

# Know Who Your Friends Are: Understanding Social Connections from Unstructured Text

Léa A. Deleris, Francesca Bonin, Elizabeth Daly, Stéphane Deparis,  
Yufang Hou, Charles Jochim, Yassine Lassoued and Killian Levacher

IBM Research - Ireland

Dublin, Ireland

lea.deleris, fbonin, edaly, stephane.deparis

yhou, charlesj, ylassoued, killian.levacher @ie.ibm.com

## Abstract

Having an understanding of interpersonal relationships is helpful in many contexts. Our system seeks to assist humans with that task, using textual information (e.g., case notes, speech transcripts, posts, books) as input. Specifically, our system first extracts qualitative and quantitative information elements (which we call signals) about interactions among persons, aggregates those to provide a condensed view of relationships and then enables users to explore all facets of the resulting social (multi-)graph through a visual interface.

## 1 Introduction

The social network of a person plays a vital role in their well being providing access to assistance, resources, support (Wellman and Wortley, 1990) and even influencing health (Christakis and Fowler, 2007). Understanding the quality of the social relationships, beyond the simple existence of a relationship or its demographic nature (family or not), provides a better perspective on the context and is essential in a variety of situations that extends beyond social media such as criminal investigations (identifying suspects and acolytes), sales (connecting to the right persons in target companies), political analysis (understanding the evolution of alliances) and human resources (improving team dynamics). For instance, Yose et al. (2018) analysed a medieval text to get novel insights into the hostilities during the battle of Clontarf of 1014.

The initial motivation for this work comes from the domain of social care. One essential task for care workers is to identify who plays a supportive or disruptive role in a patient’s environment at a given time. Problems emerge when details get lost within notes of the multiple persons composing the care team. Significant details may have been recorded but are locked inside free-text narratives,

requiring other parties to invest a great deal of time to gain a full understanding of the situation.

In this paper, we present a system designed to assist humans with the understanding of interpersonal relationships. Specifically, the system takes as input a collection of texts and automatically extracts a multigraph representation of the relationships (and their quality). From that information, we provide an interface to enable users to gain insights into the relationships, from aggregated information to fine-grained analysis of temporal patterns and emotions.

## 2 Background

Our work builds upon two areas: social graph identification and qualitative relationship analysis. Most of the research within the former attempts to build ontologies of relations between individuals based on information extracted from text (e.g., social networks, dialogues, novels) (Hauffa et al., 2011). Mori et al. (2006) try to enhance social network extraction from the web, considering a range of entities beyond just persons. They identify the nature of entities (e.g., person, firm, geopolitical entity) and their relations (e.g., mayor of).

The second area instead explores how to construct a qualitative representation of the relationship between individuals. Bracewell et al. (2011) for example determine the collegiality between two persons, as exhibited in a text. Srivastava et al. (2016) attempt to identify the polarity of relations (cooperative or conflictual) between characters in narrative summaries, using textual and structural features. Iyyer et al. (2016) model how the fictional relationship between two characters changes over time using deep recurrent autoencoders.

Altogether, while the first field enables the extraction of complex multi-party social graphs, links representing interpersonal relationships within

these graphs often lack a more nuanced representation (Hauffa et al., 2011). On the other hand, techniques from qualitative relationship analysis lead to a more extensive understanding of relationships yet they are typically applied to minimal settings consisting of two or three individuals at most. Our work brings both fields together by producing larger multi-party social graphs with qualitative links identified between the individuals.

### 3 Modeling Relationships

We model *interpersonal relationships*, or simply *relationships* between two *entities* by analyzing a list of associated *relationship signals* (hereafter *signals*) which are extracted by an NLP system. At this point, we focus on four kinds of *signals*:

- Direct Speech - A person addressing another one without mentioning him/her, e.g., *Phoebe* (to *Monica*): “*The weather is nice today*”.
- Direct Reference - A person addressing another one and mentioning him/her explicitly, e.g., *Phoebe* (to *Monica*): “*I hate you*”.
- Indirect Reference - A person mentioning a third party, i.e., someone who is not present, e.g., *Phoebe*: “*I like Rachel*”.
- Third-Party Reference - The description, by a third-party, of any kind of relation between two entities, e.g., *Phoebe*: “*Ross has been in love with Rachel forever*”.

While we focus in our examples on humans, entities could also include corporations, governmental organizations, products, brands and animals (e.g., “*Toyota allies with Intel in bid to overtake GM*”).

