis of 213 artificial intelligence policies, ranging from national AI plans to intra-national plans. Given the expanse of this analysis much can be said about the universe of AI policies and documents and thus how countries are developing AI and implementing AI technologies into a variety of government and societal systems. This report illustrates that there are convergences and divergences in the values and goals for national AI infrastructures around the world. The research illustrates that there is a wardrobe (or set) of policy mechanisms which are commonly used in AI infrastructures and that countries choose their own combinations of these mechanisms to best fit their needs. This is why we see distinct clusters of countries when it comes to the values embodied in their AI policies. To wrap up this report we provide three major takeaways from this research, which will be expanded upon in the coming years.

Three major takeaways:

- There are common elements among all national AI infrastructures, and overlaps can be seen across regional boundaries. These common elements illustrate that there are different levels of AI infrastructure development and aspects of diffusion – where countries may have learned from the experiences of other countries who have developed AI infrastructures before them.

- Despite this overlap, there are key differences among national AI plans – there are those which heavily emphasize matters of privacy, transparency, and accountability – while others many not articulate the importance for these protections and considerations. Some of these differences are due to the stage at which a country is in developing their AI infrastructure, while others may relate to political systems or aspirations. Despite this though, there is a need to focus on policy depths and regulatory or governance capabilities rather than unnecessary fear mongering about the state of AI around the world.

- Finally, while studying these national AI infrastructures it is helpful to utilize a comparative approach which allows for a deeper understanding of the universe in which AI technologies are being developed and deployed. Different approaches have influences on future iterations of AI policies and the use of AI in service provision. Therefore, utilizing methodologies which thoroughly analyze the universe of AI infrastructures and the values, goals, and implementation of them provide a more complete and realistic view of the state of AI as it is being innovated and harnessed in a variety of country contexts.

8 Team Biographies

J.P. is a Distinguished University Professor and Co-Director of the Center for Advancing Human-Machine Partnership at George Mason University. He is a Loomis Council member with the Stimson Center, and a Richard von Weizsäcker Fellow with the Robert Bosch Academy, Berlin. He works at the intersection of technology, culture and political economy in global contexts examining transformative impacts from provision of telephone service in poor countries, to the use of AI in global value chains in cutting-edge industries. J.P. has consulted or advised international organizations such as the British Council, UNESCO, the World Bank, and the World Trade Organization, and conducted field research in 36 countries. His current book project explores AI and innovation in Germany, India, and the United States. A winner of numerous research awards and fellowships, he has written 10 books and over 100 scholarly articles.

J.P. Singh
(Principal Investigator)

Amarda is a Professor of Computer Science, Associate Dean for AI Innovation, and Associate Vice President of Research for Institute of Digital InnovAtion. During 2019-2022, Amarda co-directed the Trans-disciplinary Center for Advancing Human-Machine Partnerships and served as an NSF Program Director in the CISE Directorate. Amarda considers herself an interdisciplinary scientist. Her research record includes foundational advances in AI, Machine Learning, and Algorithmics, and purposeful research that pushes the barriers of our understanding of physical world. Amarda has a galvanizing view of computing and relentless energy and advocacy for advancement of knowledge across scientific disciplines and improvement of human condition.

Amarda Shehu
(Co-Principal Investigator)

Jesse is a Research Associate Professor and the Acting Director of the Institute for Philosophy and Public Policy at George Mason University. He is also an International Security Fellow at New America and serves as a consultant for numerous organizations. His research is interdisciplinary, cutting across such fields as Philosophy, Political Science, Public Policy, and the Life and Computer Sciences. At its core, it aims to explore two central, interrelated themes: (1) how a suite of technologies, singularly and in convergence, impact peace and security, and (2) what the ethical, social, and policy implications of these impacts may be.

Jesse L Kirkpatrick

Michael is an associate professor at George Mason University's Schar School of Policy and Government, the associate director of the Schar School's Center for Security Policy Studies, and a Senior Non-Resident Fellow at the Center for Strategic and Budgetary Assessments. A former Marine Corps officer, he works on conventional deterrence, defense reform, and organizational learning.

Michael Hunzeker

Antonios is an assistant professor at George Mason University's Computer Science Department. He is a part of the George Mason Natural Language Processing Group and is interested in various aspects of multilingual Natural Language Processing and Machine Learning, with a main focus on Machine Translation and Speech Recognition for endangered languages and low-resource settings in general.

Antonios Anastasopoulos

Caroline is a PhD candidate and Presidential Scholar studying Political Science at George Mason University. She is currently working on her dissertation which focuses on science and innovation policy, with a focus on innovation clusters in different global contexts. Caroline holds a B.S. and M.S. in International Affairs from the Georgia Institute of Technology and has worked in research positions at RAND, the Center for Strategic and International Studies (CSIS), and the World Bank.

