

Automatic Identification of Character Types from Film Dialogs

Marcin Skowron, Martin Trapp, Sabine Payr, and Robert Trappl

Austrian Research Institute for Artificial Intelligence OFAI, Vienna, Austria

ABSTRACT

We study the detection of character types from fictional dialog texts such as screenplays. As approaches based on the analysis of utterances' linguistic properties are not sufficient to identify all fictional character types, we develop an integrative approach that complements linguistic analysis with interactive and communication characteristics, and show that it can improve the identification performance. The interactive characteristics of fictional characters are captured by the descriptive analysis of semantic graphs weighted by linguistic markers of expressivity and social role. For this approach, we introduce a new data set of action movie character types with their corresponding sequences of dialogs. The evaluation results demonstrate that the integrated approach outperforms baseline approaches on the presented data set. Comparative in-depth analysis of a single screenplay leads on to the discussion of possible limitations of this approach and to directions for future research.

Introduction

The research outlined in this paper aims at the modeling of conversational characters in virtual agents and social robots. As of today, the character of these new kinds of service automata is little developed. They are generally designed to give an impression of friendliness and helpfulness (footnote: companion videos) without expressing any personality of their own. In the research project *Social Engagement with Robots and Agents* (Payr 2013), the rabbit-like robotic frontend *Nabaztag* was used in a long-term field experiment. This simple robot was scripted in a way that the information, e.g., on potentially interesting events that it gave to subjects was formulated from a rabbit's viewpoint (e.g. "... and don't forget to bring me carrots from there!" or "... I as a rabbit would be afraid of the height, but maybe you'll like it!"). Subject liked these expressions of a rabbit persona. In this experiment, they had no means to explore this persona in any more depth, but it increased their attachment and the readiness to "play along" with the setup (Payr

CONTACT Robert Trappl  robert.trappl@ofai.at  Austrian Research Institute for Artificial Intelligence OFAI, Vienna, A-1010, Austria

Color versions of one or more of the figures in the article can be found online at www.tandfonline.com/UAAI.

© 2016 Marcin Skowron, Martin Trapp, Sabine Payr, Robert Trappl. Published with license by Taylor & Francis. This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives License (<http://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited, and is not altered, transformed, or built upon in any way.

2013). Robots that live and work with humans in varying situations and over long periods of time need both flexibility and consistency in their interaction in whatever role they take or task they carry out (Trapp et al. 2011). This requires that conversation “in character” is generated from a coherent character model. The intermediate goal of the research is therefore to study how character is revealed in dialog. Fictional works that present character through direct speech, especially dramas and screenplays, are therefore the primary source for our studies.

Fictional characters have been studied for a long time already, starting with Aristotle’s *Poetics* (Kuchenbuch 2005), from cultural, narrative, psychological, and literary perspectives. Even today, they are often based on cultural archetypes: the hero, the villain, the helper/teacher or angel, the skeptic, the sidekick, etc. (Truby 2007; Vogler 2007) enacting “core relational themes” (Lazarus 1991). From an author’s or screenwriter’s viewpoint, characters are built not so much as a collection of personality traits but as growing out of strengths, weaknesses, values, desires, as well as their social relationships (Truby 2007). The viewer or reader has to be able to infer this character, as well as its development (character arc) from appearance, actions, non-verbal expression, and talk.

The role of dialog in film

The first question for our line of research is whether film dialog (our primary source) actually provides enough clues to infer and distinguish characters from it. The answer is by no means clear-cut among film makers and film scholars. The role of dialog in film has undergone major changes throughout the history of cinema (Price 2013; Tieber 2008): with the advent of sound, dialog became dominant for a long time, so that screenplays resembled dramas, a tendency which, in turn, was heavily criticized by proponents of cinema as an independent art form. This conflict about dialog led many famous and influential filmmakers, such as e.g. Alfred Hitchcock (Kozloff 2000) to downplay dialog and to underline its insignificance. This neglect has been handed down to the present day, so that e.g. (Field 2005) in his influential manual on screenwriting can affirm that “... what a person does is who he is, not what he says.” (p. 69) “The logophobia of film studies is reactive: to avoid being a poor relation of theater, cinema needed to establish its own aesthetic credentials on very different terms... There is now a recognition that the logophobia was an over-reaction: with the status of film as an art form thoroughly secured, the discipline can admit that language-as-speech does contribute to the medium’s aesthetic effectiveness.” (Richardson 2010). As an additional factor that may have led to the denigration of talk in film, (Kozloff 2000) identifies the association of wordiness with femininity, and hence its incompatibility with the film industry’s

stereotypical male heroes. Without ignoring the role of image and sound, she puts film dialog back in its place as an important element in film, and distinguishes six primary functions, presented below:

Anchorage

Icons and symbolic images are used to signal time and place of a movie, but (additional) dialog is often necessary to create the story world (diegesis) for the audience, as images are basically polysemous. Anchoring dialog is frequent, often in the opening scenes: places are mentioned, and the main characters are presented in meetings, with greetings, often repeatedly. One can find side characters, especially in the opening scenes, whose only function is to talk to and/or about the main characters to introduce them to the audience. The power of dialog becomes evident in what are called “dialog hooks” (Bordwell 2007), e.g. a place is mentioned and the next shot shows a place that the viewer spontaneously identifies with the one mentioned before, although the actual location may be somewhere completely different, or just the studio.

Narrative causality

Dialog also plays a crucial role in establishing the logic of the film sequences and of the depicted events. Narratives unfold through a series of events, linked together by succession and causality, and the ulterior motive of much of film dialog is to communicate the Why and How of this succession, transforming a sequence into a believable narrative. Dialog interprets what is shown and orients the perception of the audience toward what they should see in a scene. “The dialog paves the way for us to understand the visuals, repeats their information for emphasis, interprets what is shown, and explains what cannot be communicated visually.” (Kozloff 2000) (p. 39) For example, the sense of urgency that is often introduced early in a film to create tension can only rarely be communicated visually.

Verbal events

The number of instances where an event in a film is actually verbal should not be underestimated. A secret is disclosed or a revelation is made in dialog. Declarations of love, verdicts, confessions, etc. are what (Searle 1969) called “declarative speech acts”: they are themselves actions insofar as the utterance of certain, mostly conventionally determined, words have the power to effect a change in the world.

Character revelation

A character is a construct from the very many different signs deployed in a film, among them the behavior in dialogs. It is not only the way of speaking, where verbal mannerisms, dialect, or accent give the audience clues about the

character and its background, but also what is said. Not only the character's own dialog lines, but also what other characters say to or about him/her, reveal personality, beliefs, and desires.

Code of realism

A proportion of dialog in every film serves primarily as a representation of ordinary conversational activities, e.g. greetings, ordering food, small talk. Their function is to underline a realistic everyday setting, sometimes to intentionally contrast it with extraordinary parallel or impending events.

Control of viewer evaluation and emotion

Dialog can be used in films to control pace (create urgency, see above, or slow down action), to draw attention, to guide the viewers' interpretation of what is seen, and to evoke emotional reactions (see also (Lee and Marsella 2011)).

(Kozloff 2000) adds to these aesthetic, persuasive, commercial, and ideological functions such as:

- exploitation of the resources of language (e.g. humor, irony, story-telling, poetics);
- opportunity for “star turns”: Star turns can be identified by their length, and they call for an extraordinary emotional expression and showcase performing skills (e.g. speed, tongue twisters);
- thematic messages/authorial commentary/allegory: dialog in a film's last scenes carries particular thematic burdens, either reinforcing/extending the film's moral or resisting closure.

Especially the last mentioned function of dialog, i.e. the overt message or commentary addressed to the audience, has been criticized heavily for breaking the fundamental principle of film dialog: characters speak as if they were talking to one another, while the ultimate addressee of each line is the audience. (Kozloff 2000) reflects this particularity in the title of her book, *Overhearing Film Dialogue*, and explicitly warns against mistaking film dialog as a source for studies on natural conversation, because this principle of constantly taking into account the viewer as an overhearer shapes film dialog in decisive ways. Film dialog has often been presented as the “poor relation” of rich natural conversation, tidied up, and thinned down in comparison to the complex processes of talk-in-interaction where meanings, identities, and relationships are constantly negotiated on the fly (Richardson 2010). On the other hand, the impoverishment of film dialog is counterbalanced by the specific functions that spontaneous talk has not. It has to be clear, therefore, that any study of film dialog, including the one presented here, are just that,

i.e. studies of non-natural conversations in works of fiction with their specific functions, instruments, conventions, and constraints.

In the toolbox of film studies, “logocentric” approaches, as (Richardson 2010) calls them, such as textual analyses of screenplays, have their legitimate place on the simple grounds that there is, at least as yet, no single method to study multimedial and multimodal productions comprehensively in their full range of media and modes simultaneously, so that each approach, if anchored in and controlled by the others, has its value and contributions to make.

Screenplay as text genre

What is more, film dialogs for text-based studies as the present one are mostly extracted from screenplays (some authors make and use transcripts from the screened film, but the effort required limits the amount of data that can be collected with this method). Screenplays are a very special text genre closely bound to the production and distribution of films (Price 2013), and the particularities arising from their place and role in filmmaking have to be taken into account. The specific context of production and use of screenplays also implies that it is a fluctuating, uncertain type of text that is hardly ever final and fixed.

“Regarded as works-in-progress, screenplays typically do not command the same respect during production processes as playscripts (especially playscripts penned by authors of reputation). Different kinds of scripts are generated during the course of these processes for the use of different participants; distinctions are made between screenplays, shooting scripts, lined scripts, and continuity scripts; actors will annotate their personal copies. Lines may be changed on the fly during rehearsal and filming; only certain copies of the “original” scripts will show those changes, at best. “The script” takes on a very uncertain, unstable character by comparison with the finalized production.” (Richardson 2010)

With a digital repository of screenplays like IMSDb,¹ the screenplays stored online are attributed a false finality: there is only one version of the screenplay for each film available, and its state and purpose are not clear. Shooting scripts (i.e. near-final versions with detailed instructions for camera position and cuts) co-exist on the Internet with early drafts where neither characters nor plot matches the film as distributed, as well as with a range of versions in-between.

