Diplomatic relations as a social network
Graph data analysis and female ambassador appointments

Bas Ernst

Table of Contents

Preliminary remarks

This is an elaboration of the final research project I completed for my BSc Computing and IT degree at the Open University in the UK. It contains a data analysis on diplomacy and gender. The purpose of the final module was to demonstrate the ability to independently conduct a project.

In this document, I wanted to highlight the process and results of the research itself; the OU thesis primarily requires an extensive reflection on the learning process. Consequently, all references to previous OU modules and how my research builds on and extends this knowledge have been omitted. Other required elements, such as critical reflections on time management, risk analysis, and a skills audit, have been omitted for the same reason. Since the OU enforces a strict word limit, I had to present the research in a concise manner.

Introduction

In summer 2024, I read about the GenDip dataset from the University of Gothenburg (Niklasson and Towns, 2023). This is a dataset of all diplomatic relations divided by gender in the period 1968–2021. The dataset was published with an invitation for scholars to analyse the data and enrich it with other sources.

Do states use female ambassadors' appointments as an instrument of diplomacy itself? This is the central question posed by the GenDip team. Some other questions that the researchers from Gothenburg formulated are:

- Are more women appointed as ambassadors to gender equality promoting countries?
- Do Greater Middle East countries send more female ambassadors to gender equality promoting countries?
- Did Hillary Clinton's tenure as Secretary of State attract female ambassadors to the USA?

This research aims to address these and related questions by analysing a dataset of diplomatic relations. It further explores alternative scenarios beyond the GenDip hypothesis and examines state behaviour through the structure of the diplomatic network.

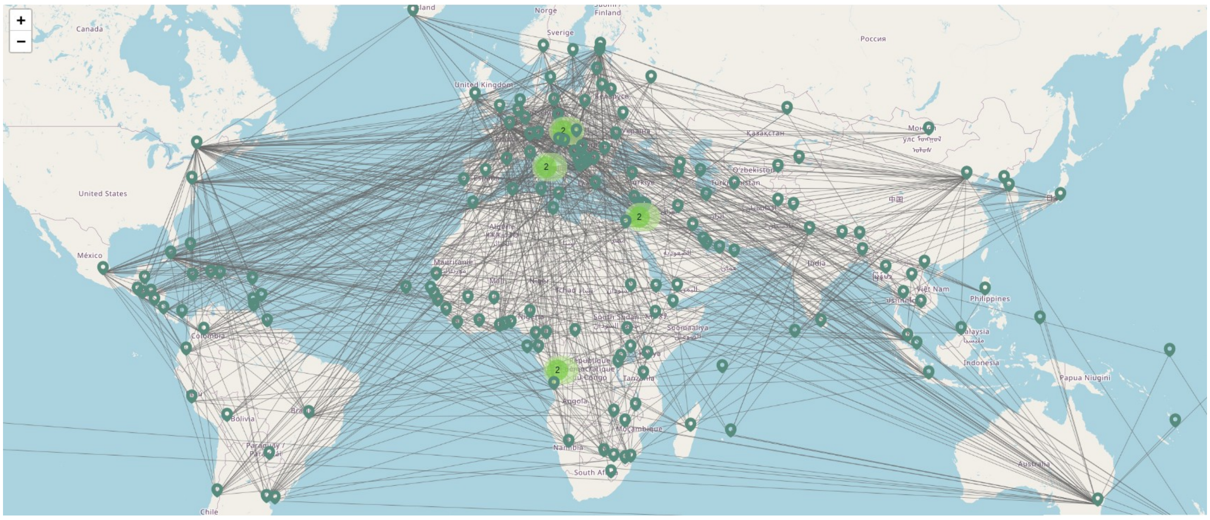

Figure 1: Part of the female diplomatic network 2021.

For this research, I will build a graph database for network analysis, and use graph data analysis to investigate factors related to the appointment of female ambassadors in hosting country. The process and findings of my research will be documented in this report.

This study consists of three phases and aims to answer 10 research questions:

A. Partial replication of the GenDip Research

- Prepare and clean data and setting up a graph database
- Analyse the data in order to answer three research questions of the GenDip research team.

B. Integration of External Datasets

- Create my own datasets to analyse the presence of female political leaders in foreign policy roles and ambassadorial postings.
- Incorporate external data to examine additional factors influencing female ambassador appointments: hardship, gender inequality index and religion.
- Use hypothesis testing to validate the outcomes of the analysis.

C. Graph Data Science (GDS) Analysis

- Apply Neo4j's Graph Data Science tools to analyse the structure of the diplomatic network and its behaviour.
- Apply relevant centrality and community detection algorithms to identify influence of countries and existence of communities in diplomatic networks.
- Detect possibly interesting anomalies to explore strategic ambassadorial appointments.

Methods

Understanding the problem domain: Gender and Diplomacy

The problem domain's key sources for understanding female ambassador appointments include the GenDip research, a book on research methods in international relations, and peer-reviewed articles on diplomacy and gender, all of which are available at the OU library.

These are some of the main insights from the literature that have guided this research:

- Towns and Niklasson (2017) examined the under-representation of female diplomats in host countries, considering factors such as economic and military power. They concluded that many studies show women's representation falls as the prestige of a position increases.
- However, their later publication stated that in diplomacy, a state's decision to appoint women as ambassadors may depend on how strongly the receiving state values gender equality. (Niklasson and Towns, 2023)
- Duque (2018) argues that status in international relations is a function of social recognition. In her research she prioritises the study of the network dynamics over analysing state's attributes as drivers for international recognition.
- International relations can be considered as a social network and diplomatic relations can be modelled as a graph, with countries as nodes and their diplomatic ties as edges forming the diplomatic network (Menninga and Goldberg, 2022).

Therefore, this study focuses on gender and diplomacy, intending to analyse socio-economic factors, as well as the structure and behaviour of the diplomatic network, and paying particular attention to states that strategically appoint female ambassadors.

Supporting the Project Work: Graph Data Analysis

Most of this research relates to the IT technology required for the project work. My sources are mainly recent textbooks and documentation written by experts in the field. Most of these sources are taken from the OU library, while software documentation is used from the relevant websites (for example Neo4j and Wikidata) and tutorials from renowned blogs, such as Medium and Towards Data Science.

Relevant algorithms and social network analysis

Redmond and Wilson (2012) describe graph databases as being particularly well-suited to analysing linked data and social networks. They are designed to store interconnected data using nodes (representing entities such as persons or countries) and edges (representing relationships between the nodes), both of which can hold properties. Although graph theory provides a systematic, quantitative foundation for studying social relationships (Golędzinowski & Błocki, 2023), this project only applies selected graph concepts.

Networks can take on a few archetypal forms. This study focuses on a small-world network. Such networks are characterised by two features (Needham & Hodler, 2019; Marchiori & Latora, 2000; Sung, 2025):

- Small diameter. The diameter is the longest shortest path between nodes (Meghanathan, 2017).
- High clustering coefficient: This describes the likelihood that the neighbours of a country node are also connected (Neo4j documentation, n.d.). The higher the number, the denser the graph.

Similar small-world examples include social media networks and transport hubs, such as airports. The diameter of these networks is, on average, between 5 and 6 steps. (Needham & Hodler, 2019).

Two further concepts are important to understand the structure of a network:

Centrality
Not all nodes are equal; some are more important than others, acting as bridges or disseminators. Centrality, understood as the influence of nodes, can be analysed in several ways (Meghanathan, 2017):

- Degree centrality: The number of directly connected nodes, incoming or outgoing.
- Betweenness centrality: Prioritises nodes that are on connecting paths between other nodes.
- Closeness centrality: Prioritises nodes that are close to other nodes.
- PageRank: An algorithm that considers the importance of countries based on the importance of neighbours and their neighbours in the network.

Figure 2 provides a visual summary of the these concepts.

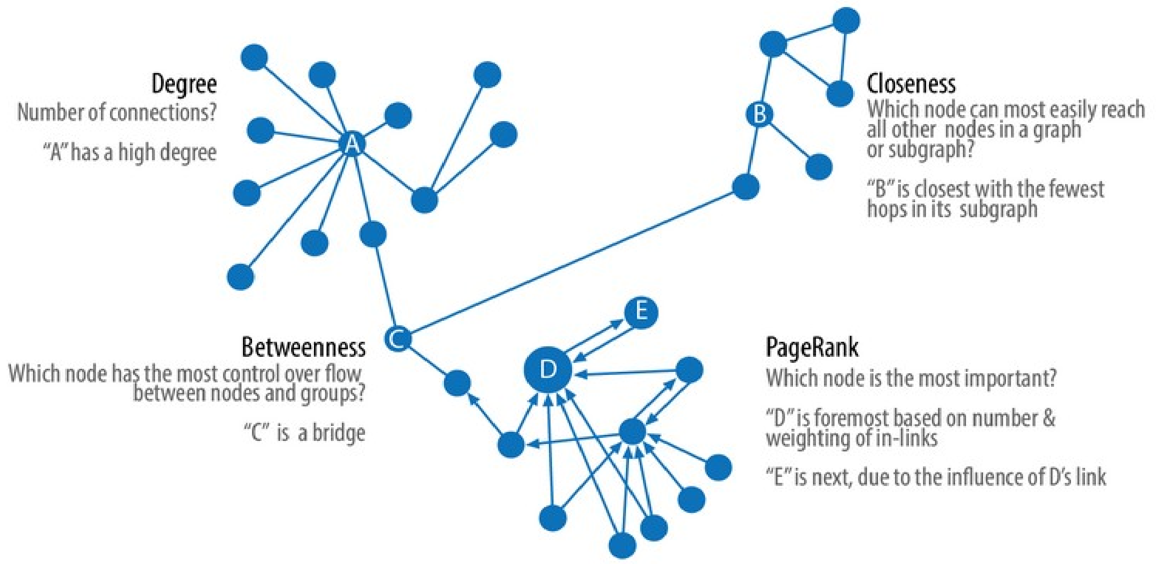

Figure 2: Summary of centrality algorithms. Taken from Needham and Hodler, figure 5-1 (2019)

Communities
Networks can contain groupings of nodes that are more closely connected to each other than to other nodes. There are various ways of analysing the presence and composition of these communities. Here, I will adopt the approach of Needham and Hodler (2019). They focus on:

- Triangles. A triangle has three fully connected nodes (Scifo, 2023). A small-world network will have a high number of triangles.

- Connected components. These are clusters of nodes only connected to themselves and not to any other nodes. In a small-world scenario, there will be just one or a limited number of connected components (Needham and Hodler, 2019).
- Louvain algorithm. The Louvain algorithm is based on the concept of modularity which means it iteratively analyses whether nodes have more connections with similar nodes than would be expected in a random graph (Scifo, 2023; Smith, 2024).

Graph query languages

Nodes and edges are stored in a dedicated graph database. Graph databases fall under the category of NoSQL systems, offering greater flexibility than traditional relational databases. Mastering query languages is essential for graph databases. I studied Ducharme (2011) on SPARQL and Scifo (2020) on Cypher, and used Wikidata and Neo4j documentation, plus tutorials and expert blogs (e.g., Webber, 2024), to explore practical applications.

SPARQL is a query language for RDF (Resource Description Framework) databases, which underpin the Semantic Web and Linked Open Data. In RDF, entities like countries or ambassadors are connected through defined properties (e.g. ‘has diplomatic relations with’). A central element is the ontology, which formally defines the structure of these relationships and the meaning of classes. Understanding the ontology is essential for effective SPARQL queries (Webber, 2024).

Cypher, by contrast, is the query language for Neo4j, a graph database using the labelled-property graph (LPG) model (Webber, 2024; Scifo, 2020). Unlike RDF, it does not rely on an ontology. Nodes and edges are labelled and can have properties, and queries use pattern matching, allowing direct exploration of node relationships.

Hypothesis testing

This research examined the correlation between the appointment of female ambassadors and various socio-economic indicators. I used hypothesis testing to determine the statistical significance of these relationships. Different statistical methods are used depending on the nature of the variables involved. I studied the handbook by Cohen et al. (2018) and used blogs like Statology to find practical examples of applying statistical methods in Python. For this research, I will use two categories of data:

- Nominal data: data with no inherent order (e.g. gender of ambassadors, regions, religions).
- Ratio-scale data: ordered data with equal intervals and a meaningful zero point (e.g. in-degree, percentages of female ambassadors).

The combination of these types of data determines which hypothesis test must be used:

- Fisher’s Exact Test is used to analyse the statistically significant relationship between two nominal variables, for example gender versus region.
- For ratio-scale variables (e.g. in-degree vs. hardship scores), Pearson's r is used to measure the correlation.
- The Kruskal–Wallis test is appropriate for comparisons of continuous and categorical data, for example religion versus in-degree (Bobbit, 2022).

Analysis begins with a null hypothesis, assuming no relationship between variables, with a significance threshold (α) of 0.05 (Cohen et al., 2018). Values below α indicate a statistically significant relationship.

Resources

Data

The GenDip dataset was published in 2023[1]. It is a CSV file containing data on countries, regions, gender, main or side postings, country codes and types of head of mission for 10 individual years spanning the period from 1968 to 2021 (figure 3).

```python
df = pd.read_csv("/home/bas/Nextcloud/Bas/gender_diplomacy/gender_diplomacy_to_clean.csv")
df.head()
```

Python

	year	cname_send	main_posting	title	gender	cname_receive	ccode_send	ccodealp_send	ccodeCOW_send	region_send	GME_send	v2lgfemleg_send	FFP_send	ccode_receive	ccodealp_receive	ccodeCOW_receive	region_receive	GME_receive	FFP_receive
0	1968	Afghanistan	1	3	0	China	4	AFG	700	1	1	2.0	0	156	CHN	710	1	0	0
1	1968	Afghanistan	1	3	0	Czechoslovakia	4	AFG	700	1	1	2.0	0	200	CSK	315	3	0	0
2	1968	Afghanistan	1	3	0	Egypt	4	AFG	700	1	1	2.0	0	818	EGY	651	4	1	0
3	1968	Afghanistan	1	3	0	France	4	AFG	700	1	1	2.0	0	250	FRA	220	3	0	0
4	1968	Afghanistan	1	3	0	Germany, Federal Republic of	4	AFG	700	1	1	2.0	0	280	DEU	260	3	0	0

```python
# remove leading and trailing whote spaces
df['ccodealp_receive'] = df['ccodealp_receive'].str.strip()
df['ccodealp_send'] = df['ccodealp_send'].str.strip()
df.shape
```

Python

```
(94509, 19)
```

Figure 3: First records of the GenDip dataset

In Part B, I built datasets and merged external open-source datasets into the database. The data-sources are from renowned international research institutes and international organisations. The geographical data is taken from a data scientist's blog:

1. Heads of State and Prime Ministers, *Wikidata*
2. Ministers of Foreign Affairs, *Wikipedia*, "List of foreign ministers in 2021"
3. Hardship data, *U.S. Department of State* – Hardship Post Allowance page
4. Coordinates Capital Cities, *Jasom Dotnet (this geo-data is used for visualisations on maps)*
5. Gender Inequality Index (GII), *United Nations Development Programme (UNDP)*
6. Religious Affiliation, *Pew Research Center*

The links to the datasets and information about the licenses are added to the reference list.

Software

I used the following software solutions for the research:.

- Neo4j: Cypher for queries and interactions with the database, Bloom and especially NeoDash for visualisations.
- Python: Pandas for data analysis, BeautifulSoup for web-scraping and Seaborn for visualisations.

[1] Although the research has been mentioned in a few articles, I am not aware of any other published analysis of this dataset.

