



HAL
open science

Narrative smoothing: dynamic conversational network for the analysis of TV Series plots

Xavier Bost, Vincent Labatut, Georges Linarès

► To cite this version:

Xavier Bost, Vincent Labatut, Georges Linarès. Narrative smoothing: dynamic conversational network for the analysis of TV Series plots. 2016. hal-01276708v3

HAL Id: hal-01276708

<https://hal.science/hal-01276708v3>

Preprint submitted on 25 Feb 2016 (v3), last revised 24 Jan 2020 (v6)

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.



Distributed under a Creative Commons Attribution - NonCommercial - ShareAlike 4.0 International License

Narrative smoothing: dynamic conversational network for the analysis of TV Series plots

Xavier Bost
Laboratoire Informatique
d'Avignon, EA 4128, Avignon,
France
xavier.bost@univ-
avignon.fr

Vincent Labatut
Laboratoire Informatique
d'Avignon, EA 4128, Avignon,
France
vincent.labatut@univ-
avignon.fr

Georges Linarès
Laboratoire Informatique
d'Avignon, EA 4128, Avignon,
France
georges.linares@univ-
avignon.fr

ABSTRACT

Modern popular TV series often develop complex storylines spanning several seasons, but are usually watched in quite a discontinuous way. As a result, the viewer generally needs a comprehensive summary of the previous season plot before the new one starts. The automation of such a task requires to identify and characterize the dynamics of the series sub-plots. One way of doing so is to study the underlying social network of interactions between the characters involved in the narrative.

The standard tools used in the Social Networks Analysis field to extract such a network rely on an integration of time, either over the whole considered period, or as a sequence of several time-slices. However, they turn out to be inappropriate in the case of TV series, due to the fact the scenes showed onscreen alternatively focus on parallel storylines, and do not necessarily respect a traditional chronology. This makes existing extraction methods inefficient to describe the dynamics of relationships between characters, or to get a relevant instantaneous view of the current social state in the plot. This is especially true for characters shown as interacting with each other at some previous point in the plot but temporarily neglected by the narrative.

In this article, we introduce *narrative smoothing*, a novel, still exploratory, network extraction method. It smooths the relationship dynamics based on the plot properties, aiming at solving some of the limitations present in the standard approaches. In order to assess our method, we apply it to a new corpus of 3 popular TV series, and compare it to both standard approaches. Our results are promising, showing narrative smoothing leads to more relevant observations when it comes to the characterization of the protagonists and their relationships. It could be used as a basis for further modeling the interleaved storylines constituting TV series plots.

Keywords

TV series, plot analysis, dynamic social network.

1. INTRODUCTION

TV series became increasingly popular these past ten years. As opposed to older TV anthologies containing standalone episodes with self-contained stories, modern series tend to develop continuous, possibly multiple, storylines, spanning several seasons. However, the new season of a series is gen-

erally broadcast on a relatively short period: the typical dozen of episodes it contains is usually aired over a couple of months. In the most extreme case, the whole season is even released at once. Furthermore, modern technologies, like streaming or downloading services, tend to free the viewers from the broadcasting pace, often resulting in an even shorter viewing time (e.g. binge-watching). In summary, *modern* TV series are highly continuous from a narrative point of view, but are usually watched in quite a discontinuous way: no sooner is the viewer hooked on the plot than he has to wait for almost one year before eventually knowing what comes next.

The main effect of this unavoidable waiting period is to make the viewer forget the plot, especially when complex. Since he fails to remember the major events of the previous season, he needs a comprehensive recap before being able to fully appreciate the new season. Such recaps come in various flavors: textual synopsis of the plot illustrated by key-frames extracted from the video stream; extractive video summaries of the previous season content, such as the “official” recap usually introduced at the beginning of the very first episode of the new season; or even videos of fans reminding, when not commenting, the major narrative events of the previous season. Though quite informative and sometimes enjoyable, such content-oriented summaries of complex plots always rely on a careful human expertise, usually time-consuming. The question is therefore to know how this task can be partially or even fully automated.

To the best of our knowledge, few works in the multimedia processing field focused on automatically modeling the plot of a movie. In [7], the authors attempt to automatically detect the typical three-act narrative structure of Hollywood full-length movies: each of these three narrative segments is claimed to be characterized by specific stylistic patterns based on film grammar; by combining low level features extracted from the video content, the authors automatically exhibit the boundaries separating the three typical consecutive acts. Nonetheless, such a style-based approach does not provide any insight into the story content and focuses on a fixed narrative structure that generalizes with difficulty to the complex plots of *modern* TV series.

The benefits of Social Networks Analysis (SNA) for investigating the plot content of fictional works have recently been emphasized in several articles. Most focus on literary works: plays [10], novels [4, 1], comics [2, 6], mythology [9], etc. Some deal with the pre-processing steps needed for identifying the characters involved in the plot – a much trickier

task for novels than for theater or movie screenplays. By projecting the dynamics of the characters' interactions onto a graph, SNA helps unveiling some underlying patterns invisible to a closer reading. SNA-based approaches are even more recent and sparse in the context of multimedia works. Weng *et al.* presented a method to automatically analyze the plot of a movie [12, 13]; Ercolessi *et al.* [5] adopted a similar approach to cluster the scenes of two TV series into separate storylines. However, such works focus either on full length-movies or on standalone episodes of TV anthologies, where character interactions are often well-structured into stable communities. Their approaches consequently do not necessarily translate well when applied to *modern* TV series.

In this paper, we present an SNA-based method aiming at automatically providing some insight into the complex plots of TV series, while solving the limitations of the previous works. For this purpose, we not only consider standalone episodes or full-length movies with stable and well defined communities, but the complex plots of TV series, as they evolve over dozens of episodes. In this case, no prior assumption can be made about a stable, static community structure that would remain unchanged in every episode and that the story would only uncover, and we have to deal with evolving relationships, possibly temporarily linked into dynamic communities. In this case, we are left with building the current state of the relationships upon the story itself, which, by focusing alternatively on different characters in successive scenes, prevents us from monitoring instantaneously the full social network underlying the plot. We thus propose to address this problem by smoothing the sequentiality of the narrative, resulting in an instantaneous monitoring of the current state of any relation at some point of the story.

Our main contributions are the following. The first is *narrative smoothing*, the method we propose for the extraction of dynamic social networks of characters. The second is the annotation of a corpus of 61 TV series episodes from 3 popular TV shows: *Breaking Bad*, *Game of Thrones*, and *House of Cards*. The third is a preliminary evaluation of our framework on these data, and a comparison with existing methods. Both tasks rely on the fictional social networks extracted from our corpus, which are publicly available online (see Section 4).

The rest of the article is organized as follows. In Section 2, we review in further details the previous works related to SNA-based plot identification. Then, in Section 3, we describe our proposed method. We first focus on the way the verbal interactions between characters are estimated, before detailing the way a dynamic view of the relationships in the plot of a TV series can be built independently from the narrative pace. In Section 4, we illustrate how our tool can be used by applying it to the three mentioned TV series, and we compare the obtained results to existing methods.

2. PREVIOUS WORKS

In our review, we distinguish works considering a static network corresponding to the temporal integration over the whole considered period, which we call *complete aggregation*, and those extracting and studying a dynamic network constituted of a sequence of smaller integration called *time-slices*.

2.1 Complete Aggregation

Cumulative networks were widely used when attempting

to apply SNA for analyzing the plot of fictional works. The interactions are iteratively inserted as edges in the network of characters. They are possibly weighted and even directed, resulting in a static graph agglomerating every past relationship, whatever their time ordering.

In [10], Moretti underlines and illustrates the light SNA can shed on literary works, either plays or novels. By projecting the time of the character interactions onto the plane of a graph, SNA helps unveiling some underlying patterns invisible to a closer reading. Moretti agglomerates the conversational interactions between characters in Shakespeare's *Hamlet*, and exhibit the contrast between the Court, densely connected and the emerging modern State, weakly connected around Horatio.