The screenplay began its existence as short sketches for silent films that were limited in length by the constraint of the film reel (i.e., a maximum of 15 minutes). During the epoch of the studio system (Tieber 2008), the first

¹<http://www.imsdb.com>

version was still a short outline or storyline, a sketch of the film limited to a few pages. On approval, this storyline was extended by in-house screenwriters to a full screenplay, often by a team where each writer contributed his/her special experience and talents (e.g. for dialogs or jokes). With the end of this era, screenwriting became a more independent activity. Freelance screenwriters, as well as an army of writers who aspired to become ones, submitted complete screenplays to the studios, to be critically read and (maybe) selected for production. This was the time when screenwriting manuals began to flourish. Would-be screenwriters sought advice on how to build stories that would find grace before the rigorous eyes of producers and on how to ensure that their text would be preferred to those of their competitors (who, of course, had read the same manuals). Ironically, the uniformity of screenplays resulting from the adherence to the manuals was criticized by their very authors (e.g. (Field 2005)). For example, the three-act-structure became the norm, with proportions of roughly 1:2:1. Together with the rule-of-thumb that one page of screenplay approximately equals 1 minute of film, one can almost be sure to find the inciting incident (the first plot point) that leads from the expository first act to the dramatic complication and development of the second, around page 30 of any screenplay for a 2-hour film. Not even alternative suggestions stray very far from this principle: a four-act-structure (Bordwell 2007) only cuts the long second act into two distinctive parts, and a structure with eight sequences² can equally be accommodated by the same structure (2:4:2).

The available screenplays therefore do not vary widely with regard neither to their structure nor to their format. The conventions of the text genre are being respected widely and are fast becoming quasi-standards through the increasing use of special screenwriting software (such as *Final Draft*), which supports the production of just these formats. For instance, character names, stage directions, and dialog lines are centered in a narrow column, while slug-lines (beginning of a new scene with indications of location and time of day (Turetsky and Dimitrova 2004), descriptive text, and camera positions (where present) are left-aligned. The format, then, does not help to identify the type and stage of screenplay. Even the existence of shooting (e.g. camera position and angle) and editing (e.g. “dissolve to”) directions does not allow to conclude that the screenplay is close to the final filmed version, because experienced screenwriters think their screenplay in filmic categories and may already envision the scene as shot even in their first draft. Nor can one be sure to have an early version in hands if these instructions are missing: where a director participates in the writing, such instructions may be only in his head or in some handwritten notes.

²<http://thescriptlab.com/screenwriting/structure/the-sequence/45-the-eight-sequences>

Only the comparison between the film as released and the available screenplay may give some hints to the stage in the production process from which the screenplay originates. Differences in names, locations, and in whole plot lines between the two are evidence of the repeated revisions and modifications (by different people) the original screenplay can undergo before the film is actually shot. In sum, the digital text genre of screenplay only appears to be a uniform body of texts, while, in reality, it is a family of more or less closely related types of texts at different stages of development and with different uses and readers.

Notwithstanding these observations, the form of most screenplays is similar. Beside slug-lines, shooting and editing directions (in capital letters), most scenes contain both descriptive text (about location, atmosphere, characters, and actions), and dialog lines with leading character names (again in capital letters), often accompanied by stage directions, e.g. “O.S.” = off screen, or “whispers.” These formal elements make it possible to extract elements from screenplays automatically, e.g. scene boundaries or dialog lines, and to create corpora out of digital screenplay texts.

Following an overview of Related Work, the paper presents the data and annotations used before discussing the methods. The Section “Methods” presents and discusses the methods, both computational and manual, applied to the detection of character types from these data: computations on linguistic annotations and semantic graphs over a corpus of action movie dialogs, as well as time-dependent and comparative analysis of characters in a single screenplay.

Related work

The understanding of fictional character is rich, but systematic descriptions of the linguistic and conversational devices that contribute to construe character are scarce (Quaglio 2009). Discourse studies have started to approach fictional dialogs as conversations rather than as the work of a single author, but until recently, discourse analyses of film and drama dialog were done manually and limited to single works or authors (Culpepper, Short, and Verdonk 1998). The open access to large numbers of digital screenplays is a relatively recent phenomenon, and took its time to be recognized as a source for research. In 2010, Richardson (Richardson 2010) could still complain about the scarce availability of screenplays for research, although the earliest mention of IMSDb that we found dates back to 2006,³ when the database seemed to have been still in its infancy so that its success in reaching a critical mass of data was uncertain. We may therefore be in a

³<http://www.cineblog.it/post/1398/the-internet-movie-script-database-imsdb>

phase where the use of such resources is still in its infancy, especially with regard to the purpose for and the methods with which they are exploited.

In recent years, several methods have been developed and evaluated on non-fictional texts, with a focus on online content. Some of these methods were also applied to the analysis of fictional dialogs (Danescu-Niculescu-Mizil and Lee 2011; Gorinski and Lapata 2015; Karsdorp et al. 2015; Kundu, Das, and Bandyopadhyay 2013; Shen 2011). Although limited to linguistic style analysis, (Walker et al. 2011) showed that the detection of some of the communicative patterns and linguistic features of a character can be achieved (Walker et al. 2011). Specifically, the published approach relied on a set of linguistic features, aggregated and extracted jointly from all dialog lines of a given character, without differentiating conversational partner, point in the dramatic arc, or narrative relevance. Table 1 gives an overview of studies on digital dialogic sources with relevance to character modeling.

Mathematical social network theory has been applied to fictional works, e.g. by (Mac Carron and Kenna 2012) in their study of mythological narratives, or by (Alberich, Miro-Julia, and Rosselló 2002; Gleiser 2007) on the universe of Marvel Comics characters. (Gleiser 2007) was able to show the difference of network structures around “hero” and “antagonist” characters. (Suen, Kuenzel, and Gil 2013) extracted character interaction networks as weighted graphs to classify their data (both dramas and screenplays), e.g. for: drama/film distinction and classification, period (of dramas), genre, and authors. In (Agarwal et al. 2015), parsing and social network analysis methods with lexical and semantic methods are used to assess the presence of female characters in screenplays (“Bechdel test”). Social network analysis and

Table 1. Overview of digital film studies with relevance to character modeling.

Study	Primary source	Secondary source	Goal	Focus
(Bednarek 2010)	TV series transcripts: dialog lines	??	Stylistics	Talk
(Walker et al. 2011)	Screenplays (IMSdb): dialog lines	IMDB	Generation, stylistics	Talk
(Bamman, O'Connor and Smith 2014)	–	Plot summaries (Wikipedia)	Latent personae (char. categories)	Action
(Agarwal et al. 2014)	Screenplays (IMSdb): scene	??	Narrative structure	Relationship, plot
(Suen, Kuenzel, and Gil 2013)	Drama, screenplays: non-dialogic text	–	Genre, author, period	Relationship
(Weng, Chu, and Wu 2009)	Film	–	Character clustering	Relationship
(Gleiser 2007)	Film	–	Character classification	Relationship (network)
(Agarwal et al. 2015)	Screenplays	Metadata: cast (IMDB)	Gender studies	Relationship
Danescu-Niculescu-Mizil and Lee 2011)	Screenplays	Metadata: Cast (IMDB)	Linguistic: Priming	Talk

related methods have been applied to a variety of research questions on fictional texts, but not to the elicitation of *character* considered as both a position in a social network and as agent of relationships.

(Traum 2003) discusses the representation and modeling of naturally occurring multiparty dialogs and participants' roles: conversational role, speaker identification, addressee recognition, and other participant roles, and considers issues in dialog management related to the transition from two-party to multiple participants. (Lee and Marsella 2011) provide an analysis of nonverbal behaviors in interactive drama associated with side participants and bystanders in multiparty interactions and construct an analysis framework incorporating characters' interpersonal relationships, conversational roles, and communicative acts. Their set of conversational roles includes: speaker, addressee, side participant, and bystander. The relationships between characters are described in terms of dominance and friendliness, and the communicative acts extend beyond the dialog acts of character utterances to include events that may evoke emotional responses from the characters.

In (Bamman, O'Connor, and Smith 2014), authors take as data the plot summaries of movies, plus metadata about films (genre) and actors (gender, age) from wikipedia.org. Using an extension of the Latent Dirichlet Allocation (LDA), (Bamman, O'Connor, and Smith 2014) obtain shared multinomial distributions estimated on actions of characters, actions done to characters, and attributes. The assignment of a character's actions, received actions, and attributes are used to describe the latent persona of a character. In (Bamman, Underwood, and Smith 2014), narratives extracted from a large collection of novels are used for the modeling of characters. In both cases, character is elicited from narrative or descriptive texts from a third-person view, while the approach presented here aims at learning character from the character's own (verbal) actions, i.e. from a first-person viewpoint.

To answer the research question of the present paper, namely whether character can be detected from verbal behavior, we focused on dialogs extracted from screenplays as a data source.

Data and annotation

We used an existing corpus of movie scripts (Walker, Lin, and Sawyer 2012), from which we automatically extracted the following information from 777 screenplays: scene number; utterance number; character name; and character's utterance. In order to find the *scene boundaries* within a script, we re-aligned the extracted utterances with the scene descriptions provided by (Walker, Lin, and Sawyer 2012). A randomly sampled set of 1000 utterances was manually checked to provide a measurement for the accuracy of the applied preprocessing steps. In the annotated set, we found 2.5% of cases

where descriptive text was falsely attributed to characters' utterances. The reverse error, i.e. dialog lines falsely labeled as descriptive text, could only be checked informally by comparing randomly sampled scenes with original screenplays. This comparison lets us assume that the error rate is roughly the same as the other way round. The average number of scenes per screenplay in our corpus is 214, where a single scene of a script contains up to eleven, and typically two characters. In the corpus, a character has on average 6.35 utterances per scene, with a maximum of 350 utterances. In total, the corpus contains 136,432 utterances. The whole corpus of 777 extracts was used for linguistic annotation (see 3.2), while a subset of 212 screenplays was used for character type annotation.

