Co-funded by the Horizon 2020
programme of the European Union

h2020mirror.eu

MIRROR

Migration-Related Risks caused by
misconceptions of Opportunities and Requirements

Grant Agreement No. GA832921

Deliverable D6.1

Work-package	WP6: Methods for Cross-media Network Analysis
Deliverable	D6.1 First Release of Cross-media SN Analysis Technologies
Deliverable Leader	LUH
Quality Assessor	SAIL
Dissemination level	Public
Delivery date in Annex I	M12, May 31, 2020
Actual delivery date	May 31, 2020
Revisions	2
Status	Final
Keywords	cross-media analysis, network analysis, media bias, information diffusion, users role

Disclaimer

List of Authors

Partner Acronym	Authors
LUH	Erick Elejalde, Miroslav Shaltev, Koustav Roudra, Tuan-Anh Hoang, Claudia Niederée
FOI	Joel Brynielsson, Johan Fernquist

Table of Contents

List of Figures

List of Tables

Executive summary

Within the MIRROR project, the Cross-media Network Analysis Technologies cover the exploration and development of analysis methods for migration-related networks of multi-source multi-modal data. These methods will provide the end-users (such as border control agencies) with tools and frameworks for exploring, understanding, and inspecting the possible sources of misperceptions that can result in threats.

We designed and implemented our Cross-Media Network Contruction (CNC) component following the use cases identified in WP2 during the end-users' requirements analysis. Based on these scenarios, we recognize the following three groups of sources: news media, social media, and an additional group of miscellaneous sources, which do not fit in the first two categories. Besides sources of data, for each scenario, we identify the most important entities, which are subjects of interest. The design of the CNC also incorporates the analysis made about to the Migration Related Semantic Concept (MRSC). We use a multigraph as the underlying structure for the representation of the unifying graph. By applying filters on the nodes and edges, we can retrieve subgraphs, and through aggregation and network analysis, we generate knowledge according to the user needs. We establish the relation between the MIRROR Information Model (introduced in Deliverable 7.2) and our graph nodes and edges' properties. Finally, the Cross-Media Network Construction is implemented as a web service with four main subsystems: CNC-parser, CNC-gate, CNC-analyser, and CNC-gui. The CNC-parser and the CNC-gate components are written around Janusgraph. Following the microservice architecture of the MIRROR system, these components are isolated and attached to the whole system via Docker containers. The communication within the system is done via the REST API.

When trying to understand the formation of immigrants' perception, the appropriate tools for detecting and investigating information bias is of utmost importance. Our perception of the situation in a country or a region is strongly influenced by the reflection of this situation in mass and social media channels. With this in mind, we perform a detailed analysis of the representation of Europe in the news across multiple geographical regions. For a deeper understanding of coverage, we analyze the news, which are directly adopted from the European news sources (i.e., overlapping of content with the EU internal media). We carry out our evaluation by exploring the differences and possible bias not only in the general news context but also by focusing on specific topics. In particular, we are interested in subjects that contribute to creating an image of Europe and may act as pull and push factors according to migration-related literature. Notably, we can see that the image of Europe across the world is not uniform and significantly differs from its home representation. This trend holds relatively uniformly across the overall news as well as for topic-specific news. Only 'Science and Technology' gets some appraisal in Asia and Africa. We also found that America has an under-representation of positive news compared to the EU, paired with an over-representation of negative stories. Interestingly, other regions tend to cover many EU-related events that are absent in the media outlets of Europe. These news that are skipped by EU-media are covered by most of the other zones positively. This study is an essential first step towards better understanding the external perception of Europe by audiences across the globe.

Complementary to the analysis of representation bias, we worked on the evolution of migrants' networks and their communities over time. The methods being developed incorporate and leverage results and findings of textual and multimedia analysis so that to better profile the communities using multi-modal data. In our first year, we have focused on the identification and characterization of several relevant communities (such as refugee and silent users) for the study of migrants' perception. In particular, we have focused on communities among Twitter users who are more engaged with refugee and immigration-related topics. Furthermore, we are interested in identifying communities among the refugees and migrants and how they interact with the communities in the host countries. The understanding of these communities' structure

and their intentions would be highly useful for a wide range of important applications (e.g., social sensing and social surveying). This has not been well studied in previous work for our domain. One particularly relevant subgroup of Twitter users is that of less active accounts, which often do not directly express their opinions online. Twitter has long been the focus of great interest as a lens through which to observe large swaths of society. However, such analyses have relied exclusively on the side of active participation and content production. The main issue with this is that content-based inferences may fail to generalize to real-world populations, or even to less active users. Although silent users rarely post on social media, they still have opinions towards different topics. If we don't take their hidden views into account, the predictions could be far from the truth. In MIRROR, we are developing a methodology to predict the opinions of the users who are "silent" in social networks like Twitter. At present, we have collected a large dataset of influential users and their followers and related tweets. Our next step will be examining the aspects and sub-topics discussed in the extracted tweets. These include analyzing and visualizing the aspects' and sub-topics' temporal dynamic, across countries and languages.

Following, we will continue with the implementation of the CNC component. Methods developed for the multi-modal analysis of the network will be continuously integrated and validated for performance and relevance with technical and application partners.

We will extend our analysis to integrate all the data sources considered in MIRROR. We will also check the differences between regional and English channels of a specific region. One is mostly used to reach global audiences, while local audiences usually consume information posted in their native languages. Hence, it is essential to analyze such representation gaps, if any. We also intend to extend and refine the topical coverage of our analysis. Likewise, so far, we have considered regions that cover large geographical areas, but they might internally follow different patterns of information consumption. We will investigate the optimal level of aggregation for different dimensions and the rate of information loss as an additional input for end-users.

The next steps in community analysis involve the development of methods to improve our understanding of how these identified communities evolve or react to specific events.

1 Introduction

1.1 Objectives and positioning in the MIRROR system architecture and scenarios

The goal of the work being developed in the Cross-media Network Technologies for the MIRROR project is to explore, examine, adapt, and develop analysis methods for migration-related social networks collected from multiple available sources. These methods will provide to the end-users (such as border control agencies) the tools and frameworks for exploring, understanding, and inspecting the possible sources of misperceptions that can result in threats. Additionally, it will provide methods for analysis of the misperception evolution across the geographical areas of interest as well as along the migration paths.

During the first stages, we have focused our efforts in the following aspects:

- Design and implementation of a Cross-Media Network Construction component.
- Understanding of some possible sources for information bias and echo chamber effects.
- Identification of relevant networks and communities and their affinity to migration related topics.
- Development of explanatory bot detection methodology for understanding automated communication about migration.

The Cross-Media Network Construction translates the input coming from the Multimedia System (MMS), as well as WP4 and WP5, into data structures suitable to be written into the unified graph model in which the component is based. Figure 1 presents a simplified view of the full MIRROR system architecture and the positioning of the Cross-Media Network Analysis container within the system. This container represents the technologies developed in WP6. The payload of the HTTP POST methods have been design in accordance with the MIRROR data model presented in deliverable D7.2.

The output of the analysis in this work package also give feedback to the text and multimedia analysis in WP4 and WP5 and the results will be evaluated in WP10 with the end users. Furthermore, links to the results of the threat analysis in WP9 will be created.

All the analysis and results included in this report have been developed within the MIRROR project.

1.2 General data protection considerations

We have paid special attention to the data protection aspect during collection and manipulation of all the datasets used in MIRROR. All of them are stored in a LUH's servers dedicated to the MIRROR project. Restriction measures are implemented in this server to grant selective access to the datasets according to the requirements and the members of each group within MIRROR. These datasets are specified and handled based on the methodology presented in MIRROR Data and Knowledge Management Plan (Deliverable D1.2).

We also use some pre-existing datasets. They are maintained and are distributed to the scientific community by third-parties (outside the MIRROR consortium). The methodology, standards and metadata used in these datasets are those defined by their respective owners. In MIRROR, we use these datasets for developing and evaluating our analysis methods and comparison against our own datasets, all in accordance with the licenses that accompany each of these datasets.

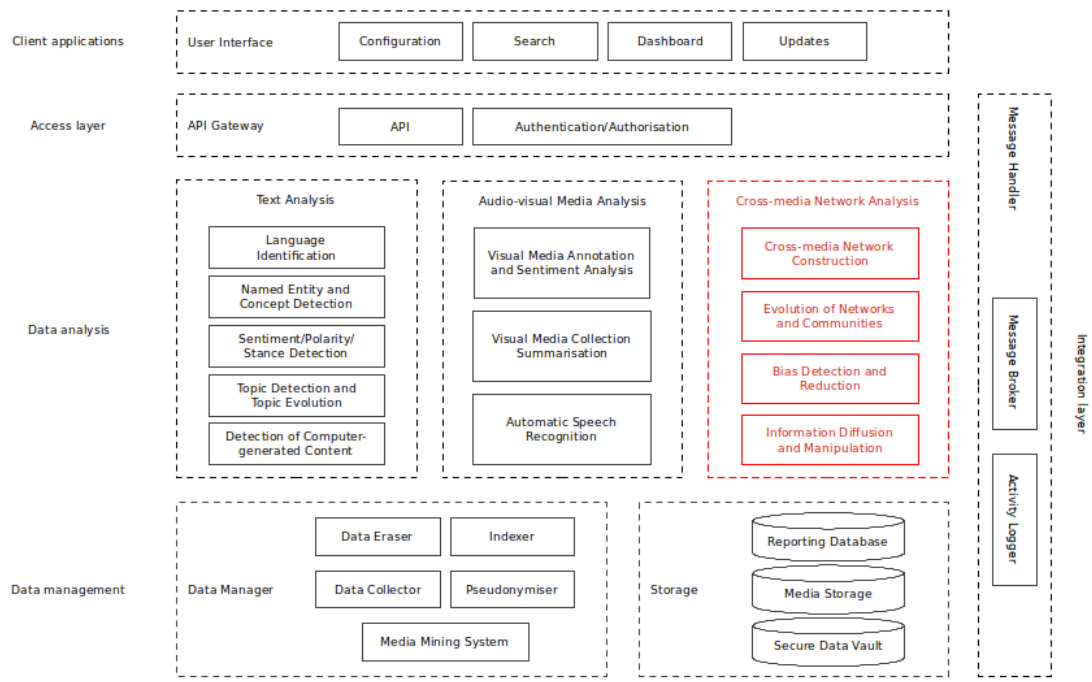

Figure 1: MIRROR system architecture overview. The red rectangles indicate the container and its constituent components that are described in the present document.

The GDELT dataset used in some of the experiments described in this document is not included in deliverable D1.2 as it was considered after the submission of the referred document. This dataset does not contains any personal data. It is only composed of news articles generated by public news outlets. The description of its content and methodology for sub-sampling are described in Section 3.1.4.

Regarding the Cross-Media Network Construction component, we have followed the security-by-design principle recommended in the MIRROR Architecture and Integration Plan (deliverable D7.1). In Section 2.3, we give the details of the multi-layer architecture of the component and security considerations.

1.3 Structure of the document

We start in Section 2 with the discussion of the design and construction of the networks in accordance with the use cases identified in WP2 during the end-users' requirements analysis. Specifically, in 2.1 we take a closer look on the Identified information sources. The formal definition of the graph and its design is given in subsection 2.2. In 2.3 we are concerned with the implementation of the cross-media component. Finally, we discuss the integration of the component in the MIRROR system in 2.4.

In Section 3 we present a detailed analysis of the representation of Europe in the mass media across multiple geographical regions. We make the analysis in two steps. In 3.1 we deal with a scalable computational methodology to detect the variations of news in other media outlets in terms of coverage and polarity. Here we use all the news from all media outlets in our dataset, regardless of the origin. However, most of the news reported by the EU home-media get filtered in other regions. Hence, in 3.2, we focus our analysis to a subset of documents where the events is covered by at least one of the media outlets in Europe and and the source country of the news document is in another zone.

We continue in Section 4 with the description of the initial steps that have been conducted in MIRROR for the analysis of social networks communities. This analysis has also been divided in two relevant and interdependent subsections. First we present progress made in the identification and characterization of Twitter communities around refugee and immigration topics in 4.1. This analysis is complemented in 4.2 with a dive into the study of silent Twitter users modeling for opinion prediction.

We offer our conclusions of the work done in Section 6.

1.4 History of the document

Section of the Document	**Date**	**Change**
Overall	10.05.2020	Preparetion of document V1 with adequate formatting
Section 5	23.05.2020	Adding information about infor-mation diffusion
Overall	28.05.2020	Created V2 to include suggestions from QA reviewers

2 Cross-media Network Construction

The objective of this task is the construction of networks in accordance with the use cases identified in WP2 during the end-users' requirements analysis.

2.1 Identified information sources

In accordance with deliverable D2.1, the four identified scenarios can be summarized as follows. Scenario 1 is concerned with detection of hybrid threats driven by perception manipulation. In Scenario 2 the objection is to counteract threats created by misperception. The support for border agents with targeted information is subject of Scenario 3. In Scenario 4 we are focused on the help for the cyber intelligence analyst when it comes to an ordinary day at work. Based on these scenarios we identify the following three groups of sources, namely: news media, social media and an additional group of miscellaneous sources, which do not fit in the first two categories. The news media is broadly represented by the online versions of printed news media and online only resources delivering news content. The social media on the other hand is dominated by the famous Facebook and Twitter social networks. The last group contains some special sources, like the Tor network, for example. For each scenario, we could specify the possible sources, as well as the most important entities which are subjects of interest. These are given in Table 1. The choice of technology will be discussed later in the section. We first focus our attention on the unified graph design.

Scenario	Sources	Main subjects of interest
1. Perception Manipulation	Social Media	Source, Country, Concept, Topic
2. Misperception	News and Social Media	Source, Country, Concept, Topic
3. Targeted information	Social and News Media	Entities, Topic
4. Cyber Intelligence	Additional Sources	Entities, Topic

Table 1: Identified information sources.

2.2 Unified graph design

Even though we identified different possible scenarios for the application of the MIRROR system, they all share information sources. With this, it is natural to introduce a common underlying structure, which is suitable for our analysis: the multigraph. Further, the fusion of the information derived from the analysis in WP4 and WP5 represents the main input for this component. This information flow will allow us later to create nearly aligned networks. Then, by applying filters on the nodes and edges, we can retrieve subgraphs, and through aggregation and network analysis, we generate knowledge according to the user needs.

Following the definitions of [Bollobás, 1998], we define our graph G as:

$$G = (V, E)$$

with V being the vertex set of G, and E being the edge set. A multigraph is then a graph with multiple edges and multiple loops. An example of a subgraph is given in Fig. 2. In the following we idenify different objects with the vertices of the multigraph and refer to these as the nodes of our unified graph.