- RDF: SPARQL for queries

Results

Acquiring and preparing the GenDip and additional datasets

Before analysing the data, I did some data cleaning:

- converting datatypes, eliminating null values (e.g. figure 4)
- renaming values for practical (reducing long names) or ethical reasons (politically neutral naming)
- eliminating irrelevant data but ensuring no crucial nodes or attributes were lost
- checking for outliers
- ensuring that relevant properties could be used as IDs when merging datasets

Region	Country	Year	Population	Christians	Muslims	Religiously_unaffiliated	Buddhists	Hindus	Jews	Other_religions	Countrycode
Asia-Pacific	Afghanistan	2020	39,070,000	< 10,000	39,020,000	< 10,000	< 10,000	< 10,000	< 10,000	40	4
Europe	Albania	2020	2,870,000	510	2,140,000	220	< 10,000	< 10,000	< 10,000	< 10,000	8
Middle East-North Africa	Algeria	2020	44,040,000	130	43,330,000	560	< 10,000	< 10,000	< 10,000	20	12
Sub-Saharan Africa	Angola	2020	33,450,000	31,120,000	90	2,050,000	< 10,000	< 10,000	< 10,000	190	24
Latin America-Caribbean	Argentina	2020	45,190,000	39,970,000	420	4,170,000	10	< 10,000	170	440	32
Asia-Pacific	Armenia	2020	2,890,000	2,810,000	< 10,000	30	< 10,000	< 10,000	< 10,000	40	51
Latin America-Caribbean	Aruba	2020	110	90	< 10,000	< 10,000	< 10,000	< 10,000	< 10,000	10	533
Asia-Pacific	Australia	2020	25,740,000	12,040,000	900	10,900,000	670	760	110	360	36
Europe	Austria	2020	8,920,000	6,080,000	740	2,000,000	30	10	< 10,000	60	40
Asia-Pacific	Azerbaijan	2020	10,180,000	40	9,640,000	480	< 10,000	< 10,000	< 10,000	< 10,000	31
Latin America-Caribbean	Bahamas	2020	400	390	< 10,000	< 10,000	< 10,000	< 10,000	< 10,000	< 10,000	44

Figure 4: Example of a dataset that needed a lot of cleaning: the PEW research on religion.

Datamodel

The key to efficiently analysing data is finding the right approach to storing it. This provided guidance for useful labels and properties for the nodes and edges (Alexopoulos, 2020). My aim was to design a data model that represented the problem domain and its context. I started with a detailed model of country nodes, diplomatic relations and additional nodes to model various variables. However, working with this model required rather complex queries. It quickly became clear that a much simpler model would improve efficiency when querying the data (figure 11). Rather than adding nodes, I added properties to the nodes and edges. This suited my needs better.

Some insights into the dataset

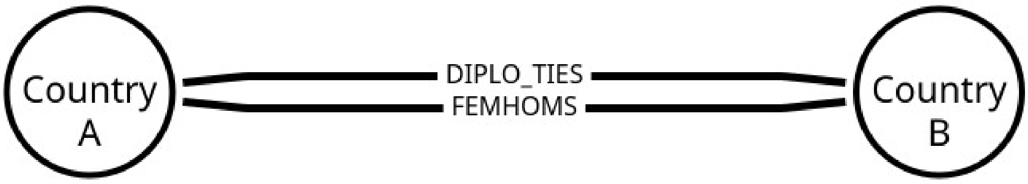

Figure 5: Updated and simpler data model, 208 country nodes, and two types of edges: DIPLO_TIES (>94k edges) and FEMHOMS (>4k). Both nodes and edges have various properties (variables).

This is an overview of the situation regarding female ambassadors. The table below shows how female, male and gender-unknown ambassadors are distributed in the dataset (counting only main postings).

Table 1: Distribution ambassadors based on gender classification

Female	Male	Gender unknown
8,262	61,907	3,873

As shown in figure 6, when we consider the male, gender-unknown and female ambassadors together, the under-representation of women becomes evident.

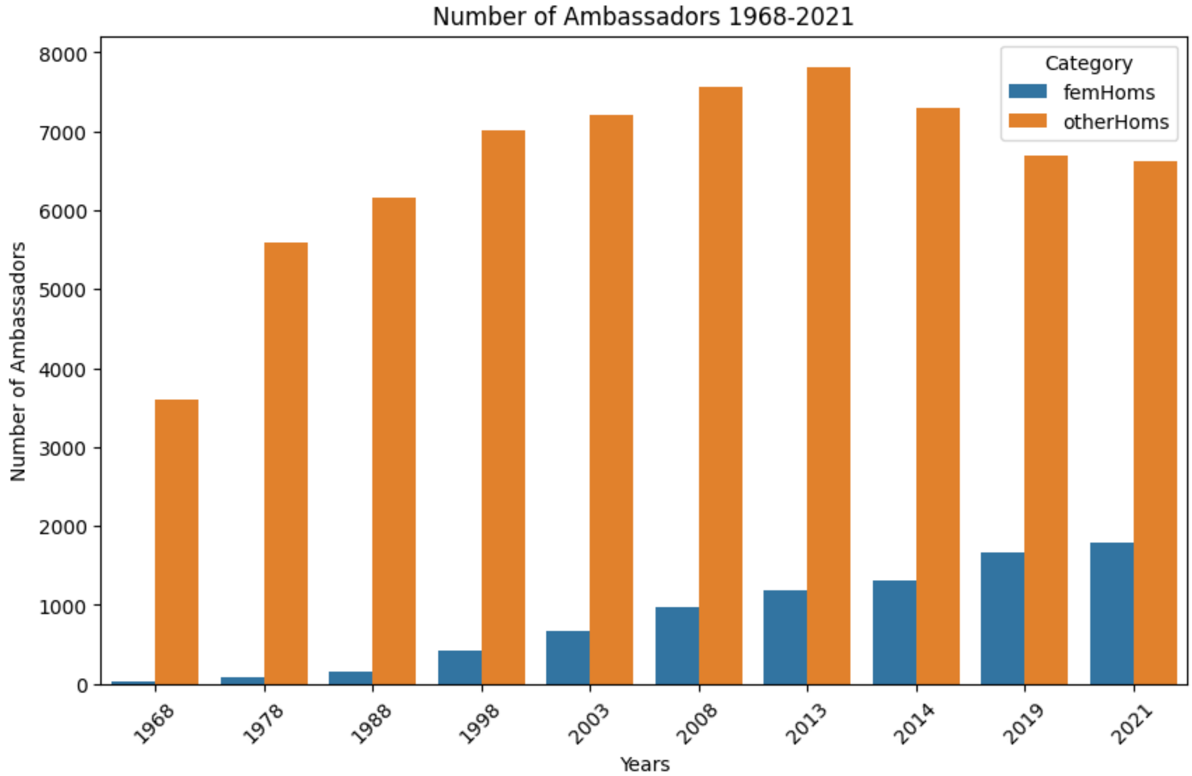

Figure 6: Systematic under-representation of women: an overview of the male and gender unknown ambassadors in orange and the female heads of missions in blue.

Another way to consider female representation is to count the number of female diplomatic postings per year for each diplomatic relationship.

There are 4,261 unique female diplomatic ties between the countries. The graph below shows that 76% of these ties were established just once or twice. (figure 7)

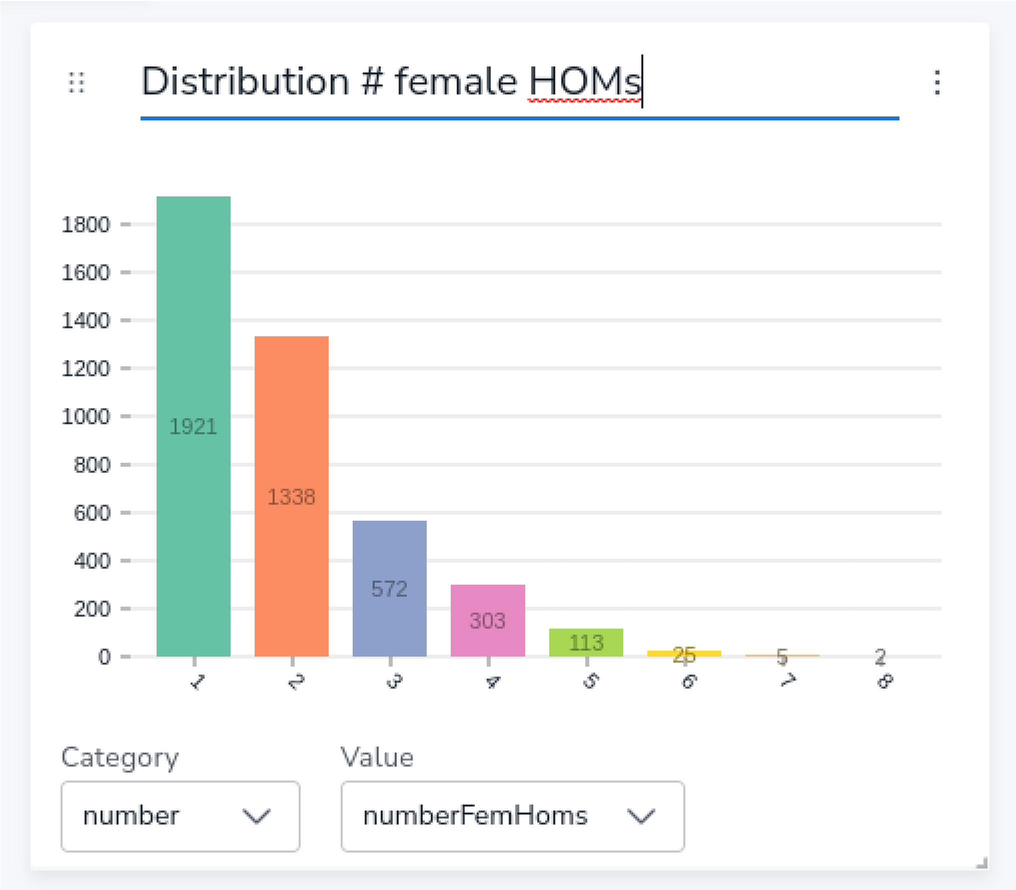

Figure 7: Frequency of female ambassadors per diplomatic relation.
HOMs = Heads of Mission

Figure 8 below shows the countries that had most consistent diplomatic relations with a female ambassador during the research period.

s.name	r.name	f.femHoms
"San Marino"	"Italy"	8
"Dominican Republic"	"Sweden"	8
"Cuba"	"Sri Lanka"	7
"Mozambique"	"Belgium"	7
"Finland"	"Namibia"	7
"Canada"	"New Zealand"	7
"Jamaica"	"Germany"	7
"Congo, Democratic Republic of the"	"Côte d'Ivoire"	6
"Australia"	"Croatia"	6
"Canada"	"Poland"	6

Figure 8: The diplomatic relations with the highest continuous female presence.

I can conclude that, although the number of female ambassadors is growing, women remain structurally under-represented in diplomatic top positions and there are just a few diplomatic relations with a continuous female presence. This confirms the findings by Towns and Niklasson (2017), which were presented in my introduction.

I will now address the ten research questions. I following two approaches:

- I performed Cypher queries directly in the Neo4j browser to explore the data and to ensure that the expected results were returned.
- I used Pandas with the Neo4j library. This enabled me to import the results of a Cypher query directly into a Pandas dataframe for further processing and visualisation. This combination proved to be particularly powerful.

Part A

Q1. Do women make up a larger share of ambassadors posted to known international G*ender Equality Promoters GEP*[2] than to other states?

This question requires the calculation of the in-degree of female ambassadors per country, this represents the number of incoming ambassadors. The bar chart below shows the top 25 countries with the highest number of female ambassadors. Eight of these countries are GEP, while the others are non-GEP (figure 9).

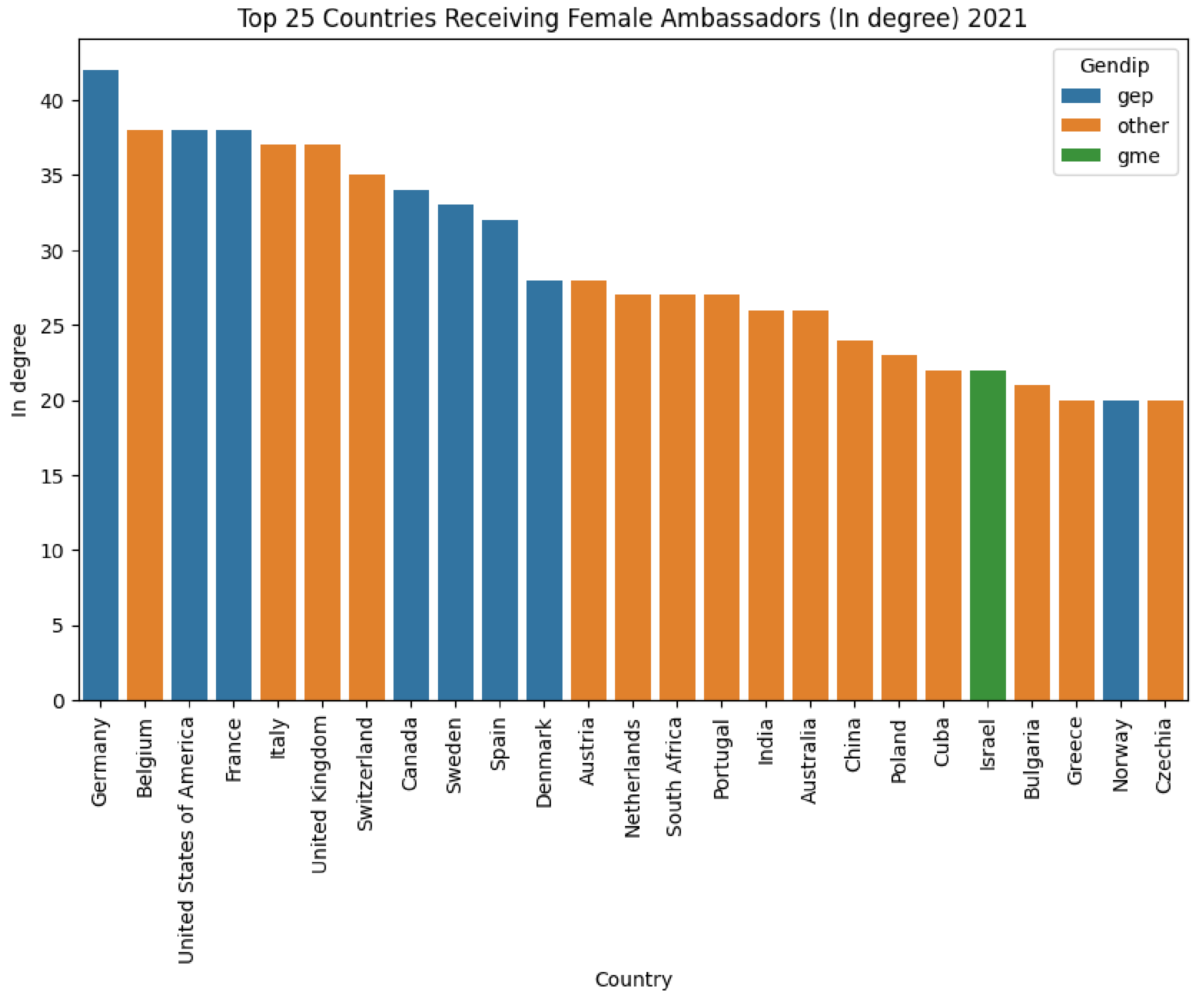

Figure 9: Top 25 countries receiving female ambassadors (in-degree) 2021.

Then, I calculated the average percentage of female ambassadors in both GEP and non-GEP countries. (Table 2).

