



HAL
open science

When face-tracking meets social networks: a story of politics in news videos

Benjamin Renoust, Tetsuro Kobayashi, Thanh Duc Ngo, Duy-Dinh Le,
Shin'Ichi Satoh

► To cite this version:

Benjamin Renoust, Tetsuro Kobayashi, Thanh Duc Ngo, Duy-Dinh Le, Shin'Ichi Satoh. When face-tracking meets social networks: a story of politics in news videos. Applied Network Science, 2016, 1, 10.1007/s41109-016-0003-2 . hal-01454830

HAL Id: hal-01454830

<https://hal.science/hal-01454830>

Submitted on 3 Feb 2017

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

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

RESEARCH

When face-tracking meets social networks: a story of politics in news videos

Benjamin Renoust^{1,2*}, Tetsuro Kobayashi³, Thanh Duc Ngo⁴, Duy Dinh Le¹ and Shin'Ichi Satoh¹

*Correspondence:
renoust@nii.ac.jp

¹National Institute of Informatics,
2-1-2 Hitotsubashi, Chiyoda-ku,
101-8430 Tokyo, Japan

Full list of author information is
available at the end of the article

Abstract

In the age of data processing, news videos are rich mines of information. After all, the news are essentially created to convey information to the public. But can we go beyond what is directly presented to us and see a wider picture? Many works already focus on what we can discover and understand from the analysis of years of news broadcasting. These analysis bring monitoring and understanding of the activity of public figures, political strategies, explanation and even prediction of critical media events. Such tools can help public figures in managing their public image, as well as support the work of journalists, social scientists and other media experts. News analysis can also be seen from the lens of complex systems, gathering many types of entities, attributes and interactions over time. As many public figures intervene in different news stories, a first interesting task is to observe the social interactions between these actors. Towards this goal, we propose to use video analysis to automatise the process of constructing social networks directly from news video archives. In this paper we are introducing a system deriving multiple social networks from face detections in news videos. We present preliminary results obtained from analysis of these networks, by monitoring the activity of more than a hundred public figures. We finally use these networks as a support for political studies and we provide an overview of the political landscape presented by the Japanese public broadcaster NHK over a decade of the 7PM news archives.

Keywords: news analysis; social network analysis; face detection and tracking; multiplex; dynamic; politics; Japan

1 Introduction

In the recent years, the publication of news information has migrated from the traditional means of newspapers, radio, and television to the wider audience offered by the Internet. With the rise of the data-intensive science [1] the analysis and monitoring of news information has given birth to the discipline called topic detection and tracking [2] which aims at segmenting, identifying, and following information, mainly from raw textual information. News analysis is now going beyond, with image information also being investigated across all varieties of media [3].

The analysis of news information is key to a wide variety of tasks, from sociology and journalism to politics and economy [4]. It could help the comprehension of user behavior such as what kind of information a category of users can be exposed to [5]. It could also bring new quantitative tools to overcome the limitations of technocratic measures in the investigation of freedom of information [6]. Even if we know that media competition can lead to a lower quality of information [4], we can hope that

public broadcasting services tend to convey an official character, and be a reliable baseline for social analysis.

The contribution of this paper is social networks analysis directly delivered from video content. The social networks are constructed from face detection and tracking of video content from the NHK News 7 broadcast, and enriched with segmentation and domain knowledge. Finally this enables a social analysis of the media representation of politics in NHK over a decade. After discussing related work in the next section, we will present our data in Section 3, with characteristics and preprocessing. Section 4 will then introduce the networks we have extracted. We present in Section 5 social analysis insights of the politico-media landscape of the 12 years of NHK supported by the social networks derived. Because this paper presents preliminary results, we will discuss our observations and future works in Section 6 before concluding.

2 Related work

Our system focuses on faces detected in news video. for which new deep learning approaches are very promising: [7] even reaches better-than-human levels of precision in face recognition, although face tracking in videos is not addressed. We use a simpler detection approach inherited from [8] available off-the-shelf but provide face tracking in return.

Many interesting works approach news analysis in a data intensive way, from text analysis. One of the most impressive approach on exploiting news data comes from [9] in which the authors combine news topic threads and demoscopic information to retrieve videos and generate a new summary video to explain prime ministers' resignations. A Natural Language Processing framework is designed in [10] to characterize news providers, programs, and newsmakers over many channels. The work from [11] is a notable effort in creating networks from news data. They generate *actor-action-object* networks over years of news, with great potential for building narrations and understanding of a news landscape.

News and media have been also material for social studies and especially political studies [12]. Questions such as the media-induced "presidentialization" effect are widely studied [13] and the case of Japan was raised from the Koizumi's unusual popularity [14]. This question is still under investigation [15], but none of these works approach the matter from a network perspective.

The relevance of network modeling for social and political studies does not need to be proven anymore [16], it is actually an old practice [17]. Beyond classical metrics [18], networks have been shown to be efficient for topic and concept analysis [19], and multiplex networks have been explored to analyse news data [20]. In particular, character networks have been broadly analyzed from literature [21], from TV dramas [22], and there is even a website dedicated to the social analysis of *Game of Thrones* [23].

News data has been one main target for visual analytics applications. Although we do not yet address visualization in this paper, the visual analysis of the networks and timelines was essential to conduct our social analysis. The following examples are all inspiring models to orient our analysis. Visual analysis is brought to help exploring large trans-media news as in [3] and [24] from which not only text but

also visual information is used. Faces are also used in the case of [25], which fuses many criteria and modalities to support user’s exploration of stories in the corpus, and introduces a network of topics, similarly to [26]. Analysis derived from large scale data [27] also includes political figures co-occurrence analysis represented as networks.

3 Face detection and tracking

To better interpret the data, we need to draw an accurate picture of what we are looking at. This section details all the preprocessing that is done before computing any social network. After describing the data, we introduce the segmentation of news, the face detection and tracking, and some domain knowledge.

3.1 Description of the data

Our video dataset consists in the daily-collected NII-TVRECS archive from [28]. The capture covers a period between March 17, 2001 and February 27, 2013; of the 4,366-day long period, 4,259 news programs have been collected cumulating about 2,102 hours (6.7 TB of video) from the Japanese public NHK channel’s daily News 7 broadcast. The few missing captures concerning mostly the beginning of the time period are due to the capture system setup. Most of the programs usually last 30mn and only a few of them fall below or beyond this format (news programs may be shorter on Sundays, or longer during commemorative events, see Figure 1(a)).

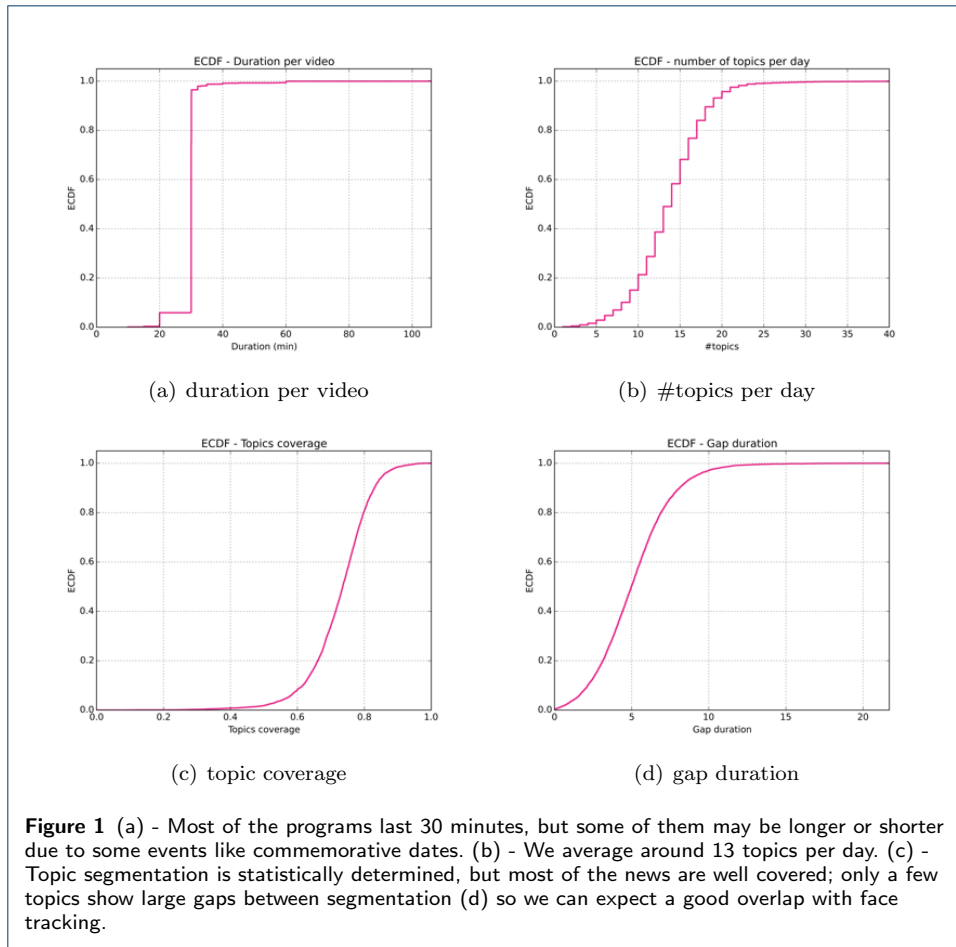
3.2 News segmentation

News are specific programs that can be segmented in different news topics. We thank the authors of [29] for providing us data in which news topics are derived from textual information (captions) synchronized with the news. In summary, a topic boundary corresponds to a point between sentences where the keyword distribution is distinct between preceding and following windows of sentences.

As a result, we have a segmentation of the news by topic, based on semantics analysis (segments in *red* in Figure 2-1). To one *topic* then corresponds what we would commonly call a *news segment* (both terms may be used interchangeably in this paper). Although we do not have the semantic information of these topics (yet), this gives us time boundaries for news segments which will turn useful for analyzing people’s apparition on screen. Overall, taking into account the differences of lengths among programs, this summarizes in a distribution of an average 13.7 news topics per day ($\sigma = 4.3$) as illustrated in Figure 1(b).

However, this segmentation implies that topics are not consecutively segmented, and gaps may occur between two consecutive topic detections. So topic detections cover on average 72.6% of the shows ($\sigma = 0.10$), because the beginning (head news summary) and end of the news (weather reports) are ignored by the topic segmentation (Figure 1(c)). In addition gaps between news segments average 5.1 seconds ($\sigma = 2.4$) and can stretch up to 20,7 seconds (see Figure 1(d)), in which faces may still be detected.

To help the face-tracking process, a standard shot segmentation derived from color histogram thresholding is provided (segments in *blue* in Figure 2-1). This is a contiguous segmentation, without in-between gaps corresponding to video cut editing.



3.3 Face-track extraction

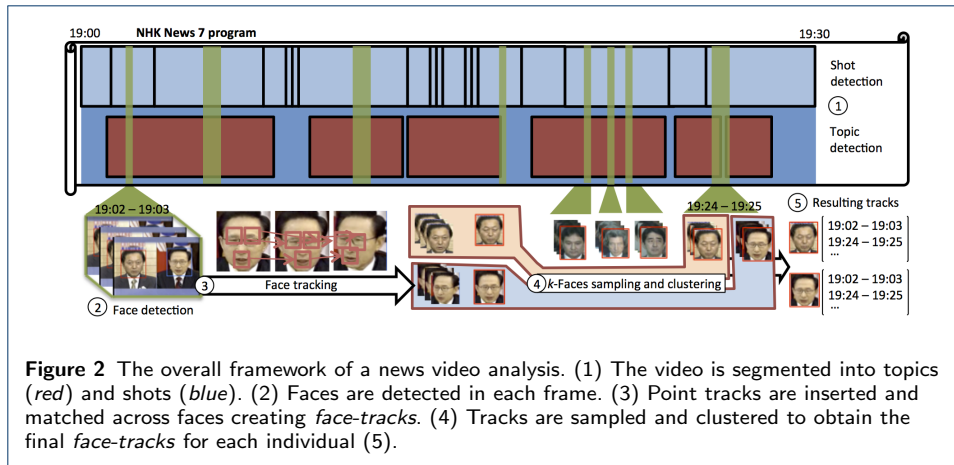
Now we can extract faces from the video shots as in [30], designed to be applied on large dataset. The extraction is done in the following steps (illustrated in Figure 2, steps 2 to 5):

Detection. We first apply a detection of faces in all images, using off-the-shelves techniques such as the Viola-Jones face detector [8]. Multiple faces can be detected per image, different detections of a same face may also overlap.

Tracking. We now need to group together the detected faces of a same individual into one single face-track. This is done by generating feature points within detected faces – a point track is a same point identified across different consecutive frames. Point tracks from the KLT point tracker [31] are matched from a starting face with the following faces given their temporal order and this results in multiple face instances regrouped in face-tracks. There can be multiple face-tracks across the same video.

Sampling the face-tracks. For each face-track, we create a *mean face* that is a representative face in the VGG feature space [32], based on the *k-Faces* method [30]. The *mean face* is a mean point in the feature space described by *k* sampled faces.

Matching the face-tracks. Face-tracks can finally be matched based on their *mean face* euclidian distance in the feature space.



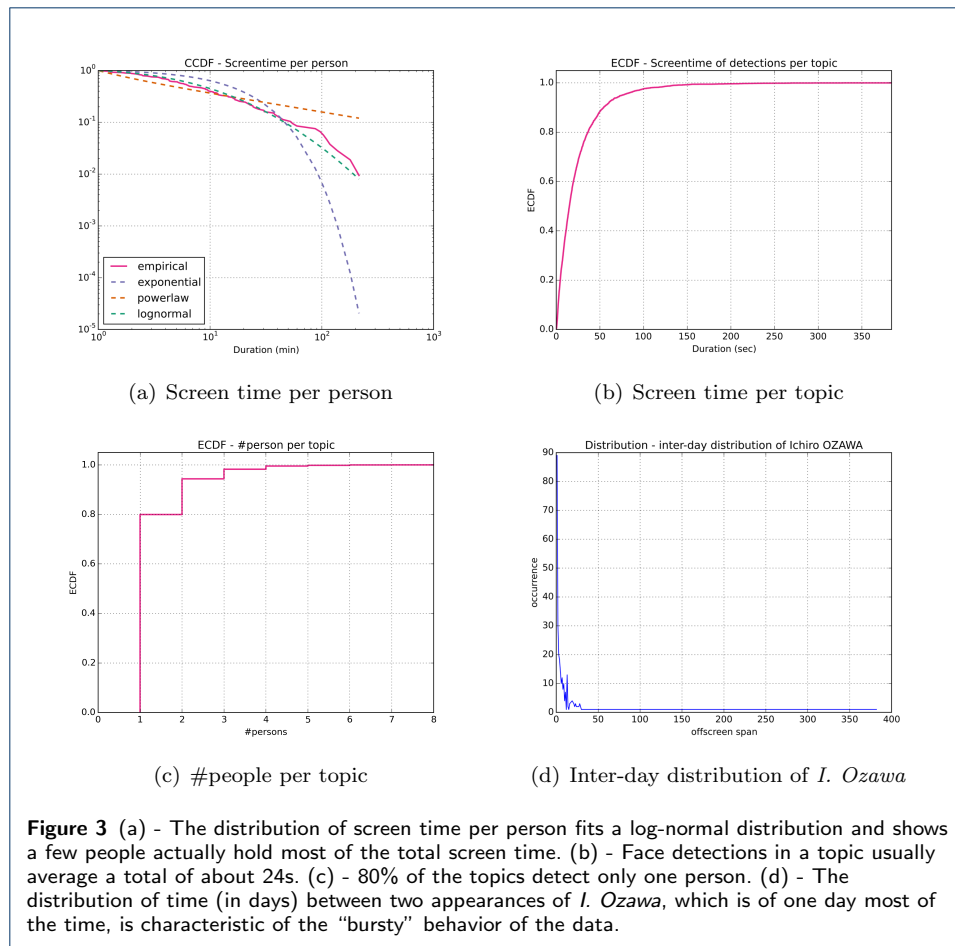
The whole process has detected over 30 million faces and 174,778 face-tracks were extracted. We now need to identify and recognize groups of face-tracks, so clustering naturally appears as the following step. However, clustering implies many new issues that we have not yet addressed this work (but we include this goal as part of our future work in Section 6). Yet, we can still use a different approach to construct our networks, that is of face retrieval.

The faces of 139 characters have been annotated during the evaluation campaign of [30] giving a ground truth for retrieving matching face-tracks. These faces are the faces of well known people among the Japanese media scene, including celebrities and politicians (Japanese and international), for which we had the highest precision of retrieval and identification. In total, over 5 thousands face-tracks were annotated, and 16,714 face-tracks of the 139 different characters were retrieved. This corresponds to 2,984 days of news program over the whole archive having matching people, covering a total of 36 hours of face tracks.

The coverage of the face-tracks averages 2.4% of a program ($\sigma = 2.6$), reaching the maximum of 38.9% of a program. The ground truth has been provided during the 2010 period, for which the face-tracks appear slightly denser (3.3% on average). The average screen apparition per person is 15.7 min ($\sigma = 32.5$) but there is a lot of variation between people (actually it fits a lognormal distribution Figure 3(a)), and a few people seem to hold most of the screen time (Table 1).

As a result, the tracks work as follows: everyday, we have a news broadcast, and every broadcast contains news segments (topics) and face-tracks of different people. We then observe an average of 23.8 seconds of cumulated detections per topics ($\sigma = 28.9$, with a maximum of 383.0 - Figure 3(b)), with in average 1.28 people detected per topic ($\sigma = 0.66$, with a maximum of 7 - Figure 3(c)). Thankfully, this shows that we can reasonably expect people to overlap across topics, although 80% of our news segments do not show more than one person detected (see Figure 3(c)).