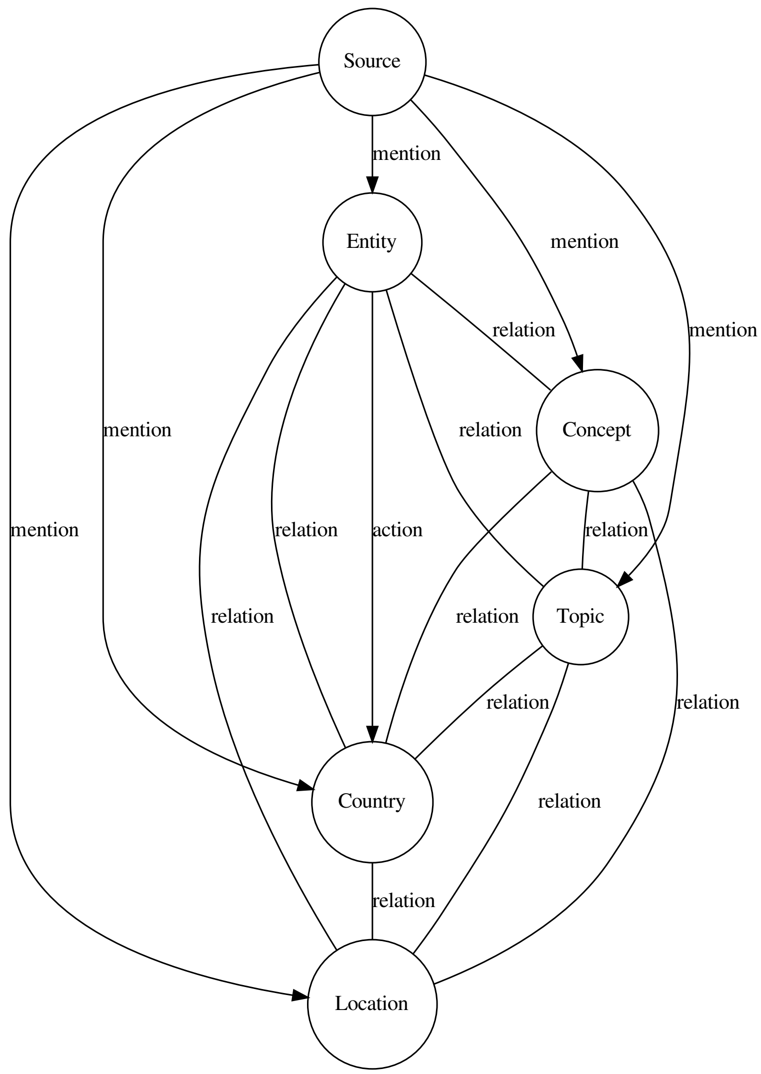

Figure 2: Property graph. Example of a subgraph of the unified graph model.

2.2.1 Node definition

We define a node simply as an object of type "node", with a set of properties. The properties taken into account at this stage of the MIRROR System are summarized in Table 2

The mandatory properties required to create a node are: id, type and name. All other properties may be presented or omitted, depending on the type of the node. The "id" is unique. It is a hash created from the name property, or for nodes of type entity, the name and entitytype properties. For nodes of type "location", the "id" is computed from the name and geolocation properties. The "type" property is one of the set given in Table 2. The "name" property could be any string. The "sourcetype" is a string of the set given in Table 2. The "sourcename" is an arbitrary string. The "entitytype" is a string from the set given in the Table 2. The "geolocation" property are the GPS location of a node of type "Location".

Property	Type	Values
entitytype	string	account, person, organization
geolocation	float	GPS coordinates
id	hash	
name	string	
sourcename	string	
sourcetype	string	news, social media, channel
type	string	source, entity, topic, concept, country, location

Table 2: Node properties.

2.2.2 Edge definition

An edge is an object of type “edge”, connecting a node with id “source” with another node with id “target”. Similar to a node object, we attach a set of properties to any edge. The properties of an edge are given in Table 3. The “type” property is a single string, one of the list “Action”, “Mention”, “Relation”. The “datetime” property is the publication date of the the source document. The “name” property is a string describing the type of the edge. The “timeframe” is a temporal information extracted from the source document. The “authenticity” property is a measure of our belief in the authenticity of the source document and with this the extracted relation. The ”credibility“ property is a floating number used to distinguish between tabloid, and broadsheet sources. The “sentiment” property is a floating number, characterizing the sentiment. Authenticity, credibility and sentiment are in the range of 0 to 1. Finally, the “docid” property is an auxiliary information identifying the full document used to extract the edge.

Property	Type	Example
annotation	string	derived, metadata
authenticity	float	0.95
credibility	float	0.80
datetime	string	20200101T120000Z
docid	string	
language	string	
mediatype	string	audio, text, image, video
name	string	commention, follows
sentiment	float	0.314
timeframe	string	two days ago
type	string	Action, Mention, Relation

Table 3: Edge properties.

The relation of the nodes and edges’ properties used in our unified graph model to the information model of the MIRROR system (introduced in Deliverable 7.2) will be given in the next section, where we discuss the implementation of the cross-media network construction component.

2.3 Cross-media Network Construction Component

The cross-media network construction component is implemented as a web service with four main subsystems, namely CNC-parser, CNC-gate, CNC-analyser and CNC-gui. The CNC-parser and the CNC-gate components are written around Janusgraph [1] , an open source distributed graph database using the Gremlin [2] graph traversal language. The CNC-analyser makes use of the NetworkX [3] library.

2.3.1 CNC-parser

The CNC-parser component translates the input coming from the Multimedia System (MMS), as well as WP4 and WP5 components, into data structures suitable to be written into the unified graph model described above.

I/O description

- Input: JSON objects generated in MMS, WP4, or WP5.
- Output: records in the CNC database. Expressed in JSON

  {“type”: “node”, “id”: “id”, “properties”: {}}.

  {“type”: “edge”, “source”: “id”, “target”: “id”, “properties”: {}}.

Technical Details - Instructions of Use

- HTTP POST request on /process, which initiates the transformation of data. The body of the POST request must be a JSON with the following fields:
  - “payload”: a JSON containing the actual data to parse.

In accordance with the MIRROR data model presented in deliverable 7.2, the payload of the HTTP POST method should provide the following fields:

- annotation - mapped to the annotation property of an edge.
- authenticity - mapped to the authenticity property of an edge.
- concept - list of concepts, with each entry mapped to the name property of a node of type “Concept”.
- country - list of countries, with each entry mapped to the name property of a node of type “Country”.
- credibility - mapped to the credibility property of an edge.

[1] https://janusgraph.org/
[2] https://tinkerpop.apache.org/gremlin.html
[3] https://networkx.github.io/

- datetime - mapped to the datetime property of an edge.
- geolocation - mapped to the geolocation property of an edge.
- language - mapped to the language property of an edge.
- location - listed of locations, with each entity mapped to the name property of a node of type “Location”.
- mediatype - mapped to the mediatype property of an edge.
- name - mapped to the name property of a node of type “Source”.
- sourcename - mapped to the name property of a node of type “Source”.
- sourcetype - mapped to the sourctype property of a node of type “Source”.
- origin - mapped to the origin property of an edge.
- sentiment - mapped to the weight property of an edge.
- timeframe - mapped to the timeframe property of an edge.
- topic - list of topics, with each entry mapped to the name property of a node of type “Topic”.
- entities - list of tuples, entity name, entity type, with each entry name mapped to the property name of a node of type entity type.
- relations - list of triplets, entity name 1, relation, entity name 2, with each relation mapped to the name property of an edge of type relation between the node with entity name 1 and the node with entity name 2.
- actions - list of triplets, entity name 1, action, entity name 2, with each action mapped to the name property of an edge of type action between the node with entity name 1 and the node with entity name 2.
- docid - unique id useful to retrieve the original document.

The processing of the payload is as follows.

- Create a node of type “Source”.
- For each item in the country, location, entity, topic or concept list, create a node of the corresponding type. For an item from an entity list, the type could be one of “Account”, “Organization” or “Person.
- Create an edge of type “Mention” between the “Source” node and each of the other created nodes.
- Creates an edge of type “Relation” with “name” property set to “co-mention” between each of the node, except the “Source” node.
- Creates an edge of type “Relation”, resp. “Action”, according to the relations, resp. actions list.

2.3.2 CNC-gate

The CNC-gate is the component serving queries to the CNC database. It has two basic modes of operation, an exploratory free search and an advanced search based on filters.

I/O description

- Input: a JSON objects.
- Output: a JSON object, including a subgraph in the JSON format.

Technical Details - Instructions of Use

- HTTP POST request on /search, which initiates a search in the data.
- HTTP GET request on /result, which returns the data found for a POST request.

The body of the POST request must be a JSON with the following fields:

- “type”: one of “freesearch” or “advancedsearch” indicating the type of the search.
- “begin”: earliest value of the “datetime” property of the edges.
- “end”: latest value of the “datetime” property of the edges.
- “filter”: a string for type “freesearch” or a JSON of filters for type “advancedsearch”.

For processing requests of type “freesearch”, the submitted string in the “filter” field undergoes natural language processing. A query is then constructed and send to the graph database. For processing requests of type “advancedsearch” the query of the database is constructed from the posted list of filters. The structer of each filter is of the form “property:value”, e.g.,

$\{$”type”:”country”,”name”:”Malta,Greece,Turkey,Syria”$\}$

A logical ”or” operator is applied between each filter and the properties of the filter.

Each POST request returns a processing “id”, which should be used in a GET request. A GET request returns a JSON with the following fields:

- ”status“: current status of the processing.
- ”result“: the subgraph in the JSON format.

In case of an error, or if the subgraph is not retrieved yet from the database, the result field of the response contains an empty subgraph.

2.3.3 CNC-analyser

The CNC-analyser performs aggregation on the subgraph and provides basic analysis results for a given subgraph input. The temporal evolution of the data is tracked through the datetime property of the edges, which allow us to compute frames of the graph.

I/O description

- Input: A subgraph as a JSON.
- Output: A JSON with analysis results.

Technical Details - Instructions of Use

- HTTP POST request on /process, which initiates analysis of the posted data.
- HTTP GET request on /process, which returns the result of the analysis.

The body of the POST request must be a JSON with the following fields:

- "timestep": indicating the requested time granularity.
- "graph": the graph to be processed.

Based on the value of the "timestep", the "graph" is split into subgraphs using the "datetime" property. Then each subgraph gets aggregated and written into a NetworkX structure, in order to apply standard algorithms and obtain measures. The aggregated subgraphs and corresponding computed values are returned in JSON format for further processing by the MIRROR system.

2.3.4 CNC-gui

The CNC-gui component is the graphical user interface complementary to the CNC-gate service and the CNC-analyser. It gathers input from users, queries the graph database through CNC-gate and visualizes the obtained results after analyses with the CNC-analyser.

2.4 Integration into MIRROR system

Following the microservice architecture of the MIRROR system, the CNC-parser, CNC-gate and CNC-analyser components are isolated and attached to the whole system via Docker technology. The communication within the system is done via the REST API described in the previous subsection.

3 Bias Detection and Reduction

Our perception of the situation in a country or a region is strongly influenced by the reflection of this situation in mass and social media channels. To avoid information overload, news outlets typically filter the available news from foreign countries based on the expected interest of the target audiences. The objective of this section is to identify two different kinds of biases: (i). social biases introduced by data seeking and retrieving methods or biases originally introduced by social selection and alignment of data sources, (ii). biases introduced by data processing and analysing methods. In this section, we systematically analyze the bias created in EU-related news reports due to editorial policies (bias type ii).

form a detailed analysis of the representation of the EU in the news across multiple geographical regions (EUROPE(NON-EU), ASIA, AFRICA, MIDDLE-EAST, AMERICA, and OCEANIA) over one year (January-December 2018). For a deeper understanding of coverage, we also analyze the news, which are directly adopted from the EU news sources (i.e., overlapping of coverage with EU internal media).

In general, not all the news reported in other zones are also published in EU internal media and vice-versa. Interestingly, other regions tend to cover many EU-related events that are absent in the media outlets of the EU(EU-MISSING news). That is why we define the following two categories of EU-related news based on their presence on EUinternal-media: **1. EU-TOTAL:** All the EU-related news (i.e., events linked to any of the 28 EU countries) over all the media outlets in our dataset regardless of the origin. **2. EU-COVERED:** This is a subset of EU-TOTAL where *the events must be covered by at least one of the media outlets of the EU region, and the source country of the news document must be in another zone.* For example, news about 'Eu's decision about Iran's nuclear deal were reported both in Germany and the US, however, in EU-COVEREDwe only include the documents originated in US. On the other hand, Jamaica's deal with Ireland for potatoes was covered in Jamaican but not in Irish media. In this last case, this event is not included in EU-COVERED.

We perform a detailed analysis of EU-TOTAL and EU-COVERED news across different regions to get an initial understanding about the variation in the representation of the EU in different parts of the globe.

3.1 Coverage of EU-TOTAL news in other geographic regions

3.1.1 Problem statement

In the international sphere, the problem of a distorted external image is of more direct concern for the decision policy makers (DPMs). For example, in cases where it endangers their influence, bilateral economic and cooperative treaties, or the effectiveness of public diplomacy towards conflicted neighbors [Chaban et al., 2019, Kelly, 2013]. A conflict between the internal and external views of an issue can lead to precarious development. For example, according to a report on media coverage of the "refugee crisis" in Europe [Georgiou and Zaborowski,], the press played a central role in shaping the public opinion and declaring the "refugee crisis". As Europe still faces multiple challenges from this crisis [UN Refugee Agency, 2019], it is of great interest to continue studying its image. In this section, we develop a scalable computational methodology to detect the variations of EU-TOTAL news in other media outlets in terms of coverage and polarity.

3.1.2 State of the art

The perception of Europe, both from inside and abroad, has been studied in a qualitative manner in several studies[Machill et al., 2006, Chaban and Holland, 2013, Chaban et al., 2019, Kelly, 2013, Timmerman et al., 2014]. Typical for the qualitative approach, the focus is on interviews and the analysis of small data samples. In contrast, quantitative research, as it is used in our work, offers the opportunity of more large-scale analysis. This is facilitated by automated analysis methods relying, e.g., on machine learning techniques, for inspecting higher volumes of data. Despite the current availability of massive amounts of digital records of news reporting, quantitative analysis of the reflection of Europe in the mass media are scarce. In this work, we introduce a methodology for the quantitative investigation of the reflection of European topics in the news as a foundation for better understanding perception of this region. Notwithstanding, we try to leverage some of the valuable insights produced by the qualitative literature on the subject.