[2] Norway, Denmark, Iceland, Finland, Sweden, United States of America, Canada, France, Mexico, Spain, Germany, Luxembourg

Table 2: Percentage of female ambassadors to GEP countries 2021

	GEP countries	Non-GEP countries
Percentage female ambassadors	30.5%	23.7%

This seemed to confirm the initial question. However, when testing the hypothesis I found a t-value of 1.55, indicating that the means for the two groups are fairly close (figure 8). The p-value is 0.123, meaning that the null hypothesis cannot be rejected. Unfortunately, there is no significant statistical difference between female ambassadors sent to GEP and non-GEP countries.

```
# Perform the two sample t-test with equal variances
stats.ttest_ind(a=(gep_df["percent"]), b=(nongep_df["percent"]), equal_var=True)

[6]  ✓ 0.0s                                                                 Python
...  TtestResult(statistic=np.float64(1.5491167520099258), pvalue=np.float64(0.12309078191458035), df=np.float64(182.0))
```

Figure 10: Hypothesis testing Question 1

As a reflection, I had hoped for statistical confirmation of the results obtained. I fear there is too much variation within the two groups to consider them as distinct communities. These findings cover 2021 only; in the conclusions of Part B, I will also include an analysis of the entire dataset.

Q2 Do states in the *Greater Middle East GME*[3] post a larger share of their female ambassadors to host states with a gender equality profile than do other states?

After querying the data in Cypher, I used Pandas to calculate the percentage of female ambassadors from different regions who were sent to GEP countries. In contrast with the GenDip team I included the Nordics also in Europe and GME countries in respectively Asia and Africa.

Figure 11 shows the regional differences:

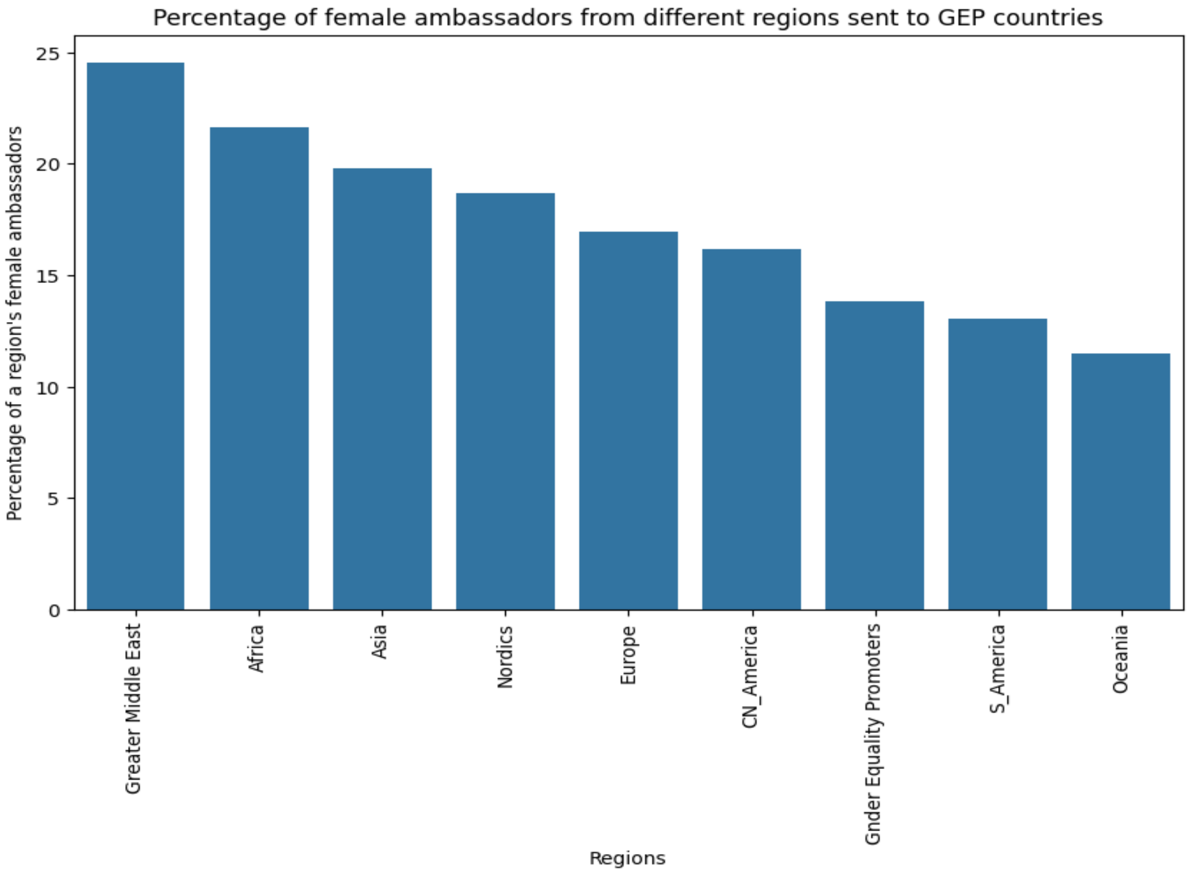

Figure 11: Percentage of female ambassadors sent from various regions to GEP countries in 2021

It seems that the largest share (24.55%) of female ambassadors sent to GEP countries came from the Greater Middle East. However, I calculated a Kruskal-Wallis H-test score of 8.0 (Figure 12), which is relatively low for 9 groups and a p-value of 0.433. Therefore, it is not possible to reject the null hypothesis, meaning that there is no statistically significant difference in the percentages of female ambassadors across regions.

[3] Afghanistan, Algeria, Armenia, Azerbaijan, Bahrain, Egypt, Georgia, Iraq, Iran, Israel, Jordan, Kazakhstan, Kuwait, Kyrgyzstan, Lebanon, Libya, Morocco, Oman, Pakistan, Palestine, Qatar, Saudi Arabia, Somalia, Sudan, Syria, Tunisia, Turkey, Turkmenistan, Tajikistan, United Arab Emirates, United Arab Republic, Uzbekistan, Yemen.

```
gme = result.loc[result['region'] == 'Greater Middle East', 'percent']
gep = result.loc[result['region'] == 'Gnder Equality Promoters', 'percent']
nordics = result.loc[result['region'] == 'Nordics', 'percent']
africa = result.loc[result['region'] == 'Africa', 'percent']
asia = result.loc[result['region'] == 'Asia', 'percent']
europe = result.loc[result['region'] == 'Europe', 'percent']
cn_america = result.loc[result['region'] == 'CN_America', 'percent']
s_america = result.loc[result['region'] == 'S_America', 'percent']
oceania = result.loc[result['region'] == 'Oceania', 'percent']

kruskal(gme,gep,nordics,africa,asia,europe, cn_america, s_america,oceania)
```

✓ 0.0s

```
KruskalResult(statistic=np.float64(8.0), pvalue=np.float64(0.43347012036670896))
```

Figure 12: Kruskal's H-score regions sending female ambassadors to GEP countries

Unfortunately, regional differences cannot explain ambassadorial decisions. This suggests that other variables must be analysed to explain the under-representation of women.

Q3 Is there a presumed “Hillary Effect”, influencing the nomination of incoming female heads of mission?

Answering this question is relatively straightforward. It requires just counting the number of incoming female ambassadors to the USA for each year present in the dataset (figure 13).

```
query= "MATCH (s:Country)-[d:DIPLO_TIES]->(r:Country{code: 'USA'}) WHERE d.main_posting = 1 AND d.gender = 1 RETURN d.year AS Year, count(d) AS femaleHoMs ORDER BY Year ;"

with driver.session() as session:
    result = session.run(query)
    clinton_era = DataFrame(result)

clinton_era.columns = ["Year",  "totalFemHoms"]
clinton_era
```

	Year	totalFemHoms
0	1978	1
1	1988	4
2	1998	13
3	2003	18
4	2008	27
5	2013	26
6	2014	24
7	2019	28
8	2021	38

Figure 13: Calculating total female ambassadors in the USA per year

After that, I created with Seaborn a line chart showing the increasing number of female ambassadors over the years (Figure 14). Although the number stalled during Clinton’s tenure (2008–2013) and declined slightly after she left office, significant increases occurred during the terms of Madeleine Albright (1997–2001) and Condoleezza Rice (2005–2009).

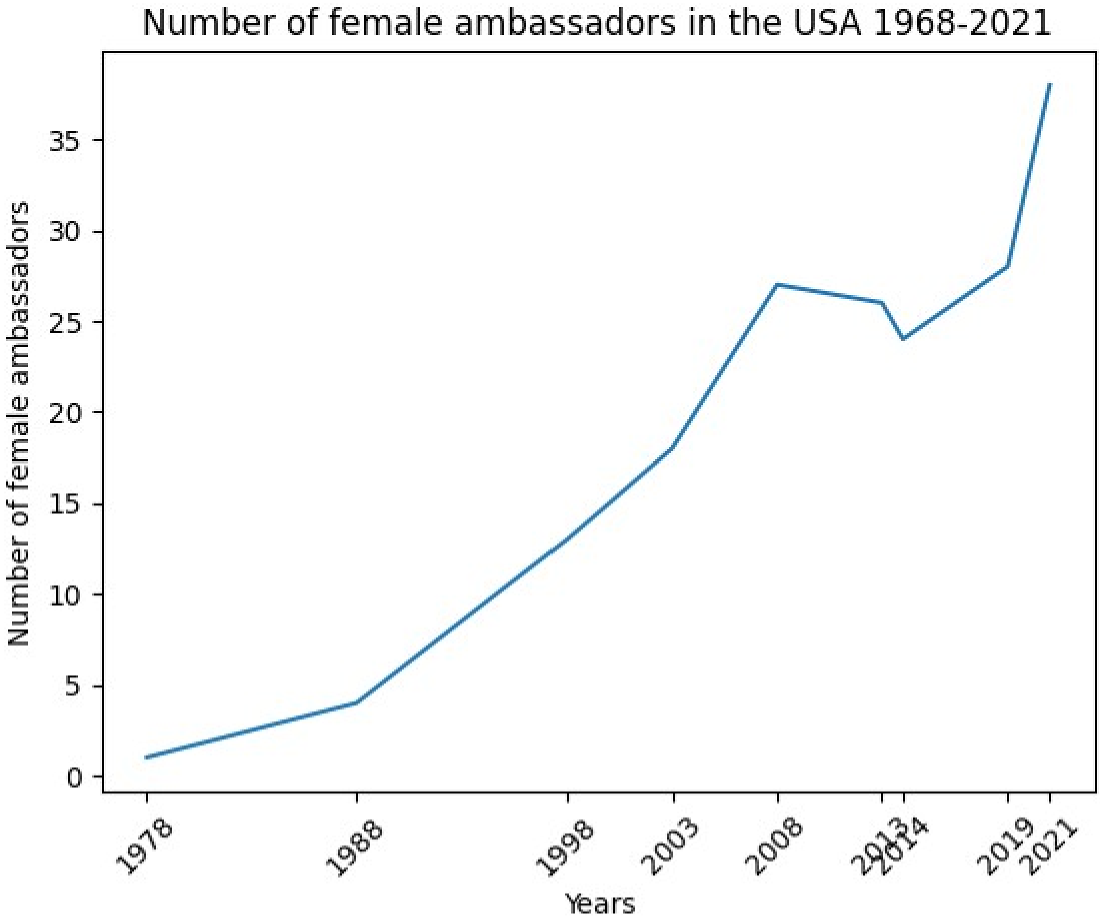

Figure 14: Line graph of incoming female ambassadors in the USA.

As Niklasson and Towns (2023) anticipated, there was no Hillary effect, but the results raise the question whether the presence of female politicians in the receiving country influences the appointment of female ambassadors. This inspired me to create datasets on female politicians. This is the subject of question 4.

Conclusions Part A

My analysis confirms the GenDip team's findings regarding the under-representation of women and demonstrates that countries make decisions where to send female ambassadors.
However, the data analysis did not find any statistical evidence to support the first two questions relating to regions, as formulated by the GenDip team. There is not enough statistical evidence to confirm:

- that more women are appointed as ambassadors to gender equality promoting countries.
- nor do Greater Middle East countries send more female ambassadors to gender equality promoting countries.

Therefore, I will investigate further whether additional data can provide alternative explanations (Part B) or whether the network structure itself can explain the findings (Part C).

Part B

Q4. Are countries sending more female than male ambassadors to countries that have a female head of state, prime minister or minister of foreign affairs?

To answer this question, I applied two techniques:

- RDF to query Wikidata and find female heads of state and prime ministers
- Web-scraping with BeautifulSoup to find ministers of foreign affairs in Wikipedia

Both operations returned a list of countries with a distinct value indicating the presence of one of the three possible female politicians. I then merged the datasets into the database by adding a property to each country node to show whether it had at least one woman in any of the three roles in 2021. I analysed whether countries send more female than male ambassadors to host countries with a female head of state, prime minister, or minister of foreign affairs.

Summarising the results in one table (figure 15) reveals a clear relationship between women in leading roles in foreign affairs and a higher presence of female ambassadors.

Role in receiving country	Gender incoming ambassadors	Percentages
Male PM	5184 male HOMs	79.4%
	1344 female HOMs	20.6%
Female PM	538 male HOMs	73.0%
	199 female HOMs	27.0%

Role in receiving country	Gender incoming ambassadors	Percentages
Male MFA	4769 male HOMs	79.6%
	1203 female HOMs	20.4%
Female MFA	1737 male HOMs	75.8%
	552 female HOMs	24.2%

Role in receiving country	Gender incoming ambassadors	Percentages
Male HOS	5184 male HOMs	79.4%
	1352 female HOMs	20.6%
Female HOS	524 male HOMs	75.6%
	188 female HOMs	24.4%

Figure 15: Overview of male and female ambassadors sent to countries with or without female politicians in key foreign affairs roles. HOS = head of state, PM = prime minister, MFA = minister of foreign affairs, HOM = head of mission

Finally I performed a Fisher's exact to test the hypothesis. The results show that, in all three cases, the null hypothesis is not supported (figure 16). Therefore, I can conclude that female ambassadors are more likely to be sent to countries where female politicians hold key roles in foreign affairs. However, this does not imply a causal relationship.

```
In [41]: # first MFA

         # select data
         gender_effectsMFA = [gender_effects.iloc[9,1], gender_effects.iloc[11,1]],\
                             [gender_effects.iloc[8,1], gender_effects.iloc[10,1]]
         # gender_effectsMFA

         # performing fishers exact test on the data
         odd_ratio, p_value = stats.fisher_exact(gender_effectsMFA)
         print('odd ratio is : ' + str(odd_ratio))
         print('p_value is : ' + str(p_value))

         odd ratio is : 1.25953395153452
         p_value is : 9.308834386247624e-05

In [42]: # PM's


         # select data
         gender_effectsPM = [gender_effects.iloc[5,1], gender_effects.iloc[7,1]],\
                            [gender_effects.iloc[4,1], gender_effects.iloc[6,1]]
         gender_effectsPM

         # performing fishers exact test on the data
         odd_ratio, p_value = stats.fisher_exact(gender_effectsPM)
         print('odd ratio is : ' + str(odd_ratio))
         print('p_value is : ' + str(p_value))

         odd ratio is : 1.4264374778721898
         p_value is : 9.466985343145293e-05

In [43]: # Heads of state

         # select data
         gender_effectsHOS = [gender_effects.iloc[1,1], gender_effects.iloc[3,1]],\
                             [gender_effects.iloc[0,1], gender_effects.iloc[2,1]]
         gender_effectsHOS

         # performing fishers exact test on the data
         odd_ratio, p_value = stats.fisher_exact(gender_effectsHOS)
         print('odd ratio is : ' + str(odd_ratio))
         print('p_value is : ' + str(p_value))

         odd ratio is : 1.375406522426487
         p_value is : 0.0006015852012563941
```

Figure 16: Hypothesis tests confirm there are more female ambassadors in all three scenarios: The odds-ratio heads of state: 1.26, prime ministers: 1.43, ministers of foreign affairs: 1.38 and p-values: HOS: 0.000093, PM: 0.000095, MFA: 0.00060. The null hypothesis is not supported.