Looking closer at the distribution of inter-day occurrences of people in topics, we can see that most of them appear on screen on a daily basis, with bigger gaps then (Figure 3(d) is an example). This is sometimes referred as a characteristic of “bursty” data [33], meaning that, over the whole period of time, there is a lower probability for two people to be detected together than random, making these links

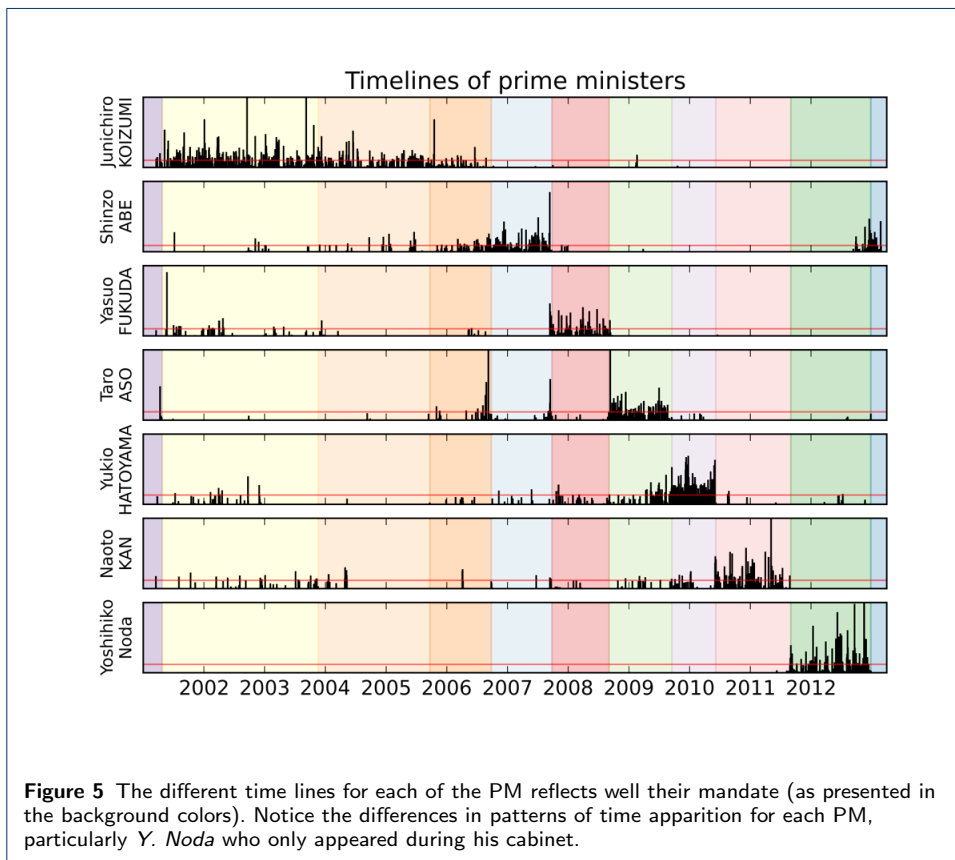
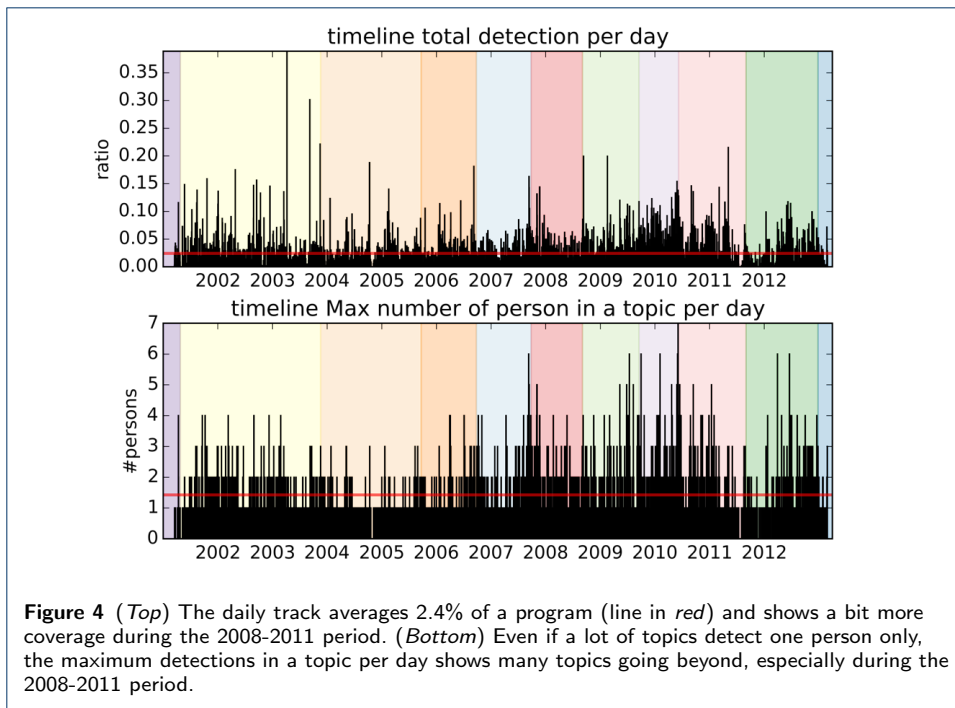


especially interesting. We can also notice that most of the people we are tracking seem to take part in similar topics during the 2008-2011 period.

3.4 About the 139 people

Some background information is necessary for a good understanding of this news data. With a little domain knowledge, we have classified the 139 characters identified into 9 categories depending on what brought them under the light of news: *Politics* (71), *Sports* (27), *Culture* (11), *Business* (7), *Imperial family* (5), *Journalism* (4), *Religion* (3), *Law* (3), and *Other* (5). Additionally, we have enriched them with their country: 96 individuals are from Japan, among which 39 political figures and 22 athletes.

Incidentally, the *Politics* class includes 23 international leaders (presidents, prime ministers...). The Japanese prime ministers – hereafter referred as PM – governing during the whole period of capture are of course represented, allowing us to create time frames covering their cabinet(s). *Yoshiro Mori* was the first PM in the timeline, but his mandate only covers a few weeks from the beginning of the capture, so himself is not included in the people subset. Finally, we obtain 11 time periods (Figure 9 details them in chronological order, note that the numbers following a PM’s name represent different cabinets formed by the same PM). Figure 5 and Figure 4 reflect these different periods as colored section in their background.



Although Pearson’s correlations between the three measures (*Screentime* S , *#days* D , and *#news segments* T – per person) are very high ($S - D = 0.95$, $S - T = 0.96$,

and $D - T = 0.97$), we can use the ranking of the top 10 percentile to extract people of interest (as presented in Table 1).

Person	Screentime	#Days	#News segments
Junichiro KOIZUMI	215	516	523
Yukio HATOYAMA	179	330	368
Ichiro OZAWA	138	294	304
Naoto KAN	118	243	249
Shinichi TAKEDA	110	337	491
Shinzo ABE	103	281	298
Yoshihiko NODA	97	195	245
Taro ASO	87	181	187
Yasuo FUKUDA	60	144	139
Seiji MAEHARA	55	116	126
Takafumi HORIE	52	116	121
George BUSH	45	86	93
Kazuo SHII	42	116	98
Sae NAKARAI	40	317	190
Katsuya OKADA	37	115	109
Hideki MATSUI	26	126	73
Wataru ABE	22	84	104

Table 1 The top 10% people among the different criteria (in bold, the Prime Ministers, and top 5 scores of each criterion). Ichiro Ozawa is the only person taking a top 5 position who has neither been a Prime Minister, nor is a news presenter.

A background checking gives us supplemental information explaining their occupation of screen space. We can already discard of our analysis the journalists of NHK (*S. Takeda* and *S. Nakarai*) who have an obvious aggregating force due to their role on TV. Out of the obvious known figures and the aforementioned Prime Ministers, *I. Ozawa*, *S. Maehara*, *K. Shii*, and *K. Okada* are famous politicians, and we will talk about them in greater details in Section 5. *T. Horie* is a businessman, *H. Matsui* is a baseball player, and *W. Abe* is active on the music scene.

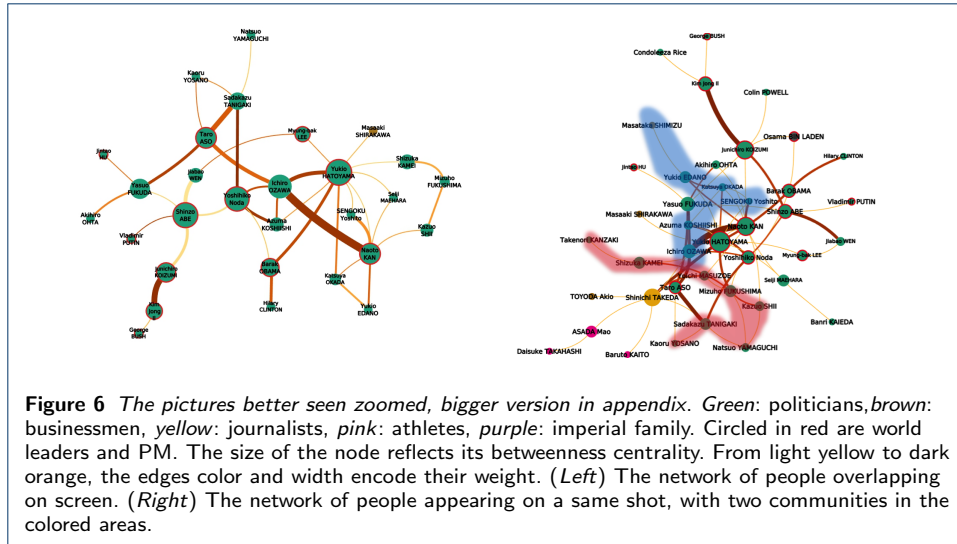
4 Different networks

We will now define and present our different networks with their preliminary analysis. We are focusing in this sections on the different ways we can create links of social networks directly derived from the analysis of the video data. Most of the following networks use the people as the same set of nodes, but with different families of ties.

4.1 Network of people overlapping on screen

Our first network connects two people when two face-tracks overlap in time. This means that we create a link between two people when they have been detected *simultaneously* on screen. These links are enriched with the screen duration of the overlapping of tracks as weights.

This network presents 35 nodes and 44 edges, with a main connected component of 29/41 (Figure 6, left). This connected component is only composed of politicians, with one business person (*M. Shirakawa*, connected to *Y. Hatoyama*). It is worth noting that *J. Koizumi*, the top individual among all other metrics, only presents here a degree of 2. Four nodes stand out in terms of betweenness centrality (*S. Abe:0.16*, *I. Ozawa:0.14*, *Y. Hatoyama:0.18*, and *Y. Noda:0.15*, with the rest of the dataset below 0.09), and 2 nodes in terms of degree (*Y. Hatoyama:10* and *Y. Noda:7*), however no clear convincing cut of communities is shown by Louvain's algorithm [34].



A few links stand out in terms of screen duration (over 1000), connecting: *Y. Noda* and *S. Tanigaki*, in 2012, *I. Ozawa* and *N. Kan* in 2003, 2006, and 2010, *Y. Hatoyama* and *I. Ozawa* in 2006, 2010, and 2012, *Y. Hatoyama* and *B. Obama* in 2009, *J. Koizumi* and *Kim Jong Il* 2002, *V. Putin* and *S. Abe* in 2012, *T. Aso* and *Y. Fukuda* in 2009. When looking at the number of days in which two different people appear together, we can notice stronger links between: *S. Tanigaki* and *T. Aso* in 2006, *J. Koizumi* and *S. Abe* in 2002, and *H. Clinton* and *B. Obama* in 2008.

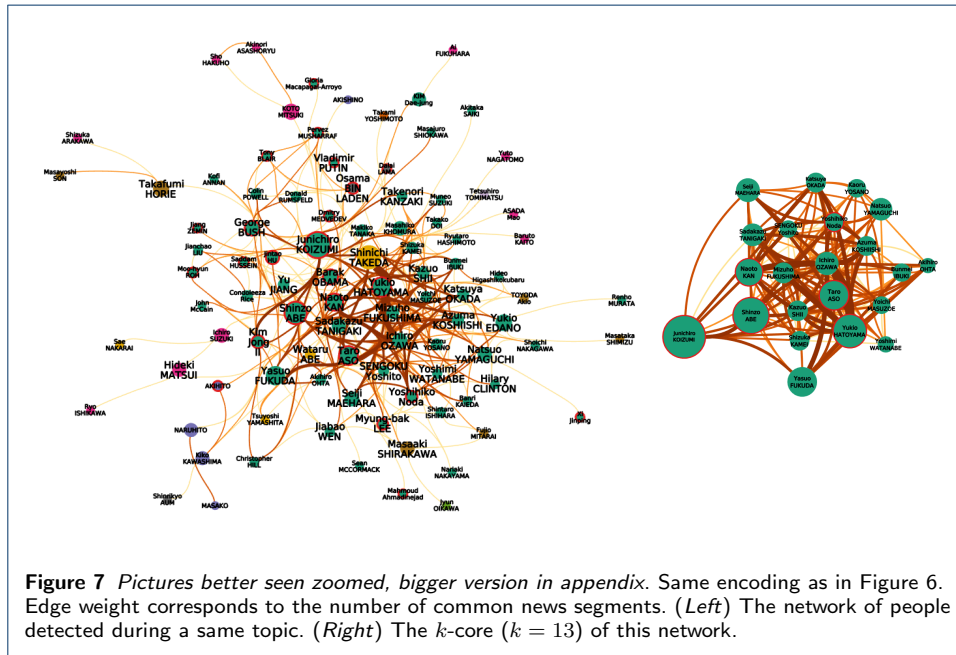
4.2 Network of people appearing in a same shot

This second family of ties defines links between people appearing in a same shot (*i.e.* an uncut segment of video). This network roughly extends the previous network, with the difference that people do not need to appear on screen together. Because shot duration greatly varies depending on the cut of the video, we cannot use it as a meaningful metric to weigh edges, instead, we will consider the number of different days that include these shots.

The network (Figure 6, center) presents 49 nodes for 75 edges with a main component of 41/71. The maximum k -core ($k = 3$) [35] presents a very intricate subnetwork of 18 nodes (Figure 6, right). It includes the PM, and the main anchorman (*S. Takeda*), later referred as the ‘main actors’. All the other nodes are politicians, including *I. Ozawa*. Getting their full list and description may go beyond the scope of this paper, but it is interesting to notice that *N. Yamaguchi* stands out as the only politician not directly connected to any of the PM. The main component presents a wider range of types of people, including 3 athletes, 3 business people, and *O. Bin Laden*. A Louvain segmentation does not present a clear cut of denser subgroups in this network. If we remove the ‘main actors’, we can interestingly observe two communities of politicians (the colored areas in Figure 6, center), one centered on *M. Fukushima* and *N. Yamaguchi*, and the other one on *K. Okada*. However one should carefully interpret the meaning of these links given the low amount of common shots (at most three).

Three edges stand out with links displaying between 5 and 8 days of connections, *T. Aso* and *S. Tanigaki*, *N. Kan* and *Y. Hatoyama*, *J. Koizumi* and *Kim Jong Il*. If we consider links connecting two people over one day only as ‘casual’ and discard them, we can reveal a network of stronger ties of people with ‘recurrent’ interactions (23/26). In this network, *I. Ozawa* displays the highest betweenness centrality, followed then by the different PM.

4.3 Networks of people appearing during a same news segment

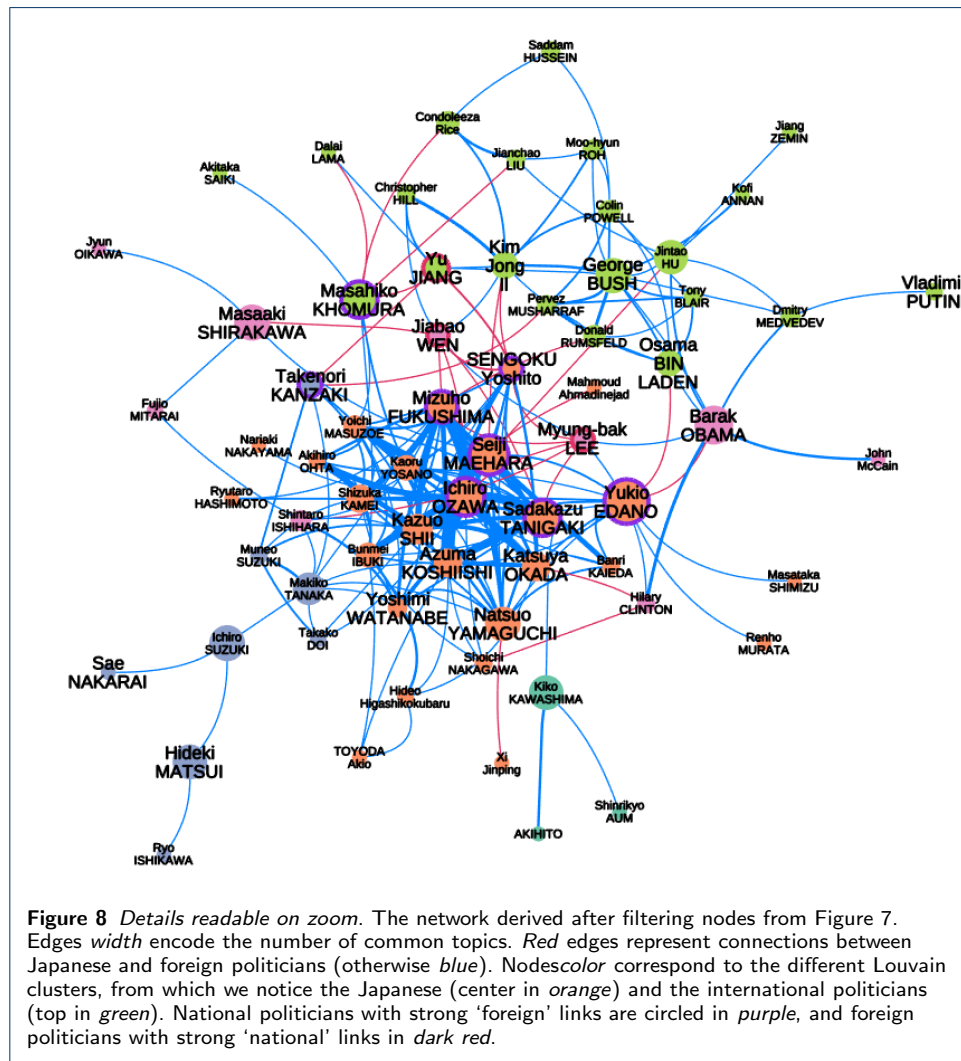


The following network connects individuals when they have been detected during a same news segment, based on the topic segmentation described in Section 3.2. This means that two people are connected when they took part of a same media event. The graph connects 107 people over 507 links with a main connected component of 96/499 (Figure 7, left). This graph presents characteristics closer to complex networks with a long tail distribution of node degrees (actually fitting a lognormal distribution).

Knowing that co-detection during a news segment is the reason linking nodes, we should first remove the journalists – occurring a lot in the dataset, in order to focus on other people’s interactions. The resulting graph presents a maximal k -core ($k = 12$) gathering 15 Japanese politicians and the 7 PM in a subgraph $G'_{k=12}$ with a density $D_{G'_{k=12}} = 0.79$ (Figure 7, right).