An influential set of qualitative studies dealing with the image of Europe and of the European Union (EU) is presented in [Chaban and Holland, 2014]. When combined, the studies in this book cover data and events recorded in more than a decade of EU transformation and evolution (i.e., 2000's). How essential events, such as the Lisbon Treaty and Eurozone debt crisis, impacted the role and image of Europe as a global actor, is analyzed in detail. Similar to these studies, in our analysis we understand '*image*' as 'the total cognitive, effective, and evaluative structures of the behavior unit, or its internal view of itself and its universe' [Boulding, 1959].

All the previous findings show the need for further analysis of the media agenda in general. In recent years, a large number of quantitative studies addressed the behavior of the media from a 'bias' perspective. Recent studies [Groseclose and Milyo, 2005a, McCluskey and Kim, 2012, Ribeiro et al., 2018, Budak et al., 2016] tried to estimate the ideological score for several major news outlets. Keywords present in the news titles have been used in [Dallmann et al., 2015] to measure the relative bias between four leading newspapers in Germany. In [Elejalde et al., 2018], the authors tried to identify sentiment expressed across various topics to measure the leaning of outlets with respect to political, social, and economic issues. In this deliverable, we also measure polarity of EU-related news but at a different level of aggregation. Our objective is to expose possible discrepancies in the narrative favored in reporting EU-related news in different regions.

Topic modeling has also been used to compare public versus media agenda [Pinto et al., 2019]. A scalable approach to measure the ideological leaning of different media through their Facebook outlets has been present in [Ribeiro et al., 2018]. Recently, different forms of media bias have been reported [Hamborg et al., 2018], namely, 'event selection [Bourgeois et al., 2018, Gruenewald et al., 2009, Saez-Trumper et al., 2013]', 'source selection [Groseclose and Milyo, 2005b, Sanderson, 1997, Agirre et al., 2016]', 'labeling and word choice [Papacharissi and de Fatima Oliveira, 2008, Bhowmick et al., 2009, Grefenstette et al., 2004]' etc. All such models aim at identifying and predicting different forms of media bias. However, our goal is to analyze the post-selection phase of news outlets. In this report, we try to measure the coverage and polarity variation of EU-related news in the outlets of different geographic zones which primarily consists of post-event news analysis and effect of wording.

3.1.3 MIRROR approach

In this section we elaborate on the method employed in our analysis. Since we are interested in analyzing how the mass media portray Europe (EU) from outside its borders, our global strategy is to compare the coverage of European issues across different geographic zones. We use the internal European coverage as a baseline, assuming they will constitute the origin for most of the reporting. To some degree, this should also show how the image that Europe tries to present of itself morphs based on different geopolitical interests.

We collected event-centric data for each of the EU countries for Jan-Dec 2018 (details are mentioned in Section 3.1.4). Although individual countries' imagery of Europe is of high interest, we will further restrict our analysis to bigger geographic areas. Previous works have shown that neighboring countries and those sharing strong cultural and economic ties will cover issues more similarly [Flaounas et al., 2010]. Here, we check the image of EU across seven different geographic zones: (i). Eu, (ii). Europe(non-EU), (iii). Africa, (iv). Asia, (v). Middle-East, (vi). America, and (vii). Oceania (see Table 4 for more details).

We carry out our evaluation by exploring the differences and possible bias not only in the general news context but also by focusing on specific topics. In particular, we are interested in indicators that contribute to creating a perception of the Quality of Life in Europe [Europe Durect, 2019] (e.g., economy, health, education, leisure, security, etc.). However, we also include other general topics that may act as pull and push factors according to migration-related literature [Borjas, 1989a, Castles et al., 2013]. As European countries are a frequent target for immigration, we expect these topics to be of interest to the international press. In this report, we consider eleven broad category of wiki-topics [Hoang et al., 2018] (detailed statistics are mentioned in the next section).

3.1.4 Experiments, datasets, results and future work

Dataset: In this section, we take a closer look at the dataset used in this deliverable. A comprehensive and diverse dataset covering a large number of countries is required for our analysis. For this purpose, we start from an extensive collection of almost 200K news outlets compiled by GDELT Technology[4]. This collection contains outlets from all continents and provides one of the broadest samples of the news media global landscape available.

To map news sources to their host countries, GDELT bet on the strong geographic bias ingrained in most news institutions' editorial policies. News outlets work on an economy of scale with a substantial first copy cost. Thus, outlets will give priority to stories where their reporters can get quickly and easily (again, to minimize the cost of the piece of news). According to Zipf's Gravity Model [Zipf, 1946b, Zipf, 1946a], the interest of a piece of news decreases as we move farther away from the source of the event. This behavior has also been tested for online media [Elejalde, Erick et al., 2019]. Hence, outlets will be most probably located in their "primary" country of focus. In general, GDELT has defined an affinity-first approach that maps news outlets to their host country based first on their top-level domain in the Domain Name System of the Internet (e.g., news sources with **.at** domain are assigned to **Austria**), then to their primary country of focus, and finally to the country where they are incorporated, or the entity that controls their domain is registered[5].

We have used the following three types of GDELT collections: 1. **Event:** This collection contains information about events found in the world's news media. Along with standard attributes (eventid, date, etc.), it also contains event-specific attributes such as entities involved in the event, location of the event, ethnic and religious communities involved in the event etc, 2. **Global Knowledge Graph (GKG):** This collection contains lower level multi-modal information extracted from different source identifiers (news urls, blogs, etc.) and is represented as a knowledge graph. It also provides information about named entities (person,

[4]https://www.gdeltproject.org/
[5]https://blog.gdeltproject.org/mapping-the-media-a-geographic-lookup-of-gdelts-sources/

organization), tones, topical themes, embedded images and videos, etc, 3. **Mention:** This collection contains information about all the references of an event present in the Event collection. In other words, Event includes information on the first appearance of an event, and all the subsequent mentions of that same event are added to the Mention dataset.

As mentioned in Section 3.1, the primary objective of this work is to capture the variations in the representation of the EU in different regions of the world. Hence, we extract information about all the events which happened in any of the 28 constituent EU countries over the year 2018. Our data creation procedure is described in details below:

1. First, we identify all the EU-centric events from the event database. We extract information from the following three fields: (i). **Actor1Geo_CountryCode:** contains geographic information (country code) about the first actor involved in the event. It captures the closest geographic reference to the Actor1 mentioned in the document, (ii). **Actor2Geo_CountryCode:** contains geographic information (country code) about the second actor involved in the event, and (iii). **ActionGeo_CountryCode:** country code of the location where the action is performed by actor1 on actor2. In some cases, there might be only one actor; in those cases, the actor2 related field is empty. For example, 'French Assistant Minister Smith was in Berlin last week.' marks 'Actor1Geo_CountryCode' as 'GM(Germany)'. Basically, an EU-event is one that took place in any of the 28 EU countries. We collect all the events monthwise for each of the 28 EU countries for the year 2018. Each event is uniquely identified by 'GlobalEventID' field.

2. As mentioned before, Event database contains only the first mention of an event. All the followup mentions of the events are present in the Mention dataset. Hence, we extract all the mentions of an event within a lookahead period of three months by querying this collection based on the 'GlobalEventID' field. We decide three months lookahead period because the number of mentions of an event almost becomes zero after a period of three months. From the retrieved entries, we get the identifiers of the documents that contain the mentions of the desired events.

3. Finally, using these document identifiers collected in the previous step, we query the knowledge graph database (GKG). For each document, we extract all its metadata, such as themes, organizations, locations, and other content analysis measures. After this step, we have country-wise information about all the events, their mentions, and multi-modal information extracted from the documents containing mentions of the events.

As mentioned earlier, this dataset also contains information about host countries of different news channels. For example, 'zznews.cn' and 'edgehospitality.ca' are hosted in China and Canada respectively. The GDELT repository contains host country mappings for around 0.19M source URLs. The motivation behind this country-based mapping of news channels is two-fold. First, we perform the whole analysis under the assumption that news channels hosted in a country influence the perspectives in that country towards local and global (worldwide) issues. In other words, European issues presented by a news channel in a country influence the perspectives in that country's audience about Europe. Second, this strategy allows to aggregate news channels at different geographic levels based on their hosting countries (e.g., German, European, etc.).

Further, these host countries are mapped to one of six geographic regions. We consider the following six geographical regions — (i). EU, (ii). EUROPE(NON-EU), (iii). ASIA, (iv). AFRICA, (v). MIDDLE-EAST, (vi). AMERICA, and (vii). OCEANIA. Some regions are further divided into sub-regions (e.g., AFRICA is divided into north, east, and south sub-regions). The region-wise country distribution is obtained from the United Nations

Zone	# Sub-regions	# Countries	# Distinct News sources
EU	1	28	58264
EUROPE(NON-EU)	5	26	16934
ASIA	5	35	18868
AFRICA	5	60	6695
MIDDLE-EAST	1	14	3612
AMERICA	4	57	75102
OCEANIA	4	29	10045

Table 4: Data statistics of news sources over different geographic regions.

database[6]. Antarctica is dropped from this list due to unavailability of the news outlets. In Table 4 we show the distribution of sub-regions, countries, and news sources for each region.

In Table 5 we show the detailed region-wise dataset distribution of the EU countries cumulated over the period of January to December, 2018. Each event can have multiple mentions i.e., different news articles might cover the same event. Each of them form a separate piece of news. Hence, *Events* report the number of unique events captured in that period. For some of the news urls, either the host country is missing or it is not possible to map that country to any region. Such entries are dropped from this deliverable. This set contains both global (English) and local (regional language) news outlets. Non-English news documents are translated using machine translation scheme[7].

After this step, we have the collection of news covering EU-events over the period Jan-Dec 2018. In total, we were able to gather 30M news documents. This final dataset of news constitutes our corpus for all further analysis presented in this report.

For each document, GDELT provides further annotations of sentiment-related attributes such as *tone*, *positive score*, *negative score*, *polarity*. These scores are identified by the Global Content Analysis Measures (GCAM) system. Each document is also associated with categorical themes (e.g., TAX_FNCACT, HUMAN_TRAFFICKING, HEALTH_VACCINATION). The system recognizes 284 general set of themes[8]. Apart from general themes, the current GDELT system also identifies several specific themes (e.g., TAX_FNCACT_CARTEL is a special case of TAX_FNCACT). Altogether the system identifies a total of 56,840 themes combining general and specific ones[9].

For each document, we also have a list of associated named entities such as persons, locations, organizations etc. Such information is automatically extrcated from the text using Leetaru algorithm[10]. This name recognition system performs well in recognizing Asian, African, Middle Eastern names along with standard English names. For locations, GDELT gives different set of information such as country code, city, state, longitude, latitude etc.

The documents are categorized into 284 general themes by GDELT system. We further annotate these themes into 11 broad wiki topics as proposed in [Hoang et al., 2018]. For this mapping, two independent coders (two of the authors) annotated each of the GDELT themes as belonging exclusively to one of the 11 topics. In a second stage, coding disagreements were solved through a negotiation among coders in order

[6] https://unstats.un.org/unsd/methodology/m49/
[7] http://www.statmt.org/moses/
[8] http://data.gdeltproject.org/documentation/GDELT-Global_Knowledge_Graph_CategoryList.xlsx
[9] http://data.gdeltproject.org/api/v2/guides/LOOKUP-GKGTHEMES.TXT
[10] http://www.dlib.org/dlib/september12/leetaru/09leetaru.html

Countries	# Distinct Events	# News	Zones							
			Eu	Europe(non-EU)	Asia	Africa	Middle-East	America	Oceania	NA
AU	172	466	61.1	22.2	50.9	11	38.6	265	10.7	5.41
BE	328	1131	333	27.8	87.6	31.1	64.4	550	28.9	7.08
BU	100	258	82.4	12.1	22.3	3.66	14.3	115	4.32	3.76
CY	73.8	199	43.5	7.87	16.5	5.44	9.65	112	2.61	1.81
EI	622	1681	1109	7.88	39.9	15.1	30.8	429	43.1	5.44
EN	51.3	132	32.2	5.04	10.7	2.32	3.02	74.1	3.69	1.69
EZ	101	256	61.4	8.04	22.5	5.07	13.5	135	7.29	3.42
FI	125	412	49.1	10.7	36.7	7.12	12.1	281	13.1	2.76
FR	1199	4090	596	80.8	417	130	224	2430	175	35.8
GM	1006	2984	431	83.2	268	82.3	169	1824	104	20.6
GR	296	787	151	31.4	44.9	13.5	26.5	487	24.7	8.01
HR	42.9	112	17.6	5.86	8.37	1.53	3.70	67.7	3.41	4.53
HU	122	320	55.9	13.5	21.3	5.13	16.9	197	5.62	4.49
IT	604	1815	284	30.5	141	54.7	70.8	1157	62.6	13.1
LG	52.7	141	40.8	5.35	10.8	1.77	4.54	73.8	2.50	1.67
LH	57.3	163	39.5	6.23	11.7	2.34	7.89	91.4	2.03	1.78
LO	37.6	126	27.7	3.12	6.81	1.37	1.94	78.7	1.95	4.42
LU	34.5	86.6	22.4	1.92	10.7	2.50	4.52	41.1	1.56	1.93
MT	106	256	88.7	3.87	15.5	5.21	7.28	128	3.85	4.13
NL	271	664	127	13.7	66.9	25.5	30.6	366	26.8	6.72
PL	232	697	114	20.4	39.2	11.3	35.9	452	16.5	6.91
PO	94.4	226	46.7	4.91	21.1	9.22	5.77	126	8.67	2.85
RO	122	360	110	8.75	19.6	4.13	12.9	196	4.51	4.48
SI	16.8	44.1	9.63	1.54	2.36	0.83	2.29	26.1	0.68	0.59
SP	425	1174	223	20.6	80.5	31.3	38.5	742	29.7	8.47
SW	194	566	84.3	12.8	65.3	15.9	35.5	325	22.9	4.32
UK	3540	11821	5647	118	737	283	285	4221	487	40.4

Table 5: Data statistics(x1K) of news collection of EU country-centric events over different geographic regions from Jan-Dec 2018. Countries are represented in FIPS-04 coding system. (No events are found for DK.)

to improve inter-rater reliability. The final mapping was established with high agreement, Cohen's Kappa κ=.74.

Notice that a news article may contain more than one categorical theme. Hence, same news might be part of different wiki topics (i.e., topic sets are overlapping in nature). For example, news about 'Imprisonment order of a French woman issued by Iraq court due to IS membership' is part of both '**armed conflicts**' and '**international relations**' topics. Since we only map categorical GDELT themes to their corresponding topic, we translate specific themes in the documents to the matching categorical theme (contained as a substring). However, some of the newly introduced themes such as 'WB_2167_PANDEMICS', 'UNGP_FORESTS_RIVERS_OCEANS' do not contain any general themes; hence, such themes are not mapped to any of the broad category of topics. We put such themes and corresponding news events under 'missing' category.