Q5 Do living conditions matter as a pull factor for female ambassadors?

For this analysis, I used the US State Department's hardship classification, which takes into account factors such as security, healthcare, and living conditions abroad (U.S. Department of State, n.d.). This dataset, provide national but also regional scores. Given that local conditions vary from country to country, I calculated national averages to provide a comparable metric for each one.

Then I added the hardship values to the country nodes in Neo4j. In NeoDash I used a query to visualise the level of hardship for the countries in the dataset (figure 17).

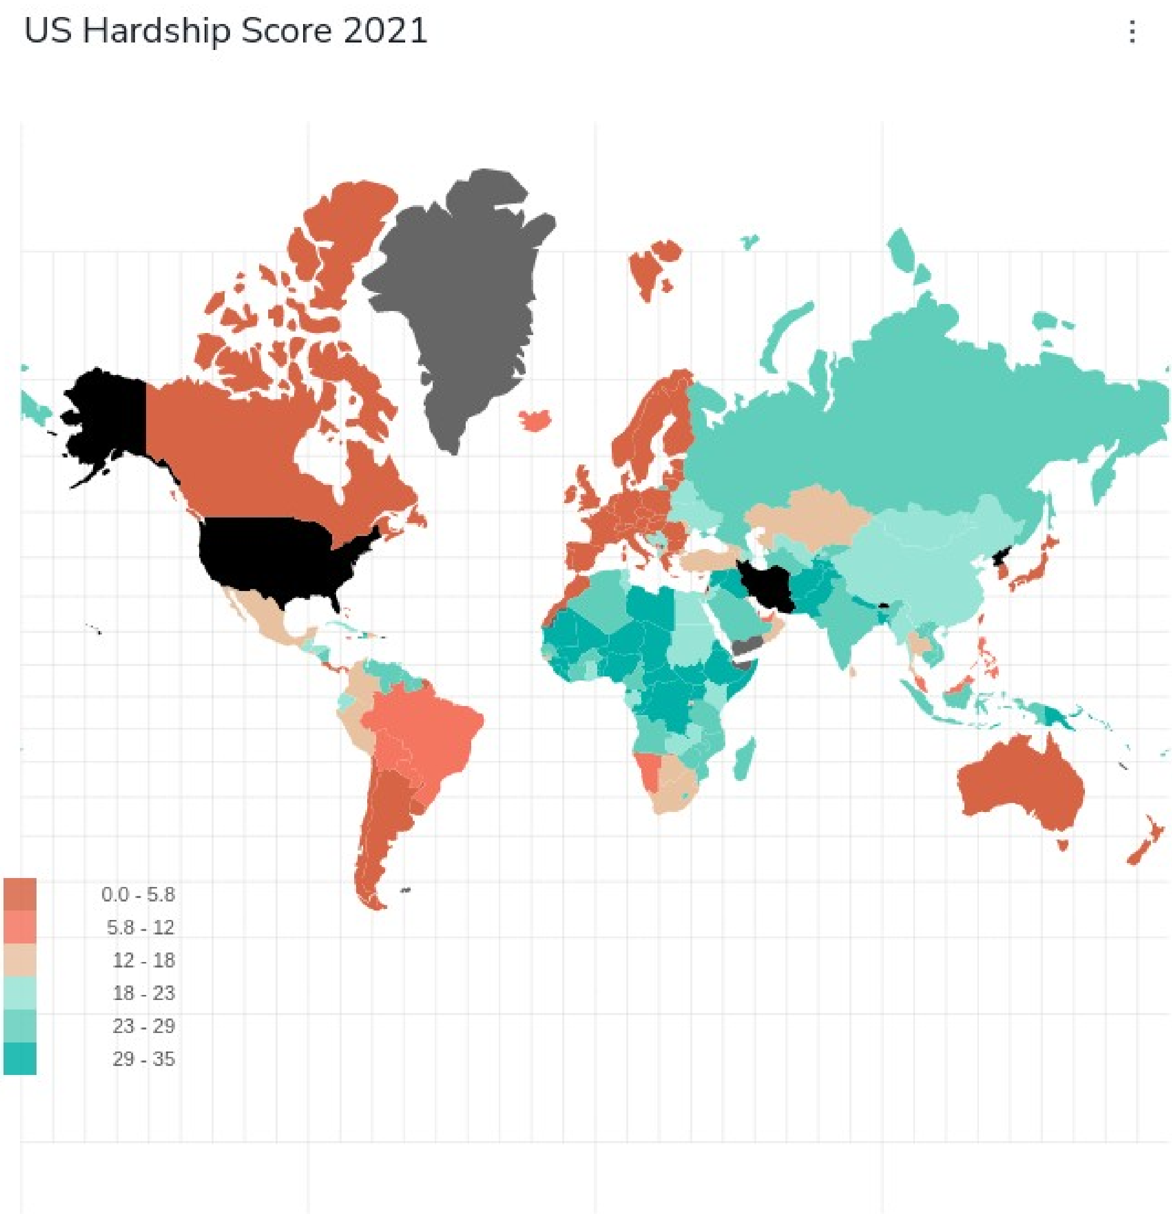

Figure 17: Average Hardship scores in the world based on US State Department's data. The black countries do not appear in the US hardship data.

To analyse the relationship between the in-degree of female ambassadors, I created a scatterplot (figure 18), which clearly shows the relationship between the living conditions in a country, the posting of female ambassadors, and the three key regions.

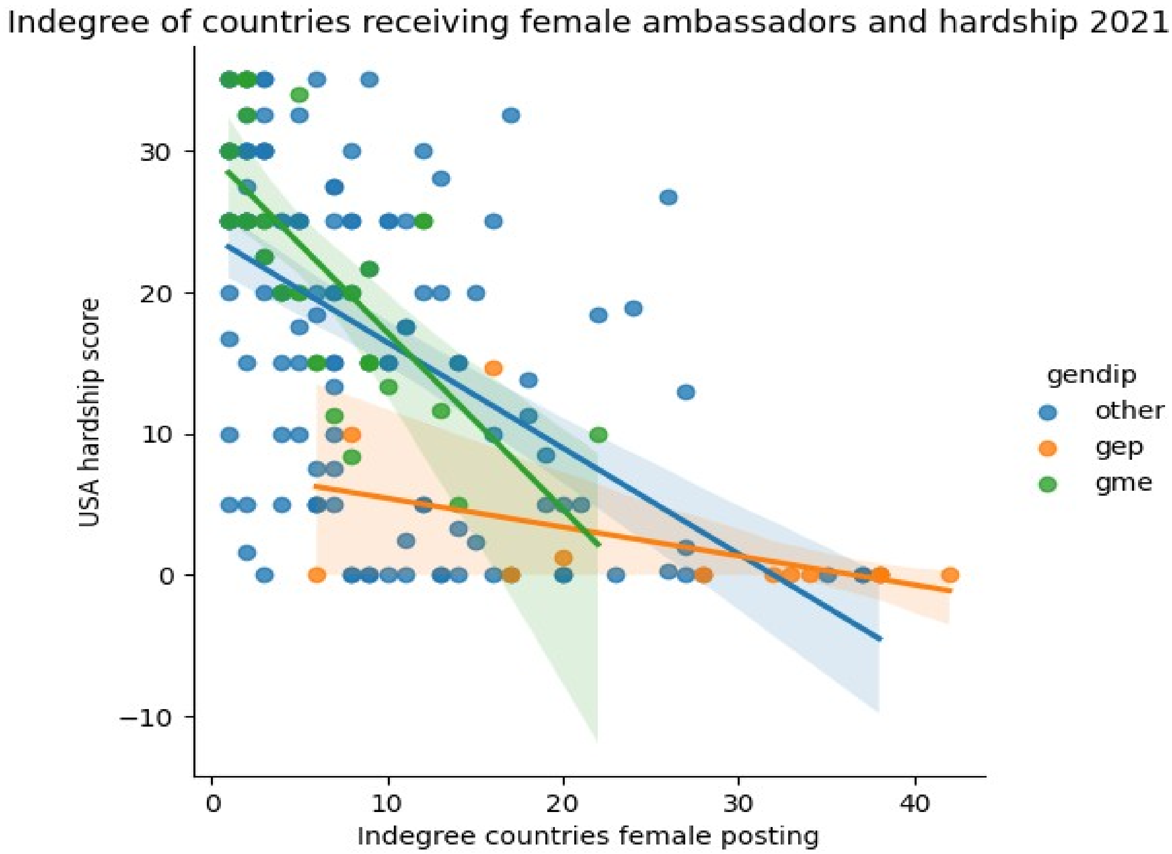

Figure 18: Scatterplot of hardship and in-degree of female ambassadors in 2021. The correlation for GME countries seems to be very strong, while the relation for GEP is weak.

I calculated Pearson's r (-0.61) which indicates a clear negative relationship between the hardship score and the in-degree. The p-value ($3.0 * 10^{-19}$) indicates that the null-hypothesis is rejected. Therefore, I can conclude that the more difficult the living conditions are in a country, the fewer women are appointed as ambassadors (figure 19).

```
# delete NaNs.

hardship_df.dropna(inplace=True)

# and retry
x = hardship_df["inDegree"]
y = hardship_df["hardshipScore"]


#calculate Spearman Rank correlation and corresponding p-value
r, p = pearsonr(x, y)

#print Spearman rank correlation and p-value
print(r)
print(p)
```

```
-0.607585622682603
3.054993790117951e-19
```

Figure 19: Hypothesis testing Hardship and in-degree

Q6 Is gender inequality a factor for female ambassadors' appointments?

Combining the UNDP gender inequality data with the GenDip database enables analysis of the relationship with in-degree for female ambassadors. As with the previous question, I queried the database and added the code to NeoDash to produce a map showing the distribution per country.

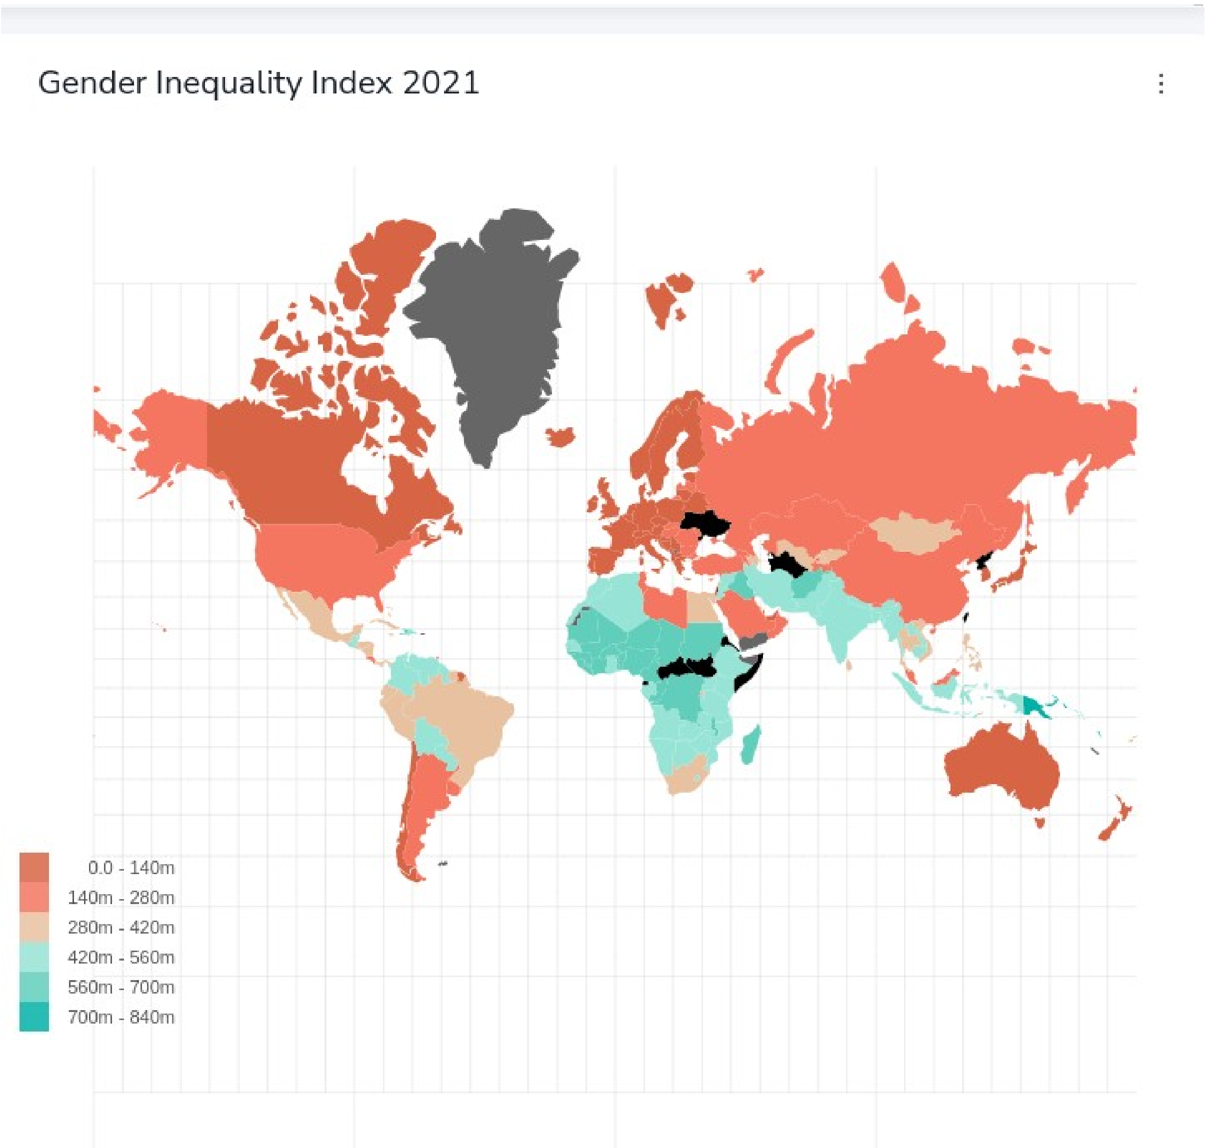

Figure 20: Gender Inequality Index 2021 (UNDP data). NB. NeoDash multiplied the GII scores x1000

Then, I created a scatterplot (figure 21) which clearly shows a negative relationship between the GII score and the number of female ambassadors in host countries.