A degree and centrality analysis will bring focus to the same people identified in the previous networks. To go beyond, we will look at the graph without the ‘main actors’, leaving 67 nodes for 221 edges. This graph clearly presents community structures, and by running a Louvain algorithm, we obtain a very interesting clustering result. The two main partitions (in *light green* and *orange* in Figure 8) clearly present international politicians and national politicians (respectively). We are now able to spot the non-PM Japanese politicians who played an active role in

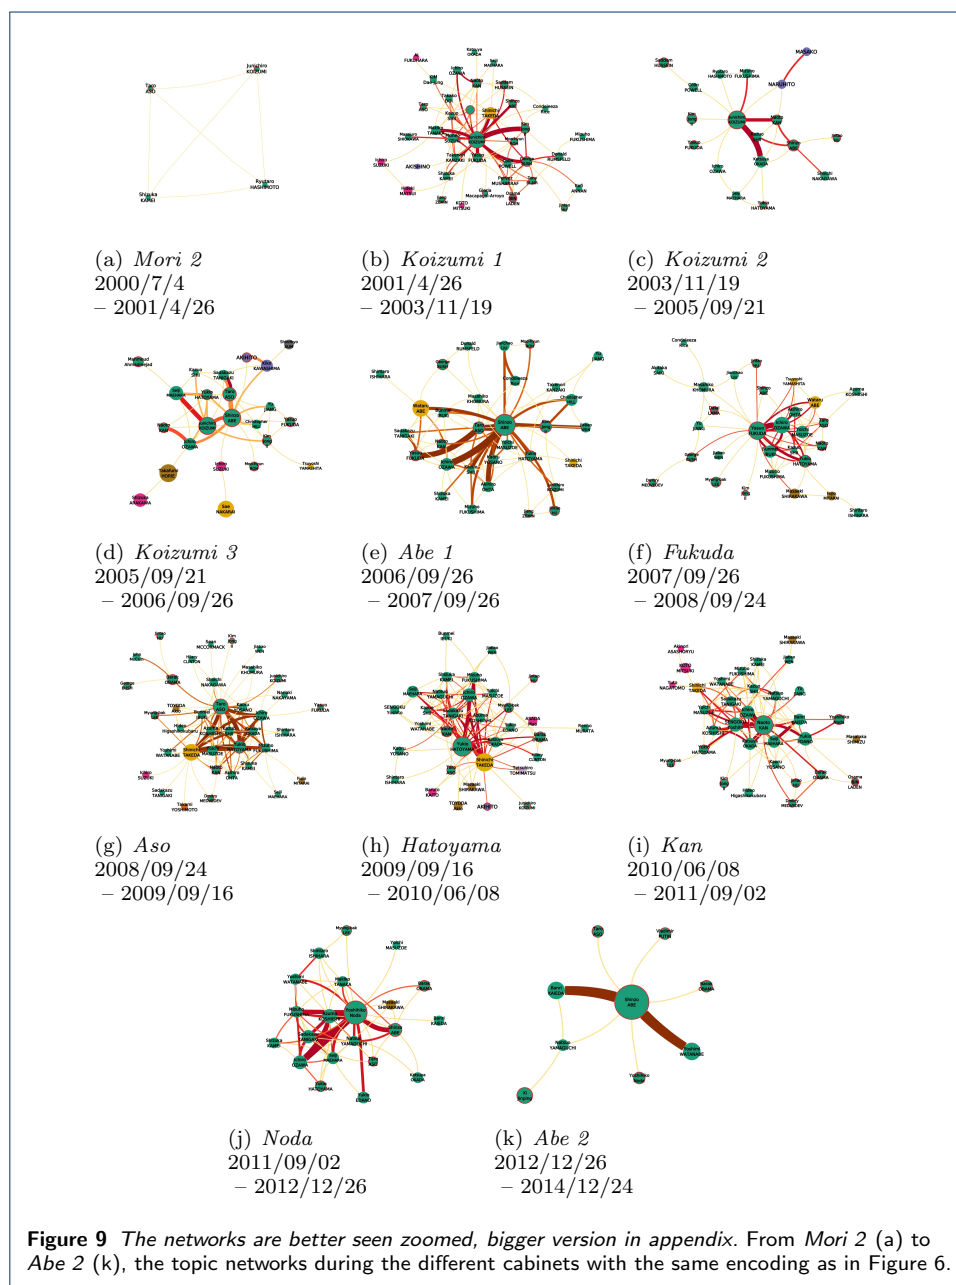
international matters by highlighting them (circled in purple in the Figure 8, mostly at the right frontier of the orange community). We do so by counting the number of their ties with international representatives and threshold them based on their cumulative probability distribution [36]. As a result, we find *Y. Edano, S. Tanigaki, S. Maehara, M. Fukushima, Y. Sengoku, I. Ozawa, T. Kanzaki, M. Khomura*. With the same process on the other side, we can identify (circled in red) *Yu Jiang, Jiabao Wen, and Lee Myung-bak* as having redundant apparition on topics with national politicians. The case of *Lee Myung-Bak* seems to have particularly raised a great interest among national politicians, totalizing 5 connections.



4.4 Time slicing the common segment network

The topic network has a dynamic multiplex characteristic – although we do not fully exploit this characteristic in this paper. A link is established between two people when matching in different points in time, which means we have virtually an individual link for each matching at different moment in time. Looking at the network in different timeframes will result in different arrangements of the links

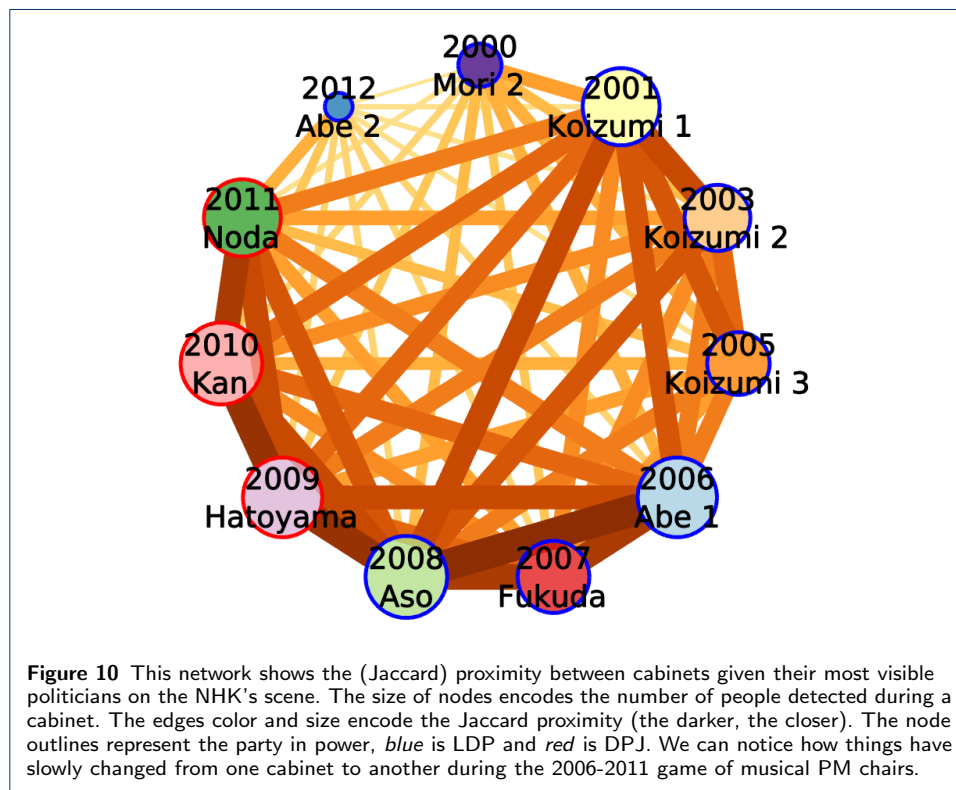
between nodes and different weights if these multiple links were to be collapsed in one single interaction. Thus, thanks to well defined periods of time corresponding to PM cabinets, we can use topic segmentation as a support to observe not the overall network but each slice involving the people's interactions over the different cabinets (Figure 9).



Before we tackle the political analysis in Section 5, we can quickly compare the political landscapes of each cabinet. To do so, we pick out the top 3 Japanese politicians in ranking of centrality and number of news segments, who are neither a PM nor have been detected during the preceding cabinets. In total we have collected

21 prominent politicians, which will be used to compare cabinets one to another. This creates a vector of all politicians per cabinet.

Based on these vectors of 21 (+ 7 PM) politicians, we can finally estimate a rough (Jaccard) proximity between cabinets as shown in Figure 10. The periods from *Abe 1* to *Noda* known for the series of resigning PM, shows the highest proximity one to another, and interestingly to *Koizumi 1*. However, Koizumi's two following cabinets appear very different, suggesting that he set a very different media/politics scene during this time.



4.5 Preliminary observations

Before we solely focus on a political analysis of the news landscape, the exploration of these networks led us to some understanding of the media/politics scene presented by NHK News 7. Based on this data, together with the knowledge of people, we can confirm that the different PM stand out like no one else in the NHK news. They can be directly identified in all aspects of the data: first, purely quantitatively speaking, they occupy most of the media scene during their own cabinets; then, in the different networks, they also occupy a very central place; the different time-related analysis makes it especially obvious during their cabinets.

We also learn by looking at individual PM: most of them show some level of activity before their mandate and we can observe two opposite cases. On one side, Abe is actually more central than Koizumi himself during *Koizumi 3* (Figure 9(d)). On the other side, Noda came 'out of nowhere' before becoming PM (Figure 5). Additionally, despite of Hatoyama and Aso appearing quite strong nodes in the

different networks, they have never been detected on screen together (Figure 6) even if they were heading two consecutive cabinets in period of time where the media/politics scene of consecutive cabinets is very similar – maybe because they are the leaders of two opposite parties.

Strikingly, one very particular politician comes out all along this study, *I. Ozawa*, who is (in)famously known as the “Shadow Shogun”. Getting into the details of Ozawa’s role in the Japanese politics is a fascinating work on its own [37], but put in short, after being leader of the opposition, he is known for all the connections and roles he has played behind the scene, building alliances and often changing side – although never he became PM.

Another very interesting point which is worth noting concerns the Imperial family. The Japanese Constitution forbids the Imperial family to take any part in politics, and observing the links surrounding the members of the family are of high interest to survey their actions. Our system finds very little connections (*purple* nodes in Figure 9(c), (d), and (h)): they mostly concern the revision of the Imperial Household Law because of the issue concerning the succession to the Imperial Throne.

5 A tool for social analysis

We want now to investigate this data with the sociologic perspective of studying the political relationships on the media scene over these 12 years. After focusing our data on political activities over time-sliced networks, we introduce political affiliation data, revisit our networks and present our insights.

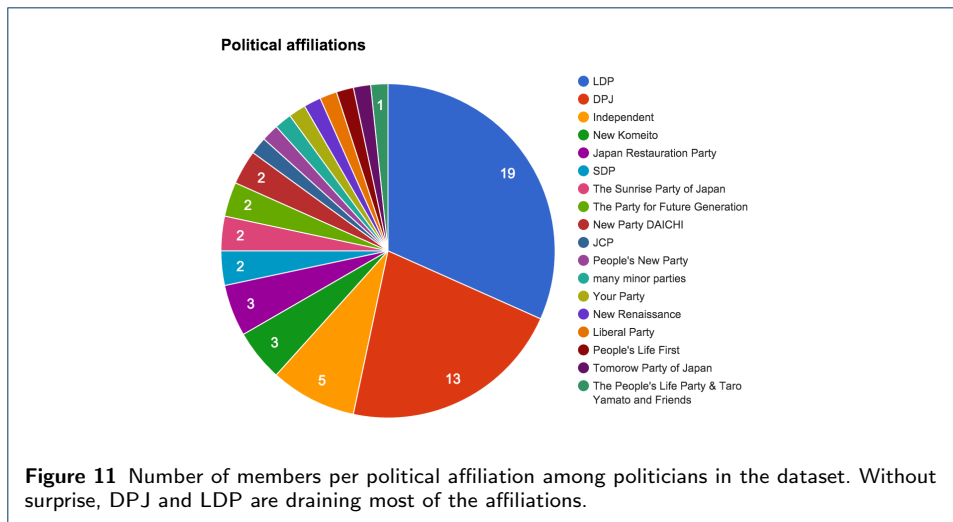
5.1 Political affiliations

We first need to subset the data: because the *Mori* and *Abe2* cabinets have only partial data, we will discard those from our analysis. We will also focus only on the national and international politically involved actors, reducing to 75 tracked individuals (including 40 Japanese politicians).

For each of the 40 Japanese politicians, we have manually collected their political affiliations over time. We collected in total 18 different parties (including *independent* politicians and the many *minor parties*) represented among the 40 politicians. As illustrated in Figures 10 and 11, the two main parties of the Japanese politics are the *Liberal Democratic Party (LDP)* – ruling the cabinets from *Koizumi1* to *Aso* – and its main opposition the *Democratic Party of Japan (DPJ)* – ruling from *Hatoyama* to *Noda*.

As we can see from Figure 11, the LDP and DPJ accumulate most of the affiliations and many small parties only have one or two members tracked in our dataset. A few politicians have created independent parties at different moments in time. Different coalitions of opposing politicians appeared especially in the end of Noda’s cabinet when it was foreseen that DPJ would not win the next election. This explains why we obtain in total more affiliations than we have politicians.

We are interested in studying how the balance of power reflects in the NHK news. In modern democracies, there are often two main parties: a party in power and its main opposition. In the case of Japan, the leader of the party in power is the Prime Minister, and the leader of the main opposition (hereafter *OL*) is the president of the main party not in power (which is here either DPJ or LDP). However, this



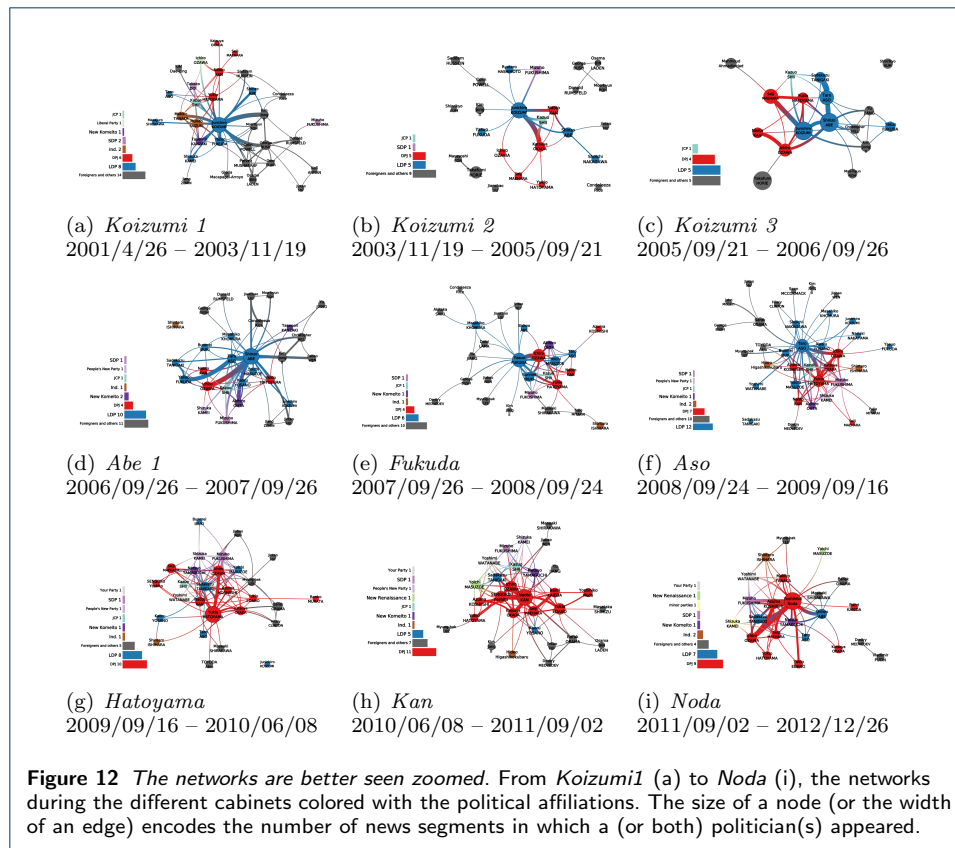
leader changes over time such as more than one politician may lead the opposition during a cabinet. To study the opposition, we created an artificial node entitled “*Opposition leader*” (visible in Figure 13). We extracted from the different persons representing the OL all their detections during their leading mandate and unified them altogether in this abstract entity. This allows us to question the place of the *Opposition leader* in different the political networks.

5.2 Political networks

Given this information, we can color the nodes of the previous networks using their political affiliations Figure 12 – the category *foreigners and others* includes Japanese business people and religious figures who could be involved in political matters. The coverage (number of news segments in which a politician has been mentioned) encodes the size of a node and of a link (for which the co-occurrence is measured). A longitudinal analysis is only permitted using the time sliced topic networks. We can notice that the editorial choice of the NHK seems to well follow the governmental trend, *i.e.* we detect more LDP members when LDP is ruling, and DPJ members when DPJ is ruling.

A quick look at the different networks bring us instantly some interesting insights:

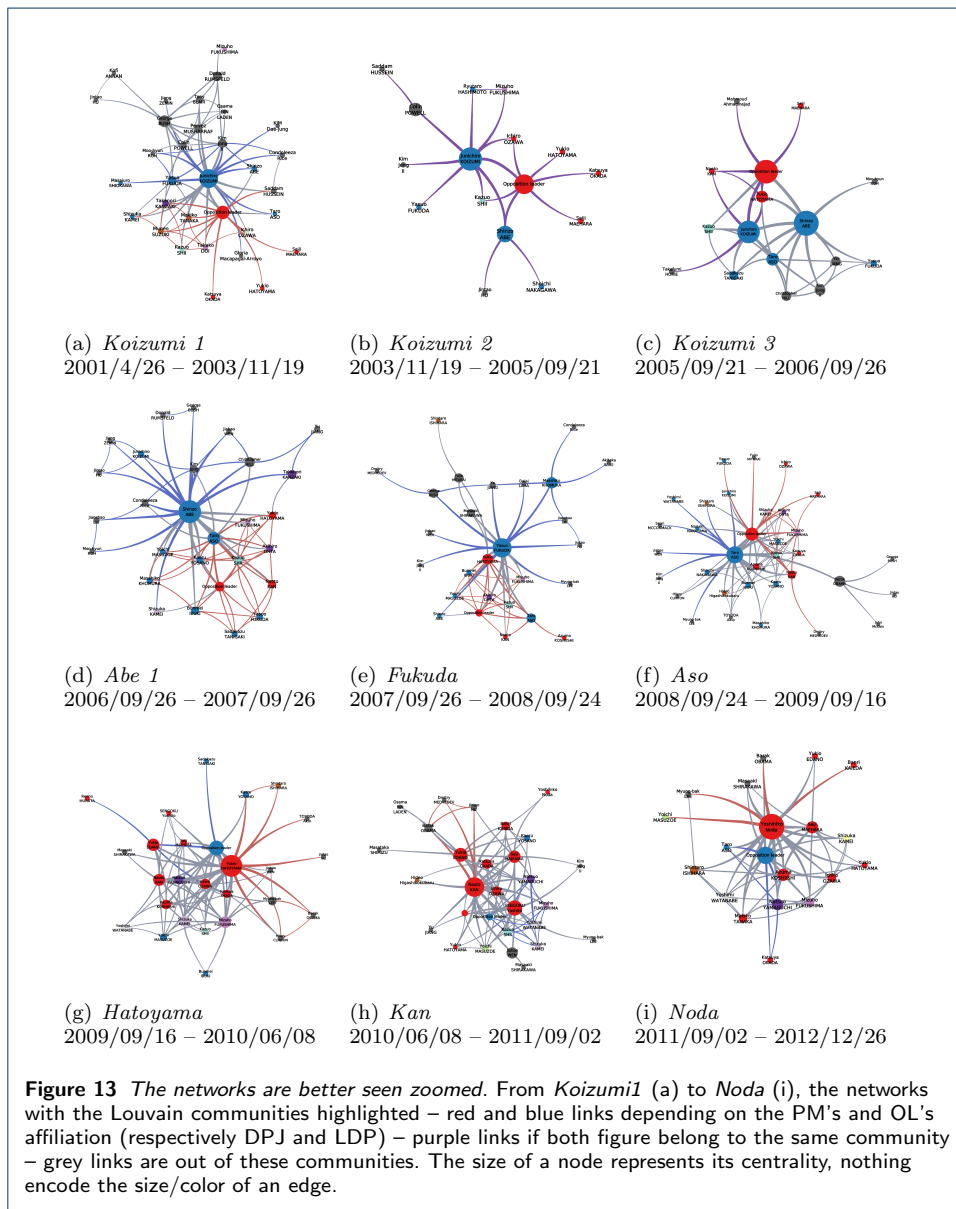
- During *Koizumi1*, foreign affairs appear to show important matters as the PM himself displays many connections with foreign politicians. It is not so much the case during Koizumi’s two following cabinets, in which opposition members appear with more intensity (*i.e.* more topic coverage as nodes and links). The second cabinet of Koizumi shows a smaller network, but with a strong coverage of Koizumi himself. The third cabinet of Koizumi clearly puts Abe in a stronger position in the network. He shows strong links with opposition figures and even more connections with foreign leaders than the PM himself.
- Abe’s first cabinet shows a real dominance of LDP members with many connections from the PM himself with foreign figures. For the first time, we can notice a rising position of the opposition leader I. Ozawa (and his strong connection with the PM).