Caroline Wesson
(Project Manager)

Manpriya is a PhD student in the Computer Science Department at George Mason University. Her research centers Machine Learning, focusing on Deep Learning architectures for Natural Language Processing (NLP). She holds a Master's in Science degree in Computer Science from George Mason University, and has worked with Palo Alto Research Center on a DARPA project investigating use of NLP methods for military defense strategy.

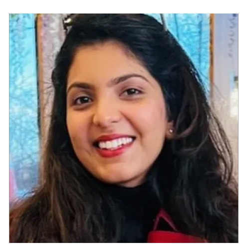

Manpriya Dua

Vasilii is a PhD in Public Policy student at Schar School and a Graduate Research Assistant at the Center for Advancing Human-Machine Partnership. His research interests lie in the intersection of economics, international relations, and institutional and technological transfer. Vasilii holds a Master of Science degree in Economics from Tufts University.

Vasilii Nosov

Webby is an MA student in Philosophy at George Mason University and a Graduate Research Assistant for the Institute for Philosophy and Public Policy. His research is mainly focused around two questions: how to make good decisions; and how to run a country well. Webby holds a BA in Philosophy and a minor in Economics from Wittenberg University. In addition to his research position, Webby has experience in multiple startup companies and is currently working on his thesis on philosophy and uncertainty.

William "Webby" Applegate

References

Aaronson, Susan Ariel. 2023. "Building trust in AI: a landscape analysis of government AI programs.".

Bareis, Jascha and Christian Katzenbach. 2022. "Talking AI into being: The narratives and imaginaries of national AI strategies and their performative politics." *Science, Technology, & Human Values* 47(5):855–881.

Blei, David M., Andrew Y. Ng and Michael I. Jordan. 2003. "Latent Dirichlet Allocation." *J. Mach. Learn. Res.* 3(null):993–1022.

CIFAR. 2017. "CIFAR pan- Canadian Artificial Intelligence Strategy 2017-20." `http://www.cifar.ca/ai/pan-canadian-artificial-intelligence-strategy`.

European Commission. 2018. *Artificial Intelligence for Europe.* Communication from the Commission to the European Parliament, the Council, the European Economic and Social Committee and the Committee of the Regions, European Commission, COM(2018) 237 final.

Foffano, Francesca, Teresa Scantamburlo and Atia Cortés. 2022. "Investing in AI for social good: an analysis of European national strategies." *AI & society* pp. 1–22.

German Federal Government. December 2020. "Artificial Intelligence Strategy of the German Federal Government.".

Hine, Emmie and Luciano Floridi. 2022. "Artificial intelligence with American values and Chinese characteristics: a comparative analysis of American and Chinese governmental AI policies." *"AI & Society"* pp. 1–22.

Nikulin, M. S. 2001. *Hellinger distance.* Encyclopedia of Mathematics EMS Press.

Niti Aayog. 2018. *Discussion Paper, National Strategy for Artificial Intelligence #aiforall.* Discussion Paper, Niti Aayog, Government of India.

Robinson, Stephen Cory. 2020. "Trust, transparency, and openness: How inclusion of cultural values shapes Nordic national public policy strategies for artificial intelligence (AI)." *Technology in Society* 63:101421.

Select Committee on Artificial Intelligence. 2016. *The National Artificial Intelligence Research and Development Strategic Plan.* National Science and Technology Council, Executive Office of the President of the United States.

"Select Committee on Artificial Intelligence". 2023. *The National Artificial Intelligence Research and Development Strategic Plan.* Select Committee on Artificial Intelligence, National Science and Technology Council, Executive Office of the President of the United States, 2023 Update.

State Council. 2017. *State Council Notice on the Issuance of the Next Generation Artificial Intelligence Development Plan (Trans. Rogier Creemers, Graham Webster, Paul Tsai, Paul Triolo, and Elsa Kania)*. State Council, China.

Wilson, Christopher. 2022. "Public engagement and AI: A values analysis of national strategies." *Government Information Quarterly* 39(1):101652.

A PRODUCT OF THE AI STRATEGIES TEAM AT
GEORGE MASON UNIVERSITY
HTTPS://WWW.AISTRATEGIES.GMU.EDU/

THIS RESEARCH IS SUPPORTED BY A $1.389 MILLION GRANT FROM THE MINERVA RESEARCH INITIATIVE. PLEASE SEND QUERIES ON THIS REPORT TO AIPOLICY@GMU.EDU.

Suggested citation:

Suggested citation: Singh, J.P., Amarda Shehu, Caroline Wesson, and Manpriya Dua. The 2023 Global Artificial Intelligence Infrastructures Report. With a Foreword from David Bray. AI Strategies Team and the Institute for Digital Innovation, George Mason University, and the Stimson Center, Washington DC. August 2023.