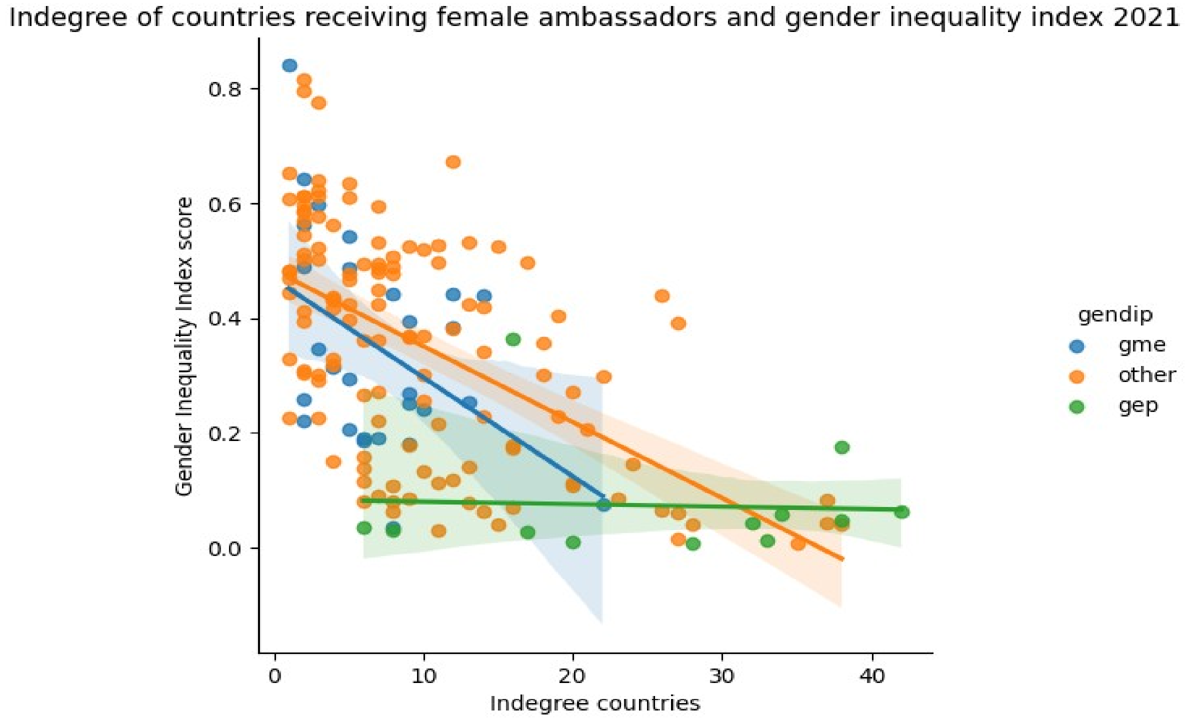

Figure 21: Scatterplot of the gender inequality index and in-degree for female ambassadors 2021. Clear negative correlation for GME countries and a weak relation for the GEP countries.

Using Python, I calculated Pearson's r = -0.58, which indicates a significant negative correlation between the Gender Inequality Index and the in-degree. The p-value = $6 * 10^{-16}$ rejects the null hypothesis. This suggests that female ambassadors are posted more often to countries with low Gender Inequality Index scores. Therefore, gender inequality also offers an additional explanation for decisions regarding female ambassadorial postings.

Reflecting on this, I believe that this index just as hardship provides a more effective classification metric than a selection countries based on the presence of certain policies, without taking in account local conditions. It seems that the GEP countries are more heterogeneous than the top countries in the Gender Inequality Index ranking, which makes such a grouping of countries less useful as a variable.

Q7 Is there a relationship between the share of female ambassadors sent and prevailing religion in the countries of origin?

I analysed the prevailing religion in relation to the in-degree of female ambassadors. The prevailing religion dataset contains categorical data, which is also extremely skewed, whereas the in-degree data is continuous. Due to their limited number, I ignored religions with fewer than seven occurrences (figure 22).

```
# Count values

sorted['prevailingReligion'].value_counts()
✓ 0.0s

prevailingReligion
Christians                 106
Muslims                     41
Religiously_unaffiliated     9
Buddhists                    7
Hindus                       3
Jews                         1
Other_religions              1
Name: count, dtype: int64
```

Figure 22: Distribution of prevailing religions in the world 2020. PEW research

For this scenario, I evaluated different visualisation techniques and opted for a box plot, as this is best suited to visualising the relationship between categorical and continuous data (Figure 23). This clearly distinguishes between Christian and religiously unaffiliated countries, which report the highest in-degree scores, and Muslim and Buddhist countries, which report the lowest.

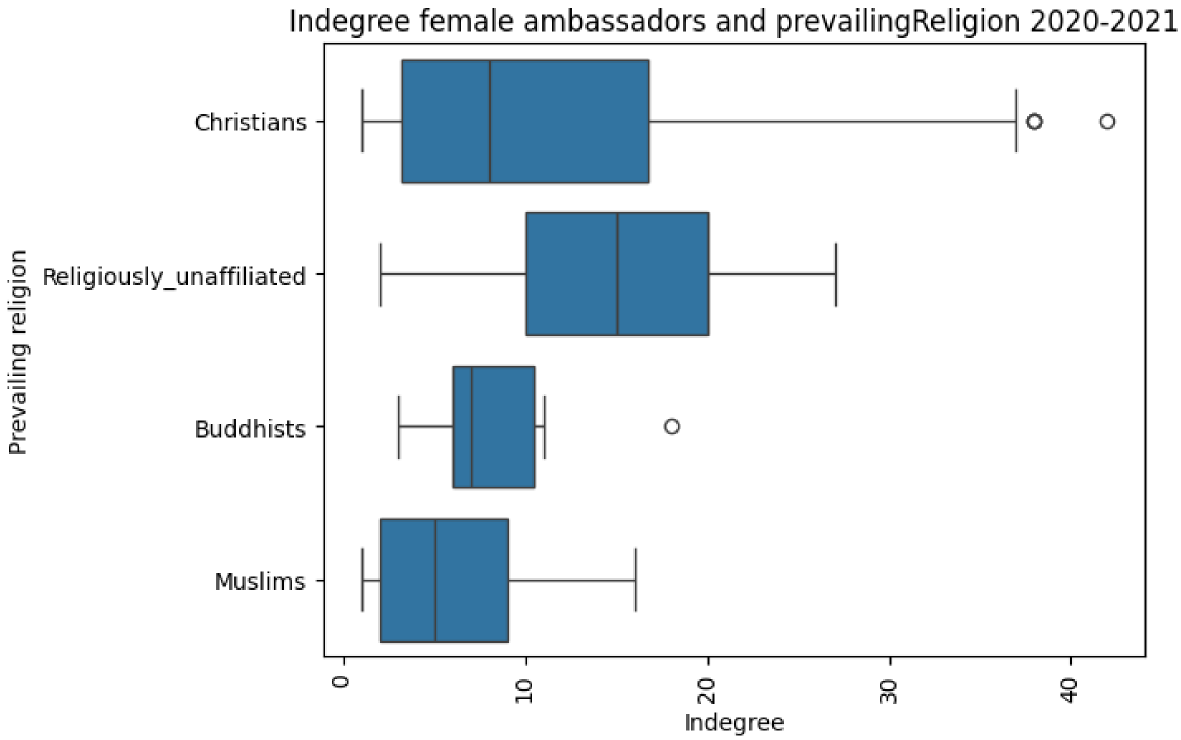

Figure 23: Boxplot of religion (2020) and in-degree of female ambassadors sent to GEP countries in 2021

Although the absolute numbers are low, on average, 27.5% of female ambassadors from the seven Buddhist countries are posted to GEP countries. This is a higher percentage than those from Muslims states (23.8%) or the GME region (24.6%).

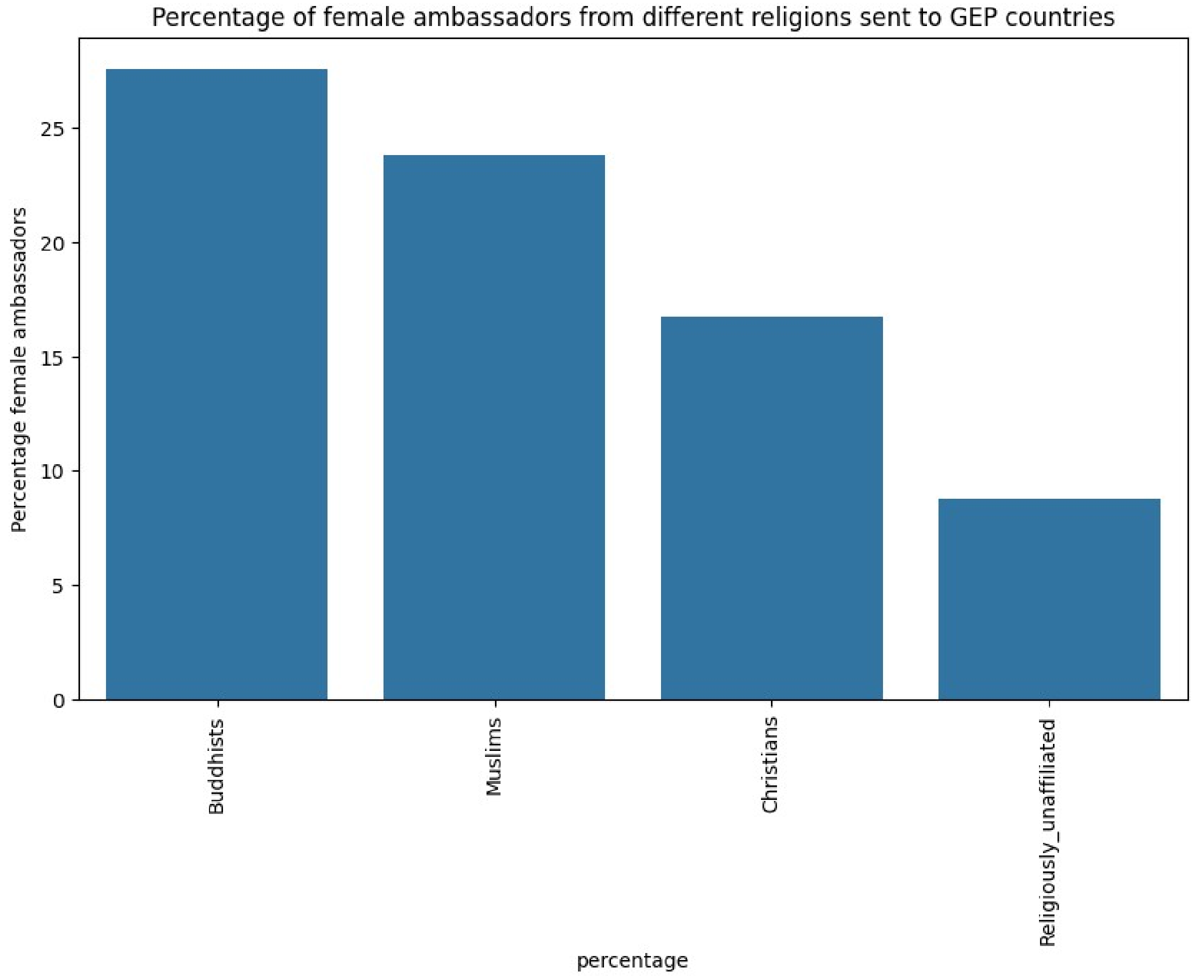

Figure 24: Percentage of female ambassadors sent to GEP countries per religion. Religion is 2020 data, diplomats 2021.

I calculated a Kruskal-Wallis H-score of 13.37, which indicates a moderate relationship between religion and in-degree. The p-value is 0.0057. We can therefore reject the null-hypothesis. There is a relationship between religion and the number of female ambassadors sent to GEP countries. Although the skewed dataset has its limitations, religion remains a variable of particular interest when identifying countries for further analysis.

My analysis does not go into the complexity of religion as an explanatory factor, and I limited it to find the largest religious groups. The skewed data also made it difficult to use some of the groups in a meaningful way. With these limitations in mind, it was interesting to observe the relatively high proportion of female ambassadors appointed by Buddhist countries.

Conclusions Part B

The following table summarises the relationships between the variables that we have analysed so far. The table illustrates some of the key findings. For the calculations I used the entire dataset, although the specific data for hardship and GII is based on 2021 research only. The percentages refer to all female ambassadors sent to GEP countries, not just those from GME countries. For the other two comparisons, I considered the top and bottom 20% of countries.

Table 3: Percentages of female ambassadors related to various variables

		T -test (< 0.05)
%women GEP n = 12	%women other countries n = 163	
16.0	12.6	0.19
%women low GII n = 34	%women high GII n = 34	
15.1	8.44	$1.44 * 10^{-7}$
%women low hardship n = 36	%women high hardship n = 38	
15.6	7.8	$2.9 * 10^{-8}$

There is a higher percentage of female ambassadors in GEP countries, but the null hypothesis cannot rejected, therefore it is impossible to base conclusions on this finding.

More robust statistical evidence was found for alternative explanations of female ambassadorial postings:

1. Countries that have female politicians active in key foreign policy areas, attract relatively more female ambassadors.
2. Living conditions and the inequality scores are significant indicators: women are more often sent to countries with better conditions.
3. Despite its limits, the religion dataset demonstrated that Buddhist and Muslim countries send a higher share of their female ambassadors to GEP countries, suggesting religion is a factor worth further exploration.

Part C

Analysing the diplomatic network

This part of the research starts with an analysis of the diplomatic network. I followed the procedure described by Needham and Hodler (2019), I first analysed the characteristics of the diplomatic network.

The female diplomatic network has a diameter of 5 (figure 25). This means that the *longest shortest path* takes five steps from one country to another. The average diameter is 2.1.

```
neo4j$ CALL gds.alpha.allShortestPaths.stream('diploFem') YIELD sourceNodeId, targetNodeId, distance RETURN max(distance) AS maxDistance, min(distance)
```

	maxDistance	minDistance	avgDistance	pathCount
1	5.0	1.0	2.107233486595816	39801

Figure 25: Query of the diameter in Neo4j

The second metric we can consider is the *local clustering coefficient* which is 0.77 (figure 26). This means that, on average, 77% of countries to which a country sends a female ambassador also send female ambassadors to each other.

```
CALL
gds.localClusteringCoefficient.stream('diploF')
YIELD localClusteringCoefficient
RETURN avg(localClusteringCoefficient)
```

	avg(localClusteringCoefficient)
1	0.7661147923910273

Figure 26: Query of clustering coefficient in Neo4j

Based on these two findings, I can conclude that the diplomatic network has the characteristics of a small-world network. Knowing this, I drew on the research approaches to small-world networks outlined in Sung (2025) and Smith (2020), using them as examples of a possible methodology.

Q8. Identify the prestigious countries and relations: centrality

Is it possible to use graph tools to identify which countries can be considered influential? I considered the following:

1. Out-degree: dimension of a country's diplomatic network
2. In-degree: importance of a hosting country
3. PageRank: influence of a country based on its connected neighbours
4. Closeness: relevance of country as an intermediary in international relations.

The out-degree is a centrality measure that counts the number of outgoing relations. Considering the costs and organisation involved in maintaining a diplomatic network (Duque, 2018), the out-degree can be considered as an indication of a country's diplomatic power.

When applying this function on our entire dataset, we find that the top countries are the economic and military superpowers in world politics (figure 27).

Figure 27: Top 20 countries based on out-degree

Next, I created an overview of all centrality algorithms by calculating the scores for female ambassadors and assigning them to country nodes. Using Pandas, I created a dataframe with the four measures and sorted them to show the top 20 countries for each (figure 28).

```
ranked = {}
for col in ["out-degree", "in-degree", "PageRank", "closeness"]:
    ranked[col] = centrality.sort_values(col, ascending=False)["country"].reset_index(drop=True)

ranking_df = pd.DataFrame(ranked)
ranking_df.index = ranking_df.index + 1
ranking_df.head(20)
```

✓ 0.0s

	out-degree	in-degree	PageRank	closeness
1	United States of America	United States of America	Belgium	Belgium
2	United Kingdom	United Kingdom	United States of America	France
3	Sweden	France	France	United States of America
4	Canada	Sweden	Sweden	Germany
5	Cuba	Canada	United Kingdom	Italy
6	France	Cuba	Germany	Sweden
7	Finland	Netherlands	Switzerland	Canada
8	South Africa	Germany	Italy	United Kingdom
9	Australia	Australia	Canada	Switzerland
10	Netherlands	Austria	Netherlands	Netherlands
11	Norway	Finland	Austria	Austria
12	Austria	India	Spain	Spain
13	Germany	South Africa	China	South Africa
14	Venezuela, Bolivarian Republic of	Norway	Cuba	Cuba
15	Philippines	Venezuela, Bolivarian Republic of	Denmark	Brazil
16	Colombia	Türkiye	Brazil	Portugal
17	Romania	Switzerland	Australia	Denmark
18	Türkiye	Romania	South Africa	Australia
19	Switzerland	Philippines	Portugal	India
20	India	Czechia	India	China

Figure 28: Overview of centrality scores countries based on presence female ambassadors.