For each detected signal, our system seeks to present a qualitative description including **sentiment**, **emotion** such as anger, disgust, joy, etc., other qualitative aspects such as **intensity**, **formality**, **cooperative vs. adversarial** along with information about the context such as **geographical location**, settings (**one-on-one vs. group discussion**), whether **face-to-face** or **remote**, whether **synchronous** or **asynchronous** interactions.

In turn, through a model to aggregate those signals together and over time, this enables us to attach qualitative metrics to *relationships* for instance **sentiment and emotions** – so as to differentiate between supportive relationships and negative relationships – but also **volatility** of sentiment and

emotions, **intensity/frequency**<sup>1</sup> and **recency**.

The logic underneath the aggregation model can vary by situation, but is typically based on weighted counts of the various signals. For instance, in the sentence “*You, Frankie, you are a liar*”, the system would extract at the atomic level one direct speech signal from the speaker towards Frankie and three direct references as well. When aggregating, it may be that the three direct references are counted as one rather than three, being all from the same utterance and that the direct reference takes priority over the direct speech signal which could be discarded. However, in some scenario, it may be relevant to keep track at the level of each atomic signal. Our system provides the flexibility to design bespoke models of how signals map to relationships. From an engineering perspective, we rely on Solr<sup>2</sup> to index each signal with all inferred facets such as entities, sentiment and type.

### 4 NLP System

Our system combines NLP tools to perform *entity mention resolution* and then extracts signals between these entities.<sup>3</sup> Our entity mention resolution borrows from named-entity recognition, mention detection, coreference resolution, and entity linking. We address more than just named entities (i.e., we are interested in all (person) entity mentions), and look at not only coreferenced ones. To get a complete picture of the connection for a person, we want to identify all possible references to them, whether named entities (e.g., Ross Geller, Ross) or not (e.g., him, you, that guy, his son, the professor).

**Entity Mention Detection.** We use named-entity recognition and coreference resolution to identify mentions of entities (people) in the text and supplement that with tools for identifying common roles (e.g., professions), titles, or relations (e.g., brother or neighbor).

**Entity Linking and Resolution.** We leverage social graph information to resolve the detected entity mentions. This involves, for example, making use of family relationships to find or disambiguate new mentions of entities in the text. When no social graph is provided, this step includes building the graph from a knowledge source or to bootstrap social graph creation using the output of the

<sup>1</sup>Granovetter observes that “the more frequently persons interact with one another, the stronger their sentiments of friendship for one another are apt to be” (Granovetter, 1973)

<sup>2</sup><http://lucene.apache.org/solr/>

<sup>3</sup>As mentioned, we focus primarily on person entities.

entity mention detection. For those cases, further named-entity disambiguation may be needed (Pabico, 2014; Han and Zhao, 2009).

**Relationship Signal Detection.** Using the resolved entity mentions we detect relationship signals between entities and characterize the signal type (Section 3). In addition, we perform targeted sentiment and emotion analysis between entities.

Figure 1 illustrates how the different modules are articulated. The dashed arrow going from the graph-based mention resolution towards the text-level mention detection module indicates that information from the graph can further contribute in disambiguating and enriching the text-level mention detection steps (for instance pronoun resolution).

The output of the NLP extraction pipeline comprises a list of entity mentions and a list of relationship signals between those entities. These contribute to enrich our understanding of the social graph among the entities (or create one if none were provided as input). From this combined information we draw the relationship graph.

## 5 Visual Interface Features

Our current system enables users to (i) Aggregate information (over time, over sentiment categories), (ii) Visualize temporal relations over time and (iii) Scrutinize each atomic signal underneath each relationship (with the possibility of editing or correcting it). Another task that we are considering enabling is the ability to reason over the graph (knowledge propagation e.g., if A “admires” B and B “admires” C, it is likely that A “admires” C).

Our current interface hinges on three primary views. **Network overview** (Fig. 2) which provides a snapshot of the social relationships using visual cues to summarise them graphically. The average sentiment is indicated through the color of the links, red for negative, green for positive where thickness gives an indication of the strength, the thicker the link, the more positive or negative the sentiment. The UI supports visualising the network as an overview or over specific time intervals. Next version will include a representation of intensity (through dashes) along with an indication of volatility by replacing the line with a sinusoidal curve. Recency will be captured through the placement relative to the ego node, people who more recently interacted with the person are closer.

The **Personal network view** (Fig. 3) supports a view of the two-hop interpersonal relationships

of the individual along with a summary view of the person in terms of general emotions expressed (donut chart in the bottom right quadrant of the screen). The view can show the overall sentiment and intensity of how they interact with others.