We perform a detailed analysis of EU-Total news across different regions to get an initial understanding about the variation in the representation of the Eu in different parts of the globe. We discuss our observations and some of the important straightforward implications of this analysis in the next part.

Experimental Results: As mentioned before, we want to examine the way the Eu is represented in other geographic areas. For this, we consider the whole set of Eu news (Table 6) from different regions, and we inspect the variation in their sentiment polarity distribution. We conduct a comparative analysis using the reported view of Eu in the internal media as the baseline.

Topic	EU	EUROPE(NON-EU)	ASIA	AFRICA	MIDDLE-EAST	AMERICA	OCEANIA
business and economy	823	33.2	186	64.9	70.5	881	68.3
health and medicine	881	12.5	95.4	46.1	50.6	809	80.1
international relations	582	48.8	172	46.3	157	1006	43
arts and culture	3121	268	1106	281	805	6791	457
armed conflicts	549	36	164	50.5	119	1317	106
law and crime	2197	79.5	420	155	243	2965	212
disasters and accidents	919	27	157	64.5	112	1477	108
politics and elections	3946	208	1052	367	602	6518	368
science and technology	681	10.8	92.9	39.5	48.3	744	63.7
environment	235	7.20	59.1	23.8	19.1	426	37.7
others	6693	265	1465	505	790	10057	785
missing	9087	443	2114	711	1133	13892	987
Total	29720	1440	7088	2356	4151	46888	3318

Table 6: Data statistics(x1K) of documents per topic over different geographic regions. For each topic we list EU-TOTALnews covered in different regions.

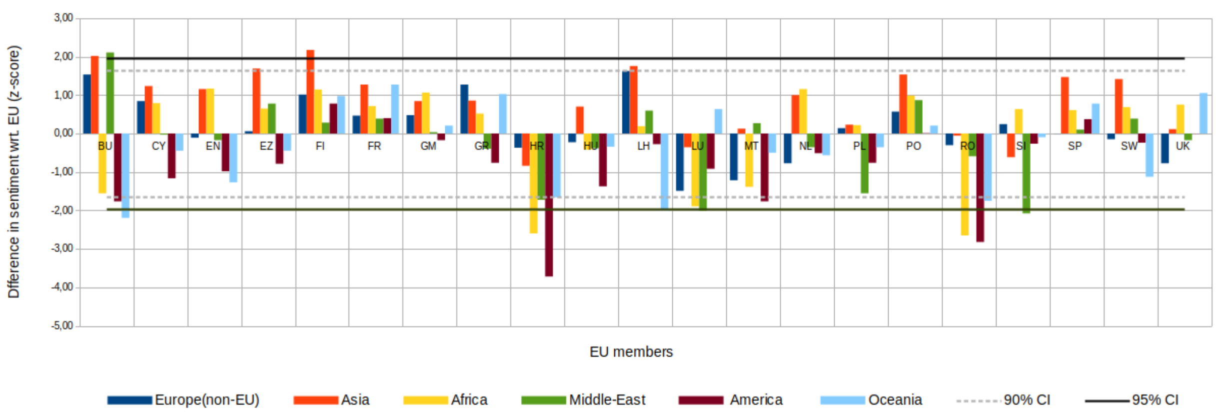

Figure 3: Divergence of each region's average tone over news in EU-TOTAL w.r.t. the EU coverage. News are aggregated for EU member states.

First, we compare the average sentiment/tone of the reported news from different regions towards each of the EU member states aggregated over the year 2018. Figure 3 represents the difference in the sentiment score of other regions compared to the EU for the EU-TOTAL dataset. To have a fair comparison among scores from different regions for the same EU country, we compute $z-score$. Figure 3 shows the standard score (z-score) representing the average sentiment towards each EU-member on the corresponding region's mean for EU-TOTAL.

By using the average sentiment of the coverage from EU as our baseline, we can evaluate to what extent the internal view of different country members of the EU might change when reported in other geographic regions. For example, we can see that the 'image' that is presented in the media from outside the EU about France (FR) is, on average, more favorable when compared to the EU's coverage of the same country. To have a fair comparison among scores from different regions for the same EU country, we are showing in Figure 3 the standard score (z-score) for each value on the corresponding region's mean. To illustrate, let's see the case of Finland (FI) for EU-TOTAL news. Although this country receives a more positive (or less negative) coverage from every other zone compared to the EU's, the graph shows that only ASIA gives significantly favorable coverage to Finland even when taking into account this regions' standards. For clarity, we are only showing in the graph EU countries for which at least one region is over one standard deviation above or below the region's mean.

To complement our analysis, we also measure the distribution of polarity values of EU-TOTAL news for all EU countries from each of the regions. This will give an intuition about the viewpoint of the different regions towards the EU. We also check the cumulative distribution frequency (CDF) of polarity values of the EU countries from each of the six geographic regions. AMERICA(96%), MIDDLE-EAST(90%), OCEANIA(80%), and EUROPE(NON-EU)(75%) itself present most of the EU countries in a relatively negative way. Meanwhile, ASIA represents more than 60% of the EU countries in a more positive way (including countries such as Bulgaria (BU), Czech Republic (EZ), Lithuania (LH), Portugal (PO), and Spain (SP)). Moreover, those countries covered with a comparatively more negative press in Asia are closer to the average sentiment used in the EU than in any other region. On the other hand, Africa follows a mixed trend; despite covering more than 40% of the EU countries in a relatively positive way, they do not deviate significantly in this direction from the European coverage. Yet, when we look at the negative side of the chart, we can see that African media give a heavily negative representation of some of the EU-countries (e.g., Romania (RO), Croatia (HR), Luxembourg (LU)). It is interesting to note that European countries (non-EU members) tend to present EU members from a more negative point of view compared to the EU. Although, we should also remark that this is the only region that does not deviate significantly in the negative nor positive direction from the EU coverage for any of the country members.

This leads to another interesting question: "**Do different regions favor EU countries in the same way?**" We rank the EU countries based on their sentiment scores for each of the regions and measure Spearman rank correlation between each pair of regions. The pairs <EU-EUROPE(NON-EU)>, <EU-ASIA>, <EUROPE(NON-EU)-ASIA>, <MIDDLE-EAST-ASIA>, <AFRICA-AMERICA>, <AFRICA-OCEANIA> and <AMERICA-OCEANIA> show strong positive correlation ($\rho(27) > .60, p < .001$) [Bujang and Baharum, 2016]. This indicates that there is affinity in how these regions promote individual EU-countries. The ranked coverage of EU shows only a relatively weaker correlation with all of the other regions. We assume that a region will prioritize the coverage of individual countries, or put them in a more positive light, based on their political and economic objectives. For example, some regions' sentiment-ranking of the EU-countries align with their business deals/foreign trade alliances, whereas some of them give preference to health issues and historical international relations.

So far, we analyzed the average sentiment values for each of the EU countries and other regions. We also perform a detailed statistical analysis over all the sentiment values collected from all the news articles of different regions. This region-wise analysis is performed over all the news reported in Table 5. We perform two-sample Welch's t-test [Welch's unequal variances t-test, 2020] to check the difference between regions in their calculated sentiment for the EU news coverage. For the EU-TOTAL dataset, most other regions' coverage of EU is significantly more negative ($\alpha = .001$) than the internal EU news. The only exception is Asia $(M = -1.85, SD = 3.5), t(3M) = -0.334, p = .73, d = 3.52$, which seems to represent very closely - although slightly more positive - the domestic sentiment of the EU media $(M = -1.81, SD = 3.5)$.

All these analyses were done over the whole set of news without considering their themes/topics. However, sentiment of news is highly correlated with the topic and topic has a huge influence in the representation of a news article [Frech, 2008]. Hence, we inspect the impact of topics on the news sentiment.

Table 7 reports the topic-wise difference in the EU-TOTAL news sentiment polarity between a region and EU. Notably, all the regions have a tendency to represent EU in a more negative way. This trend holds quite uniformly across the overall news as well as for topic-specific news. Only 'Science and Technology' gets some appraisal in the ASIA and AFRICA. The Figures show that AMERICA has an under-representation of positive news compared to the EU, paired with an over-representation of negative stories. These two regions are

Topic	ER	AS	AF	ME	AM	OC
business and economy	-0.44	-0.23	-0.04	-0.42	-0.71	-0.25
health and medicine	-1.57	-0.45	-0.09	-0.77	-0.92	0.03
international relations	-1.19	-1.46	-1.49	-1.37	-1.81	-1.19
arts and culture	-0.84	-0.27	-0.77	-0.81	-1.23	-0.50
armed conflicts	-1.05	-0.87	-1.14	-1.21	-1.06	-0.09
law and crime	-0.05	-0.25	-0.27	-0.46	-0.59	-0.35
disasters and accidents	-0.92	-0.83	-0.81	-0.61	-0.69	0.07
politics and elections	-0.35	-0.05	-0.09	-0.71	-0.76	-0.45
science and technology	-0.63	0.31	0.26	-0.76	-0.51	0.11
environment	-0.21	0.16	-0.31	-1.09	-0.69	-0.07
others	-0.42	-0.07	-0.14	-0.45	-0.70	0.03
Average	-0.53	-0.18	-0.31	-0.67	-0.85	-0.22

Table 7: Topic-wise sentiment polarity representation of EU-TOTAL events among different zones relative to EU.

important global actors in different aspects such as trading, health-services, education, etc. A possible reason for advancing a rather negative image of the EU zone might be such competing interests in the political and economic fields. It is also evident from Table 7 that the negative representation is highest for the topic 'international relations'. We also perform statistical Welch's t-test [Welch's unequal variances t-test, 2020] between the sentiment distribution of EU and each of the other regions. The distributions turn out to be significantly different from EU which indicates that each of these regions follows different editorial and word selection strategies to present the EU-news.

In this section, we make a thorough analysis of EU-TOTAL news and their editorial customization across different media channels all over the globe. It is quite evident that the image of EU is not uniform and significantly different than its home representation. In the future, we will focus on migration-related topics and their variations in representation among different global outlets.

3.2 Coverage of EU-COVERED news in other geographic regions

Most of the news reported by EU home-media get filtered in other regions. Hence, analysis of EU-TOTAL news is not enough. In this section, we extend our analysis to the EU-COVERED news.

3.2.1 Problem statement

According to [Chaban and Holland, 2013, Chaban and Holland, 2014], three main elements should be considered in the analysis of EU external perception. The first one is the inspection of EU imagery in the national (and regional) news media (e.g., press and television). The other two elements identified by the authors are the gaining of insights into the public's opinion and making an accurate assessment of the views of the national (and regional) DPMs. In this report, we aim to tackle the first element by developing an advanced and scalable computational methodology. Our approach leverages the automatic processing of large collections of documents, which also allows for spatial scaling.

We perform a detailed analysis of the representation of EU in the news across multiple geographical regions

(EUROPE(NON-EU), ASIA, AFRICA, MIDDLE-EAST, AMERICA, and OCEANIA) over one year (January-December 2018). For a deeper understanding of coverage, we also analyze the news, which are directly adopted from the EU news sources (i.e., overlapping of coverage with EU-internal media — EU-COVERED). Furthermore, we compare the coverage of EU-COVERED news across eleven broad topics in three dimensions, namely, volume, sentiment polarity, and editorial strategies.

3.2.2 State of the art

Most of the comparative studies done so far are based on the total news set. We mentioned some of them in Section 3.1.2. A complementary line of research to the bias analysis is the analysis of opinion-shaping or setting the agenda [McCombs and Shaw, 1972] through mass media. In [Pinto et al., 2015] a model capable of imposing the view of a news outlet to a large number of consumers is analysed. Recently, the difference between objective and perceived bias has been pointed out [Eberl, 2019]. Here, we focus only on the objective analysis of the media "coverage bias". However, the role of perceived bias in the communication of the EU image is an important aspect that we will pursue in future work.

In recent times, a few computational approaches were proposed to analyze the selection of political speeches, conventions, media highlights, propagation pattern of different news, effect of bias in opinion shaping etc. [Niculae et al., 2015, Tan et al., 2018, Romero et al., 2011, Abebe et al., 2018, Rotabi et al., 2017]. However, to the best of our knowledge, none of the prior studies focused on a detailed computational analysis of the image of the EU in EU and other regional media across the globe. In this section, we perform a detailed computational analysis of EU-COVERED news and report some of its potential implications.

3.2.3 MIRROR approach

Another important aspect of EU-news is the overlap in the selection of events by different regions. As mentioned in Section 3.1, the same piece of news may be reported in different news outlets of another region. For example, the news about '**resignation of Brexit secretary David Davis**' was published by *India TV News*[11] and *vnewsbd.com*[12] articles of ASIA. In general, not all the news reported in other zones are also published in EU and vice-versa. Interestingly, other regions tend to cover many EU-related events that are absent in the media outlets of the EU. That is why we also analyze EU-COVERED news to get a deeper understanding of the representation of EU-news.

For EU-COVERED, we consider events ('GLOBALEVENTID') covered by media outlets of the EU-zone and select only those event-related news from media outlets of other zones. Table 8 shows the detailed topic-wise statistics about EU-COVERED news. As expected, a majority of the European news gets filtered and does not appear in the news outlets of other zones. In their role as gatekeepers, journalists, editors, and other involved parties have to decide and pick from a massive pool of stories what to cover. Their selection is constrained by a combination of organizational factors, news norms, and audience interests [Shoemaker et al., 2001, Soroka, 2012]. On the other side, what they do select to report could be very telling on their editorial strategy and underlying forces shaping their framing of real-world events [Soroka, 2012]. From the **Total** row in Table 6 and Table 8, we see that the proportion of EU-COVERED

[11] https://www.indiatvnews.com/news/world-brexit-secretary-david-davis-resigns-over-policy-differences-with-uk-pm-theresa-may-451790
[12] http://vnewsbd.com/2018/07/09/brexit-secretary-david-davis-resigns/

Topic	EU	EUROPE(NON-EU)	ASIA	AFRICA	MIDDLE-EAST	AMERICA	OCEANIA
business and economy	823	5.05	36.2	13.7	12.1	239	14.7
health and medicine	881	3.25	28.2	13.7	9.03	247	22.5
international relations	582	8.98	45.9	16.4	26.2	283	10.7
arts and culture	3121	42.2	251	93	117	1942	102
armed conflicts	549	6.81	46.1	19.4	24.1	370	17.4
law and crime	2197	15.3	115	48.8	51.1	911	50.7
disasters and accidents	919	6.01	44.1	23.6	21.8	385	24.5
politics and elections	3946	37.6	247	99.5	104	1893	90.7
science and technology	681	1.65	17.4	7.18	5.48	161	12.9
environment	235	1.49	11.2	6.52	3.66	105	7.51
others	6693	47.9	358	143	140	2970	195
missing	9087	75.1	498	200	201	3975	232
Total	29720	251	1700	686	716	13486	781

Table 8: Data statistics(x1K) of documents per topic over different geographic regions. For each topic we list the number of news items from other zones which are also covered in the media of EU(EU-COVERED).

to EU-TOTAL for each region varies from 17% (EUROPE(NON-EU) and MIDDLE-EAST) to 29% (AFRICA and AMERICA). *Note that, these values (specially the EU-TOTAL news) differ from the ones reported in Table 5 because most of the news contain multiple topics and same news might be part of different topics.* However, the ratios remain quite consistent, even across the topics. As mentioned above, this might be due to multiple factors such as the different editorial policies followed by the news outlets of different areas based on their target audiences.