Finally, I analysed the percentage of female ambassadors and the prestige of host countries based on their overall PageRank (figure 29).

```
[62]  print(np.mean(highPrAll["percent"]), np.mean(lowPrAll["percent"]))
      ✓ 0.0s
...   12.174080837052804 18.453561853523155

[63]  stats.ttest_ind(a=(highPrAll["percent"]), b=(lowPrAll["percent"]), equal_var=True)
      ✓ 0.0s
...   TtestResult(statistic=np.float64(-2.142525137583), pvalue=np.float64(0.03861930049288016), df=np.float64(38.0))
```

Figure 29: Calculation relation in-degree and PageRank

If we compare the top 20 countries on PageRank and the bottom 20 and compare the average presence of incoming ambassadors, the results are as follows (table 4):

Table 4: Portion of female ambassadors based on PageRank

% in high PageRank countries n = 20	% in low PageRank countries n = 20	T -test (< 0.05)
12.2	18.5	0.04

Since the null hypothesis is not supported, I conclude that influential countries receive fewer female ambassadors. This supports Duque's thesis (2018) on status in networks and aligns with Niklasson and Town's 2017 findings on women's under-representation in prestigious postings.

Also interesting is that high-ranking countries on PageRank and closeness include regional hubs (Cuba, Brazil, Australia) and hosts of multilateral organisations (Belgium, Italy, Switzerland).

The centrality analysis was able to address the core research questions. These graph metrics effectively link theoretical assumptions, insights from additional datasets, and the network structure, providing evidence why women are sent to certain countries, especially less prestigious nations.

Q9: Identify meaningful communities

The Graph Data Science library contains various algorithms for analysing the presence of communities in the dataset. Again, I will follow the approach described by Needham and Hodler's (2019) to better characterise the network.

Triangles

The first step in identifying the presence of communities is to calculate the number of triangles. Neo4j's GDS library makes it easy to perform the calculations. There are 78,003 triangles (figure 30).

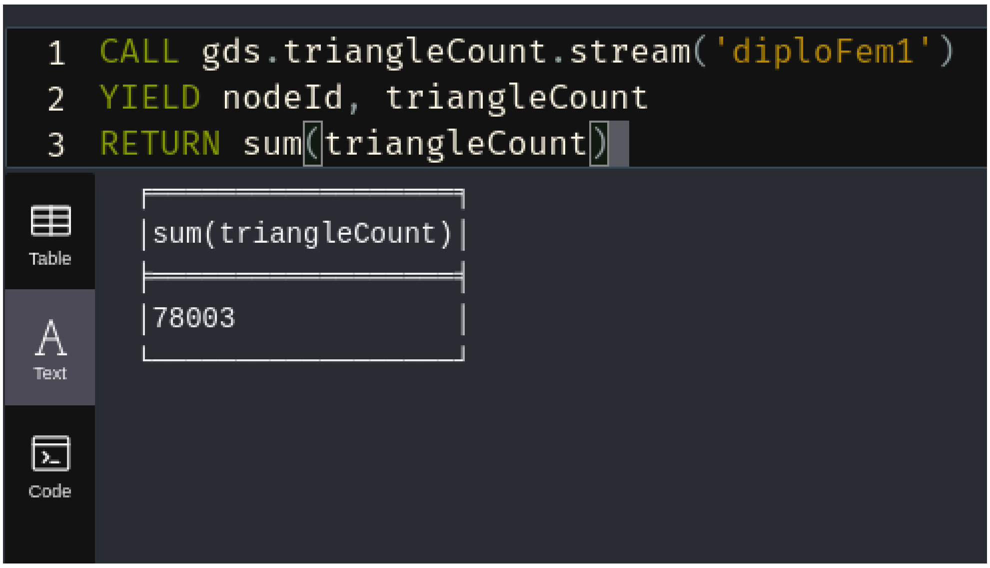

Figure 30: Calculation of triangles

This absolute number should be understood in relation to the size of the network. I have already calculated the average clustering coefficient, which indicates what proportion of the nodes are connected to each other. Figure 26 shows a value of 0.77, meaning that, on average, countries are connected to 77% of the network. This is a high score.

Weakly connected components
The next step is to analyse whether there are any connected components. We already know that we are dealing with a small-world network, so it is unlikely there will be lots of connected components.

The algorithm returns just one cluster (figure 31). This means that we are dealing with a network that has a so-called Giant Connected Component that contains a significant proportion of the entire graph's nodes (and therefore the countries).

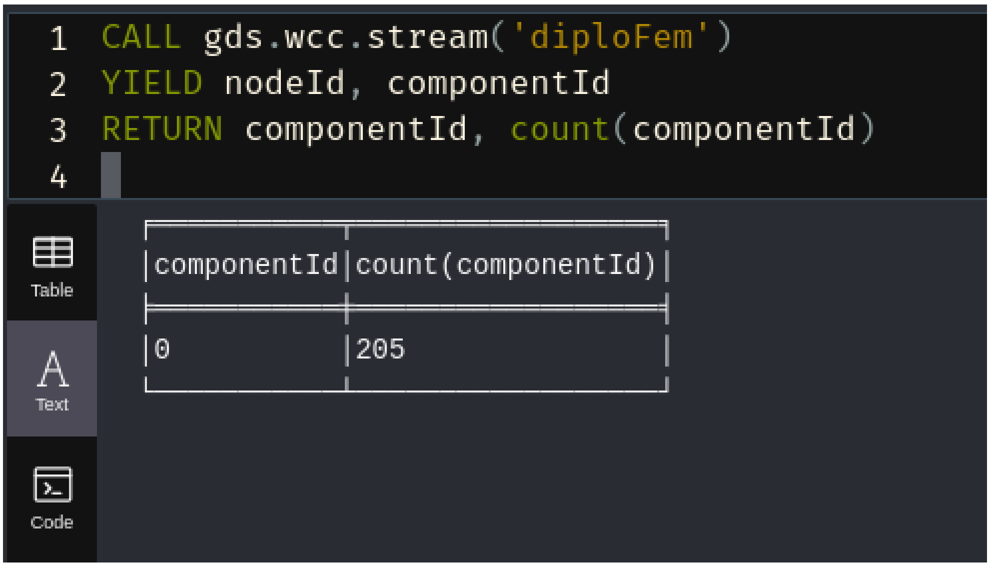

Figure 31: Calculation of connected components

Louvain

Finally, I applied the Louvain algorithm to identify clusters. It identified 4 communities that can be visualised on a map (figure 32). While the distribution may not be immediately obvious, analysing how the Louvain algorithm groups the properties used so far allowed me to describe the four groups and gain a better understanding of the types of countries present in the dataset.

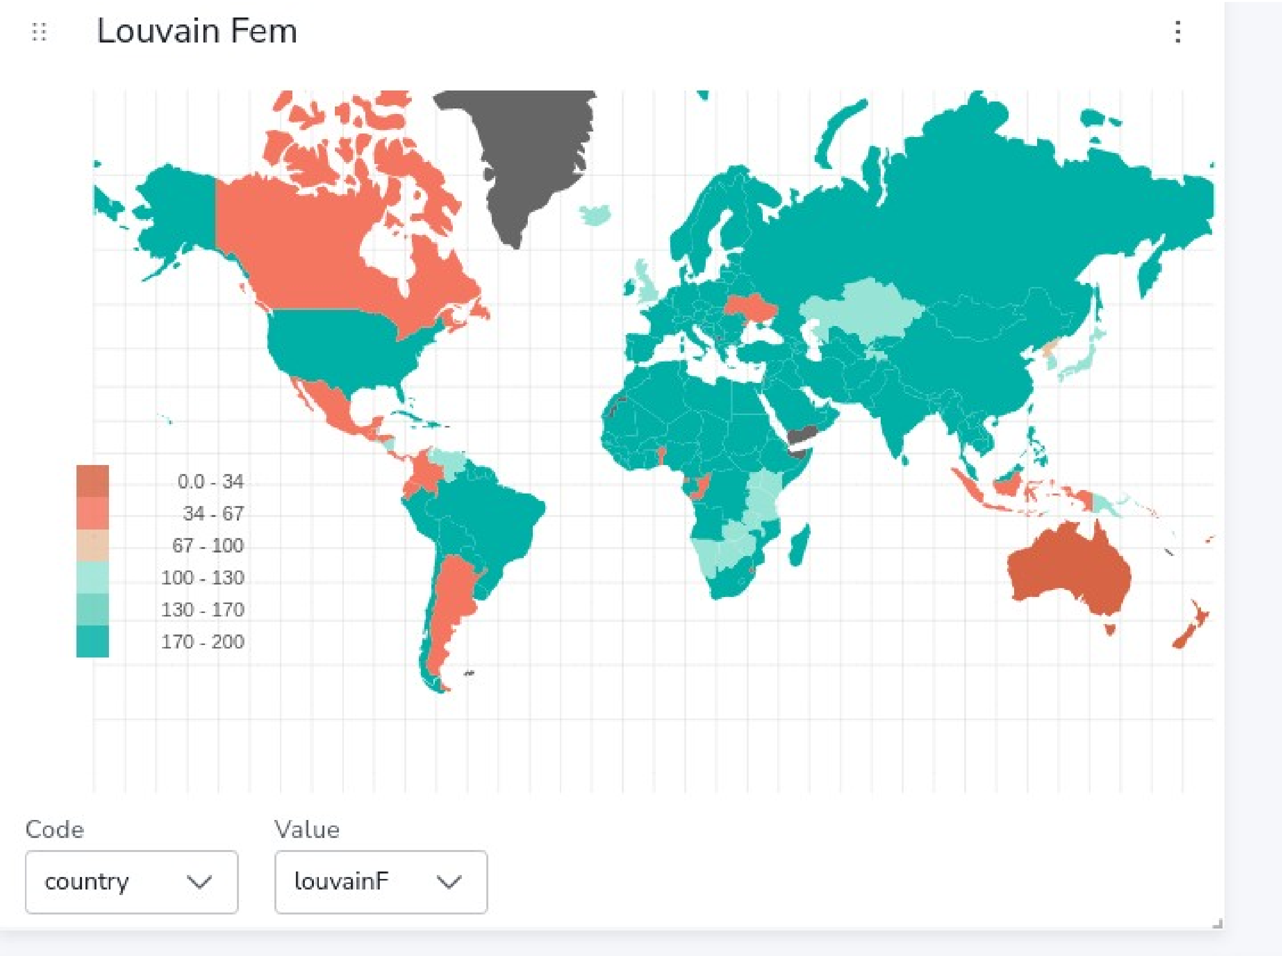

Figure 32: Countries communities detected with Louvain algorithm

The following page includes the output of three queries. The first contains the four communities and their scores on the various variables used in this research (figure 33). The second is a calculation of the modularity (Figure 34) and the final contains the names of the countries grouped by community (figure 35).

community	community_size	s_outDegree	r_outDegree	s_avgPRAll	r_avgPRAll	avgDistance	s_avgGII	r_avgGII	s_avgHardship	r_avgHardship	s_avgGDP	r_avgGDP
169	75	58.9675107056412	62.93317299834223	1.450260772670153	1.4794295251658744	5305.042107761703	0.192574728959156	0.26726078392260777	10.087176591782537	13.089538622073881	38895.215795438955	32407.877367539124
94	65	80.3320196608067	65.21206337934439	1.6387029013649879	1.5448682419042383	6155.300310729316	0.39260906788247296	0.29451275377407593	17.255219024160283	14.089714840058488	22121.382945392845	30144.63170787454
138	22	84.55492755112549	74.13125386040785	1.1869057653756614	1.589605827595504	7812.869754429599	0.3397516185386903	0.2559295158999026	13.04659629754687	10.930057316420987	22931.983322497883	33354.220690094255
201	41	59.92574678430942	60.67165849499051	1.774292227896591	1.4383389898998538	7655.816066630837	0.2515378817269854	0.2868557965254902	13.647833345720487	14.173689753971294	33634.64699658179	30758.356923700347

Figure 33: Average scores of the communities of some of the variables analysed in this study

Louvain is based on the calculation of modularity.

```
CALL gds.louvain.write('diploAll',
    {maxIterations:20,
     includeIntermediateCommunities:true,
     writeProperty:'louvain'})
YIELD ranLevels, communityCount,modularity,modularities
```

Table
Text

	ranLevels	communityCount	modularity	modularities
1	2	4	0.13940780043774997	[0.1268667710661192, 0.13940780043774997]

Figure 34: Calculation of modularity of the community algorithm

The communities are composed as follows:

```
MATCH (c:Country)
WHERE c.louvain[-1] IS NOT NULL
WITH c, c.louvain[-1] as community
ORDER BY c.prAll DESC
RETURN community, collect(c.name)
```

community	collect(c.name)
94	["United States of America", "United Kingdom", "France", "Brazil", "Cuba", "Türkiye", "Ethiopia", "Nigeria", "Kenya", "Holy See", "Senegal", "Congo, Democratic Republic of the", "Tanzania, United Republic of", "Ghana", "Côte d'Ivoire", "Zimbabwe", "Zambia", "Angola", "Mozambique", "Uganda", "Cameroon", "Gabon", "Guinea", "Yugoslavia", "Bahrain", "Mali", "Namibia", "Congo", "Taiwan", "Liberia", "Botswana", "Burkina Faso", "Haiti", "German Democratic Republic", "Malawi", "Benin", "Central African Republic", "Niger", "Chad", "Tajikistan", "Burundi", "Djibouti", "Iceland", "Togo", "Sierra Leone", "Eritrea", "Equatorial Guinea", "Guinea-Bissau", "South Sudan", "Gambia", "Cabo Verde", "Seychelles", "Eswatini", "Bahamas", "Saint Lucia", "Comoros", "Micronesia, Federated States of", "Grenada", "Saint Vincent and the Grenadines", "Antigua and Barbuda", "Palau", "Dominica", "Saint Kitts and Nevis", "Andorra", "Tuvalu"]
169	["Belgium", "Germany", "Italy", "Austria", "Spain", "Sweden", "Saudi Arabia", "Netherlands", "Malaysia", "Poland", "Algeria", "Indonesia", "Greece", "Morocco", "Portugal", "Israel", "United Arab Emirates", "Romania", "Kuwait", "Pakistan", "Denmark", "Thailand", "Hungary", "Bulgaria", "Lebanon", "Chile", "Czechia", "Libya", "Viet Nam", "Tunisia", "Jordan", "Finland", "Qatar", "Peru", "Iraq", "Ukraine", "Ireland", "Sudan", "Syrian Arab Republic", "Uruguay", "Kazakhstan", "Oman", "USSR", "Serbia", "Cyprus", "Croatia", "Myanmar", "Yemen", "Azerbaijan", "Afghanistan", "Slovakia", "Albania", "Bosnia and Herzegovina", "Belarus", "Uzbekistan", "Slovenia", "Latvia", "Luxembourg", "Lithuania", "Estonia", "Georgia", "Mauritania", "Turkmenistan", "North Macedonia", "Malta", "Somalia", "Moldova, Republic of", "Montenegro", "Kosovo", "Serbia and Montenegro", "Maldives", "Sao Tome and Principe", "Palestine, State of", "San Marino", "Liechtenstein", "Yemen"]
201	["China", "Japan", "India", "Egypt", "Russian Federation", "Switzerland", "South Africa", "Australia", "South Korea", "Iran, Islamic Republic of", "Norway", "Philippines", "Singapore", "New Zealand", "Bangladesh", "Sri Lanka", "Fiji", "Lao People's Democratic Republic", "North Korea", "Rwanda", "Nepal", "Brunei Darussalam", "Cambodia", "Mongolia", "Madagascar", "Papua New Guinea", "Kyrgyzstan", "Mauritius", "Timor-Leste", "Solomon Islands", "Lesotho", "Samoa", "Vanuatu", "Tonga", "Kiribati", "Marshall Islands", "Bhutan", "Monaco", "Nauru", "Cook Islands", "Niue", "Viet Nam"]
138	["Canada", "Mexico", "Argentina", "Venezuela, Bolivarian Republic of", "Colombia", "Panama", "Ecuador", "Guatemala", "Costa Rica", "Dominican Republic", "Jamaica", "Nicaragua", "Bolivia, Plurinational State of", "El Salvador", "Trinidad and Tobago", "Paraguay", "Honduras", "Armenia", "Guyana", "Suriname", "Barbados", "Belize"]

Figure 35: Composition of the communities

The modularity score of 0.139 is considered low (Smith, 2024), which may cause inconsistent community assignments, with many nodes linked across communities.