Finally, the **Relationship view** (Fig. 4) consists of a stream graph presenting the five primary emotions, anger, disgust, joy, fear and sadness as they change over time. It can be used to see the intensity of the relationships either by type, sentiment or emotions. Drill down support is provided where the user may hover over peak areas of interest in the graph to inspect the snippets of text from the input corpus that support the inferred emotions. The emotions are directional where the upper part of the stream graph represents the signals from Entity X to Entity Y and the lower part of the graph represents Entity Y towards Entity X.

## 6 Demo Flow

Ideally, we would have relied on anonymized care-worker notes to demonstrate our system. However, privacy restrictions made access to such notes challenging and for our proof of concept, we have used the transcripts from the Friends TV series, whose theme, i.e. the interactions among a core group of friends is well aligned with our objective. We specifically focus on the first two seasons and leverage the character identification corpus created by Chen and Choi (2016), which includes entity annotation of personal mentions<sup>4</sup>.

A natural starting page for an analysis is the (messy) social graph for all characters over all episodes of Season 1. Fig. 2 shows the overall social network when narrowing down to Season 1 Episode 2. From this place, a user can further dive into one specific character (selecting a node), or rather explore a relationship (selecting a link).

Assuming the user selects to focus on *Ross* then his personal network view (Fig. 3) would be displayed. The dominating emotion conveyed in this episode is joy, both in quantity and in strength (represented through length of the arc and radius respectively), followed by sadness and fear. In this episode, *Ross* learns from his now-lesbian ex-wife *Carol* that she is pregnant with his child. In terms of relationship, the graph shows mostly negative sentiments towards other characters, especially *Barry*, though positive ones towards *Rachel*.

<sup>4</sup><https://github.com/emorynlp/character-mining>

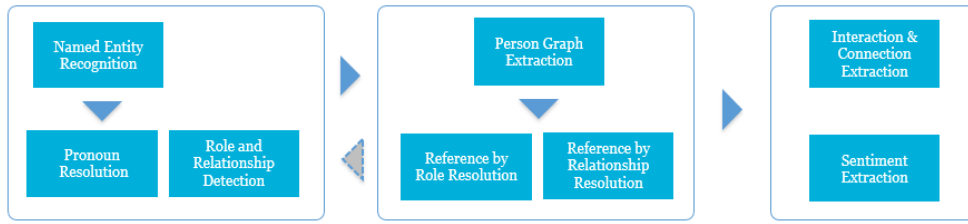


Figure 1: Overview of the NLP Interaction Extraction System

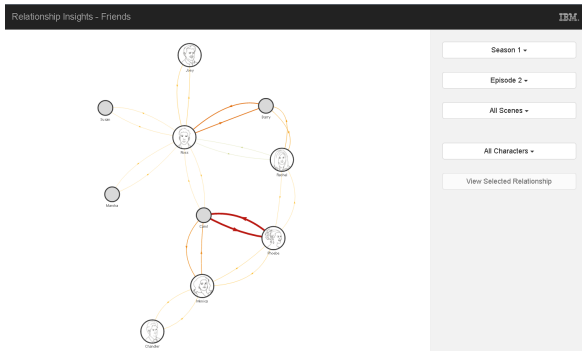


Figure 2: Overview of Social Relationships [S1E2]

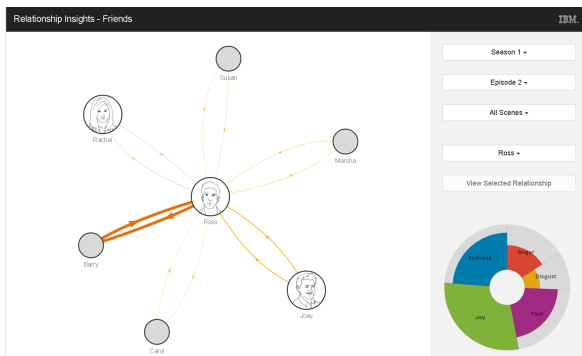


Figure 3: Personal Network for Ross [S1E2]

Selecting a relationship (eg. the one between Ross and Rachel) the user will be shown a stream graph as in Fig. 4. The graph shows the five primary emotions as they change over time during the series. The x axis represents the progressive episodes' numbers and the y axis the intensity of the emotion. Above the x axis the relation Ross towards Rachel is shown, below the relation Rachel towards Ross. We can observe the distinctive asymmetry of the relationship. For instance, in Episode 11, there is no specific emotion from *Rachel* towards *Ross* though quite a mix of emotions expressed by *Ross* towards *Rachel*. In this episode, *Ross* is jealous of *Rachel's* boyfriend *Paolo* and confides his feelings and love troubles towards *Rachel* to *Chandler's*