Annotation of character types

In the presented work, we focus on action movies, as this genre most reliably adheres to traditional narrative principles (Vogler 2007) translating into casts of recognizable character types. The action movies were selected using the genre labels in the corpus (Walker, Lin, and Sawyer 2012). From the relevant set of 212 screenplays, 966 characters with at least 10 utterances in the scripts were extracted. By comparing them to the cast list in IMDB,⁴ the lineup of characters was extended to 2010 characters and aligned with additional data, including actor's name, age, and gender. The final list of characters was manually annotated by scanning information about the movie. Summaries and synopses in IMDB were used as primary source. Where insufficient, user-generated reviews on IMDB and Rotten Tomatoes⁵ were added for more information. Additionally, popular rankings of movie characters, e.g. "50 Best Movie Mentors,"⁶ were consulted for role attributions. The following character types were labeled: hero (H); antagonist (A); spouse/partner/lover (of either) (H-L) (A-L); sidekick (of either) (H-SK) (A-SK); supporting role characters for any of these (H-SR) (A-SR); mentor (M); the power in the background (e.g. businessman, government representative) (BUS); and representatives of the law (e.g. police detective, judge) (LAW). Some groups are broader than others, such as H-SR and A-SR, which contain a wide spectrum of supporting roles—e.g. parents, teachers, secretaries, taxi drivers—in contrast to character types like antagonist (villain) or hero, which are more homogeneous. Some of the scripts include multiple heroes, e.g. Marvel's *Avengers* where there is typically a group of (super) heroes, so that the final number of heroes exceeds the number of movies.

⁴<http://www.imdb.com>

⁵<http://www.rottentomatoes.com>

⁶<http://www.shortlist.com/shortlists/best-movie-mentors>

Table 2. Summary of basic statistics for character types in the corpus. NA denotes characters that were not annotated, but who had at least 10 utterances.

Linguistic Marker	No. Classes	F1	F1	F1
		Avg. weighted	Majority class	Minority class
Dialog acts	5	0.91 \pm 0.01	0.94 \pm 0.01	0.92 \pm 0.07
Sentiment	3	0.95 \pm 0.00	0.98 \pm 0.00	0.36 \pm 0.22
Address	2	0.90 \pm 0.02	0.95 \pm 0.01	0.63 \pm 0.07
Conditional	2	0.91 \pm 0.02	0.97 \pm 0.01	0.15 \pm 0.15
Discontinuity	2	0.94 \pm 0.02	0.97 \pm 0.01	0.81 \pm 0.08
Elliptic	2	0.84 \pm 0.02	0.90 \pm 0.01	0.64 \pm 0.04
Hedge and filler	2	0.88 \pm 0.01	0.94 \pm 0.01	0.67 \pm 0.02
Interjection	2	0.92 \pm 0.02	0.97 \pm 0.01	0.47 \pm 0.13
Location	2	0.95 \pm 0.00	0.98 \pm 0.00	0.00 \pm 0.00
NSE lexicon	2	0.89 \pm 0.01	0.94 \pm 0.01	0.65 \pm 0.03
NSE syntax	2	0.91 \pm 0.01	0.97 \pm 0.00	0.15 \pm 0.14
Organization	2	0.98 \pm 0.00	0.99 \pm 0.00	0.00 \pm 0.00
Person	2	0.88 \pm 0.03	0.93 \pm 0.02	0.68 \pm 0.07
Polarity	2	0.96 \pm 0.01	0.98 \pm 0.00	0.82 \pm 0.05
Politeness	2	0.96 \pm 0.02	0.98 \pm 0.01	0.50 \pm 0.21

Table 2 summarizes basic statistics about the annotated character types: total count of character instances and scripts they appear in; percentages of scenes a character appears in; percentage of all characters from a script a character interacts with; percentage of all utterances from a script a character utters or is an addressee of, i.e.: utterances of other characters that directly precede a given character turn.

Annotation of linguistic markers

A data set consisting of 1000 utterances was extracted randomly from the complete corpus of 777 screenplays. The set was annotated for named entities, forms of address, non-standard English usage, discontinuities, ellipsis, discourse markers, interjections, sentiment, polarity, and dialog act type. The set was annotated by a single annotator with linguistic background. The use of a single annotator required the definition of explicit rules and word lists so that the annotations were comprehensible and reproducible. The annotated features roughly fall into two groups, related to the building blocks of fictional characters (Truby 2007): one is about the presentation of a character's personality through expressivity, the other about the rendering of their social relationships. Features of *expressivity* relate to the character-revealing function of film dialog (Field 2005; Kozloff 2000), and include:

- *Sentiment*: explicit expressions of sentiments with positive or negative valence; and a neutral sentiment. Examples: “Oh, my God!,” “Damn you—why won’t you stop ?!” for negative valence, “I love this place!,” “Great!” for positive valence.

- *Interjections*: primary (“Oh” etc.) and secondary (“Damn” etc.) as well as combinations thereof (“oh, God!”), considered as markers of intensity of sentiment expression (Quirk et al. 1985).
- *Ellipsis*: a potential marker of terse, contracted talk that film critics associate with the prototypical male hero of action movies (Peberdy 2013). Example: “Funny in the head. High suicide rate in the medical profession.”

Social-relational features are mostly signals of (static) status or power (Kemper 1990) and of (situational) dominance or initiative (Payr 2007). We annotated our data set for the following features, which we assume to be relevant for expressing social relationships:

- *Dialog act types*: coarse-grained annotation of declaratives/statements, wh- and yn-questions, yn-answers, imperatives, with the rest (e.g. interjections) attributed to a class “other.” Examples: “Who’s he?” (wh-question), “Are you enjoying yourself?” (yn-question), “Don’t go yet, Mama.” (command), “Oh my” (other), “Okay” (yn-answer), “This is wrong, Jode.” (statement).
- *Address*: (Berliner 2013) notes that addressing the dialog partner(s) by name, pet names, or honorifics is far more frequent in screenplays than in naturally occurring conversations. Examples: “dear,” “sir,” [name] (as identified by Named Entity Recognition).
- *Discontinuities*: (Prabhakaran and Rambow 2013) state that written texts do not contain interruptions; dialogs in screenplays are different, in that discontinuities such as pauses, fade-out, or barge-in are represented. As the writer’s intention is not to mimic natural conversation (Kozloff 2000), such discontinuities are considered markers of uncertainty, shyness, and/or lower social status of the speaker. Examples: cut-off sentence + “...” anywhere in the turn, or “—” at end of turn, e.g. “Stop thinking that I will be the”
- *Discourse markers*, e.g. hedges and fillers, are frequent in spoken, but rare in written language. Like discontinuities, they can be considered deliberately placed markers in screenplay dialog, in particular where they are contrasted with the conventional efficiency and flawlessness of e.g. the action hero’s lines (Kozloff 2000). Examples: “I mean, I’d be tempted to eat them myself. So I guess, just, if you could get it over with quick.” In this dialog turn, we annotated: “I mean,” “so,” “I guess,” “just.”
- *Non-standard English usage (NSE)*: while the borders between standard and non-standard English expressions and grammar are fluid in natural conversation, their use in written dialogs is here assumed to be intended by the author (Taavitsainen 1999) to reveal social status, education, or origin of a character. NSE lexicon examples: “gotcha,” “wanna,” “motherfucker.” NSE syntax examples: “I been walking around” (zero copula), “Might end up cutting your throat.” (dropped subject pronoun).

- *Preterite modals* (“could,” “would,” “should,” “might”) are used in 7% of the sample utterances and at present are annotated without consideration of use context, e.g. politeness or hedging (Shen 2011).
- *Conventional politeness* (Watts 2003): both greetings and conventional polite expressions (e.g. “sorry,” “please,” “thank you”) were annotated, but turned out to be infrequent in the sample.
- *Entities*: annotation of references to person (numerous: see Address, above), location and organization (infrequent).

Methods

Linguistic markers

In order to allow for semantic reasoning at the sentence level, we explored use of the skip-thought model (Kiros et al. 2015), which learns unsupervised representations of text. The skip-thought model is motivated by the distributional hypothesis that consecutive sentences that share a contextually similar surrounding are likely to be semantically similar. We use a model trained on the corpus of books (Zhu et al. 2015) to obtain fixed-length vector representations. Based on those high-dimensional representations, we trained classifiers for the set of linguistic markers presented above. The classifiers are trained independently of character types. All the results are computed using Support Vector Machines (Vapnik 2013) with linear kernel, optimized using exhaustive grid search. The F1 score, averaged over all 15 types of linguistic markers, stratified split, fivefold cross-validation, is 0.91, with baseline scores obtained with different representations: 0.87 (unigram with tf-idf scores), and 0.85 (unigrams and bigrams with tf-idf scores).

Table 3 presents the classifiers’ performance—average weighted F1 scores obtained with the selected approach: linear kernel Support Vector Machines (SVM) and skip-thought vectors representations. The presented results were computed using a stratified split in fivefold cross-validation. Further, the table also includes information on the number of classes for each type of the

Table 3. Performance of the linguistic marker classifiers: number of classes; F1 scores along with the standard deviation for the average weighted, and majority and minority class.