- Fukuda's cabinet also presents a remarkable connection of himself with foreign leaders. M. Khomura is in a central position with foreign affairs (which is expected since he is the Minister of foreign affairs during this period). Foreign affairs apart, we can observe I. Ozawa holding an equally central position as the PM among Japanese politicians.
- The central position of PM Aso during his cabinet cannot be doubted. We can witness many links with foreign leaders, LDP politicians, and the multiple members of the opposition. The opposition gives a large role to DPJ, but is not focused only on I. Ozawa – having a less central position – to the benefit of Y. Hatoyama and some other members.
- Hatoyama's cabinet bring a clear dominance of the DPJ members, with many other parties represented. I. Ozawa seems to play an important role during this cabinet but the overall topology of this network is denser than the previous one, even in relationship with the foreign figures. This means that the discussion of foreign matters is not focused on the PM, but discussed by all politicians.
- Kan's cabinet presents similar features than Hatoyama's with a denser structure between DPJ members, and very limited LDP members (some changed party during this period). Interestingly, there seems to be very limited links between the OL and the PM during this cabinet. We should not forget that Kan's cabinet has been witnessing the tragic Great East Japan Earthquake

and its consequences became the focus of the media scene, rather than the political games.

- Finally, Noda’s cabinet resembles more of Abe’s or Koizumi’s, with the PM in a very central position. We can only notice that the two main opposition figures (S. Tanigaki and S. Abe) have quite separated positions in the network.



5.3 A small analysis of the political networks

With all uncertainties taken into account (which will be discussed in Section 6), we want to use the generated networks as support for social analysis. Inspired by recent works in social science [15], general questions drive our study: *How is the opposition represented during the different cabinets? Can we find traces of “presidentialization” of the politics in the media?* These questions can of course be tackled quantitatively

by studying the coverage of the different actors, but can we learn something else from a social network analysis perspective? To investigate these questions, we propose to observe two phenomena: firstly we can study the place of political leaders in their social networks, secondly we can have a thorough look at the links between the two political parties.

We can then start by focusing the study on the different metrics of the two main nodes (the PM node and the OL node) and their evolution in time (although we will juxtapose different time frames corresponding to cabinets). Relative to their position in the whole network, three quantities especially interest us: their degree centrality (*i.e. how many politician are directly connected to these individuals?*), their betweenness centrality (*i.e. how many people need to go through this politician to access others?*), their closeness centrality (*i.e. how easy is the access of these politicians to the whole network?*). Table 2 summarizes these quantities. We applied the Louvain community detection to each time slice. In most of cabinets, the PM and the OL belong to different partitions (Figure 13). The size of their respective communities is not a meaningful metric due to the varying size of the networks, but we can notice that both PM and OL reach usually high betweenness centrality in their local community.

Although we cannot measure a clear trend such as a growing importance of the OL over time, we can notice a couple of facts. Most often, PM and OL get the top structural positions in the different networks. Except maybe for the PM, more coverage for a politician does not mean higher position in terms of structural measures. We also have two remarkable cases. First is the case of *Koizumi3* in which the OL reaches higher centrality than the PM himself. This is due to the inclusion of a very broadcasted and connected S. Abe who tops all metrics. Then comes the cabinet of Kan: the main opposition does not seem very represented so the OL does not hold an important role. As we mentioned earlier, during this cabinet, NHK was probably more focused on reacting to the tragedy, probably displaying alignment of ideas rather than a diversity of point of views. Figure 14(b) reports the timeline of the news segments attributed to the PM and the OL during Kan's cabinet, and the last detection of the OL in our dataset only happened a few days before the tragedy.

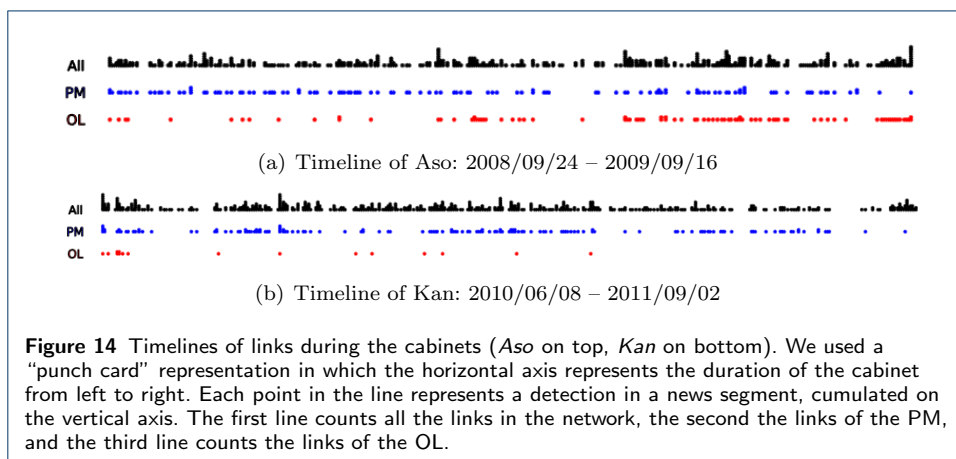
To study the interactions between the two parties, we can plot the share of news segments that involves each of the two main parties, alongside with the ones that involves both of them together (Figure 15(a)). It is expected to see most of the cabinets having the interaction of their own members put forward in the media during their own cabinet. This is especially true during all DPJ cabinets, Koizumi and Abe's first cabinets. Aso's cabinet displays the opposite behavior. Indeed, it became clear at some point that DPJ will win the following election. At this moment, the media increasingly reported the activity of the opposition (as we can see from the timeline in Figure 14(a)). This same behavior is confirmed by looking only at the activity of the main figures of both parties (Figure 15(b)) suggesting that most activities in the political networks are supported by these strong characters.

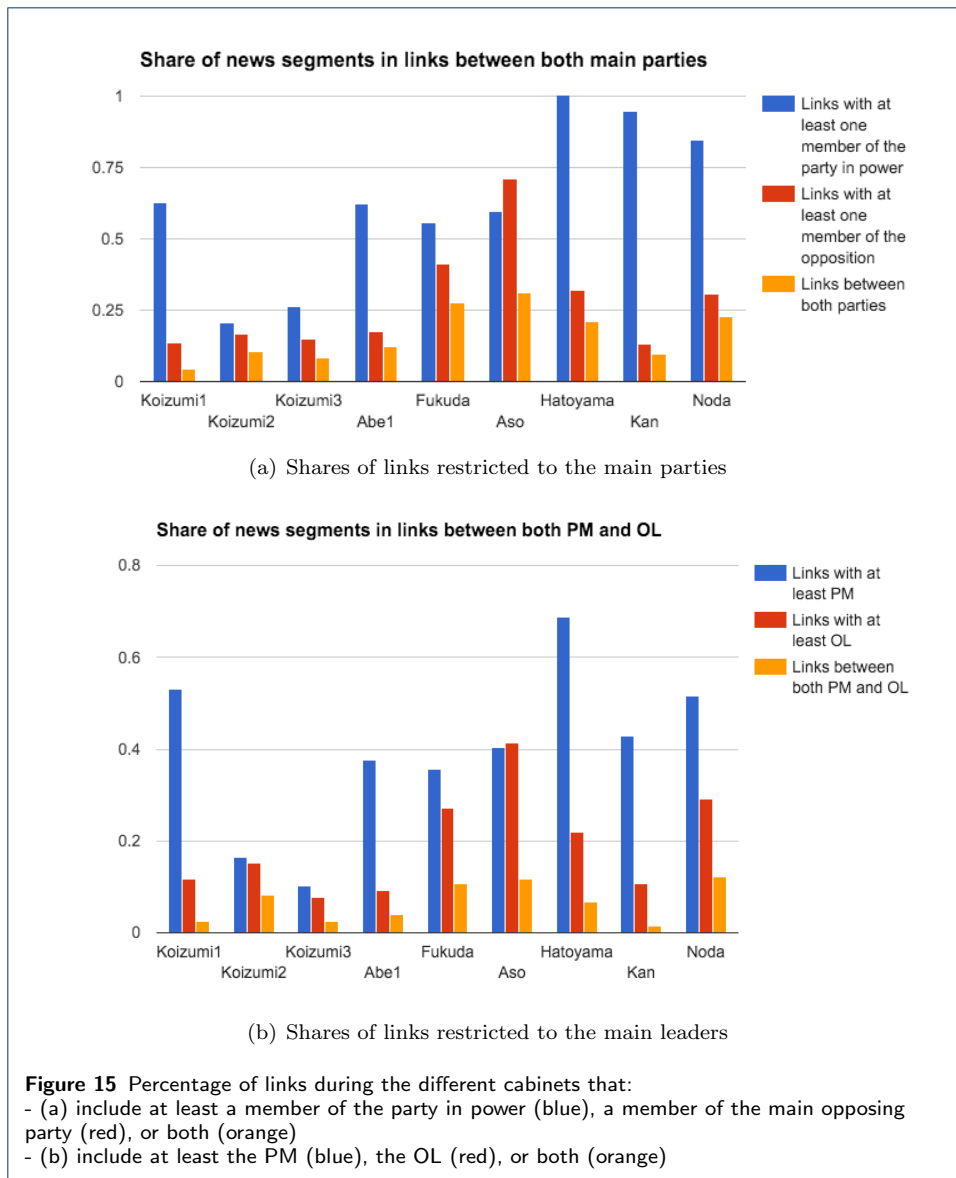
6 Discussion and future work

The different network views provide a powerful tool to understand the media situation, but we also need to draw the limits of the definition of these networks. As

Table 2 Structural measures of the Prime Minister (PM) and the main Opposition Leader (OL) during the different cabinets - *Coverage* is measured in number of news segments.

		<i>Betweenness</i>	<i>Degree</i>	<i>Closeness</i>	<i>Coverage</i>
Koizumi1	PM	0.172 (1)	25 (1)	0.026 (1)	346 (1)
	OL	0.067 (2)	18 (2)	0.019 (4)	44 (5)
Koizumi2	PM	0.079 (1)	10 (2)	0.048 (1)	118 (1)
	OL	0.051 (2)	11 (1)	0.042 (2)	42 (3)
Koizumi3	PM	0.065 (3)	10 (2)	0.048 (2)	46 (3)
	OL	0.073 (2)	10 (2)	0.043 (3)	36 (4)
	Abe	0.085 (1)	11 (1)	0.050 (1)	48 (2)
Abe1	PM	0.164 (1)	24 (1)	0.033 (1)	150 (1)
	OL	0.011 (4)	10 (3)	0.022 (3)	24 (2)
Fukuda	PM	0.202 (1)	20 (1)	0.030 (1)	75 (1)
	OL	0.008 (8)	10 (2)	0.022 (3)	63 (2)
Aso	PM	0.176 (1)	26 (2)	0.024 (1)	119 (1)
	OL	0.087 (2)	28 (1)	0.020 (2)	88 (2)
Hatoyama	PM	0.099 (1)	22 (1)	0.032 (1)	213 (1)
	OL	0.035 (2)	17 (2)	0.027 (3)	30 (4)
Kan	PM	0.090 (1)	22 (1)	0.029 (1)	150 (1)
	OL	0.003 (12)	11 (6)	0.021 (7)	16 (8)
Noda	PM	0.080 (1)	19 (1)	0.050 (1)	230 (1)
	OL	0.027 (2)	17 (2)	0.040 (2)	56 (2)





for now, the news segment association brings the most meaningful construction of links, even if no actual semantics has yet been introduced into the system. It is equally important to understand how the different pre-processing parts may have strong influences in later interpretations.

The screen overlap network has the strongest family of links in terms of social ties, but it is also the most subject to controversy in two ways. First, because of the bursty characteristics of our data, the limited but reliable subset of people, and parameters of our face detector make the amount of screen co-detections limited. Then, because many detections concern split screens, which in turn often means an opposition of ideas on a same subject, hence defining a sort of *negative* link – something we would like to investigate in the future. We want to distinguish this case from the screen co-occurrence, which holds the different meaning of people standing in the same room at the same time (Figure 16).

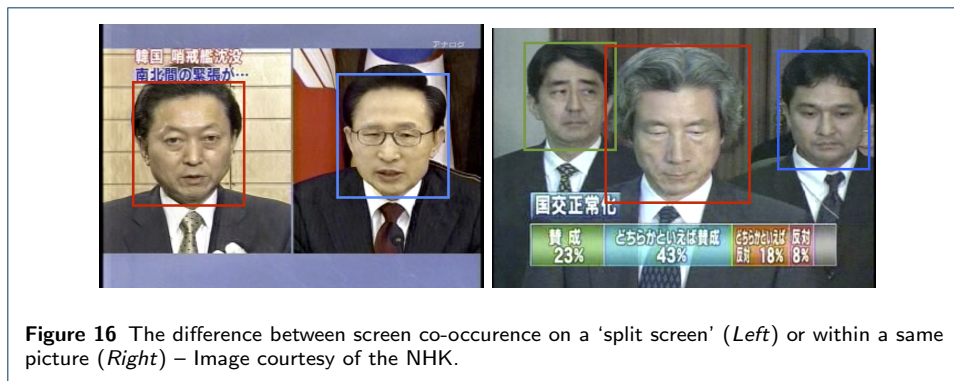


Figure 16 The difference between screen co-occurrence on a ‘split screen’ (Left) or within a same picture (Right) – Image courtesy of the NHK.

The same comment may also be made on the shot co-occurrence network, which finally extends the latter with a lighter meaning. For example, some shots occur behind the anchorman switching from one segment to another, sometimes leading to false positive links. Besides the system showing a good accuracy [30], some face occurrences may remain untracked, but we can still draw our conclusions thanks to the large period of time we observe.

This leads us to discuss the effectiveness of our system. Our contribution solely focuses on face-tracking because we have a long-term goal to use the screen time as a reliable measure. Our current experiment shows that measuring screen time is yet limited, and we need to discuss the accuracy of our system to understand this limitation. Our face-tracking reaches about 61% precision on the NII-TVRECS dataset [28], for which we estimate an average of 37% recall [38, 30]. This measures the quality of the *retrieval* of people’s face-tracks, keeping in mind that we had performance constraints so we could applied face-tracking on the whole database. Because we are focusing on annotated face-tracks, we made sure that the subset we are exploring for this paper reach 100% precision. However, our system still misses a lot of face-tracks due to our low recall.

The recall depends on the multiple steps described in the first part of this document.

- It first depends on the performances of our face detector, the Viola-Jones detector [8], which performs up to 88% accuracy depending on conditions and parameters. We chose this detector because of it is rather fast to execute (on our large database) and it is readily available from openCV [39]. The main drawbacks of this detector are its sensibility to posture and lighting conditions, and the high number of false positive among small size faces (so we actually set a minimum face size of 60×60 pixels). Nowadays, deep learning approaches such as in [40] reaches over 97% of accuracy in detection, with far less sensibility to orientation and lighting.
- Our tracking phase also filters out some detected faces because we use a point tracker, which, by definition, can be very sensible to lighting conditions and occlusions (although we made efforts to improve its robustness [30]). Indeed, we detected some face instances without reliable tracking, so we could not keep them as face-tracks and discarded them. Optical flow approaches, such as in [41], seem to be very promising improvements.

- The last step concerns of course the face-track recognition, for which we apply a k -Faces sampling [30] to match *mean faces* described using VGG features [32]. Once more, deep learning approaches, such as in [7] and [42], bring higher-than-human performances (over 97% of accuracy). We can expect excellent results by matching *mean faces* using these deep features, and at least retrieve more face-tracks of a same individual.

Additionally, the networks we exploited to support social analysis were derived from news segment associations, *i.e.* networks built when two individuals appeared in a same news segment. One can argue that we would reach much higher accuracy by textual information instead of visual information. Indeed, named entity recognizers (NER) can reach high recall, sometime over 90% for the Stanford NER [43] on “Person” detection (see [44] for comparisons). This depends of course of the dataset, training, and above all, language models. Although results could reach up to 62% accuracy for “Person” detection in Japanese [45], we want to keep in mind that actually *seeing* somebody on TV and *hearing* a person’s mention has a different impact on viewers. However, we will definitely keep the analysis of the textual source and construction of networks from mentions as a future work. The comparison between both sources of information will actually be even more interesting.

Since we are discussing the data itself, our future work will extend the set of people to all faces detected in the dataset, not limited to the tagged individuals. We also put effort in enhancing the recall and the precision of the detection; and the addition of semantic information derived from the topic detection will be a great improvement.

This paper only scratches the surface, but the analysis of news data craves for application of many network analysis techniques. For example, the different families of links (screen, shot, segments, semantics, parties, *etc.*) also form a multiplex network as in [46], for which we can search for metrics [47]. We can draw multiplex networks as in [20] with people interacting through cabinets and hopefully find cohesive groups of politicians. The dynamic of links is also of great interest and Δ -cliques [48] (cliques over time in a stream of links) is a promising lead. In addition to finding outliers, we will be interested in groups of political actors who regularly appear together among similar subjects.

As for the social analysis, these networks provide a fertile ground to completely investigate political questions. The definition of the links is extremely important though. To provide convincing conclusive analysis, links and weights need to be clearly interpreted - the unambiguous non-split screen co-occurrence being the best. Using the screen co-occurrence will require to measure exactly the uncertainty of the detection and of the matching, so screen-time may be turned into a reliable variable for regression.

We have not studied the diversity of political opposition, but there are many other parameters we would like to investigate: for example, we could create meta-networks of political parties and observe who bridges them and when do these links occur, we could also focus on the role of the opposition with foreign leaders. We believe that such analysis will ultimately bring quantitative evidence of social phenomenon, such as the “presidentialization” of the Japanese politics [15].

7 Conclusion

This work has introduced the production and analysis of face detection and tracking data over twelve years of news broadcast. We have detailed the data's characteristics and brought a few outliers. Together with the use of topic segmentation and limited domain knowledge, we have derived many networks, each presenting a different point of view on the data, conducting to the understanding of the politico-media situation displayed by NHK during this period of time.

The combined views of these networks show interesting insights on the story behind the data, an arguably clear picture of the media/politics landscape during the different cabinets, also isolating key players at different levels. That is what the general reader may take away: thanks to the networks, even those knowing nothing of the Japanese media landscape can quickly get an idea of who the main actors are, and their relative importance. Of course, the level of comprehension will increase as we improve the precision of our detectors, and the semantics of our links.