In [12, 13], relying on similar observations, Weng *et al.* make use of SNA to automatically analyze the plot of a movie. The social network of characters (denoted "RoleNet") is built as follows. They first manually characterize the scenes by their boundaries and the characters they involve. They then hypothesize an interaction between two characters as soon as they both appear within the same scene. The network is obtained by representing characters as nodes and their interactions by links. These links are weighted according to the number of scenes in which they co-appear, resulting in a *cumulative* representation of time. The authors analyze this network through community detection (i.e. partition of the node set). They apply this approach to so-called "bilateral movies", which involve only two major characters, each of them central in his own community. In [12], the *RoleNet* is used for further investigating the plot, by classifying scenes into one of the two storylines constituting a bilateral movie. In [13], an extended version of the network, without any prior assumption about the number of communities involved, is used as a basis for automatically detecting breakpoints in the story: a narrative breakpoint is assumed if the characters involved in successive scenes are socially distant in the network of characters accumulated over the whole story.

In [5], a similar network of interacting speakers is used, among other features, for clustering scenes of two TV series episodes into separate storylines, defined as homogeneous narrative sequences related to major characters. A standard community detection algorithm is applied to the network of speakers, as built upon each episode, before the social similarity between any pair of scenes is computed, as a relevant high-level feature for clustering scenes into sub-stories.

In summary, cumulative networks can be used as a reliable basis for automatically or manually analyzing the plot of fictional works with well-defined communities, as in dramas, full length-movies or standalone episodes of TV anthologies¹. But for TV series with complex, evolving and possibly parallel storylines, such a static approach is not appropriate. Indeed, a cumulative network built over a long period of time, as in modern TV series, is relatively dense and does not allow extracting meaningful information. Moreover, communities in the final cumulative network undoubtedly always correspond to sub-stories, partially disconnected in the narrative, but the opposite does not generally stand. Some individuals may have been strongly connected to each other at some point of the story, before some of them interact with other

¹The website <http://www.moviegalexies.com> provides a way of visualizing such static, cumulative interaction networks for a database of about 700 movies.

people for some time, resulting in a second sub-story. Once agglomerated in the cumulative network, such changes in the interaction patterns may be obscured. In some extreme cases, distinct narrative sequences may even result in a complete cumulative graph, for instance in the interaction pattern that follows:

$$s_{12}^{(1)} \dots s_{12}^{(2)} s_{13}^{(3)} \dots s_{13}^{(4)} s_{23}^{(5)} \dots s_{23}^{(6)}$$

where $s_{ij}^{(t)}$ denotes the fact that the i^{th} and j^{th} characters are the only interacting speakers in the t^{th} episode. The three consecutive interaction sequences result in a triangular interaction pattern unable to reflect the three corresponding sub-stories.

2.2 Time-slices

Some works attempt to take into account the evolution of the social network of the characters when analyzing the plot of fictional works. In [1], the authors emphasize the limitations of the static, cumulative graph when analyzing the centrality of the various characters of the novel *Alice in Wonderland*. A dynamic view of the social network is then introduced, by building successive static views of the network in every chapter, before standard centrality measures are separately computed in each of them and traced over time for some major characters. Each view corresponds to a so-called *time-slice*.

Though widely used when considering the evolution of general networks (i.e. not necessarily narrative ones) over time [8], time-slice networks, as resulting from the differentiation over some time step of the cumulative network, may still be problematic. In [3], the authors focus on the critical issue of the time-slice duration, called “snapshot rate”. It must be chosen carefully to allow capturing a sufficient amount of interactions, but not too many, otherwise one may obtain irrelevant network statistics. The authors then describe a way of automatically estimating the natural time-slice for monitoring over time the evolution of a network of daily contacts in a professional context.

In order to monitor the plot of TV series and allow further analysis, such a time-slice should be short enough to capture punctual narrative events related to the social network of characters, but long enough to provide a comprehensive view of the relationships state at any point of the story. Unfortunately, getting such a snapshot of the current state of the relationships between the protagonists within the plot turns out to be particularly challenging.

As a smoother alternative for monitoring the state of the network over time, [11] applies temporal decay to the past occurrences of the relationships between the characters of Shakespeare’s plays for monitoring their evolution over time.

Unlike the network of physical contacts described in [3], the state of the relationships within a story is not fully monitored at any moment, but has to be inferred from the story itself. The narrative usually focuses alternatively on some relationships, possibly belonging to parallel storylines, and only provides a partial view on the network’s current state. Some relationships may even take place at the same moment in different places, but will be shown sequentially in successive scenes. Fig. 1 illustrates the typical sequential nature of the story as being narrated: three disjoint sets of interacting speakers, possibly at the same time but in different places, are shown sequentially in the story in three

successive scenes.



Figure 1: Three different sets of interacting characters from three consecutive scenes.

As a consequence, the temporalness of the narrative may be quite different from the temporalness of the underlying network: in particular, the mere fact that a group of mutually interacting characters temporarily disappear from the story does not imply that the corresponding relationships disappeared from the network. The narrative focus on those relationships may only have been postponed by the filmmaker. Furthermore, the pace of activation of the relationships occurring in different regions of the interaction network remains largely unpredictable, especially when multiple, disjoint storylines take place in parallel within the narrative. Fig. 2 plots the scene occurrences of three major character-based storylines in the first two seasons of *Game of Thrones*. Except in the very beginning of the first season, where Jon and Tyrion meet each other, the three characters interact within well separated communities.

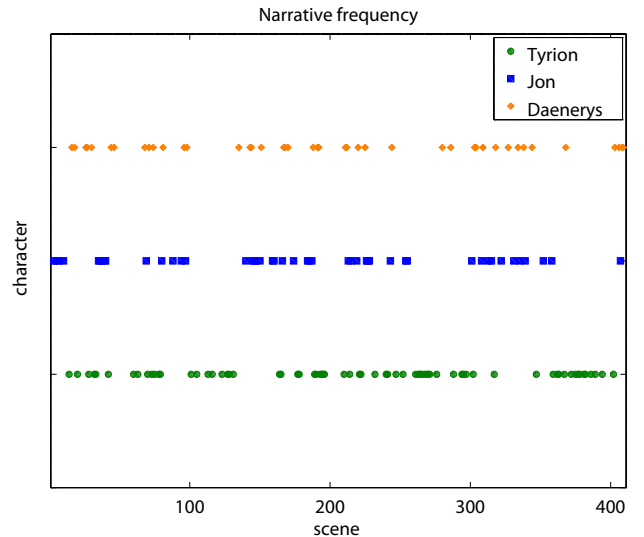


Figure 2: Narrative frequency of three character-based storylines in the first two seasons of *Game of Thrones*.

As can be seen, the way the story alternatively activates these three major storylines does not seem to follow some regular patterns. In such a case, the “ideal” time-slice may be tricky to set. If too large, it will possibly mask the fast changes usually occurring in the most frequently activated storyline, i.e. the story centered around Tyrion. If too narrow, it would lead to irrelevant interpretations of the narrative disappearance of some groups of relationships: the absence of Jon Snow’s storyline from the scene 250 up to the scene 300 does definitely not imply that he does not remain socially active in the meantime in his own community. Therefore, the sequential nature of the story should prevent

us to identify the time of the narrative to the time objectively affecting the social network it sequentially unveils.

In the rest of this paper, we introduce a novel way of building the dynamic network of interactions between the characters of TV series that allows to fully capture the instantaneous state of every relationship at any point of the story, whatever the pace of activation of each storyline in the narrative.

3. METHODS

We now describe the two steps constituting our method. First, we explain how we identify and characterize interactions between TV series characters. Second, we describe how we extract a smoothed dynamic network from a set of interactions.

3.1 Estimating Verbal Interactions

In this work, we focus on relationships defined in a *strong sense*, as based on personal, verbal, interactions between characters. The resulting network can thus be considered as a *conversational network*, in contrast to the co-occurrence network of characters described in [12, 13] and used in [5].

We first manually annotated the scenes boundaries: similarly to the rule of the three unities classically prescribed for dramas, a scene in a movie is defined as a homogeneous sequence of actions occurring at the same place, within a continuous period of time. The characters co-appearing in a single scene are thus supposed to interact with one another. However, if being at the same place at the same time is usually required to consider that several persons interact, it is rarely sufficient. Fig. 3 shows two consecutive dialogues extracted from the TV series *House of Cards*, and belonging to the same scene. Three speakers are involved, but without any interaction between the second (*D. Blythe*) and the third (*C. Durant*). The first speaker (*F. Underwood*) talks to *D. Blythe* in the first sequence, then moves to *C. Durant* and starts discussing with her.



Figure 3: Two consecutive dialogue sequences within the same scene.