Role	H	A	M	H-L	A-L	H-SK	A-SK	LAW	BUS	H-SR	A-SR	NA
# Instances	228	161	31	88	6	117	11	38	38	204	44	6438
# Scripts	182	130	27	83	5	85	11	33	34	105	30	204
% Scenes	53.8	15.2	3.3	13.8	0.4	17.1	1.3	2.9	2.3	16.4	2.8	69.7
% Ch. Inter.	34.7	18.5	16.8	19.0	10.6	21.0	16.6	14.1	15.0	14.9	12.1	3.1
% Utt. (S)	27.4	6.2	1.3	6.1	0.2	7.0	0.4	1.3	1.1	7.3	1.2	40.4
% Utt. (A)	27.2	6.4	1.4	5.9	0.2	6.8	0.4	1.3	1.1	7.1	1.2	33.9

For the linguistic markers with more than two classes, namely sentiment and dialog acts, these are, respectively: neutral and positive classes; statement and imperative classes.

proposed linguistic markers as well as the corresponding F1 scores for the majority and minority classes and standard deviation.

The ability to accurately classify the selected linguistic markers varies between different types of markers. While the weighted average F1 scores indicate a reasonable performance in all the classification tasks, a closer inspection of the F1 scores for the minority classes reveals deficiencies in some of the classifiers. This is particularly the case for the *location*, *organization*, *conditional*, and *NSE syntax* linguistic markers, for which only a relatively small number of instances in a minority class exists. However, the other linguistic markers classifiers demonstrate both high weighted average F1 scores and the ability to predict the instances of minority class. This validates their application in the analysis of the fictional dialogs and in the remainder of experiments presented below.

Semantic dialog graph features

Investigations of communication characteristics and social roles of actors in social networks through descriptive analysis have gained increasing consideration (Myers et al. 2014; Otte and Rousseau 2002). Descriptive analysis of networks is done through computation of measures of the graph structure (Kolaczyk and Cs'Ardi 2014) and can help to understand and characterize the flow of information and relationship between actors in such networks.

We use descriptive analysis of graph structures to represent the interactive and communication characteristics of character roles in screenplays. To this end, we extract weighted directed graphs $G = (\mathcal{V}, \mathcal{E}, \mathcal{W})$, with N vertices $v \in \mathcal{V}$ and M edges $e \in \mathcal{E}$, each associated with a weight $w_e \in \mathcal{W}$. In the following, we denote v, x, y, s and t to be nodes in a graph where s refers to the source node of a path in the network and t to the target or sink node.

Each vertex in the graph can be understood as a character in a movie screenplay and each edge as the characterization of the interaction between two characters in the screenplay. To allow rich characterizations, we construct several weighted graphs, each reflecting the strength a specific linguistic marker carries in the interactions. We call such graphs *semantic dialog graphs*. Application of these graphs allows us to capture social aspects and relationships of character roles inside a screenplay. In particular, we characterize a character v in terms of the following measures computed on each graph:

- *Closeness centrality*: How easy can v reach out to other characters.
- *Betweenness centrality*: How easy can two characters (s and t) reach out to each other by communicating via v .
- *Clustering coefficient*: How likely is it that two characters (x and y) who interact with v , also interact with each other.

- *Squared clustering coefficient*: How likely is it that two characters (x and y) who interact with v share a common interaction partner who is different from v .

Descriptive analysis of semantic dialog graphs

The measures used in this work are formally defined as follows. The *closeness centrality* (Beauchamp 1965) for a character in a screenplay $v \in G$ is defined as

$$c_v = \frac{N - 1}{\sum_{x \in V \setminus v} d(v, x)} \quad (1)$$

where $d(v, x)$ denotes the shortest path distance from v to x . In the case of semantic dialog graphs, the shortest path distance can be obtained using the algorithm by (Dijkstra 1959) for weighted directed graphs. Note that we use the inverse of the weights in all computations of the shortest path. The measure of *betweenness centrality* (Freeman 1977) gives information on how well a v connects other characters and is defined as

$$b_v = \sum_{s, t \in V \setminus v} \frac{\sigma(s, t)}{\sigma_v(s, t)} \quad (2)$$

where $\sigma(s, t)$ denotes the number of *shortest paths* from s to t and $\sigma_v(s, t)$ denotes the number of shortest paths from s to t passing through v . Since the measure of betweenness favors vertices that join dense subgraphs and expresses the influence of v on the transport of information throughout the graph, this measure is a possible indicator for central roles, e.g. *Heroes* in action movies. We found that in 77% of the cases, characters with the highest degree of betweenness in a movie are *Heroes*, while the second most frequent group, with 8.7%, consists of characters without annotation. Note that this group contains potential *Hero* type characters that have not been annotated.

In addition to measures of centrality, we analyze the topology of each graph by computing the *clustering coefficient* and the *squared clustering coefficient*. Previous work (Watts and Strogatz 1998) suggests that grouping tendencies of social networks can be expressed by the clustering coefficient, which is computed for weighted graphs (Saramäki et al. 2007), as follows:

$$l_v = \frac{1}{k_v(k_v - 1)} \sum_{x, y} (\tilde{w}_{x,v} \tilde{w}_{y,v} \tilde{w}_{x,y})^{1/3} \quad (3)$$

where $\tilde{w}_{x,y} = \frac{w_{x,y}}{\max_{w \in \mathcal{W}} w}$ is the normalized weight of the edge connecting x with y and k_v denotes the number of neighbors of v . As the *clustering coefficient* assumes that the network contains cycles of size three, however, depending

on the screenplay and the linguistic marker, this assumption might not be met. We therefore additionally compute the *squared clustering coefficient* by (Lind, González, and Herrmann 2005), which is still well defined in such cases. To capture the listening and speaking behaviors of a character, we compute the *clustering coefficient* and the *squared clustering coefficient* for the listening and the speaking behaviors independently. We construct independent undirected semantic dialog graphs, which are weighted with the normalized frequency at which a linguistic marker is sent or received.

In addition, we further capture simple measures of listening and speaking behavior in terms of the *in-* and *out-degree* in a semantic dialog graph. We also extended the set of graph-based features with a binary feature indicating if character v is the most prominent character in the semantic dialog graph. In this analysis, we consider the character with the highest betweenness to be the most prominent.

Examples of semantic dialog graphs

We analyzed to which extent different characters are captured using *closeness centrality*, *betweenness centrality*, and *clustering coefficient* in semantic dialog graphs. We found that *closeness centrality* is slightly linearly correlated with main characters of the plot as well as with supporting roles. As illustrated in Figure 1, in the case of the movie *The Matrix*, the character *Morpheus*, which acts as mentor and supporting character of the hero *Neo*, obtains the largest *closeness centrality* score in the semantic dialog graph weighted by the

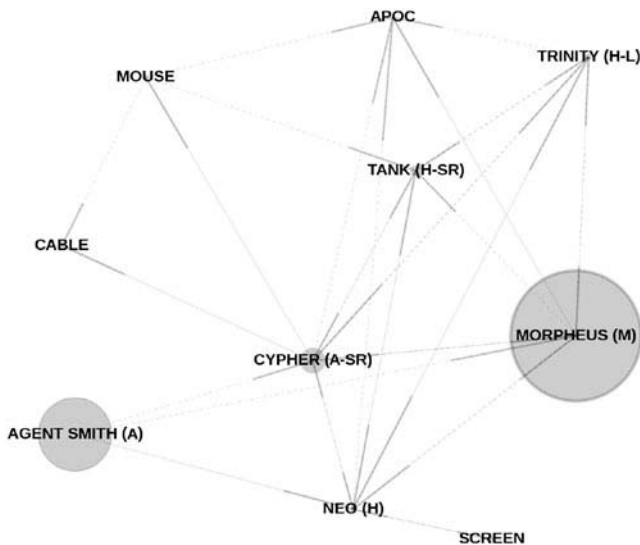


Figure 1. Illustration of closeness centrality for the screenplay of *The Matrix* in the semantic dialog graph weighted by the address marker, i.e., larger circles indicate higher closeness centrality values. The plot indicates that the mentor character *Morpheus* can address characters better than other characters in the plot.

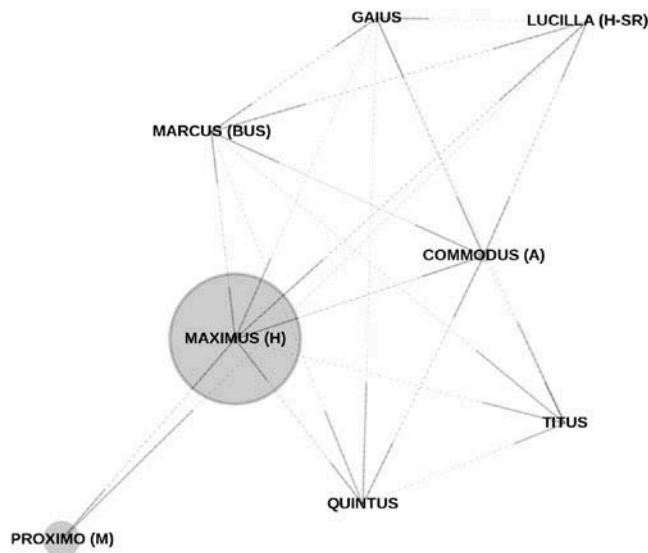


Figure 2. Illustration of betweenness centrality for the screenplay of *Gladiator* in the semantic dialog graph weighted by the address marker. Betweenness centrality in an address marker weighted graph is a measure of control over addressing messages passed between characters. For the screenplay of *Gladiator*, the character with the strongest betweenness centrality is the hero *Maximus*, indicating that this character has a central role.

address marker. Figure 2 illustrates the effect of *betweenness centrality* for the screenplay of *Gladiator* and Figure 3 the effect of the *clustering coefficient* for the screenplay of *Mission Impossible* in the undirected semantic dialog graph weighted by the interjection marker for the listening behavior.