3.2.4 Experiments, datasets, results and future work

Dataset: Table 8 shows the distribution of EU-COVERED news.

Experimental Results: As mentioned before, we want to examine the way the EU is represented in other geographic areas. For this, we consider the whole set of EU news (Table 8) from different regions, and we inspect the variation in their sentiment polarity distribution. We conduct a comparative analysis using the reported view of EU in the internal media as the baseline.

To have a fair comparison among scores from different regions for the same EU country, we compute $z-score$. Figure 4 shows the standard score (z-score) representing the average sentiment towards each EU-member on the corresponding region's mean for EU-COVERED news.

For EU-COVERED, EUROPE(NON-EU), MIDDLE-EAST, AMERICA,and OCEANIA follow a quite similar trend to EU-TOTAL news. However, it is interesting to note that there is a difference in the pattern for ASIA and AFRICA. For most of the countries such as 'Estonia (EN)', 'Lithuania (LH)', and 'Portugal (PO)' etc. they reduced their positive tone and in some cases ('Malta (MT)') they took a reverse standpoint.

In an attempt to explain this behavior, let us first divide the non-EU Countries into former west and east block states. From a social and economic point of view, the former west block countries are closer in nature and standard to the EU, e.g., Norway and Switzerland. These countries are typically in a tight relationship with the EU. Without clear benefits of possible EU membership, it is reasonable to prefer and defend the status quo also by depicting a more negative image of the EU. The former east block countries, on the other hand,

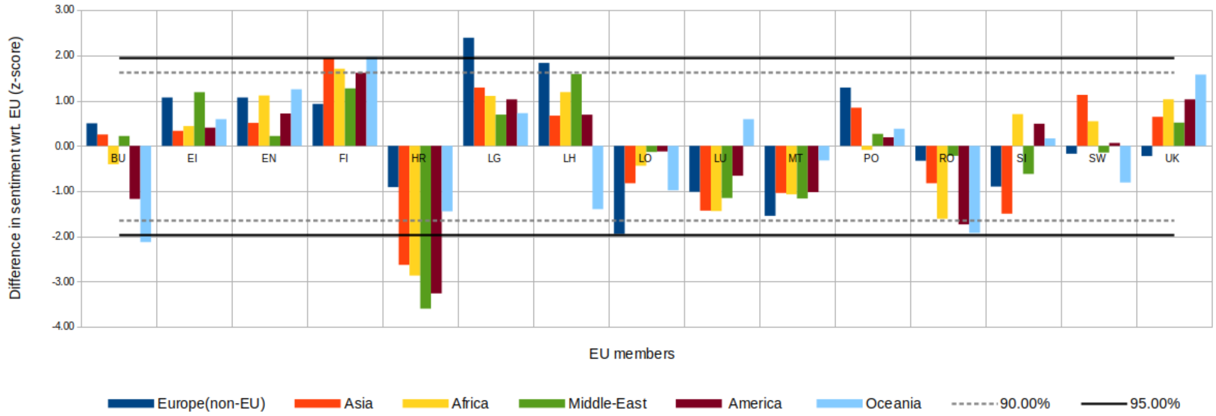

Figure 4: Divergence of each region's average tone over news in EU-COVERED w.r.t. the EU coverage. News are aggregated for EU member states.

still encounter tremendous economic difficulties and geopolitical interests and influence of third parties, most notably Russia (Armenia, Georgia, Moldova, Serbia) and US(Albania, Bosnia, Macedonia, Montenegro, Kosovo). Country-Country conflicts also cannot be excluded, especially when it comes to states established after the break-up of Yugoslavia and the war on the West Balkan. Last but not least, the negative opinion of some EU members on non-EU European countries, may lead to the reciprocal negative presentation of the whole European Union in the local media, e.g., Turkey, Ukraine.

We also analyze the question **"Do different regions favor EU countries in the same way?"** for EU-COVERED news. <AMERICA> shows strong positive correlation ($\rho(27) > .60, p < .001$) [Bujang and Baharum, 2016] with all the other regions except EU. We observe several new strongly correlated pairs for EU-COVERED news. On the other hand, this correlation effect gets lost in some of the cases such as <EU-ASIA>. Similar to EU-TOTAL news, we also perform a two-sample Welch's t-test [Welch's unequal variances t-test, 2020] to check the difference between regions in their calculated sentiment for the EU news coverage. All the regions cover EU in significantly more negative way ($\alpha = .001$) than the internal EU news.

Next, we perform a topic-level sentiment analysis to capture the view-point of different regions from various aspects (business or trade, international relations etc). We also perform statistical Welch's t-test [Welch's unequal variances t-test, 2020] between the sentiment distribution of EU-COVERED news between EU and each of the other regions. The distributions turn out to be significantly different from EU and this highlights different editorial policies of other regions.

Table 9 reports another interesting trend between EU-TOTAL and EU-COVERED news. It reveals that the EU-related news that are skipped by EU-media are covered by most of the other zones in a positive way. A detailed topic-wise investigation of such news shows several interesting patterns as follows:

1. Most of the missing news (EU-TOTAL-EU-COVERED) are from 'arts and culture', 'politics and elections', and 'law and crime'. Almost all the regions follow quite similar trend; hence, the Spearman correlation between different regions based on the topic-wise coverage of EU-MISSING news is significantly high (> 0.60).

2. Distribution of sentiment values of EU-MISSING and EU-COVERED news is significantly different as per Welch t-test [Welch's unequal variances t-test, 2020]. The news not covered by EU-media are represented in more positive way than their counterparts.

Topic	ER	AS	AF	ME	AM	OC
business and economy	-1.06	-1.02	-1.37	-1.34	-1.23	-0.31
health and medicine	-1.94	-1.26	-0.92	-1.99	-1.25	0.25
international relations	-1.44	-1.81	-2.10	-1.89	-1.71	-1.10
arts and culture	-1.59	-1.45	-1.67	-1.71	-1.64	-0.69
armed conflicts	-1.28	-1.19	-1.02	-1.57	-1.35	-0.63
law and crime	-0.93	-1.01	-1.17	-1.29	-1.17	-0.67
disasters and accidents	-1.19	-1.37	-1.31	-1.47	-0.98	-0.06
politics and elections	-0.95	-0.89	-1.16	-1.37	-0.96	-0.48
science and technology	-1.09	-0.67	-0.83	-1.68	-0.84	-0.11
environment	-0.95	-0.97	-1.68	-1.79	-1.19	-0.53
others	-1.18	-1.04	-1.23	-1.45	-1.08	-0.01
Average	-1.23	-1.15	-1.35	-1.56	-1.23	-0.32

Table 9: Topic-wise sentiment polarity representation of EU-COVERED events among different zones relative to EU.

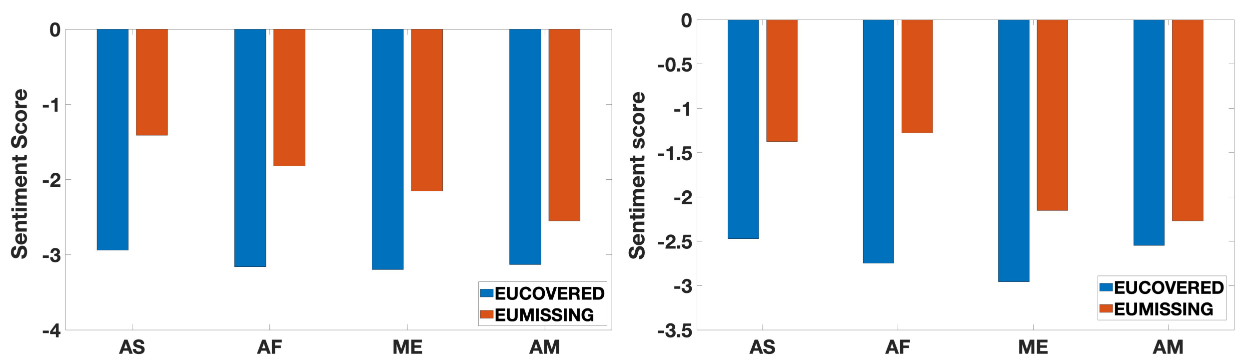

(a) **Arts** (b) **Politics**

Figure 5: Mean polarity value of EU-COVERED and EU-MISSING news for arts and politics across ASIA(AS), AFRICA(AF), MIDDLE-EAST(ME), and AMERICA(AM). EU-MISSING news are presented in a positive than the EU-COVERED news.

3. Figure 5 shows the distribution of news polarity for the EU-MISSING and EU-COVERED news over four different regions for topics 'Arts and Culture', and 'Politics and Elections'. It is quite evident from the figure that self reported EU-related news of different regions get more positive response than the EU-COVERED news.

4. Most of the news not selected/covered by the media outlets of other zones are either local to a EU country i.e., do not have global importance or controversial in nature.

5. We also observe a big difference in the polarity for the EU-COVERED news between EU and other zones. For example, Table 10 shows some topic-wise examples where same news is covered with different sentiments in different zones.

6. The distribution of news channels involved in broadcasting EU-MISSING news follows a power-law distribution i.e., some specific news channels from each region are primarily involved in circulating such EU-related news that are not covered in EU-media.

Topic	News	Zone	Polarity Difference
politics and elections	Brexit Deal of Theresa May	Oceania	2.3
armed conflicts	US France call for action against Syria chemical attacks	Asia	3.75
business and economy	Riots in Paris	Oceania	-12.3

Table 10: Examples of big difference in the polarity for the EU-Covered news between Eu and other zones.

The present deliverable also has some limitations as follows:

1. We consider the news posted in both English and regional news outlets of different regions and the non-English contents are translated into English for the analysis purpose. However, we did not check the differences in the news coverage between the regional and English outlets of a specific region. English outlets are mostly used to reach global audiences, not local people. Local audiences usually consume news posted in their native languages. Hence, it is essential to analyze such representation gaps, if any.

2. All the results reported in this deliverable are based on the GDELT dataset. We consider all the sentiment values, polarities as golden truth. However, the sentiment measures are based on lexicons. Hence, some noise is involved in this analysis. In future, we have a plan to extend this analysis to other channels and use deep-learning based sentiment detection tools (developed in WP4) to reduce the noise.

3. We consider eleven broad categories of topics in this report. However, some of the sub-topics such as 'Human Trafficking', 'Immigration' are of great concern and this news representation and its perception might have a great impact on such sub-topics. A detailed and thorough understanding of such sub-topics is required to exploit these influences.

4. We just explore the differences in the news representation across different regions. However, it will be interesting to check the differences in the perceptions generated due to such representations.

5. We consider Asia, Africa, and Americaas a single unit. However, these regions cover large geographical area and they might follow different policies. For example, Central-Asia and South-East-Asia might present different set of news to their audiences. In the future, we will explore sub-regional analysis.

6. We will specifically investigate and compare regions with COOs, COTs, and CODs.

Overall, we perform a detailed analysis to check the view of Eu in the media of other geographical regions and their differences with the internal reporting inside Eu. We were able to observe how Eu is covered differently in the news of other areas. We found that the sentiment polarity for Eu-topics differs from that of Eu-media and that only a small subset of the Eu-related news is directly adopted. As reported earlier, we performed a detailed analysis of **representation** of Eu-related news. We further extended this analysis over different topics and sub-regions of Asia, Africa, and America. The results of this deliverable are an important first step towards better understanding the **perception** of Eu among the audiences across the globe. An open interview with the audiences might be helpful in exploiting the correlation between **representation-perception**. It is also interesting to check their influence on migration (possibly by establishing correlations between different editorial strategies and migration values as reported by official sources). By understanding the differences in editorial strategies in other regions and their impact on migration, we can better process the content from other media types (e.g., social network) and adjust our models to account for these differences. By doing this, we reduce possible bias introduced to our models by otherwise out-of-context information.

Digital availability of news from around the world allows the analysis of perception at a large and global scale. We have seen how the narrative (through selection of events and sentiment polarity) of EU-related issues may change when reported from outside its borders. Quickly identifying subjects of greater interest and potential threats by misconceptions will empower decision and policy-makers to better understand the global context and react accordingly in a timely manner.

4 Evolution of Networks and Communities

The main goal of this task deals with the evolution of migrants' networks and their communities over time. We aim at the development of novel methods to improve the modeling performance, given the heterogeneity of the cross-media networks constructed in the previous task. The methods being developed incorporate and leverage result and findings of tasks 4.3 and 4.4 of WP4 and tasks 5.2 and 5.3 of WP5 so that to better profile the communities using multimodal data. The output of this task will includes a wide range of statistical observations on (i) the networks' growth and shrinkage, (ii) the communities' formation and characteristics, and (iii) interactions among the communities. These will give insights into how migrants connect and interact and the factors which affect their behavior and perception.

In our first year, we have focused on the identification and characterization of several relevant communities for the study of migrants perception. The next steps involve the development of methods to improve our understanding of how these identified communities evolve over time or react to specific events.

4.1 Identification and characterization of Twitter communities around refugee and migration topics

4.1.1 Problem Statement

As refugee and migration have been controversial topics, especially in the last five years, they have attracted much interest and discussion worldwide. In this context, Twitter has been heavily used for seeking and sharing information toward the topics [Aswad and Menezes, 2018, Hadgu et al., 2016]. People have also used Twitter to express their emotion and opinion on the topics [Kreis, 2017, Siapera et al., 2018]. Through these interactions, Twitter users may form communities whose members share common viewpoints or stances toward aspects of the topics. The understanding of these communities' structure and their intention would be highly useful for a wide range of important applications, e.g., social sensing [Gualda and Rebollo, 2016] and social surveying [Righi, 2019]. This has however not been well studied in previous work.

In this work, we therefore would like to uncover the communities among Twitter users who are more engaged with refugee and migration - related topics. We also want to examine the interactions among the communities and their evolution over time. Furthermore, we are interested in identifying communities among the refugees and migrants and how they interact with the communities in the host countries. Insights from these interactions would bring us a better understanding of the refugees' and migrants' perception of and how they integrate with the host countries [Lamanna et al., 2018].