Communities 169 and 201 show the smallest female diplomatic networks (low out-degree) and post ambassadors to countries with larger networks, though not the most influential (PageRank 1.36 and 1.31). Differences exist: Community 169 sends women to states with better Gender Inequality Index scores, less hardship, and shorter distances, covering much of Europe and the Greater Middle East. Community 201, mainly Asian and Oceanian states, shows less favourable patterns.

Communities 94 and 138 have larger networks and send ambassadors to smaller ones. Community 138 appoints women to the most influential countries, with high PageRank, good GII, and favourable living conditions, comprising Latin America, Armenia, and Canada. Community 94 posts women to less favourable states, with poorer GII and high hardship, including the USA, UK, France, and much of Africa.

Conclusions

The Louvain algorithm identifies four communities with some distinctive features. These seem to reflect geographical and historical colonial ties, offering a starting point for further analysis to discover other features.

For this research, however, the usefulness is limited. First, the small-world nature of the network may have made it difficult for the algorithms to separate clusters clearly. Second, these methods group sending and receiving countries together, whereas my focus lies on similarities within either sending or receiving groups. This makes the results difficult to interpret for my use case.

I had hoped to find more useful results. This may be due to inexperience, perhaps in the selection of properties when building the projected graph in Neo4j, or in the choice of community detection algorithm. I did re-check procedures and even applied another algorithm designed to identify overlapping communities, but this too produced a giant connected component. I tend to conclude that the results found reflect the nature of diplomatic networks, where states maintain relations widely out of interest rather than within small circles of partners.

The most meaningful outcomes so far come from centrality measures, and I will therefore focus the final analysis on centrality algorithms.

Q10. Are there countries that use their female ambassadorial appointments strategically?

The GenDip paper indicated Saudi Arabia appointing female ambassadors strategically, sending them to the USA, Norway, Sweden, and Myanmar (figure 36). Three are GEP countries, and Myanmar then had prominent opposition leader Aung San Suu Kyi.

Based on the diplomatic network findings, I examined anomalies to identify strategic female ambassador nominations. Combining previous results may reveal countries using such appointments as signals to hosts, as hypothesised by Niklasson and Towns (2023).

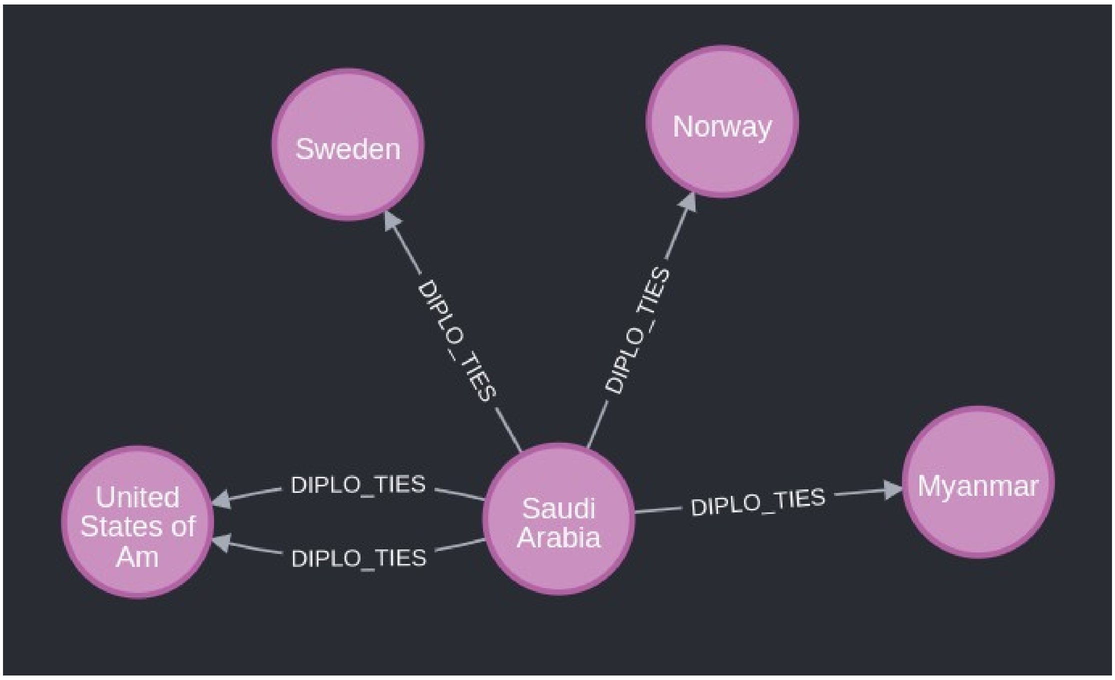

Figure 36: Saudi Arabia's female ambassador postings

GME countries and female PageRank

I analysed whether similar situations to Saudi Arabia exist in the dataset, starting with an analysis on postings for female ambassadors based on the network structure, in specific the countries with high female PageRank scores. My approach involved:

1. Identifying GME countries with a limited number of female ambassadors (out-degree < 11).
2. Determining to which GEP countries (with high PageRank) these female ambassadors were posted to.
3. Calculating the percentage of female ambassadors sent to each target country.
4. Considering the cumulative number of female ambassadors over the years (weighted).

This analysis highlighted Saudi Arabia and Qatar, both sending over 50% of their female ambassadors to these countries. In comparison, other countries had much lower percentages (figure 37).

```
MATCH (s: Country)-[d:FEMHOMS]→(r:Country)
// WHERE s.degreeFem ⩾ 3 AND s.degreeFem <11 AND r.prFem > 1.95
WHERE s.gme = 1 AND r.gep = 1
RETURN s.name, COUNT(d) AS femHomsSent, s.degreeFem AS totalFemHoms, round(100.0 * COUNT(d)/s.degreeFem, 2) AS perc, sum(d.femHoms) AS weight, collect(r.name) AS hostCountries ORDER BY
perc DESC, weight DESC
```

	s.name	femHomsSent	totalFemHoms	perc	weight	hostCountries
1	"Saudi Arabia"	3	5.0	60.0	4	["Norway", "Sweden", "United States of America"]
2	"Qatar"	1	2.0	50.0	2	["Sweden"]
3	"Afghanistan"	3	9.0	33.33	5	["Germany", "Norway", "United States of America"]
4	"Iran, Islamic Republic of"	1	3.0	33.33	2	["Denmark"]
5	"Oman"	3	10.0	30.0	7	["Germany", "Spain", "United States of America"]
6	"Armenia"	6	22.0	27.27	9	["Canada", "France", "Germany", "Mexico", "Spain", "United States of America"]
7	"United Arab Emirates"	6	23.0	26.09	11	["Denmark", "Finland", "France", "Germany", "Spain", "Sweden"]
8	"Palestine, State of"	3	13.0	23.08	6	["Germany", "Norway", "Sweden"]
9	"Jordan"	4	18.0	22.22	8	["Canada", "France", "Spain", "United States of America"]
10	"Azerbaijan"	2	9.0	22.22	4	["France", "United States of America"]
11	"Libya"	2	9.0	22.22	4	["Sweden", "United States of America"]
12	"Bahrain"	2	9.0	22.22	2	["France", "United States of America"]

Figure 37: Anomaly detection, GME and GEP countries

Low female out-degree and high female PageRank

Previous analysis showed that women are generally not sent to ‘prestigious’ countries (high female PageRank). The GenDip research hypothesises that appointments used as a political signal are more likely in prestigious countries with favourable female foreign policies. To identify these, I focused on countries with few female ambassadors (low out-degree) and high female PageRank countries.

This query revealed several interesting cases, including Mauritania, Dominica, the Bahamas, Bhutan, and Mongolia (figure 38). Two of these countries consistently post their female ambassadors to this group of host states, as indicated by the ‘weight’ column. And interestingly two are predominantly Buddhist and two are Muslim countries.

This is, of course, a high-level quantitative analysis, and more detailed investigation is needed to understand the dynamics within each country. Nevertheless, I would argue that the value of this research lies in providing useful insights and input for further study.

```
MATCH (s:Country)-[d:FEMHOMS]->(r:Country)
WHERE s.degreeFem >=3 AND s.degreeFem < 11 AND r.prFem >= 1.95
RETURN s.name, COUNT(d) AS femHomSent, s.degreeFem AS totalFemHoms, round(100.0 * COUNT(d)/s.degreeFem, 2) AS perc, SUM(d.femHoms) AS weight, collect(r.name) AS
hostCountries ORDER BY perc DESC, weight DESC
```

	s.name	femHomSent	totalFemHoms	perc	weight	hostCountries
1	"Mauritania"	3	3.0	100.0	3	["Belgium", "France", "Italy"]
2	"Dominica"	5	7.0	71.43	7	["Belgium", "Brazil", "Italy", "United Kingdom", "United States of America"]
3	"Bahamas"	6	10.0	60.0	9	["Belgium", "Canada", "China", "Cuba", "United Kingdom", "United States of America"]
4	"Bhutan"	3	5.0	60.0	4	["Belgium", "Canada", "India"]
5	"Mongolia"	4	7.0	57.14	4	["Austria", "Brazil", "Cuba", "France"]
6	"Bahrain"	5	9.0	55.56	9	["Belgium", "China", "France", "United Kingdom", "United States of America"]
7	"Guinea-Bissau"	3	6.0	50.0	4	["Brazil", "Cuba", "France"]
8	"Kazakhstan"	4	8.0	50.0	4	["Belgium", "Italy", "Netherlands", "Switzerland"]
9	"Djibouti"	2	4.0	50.0	3	["France", "Switzerland"]
10	"Kuwait"	4	9.0	44.44	4	["Austria", "Belgium", "Canada", "South Africa"]
11	"Serbia and Montenegro"	4	9.0	44.44	4	["Belgium", "Cuba", "Netherlands", "Portugal"]
12	"Antigua and Barbuda"	3	7.0	42.86	7	["Brazil", "United Kingdom", "United States of America"]

Figure 38: Anomaly detection, out-degree and PageRank

Low hardship and high hardship countries

This research also showed that more women are sent to countries with better living conditions. An interesting anomaly is to examine the percentages of female ambassadors from low-hardship countries who are posted to high-hardship countries (figure 39). This query highlights Denmark and the USA, both sending around 10% of their female ambassadors to countries with the poorest living conditions, with the US showing particularly consistent postings (weight of 72). It would be interesting to examine the motivation behind the specific decisions made by the US State Department and Denmark's MFA.

```
MATCH (s: Country)-[d:FEMHOMS]→(r:Country)
//   WHERE s.prFem > 1.95 AND r.iig > 0.596
WHERE s.hardship = 0 AND (r.hardship ⩾ 30 OR r.code = "IRN")
  RETURN s.name, COUNT(d) AS femHomsSent, s.degreeFem AS totalFemHoms, round(100.0 * COUNT(d)/s.degreeFem, 2) AS perc, sum(d.femHoms) AS weight, collect(r.name) AS hostCountries ORDER BY
perc DESC, weight DESC
```

	s.name	femHomsSent	totalFemHoms	perc	weight	hostCountries
1	"Denmark"	9	85.0	10.59	15	["Afghanistan", "Bangladesh", "Ethiopia", "Guinea", "Mali", "Nepal", "Pakistan", "Syrian Arab Republic", "Uganda"]
2	"United States of America"	29	296.0	9.8	72	["Benin", "Burundi", "Cabo Verde", "Central African Republic", "Congo", "Equatorial Guinea", "Eritrea", "Ethiopia", "Guinea", "Haiti", "Iraq", "Liberia", "Libya", "Mali", "Mauritania", "Micronesia, Federated States of", "N
3	"Germany"	11	126.0	8.73	19	["Bangladesh", "Congo, Democratic Republic of the", "Ethiopia", "Iraq", "Mauritania", "Nepal", "Niger", "Nigeria", "Sierra Leone", "Tajikistan", "Yemen"]
4	"Japan"	2	23.0	8.7	2	["Ethiopia", "Vanuatu"]
5	"France"	14	176.0	7.95	23	["Bangladesh", "Benin", "Cabo Verde", "Chad", "Eritrea", "Ethiopia", "Guinea", "Libya", "Nepal", "Nigeria", "Tajikistan", "Togo", "Uganda", "Vanuatu"]
6	"United Kingdom"	17	220.0	7.73	27	["Afghanistan", "Congo, Democratic Republic of the", "Eritrea", "Guinea", "Haiti", "Libya", "Mali", "Nepal", "Niger", "Nigeria", "Sierra Leone", "Somalia", "Timor-Leste", "Tonga", "Uganda", "Vanuatu", "Yemen"]
7	"Netherlands"	9	144.0	6.25	12	["Afghanistan", "Bangladesh", "Benin", "Burundi", "Ethiopia", "Iraq", "Mali", "Syrian Arab Republic", "Uganda"]
8	"Spain"	5	83.0	6.02	6	["Cabo Verde", "Ethiopia", "Mauritania", "Niger", "Pakistan"]
9	"Greece"	4	75.0	5.33	4	["Congo, Democratic Republic of the", "Ethiopia", "Libya", "Syrian Arab Republic"]
10	"Canada"	11	214.0	5.14	18	["Afghanistan", "Bangladesh", "Congo, Democratic Republic of the", "Iraq", "Libya", "Mali", "Niger", "Nigeria", "Pakistan", "South Sudan", "Syrian Arab Republic"]
11	"Switzerland"	5	102.0	4.9	11	["Bangladesh", "Ethiopia", "Haiti", "Nepal", "Papua New Guinea"]
12	"New Zealand"	4	85.0	4.71	11	["Papua New Guinea", "Solomon Islands", "Tonga", "Vanuatu"]

Figure 39: Anomaly detection, hardship

Conclusions Part C

This project demonstrated both the usefulness and the limitations of graph data analysis for describing the diplomatic network. A major strength was the use of centrality algorithms, which effectively identified influential countries. The GDS library successfully captured the network's characteristics, showing it to be very dense, with most countries highly interconnected. However, these strong connections made it difficult for community-detection algorithms to identify meaningful sub-graphs. Although I identified four communities, their usefulness for this research was limited.