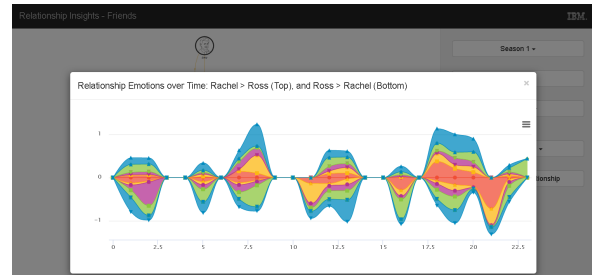


Figure 4: Stream graph of Relationship over time between Ross and Rachel [Season 1]

mother. In Episode 16, sentiments of anger, fear and disgust are present in Ross towards Rachel, but not viceversa.

As part of the actual demonstration, we will challenge users with a short information retrieval task aimed at illustrating the types of questions that our approach supports. For fact-based questions, we will give participants a set amount of time (i.e. 2 minutes) and will maintain a leaderboard of the fastest information finders (with participants' consent). Sample questions would be “*Whose mother did Ross kiss during season 1?*”, or “*How long did Rachel go out with Paolo?*”. For questions that require interpretation (e.g. “*Does Chandler like Frankie?*”), we would let participants tell us their opinion before and after using the tool and ask them if the system was helpful in gaining confidence about their final answer.

## 7 Conclusion

This paper presents a system to support the analysis and understanding of interpersonal relationships. Relationships are described along multiple quantitative and qualitative dimensions which are automatically populated from relationship signals extracted from text by a NLP system. The associated interface enables a user to quickly focus on a specific person or pair of persons and to investigate how the relationship evolves over time.

## References

- David B Bracewell, Marc Tomlinson, Ying Shi, Jeremy Bensley, and Mary Draper. 2011. Who's playing well with others: Determining collegiality in text. In *Proceedings of the 5th IEEE International Conference on Semantic Computing (ICSC 2011)*, pages 21–26. IEEE.
- Yu-Hsin Chen and Jinho D. Choi. 2016. Character identification on multiparty conversation: Identifying mentions of characters in TV shows. In *Proceedings of the SIGDIAL 2016 Conference, The 17th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 90–100. The Association for Computer Linguistics.
- Nicholas A Christakis and James H Fowler. 2007. The spread of obesity in a large social network over 32 years. *New England journal of medicine*, 357(4):370–379.
- Mark S Granovetter. 1973. The strength of weak ties. *American journal of sociology*, pages 1360–1380.
- Xianpei Han and Jun Zhao. 2009. [Named entity disambiguation by leveraging wikipedia semantic knowledge](#). In *Proceedings of the 18th ACM Conference on Information and Knowledge Management, CIKM '09*, pages 215–224, New York, NY, USA. ACM.
- Jan Hauffa, Gottlieb Bossert, Nadja Richter, Florian Wolf, Nora Liesenfeld, and Georg Groh. 2011. Beyond foaf: Challenges in characterizing social relations. In *Privacy, Security, Risk and Trust (PASSAT) and 2011 IEEE Third International Conference on Social Computing (SocialCom), 2011 IEEE Third International Conference on*, pages 790–795. IEEE.
- Mohit Iyyer, Anupam Guha, Snigdha Chaturvedi, Jordan Boyd-Graber, and Hal Daumé III. 2016. Feuding families and former friends: Unsupervised learning for dynamic fictional relationships. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1534–1544.
- Junichiro Mori, Takumi Tsujishita, Yutaka Matsuo, and Mitsuru Ishizuka. 2006. Extracting relations in social networks from the web using similarity between collective contexts. *The Semantic Web-ISWC 2006*, pages 487–500.
- Jaderick P. Pabico. 2014. An analysis of named entity disambiguation in social networks. *Asia Pacific Journal of Multidisciplinary Research*, 2(4).
- Shashank Srivastava, Snigdha Chaturvedi, and Tom M Mitchell. 2016. Inferring interpersonal relations in narrative summaries. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, pages 2807–2813.
- Barry Wellman and Scot Wortley. 1990. Different strokes from different folks: Community ties and social support. *American journal of Sociology*, pages 558–588.
- Joseph Yose, Ralph Kenna, Máirín MacCarron, and Pádraig MacCarron. 2018. Network analysis of the viking age in ireland as portrayed in cogadh gaedhel re gallaibh. *Royal Society Open Science*, 5(1):171024.