Instead of globally considering the scene unit, we chose to tackle this problem by identifying the verbal interactions upon the sequence of speech turns in each scene, once manually labeled according to the corresponding speakers. In order to estimate the verbal interactions from the single sequence of utterances, we apply four basic heuristics:

Rule (1): Surrounded speech turn. We consider that a speaker s_2 talks to another speaker s_1 if he speaks both after and before him, resulting in a speech turns sequence $s_1 s_2 s_1$, where each speech turn is labeled according to the corresponding speaker. Fig. 4a shows the subgraph resulting from the application of Rule (1) to the speech turns sequence from Fig. 3, where each edge is labeled according to the

number of times each speaker is considered as talking to another one.

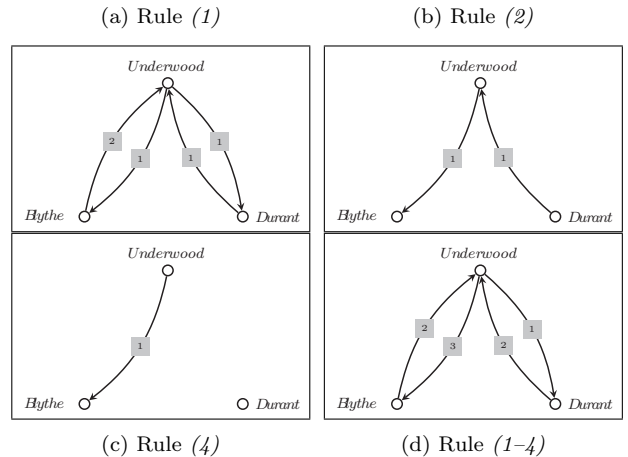


Figure 4: Verbal interactions estimated from the separate and joint application of Rules (1, 2, 4) to the speech turn sequence shown on Fig. 3.

Rule (2): Starting and ending speech turns. This rule aims at processing the first and last utterances of each sequence $s_1 s_2 \dots s_3 s_4$ of speech turns, by adding two links $s_1 \rightarrow s_2$ from the first to the second speaker and $s_4 \rightarrow s_3$ from the fourth to the third one. The network resulting from the application of Rule (2) to the sequence of Fig. 3 is shown on Fig. 4b.

The last two rules are introduced to process ambiguous sequences of the type $s_1 s_2 s_3$, where three consecutive speech turns originate in three different speakers: in such cases, the second speaker could be stated to talk to the first one as well as to the third one, or even to both of them. However, such speech turns sequences can often be disambiguated by considering speakers preceding and following the sequence.

Rule (3): Local disambiguation. We distinguish 2 variants of this rule. Rule (3a) applies when the second speaker appears before the sequence, but not after, as in $(s_2) s_1 s_2 s_3 (s_4)$. We then consider he speaks to s_1 rather than to s_3 . Symmetrically, Rule (3b) concerns the case when the second speaker appears after, but not before the sequence, as in $(s_0) s_1 s_2 s_3 (s_2)$, and is therefore supposed to speak to s_3 .

Rule (4): Temporal proximity. When the second speaker is involved in the conversation both before and after the ambiguous sequence, as in $(s_2) s_1 s_2 s_3 (s_2)$, we consider the ambiguous speech turn to be intended for the speaker whose utterance is temporally closer. In the sequence shown on Fig. 3, the fifth, ambiguous utterance would then be hypothesized as intended for the first speaker *D. Blythe*, resulting in the additional link shown on Fig. 4c. The same Rule (4) is applied when the speaker s_2 is not involved in the immediate conversational context.

Fig. 4d shows the amount of directed interactions between any two speakers involved in the scene shown on Fig. 3. For the sake of simplicity, the interactions are evaluated here as a number of interactions from one speaker to another one, but can as well be expressed in terms of interaction duration.

3.2 Extracting the Dynamic Network

We obtain the total amount of interaction $h_{ij}^{(t)}$ between the speakers s_i and s_j in the t^{th} scene by summing up the amount of speech flowing in the scene t from s_i to s_j and from s_j to s_i , resulting in an undirected local interaction amount, possibly null, expressed in seconds.

As stated in Section 1, we would like to get an instantaneous measurement of the strength of any relationship at any moment, but from the successive partial views of the underlying network that the narrative provides us. Intuitively, a particular relationship may be considered as especially important at some point of the story if the involved characters both speak frequently and a lot to each other: the time interval needed before the interaction is reactivated in the narrative is expected to be short, and the interaction time to be long whenever the relationship is active in the plot.

Four possible states have to be considered when monitoring a single relationship over time: 1) the relationship is active in the current scene; 2) it has been active in the story and will be active again later; 3) it was active before, but will no longer be active in the narrative; and 4) it has not yet been active in the narrative.

The first case is the simplest one: each time the interaction occurs, its strength can be estimated in a standard way as the duration of the interaction, expressed in seconds: at any scene t where speakers s_i and s_j are hypothesized as talking to each other, the instantaneous weight of their relationship $w_{ij}^{(t)}$ is estimated as follows:

$$w_{ij}^{(t)} = h_{ij}^{(t)} \quad (1)$$

where $h_{ij}^{(t)}$ denotes the interaction time between the i^{th} and j^{th} speakers in scene t .

The last three cases are much more tricky. Between two consecutive occurrences of the same relationship in the story, it would be tempting to consider that the relationship is still (resp. already) active if it is recent (resp. imminent) enough at each moment considered. In [11], the author applies temporal decay to the past occurrences of the interactions between characters in Shakespeare's plays in order to visualize their evolution over time. According to the time-slice framework described in Section 2, as long as the relationship is present in the observation window of the network over time, it is stated as active, and inactive as soon as no longer observed.

As emphasized in Section 2, such a way of handling the past and future occurrences of the relationships is inappropriate for many modern TV serials. Some interacting characters may be absent of the narrative for an undefined period of time but still linked in the underlying network, as confirmed by the fact that the last state of the relationship is generally used as a starting point when the characters are re-introduced in the story. Indeed, the temporalness of the narrative should affect a relationship only when at least one of the involved characters interacts with others after and/or before the relationship is active: the relationship between two characters should only get weaker if they interact separately with others before interacting again with one another.

In order to perform such a narrative smoothing, we introduce two quantities to handle the scenes where the two characters do not interact. First, $\Delta_{ij}^{(l)}(t)$ is the *narrative persistence* between speakers s_i and s_j considered at scene t . It is defined relatively to the last scene in which their

relation was active, noted l :

$$\Delta_{ij}^{(l)}(t) = h_{ij}^{(l)} - \sum_{t'=l+1}^t \sum_{k \neq i,j} (h_{ik}^{(t')} + h_{jk}^{(t')}) \quad (2)$$

This measure $\Delta_{ij}^{(l)}(t)$ corresponds to the net balance between the duration of the last interaction occurrence $h_{ij}^{(l)}$ and the conversational time (represented by the double sum) the two characters i and j have devoted separately to other characters k until their relation occurs anew.

Symmetrically, $\Delta_{ij}^{(n)}(t)$ is the *narrative anticipation* between speakers s_i and s_j considered at scene t . It is defined relatively to the next scene in which their relation will be active again, noted n :

$$\Delta_{ij}^{(n)}(t) = h_{ij}^{(n)} - \sum_{t'=t}^{n-1} \sum_{k \neq i,j} (h_{ik}^{(t')} + h_{jk}^{(t')}) \quad (3)$$

We then define the instantaneous weight $w_{ij}^{(t)}$ of the relationship between the speakers s_i and s_j in any scene t occurring between two consecutive occurrences of their relationship as:

$$w_{ij}^{(t)} = \max(\Delta_{ij}^{(l)}(t), \Delta_{ij}^{(n)}(t)) \quad (4)$$

If neither of the two characters speaks to others before they interact again with one another, $w_{ij}^{(t)} = \max(h_{ij}^{(l)}, h_{ij}^{(n)})$ and the last (resp. next) occurrence of the relation is considered as still (resp. already) fully present in the network, whatever the number of intermediate scenes the narrative introduces in-between to focus on other parts of the plot.

The weight of the relationship between the i^{th} and j^{th} speakers in any scene t occurring after its very last occurrence in the narrative is expressed as follows, provided that one of the two characters remains involved in the story by interacting with others:

$$w_{ij}^{(t)} = \Delta_{ij}^{(l)}(t) \quad (5)$$

Symmetrically, the weight of the relationship between the i^{th} and j^{th} speakers in any scene t occurring before its first occurrence in the story is computed as follows, as long as one of the two characters has already been shown as interacting with other people:

$$w_{ij}^{(t)} = \Delta_{ij}^{(n)}(t) \quad (6)$$

In the very last case, when neither of the two characters is still (resp. already) active, the weight w_{ij} is set to $-\infty$.