The project began with the question of whether countries use female ambassadorial appointments as a political signal. By combining relevant variables with centrality measures, I was able to identify countries, other than Saudi Arabia, that deviated from the majority of the network. These deviations could suggest deliberate strategies, though the available data cannot confirm individual countries' intentions. Nonetheless, this analysis provides valuable input for further research.

This part required extensive study of both theory and algorithms, along with many experiments. There is no off-the-shelf method for detecting anomalies; this requires combining domain knowledge with suitable algorithms. The research confirmed the GenDip team's initial findings on anomalies and revealed additional scenarios. While other variable combinations are possible, I focused on those that Part B indicated had the strongest relationship with female ambassadorial appointments. The key takeaway is that graph data science is a powerful tool for international relations research, and the effort invested in mastering it was worthwhile.

Overall, I am satisfied with the outcomes of this part. It was a significant learning curve, improving both my technical understanding and knowledge of the research domain. While the community analysis provided limited insight, I hope this reflects the highly connected nature of the diplomatic network rather than any limitations on my part.

Final conclusions and reflections

As Menninga and Goldberg anticipated, I found that using graph data technology in the study of international relations was effective, enabling me to answer most of the questions I had. I was able to highlight the strengths and weaknesses of the technology used, and outline steps for further analysis.

The project progressed as planned and the milestones were achieved. There were also challenges and failures, which I learnt to solve during the process.

One key learning moment occurred when I realised that I had excluded a large amount of data due to a serious error when setting up the database. Fortunately, I found online solutions for enhancing the MERGE clause in Neo4j making sure that all the records were correctly included.

Another moment was when I applied hypothesis testing to one of the research questions. Initially, I thought the first two GenDip questions had been confirmed, but hypothesis testing proved me wrong. This changed my perspective entirely, showing me that I had been focusing too narrowly on just one aspect of the data. From then on, I used hypothesis testing to validate all my findings, improving the quality of the research.

Although there was little evidence of a 'Hillary effect' in Question 3, it motivated me to create my own datasets on female politicians. These datasets revealed also the relationship with conditions in host countries.

Initially, I was disappointed by the lack of statistical evidence for the first two questions. However, this helped me to focus my research on developing alternative explanations that could support the GenDip team's thesis.

Although the community algorithms did not identify any clear sub-graphs that could shed new light on the diplomatic network, I learned that this probably confirmed my previous findings regarding the small-world network characteristics of the diplomatic network.

I have been thinking about this project for a year. Although it required considerable effort, it has been a rewarding experience. My OU studies provided me with the foundation to independently learn IT technologies and apply them to use cases of my choice. The support and guidance of my tutor was also invaluable in helping me reach this point. After eight years of part-time study at the OU, I am very satisfied to conclude my studies with research combining two subjects I am passionate about: data analysis and diplomacy.

References

Gender, diplomacy and international relations

Niklasson, B. and Towns, A. (2023) 'Diplomatic Gender Patterns and Symbolic Status Signaling: Introducing the GenDip Dataset on Gender and Diplomatic Representation', *International Studies Quarterly*, 67(4). Available at: https://doi.org/10.1093/isq/sqad089 (Accessed: 13 September 2025).

Duque, M.G. (2018) 'Recognizing International Status: A Relational Approach', *International Studies Quarterly*, 62(3), pp. 577–592. Available at: https://doi.org/10.1093/isq/sqy001 (Accessed: 13 September 2025).

Menninga, E.J. and Goldberg L.A. (2022) 'Network Analysis', in Huddleston, R.J., Jamieson, T. and James, P. (eds) *Handbook of research methods in international relations*. Cheltenham, United Kingdom ; Northampton, MA, USA: Edward Elgar Publishing. pp. 489-506. Available at: https://doi.org/10.4337/9781839101014.00038 (Accessed: 13 September 2025).

Towns, A. and Niklasson, B. (2017) 'Gender, International Status, and Ambassador Appointments', *Foreign policy analysis*, 13(3), pp. 521–540. Available at: https://doi.org/10.1093/fpa/orw039 (Accessed: 13 September 2025).

Information technology

The Open University (2024), 'Data Management and Analysis'. Module code: TM351, Available at: https://learn2.open.ac.uk/course/view.php?id=216342

Semantic web and linked data

Alexopoulos, P. (2020) Semantic modeling for data: avoiding pitfalls and breaking dilemmas. First edition. Beijing Boston Farnham Sebastopol Tokyo: O'Reilly.

Open University (no date). 'Part 24 Linked data and the Semantic Web' in *Teaching materials TM351*, Available at: https://learn2.open.ac.uk/mod/oucontent/view.php?id=2149320&printable=1 (Accessed: 30 December 2024).

Graph data theory and social networks

Bratanic, T. (2020) 'Community detection of the countries of the world with Neo4j Graph Data Science', *Medium*. Available at: https://medium.com/data-science/community-detection-of-the-countries-of-the-world-with-neo4j-graph-data-science-4d3a022f8399 (Accessed: 8 August 2025)

Golędzinowski, W. and Błocki, W. (2024) 'SOCIAL NETWORK ANALYSIS: FROM GRAPH THEORY TO APPLICATIONS', *Social communication*, 24(1), pp. 151–164. Available at: https://doi.org/10.57656/sc-2023-0012 (Accessed: 13 September 2025).

Marchiori, M. and Latora, V. (2000) ‘Harmony in the small-world’, *Physica A*, 285(3), pp. 539–546. Available at: https://doi.org/10.1016/S0378-4371(00)00311-3 (Accessed: 13 September 2025).

Meghanathan, N. (2017) ‘A Walk Through Social Network Analysis’, in *Graph Theoretic Approaches for Analyzing Large-Scale Social Networks*. United States: IGI Global.

Pei-Hsuan Sung (2025) *Community Structure and Connectivity Analysis in Social Networks Using Community Detection Methods* Available at: https://d197for5662m48.cloudfront.net/documents/publicationstatus/241987/preprint_pdf/1f13252f081515b5232dd28c99dbf002.pdf (Accessed 17 August 2025)

Smith, N. (2020), ‘Exploring Small-World Phenomena with Neo4j’, *Medium. Available at: https://medium.com/neo4j/exploring-small-world-phenomena-with-neo4j-58c767dcbafd* *(*Accessed 17 August 2025*)*

Smith, A.H., (2024) ‘Class 10: Clustering & Community Detection 1 — Traditional’ , Network Science Data, PHYS 7332 (graduate-level course) Northeastern University. Available at: https://asmithh.github.io/network-science-data-book/intro.html (Accessed 3 September 2025)

SPARQL, RDF, Neo4j, databases

DuCharme, B. (2011) *Learning SPARQL.* Available at: https://learning.oreilly.com/library/view/learning-SPARQL/9781449311285/ (Accessed: 26 December 2024).

Needham, M. and Hodler, A.E. (2019) *Graph Algorithms : practical examples in Apache Spark & Neo4j*. 1st edition. Sebastopol: O’Reilly Media. Available at: https://learning-oreilly-com.libezproxy.open.ac.uk/library/view/graph-algorithms/9781492047674/preface01.html (Accessed 13 September 2025).

Scifo, E. (2020) *Hands-on graph analytics with Neo4j : perform graph processing and visualization techniques using connected data across your enterprise*. 1st edition. Birmingham, England ; Packt Available at: https://learning-oreilly-com.libezproxy.open.ac.uk/library/view/hands-on-graph-analytics/9781839212611/946b2dda-31ed-466b-8c4b-4c02bf16f3b0.xhtml (Accessed 13 September 2025).

Scifo, E. (2023) *Graph Data Science with Neo4j: Learn How to Use Neo4j 5 with Graph Data Science Library 2. 0 and Its Python Driver for Your Project*. 1st edn. Birmingham: Packt Publishing, Limited. Available at: https://learning-oreilly-com.libezproxy.open.ac.uk/library/view/graph-data-science/9781804612743/ (Accessed 13 September 2025).

Redmond, E. and Wilson, J.R. (2012) *Seven databases in seven weeks : a guide to modern databases and the NoSQL movement*. 1st edition. Edited by J. Carter. Dallas, Texas ; The Pragmatic Bookshelf. https://learning-oreilly-com.libezproxy.open.ac.uk/library/view/seven-databases-in/9781941222829/ (Accessed 13 September 2025).

Webber, J., 2024. *RDF vs. Property Graphs: Choosing the Right Approach for Implementing a Knowledge Graph*. Neo4j. Available at: https://neo4j.com/blog/knowledge-graph/rdf-vs-property-graphs-knowledge-graphs/ (Accessed 1 Jul. 2025)

Hypothesis testing

Cohen, L., Manion, L. and Morrison, K. (2018) *Research methods in education*. Eighth edition. London: Routledge. Available at: https://ebookcentral.proquest.com/lib/open/detail.action?docID=308686&pq-origsite=primo (Accessed 13 September 2025).

Bobbit, Z. (2022) 'Kruskal-Wallis Test: Definition, Formula, and Example', *Statology*, 7 March 2022. Available at: https://www.statology.org/kruskal-wallis-test/ (Accessed: 7 July 2025).

GeeksforGeeks (2022a). How to perform Fisher's Exact Test in Python. Available at: https://www.geeksforgeeks.org/python/how-to-perform-fishers-exact-test-in-python/ (Accessed 6 June 2025)

GeeksforGeeks (2022b) How to Conduct a Two Sample T-Test in Python. Available at: https://www.geeksforgeeks.org/machine-learning/how-to-conduct-a-two-sample-t-test-in-python/ (Accessed 6 June 2025)

Data sources

Jasom Dotnet (n.d.) *List of capital cities with latitude and longitude*. Available at: https://www.jasom.net/list-of-capital-cities-with-latitude-and-longitude/ (Accessed: 7 June 2025).

Hackett, C. *et al.* (2025) *How the Global Religious Landscape Changed From 2010 to 2020* [text/html,application/pdf]. Pew Research Center. Available at: https://doi.org/10.58094/FJ71-NY11. (Accessed: 15 June 2025).

Niklasson, B., & Towns, A. E. (2023). *The GenDip dataset on gender and diplomatic representation, version June23*. University of Gothenburg: the GenDip Program. https://www.gu.se/en/gendip/gendip-data (Accessed: 1 August 2024).

UNDP (n.d.) *Gender Inequality Index (GII)*. United Nations Development Programme. Available at: https://hdr.undp.org/data-center/thematic-composite-indices/gender-inequality-index#/indicies/GII (Accessed: 21 June 2025).

U.S. Department of State (n.d.) *Hardship post allowance*. Available at: https://allowances.state.gov/web920/hardship.asp (Accessed: 14 June 2025).

Wikipedia (2021) *List of foreign ministers in 2021*. Available at: https://en.wikipedia.org/wiki/List_of_foreign_ministers_in_2021 (Accessed: 3 June 2025).

Licensing information

GenDip dataset

Source: https://www.gu.se/en/gendip/gendip-data

Licensing: There is no specific license information provided by the University of Gothenburg. The documentation asks scholars who wish to use the dataset in their research to attribute the source. Currently I am using the 2023 dataset. GenDip has recently published an updated 2025 version. Prof. Towns offered to send the most recent GenDip dataset.

Heads of State and Prime Ministers
Source: *Wikidata*, Wikidata Query Service
Licensing: All data on Wikidata is available under the Creative Commons CC0 1.0 Universal Public Domain Dedication. This licence allows anyone to copy, modify, distribute, and perform the work, including for commercial purposes, without requiring permission.
Wikidata (n.d.) *Wikidata: Licensing*. Available at:
https://www.wikidata.org/wiki/Wikidata:Licensing (Accessed: 3 June 2025).

Ministers of Foreign Affairs
Source: *Wikipedia*, specifically the page titled "List of foreign ministers in 2021"
Licensing: Content on Wikipedia is generally available under the Creative Commons Attribution-ShareAlike 4.0 International Licence (CC BY-SA 4.0), or is in the public domain. Under this licence, I am required to provide appropriate attribution, indicate any changes made, and distribute any contributions under the same licence.
Wikipedia (2021) *List of foreign ministers in 2021*. Available at:
https://en.wikipedia.org/wiki/List_of_foreign_ministers_in_2021 (Accessed: 3 June 2025).

Hardship
Source: U.S. Department of State, official website – Hardship Post Allowance page
Accessed: 14 June 2025
Licensing: According to the U.S. Department of State copyright policy, unless otherwise indicated, content on State Department websites is in the public domain and may be copied and distributed without permission. Citation of the U.S. Department of State as the source is appreciated.
U.S. Department of State (n.d.) *Hardship post allowance*. Available at:
https://allowances.state.gov/web920/hardship.asp (Accessed: 14 June 2025).
U.S. Department of State (n.d.) *Copyright information*. Available at:
https://www.state.gov/copyright-information (Accessed: 14 June 2025).

Coordinates Capital Cities
Source: Jasom Dotnet, personal website providing public datasets
Accessed: 7 June 2025
Licensing: No explicit copyright is indicated. However, the data appears to be intentionally made available for public use. The site states: “You can use this data in your Google (MapQuest, Mapbox, Yandex, Bing) maps when you want to display country with marker above the capital.” The coordinates are considered to be in the public domain.
Jasom Dotnet (n.d.) *List of capital cities with latitude and longitude*. Available at:
https://www.jasom.net/list-of-capital-cities-with-latitude-and-longitude/ (Accessed: 7 June 2025).

Gender Inequality Index (GII)
Source: *United Nations Development Programme (UNDP)*
Accessed: 21 June 2025
Licensing: All materials on the UNDP website are provided under the Creative Commons Attribution 3.0 IGO Licence (CC BY 3.0 IGO). This licence permits sharing, adaptation, and commercial use, provided appropriate credit is given, a link to the licence is included, and any changes are indicated. The attribution must not suggest endorsement by the UNDP.

UNDP (n.d.) *Terms of use*. United Nations Development Programme. Available at: https://hdr.undp.org/terms-use (Accessed: 21 June 2025).

Religious Affiliation
Source: *Pew Research Center*
Accessed: 15 June 2025
Licensing: In line with the Pew Research Center's mission to inform public debate, content may be accessed, copied, distributed, modified, or used to create derivative works, provided that all copyright and attribution notices are retained. The content must not be used in any way that implies endorsement by the Centre, or that attributes specific policy, political, or religious viewpoints to it.
Licence terms: Pew Research Center Terms of Use
Pew Research Center (n.d.) *Terms and conditions*. Available at: https://www.pewresearch.org/about/terms-and-conditions/ (Accessed: 15 June 2025).

Documentation

Gephi (n.d.) *Gephi documentation*. Available at: https://docs.gephi.org/ (Accessed: 1 September 2025).
Neo4j (n.d.) *Neo4j documentation*. Available at: https://neo4j.com/docs/ (Accessed: 1 September 2025).
Pandas (2025) *Pandas documentation* (version 2.3.2). Available at: https://pandas.pydata.org/docs/ (Accessed: 1 September 2025).
Richardson, L. (2025) *Beautiful Soup documentation* (version 4.13.5). Last modified 24 August 2025. Available at: https://www.crummy.com/software/BeautifulSoup/bs4/doc/ (Accessed: 1 September 2025).
Seaborn (2024) *Seaborn documentation* (version 0.13.2). Available at: https://seaborn.pydata.org/ (Accessed: 1 September 2025).
Wikidata (n.d.) *Wikidata: SPARQL tutorial*. Available at: https://www.wikidata.org/wiki/Wikidata:SPARQL_tutorial (Accessed: 1 September 2025).