Evaluation of an integrative approach

We conducted experiments to compare the different representations and measured their capacities to identify character types from sequences of multiparty dialogs and character interactions. The presented results are based on the data set of action movie scripts presented in Section 3.1.

Table 4 illustrates the contributions of different feature types: *N-Grams*—literal content of utterances; *ST vect.*—high-dimensional representation of utterances based on skip-thought model (Kiros et al. 2015); *LIWC*—representation of utterances based on Linguistic Inquiry and Word Count dictionary (Pennebaker, Francis, and Booth 2001); *Ling. Markers*—linguistic markers of expressivity and social role introduced in Section 3.2; *Graph*—features inferred from analysis of semantic dialog graphs described in Section 4.2. For the baseline approaches, we present only the results for the best performing type of representation, weighting scheme, and scaling. Specifically, for *N-Grams* the best results were obtained using the Term Frequency–Inverse Document Frequency (TFIDF) weighted unigrams. For

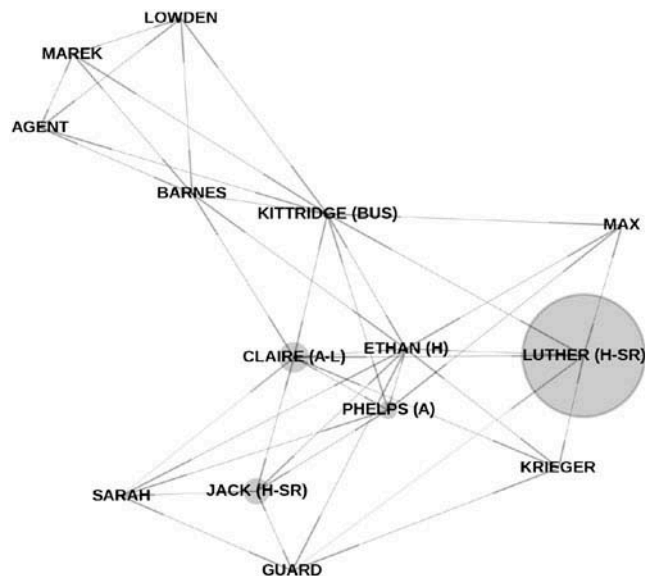


Figure 3. Illustration of the clustering coefficient for the screenplay of *Mission Impossible* in the undirected semantic dialog graph weighted by the normalized frequency of received interjection markers. As shown, the clustering coefficient can be understood as feature for side roles.

Table 4. Macro-average F1 scores for character type prediction. Best results for each character type are typed with bold font.

Role	<i>N</i> -grams	<i>ST vect.</i>	LIWC	Ling. markers	Graph	Graph + <i>N</i> -grams
H	0.61	0.64	0.38	0.49	0.73	0.77
A	0.44	0.29	0.00	0.29	0.39	0.42
M	0.00	0.00	0.00	0.00	0.00	0.00
H-L	0.21	0.08	0.00	0.11	0.20	0.24
A-L	0.00	0.00	0.00	0.00	0.00	0.00
H-SK	0.19	0.17	0.00	0.00	0.25	0.26
A-SK	0.00	0.00	0.00	0.00	0.00	0.00
LAW	0.20	0.36	0.00	0.00	0.00	0.20
BUS	0.00	0.00	0.00	0.00	0.00	0.00
H-SR	0.41	0.48	0.30	0.42	0.43	0.45
A-SR	0.00	0.00	0.00	0.14	0.00	0.36
Weighted avg.	0.36	0.37	0.09	0.24	0.37	0.43

ST vect. for all character’s vectors, the average, minimum, and maximum for each of the 4800 dimensions were calculated. Based on the experimental results (best F1 average score) obtained for each of the scaling, the representation using the minimum for each dimension was selected.

The averaged F1 scores depicted in Table 4 were obtained using stratified, fivefold cross-validation computed over 5 independent runs. We used the same random seeds for all experiments. The standard deviation of the average F1 score for all the feature types was ≈ 0 ; similarly for each character type, the F1 score was below 0.05. All the results are computed

using SVM (Vapnik 2013), and linear and rbf kernels. For each feature type, only the result obtained using the better performing kernel type is reported.

The overall best performing approach, as well as the best performing for four distinct dramatic character types, integrates the *Graph* and *N-grams* features. As expected, both the number of character type instances and the activeness of a character (see Table 2) influence the classification results. The presented joint treatment of features provides an improvement over the other approaches and demonstrates capacities to represent some of the important aspects of characters' dialogs and of their interactive characteristics. As measures extracted on the graphs weighted by linguistic markers encode information on centrality of a character and the relation to the most prominent character, this feature type helps in particular to discriminate heroes from other character types. By contrast, the utterance content features are particularly important for the characterization of representatives of the law, hero's supporting roles *ST vect.* and antagonists *N-grams*. The joint analysis of measures obtained from semantic dialog graphs and utterances' content *Graph + N-grams* leads to a significant improvement in the ability to identify less central character types such as A-SR, H-SK, and H-L.

A random single run (1 out of 5) using our integrative approach [*Graph + N-grams*], with the classifier trained on 80% of data and tested on the remaining 297 instances, resulted in 192 correct predictions (65%) and 105 misclassifications. These results were more closely inspected for the types of misclassifications that occurred. It turned out that 51% of the instances (54) were misclassified only inside their group of characters, e.g. H-SR as H-SK, so that in 83% of instances, at least the side the character is on was detected correctly. The rest of misclassifications were between groups, including the ambiguous or neutral LAW and BUS type characters. There were only two major errors where antagonist and hero were confused.

The Hilbert-Schmidt Independence Criterion (HSIC) (Gretton et al. 2005) analysis of linguistic markers and semantic dialog graph features indicates similarities in the hero's team, i.e. the characters that are socially close to the hero, in particular sidekicks, romantic/marital partners, and other supporting roles (Figure 4). Most notably, for all three groups of characters, the frequency of elliptic sentences is high. The expectation that the hearer can easily infer the missing Parts of Speech (PoS) seems to be used by screenwriters to underline closeness and existing common ground among the hero and his/her group. Berliner's (Berliner 2013) informal observation that, in movies, people are much more often called by name than in natural conversation can be specified with our data: the phenomenon of explicitly addressing other characters, by name, title, or pet name, is not general, but more characteristic of the hero(ine) and his/her supporting cast than of other groups. Supporting cast are also characterized by relative frequency of interjections. From a pragmatic viewpoint, this allows the writer to keep the lines of supporting

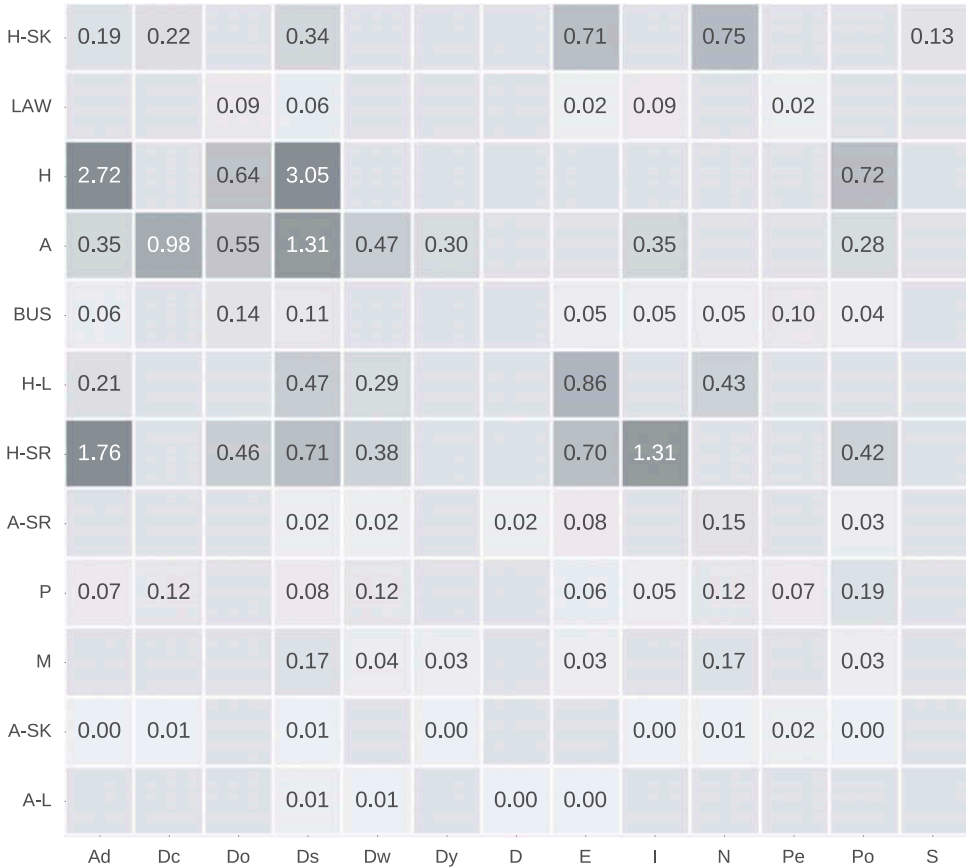


Figure 4. Selected linguistic and graph features per character type (listener, speaker) using dependence between variables and class label: **Address**; dialog act classes: **D**command, **D**other, **D**statement, **D**wh question, **D**yn question; **D**iscontinuities; **E**llipsis; **I**nterjection; **N**onstandard English lexicon; **P**erson; **P**olarity; **S**entiment neutral. The degree of dependence is stated numerically and reflected in the gray levels of cells. Blank cells indicate that the feature has not been selected.

characters short, but it can also be an indicator of their role as mediators of the audience’s intended emotional response (Lee and Marsella 2011).

Manual analysis of an action movie screenplay has shown that the hero character, contrary to expectations, delivers imperatives (commands) below average of all characters, while the antagonist (often the villain) has more than average: this finding is corroborated by the HSIC-based analysis, where the use of commands is the single most significant distinction between the antagonist and the rest of the cast.