We then normalize the weights of the interactions linking any couple of characters in some scene t . We use the following formula, resulting in an undirected graph $\mathcal{G}^{(t)}$, capturing the instantaneous state of the social network that the story sequentially unveils:

$$n_{ij}^{(t)} = \frac{1}{1 + e^{-\lambda w_{ij}^{(t)}}} \quad (7)$$

where $n_{ij}^{(t)}$ is the normalized weight of the relationship between the speaker s_i and s_j .

The choice of the sigmoid function for such a normalization purpose both allows to get weights ranging from 0 to 1 and to simulate the way the past and future states of a relationship in the narrative could influence its current state at some point t . The parameter λ is a parameter of sensitivity

to the past and future states of the network and was set to $\lambda = 0.01$ (high values imply low dependence on the future and past states).

Fig. 5 shows four excerpts of four consecutive scenes in *House of Cards*, involving five individuals. The first two of them, namely *Francis Underwood* and his wife *Claire*, interact with each other in the first and last scenes (red border) respectively during 30 and 20 seconds, whereas Claire interacts in-between 40 seconds with another person in the second scene (green border) and two other people are talking to one another in the third scene during 50 seconds.



Figure 5: Example of application of the weighing scheme to a specific relationship.

In the first and fourth scenes, Claire and Francis are interacting with each other: according to Equation 1, the weights of their relationship are then set to the corresponding interaction times, namely 30 in the first scene and 20 in the fourth one. In the second scene, the relationship between Claire and Francis is weakened by the fact Claire interacts 40 seconds with someone else, and decreases to -10 ($\max(30 - 40, 20 - 40)$) according to Equation 4. The remaining third scene involves other characters and does not impact, whatever its duration, the specific relationship between Francis and Claire, still amounting at that point to -10 . The resulting sequence of unnormalized, instantaneous weights for the relationship between Claire and Francis is then $(30, -10, -10, 20)$ at the four considered moments.

4. EXPERIMENTS AND RESULTS

In this section, we evaluate qualitatively narrative smoothing, our graph extraction method, by comparing it to both types of methods described in Section 2. For this purpose, we focus on three recent and popular TV series, and explore their plots from the dynamics of their underlying social network of characters. In the rest of this section, we first describe our corpus, then analyze the obtained networks from the perspective of protagonists (nodes) and their relationships (links).

4.1 Corpus

Our corpus relies on three very popular TV series: *Breaking Bad* (BB), *Game of Thrones* (GoT), and *House of Cards* (HoC). More precisely, we used the first two seasons of each of these TV shows. We manually annotated the scene boundaries and labeled each subtitle according to the corresponding speaker. The obtained annotations were then used to extract the social networks of characters, by first estimating the verbal interactions according to the rules described in Subsection 3.1 and then by using the existing methods presented in Section 2 as well as our own narrative smoothing approach. The resulting networks are publicly available online on FigShare², along with pictures representing the cumulative networks, and short videos showing the evolution of the three character networks over the first two seasons of each series. Table 1 reports the main features of the resulting corpus.

²<https://dx.doi.org/10.6084/m9.figshare.2199646>

Table 1: *Corpus: main features.*

Corpus	BB	GoT	HoC
# episodes	20	20	26
Total duration (hours)	$\simeq 15$	$\simeq 17$	$\simeq 19$
Speech duration (seconds)	23,403	27,616	39,175
# subtitles	11,544	14,204	21,005
# scenes	402	459	912
% spoken scenes	95.03	97.17	97.70
# speakers/scene (avg.)	2.38	2.97	2.49
# speakers/scene (std. dev.)	1.16	1.53	1.12

Speech is uniformly distributed over the scenes, with in average more than 95% of the scenes containing at least one subtitle, which suggests that most social interactions are expressed verbally in these three TV series. Furthermore, the average number of speakers by scene remains quite low (ranging from 2.38 to 2.97 depending on the TV series), often resulting in the simple patterns of verbal interactions properly handled by applying the basic heuristics described in Subsection 3.1.

4.2 The protagonists

We first base our analysis on the protagonists of the considered TV series, i.e. the nodes in the corresponding extracted social networks. We present only a small number of results, which concern characters of particular interest. We characterize them using the *node outgoing strength*, a generalization of the node degree defined as the sum of the weights of the links originating from the considered node. In our case, weights correspond to spoken interaction durations, so the strength of a character is related to how much he speaks to others.

We first focus on Walter White, the main character of *Breaking Bad*, and Tuco Salamanca, one of the drug dealers with whom he is in business. The cumulative network (as described in Section 2.1), i.e. the temporal integration over the first 20 episodes is represented in Figure 11 (Appendix). In this network, the strength of Walter White (his total interaction time with others) is about twenty times as large as the strength of Tuco: 12,332 seconds for Walter (rank 1) vs. 590 for Tuco (rank 11). By comparison, Fig. 6 displays the evolution of their instantaneous strengths, obtained with our narrative smoothing method, as a function of the scenes ordered chronologically. This leads us to a completely different vision of Tuco’s role in the plot. As Fig. 6 shows, from scene 100, his importance is constantly increasing and even overcomes the importance of the main protagonist for some time, before suddenly decreasing after scene 130. This clearly corresponds to a subplot, or a short narrative episode, ending with Tuco’s death, at the end of scene 135 (vertical line on Fig. 6).

We now switch to Daenerys Targaryen and Tyrion Lannister, two major protagonists of *Game of Thrones*. Fig. 7 shows how their strengths evolve over the first two seasons of the series, again as a function of the chronologically ordered scenes. The first two plots were obtained through the use of fixed-size observation windows, set to 10 scenes (around half an episode) for the first and 40 (about two episodes) for the second. The last plot relies on our narrative-smoothing method. The appearance of Daenerys’ storyline has a relatively slow pace in these seasons (Fig. 2). When the window

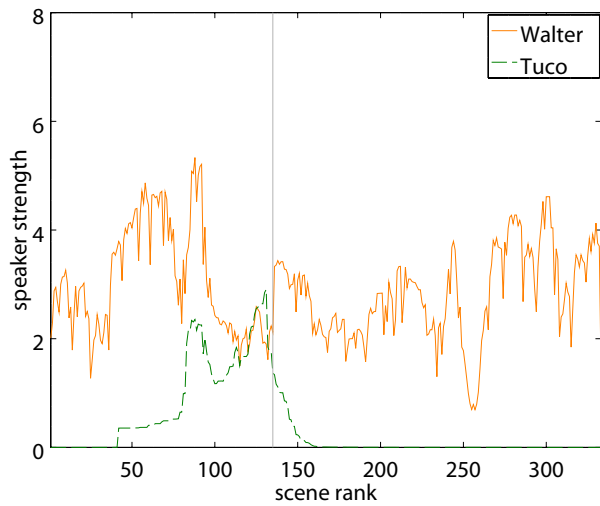


Figure 6: Strengths of two important characters in Breaking Bad, plotted as a function of the scenes.

is too narrow, this creates noisy, irrelevant measurements of her narrative importance (first plot on Fig. 7). It appears very unstable because her storyline alternates with many others on the screen. A wider observation window (second plot of the same figure) is more likely to cover successive occurrences of Daenerys in the narrative, but, unlike our narrative smoothing method, prevents us from locating precisely the scenes responsible for Tyrion's current importance: for instance, a local maximum in Tyrion's strength is reached at scene 220 (third plot on Fig. 7), just after a major narrative event took place – the nomination of Tyrion as the King's Counselor (vertical line). Such an event remains unnoticed when accumulating the interactions during too large time-slices (second plot on Fig. 7).

Fig. 7 also reveals an important property of our way of building the dynamic network. Because the past (resp. future) occurrences of a particular relationship are still (resp. already) active as long as the involved characters do no interact with others in the meantime, the respective strengths of the main characters of the story appear remarkably balanced. Whereas Tyrion looks much more central than Daenerys in the time-slice based dynamic networks, whatever the size of the observation window, Daenerys is nearly as central as Tyrion in the network based on our narrative smoothing method: few of her acquaintances are shown onscreen as interacting with others, whereas the story, by focusing more frequently on Tyrion, also unveils more extensively his social network, even when active without him (especially after scene 220).