We have demonstrated in this small use case an entire pipeline: how – from raw video data – we can create networks that can tell us a story of the politico-media scene in Japan, with the help of their structural properties. Of course, to support a full political analysis, we need to assess the uncertainty of links derived from the face detection, since we want to test hypothesis with these tools.

Rather different to the classical topic detection and tracking approaches of news data, this work brings up if not confirms the relevance of network analysis derived from news data. By itself, this is also an interesting framework for many potential contributions to the current challenges of social network analysis – including, but not limited to, multiplex and multi-attributed network analysis, dynamic networks, and their combination.

Finally, this work has given us useful directions that will help us design visual analytics tools, which we wish to put in the hands of domain experts, sociologists and journalists, to conduct the in-depth analysis of over 12 years of news.

Figures

To ease the reading flow of the paper, we propose a larger format of all small figures in the Appendix section at the end of this document. Additionally, note that all figures are zoomable on the digital copy of the document.

Competing interests

The authors declare that they have no competing interests.

Acknowledgments

We would like to thank the Japanese Society for the Promotion of Science for funding this research (as part of the FY2015 *Grant-in-Aid for Scientific Research*). We thank the NHK for their courtesy of the video data. All the network illustrations have been realized with the Tulip software [49].

Author details

¹National Institute of Informatics, 2-1-2 Hitotsubashi, Chiyoda-ku, 101-8430 Tokyo, Japan. ²JFLI CNRS UMI 3527, 7-3-1, Hongo, Bunkyo-ku, The University of Tokyo, 103-0033 Tokyo, Japan. ³City University of Hong Kong, 18 Tat Hong Avenue, Kowloon Tong, Hong Kong, RPC. ⁴University of Information Technology, Vietnam National University, Quarter 6, Linh Trung Ward, Thu Duc District, Ho Chi Minh City, Vietnam.

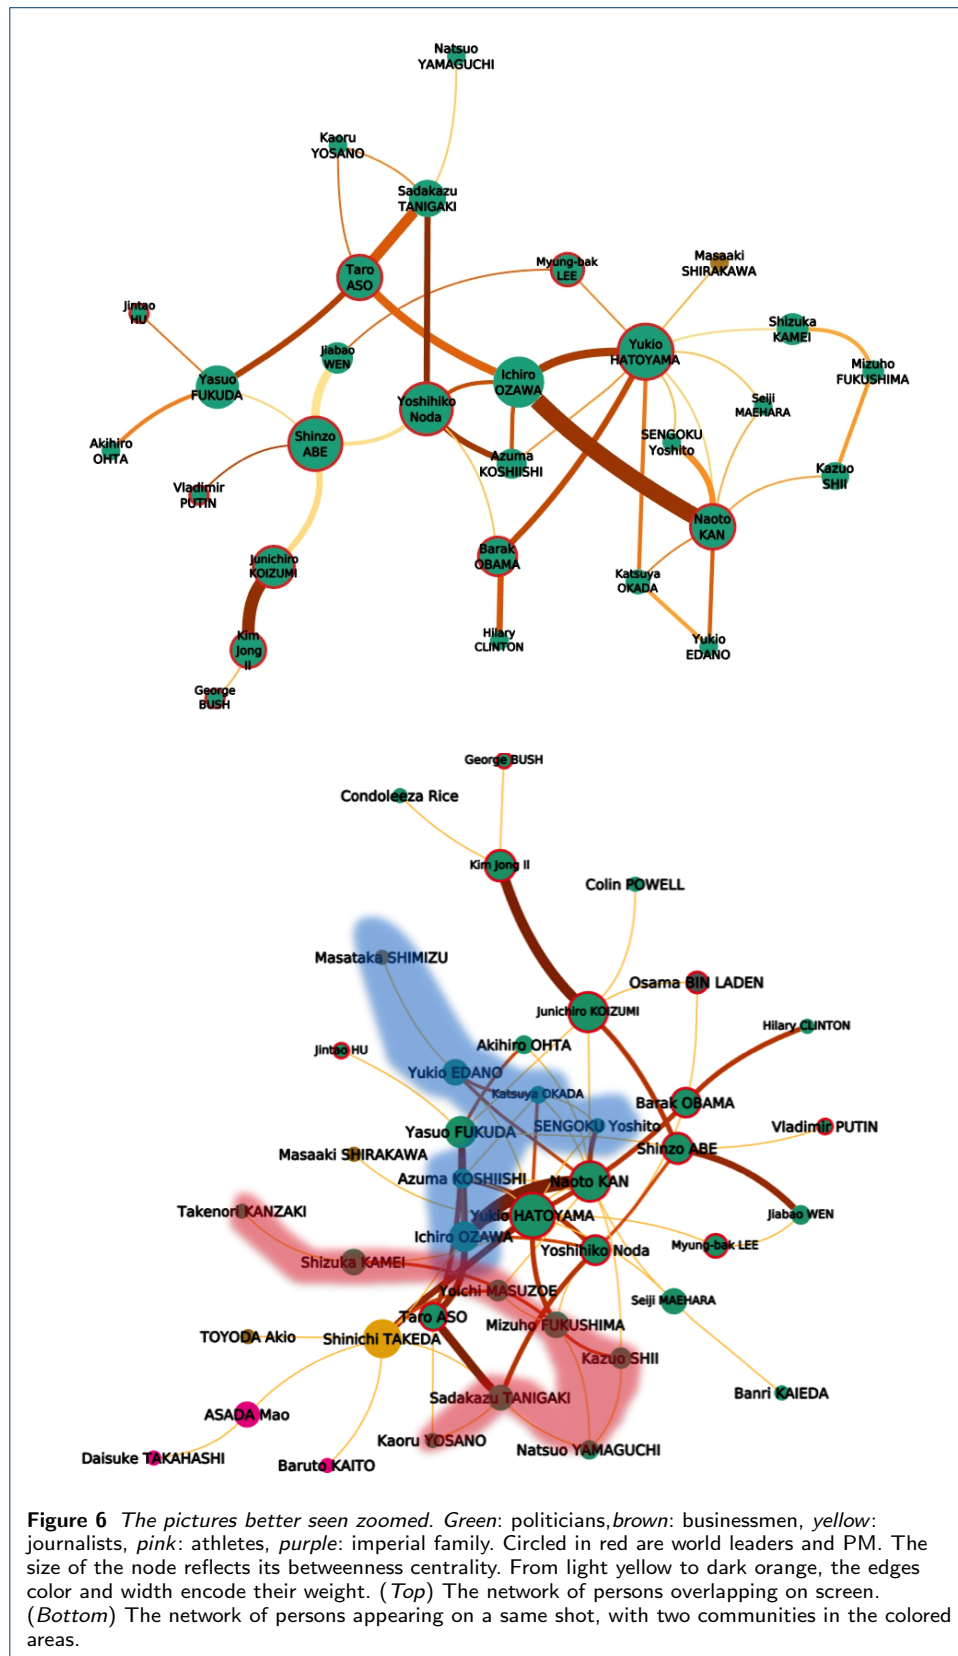
References

1. Hey, A.J., Tansley, S., Tolle, K.M.: The fourth paradigm – data-intensive scientific discovery. *E-Science and Information Management* **12**, 1–1 (2009)
2. Allan, J.: Topic detection and tracking: Event-based information organization. *Kluwer Academic Publishers* **12** (2012)
3. Hervé, N., Viaud, M.-L., Thièvre, J., Saulnier, A., Champ, J., Letessier, P., Buisson, O., Joly, A.: Otmedia: the french transmedia news observatory. In: *Proceedings of the 21st ACM International Conference on Multimedia*, pp. 441–442 (2013). ACM

4. Cagé, J.: Media competition, information provision and political participation (2014). Unpublished manuscript, Harvard University
5. Resnick, P., Garrett, R.K., Kriplean, T., Munson, S.A., Stroud, N.J.: Bursting your (filter) bubble: Strategies for promoting diverse exposure. In: Proceedings of the ACM 2013 Conference on Computer Supported Cooperative Work Companion, pp. 95–100 (2013)
6. Hazell, R., Worthy, B.: Assessing the performance of freedom of information. *Government Information Quarterly* **27**(4), 352–359 (2010)
7. Schroff, F., Kalenichenko, D., Philbin, J.: Facenet: A unified embedding for face recognition and clustering. arXiv preprint arXiv:1503.03832 (2015)
8. Viola, P., Jones, M.J.: Robust real-time face detection. *International journal of computer vision* **57**(2), 137–154 (2004)
9. Ide, I., Nack, F.: Explain this to me! *ITE Transactions on Media Technology and Applications* **1**(2), 101–117 (2013)
10. Castillo, C., De Francisci Morales, G., Mendoza, M., Khan, N.: Says who?: automatic text-based content analysis of television news. In: Proceedings of the 2013 International Workshop on Mining Unstructured Big Data Using Natural Language Processing, pp. 53–60 (2013). ACM
11. Sudhahar, S., Franzosi, R., Cristianini, N.: Automating quantitative narrative analysis of news data. In: WAPA, pp. 63–71 (2011)
12. Street, J.: *Mass media, politics and democracy*. Palgrave Macmillan (2010)
13. Mughan, A.: Media and the presidentialization of parliamentary elections. *American History in Depth* (2000)
14. Krauss, E.S., Nyblade, B.: 'presidentialization' in japan? the prime minister, media and elections in japan. *British Journal of Political Science* **35**(02), 357–368 (2005)
15. Jou, W., Endo, M.: Presidentialization of japanese politics? examining political leader evaluations and vote choice. *Japanese Journal of Political Science* **16**(03), 357–387 (2015)
16. Lazer, D.: Networks in political science: Back to the future. *PS: Political Science & Politics* **44**(01), 61–68 (2011)
17. Davis, A., Gardner, B.B., Gardner, M.R.: *Deep south: A social anthropological study of caste and class*. University of South Carolina Press **1** (2009)
18. Adamic, L.A., Glance, N.: The political blogosphere and the 2004 us election: divided they blog. In: Proceedings of the 3rd International Workshop on Link Discovery, pp. 36–43 (2005). ACM
19. Martin, M.K., Pfeffer, J., Carley, K.M.: Network text analysis of conceptual overlap in interviews, newspaper articles and keywords. *Social Network Analysis and Mining* **3**(4), 1165–1177 (2013)
20. Renoust, B., Melançon, G., Viaud, M.-L.: Entanglement in multiplex networks: understanding group cohesion in homophily networks. *Social Network Analysis - Community Detection and Evolution* **1**(1), 89–117 (2014)
21. Waumans, M.C., Nicodème, T., Bersini, H.: Topology analysis of social networks extracted from literature. *PLoS one* **10**(6), 0126470 (2015)
22. Nan, C.-J., Kim, K.-M., Zhang, B.-T.: Social network analysis of tv drama characters via deep concept hierarchies. In: Proceedings of ASONAM 2015 (2015)
23. Mish, B.: *Game of Nodes: A Social Network Analysis of Game of Thrones* (2015). <https://gameofnodes.wordpress.com/>
24. Itoh, M., Toyoda, M., Zhu, C.Z., Satoh, S., Kitsuregawa, M.: Image flows visualization for inter-media comparison. In: Pacific Visualization Symposium (PacificVis), 2014 IEEE, pp. 129–136 (2014). IEEE
25. Luo, H., Fan, J., Yang, J., Ribarsky, W., Satoh, S.: Analyzing large-scale news video databases to support knowledge visualization and intuitive retrieval. In: Visual Analytics Science and Technology, 2007. VAST 2007. IEEE Symposium On, pp. 107–114 (2007). IEEE
26. Viaud, M.-L., Thièvre, J., Goëau, H., Saulnier, A., Buisson, O.: Interactive components for visual exploration of multimedia archives. In: Proceedings of the 2008 International Conference on Content-based Image and Video Retrieval, pp. 609–616 (2008). ACM
27. Seifert, C., Sabol, V., Kienreich, W., Lex, E., Granitzer, M.: Visual analysis and knowledge discovery for text. *Large-Scale Data Analytics* **1**(1), 189–218 (2014)
28. Katayama, N., Mo, H., Ide, I., Satoh, S.: Mining large-scale broadcast video archives towards inter-video structuring. *Advances in Multimedia Information Processing-PCM 2004* **1**(1), 489–496 (2005)
29. Ide, I., Mo, H., Katayama, N., Satoh, S.: Topic threading for structuring a large-scale news video archive. *Image and Video Retrieval* **1**(1), 123–131 (2004)
30. Ngo, T.D., Vu, H.T., Duy-Dinh, L., Satoh, S.: Face retrieval in large-scale news video datasets. *IEICE TRANSACTIONS on Information and Systems* **96**(8), 1811–1825 (2013)
31. Shi, J., Tomasi, C.: Good features to track. In: Computer Vision and Pattern Recognition, 1994. Proceedings CVPR'94., 1994 IEEE Computer Society Conference On, pp. 593–600 (1994). IEEE
32. Everingham, M., Sivic, J., Zisserman, A.: Hello! my name is... buffy"—automatic naming of characters in tv video. In: *BMVC*, vol. 2, p. 6 (2006)
33. Wang, X., Zhai, C., Hu, X., Sproat, R.: Mining correlated bursty topic patterns from coordinated text streams. In: Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 784–793 (2007). ACM
34. Blondel, V.D., Guillaume, J.-L., Lambiotte, R., Lefebvre, E.: Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment* **2008**(10), 10008 (2008)
35. Seidman, S.B.: Network structure and minimum degree. *Social networks* **5**(3), 269–287 (1983)
36. Herman, I., Marshall, M.S., Melançon, G.: Density functions for visual attributes and effective partitioning in graph visualization. In: *Information Visualization, 2000. IEEE Symposium On*, pp. 49–56 (2000). IEEE
37. Meyer, I.: The History of Japan podcast, Episode 82: The Shadow Shogun, Redux (2014). <https://historyofjapan.wordpress.com/2014/12/20/episode-82-the-shadow-shogun-redux/>
38. Nguyen, T.N., Ngo, T.D., Le, D.-D., Satoh, S., Le, B.H., Duong, D.A.: An efficient method for face retrieval from large video datasets. In: Proceedings of the ACM International Conference on Image and Video Retrieval,

- pp. 382–389 (2010). ACM
39. Bradski, G., *et al.*: The opencv library. *Doctor Dobbs Journal* **25**(11), 120–126 (2000)
 40. Farfadi, S.S., Saberian, M.J., Li, L.: Multi-view face detection using deep convolutional neural networks. *CoRR abs/1502.02766* (2015)
 41. Ranftl, A., Alonso-Fernandez, F., Karlsson, S.: Face tracking using optical flow. In: *Biometrics Special Interest Group (BIOSIG), 2015 International Conference of The*, pp. 1–5 (2015). IEEE
 42. Taigman, Y., Yang, M., Ranzato, M., Wolf, L.: Deepface: Closing the gap to human-level performance in face verification. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1701–1708 (2014)
 43. Finkel, J.R., Grenager, T., Manning, C.: Incorporating non-local information into information extraction systems by gibbs sampling. In: *Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics*, pp. 363–370 (2005). Association for Computational Linguistics
 44. Atdag, S., Labatut, V.: A comparison of named entity recognition tools applied to biographical texts. In: *Systems and Computer Science (ICSCS), 2013 2nd International Conference On*, pp. 228–233 (2013). IEEE
 45. Ichihara, M., Komiya, K., Iwakura, T., Yamazaki, M.: Error analysis of named entity recognition in bccwj. In: *Proceeding of NLP2015*, pp. 0–0 (2015). ANLP
 46. Kivelä, M., Arenas, A., Barthelemy, M., Gleeson, J.P., Moreno, Y., Porter, M.A.: Multilayer networks. *Journal of Complex Networks* **2**(3), 203–271 (2014)
 47. Battiston, F., Nicosia, V., Latora, V.: Structural measures for multiplex networks. *Physical Review E* **89**(3), 032804 (2014)
 48. Viard, J., Latapy, M., Magnien, C.: Computing maximal cliques in link streams. *arXiv preprint arXiv:1502.00993* (2015)
 49. Auber, D., Archambault, D., Bourqui, R., Delest, M., Dubois, J., Pinaud, B., Lambert, A., Mary, P., Mathiaut, M., Melançon, G.: Tulip III. *Encyclopedia of Social Network Analysis and Mining* **1**(1), 2216–2240 (2014)