While the presented results indicate that the main characters, in particular the hero and the antagonist, as well as their respective supportive roles can be detected automatically, the currently applied experimental setup does not allow for a complete evaluation of the described methods. The following primary restrictions can be identified: 1) limited number of instances available for each

character type, see Table 2, which significantly influences the classification performance; and 2) heterogeneous nature of character instances currently grouped as a single character type, e.g., compare the following mentor characters: Morpheus—*Matrix*, John Keating—*Dead Poets Society* and Red—*The Shawshank Redemption*. However, these two limiting factors contradict each other: the extension of the data set to other film genres implies an increase in the heterogeneity of character types, so that a gain in number of instances risks to be cancelled out by a further loss in categorization accuracy.

In order to explore alternative routes, we conducted an in-depth analysis of a single movie dialog, which is presented in the remainder of this section. The aim here is to gain a clearer picture of the differentiation and development of characters in dialog as keys to a finer-grained description of characters' verbal behavior patterns.

Single movie analysis

The complete dialog of an action movie, namely *Avatar* (2009), was annotated manually. The choice of this screenplay was motivated 1) by its length, which results in 1002 dialog lines; and 2) by the fact that it was written by the film director (John Cameron) and therefore is very close to the movie as shot. The purpose of this work was twofold: One goal was to gain insight into the dynamics of a character's development and relationships over the course of a film (character arc), see Section 4.5. The other goal was to answer the question if and how characters are distinguished by their dialog behavior within one film.

Annotation

The extracted dialog was first compared with the whole screenplay so as to check for completeness and correctness of the extract. Most categories for manual annotation were chosen such that, in principle, their automatic detection would also be possible:

- *Person*: the utterance can be in first, second, or third person; number (singular/plural) was not annotated.
- *Tense*: past (all past tenses), present, or future.
- *Addressee*: normally, the addressee is the other character involved in the dialog. Human annotators can infer in most cases who, out of several characters present in a scene, is the addressee, but annotation of more than one addressee was possible. "Invisible" addressees (e.g. the audience) were inferred.
- *Dialog Act*: the same coarse-grained categories were used as in the annotation of random dialog lines, i.e. s (statement), c (command), i (interjection), wh (wh-question), and yn (yes-no-question); commands were afterward extended from use of imperative to indirect commands

such as “we have to,” “I need you to.” This extension increased the overall proportion of command-like utterances from 184 to 253 (=1/4 of utterances), with the highest increase found for the antagonist character (Quaritch) with + 58%, and better reflected the informal impression of relative power status of the characters.

- *Connectedness*: we annotated separately whether an utterance is a reaction to a previous dialog line, and whether it is reacted to in the following turn. These parameters are meant to show the interconnectedness of characters. The annotation was here based on human understanding and could hardly be replicated automatically, but the intention was to explore whether they correlate reliably with other, detectable features.
- *Direction*: (Shen 2011) proposed this parameter to identify power relationships in written (e-mail) institutional conversations. The distinction between forward/backward orientation of utterances roughly corresponds to that between world-to-word and word-to-world fit in Searle’s Speech Act Theory (Searle 1969).

Analysis

Some simple statistical analyses were done on this annotated dialog, for which we only used the ten characters with the most lines. The four main characters together (Jake, Grace, Quaritch, and Neytiri) have nearly 70% of the 1000+ dialog lines. Among them, the protagonist alone has 328, hence nearly a third. The remaining six characters have between 29 and 55 lines.

- *Speakers and addressees*: The protagonist is addressed by other characters even more often than he speaks himself (ratio 1.08). By contrast, this ratio is particularly low for the antagonist (Quaritch, 0.32) with Grace and Neytiri in-between (0.64 and 0.84, respectively). Of the remaining characters, only Norm has a ratio of >1. This supporting character will be discussed separately below.
- *Person*: Only the protagonist uses the first person more often than second or third. The other three main characters use the second person most, which is in line with the frequency in which the protagonist is addressed by others. However, most supporting characters talk more in third person than in first or second. The runaway value here is, as with commands, provided by the character Max (see below).
- *Tenses*: As expected, the default is present tense, which accounts for almost 80% of utterances. The four main characters use the past tense more often than the future, although the difference is small for the antagonist. The two “natives” among the supporting characters use significantly more future than past tense. It is not clear whether this results from their specific roles in the movie or whether it is a stylistic particularity marking their foreignness.

- *Dialog act types*: Of particular interest to our analysis are commands as a direct expression of power. As mentioned above, the dialog act classification based on imperatives alone was extended to include indirect commands. In absolute numbers, the protagonist Jake utters most commands, followed by Grace and Quaritch, the antagonist. In relative terms, however, Quaritch has by far the highest proportion of commands, followed by Grace and Neytiri. Four of the six remaining characters have a higher rate of commands than the protagonist, with the character Max leading with nearly 45%.
- *Orientation*: (Shen 2011) considers the orientation of an utterance as an indicator of status in his study: the more “powerful” participants are in a position to plan, to order, to predict future events and actions, and to request information, while participants with lower status tend more to report or comment on past actions or statements, and to give answers. The manual annotation for orientation shows that, among the main characters, the antagonist has the highest proportion of forward-oriented utterances, followed by Grace. However, considering the whole group of 10 characters, the leader of the native people (Mo’at) has by far the highest proportion of forward-oriented turns. Not surprisingly, the use of future tense and of commands taken together predicts forward orientation sufficiently, so that this manually annotated power-related dimension can be derived from the automatic detection of tense and dialog act types.
- *Connectedness*: By annotating whether an utterance is a reaction (isreaction) to the previous one, or gets a reaction (getsreaction) in the following turn, we intend to reflect the social qualities of the film characters. While connectedness is the rule in natural conversation, film dialog often is not very dialogic at all, given that scenes can show only parts of conversations, or substitute images (e.g. of actions) for the verbal next turn. Such nonverbal events which are reactions or to which an utterance reacts are not represented in our data set, so that the criterion can only capture verbal connectedness. We annotated four states: an utterance isreaction, getsreaction, both, or none of these. Among the four main characters, Neytiri (the protagonist’s love interest and the mediator between the native people and the humans) has the highest proportion of utterances that are both reactions and reacted to, followed by the protagonist, while the antagonist is least connected. Among the remaining characters, only the antagonist’s supporting character (Selfridge) is similarly poorly connected. On the other hand, Norm, a protagonist’s supporting role, is the best connected of all characters. The search for automatically recognizable features that could serve as indicators of connectedness was less straightforward than for orientation. We found that the proportion of commands was the best predictor of non-connectedness, i.e. neither isreaction nor getsreaction. The proportion of isreaction utterances of a character correlates to the frequency with which a character is addressed by others. No

reliable substitutes could be found for the properties “getsreaction” and “both isreaction and getsreaction.”

Character profiling

We conducted an exercise in pairwise comparison of two characters to answer the question whether dialog behavior allows to draw a profile that provides hints to a character’s particularities and his/her roles and functions in the film. We chose two side characters, namely Max and Norm, both on the protagonist’s side. Max has 29 and Norm has 50 dialog lines out of the total of 1002 annotated utterances. Table 5 groups their dialog characteristics in comparison with the average of all annotated utterances (where applicable).

Max utters the highest proportion of commands among all characters, and consequently, the proportion of utterances in second person is also among the highest. By contrast, the proportion of 1st-person utterances is one of the lowest, and he never utters a question. Utterances in past or future tense are rare (in absolute numbers, we found only one of each). This character is firmly set in the present, makes no personal statements (e.g. of opinion or emotion), but commands others. The degree of connectedness is low: Max is not often addressed, and the ratio of active:passive addressing is low (0.58). Consequently, only a quarter of his utterances are reactions, and he gets reactions even less often. The proportion of utterances that both are and get a reaction is one of the lowest (6.9%), only the antagonists (Quaritch and Selfridge) appear as less connected and dialogic. Part of this can be due to the fact that a side character does not often get dialogs with several lines.

Norm, however, while also a side character, is the best connected. The proportion of utterances that both are and get reactions is the highest of the whole cast (18%), and nearly half of his utterances are reactions to others. Beside the hero, he is the only character who is addressed more often than he addresses

Table 5. Dialog characteristics of side characters Norm and Max as compared to total of annotated utterances.

Marker	Norm	Max	Average
Utterances	50	29	
% Commands (extended)	10.00	44.83	18.40
% Questions (yn + wh)	26.00	0.00	11.60
Is addressed	59	15	
Ratio active/passive addr.	1.11	0.58	0.82
% Is reaction	48.00	24.14	31.20
% Gets reaction	34.00	20.69	26.20
% Is and gets reaction	18.00	6.90	10.80
% 1st person	24.00	13.79	29.70
% 2nd person	24.00	44.83	33.80
% 3rd person	52.00	41.38	38.91
% Past	12.00	3.45	13.33
% Present	82.00	93.10	77.83
% Future	6.00	3.45	8.85

others (ratio of 1.11). This character utters few commands, but many questions: a quarter of his dialog lines are wh- or yn-questions, the highest rate of the cast. Norm also has the highest rate of utterances in 3rd person, while the use of second person is considerably below average: he talks about things and other persons, but rarely addresses the dialog partner directly.