4.1.2 Related Work

There is a rich literature on analyzing communities on social networks and, particularly, in online social networks [Tang and Liu, 2010, Fang et al., 2020]. To the best of our knowledge, there is however no such prior work in the refugee and migration domain. Most works on social media in this domain focus on either (i) the usage of social media for discussing topics of the domain [Aswad and Menezes, 2018, Siapera et al., 2018], (i) analyzing the discussed topics [Siapera et al., 2018, Nerghes and Lee, 2019], and the overall sentiment and emotion expressed toward the topics [Gualda and Rebollo, 2016, Kreis, 2017, Nerghes and Lee, 2018].

There are also previous works on understanding how immigrants and refugees make use of social media extensively to facilitate their migration [Dekker and Engbersen, 2014, Dekker et al., 2018], communicate and exchange information about the integration process [Alencar, 2018], as well as to strengthen bonds within their communities in host countries and maintain their relationships in original countries [Komito, 2011, Dekker et al., 2016]. These works were however conducted on small datasets and/or based on recruited migrants. In summary, most existing work on this domain are human-subject based studies or content analysis - based studies. There is a lack of comprehensive analysis on both social network and content perspectives. This may be due to the challenges in collecting and cleaning data from social media at large scale.

4.1.3 MIRROR approach

In this work, we would like to fill the gap in the literature on social media studies as discussed above by first addressing the absence of comprehensive datasets that contain both social networks among Twitter users who are interested in refugee and migration topics as well as their content and interaction regarding the topics. These collected datasets would enable us to conduct extensive empirical analysis from multiple perspectives, including social communities, topics, and affective factors. To do so, we employ a snowball sampling approach for collecting the datasets. Precisely, starting from an initial set of Twitter users who are well known to be influential in the domain (e.g., *@Refugees*, *@RefugeesOlympic* and *@TeamRefugees*), we first manually select other influential users in several EU and Middle East countries. These manual selection steps are performed iteratively by scanning through the followee lists of the previously selected users. This is a reasonable strategy for improving the recall in collecting highly influential users in the domain as influential users often follow other influential ones. Next, we collect all the followers of these selected influential users. We focus only on the followers who follow at least three selected influential users. This minimum threshold is to ensure that most the users in our dataset actually engage with the domain. Lastly, we collect all the following links among the users as well as all their tweets and interactions.

4.1.4 Datasets and future work

Leveraging the approach presented above we have collected a large dataset. In the following we briefly describe the dataset.

- Influential users: We have manually selected more than 600 influential users. Most of them are from EU countries, others are from Turkey, Middle East and North African countries. The typical selected influential users are:
  - Activists and organizations that provide information, help, service, and education for immigrants
  - Agencies that aggregate and broadcast news about refugee and migration topics
  - Institutions, forums, and conferences that are dedicated to the refugee and migration topics
- The followers: In total, the influential users have more than 3M unique followers. Among those followers, more than 290K users follow at least 3 selected influential users. We have collected all the following links among those 290K+ users as well.
- The tweets: For each of the selected influential users and selected followers, we make use of Twitter's APIs to collect the user' tweets. Due to the APIs' restriction, we can only collect the last 3200 tweets

of the users. This number actually covers the whole historical tweets of most the users. Only for a very small proportion of the users who tweet frequently have more than 3200 tweets. Nevertheless, for almost all users, we can collect all their tweets in the last 5 years. That is, our collected dataset fully covers the time span of the European migrant crisis, which is the most interesting time span of the domain.

Our ongoing work includes filtering and cleaning the collected dataset. Specifically, we would like to extract from the dataset tweets that are relevant to refugee and migration topics. This extraction step is necessary, even though our collected dataset is focused on the topics, as Twitter content is highly noisy [Balasubramanyan and Kolcz, 2013] and covers a wide range of covered topics in Twitter [Yang et al., 2014b]. Given the huge size of the collected dataset, we determine to employ a keyword - based approach for this step. That is, we will manually compile a list of keywords for the topics, and will consider a tweet relevant to the topics if it contains some key words.

Our next step would be examining the aspects and sub-topics discussed in the extracted tweets. These include analyzing and visualizing the aspects' and sub-topics' temporal dynamic, across countries and languages (e.g., English, German, and Arabian). This step will be performed in conjunction with WP4's Task 4.4 on topic detection and evolution.

Ultimately, our objective is in discovering and characterizing the communities of users in the collected dataset. To do this, we will also examine and investigate different methods for community detection in social media, particularly efficient methods for dealing with the large size of the dataset. For characterizing the detected communities, we will examine the aspects and sub-topics discussed in their users' tweets, and the sentiment expressed toward the aspects and sub-topics. Lastly„ we will employ the snapshot - based approach for highlighting the temporal dynamic of the communities overtime.

4.2 Silent Twitter users modeling for opinion prediction

4.2.1 Problem statement

As discussed in the previous section, Twitter has long been the focus of great interest as a lens through which to observe large swaths of society. However, such analyses have relied exclusively on the side of active participation and content production. It is known that a large percentage of the social media users rarely tweet - according to a Pew Research report [Wojcik and Hughes, 2019], in the USA, 80% of the tweets come from only 10% of the users. Therefore, content-based inferences may fail to generalize to real-world populations, or even to less active users. Studies on mass media information dissemination are affected by these generalization issues as they increasingly rely on social media as a proxy for audience reach [Elejalde, Erick et al., 2019]. The power of social networks to shape the informational landscape of a population and the ever more important role of Twitter as a source for news consumption make this problem very relevant in contemporary society, and in particular for the accurate analysis of migrants' perception.

Although the silent users (also called lurkers) rarely post on social media, they still have opinions towards different topics. If we don't take their hidden opinions into account, the predictions could be far from the truth. For example, a lurker might never express their political preferences on the platforms, they just read and retweet some tweets posted by some other influential user. During a decision-making process, the probability that they are inclined to favor the opinion of the followed influencer is much higher than that of changing to another set of beliefs. By looking into silent users' actions (e.g., retweets, or network connections), we can also capture the preferences of this important group.

Our primary goal in this study is to identify a methodology that predicts the opinions of the users who are “silent” in social networks like Twitter. We define “silent users” (following related bibliography) as those accounts that never or seldom express their opinions explicitly on a given topic. In some cases, they may retweet, like, or mention (@) other more active users. As mentioned before, these “silent” users are a significant part of most social networks. So, determining their real opinions or interest is very important to estimate the external perception of the EU. There are already some previous works related to opinion mining. Here we aim to adapt the methods that represent state of the art in these fields and further improve their performance in the context of silent users.

4.2.2 State of the art

As mentioned before, given the low level of content generation that characterizes silent users, we need to appeal to other features that will allow us to infer their opinion in a given matter accurately. Previous works have suggested different combinations of features that can help us in this task. One group of these features are those based on the topology of the underlying network. They work under the principle of homophily [McPherson et al., 2001], where we assume that social entities will associate with similar others. In [Tan et al., 2011], the authors use the ‘following’ relation for the user’s preference prediction. They experiment with both unidirectional relations (e.g., v_i follows v_j, but v_j may or may not follow v_i in return) and bidirectional relations (e.g., v_i follows v_j and vice versa). With that same aim, other studies have also used second-order co-following [Garimella and Weber, 2014]. In this case, two accounts that do not share a single follower can still be considered similar if their followers share many friends. Authors in [Volkova et al., 2014] divide features in three categories: Social Graph, that uses different social circles between Twitter users (e.g., follower, friend, mention, etc.); Candidate-Centric Graph, that uses following relationships between the users and different political candidates; and Geo-Centric Graph, that collect users with different political preferences (as self-reported in their biographies) from various regions. Other useful relations between users can be extracted from Twitter elements such as hashtags, replies, retweets, and the already cited mention (@) of other users. These also provide valuable insights for the detection of communities [Tan et al., 2011, Volkova et al., 2014, Yang et al., 2013], which might shed some light on the characteristics of member accounts. But, these mentioned elements are also extracted from Tweets’ content and other social interactions. In our specific domain, silent users might tend to have low participation in the debate not only in terms of original content but also in a limited number of other kinds of interactions. For example, the number of followers is usually positively correlated with the influential power of the Twitter user [Riquelme and González-Cantergiani, 2016]

Another relevant group of features is based on latent relations between users (those that are not directly observable). In [Yang et al., 2014a], the authors complement topology information with prior information constraints. Their experiments show and improve the performance of community detection algorithms when these hidden features are added to the analysis. Latent relations (sometimes also called pseudo-relations) are usually extracted from content-based analysis. However, in our case, content-based data for the targeted users is either non-existing or insufficient at best. In [Wang et al., 2019], the authors pull pseudo-friend/foe relationships defined by meta-paths from the graph of topological relations. They go on to use a combination of both observed and hidden features for the user’s preference prediction through a coupled sparse matrix factorization (CSMF) model. [Gong et al., 2015] takes a similar multi-modal features approach to preference prediction. The authors profile lurkers by complementing direct links information with other inferred features such as marital status, religion, and political orientation. They use followees’ tweets to predict marital status, followers and friends’ tweets for identifying religion, and followees’ tweets for the political class. They can achieve performance in silent users that are similar to the profiling of active users.

Based on the previously described set of relations (both topological and latent), there have been multiple efforts on modeling either the general problem of users' opinions prediction or relevant specific subcomponents such as community detection or sentiment analysis. In [Wang et al., 2019], the authors propose a Coupled Sparse Matrix Factorization (CSMF) Model for silent users' opinion prediction. With the help of a community-topic opinion matrix and the users' feature matrix, less active users' missing values in the user-topic opinion matrix are predicted (i.e., the attitude of silent users towards a given topic). By coupling these two matrices, the authors overcome the sparsity of the information on each component. They evaluate their results in terms of root mean square error (RMSE) and accuracy.

As mentioned above, community detection is another crucial element in our analysis. When a user does not personally express their opinion, a reasonable assumption is to associate this person with the predominant point of view of those closest to them. We already covered several relevant related works for community detection in the previous section (see Section 4.1). However, given the scarcity of work in our domain (i.e., migration-related topics), we review some other relevant approaches that we intend to use as base-line for our own models.

In [Yang and Leskovec, 2013], the authors present an overlapping community detection method that scales to large networks of millions of nodes and edges. They also use a non-negative matrix factorization approach that they evaluate using Ω-Index, F1 score, and recall. Another relevant study that could be considered as an extension of the work by [Yang and Leskovec, 2013] is introduced in [Yang et al., 2013]. The algorithm aims at a scalable community detection from edge structure and nodes attributes. It statistically models the interaction between the network structure and the node attributes. Similarly, in [Yang et al., 2014a], the node property representation, which is used for clustering, is determined by both topology information and pairwise prior information. The authors propose a unified graph regularized semi-supervised framework which can make use of the prior knowledge to improve the performance of community detection

In the field of sentiment analysis, the literature is vast. We are interested in those approaches specifically designed for social networks and microblog platforms like Twitter. Usually, these methods incorporate local characteristics of the systems, such as the features that we also collect for our study. For example, in [Tan et al., 2011], the authors proposed a novel approach that incorporates both textual and social-network information in a single heterogeneous graph on a given topic, where nodes can correspond to either users or tweets. Then, they create a factor-graph model for user labels and employ transductive learning algorithms [Joachims, 1999] in the model to predict the labels of all users in the graph. MacroF1 and accuracy are used for evaluation. Alternatively, in [Li et al., 2009], a term-document matrix X is approximated by three factors: a non-negative matrix F representing knowledge in the word space, a non-negative matrix S representing knowledge in the document space, and a non-negative matrix G providing a condensed view. The study introduces a semi-supervised matrix factorization with lexical knowledge (SSMFLK) model, which can be implemented using simple update rules for the matrices F, S, and G.

Given the dynamic nature of the social networks, streaming models for personal analytics that dynamically update user labels based on their stream of communications have been proposed [Volkova et al., 2014]. Such models better capture the real-time nature of the evidence being used in latent author attributes prediction tasks. We will incorporate some of these ideas in our analysis.

4.2.3 MIRROR approach

In MIRROR we define a silent user as an account that ranks in the bottom 10% of original content generation in our collection of users (still, they may retweet other users' tweets). We will try different community detection algorithms as described in the Section 4.1 above. For example, we will implement the Communities from Edge Structure and Node Attributes (CESNA) model [Yang and Leskovec, 2013] based on the matrix factorization and users' attributes. For network generation, only follower/followee info is not enough, so we will add several other relations (both topological and latent) and measure the relevance of each feature in the accuracy of the final model. Different graphs may be suitable for different purposes. For opinion prediction, at first we will use an improved version of the Coupled Sparse Matrix Factorization (CSMF) Model [Wang et al., 2019].

4.2.4 Datasets and future work

Our initial step was the collection of a test-dataset that could be used for the evaluation of the proposed models. We started from a database of news outlets with an active Twitter presence and geolocated in the same city. Using the Twitter API, we collect the followers of these outlets. Based on their profiles, we restrict the results to followers also geolocated (with relatively high confidence [Elejalde et al., 2019]) in the same city as the news outlets. Furthermore, we focus our analysis on the subset of accounts that follow at least three different newspapers simultaneously. Our final data set counts approximately 200.000 users. We then proceed to the collection of 6-months worth of tweets for each of the accounts in our set (including the news outlets).

News accounts are relatively active and popular in social networks. Moreover, different news outlets usually cover the same events from different points of view. The most controversial topic/events often generate opposing sides in the social discussion, and, in most cases, the popular camps are also represented in the media landscape. Also, we know from previous studies that people frequently restrict their news consumption to channels that amplify or reinforce their own beliefs (this phenomenon is known as an echo chamber) [Mutz, 2006, Colleoni et al., 2014]. By focusing our analysis around people that consume their news from these media systems, we have a better chance to infer their real opinions accurately. By creating redundancy of at least three sources being followed, we can more precisely estimate the correct point of view being represented in the intersection of these. Furthermore, by selecting users in the same location, we minimize possible bias introduced by geographic, cultural, and social factors. After all, the perceived direction and magnitude in which a news outlet leans concerning a given controversy is relative to the context of the observer [Elejalde et al., 2018].

Our next steps in the analysis are:

- Continue the collection of the tweets of the 200K users so that we can extract the latent features and other content-based relations like hashtags, likes and mentions. The tweets are being collected for the time range between 2019-09-01 and 2020-04-23.
- Implemented the model 'Coupled Sparse Matrix Factorization' (CSMF) from [Wang et al., 2019] as a state of the art benchmark.
- Start with sentiment analysis and community detection analysis.

5 Information Diffusion and Manipulation

This section focuses on how information is diffused and manipulated within and across migrants' networks. To study the diffusion of information on different topics and domains at network level, the rate and phases of the diffusion are examined and studied: how these deviate with regard to the networks, the information types, and the origins of the information. At individual level, users playing different roles in the diffusion process are identified. The roles include trendsetters, influencers, propagators, early adopters, etc. During the first project year the use of bots for collecting, (re-)producing, and propagating information, has been studied. Of particular interest is to be able to distinguish between different kinds of bots in general, and to detect migration bots in particular. During the forthcoming project years, methods to identify information manipulation by the use of duplicated, clone, or fake accounts on social media, will be developed. The results of these analyses help to establish better sensing and monitoring of information diffused within migrant communities, thereby providing efficient methods for detecting and preventing misinformation.