Our results confirm that cumulative networks, by neglecting the temporal dimension, tend to completely miss punctual changes in the importance of certain characters relatively to the plot. The time-slice based methods can handle the network dynamics, however our observations illustrate that they cannot properly tackle the narrative issue we described in Subsection 2.2. The choice of an appropriate time window, is a particularly sensitive point. By comparison, narrative smoothing captures the state of a relationship at any moment of the plot, using a time scale which directly depends on the narrative pace of the considered series. This

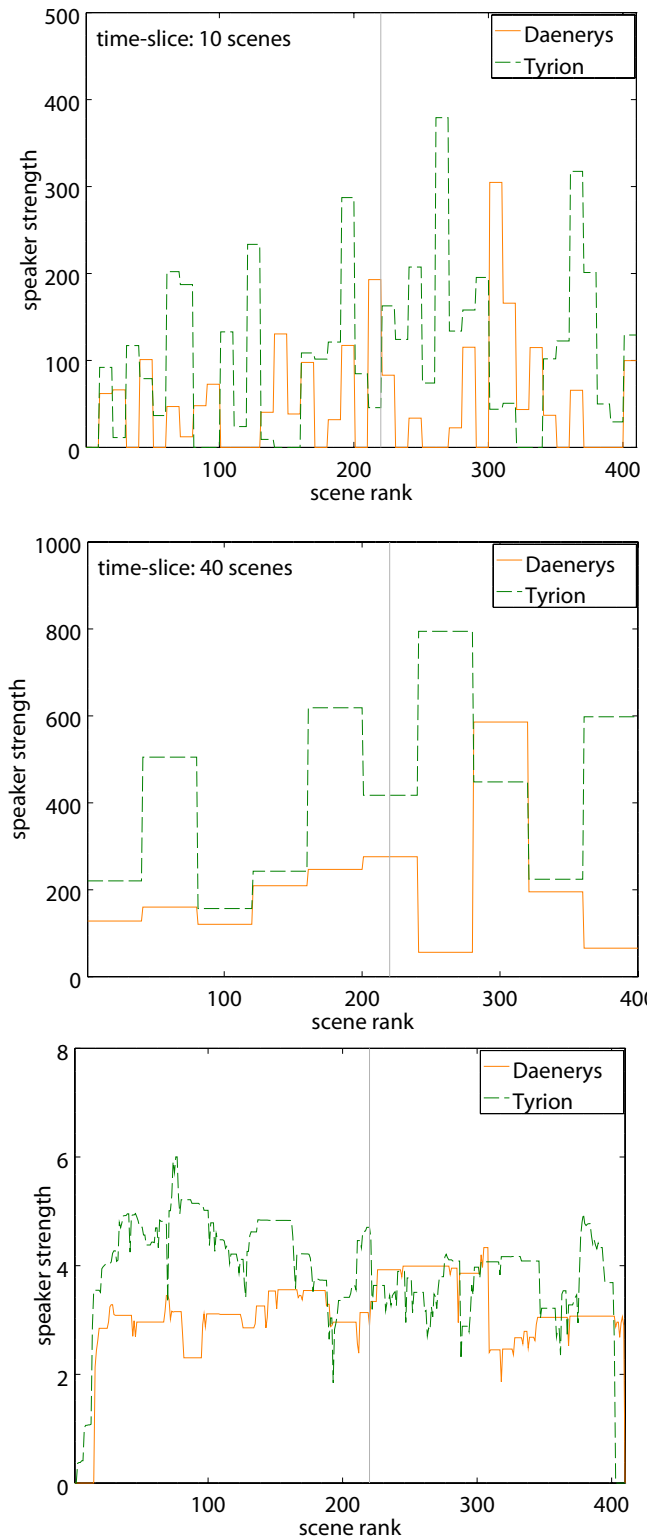


Figure 7: Strength of two major characters of *Game of Thrones* plotted as a function of the chronologically ordered scenes. From left to right: 10 and 40 scenes time-slices, and narrative smoothing.

allows to finely evaluate the degree of instantaneous involvement of any character in the plot.

4.3 The relationships

We now consider relationships between pairs of characters, instead of single individuals. We characterize each relation depending on its weight, i.e. the amount of time the characters talked to each other, either cumulated over time-slices, possibly consisting of the whole set of episodes, or smoothed with respect to the narrative. Like in the previous subsection, we focus on relationships of particular interest.

Let us consider two relationships in *House of Cards*, representative of two sub-stories: the first one corresponds to a narrative sequence in the storyline related to the main character Francis Underwood – his fight with a former ally, the unionist Martin Spinella; the second one is a similar subplot, but related to a secondary character, not as frequently present in the narrative, the journalist Lucas Goodwin, who requests the help of the hacker Gavin Orsay to investigate on Francis. Though locally important in these two sub-stories, neither of these relationships lasts long enough to be noticed in the cumulative network, as resulting from the first two seasons of the series (cf. Fig. 13 in the Appendix): the interaction time amounts to 562 seconds for the relation between Francis and Martin, and to 294 seconds for the relation between Gavin and Lucas. These total interaction times remain quite small compared to the central relation between Claire and Francis, amounting to 2,319 seconds.

Nonetheless, once plotted as a function of the chronologically ordered scenes (Fig. 8), the respective weights of these relationships in the narrative look quite different, whatever the weighting scheme. Both sub-stories, the one based on the relation between Francis and Martin and the one based to the relation between Lucas and Gavin, turn out to be locally as important as the long-term sub-story based on the relation between the two main characters Claire and Francis.

Furthermore, all three ways of monitoring these relationships over time are not equivalent: agglomerating the interactions within short time-slices (first plot on Fig. 8) makes us miss the continuity of Lucas/Gavin’s sub-story, which occurs *in the narrative* at a slower rate than the sub-stories related to Francis. Conversely, large time-slices (second plot on Fig. 8) allow to capture this sub-story, but agglomerate the two main stages of the relation Francis/Martin: before becoming an enemy, Martin is first an ally of Francis; these two parts in the relation correspond to well separated stages in the narrative, that too large time-slices tend to merge, whereas the separation remains clear when using our narrative smoothing method (materialized by a vertical line on the third plot of Fig. 8).

Our results confirm that cumulative network are inappropriate when attempting to capture punctual sub-stories supported by specific relationships. Moreover, though much more appropriate to such a task, the time-slice approaches suffer from a major drawback: once fixed, the time slice cannot adapt to the variable rates at which the sub-stories appear in the narrative. By overcoming the narrative contingencies, our narrative smoothing approach allows to monitor more accurately over time any relationship, whatever the way the narrative focuses on it.

Fig. 9 illustrates the possible insight our narrative smoothing approach can give on important sub-stories.