These observations show significant differences in dialog behavior for these two side characters, but what do they mean for their characterization? Max is shown as the technical expert and support. He is the one who deals with the (technical) problems at hand, gives orders on the basis of his competence, but does not interfere with the larger problems and conflicts in the film. As a character type, he is the loyal, competent “second-in-command” on the hero’s side. We could expect to find similar characters of this type in films where the protagonist’s group is structured hierarchically (e.g. military unit, spaceship crew) and includes a character with technical competence and executive functions, and, within these boundaries, with considerable power. The case of the Norm character is more diffuse. The analysis of the dialog alone cannot reveal the functions of this character in the film. He is introduced first as a contrast to the hero: arriving at his assignment as a fully trained scientist, qualified and ready for the job, he highlights the hero’s helter-skelter recruitment and his initial ignorance, which earns him the scorn of the research team. Norm also highlights the exceptionality of the hero’s being accepted by the native Na’vi, mostly through his jealousy, and serves as a contrast to the hero’s physical abilities and warrior qualities. Nevertheless, the character is loyal to the protagonist’s side and puts his superior knowledge at the hero’s service by instructing him. In this function, he is given a considerable part of the expository text to speak, and, through knowledgeable questions, elicits more backstory especially from Grace. The integration of these functions, namely rival, tutor, backstory narrator, and contrast character, into one character is specific to this film, so that it is harder to talk about a generalizable character type here. These functions can be distributed or grouped differently in other films.

In summary, this exercise shows that it is possible to profile a character’s specific dialog behavior using a combination of the annotated features discussed in Section 4.4.2. It has to be noted that these particularities only emerge within one screenplay, i.e. in comparison with other characters and with the average behavior. This observation means that it is hard to detect a character type across a number of screenplays on the basis of the absolute presence of the features. What can be compared across films are only relative frequencies of certain forms and behaviors.

The exercise also shows that what is revealed with these annotations is primarily the function of a character in the film together with the power status with regard to the plot and to other characters. The fictional personality, i.e. what kind of person the character is intended to represent (Eder 2014) eludes this set of annotations. Their combination with lexico-semantic

analyses would reveal more about aspects of characterization, and could answer the question whether function and character are correlated: the usual list of character types (e.g. mentor, sidekick, etc.) is, basically, a list of functions but not of character in the psychological sense. This point is illustrated by the characters chosen for profiling: they would both be classified as supporting characters to the hero, so that the distinctions in verbal behavior are neutralized when the group is considered as a whole.

Time-dependent analysis

To study the question whether the annotated dimensions are able to reflect the “character arc,” i.e. the development of a character over the course of the film, we divided the same data (*Avatar* dialog) into 12 sections. For a partitionment that reflected more or less the timing of the film, we chose the original screenplay as a basis. The partitions were based on the number of lines in the screenplay, with minor adjustments to respect scene boundaries. The section boundaries were then transferred to the extracted dialog, so that the number of dialog lines per section varies considerably, depending on the proportion of descriptive text per sequence.

Plot point extraction

Based on the manually adjusted scene boundaries, 2D histograms that reflect the interaction frequencies of characters in the screenplay were extracted automatically.

These simplified representations were used to automatically identify major incidents in the plot (plot points) using the Sobolev norm (Wilson, Baddeley, and Owens 1995). The Sobolev norm was first introduced to compare grayscale images and is a natural choice for histogram comparisons. The results are illustrated in Figure 5, and the heatmaps used in this task are shown in Figure 6. Based on the computed structural differences in the plot, the plot points were extracted using a peak detection algorithm (Cormen et al. 2001).

Discussion of time-dependent analysis

The close fit of the histogram representation to the assumed three-act structure shows that it is possible to computationally detect the form of a screenplay. This step is a prerequisite for the analysis of the dynamics of relationships and interactions in a film, and a character’s development. The heatmap shows in more detail the nature of the changes that occur in the interaction patterns especially at the turning points of the plot. For example, the group of native characters first appears at the beginning of the second act sequence 4) which, in narrative terms, is characterized by the hero’s entrance into the world of his adventure (in this case very literally a foreign planet) (Vogler 2007). The varying patterns of the sequences in the second act are

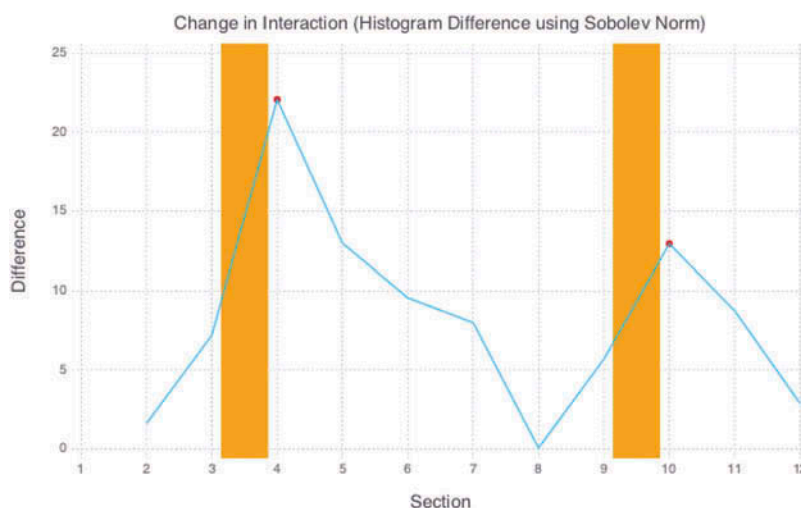


Figure 5. Structural differences in the plot *Avatar* based on differences obtained from histogram representations of 12 manually adjusted scene boundaries. The computed differences represent the change of character interaction compared to the previous scene. Automatically extracted plot points (peaks in the interaction profile) are shown as red dots. The plot points assumed on the basis of the common three-act structure are indicated by orange bars.

reflections of the dramatic complication. The transition to the third act is dramatically visible in sequence 10. The protagonist is now clearly the central figure on his side, all interaction is with him. Antagonist's and protagonist's sides are firmly separated and do not communicate with each other anymore: the stage is set for the final battle. (Truby 2007) The relevance of plot structure for characterization is illustrated with a simple example, using the characters' Jake (protagonist) and Quaritch (antagonist) behavior in giving orders (Figure 7). Although Jake utters more orders in absolute numbers (59 vs. 38), their percentage of the character's utterances is half as high (18% vs. 39%). The commands of both become more frequent toward the end of the film. A good part of Jake's orders are concentrated around the transition from second to third act, the dramatic turning point of the plot (sequences 8–10). The hero becomes not only the pivot of interaction, but also the leading character. The rate of commands uttered by Quaritch also rises in this phase, culminating in the final battle (sequence 11), when as many as 80% of his utterances have command character. This behavior is significant for the plot insofar as the refusal by some subordinates to obey his orders becomes the telltale sign of his ultimate defeat. If we consider the frequency of orders an indicator of status, the hero's character arc shows a clear increase in power toward the end. The antagonist's behavior shows a similar development so that the controversy between the “good” and the “bad” culminates. There is a difference in the style of commands between the two (roughly as between a team leader and an army officer), which have to be




Figure 6. Histogram representations of 12 manually adjusted scenes of the screenplay of *Avatar*. For better visual inspection, frequencies are colored in log space.

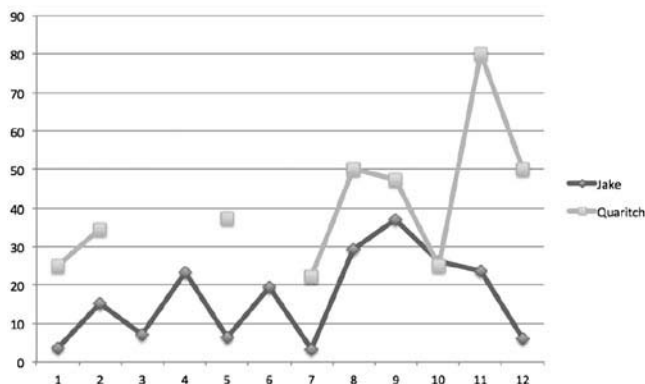


Figure 7. Percentage of commands and command-like utterances by protagonist (Jake) and antagonist (Quaritch) in *Avatar* over time.

revealed through other features of their utterances. Changes of verbal behavior over time can therefore provide insights into the development of characters in a film, and the major plot points are the moments where such changes should be looked for, in particular for the main characters who drive the plot.

Conclusions and future work

In this paper, we have described an integrated approach for the identification of dramatic character types based on sequences of dialogs extracted from action movie scripts. The presented set of features integrates cues on characters' presence, interactions, and content of utterances exchanged. Specifically, we applied an unsupervised representation of text, supervised quantization of vectors, and analysis of semantic dialog graphs incorporating social relationship and expressivity of characters. The approach was evaluated on a new data set of action movie characters.⁷ The results validate the presented multifaceted approach and a joint treatment of different aspects of characters' presence, their communication, and interactions in sequences of multiparty dialogs for the identification of main dramatic character types. The presented approach demonstrates an improvement over baseline methods, and also enables for a more in-depth analysis of features that convey characteristics of latent personae in sequences of dialogs. In-depth analysis of a single screenplay has provided insights into character profile and development that will be used in further research.

⁷The data set is available for research purposes—<http://github.com/anonymized>

Future work will therefore include the evaluation of the sequential (relative to progression of scenes) analysis of changes between different measurements inferred from the semantic dialog graphs to account for the relevant changes of the importance and roles of characters in a plot, and the introduction of distinct interpretations of interactions with other characters to represent different roles they may have in a character's development.

Funding

This research is partially funded by the Austrian Science Fund (FWF) under grant no. P 27530.