Appendix



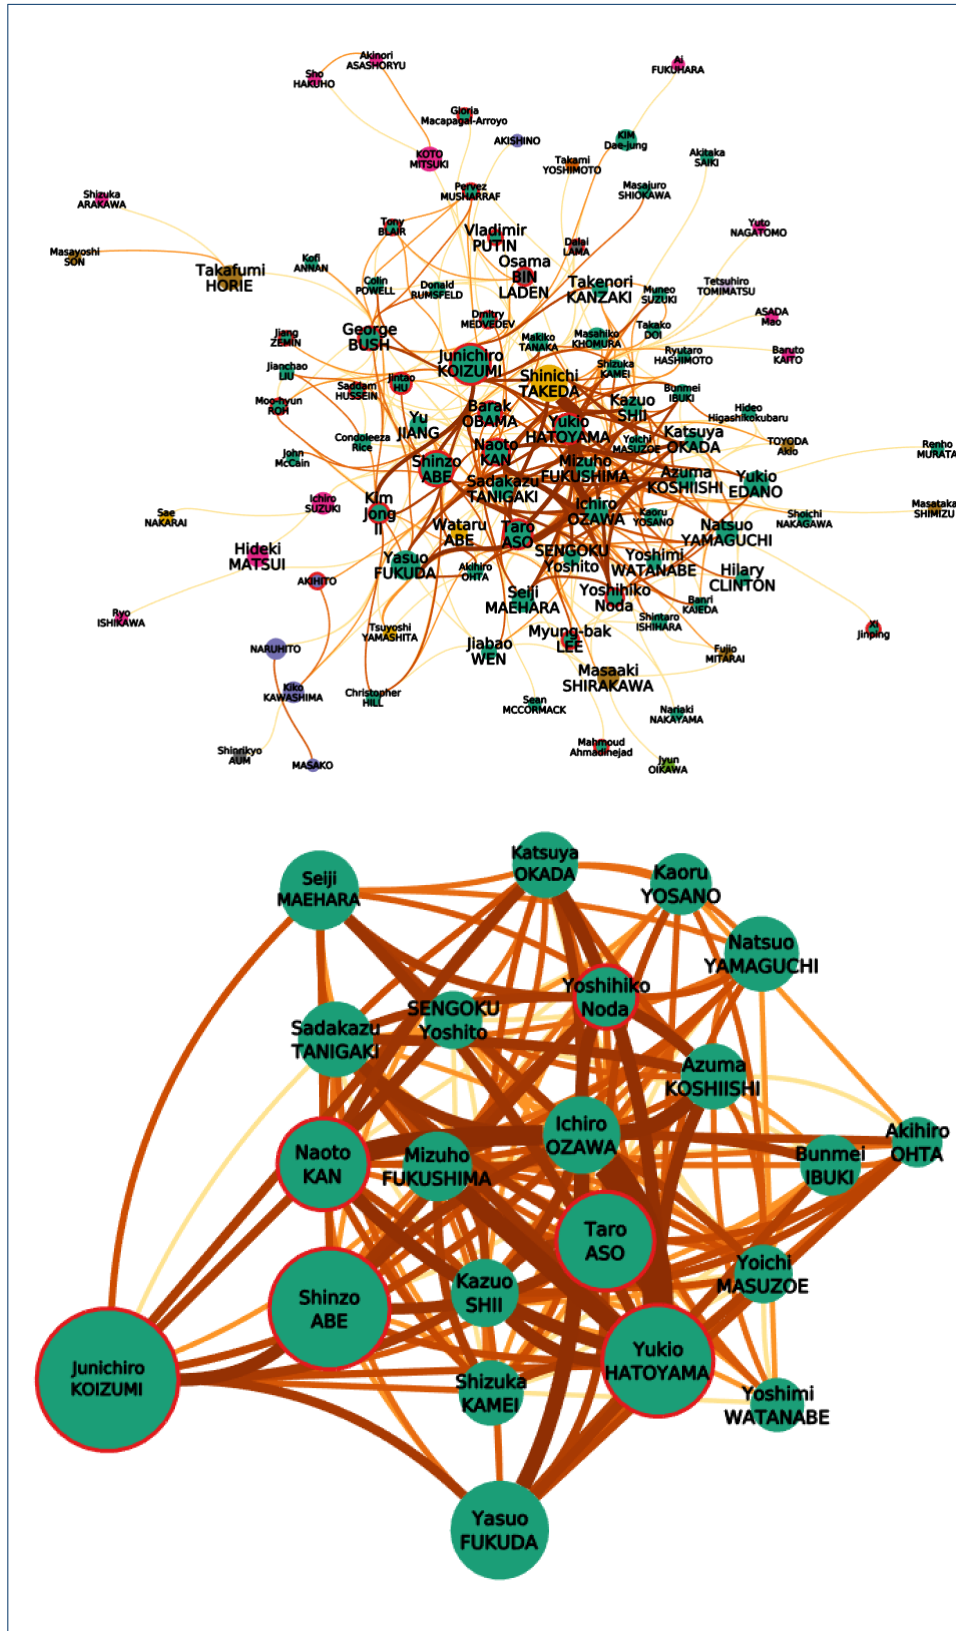
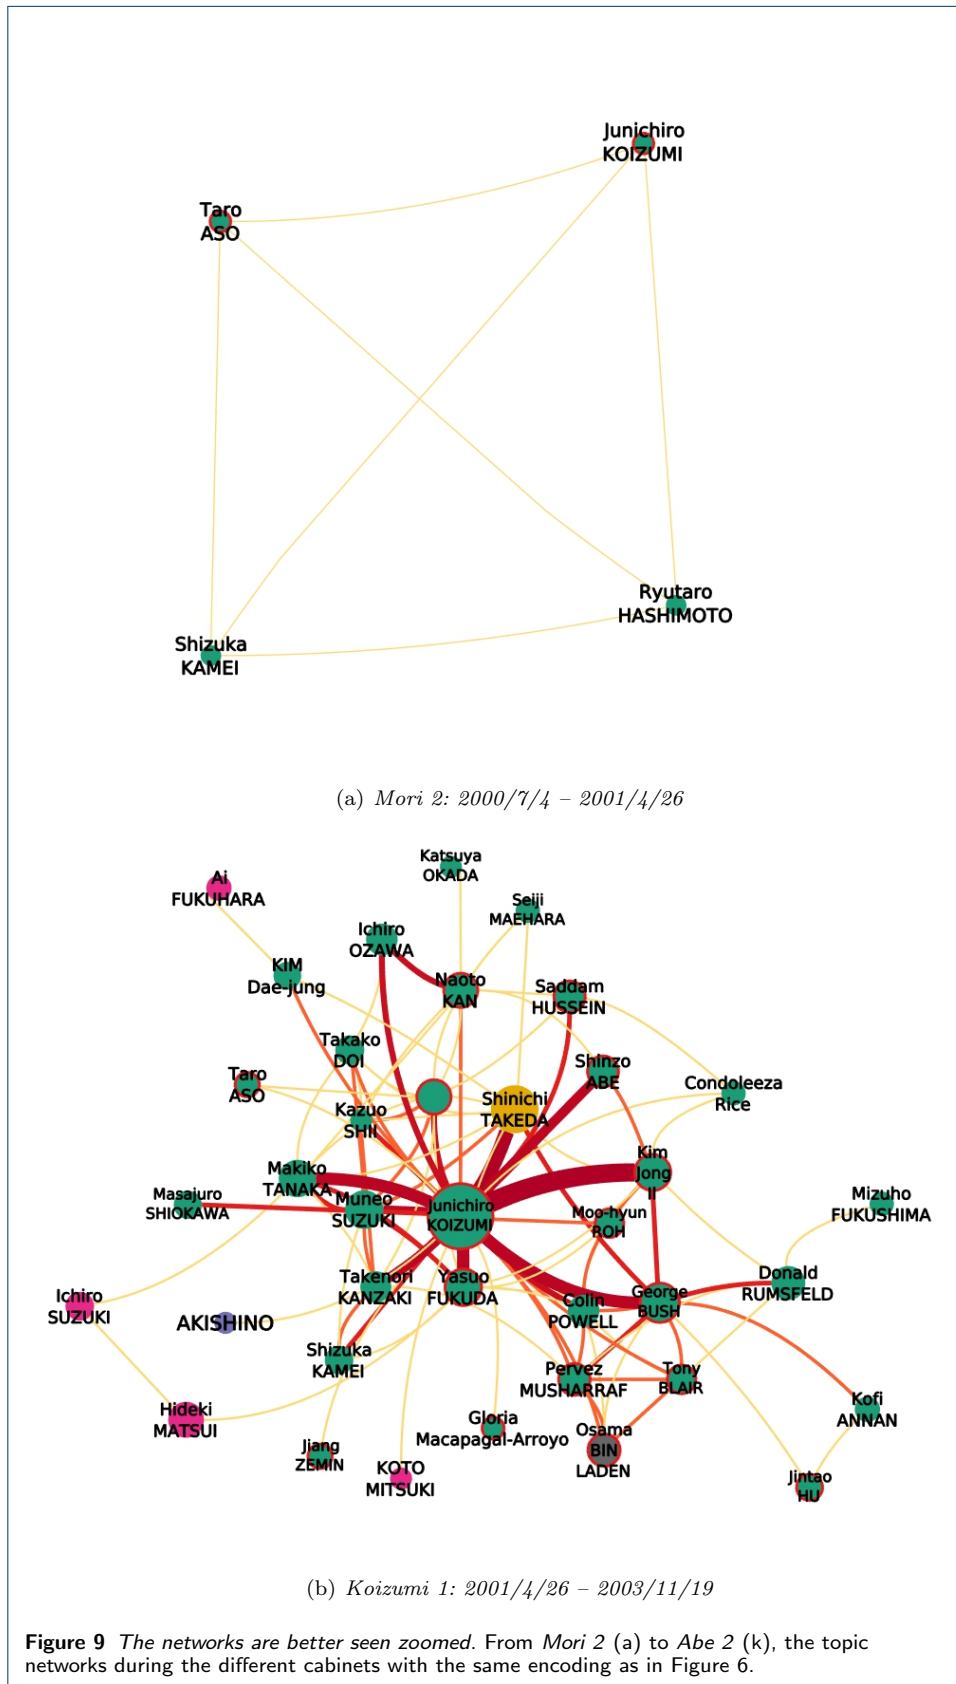
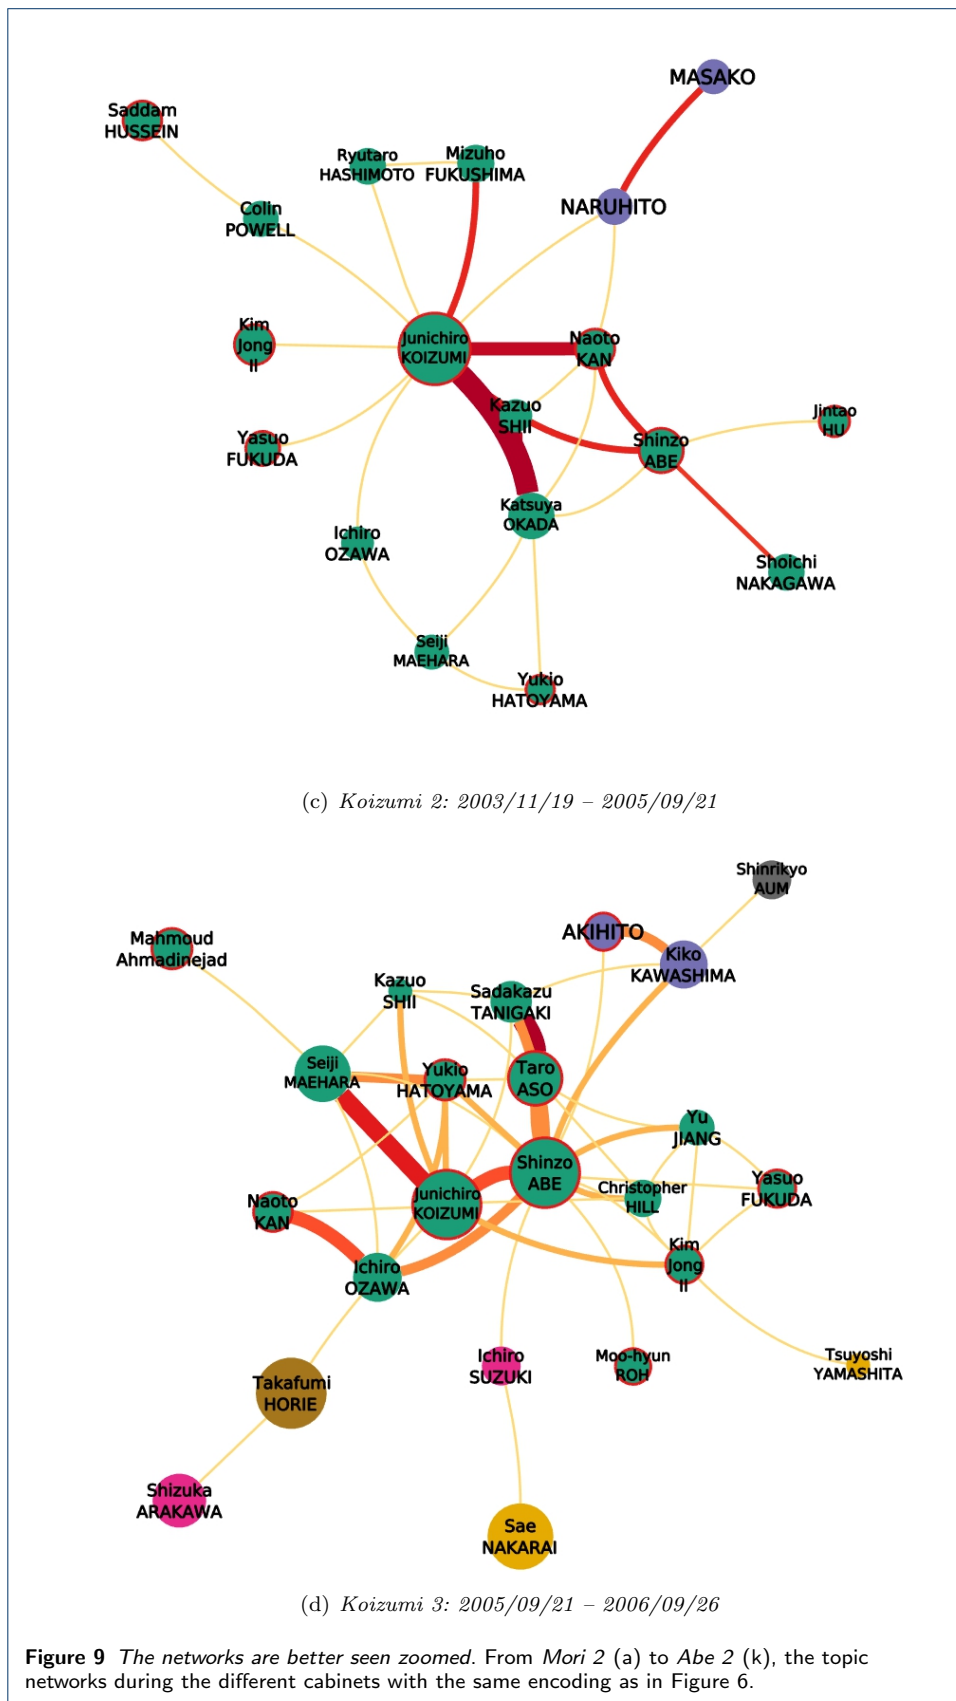


Figure 7 Pictures better seen zoomed. Same encoding as in Figure 6. Edge weight corresponds to the number of common news segments. (Top) The network of persons detected during a same topic. (Bottom) The k -core ($k = 13$) of this network.






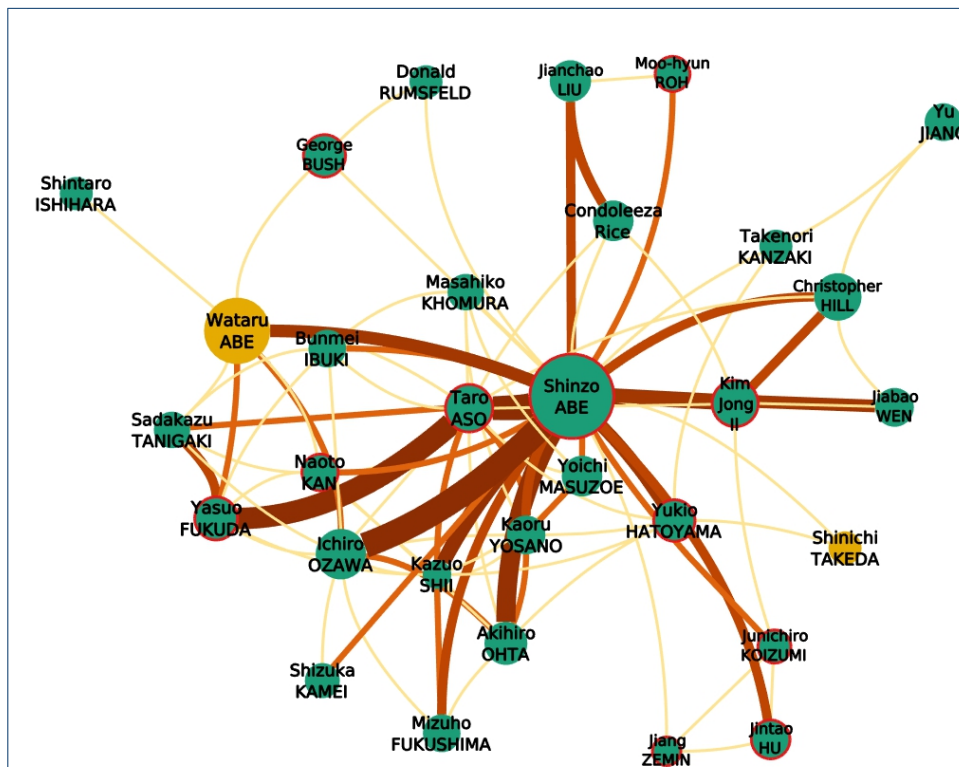
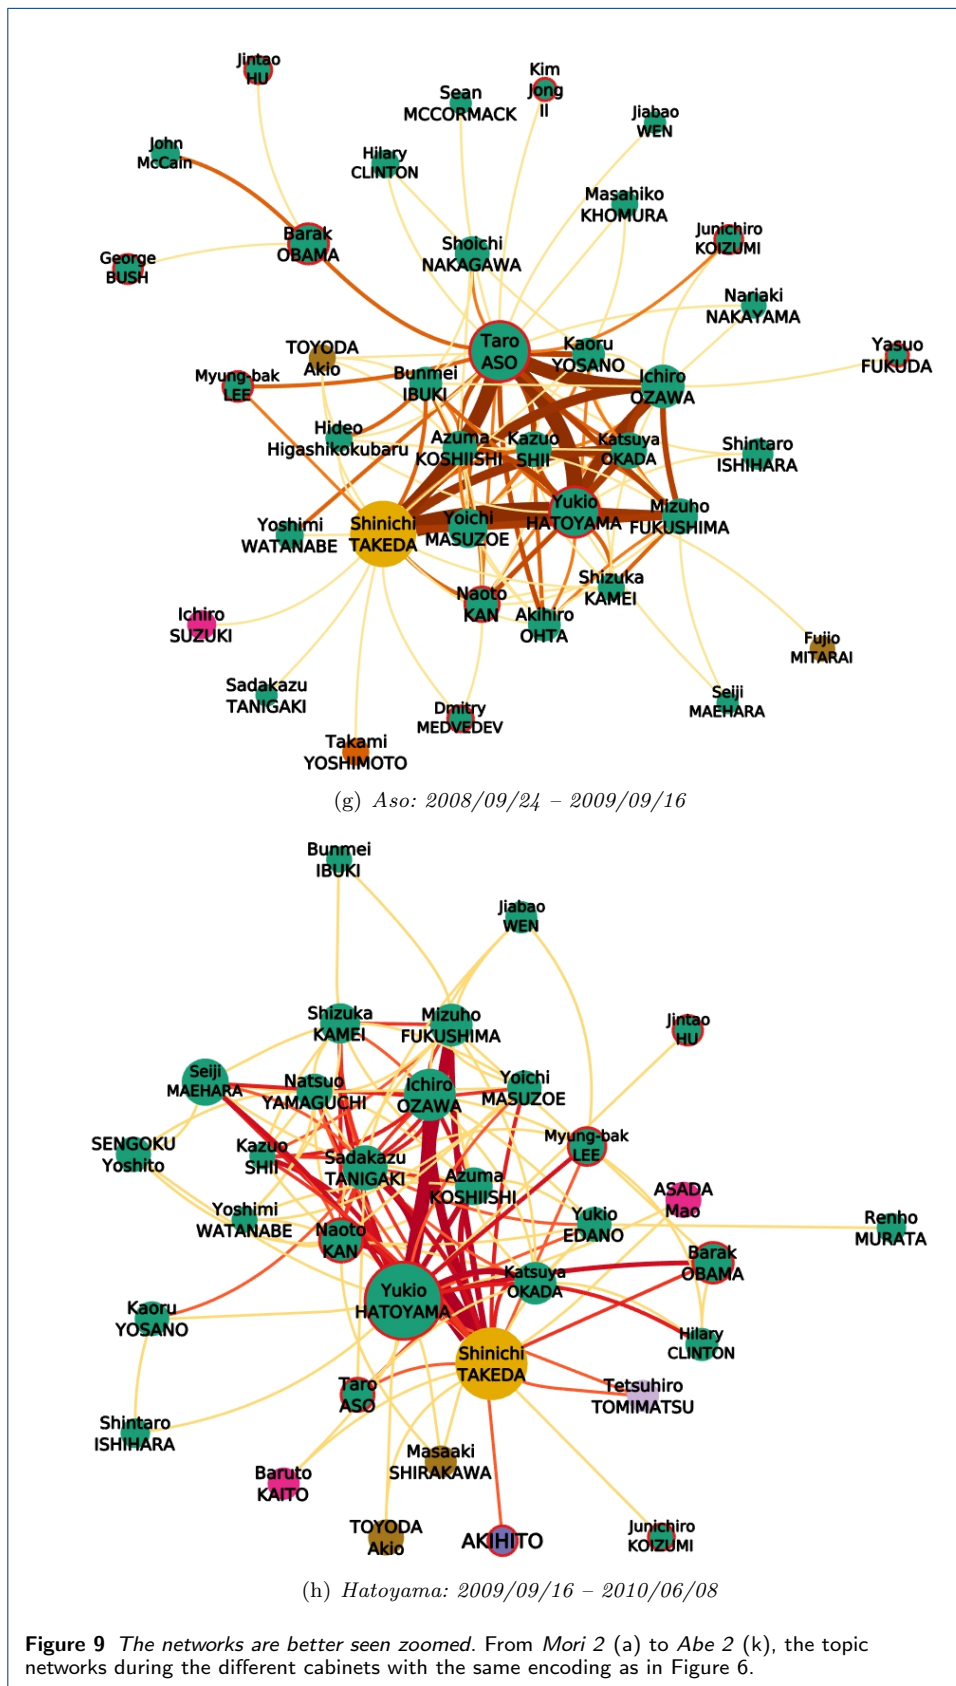


Figure 9 The networks are better seen zoomed. From Mori 2 (a) to Abe 2 (k), the topic networks during the different cabinets with the same encoding as in Figure 6.