As can be seen, two relationships between three characters

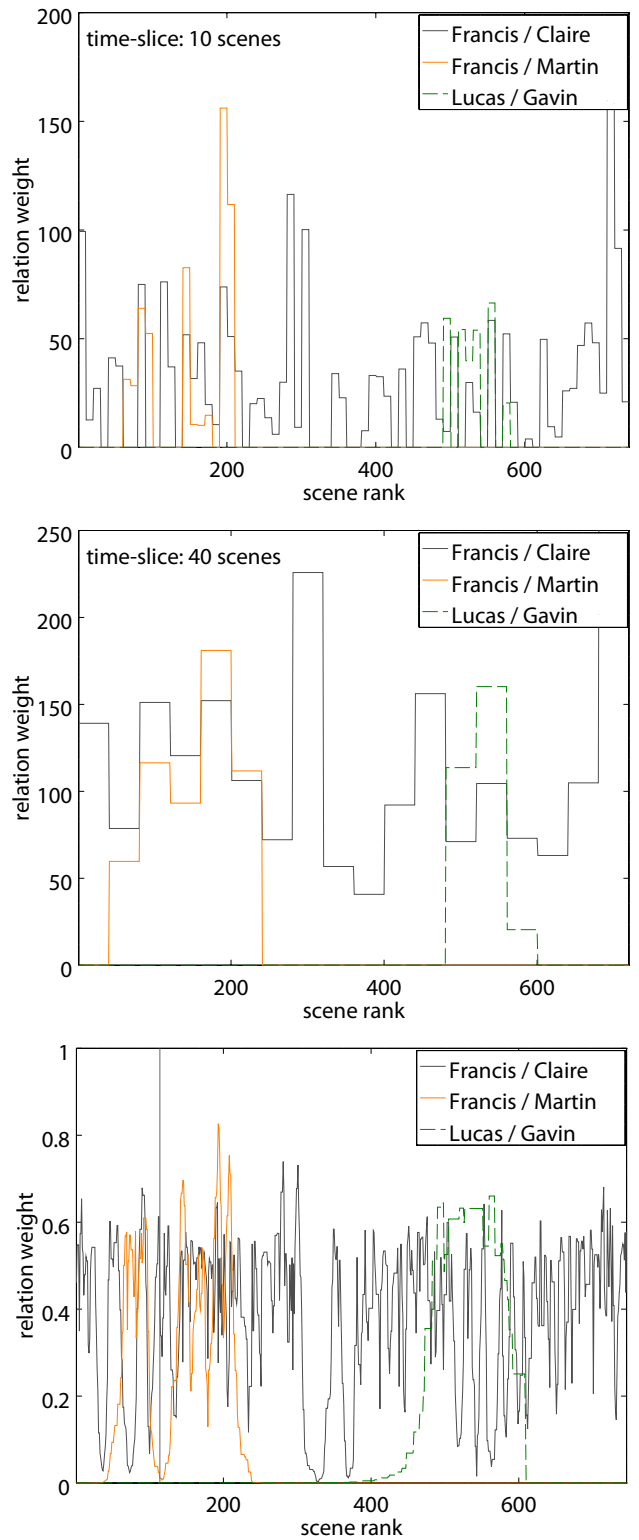


Figure 8: Weight of three relationships between five characters of *House of Cards* plotted as a function of the chronologically ordered scenes. From left to right: 10 and 40 scenes time-slices, and narrative smoothing.

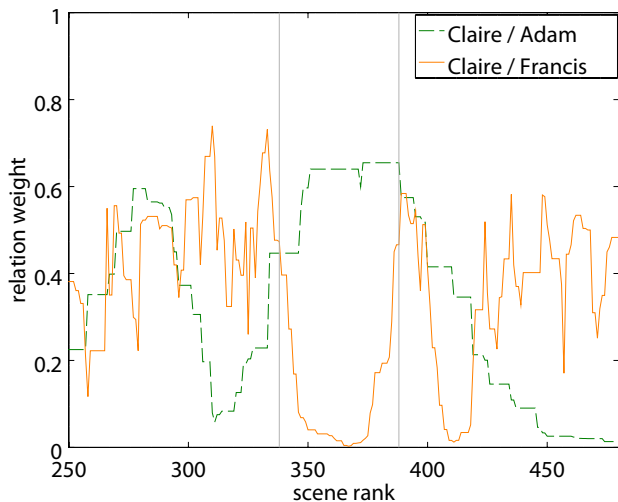


Figure 9: Weight of two relationships between three characters of *House of Cards* plotted as a function of the chronologically ordered scenes and based on narrative smoothing.

of *House of Cards* are considered: the relationship between Claire and her husband on the one hand and the relationship between Claire and her former lover Adam Galloway on the other hand. In the cumulative graph (Fig. 13), the interaction time between the two legal spouses is far more important than between Claire and Adam. Nonetheless, once considered over the narrative, it is clear that the relationship between Claire and Adam is locally much stronger than with her husband: between the scenes 338 and 388 (vertical lines on the plot), both relationships are complementary and suggest a specific sub-plot where Claire is much closer to Adam than to Francis.

Let us go back once again to *Game of Thrones* and its complex plot. Fig. 10 focuses on two relationships between three characters: Catelyn Stark and Ned Stark on the one hand, Catelyn Stark and Tyrion Lannister on the other.

Neither of these relationships would be considered as a major one from the cumulative graph at the end of the first two seasons (cf. Fig. 12 in the Appendix). Nevertheless, once dynamically considered, they both correspond to two successive sub-stories in the first season of *Game of Thrones*. As can be seen, our narrative smoothing approach even allows to separate two steps in the relationship between Catelyn and Ned: the first step of their relationships takes place in Winterfell; Ned then leaves Catelyn there and goes on his own to Kings Landing, freshly named as the King’s Counselor (around the scene 65), before Catelyn joins him there. Catelyn and Tyrion start interacting with each other after Catelyn leaves Kings Landing to Winterfell and is well preserved once monitored according to our method, though shown in the narrative at a quite slow and irregular pace.

5. CONCLUSION AND PERSPECTIVES

In this paper, we described a novel way of monitoring over time the state of the relationships between characters involved in the usually complex plots of modern TV series. The two methods previously used for this purpose are the cumulative approach, consisting in integrating every relation over the whole considered time period, and the time-

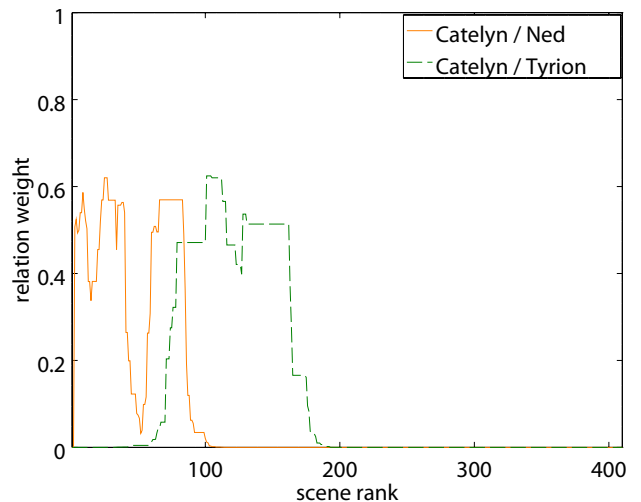


Figure 10: Weight of two relationships between three characters of *Game of Thrones* plotted as a function of the chronologically ordered scenes and based on narrative smoothing.

slice approach, consisting in breaking down the time-line into smaller discrete chunks. The first one turns out to be relatively inefficient for investigating complex storylines and a dynamic perspective is more appropriate. The second one complies with this constraint, but defining an appropriate size for the observation window is a very difficult task and constitutes a major drawback. Moreover, the current state of the interactions between any characters in a movie is not fully reachable, and consequently has to be inferred from the story itself. Due to its sequential nature, the narrative makes us alternatively focus on some particular relationships that take place in successive scenes. As a consequence, the story provides at some point a very narrow view on the underlying social network of characters, and the absence in the current scene of some past relationship can usually not be interpreted as a real disappearance, which invalidates the time-slice approach.

To address this issue, we chose to smooth the narrative sequentiality, by considering that the relation between some speakers remains active as long as neither of them interacts with others; if so, such external interactions result in a progressive dissolution of the past link. Symmetrically, the imminence of the next occurrence of the relationship has to strengthen the link. We experimentally compared our method, which we call narrative smoothing, to both mentioned approaches, on three recent popular TV series. Though exploratory and qualitative, our results show that our method, leads to more relevant results than both other methods, when it comes to instantaneously monitoring the importance of a particular character or of a specific relationship at some point of the story.

The way some characters temporarily aggregate at some point of the story in a community-like structure suggests that some narrative sequences result in the stabilization, possibly temporarily, of certain areas in the network. By automatically detecting such a narrative stabilization of some groups of relationships, it should be possible to split the whole story into sub-stories, without assuming a static, pre-defined, community structure. Finally, the statistical prop-

erties of such a dynamic network, as based on the smoothing of the narrative, have still to be studied: the relative balance between the characters stated as important suggests for instance that the traditional heavy-tailed distribution of the nodes degree or strength may not stand in this case.

6. APPENDIX

Figures 11, 12 and 13 show the cumulative networks of Breaking Bad, Game of Thrones and House of Cards, respectively, extracted over their first two seasons.