5.1 Detection and analysis of bot communication

5.1.1 Problem statement

Internet robots, or simply bots, can be used to efficiently communicate a message and quickly reach out to a large audience, and today social media platforms are flooded with bot-generated content tailored to affect people's viewpoints and opinions in one way or the other. Such content is produced automatically or semi-automatically by software applications for the purpose of shaping perceptions. Less harmless applications include, for example, advertisement of products, whilst applications that are more troublesome include, for example, interference in political elections or creation of biased perceptions of Europe related to migration.

5.1.2 State of the art

In recent years, the area of bot detection has consequently emerged as a popular and important research area. As an example, during the Swedish general election in 2018, the Swedish Twittersphere was monitored in an effort to try to detect malicious activities during the election related to discussions spurred by bots. During this project, a bot detection model was developed [Fernquist et al., 2018]. However, during the related analysis work, two critical flaws in existing bot detection methods were identified: (i) today's detection models focus entirely on the bot or not bot question, and (ii) current methods do not provide insight that can be used to guide further analysis and take appropriate action.

The focus on whether a social media account is a bot or not creates an ambiguous situation since there is no universal definition of what a bot is: what some people refers to as a bot might not be considered a bot for other people. Some people might say that only software-powered accounts such as spambots or chatbots ought to be considered bots, and that accounts controlled by persons do not qualify as bots even though they behave exactly the same way as a software bot. Others might say that it does not matter whether there is a computer program or a human behind the account, but that accounts exhibiting automatic behavior (such as sockpuppets and trolls) ought to be considered bots.

The lack of explainability in today's bot detection methods makes the bot detection result difficult to use. Existing methods often only use the username of an account as input, and provides the output as a percentage number indicating the likelihood or extent of the account being a bot. Examples of services using such methods are Twitter Audit [twi,], Bot Sentinel [bot,], and Botometer (earlier called BotOrNot) [Davis et al., 2016]. All the mentioned services works similarly: the user can select a Twitter account to see whether the account is likely to be a bot or a genuine account. From a user perspective, you get no indications on what features or attributes that have been taken into consideration for drawing the conclusion whether the account is a bot or not. During evaluations of established bot detection methods it has been shown that obviously genuine accounts operated by real people are often classified as bots, but without an explanation of why the accounts might be bots it is hard to analyze and change the methods used [Kreil, 2018].

5.1.3 MIRROR approach

In the MIRROR project, the above-described challenges are used as a starting point for developing a detection method that (i) distinguishes between different kinds of bots, and (ii) provides an explanation of why an account has been classified as bot or not bot. The approach is based on easy-to-explain attributes, where each attribute measures something that is (almost) intuitive for the user to understand rather than the complex features that are typically used in machine learning models. The idea with this approach is that different combinations of intervals of the attributes should then correspond to different kinds of accounts.

5.1.4 Experiments, datasets, results and future work

A first set of initial attributes has been developed, and some experiments applying the attributes on different account types have been conducted [Fernquist et al., 2019]. These preliminary results suggest that it is possible to distinguish between different account types using only a few different attributes by using principal component analysis (PCA) to project the accounts and their attributes in two dimensions. The next step will be to further investigate different kinds of accounts and in parallel develop the attributes further to be able to distinguish between different kinds of bots (spambots, pornbots, etc.), and make use of this knowledge to be able to single out migration bots.

A number of reference accounts representing different kinds of bot behavior provides a baseline for the work. Different types of accounts will be annotated and further desired types need to be found and harvested. Small-scale annotation of account types such as politicians, celebrities, comedians, and news sites has already been conducted. In addition, there already exist datasets with labeled accounts, which can be used as reference accounts. Examples of such datasets include accounts annotated as verified accounts [Yang et al., 2019b], purchased fake followers, celebrities, and bots advertising scam sites [Yang et al., 2019a], as well as completely manually annotated accounts [Gilani et al., 2017, Varol et al., 2017, Cresci et al., 2017]. After listing a first set of different account types, the attributes will be developed further to account for the bots' distinguishing features. When appropriate attributes have been created and labeled accounts for different account types have been gathered, experiments to evaluate the attributes will be made. Machine learning classifiers and clustering algorithms will be used to investigate the performance of the approach. Finally, the developed attributes and the gained insight will be used to construct the combination of attribute intervals best corresponding to migration and to finding migration bots. The developed migration bot classifier will be evaluated both in terms of classification accuracy and in terms of explainability.

6 Conclusions

This report presents the first release of Cross-media Social Network Analysis Technologies. It describes the methodologies and technologies used for cross media social network construction and contextualizes the CNC component within the MIRROR architecture and in connection with the Information Model of the system.

We perform a detailed analysis of EU-related news across different regions to get an initial understanding about the variation in the representation of the EU in different parts of the globe. This analysis provides a first essential step in the automatic characterization of the perception of Europe from outside its borders. We have seen how the narrative (through selection of events and sentiment polarity) of EU-related issues may change when reported from outside its borders.

Further, we addressed one critical aspect for the network analysis, namely, the identification and characterization of several relevant communities for the study of migrants perception. We describe the approaches to be followed in MIRROR and the large datasets that are being collected to validate the developed models.

In addition, for further inspecting and understanding manipulation of information, the use of bots for propagating migration related information (and misinformation) has been studied, with a view to developing a novel bot detection methodology specifically tailored to detection of automated communication related to migration.

Following, we will continue with the implementation of the CNC component. Methods developed for the multi-modal analysis of the network will be continuously integrated and validated for performance and relevance with technical and application partners.

Bias detection is crucial for our component (and project in general). We will extend our analysis to integrate all the data sources considered in MIRROR. We will also check the differences between regional and English channels of a specific region. One is mostly used to reach global audiences, while local audiences usually consume information posted in their native languages. Hence, it is essential to analyze such representation gaps, if any. We also intend to extend and refine the topical coverage of our analysis. Likewise, so far we have considered regions that cover large geographical areas, but they might internally follow different patterns of information consumption. We will investigate the optimal level of aggregation for different dimensions and the rate of information loss as an additional input for end-users.

The next steps in community analysis involve the development of methods to improve our understanding of how these identified communities evolve or react to specific events.

7 References

[bot,] Bot sentinel. https://botsentinel.com/.

[twi,] Twitteraudit. https://www.twitteraudit.com/.

[Abebe et al., 2018] Abebe, R., Kleinberg, J., Parkes, D., and Tsourakakis, C. E. (2018). Opinion dynamics with varying susceptibility to persuasion. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD '18, page 1089–1098.

[Abras et al., 2004] Abras, C., Maloney-Krichmar, D., and Preece, J. (2004). User-centered design. In Bainbridge, W. S., editor, *Berkshire Encyclopedia of Human-Computer Interaction: When science fiction becomes science fact*, volume 2, pages 763–768. Berkshire Publishing Group, Great Barrington, Massachusetts.

[Adam, 2012] Adam, L. (2012). The significance of eu topics in national media. has there been an europeanization of reporting in the national media. In *Bruges Political Research Papers No. 27*. College of Europe.

[Agirre et al., 2016] Agirre, E., Banea, C., Cer, D., Diab, M., Gonzalez-Agirre, A., Mihalcea, R., Rigau, G., and Wiebe, J. (2016). Semeval-2016 task 1: Semantic textual similarity, monolingual and cross-lingual evaluation. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 497–511.

[Alencar, 2018] Alencar, A. (2018). Refugee integration and social media: A local and experiential perspective. *Information, Communication & Society*, 21(11):1588–1603.

[Andriole, 1989] Andriole, S. J. (1989). *Storyboard Prototyping: A New Approach to User Requirements Analysis*. QED Information Sciences, Inc., Wellesley, Massachusetts.

[Aswad and Menezes, 2018] Aswad, F. M. S. and Menezes, R. (2018). Refugee and immigration: Twitter as a proxy for reality. In *The Thirty-First International Flairs Conference*.

[Bahamonde et al., 2018] Bahamonde, J., Bollen, J., Elejalde, E., Ferres, L., and Poblete, B. (2018). Power structure in chilean news media. *PLOS ONE*, 13(6):1–18.

[Balasubramanyan and Kolcz, 2013] Balasubramanyan, R. and Kolcz, A. (2013). "w00t! feeling great today!" chatter in twitter: identification and prevalence. In *2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM 2013)*, pages 306–310. IEEE.

[Bhowmick et al., 2009] Bhowmick, P. K., Basu, A., and Mitra, P. (2009). Reader perspective emotion analysis in text through ensemble based multi-label classification framework. *Computer and Information Science*, 2(4):64–74.

[Bollobás, 1998] Bollobás, B. (1998). *Modern Graph Theory*. Graduate texts in mathematics. Springer, corrected edition.

[Borjas, 1989a] Borjas, G. J. (1989a). Economic theory and international migration. *The International Migration Review*, 23(3):457–485.

[Borjas, 1989b] Borjas, G. J. (1989b). Economic theory and international migration. *The International Migration Review*, 23(3):457–485.

[Boulding, 1959] Boulding, K. E. (1959). National images and international systems. *The Journal of Conflict Resolution*, 3(2):120–131.

[Bourdieu, 1986] Bourdieu, P. (1986). The forms of capital.

[Bourgeois et al., 2018] Bourgeois, D., Rappaz, J., and Aberer, K. (2018). Selection bias in news coverage: Learning it, fighting it. In *Companion Proceedings of the The Web Conference 2018*, WWW ’18, pages 535–543.

[Budak et al., 2016] Budak, C., Goel, S., and Rao, J. M. (2016). Fair and balanced? quantifying media bias through crowdsourced content analysis.

[Bujang and Baharum, 2016] Bujang, M. A. and Baharum, N. (2016). Sample size guideline for correlation analysis. *World*, 3(1).

[Castles et al., 2013] Castles, S., De Haas, H., and Miller, M. J. (2013). *The age of migration: International population movements in the modern world*. Macmillan International Higher Education.

[Castles and Miller, 2009] Castles, S. and Miller, M. J. (2009). *The Age of Migration*. The Guilford Press, London, 4 edition.

[Chaban and Holland, 2013] Chaban, N. and Holland, M. (2013). Seeing the eu from outside its borders: Changing images of europe. *Baltic Journal of European Studies*, 3(3):3 – 12.

[Chaban and Holland, 2014] Chaban, N. and Holland, M. (2014). *Communicating Europe in times of crisis: External perceptions of the European Union*. Springer.

[Chaban et al., 2019] Chaban, N., Knodt, M., Liekis, Š., and NG, I. (2019). Narrators’ perspectives: communicating the EU in ukraine, israel and palestine in times of conflict. *European Security*, pages 1–19.

[Chang et al., 2020] Chang, J. P., Cheng, J., and Danescu-Niculescu-Mizil, C. (2020). Don’t let me be misunderstood: Comparing intentions and perceptions in online discussions. In *Proceedings of the 29th International Conference on World Wide Web*, WWW ’20.

[Cohen, 1987] Cohen, R. (1987). *The New Helots: Migrants in the International Division of Labour (Research in Ethnic Relations Series)*. Gower Pub Co.

[Colleoni et al., 2014] Colleoni, E., Rozza, A., and Arvidsson, A. (2014). Echo Chamber or Public Sphere? Predicting Political Orientation and Measuring Political Homophily in Twitter Using Big Data. *Journal of Communication*, 64(2):317–332.

[Cresci et al., 2017] Cresci, S., Di Pietro, R., Petrocchi, M., Spognardi, A., and Tesconi, M. (2017). The paradigm-shift of social spambots: Evidence, theories, and tools for the arms race. In *Proceedings of the 26th International Conference on World Wide Web Companion*, pages 963–972.

[Dallmann et al., 2015] Dallmann, A., Lemmerich, F., Zoller, D., and Hotho, A. (2015). Media bias in german online newspapers. In *HT*.

[Davis et al., 2016] Davis, C. A., Varol, O., Ferrara, E., Flammini, A., and Menczer, F. (2016). Botornot: A system to evaluate social bots.

[Dekker and Engbersen, 2014] Dekker, R. and Engbersen, G. (2014). How social media transform migrant networks and facilitate migration. *Global Networks*, 14(4):401–418.

[Dekker et al., 2016] Dekker, R., Engbersen, G., and Faber, M. (2016). The use of online media in migration networks. *Population, Space and Place*, 22(6):539–551.

[Dekker et al., 2018] Dekker, R., Engbersen, G., Klaver, J., and Vonk, H. (2018). Smart refugees: How syrian asylum migrants use social media information in migration decision-making. *Social Media+ Society*, 4(1):2056305118764439.

[Del Vicario et al., 2015] Del Vicario, M., Bessi, A., Zollo, F., Petroni, F., Scala, A., Caldarelli, G., Stanley, H. E., and Quattrociocchi, W. (2015). Echo chambers in the age of misinformation. *arXiv preprint arXiv:1509.00189*.

[Del Vicario et al., 2016] Del Vicario, M., Vivaldo, G., Bessi, A., Zollo, F., Scala, A., Caldarelli, G., and Quattrociocchi, W. (2016). Echo chambers: Emotional contagion and group polarization on facebook. *Scientific reports*, 6:37825.

[Doe et al., 2019] Doe, J., Appleseed, J., and Atkins, T. (2019). MIRROR Project Deliverable DX.Y - Yet Another Exciting Result. https://h2020mirror.eu/DX.Y.pdf. Last visited 20-10-2019.

[Eberl, 2019] Eberl, J.-M. (2019). Lying press: Three levels of perceived media bias and their relationship with political preferences.

[Elejalde et al., 2018] Elejalde, E., Ferres, L., and Herder, E. (2018). On the nature of real and perceived bias in the mainstream media. *PLOS ONE*, 13(3):1–28.

[Elejalde et al., 2019] Elejalde, E., Ferres, L., and Schifanella, R. (2019). Understanding news outlets´ audience-targeting patterns. *EPJ Data Sci.*, 8(1):16.

[Elejalde, Erick et al., 2019] Elejalde, Erick, Ferres, Leo, and Schifanella, Rossano (2019). Understanding news outlets´ audience-targeting patterns. *EPJ Data Sci.*, 8(1):16.

[Europe Durect, 2019] Europe Durect (2019). EUROSTAT - your key to european statistics.

[Fang et al., 2020] Fang, Y., Huang, X., Qin, L., Zhang, Y., Zhang, W., Cheng, R., and Lin, X. (2020). A survey of community search over big graphs. *The VLDB Journal*, 29(1):353–392.