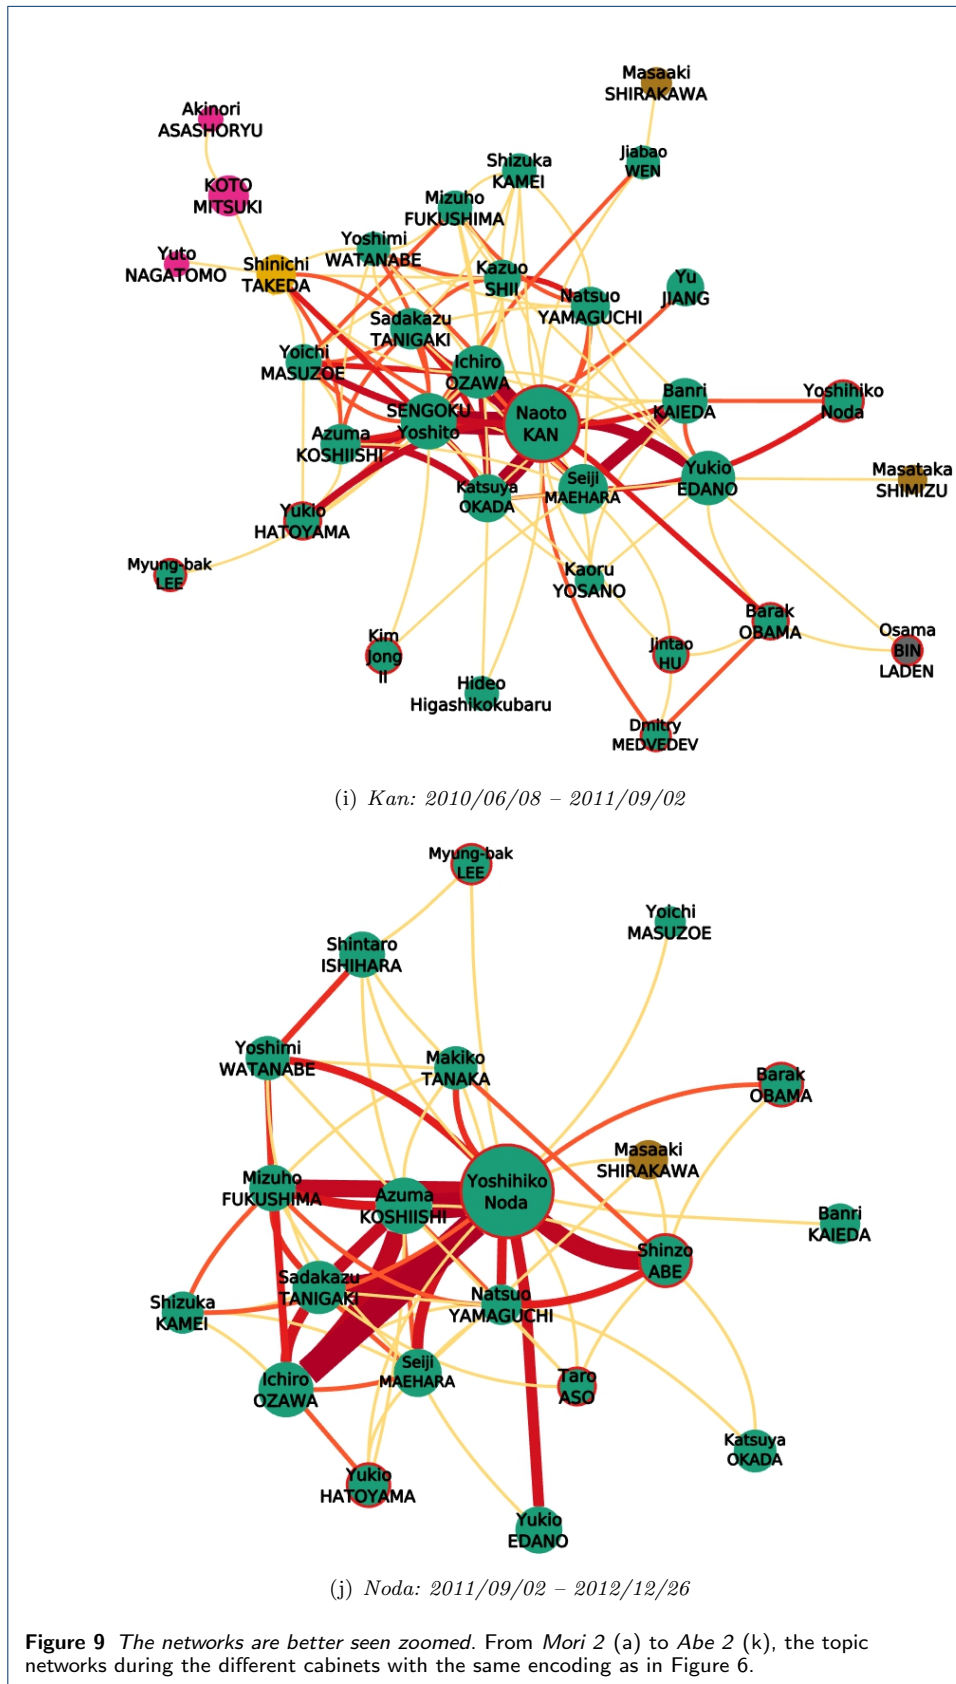
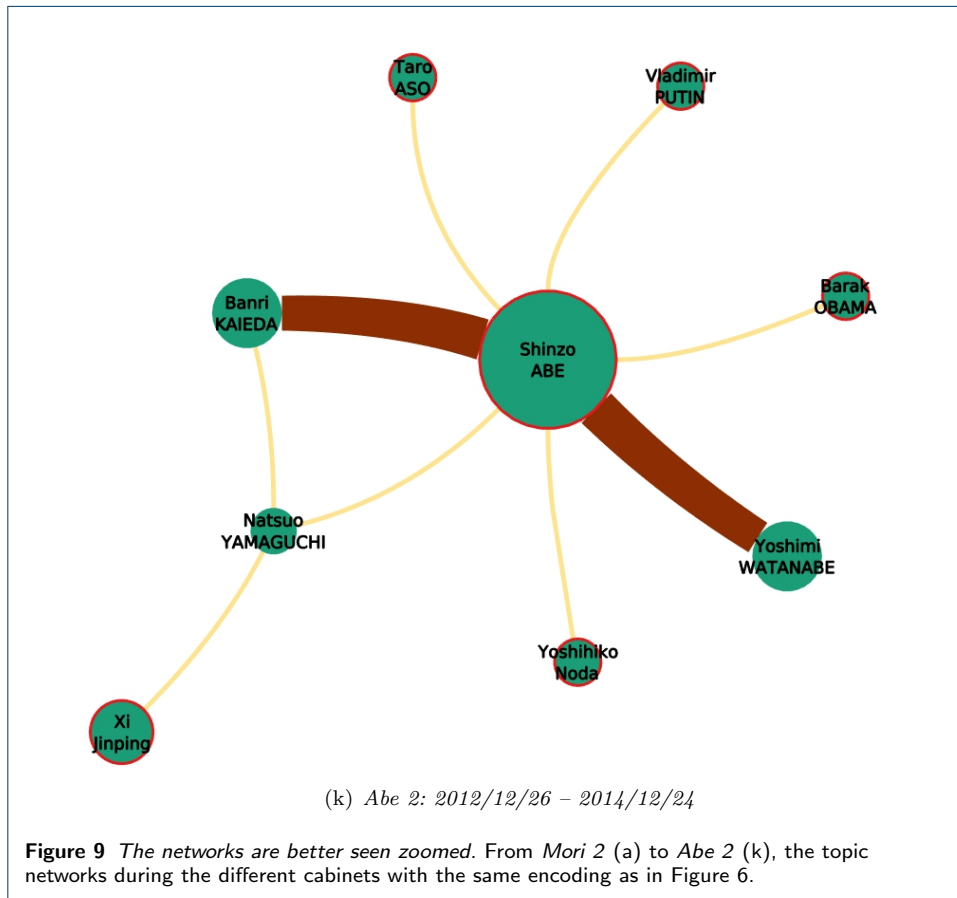





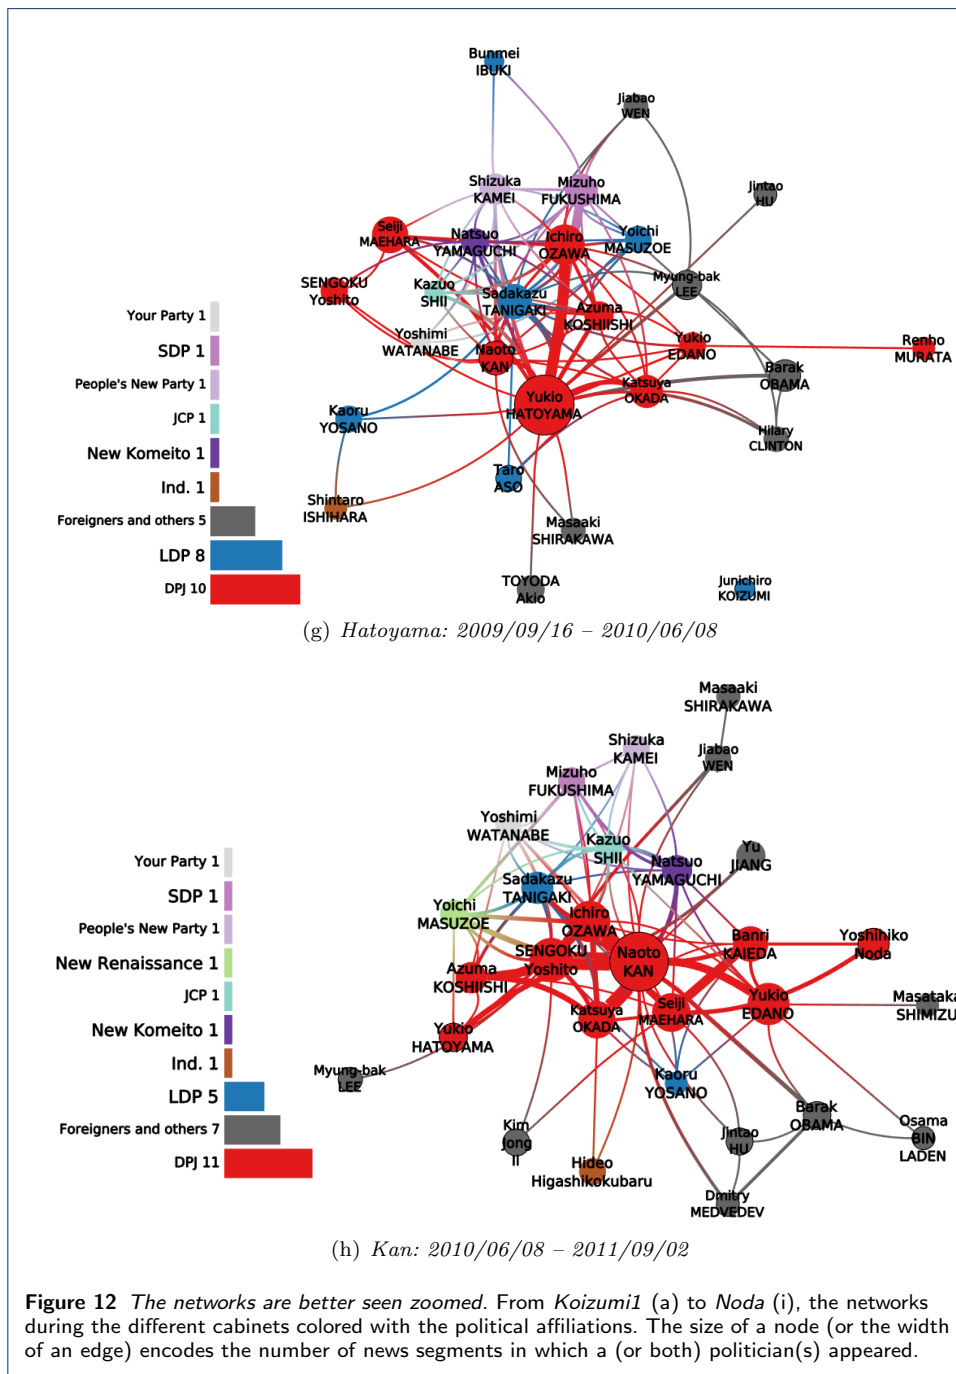
Figure 9 The networks are better seen zoomed. From Mori 2 (a) to Abe 2 (k), the topic networks during the different cabinets with the same encoding as in Figure 6.

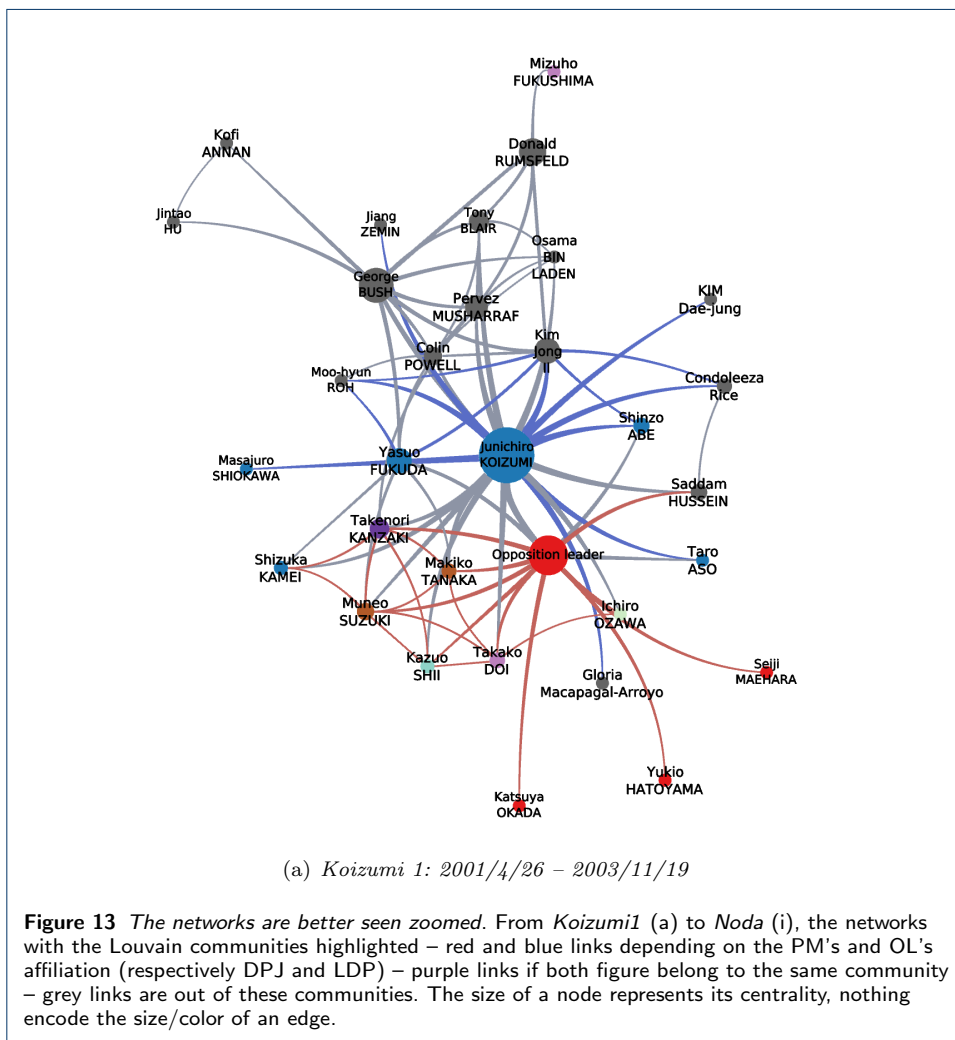
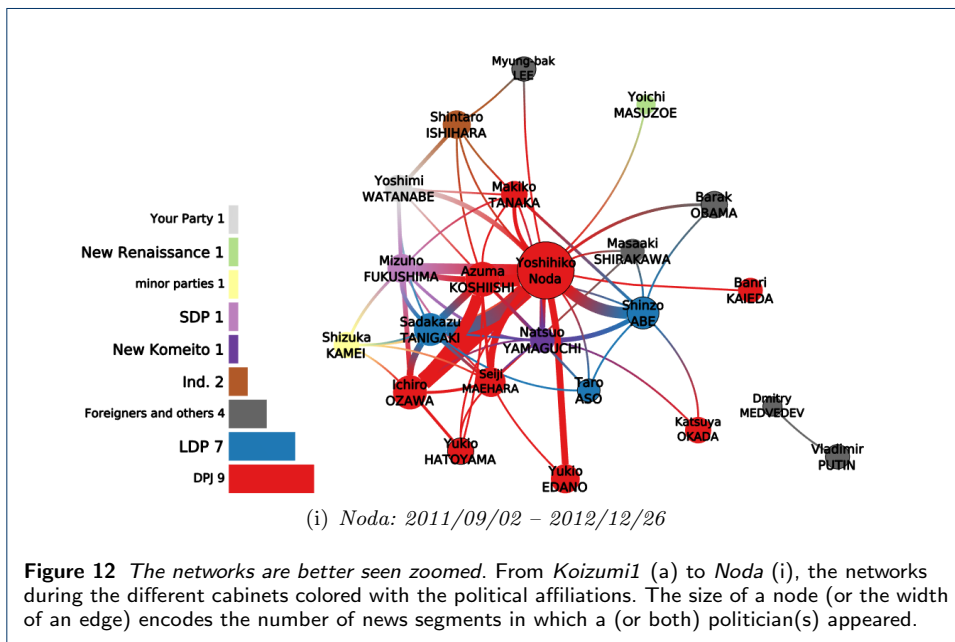