7. REFERENCES

- [1] A. Agarwal, A. Corvalan, J. Jensen, and O. Rambow. Social network analysis of alice in wonderland. In *NAACL - CLL Workshop*, 2012.
- [2] R. Alberich, J. Miro-Julia, and F. Rossello. Marvel universe looks almost like a real social network. *arXiv*, cond-mat.dis-nn:cond-mat/0202174, 2002.
- [3] A. Clauset and N. Eagle. Persistence and periodicity in a dynamic proximity network. *arXiv preprint arXiv:1211.7343*, 2012.
- [4] D. K. Elson, N. Dames, and K. R. McKeown. Extracting social networks from literary fiction. In *48th annual meeting of the ACL*, pages 138–147, 2010.
- [5] P. Ercolessi, C. Sénac, and H. Bredin. Toward plot de-interlacing in tv series using scenes clustering. In *IEEE CBMI Workshop*, pages 1–6, 2012.
- [6] P. M. Gleiser. How to become a superhero. *Journal of Statistical Mechanics*, 2007(09):P09020, 2007.
- [7] T. Guha, N. Kumar, S. S. Narayanan, and S. L. Smith. Computationally deconstructing movie narratives: an informatics approach. In *IEEE ICASSP*, pages 2264–2268, 2015.
- [8] P. Holme and J. Saramäki. Temporal networks. *Physics reports*, 519(3):97–125, 2012.
- [9] P. Mac Carron and R. Kenna. Universal properties of mythological networks. *EPL*, 99(2):28002, 2012.
- [10] F. Moretti. Network theory, plot analysis. *Stanford Literary Lab*, 2, 2011.
- [11] P. Mutton. Inferring and visualizing social networks on internet relay chat. In *ICCIV*, pages 35–43, 2004.
- [12] C.-Y. Weng, W.-T. Chu, and J.-L. Wu. Movie analysis based on roles social network. In *IEEE ICME*, pages 1403–1406, 2007.
- [13] C.-Y. Weng, W.-T. Chu, and J.-L. Wu. Rolenet: Movie analysis from the perspective of social networks. *IEEE Transactions on Multimedia*, 11(2):256–271, 2009.

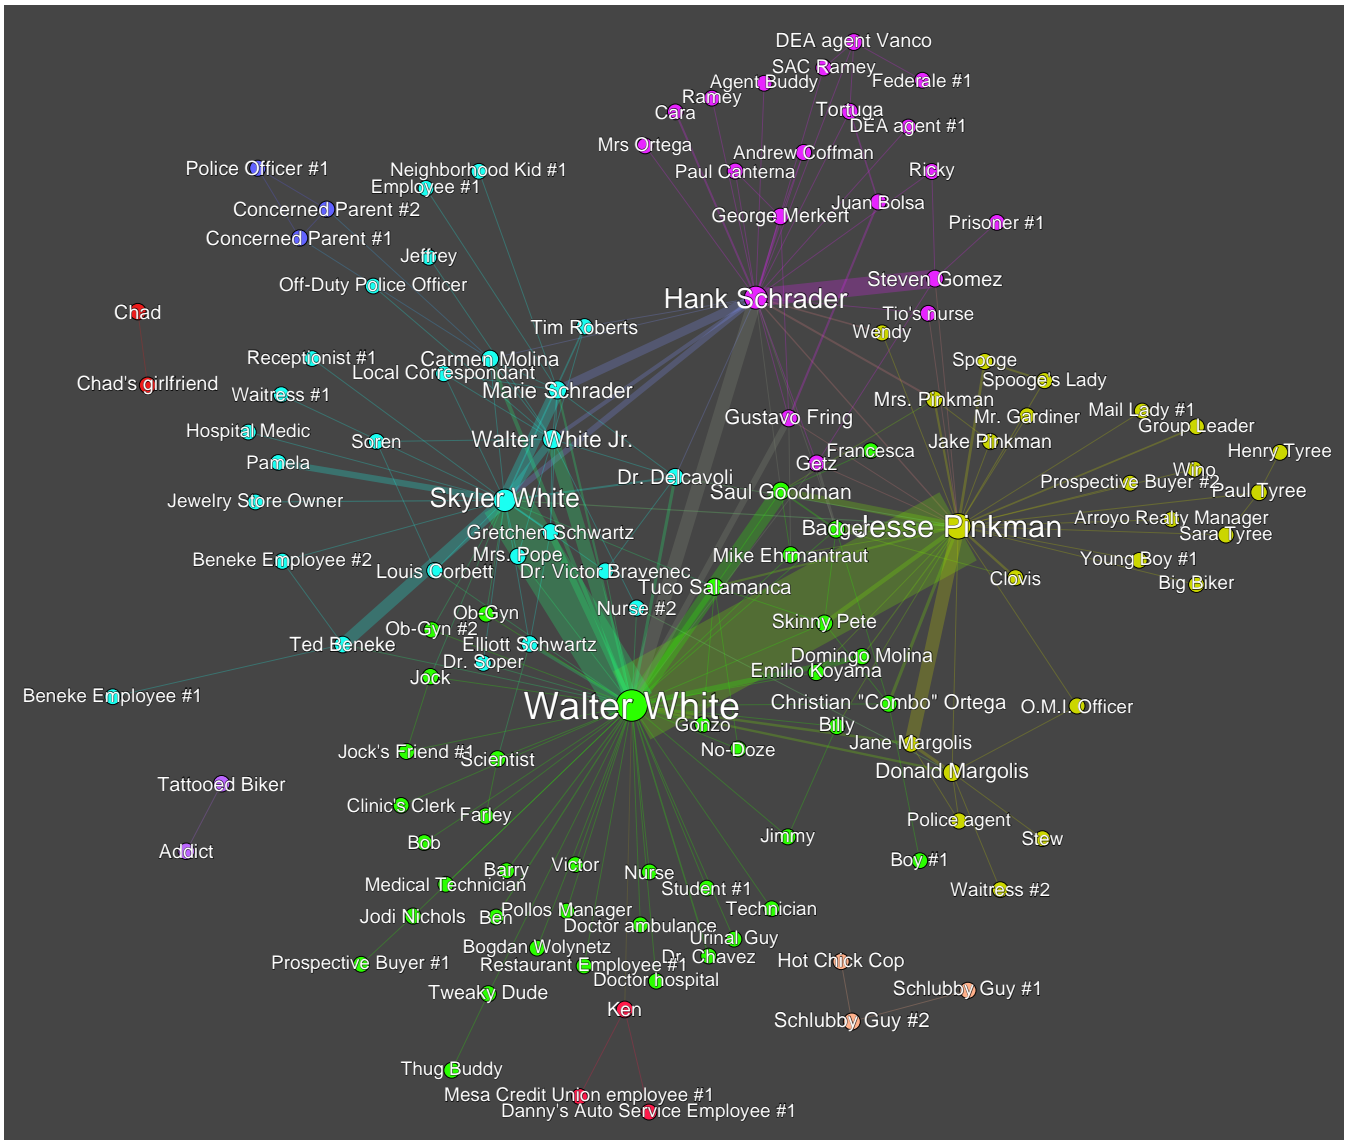


Figure 11: Cumulative network for the first two seasons of Breaking Bad.