[Fernquist et al., 2018] Fernquist, J., Kaati, L., and Schroeder, R. (2018). Political bots and the swedish general election. In *Proceedings of the 2018 IEEE International Conference on Intelligence and Security Informatics (ISI 2018)*, pages 124–129.

[Fernquist et al., 2019] Fernquist, J., Svenonius, O., Kaati, L., and Johansson, F. (2019). Extracting account attributes for analyzing influence on twitter. In Brynielsson, J., editor, *Proceedings of the 2019 European Intelligence and Security Informatics Conference (EISIC 2019)*, Piscataway, New Jersey. IEEE.

[Flaounas et al., 2010] Flaounas, I., Turchi, M., Ali, O., Fyson, N., De Bie, T., Mosdell, N., Lewis, J., and Cristianini, N. (2010). The structure of the eu mediasphere. *PLOS ONE*, 5(12):1–6.

[Frech, 2008] Frech, A. (2008). ’mass migration’, crime and ’decent people’: The portrayal of polish migrants in britisch newspapers.

[Garimella and Weber, 2014] Garimella, V. R. K. and Weber, I. (2014). Co-following on twitter. In *Proceedings of the 25th ACM conference on Hypertext and social media*, pages 249–254.

[Georgiou and Zaborowski,] Georgiou, M. and Zaborowski, R. Media coverage of the ”refugee crisis”: A cross-european perspective. *Council of Europe report (DG1(2017)03)*.

[Gilani et al., 2017] Gilani, Z., Farahbakhsh, R., Tyson, G., Wang, L., and Crowcroft, J. (2017). Of bots and humans (on twitter). In *Proceedings of the 2017 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM 2017)*, pages 349–354.

[Gong et al., 2015] Gong, W., Lim, E.-P., and Zhu, F. (2015). Characterizing silent users in social media communities. In *Ninth International AAAI Conference on Web and Social Media*.

[Grefenstette et al., 2004] Grefenstette, G., Qu, Y., Shanahan, J. G., and Evans, D. A. (2004). Coupling niche browsers and affect analysis for an opinion mining application. *Proceedings of Recherche d'Information Assistée par Ordinateur (RIAO)*.

[Groseclose and Milyo, 2005a] Groseclose, T. and Milyo, J. (2005a). A measure of media bias. *The Quarterly Journal of Economics*, 120(4):1191–1237.

[Groseclose and Milyo, 2005b] Groseclose, T. and Milyo, J. (2005b). A measure of media bias. *The Quarterly Journal of Economics*, 120(4):1191–1237.

[Gruenewald et al., 2009] Gruenewald, J., Pizarro, J., and Chermak, S. M. (2009). Race, gender, and the newsworthiness of homicide incidents. *Journal of criminal justice*, 37(3):262–272.

[Gualda and Rebollo, 2016] Gualda, E. and Rebollo, C. (2016). The refugee crisis on twitter: A diversity of discourses at a european crossroads. *Journal of Spatial and Organizational Dynamics*, 4(3):199–212.

[Hadgu et al., 2016] Hadgu, A. T., Naini, K. D., and Niederée, C. (2016). Welcome or not-welcome: Reactions to refugee situation on social media. *arXiv preprint arXiv:1610.02358*.

[Hamborg et al., 2018] Hamborg, F., Donnay, K., and Gipp, B. (2018). Automated identification of media bias in news articles: an interdisciplinary literature review. *International Journal on Digital Libraries*.

[Hoang et al., 2018] Hoang, T.-A., Vo, K. D., and Nejdl, W. (2018). W2e: A worldwide-event benchmark dataset for topic detection and tracking. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, CIKM '18, page 1847–1850.

[Joachims, 1999] Joachims, T. (1999). Transductive inference for text classification using support vector machines. In *Icml*, volume 99, pages 200–209.

[Jr and Lee, 2000] Jr, R. and Lee, M. (2000). On immigration and crime. *Criminal Justice*, 1:485–524.

[Kelly, 2013] Kelly, S. (2013). New zealand elite perceptions on the eu: A longitudinal analysis. *Baltic Journal of European Studies*, 3(3):153–174.

[Kloosterman et al., 1998] Kloosterman, R., Van der Leun, J., and Rath, J. (1998). Across the border: economic opportunities. *Social capital, and informal business activities of immigrants', Journal of Ethnic and Migration Studies*, 24:239–258.

[Komito, 2011] Komito, L. (2011). Social media and migration: Virtual community 2.0. *Journal of the American society for information science and technology*, 62(6):1075–1086.

[Kreil, 2018] Kreil, M. (2018). The social bot research of Oxford and Co. is flawed.

[Kreis, 2017] Kreis, R. (2017). # refugeesnotwelcome: Anti-refugee discourse on twitter. *Discourse & Communication*, 11(5):498–514.

[Lamanna et al., 2018] Lamanna, F., Lenormand, M., Salas-Olmedo, M. H., Romanillos, G., Gonçalves, B., and Ramasco, J. J. (2018). Immigrant community integration in world cities. *PloS one*, 13(3).

[Lee, 1966] Lee, E. S. (1966). A theory of migration. *Demography*, 3(1):47–57.

[Li et al., 2009] Li, T., Zhang, Y., and Sindhwani, V. (2009). A non-negative matrix tri-factorization approach to sentiment classification with lexical prior knowledge. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP: Volume 1-Volume 1*, pages 244–252. Association for Computational Linguistics.

[Machill et al., 2006] Machill, M., Beiler, M., and Fischer, C. (2006). Europe-topics in europe's media: The debate about the european public sphere: A meta-analysis of media content analyses. *European Journal of Communication*, 21(1):57–88.

[McCluskey and Kim, 2012] McCluskey, M. and Kim, Y. M. (2012). Moderatism or polarization? representation of advocacy groups' ideology in newspapers. *Journalism & Mass Communication Quarterly*, 89(4):565–584.

[McCombs and Shaw, 1972] McCombs, M. E. and Shaw, D. L. (1972). The agenda-setting function of mass media. *The Public Opinion Quarterly*, 36(2).

[McPherson et al., 2001] McPherson, M., Smith-Lovin, L., and Cook, J. M. (2001). Birds of a feather: Homophily in social networks. *Annual review of sociology*, 27(1):415–444.

[MIRROR, 2019] MIRROR (2019). Project Web Site. https://h2020mirror.eu. Last visited 20-10-2019.

[Moon, 1995] Moon, B. (1995). Paradigms in migration research: exploring 'moorings' as a schema. *Progress in Human Geography*, 19(4):504–524. PMID: 12347395.

[Mutz, 2006] Mutz, D. C. (2006). *Hearing the other side: Deliberative versus participatory democracy*. Cambridge University Press.

[Nerghes and Lee, 2018] Nerghes, A. and Lee, J.-S. (2018). The refugee/migrant crisis dichotomy on twitter: A network and sentiment perspective. In *Proceedings of the 10th ACM Conference on Web Science*, pages 271–280.

[Nerghes and Lee, 2019] Nerghes, A. and Lee, J.-S. (2019). Narratives of the refugee crisis: A comparative study of mainstream-media and twitter. *Media and Communication*, 7(2 Refugee Crises Disclosed):275–288.

[Niculae et al., 2015] Niculae, V., Suen, C., Zhang, J., Danescu-Niculescu-Mizil, C., and Leskovec, J. (2015). Quotus: The structure of political media coverage as revealed by quoting patterns. In *Proceedings of the 24th International Conference on World Wide Web*, WWW '15, page 798–808.

[Papacharissi and de Fatima Oliveira, 2008] Papacharissi, Z. and de Fatima Oliveira, M. (2008). News frames terrorism: A comparative analysis of frames employed in terrorism coverage in us and uk newspapers. *The international journal of press/politics*, 13(1):52–74.

[Pinto et al., 2019] Pinto, S., Albanese, F., Dorso, C. O., and Balenzuela, P. (2019). Quantifying time-dependent media agenda and public opinion by topic modeling.

[Pinto et al., 2015] Pinto, S., Balenzuela, P., and Dorso, C. (2015). Setting the agenda: Different strategies of a mass media in a model of cultural dissemination. *Physica A: Statistical Mechanics and its Applications*, 458.

[Portes and Böröcz, 1989] Portes, A. and Böröcz, J. (1989). Contemporary immigration: Theoretical perspectives on its determinants and modes of incorporation. *The International Migration Review*, 23(3):606–630.

[Ribeiro et al., 2018] Ribeiro, F., Henrique, L., Benevenuto, F., Chakraborty, A., Kulshrestha, J., Babaei, M., and Gummadi, K. (2018). Media bias monitor: Quantifying biases of social media news outlets at large-scale.

[Righi, 2019] Righi, A. (2019). Assessing migration through social media: a review. *Mathematical Population Studies*, 26(2):80–91.

[Riquelme and González-Cantergiani, 2016] Riquelme, F. and González-Cantergiani, P. (2016). Measuring user influence on twitter: A survey. *Information processing & management*, 52(5):949–975.

[Romero et al., 2011] Romero, D. M., Meeder, B., and Kleinberg, J. (2011). Differences in the mechanics of information diffusion across topics: Idioms, political hashtags, and complex contagion on twitter. In *Proceedings of the 20th International Conference on World Wide Web*, pages 695–704.

[Rotabi et al., 2017] Rotabi, R., Danescu-Niculescu-Mizil, C., and Kleinberg, J. (2017). Competition and selection among conventions. In *Proceedings of the 26th International Conference on World Wide Web*, WWW '17, page 1361–1370.

[Saez-Trumper et al., 2013] Saez-Trumper, D., Castillo, C., and Lalmas, M. (2013). Social media news communities: gatekeeping, coverage, and statement bias. In *Proceedings of the 22nd ACM international conference on Information & Knowledge Management*, pages 1679–1684.

[Sanderson, 1997] Sanderson, M. (1997). Duplicate detection in the reuters collection. *" Technical Report (TR-1997-5) of the Department of Computing Science at the University of Glasgow G12 8QQ, UK"*.

[Sassen, 1988] Sassen, S. (1988). *The Mobility of Labor and Capital*. Cambridge University Press.

[Shoemaker et al., 2001] Shoemaker, P. J., Eichholz, M., Kim, E., and Wrigley, B. (2001). Individual and routine forces in gatekeeping. *Journalism & mass communication quarterly*, 78(2):233–246.

[Siapera et al., 2018] Siapera, E., Boudourides, M., Lenis, S., and Suiter, J. (2018). Refugees and network publics on twitter: Networked framing, affect, and capture. *Social Media+ Society*, 4(1):2056305118764437.

[Simonsen and Robertson, 2013] Simonsen, J. and Robertson, T. (2013). *Routledge International Handbook of Participatory Design*. Routledge, New York.

[Soroka, 2012] Soroka, S. N. (2012). The gatekeeping function: Distributions of information in media and the real world. *The Journal of Politics*, 74(2):514–528.

[Sørensen and Gammeltoft-Hansen, 2012] Sørensen, N. N. and Gammeltoft-Hansen, T. (2012). The migration industry and future directions for migration policy. *Danish Institute for International Studies*, page 4.

[Tan et al., 2011] Tan, C., Lee, L., Tang, J., Jiang, L., Zhou, M., and Li, P. (2011). User-level sentiment analysis incorporating social networks. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1397–1405.

[Tan et al., 2018] Tan, C., Peng, H., and Smith, N. A. (2018). "you are no jack kennedy": On media selection of highlights from presidential debates. In *Proceedings of the 2018 World Wide Web Conference*, WWW '18, page 945–954.

[Tang and Liu, 2010] Tang, L. and Liu, H. (2010). Community detection and mining in social media. *Synthesis lectures on data mining and knowledge discovery*, 2(1):1–137.

[Timmerman et al., 2014] Timmerman, C., Hemmerechts, K., and Clerck, H. M.-L. D. (2014). The relevance of a “culture of migration” in understanding migration aspirations in contemporary turkey. *Turkish Studies*, 15(3):496–518.

[UN Refugee Agency, 2019] UN Refugee Agency (2019). UNHCR - the un refugee agency.

[Varol et al., 2017] Varol, O., Ferrara, E., Davis, C. A., Menczer, F., and Flammini, A. (2017). Online human-bot interactions: Detection, estimation, and characterization. In *Proceedings of the Eleventh International AAAI Conference on Web and Social Media*.

[Volkova et al., 2014] Volkova, S., Coppersmith, G., and Van Durme, B. (2014). Inferring user political preferences from streaming communications. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 186–196.

[Wang et al., 2019] Wang, L., Niu, J., Liu, X., and Mao, K. (2019). The silent majority speaks: Inferring silent users’ opinions in online social networks. In *The World Wide Web Conference*, pages 3321–3327.

[Welch’s unequal variances t-test, 2020] Welch’s unequal variances t-test (2020). Welch’s t-test.

[Wojcik and Hughes, 2019] Wojcik, S. and Hughes, A. (2019). Sizing up twitter users. *Internet & Tech. Pew Research Center*. Available from: http://www.people-press.org/2007/10/12/too-much-celebrity-news-too-little-good-news/ [Accessed 02-Jan-2018].

[Yang and Leskovec, 2013] Yang, J. and Leskovec, J. (2013). Overlapping community detection at scale: a nonnegative matrix factorization approach. In *Proceedings of the sixth ACM international conference on Web search and data mining*, pages 587–596.

[Yang et al., 2013] Yang, J., McAuley, J., and Leskovec, J. (2013). Community detection in networks with node attributes. In *2013 IEEE 13th International Conference on Data Mining*, pages 1151–1156. IEEE.

[Yang et al., 2019a] Yang, K.-C., Varol, O., Davis, C. A., Ferrara, E., Flammini, A., and Menczer, F. (2019a). Arming the public with artificial intelligence to counter social bots. *Human Behavior and Emerging Technologies*, 1(1):48–61.

[Yang et al., 2019b] Yang, K.-C., Varol, O., Hui, P.-M., and Menczer, F. (2019b). Scalable and generalizable social bot detection through data selection.

[Yang et al., 2014a] Yang, L., Cao, X., Jin, D., Wang, X., and Meng, D. (2014a). A unified semi-supervised community detection framework using latent space graph regularization. *IEEE transactions on cybernetics*, 45(11):2585–2598.

[Yang et al., 2014b] Yang, S.-H., Kolcz, A., Schlaikjer, A., and Gupta, P. (2014b). Large-scale high-precision topic modeling on twitter. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1907–1916.

[Zipf, 1946a] Zipf, G. K. (1946a). The P_1P_2/D hypothesis: on the intercity movement of persons. *American Sociological Review*, 11(6):677–686.

[Zipf, 1946b] Zipf, G. K. (1946b). Some determinants of the circulation of information. *The American Journal of Psychology*, 59(3):401–421.