References

- Agarwal, A., et al. 2014. *Using determinantal point processes for clustering with application to automatically generating and drawing xkcd movie narrative charts* USA: Academy of Science and Engineering (ASE).
- Agarwal, A., J. Zheng, S. V. Kamath, S. Balasubramanian, and S. A. Dey. 2015. Key female characters in film have more to talk about besides men: Automating the bechdel test. Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Denver, Colorado, USA, 830–40.
- Alberich, R., J. Miro-Julia, and F. Rosselló. 2002. Marvel Universe looks almost like a real social network. *Arxiv Preprint*.
- Bamman, D., B. O'Connor, and N. A. Smith. 2014. Learning latent personas of film characters. Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL), 352.
- Bamman, D., T. Underwood, and N. A. Smith. 2014. A bayesian mixed effects model of literary character. Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL), (1):370–79.
- Beauchamp, M. A. 1965. An improved index of centrality. *Behavioral Science* 10 (2):161–3. doi:[10.1002/bs.3830100205](https://doi.org/10.1002/bs.3830100205).
- Bednarek, M. 2010. *The language of fictional television: Drama and identity*. New York, USA: Bloomsbury Publishing.
- Berliner, T. 2013. Killing the writer: Movie dialogue conventions and John Cassavetes. In *Film dialogue*, ed. J. Jaecle, 103–15. London/New York: Wallflower Press.
- Bordwell, D. 2007. *Poetics of cinema*. New York: Routledge.
- Cormen, T. H., et al. 2001. *Introduction to algorithms*. 2nd ed. Cambridge, MA, USA: McGraw-Hill Higher Education. ISBN: 0070131511.
- Culpepper, J., M. Short, and P. Verdonk. 1998. *Exploring the language of drama. From text to context*. London/New York: Routledge.
- Danescu-Niculescu-Mizil, C., and L. Lee. 2011. Chameleons in imagined conversation: A new approach to understanding coordination of linguistic style in dialogs. Conference Paper.
- Dijkstra, E. W. 1959. A note on two problems in connexion with graphs. *Numerische Mathematik* 1 (1):269–71. ISSN: 0945-3245. doi:[10.1007/BF01386390](https://doi.org/10.1007/BF01386390).
- Eder, J. 2014. *Die Figur im Film*, 2nd ed. Marburg: Schueren.
- Field, S. 2005. *Screenplay. The foundations of screenwriting*. New York: Bantam Dell.
- Freeman, L. C. 1977. A set of measures of centrality based on betweenness. *Sociometry*, 40:35–41.

- Gleiser, P. M. 2007. How to become a superhero. *Journal of Statistical Mechanics: Theory and Experiment* 2007 (09):P09020. doi:[10.1088/1742-5468/2007/09/P09020](https://doi.org/10.1088/1742-5468/2007/09/P09020).
- Gorinski, P. J., and M. Lapata. 2015. Movie script summarization as graph-based scene extraction. North American Chapter of the Association for Computational Linguistics – Human Language Technologies (NAACL HLT 2015), Denver, Colorado, USA.
- Gretton, A., et al. 2005. “Measuring statistical dependence with Hilbert-Schmidt norms”. In: *Algorithmic learning theory*, 63–77. Springer-Verlag Berlin, Heidelberg: Springer.
- Karsdorp, F., et al. 2015. The love equation: Computational modeling of romantic relationships in french classical drama. Sixth International Workshop on Computational Models of Narrative, Atlanta, GA, USA.
- Kemper, T. D. 1990. Social relations and emotions: A structural approach. In *Research agendas in the sociology of emotions*, ed. T. D. Kemper, 207–36. New York: State University of New York Press.
- Kiros, R., et al. 2015. Skip-thought vectors. *Advances in Neural Information Processing Systems*, 3276–84.
- Kolaczyk, E. D., and G. Cs’Ardi. 2014. *Statistical analysis of network data with R*. Springer-Verlag, New York, USA: Springer.
- Kozloff, S. 2000. *Overhearing film dialogue*. Berkeley, CA: University of California Press.
- Kuchenbuch, T. 2005. *Filmanalyse. Theorien, Methoden, Kritik*. Wien, Koeln, Weimar: Boehlau.
- Kundu, A., D. Das, and S. Bandyopadhyay. 2013. Scene boundary detection from movie dialogue: A genetic algorithm approach. *Polibits* 47:55–60. doi:[10.17562/PB-47-6](https://doi.org/10.17562/PB-47-6).
- Lazarus, R. S. 1991. *Emotion and adaptation*. Oxford: Oxford University Press.
- Lee, J., and S. Marsella. 2011. Modeling side participants and bystanders: The importance of being a laugh track. In *IVA*, eds. H.H. Vilhjálmsón, S. Kopp, S. Marsella, and K.R. Thórisson, 240–7. Berlin/ Heidelberg: Springer.
- Lind, P. G., M. C. González, and H. J. Herrmann. 2005. Cycles and clustering in bipartite networks. *Physical Review E* 72 (5):056127. doi:[10.1103/PhysRevE.72.056127](https://doi.org/10.1103/PhysRevE.72.056127).
- Mac Carron, P., and R. Kenna. 2012. Universal properties of mythological networks. *EPL (Europhysics Letters)* 99 (2):28002. doi:[10.1209/0295-5075/99/28002](https://doi.org/10.1209/0295-5075/99/28002).
- Myers, S. A., et al. 2014. Information network or social network?: The structure of the twitter follow graph. Proceedings of the companion publication of the 23rd international conference on World wide web companion. International World Wide Web Conferences Steering Committee, 493–8. Seoul, Republic of Korea.
- Otte, E., and R. Rousseau. 2002. Social network analysis: A powerful strategy, also for the information sciences. *Journal of Information Science* 28 (6):441–53. doi:[10.1177/016555150202800601](https://doi.org/10.1177/016555150202800601).
- Payr, S. 2007. So Let’s See: Taking and Keeping the Initiative in Collaborative Dialogues. In *Intelligent Virtual Agents. Proc. IVA’07, Paris*, eds. C. Pelachaud, et al., 175–82. Heidelberg: Springer.
- Payr, S. 2013. Virtual butlers and real people: Styles and practices in long-term use of a companion. In *Virtual butlers: The making of*, ed. R. Trapp. Heidelberg: Springer.
- Peberdy, D. 2013. Male sounds and speech affectations: Voicing masculinity. In *Film dialogue*, ed. J. Jaekle, 206–20. London/New York: Wallflower Press.
- Pennebaker, J. W., M. E. Francis, and R. J. Booth. 2001. Linguistic inquiry and word count: LIWC 2001. *Mahway: Lawrence Erlbaum Associates* 71:2001.
- Prabhakaran, V., and O. Rambow. 2013. Written dialog and social power: Manifestations of different types of power in dialog behavior. In Proceedings of the Sixth International Joint Conference on Natural Language Processing (IJCNLP), Nagoya, Japan, October.
- Price, S. 2013. *A history of the screenplay*. UK: Palgrave Macmillan.

- Quaglio, P. 2009. *Television dialogue*. Amsterdam: John Benjamins.
- Quirk, R., et al. 1985. *A comprehensive grammar of the english language*. London: Longman.
- Richardson, K. 2010. Multimodality and the study of popular drama. *Language and Literature* 19 (4):378–95. doi:[10.1177/0963947010377948](https://doi.org/10.1177/0963947010377948).
- Saramäki, J., et al. 2007. Generalizations of the clustering coefficient to weighted complex networks. *Physical Review E* 75 (2):027105. doi:[10.1103/PhysRevE.75.027105](https://doi.org/10.1103/PhysRevE.75.027105).
- Searle, J. R. 1969. *Speech Acts. An essay in the philosophy of language*. Cambridge: Cambridge.
- Shen, Y. K. 2011. Identifying power relationships in dialogues. Thesis, MIT.
- Suen, C., L. Kuenzel, and S. Gil. 2013. Extraction and analysis of character interaction networks from plays and movies. Proc. Digital Humanities 2013, Lincoln, NE, USA.
- Taavitsainen, I. 1999. *Writing in Nonstandard English*. Amsterdam: John Benjamins.
- Tieber, C. 2008. *Schreiben fuer Hollywood. Das Drehbuch im Studiosystem*. Vienna: LIT Verlag.
- Trappl, R., et al. 2011. Robots as Companions: What can we learn from servants and companions in literature, theater, and film? Proc. European Future Technologies Conference 2011 (FET11), 96–8.
- Traum, D. 2003. Issues in multiparty dialogues. In *Advances in agent communication*, 201–11. Frank Dignum: Springer.
- Truby, J. 2007. *The anatomy of story: 22 steps to becoming a master storyteller*. London, England: Macmillan.
- Turetsky, R., and N. Dimitrova. 2004. Screenplay alignment for closed-system speaker identification and analysis of feature films. Proceedings ICME'04, 2004 IEEE International Conference on Multimedia and Expo, 1659–62, Taipei, Taiwan.
- Vapnik, V. 2013. *The nature of statistical learning theory*. Springer-Verlag New York: Springer Science & Business Media.
- Vogler, C. 2007. *The Writer's journey. Mythic structures for writers*, 3rd ed. Studio City, CA: Michael Wiese Productions.
- Walker, M. A., et al. 2011. Murder in the Arboretum: Comparing character models to personality models. In Proc. of the 4th International Workshop on Intelligent Narrative Technologies, AAAI Press.
- Walker, M. A., G. I. Lin, and J. Sawyer. 2012. An annotated corpus of film dialogue for learning and characterizing character style. LREC, 1373–8.
- Watts, R. J. 2003. *Politeness*. Cambridge: Cambridge University Press.
- Watts, D. J., and S. H. Strogatz. 1998. Collective dynamics of 'small-world' networks. *Nature* 393 (6684):440–2. doi:[10.1038/30918](https://doi.org/10.1038/30918).
- Weng, C.-Y., W.-T. Chu, and J.-L. Wu. 2009. RoleNet: Movie analysis from the perspective of social networks. *IEEE Transactions on Multimedia* 11 (2):256–71. doi:[10.1109/TMM.2008.2009684](https://doi.org/10.1109/TMM.2008.2009684).
- Wilson, D. L., A. Baddeley, and R. A. Owens. 1995. A new metric for grey-scale image comparison. *International Journal of Computer Vision* 24:24–1.
- Zhu, Y., et al. 2015. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. Proceedings of the IEEE International Conference on Computer Vision. 19–27, Santiago, Chile.