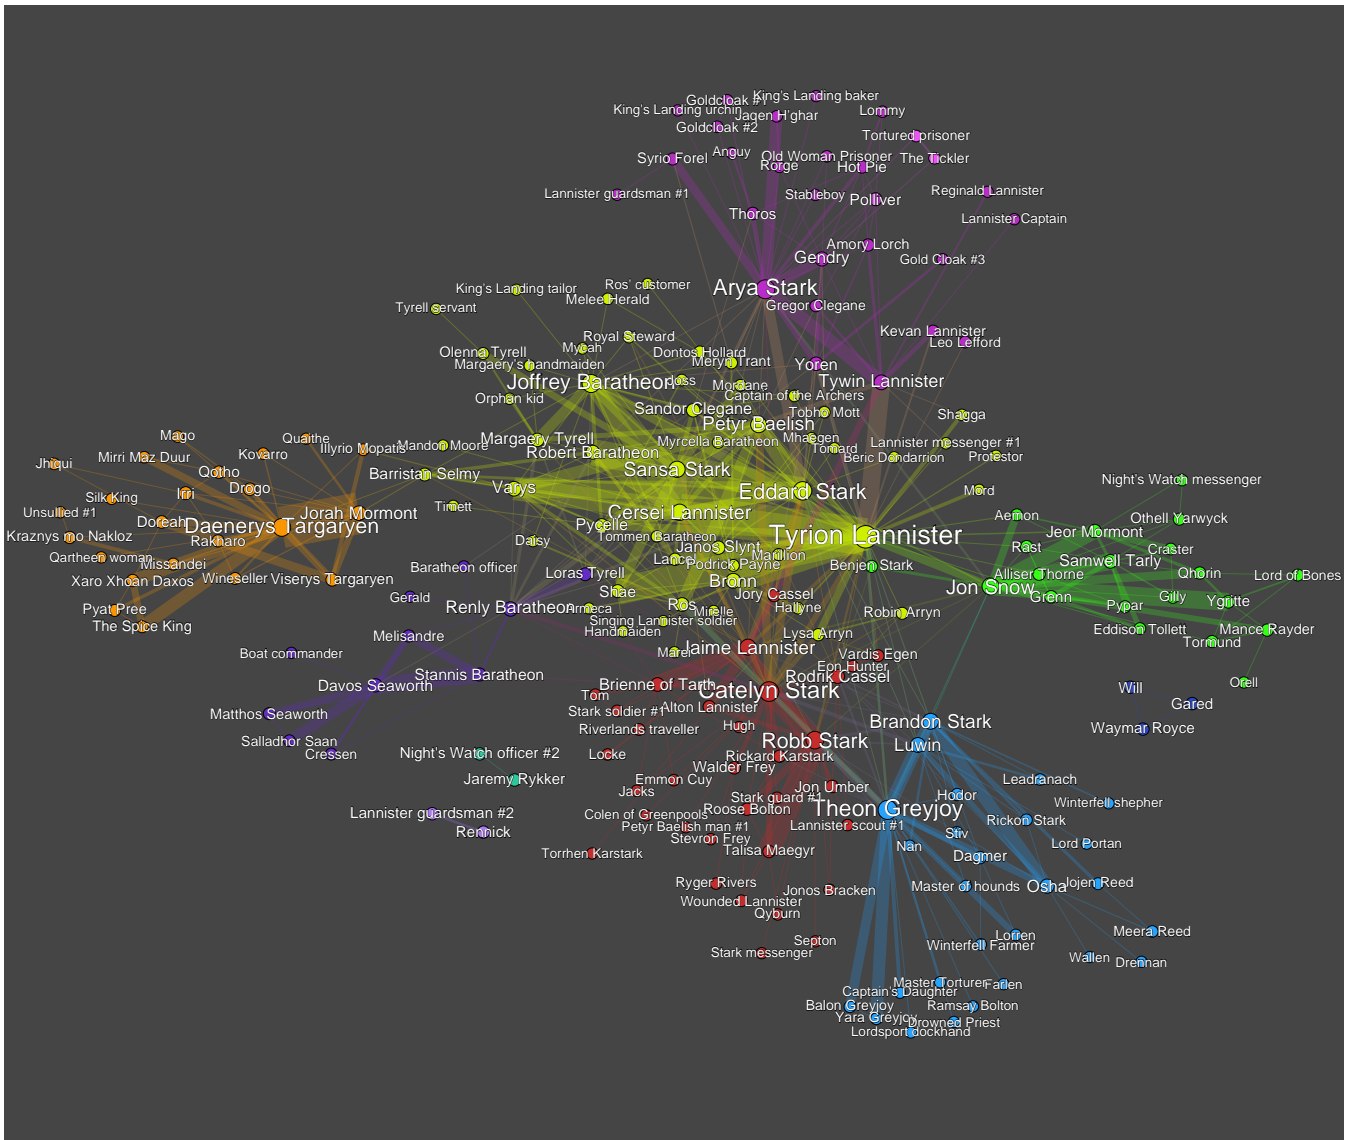


Figure 12: Cumulative network for the first two seasons of Game of Thrones.

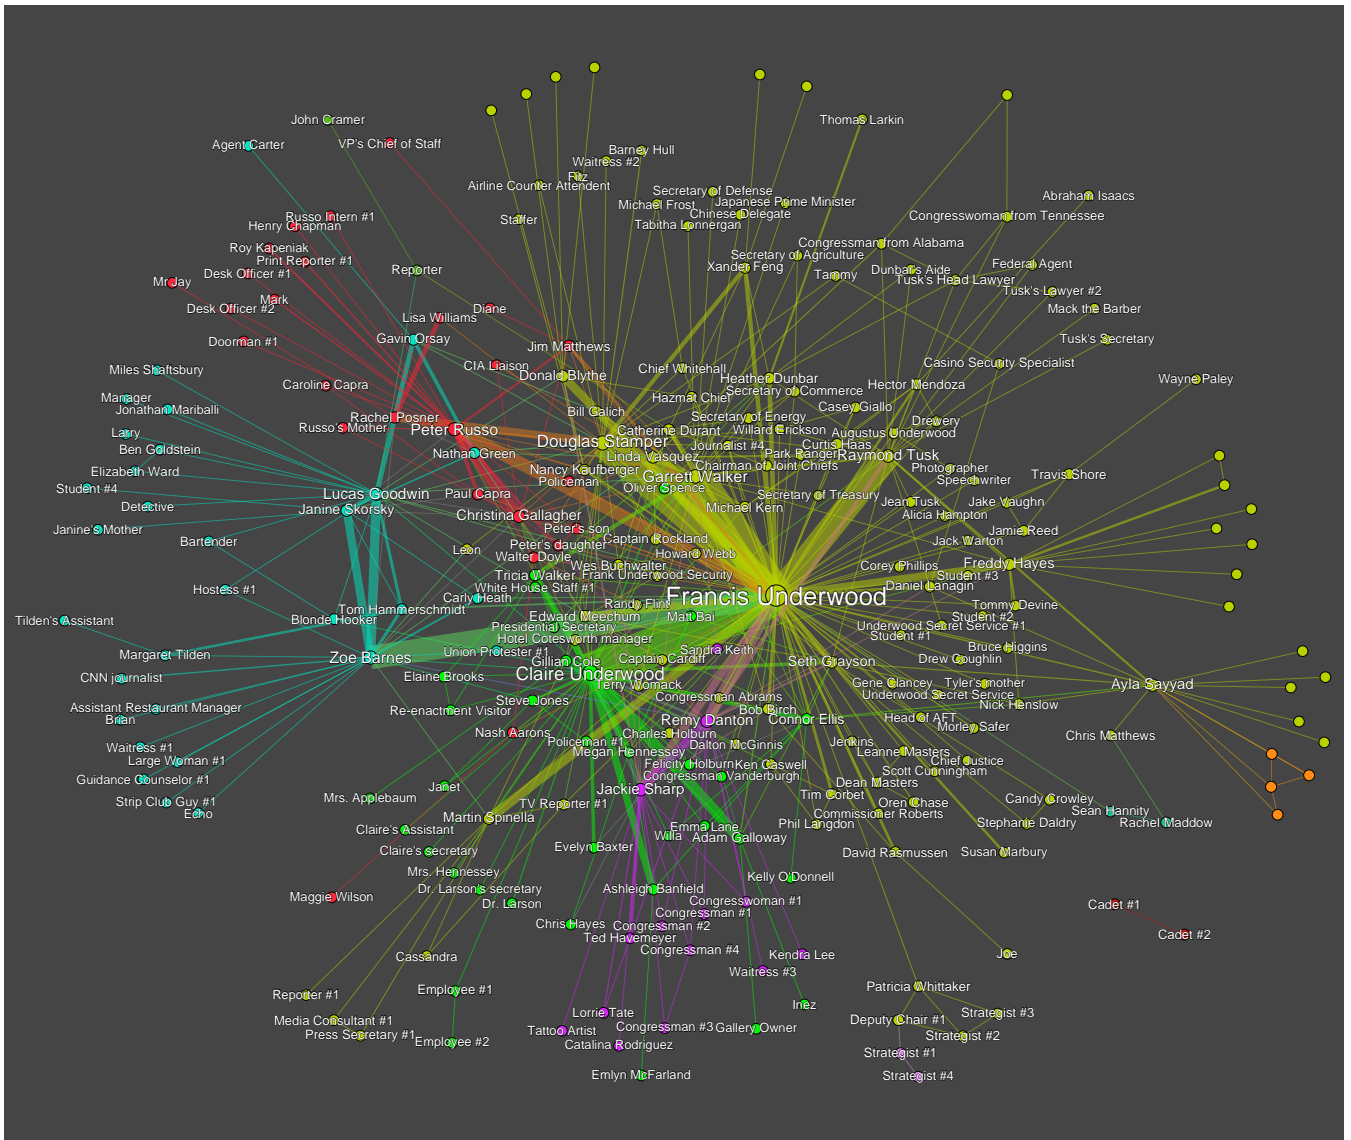


Figure 13: Cumulative network for the first two seasons of House of Cards.