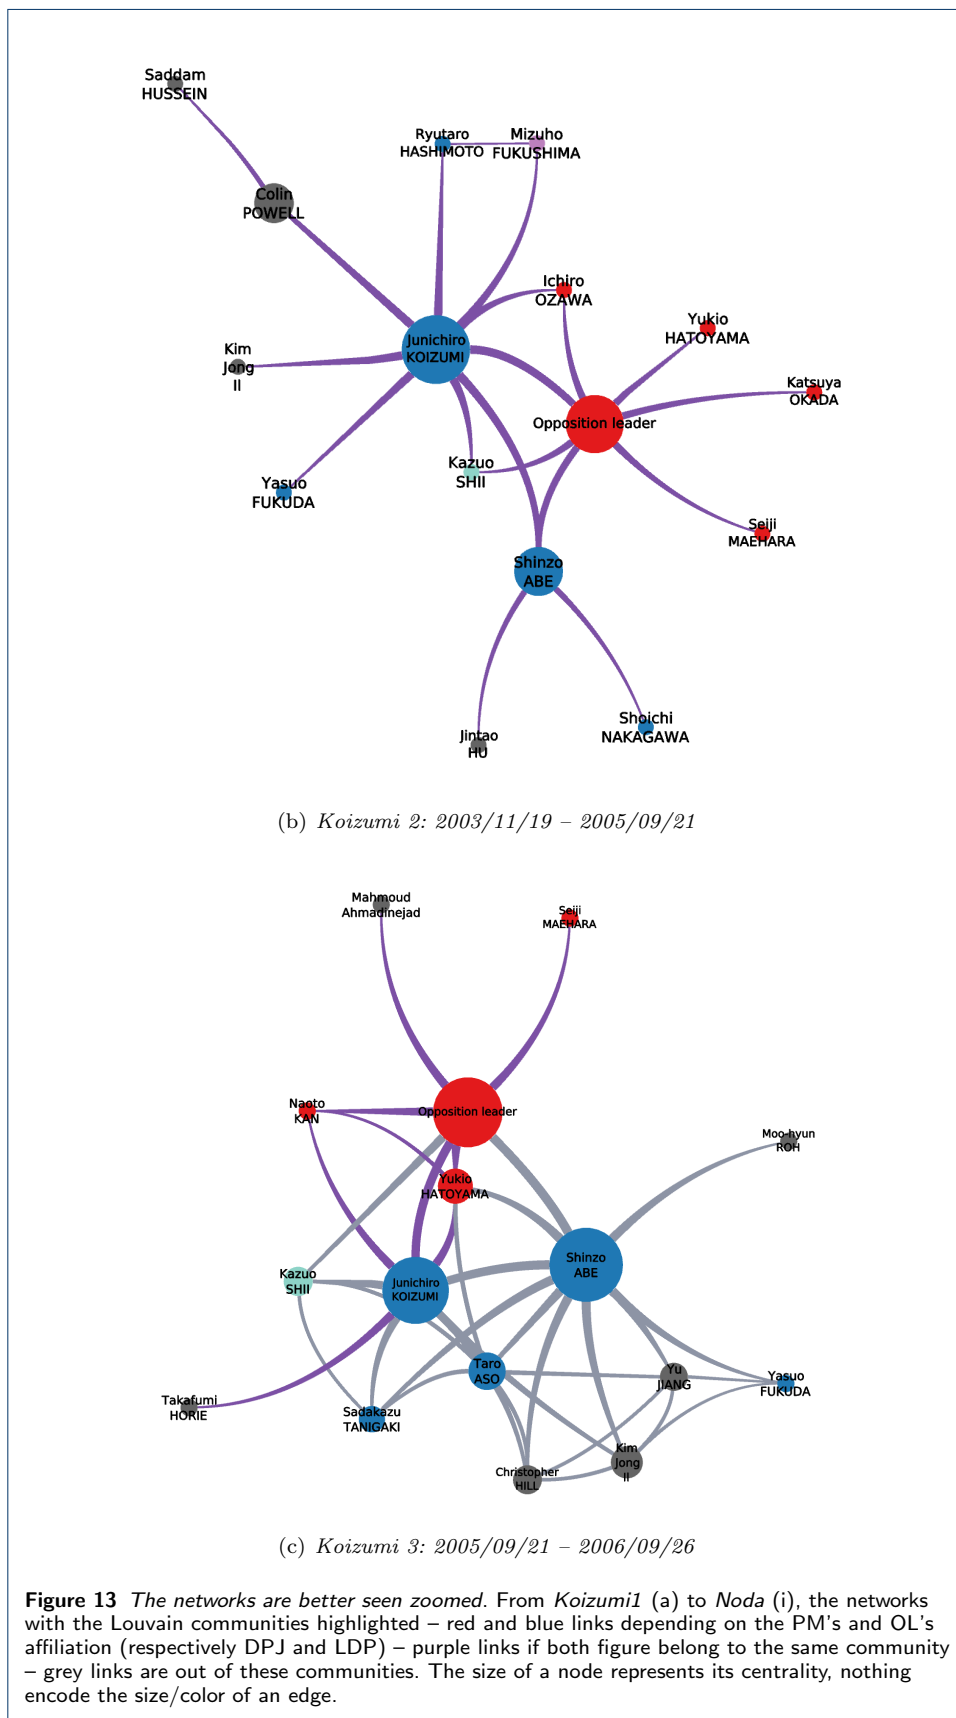


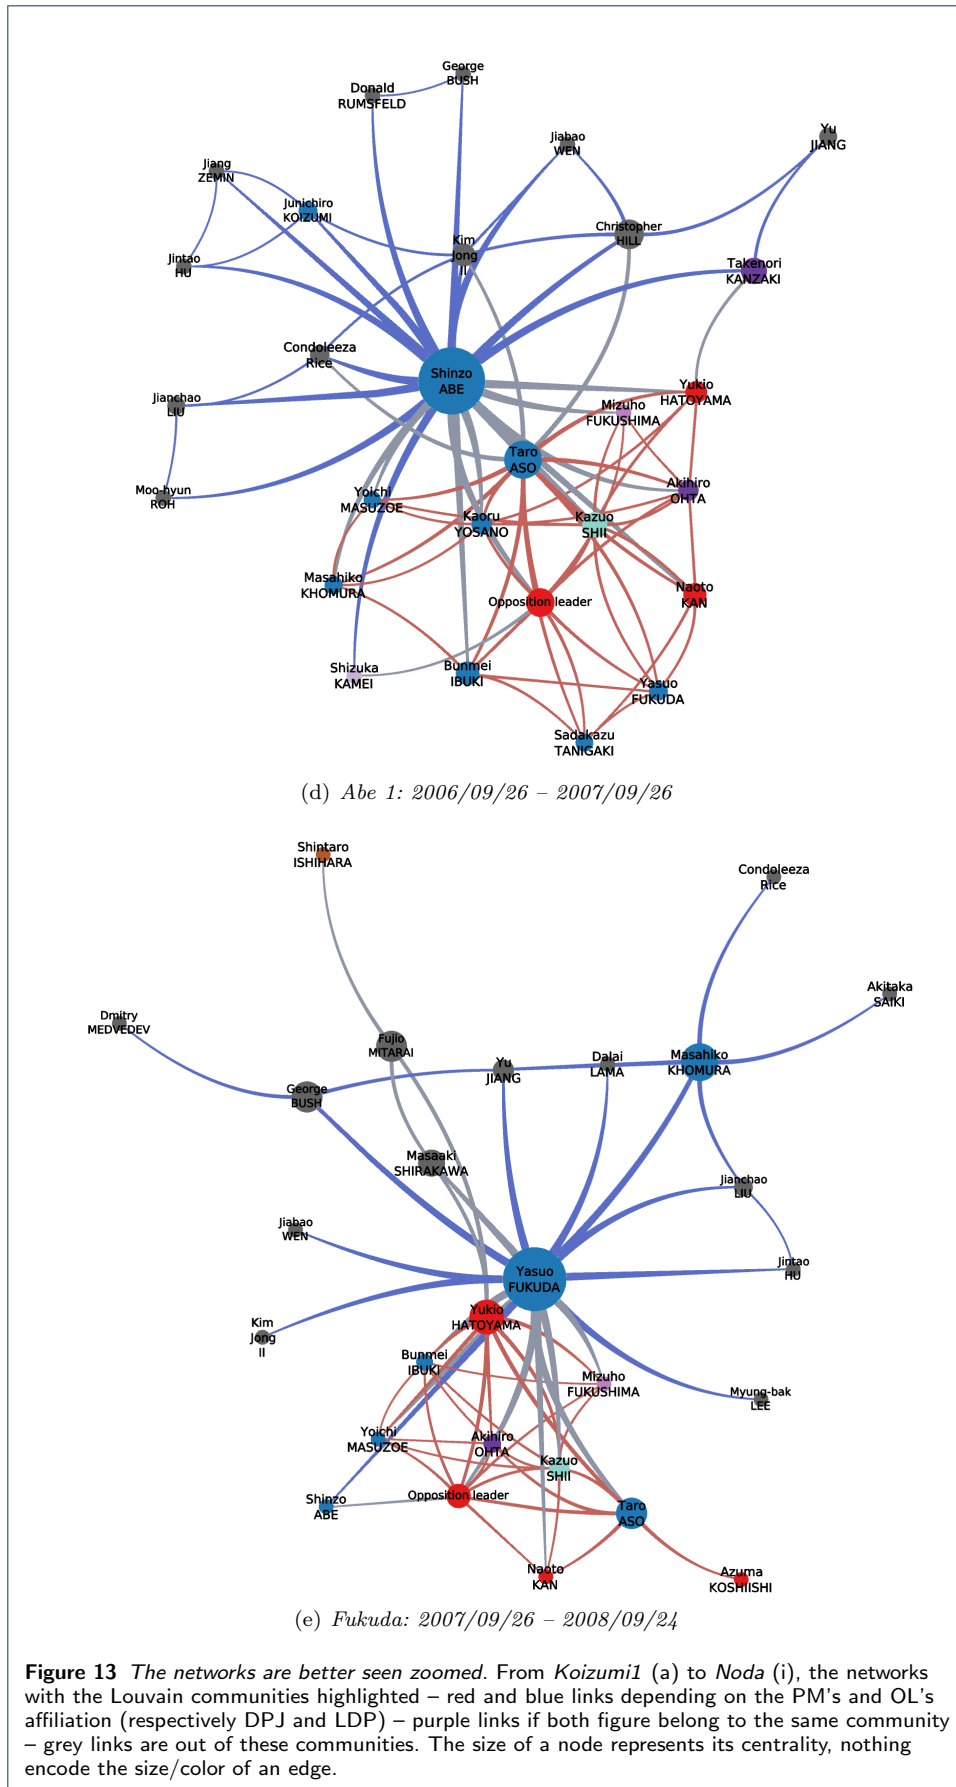


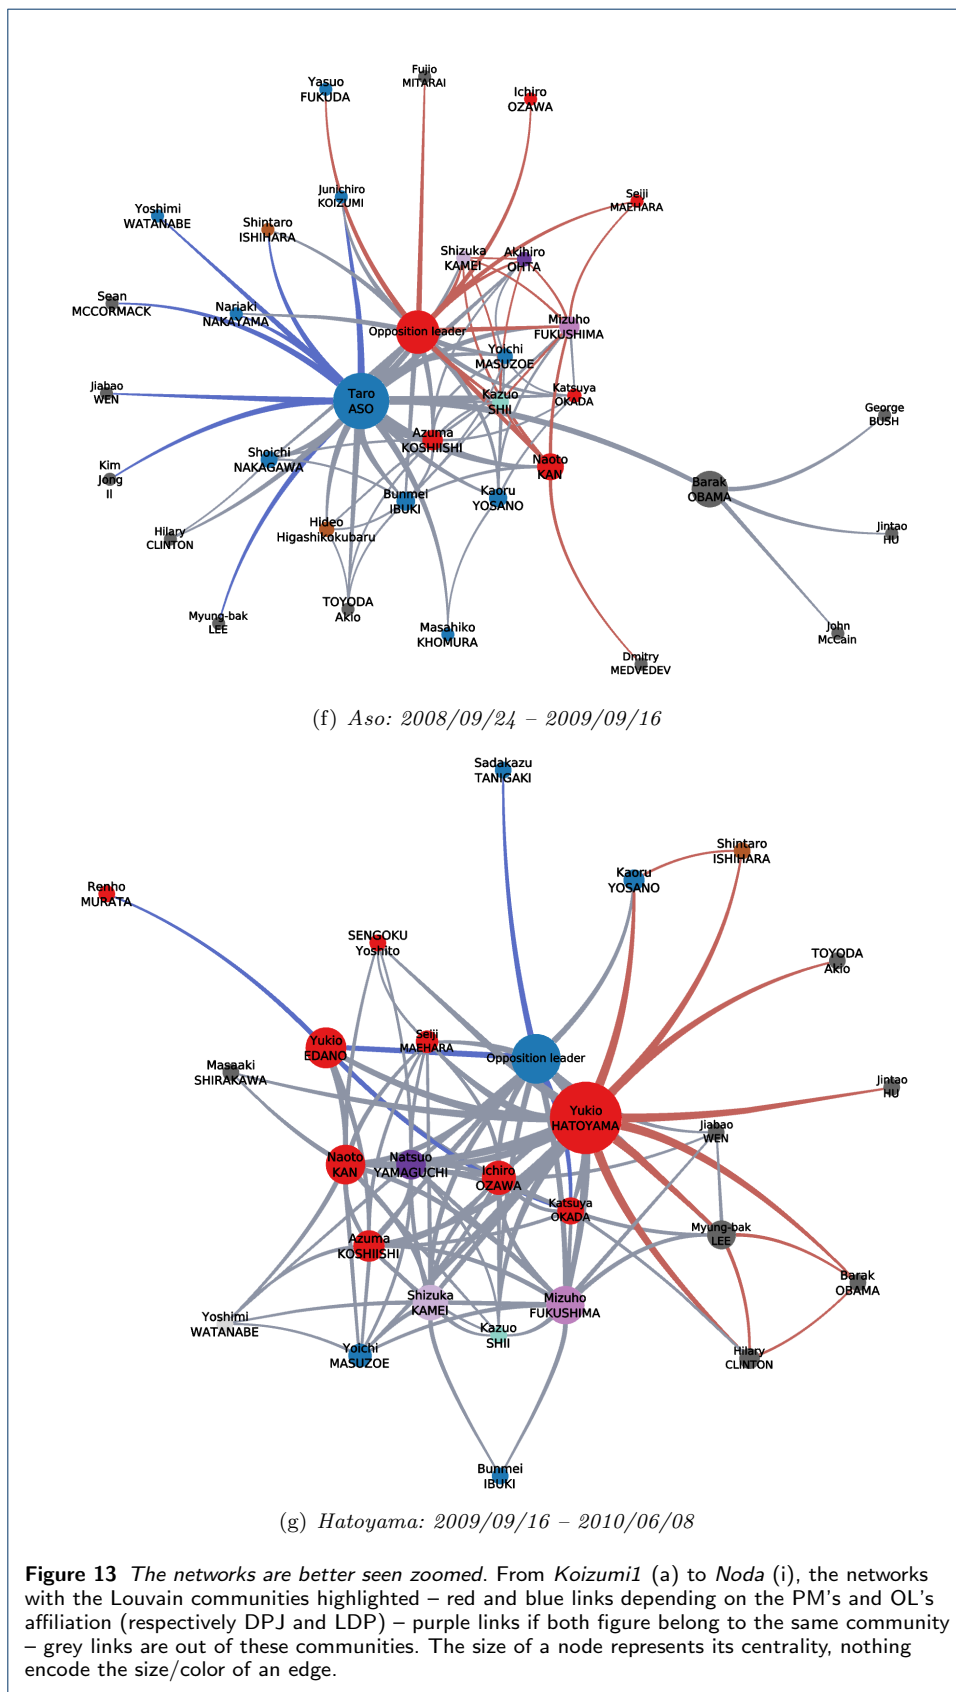


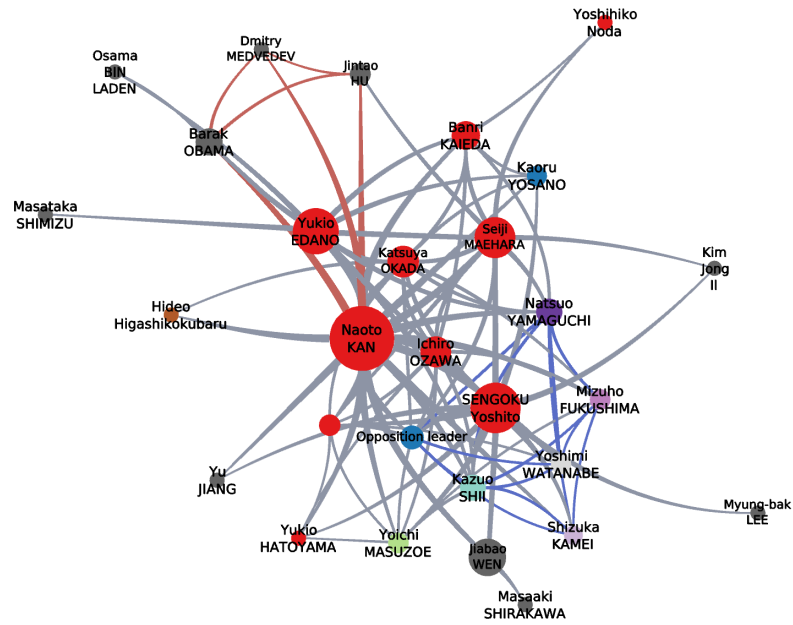




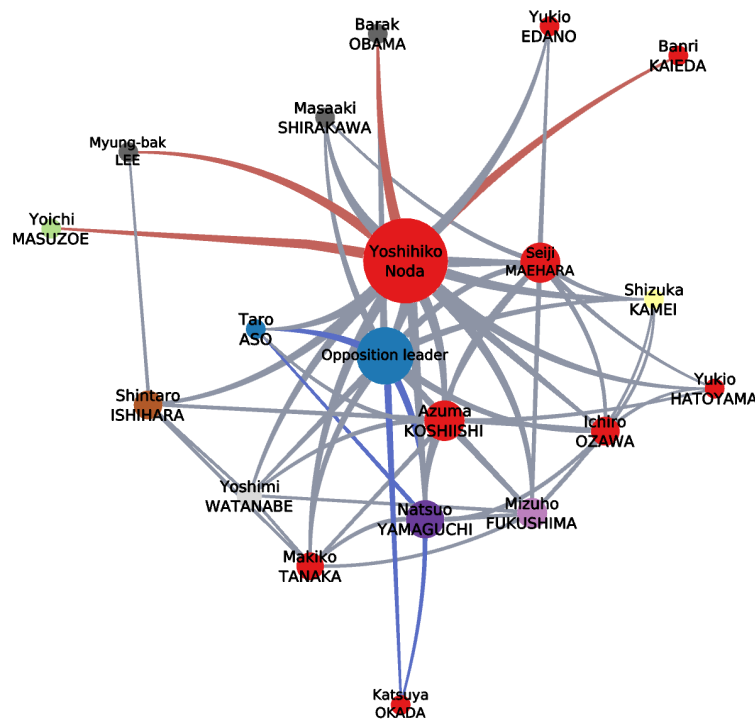








(h) Kan: 2010/06/08 – 2011/09/02



(i) Noda: 2011/09/02 – 2012/12/26

Figure 13 The networks are better seen zoomed. From Koizumi1 (a) to Noda (i), the networks with the Louvain communities highlighted – red and blue links depending on the PM’s and OL’s affiliation (respectively DPJ and LDP) – purple links if both figure belong to the same community – grey links are out of these communities. The size of a node represents its centrality, nothing encode the size/color of an edge.