

Article

A Graph-Learning Approach for Detecting Moral Conflict in Movie Scripts

Frederic René Hopp, Jacob Taylor Fisher and René Weber *

Media Neuroscience Lab, Department of Communication, University of California Santa Barbara, Santa Barbara, CA 93106, USA; E-Mails: fhopp@ucsb.edu (F.R.H.), jacobtfisher@ucsb.edu (J.T.F.), renew@ucsb.edu (R.W.)

* Corresponding author

Submitted: 14 April 2020 | Accepted: 24 June 2020 | Published: 13 August 2020

Abstract

Moral conflict is central to appealing narratives, but no methodology exists for computationally extracting moral conflict from narratives at scale. In this article, we present an approach combining tools from social network analysis and natural language processing with recent theoretical advancements in the Model of Intuitive Morality and Exemplars. This approach considers narratives in terms of a network of dynamically evolving relationships between characters. We apply this method in order to analyze 894 movie scripts encompassing 82,195 scenes, showing that scenes containing moral conflict between central characters can be identified using changes in connectivity patterns between network modules. Furthermore, we derive computational models for standardizing moral conflict measurements. Our results suggest that this method can accurately extract moral conflict from a diverse collection of movie scripts. We provide a theoretical integration of our method into the larger milieu of storytelling and entertainment research, illuminating future research trajectories at the intersection of computational communication research and media psychology.

Keywords

computational narratology; entertainment; eMFD; graph learning; MIME; moral conflict; movie scripts; network science

Issue

This article is part of the issue “Computational Approaches to Media Entertainment Research” edited by Johannes Breuer (GESIS—Leibniz Institute for the Social Sciences, Germany), Tim Wulf (LMU Munich, Germany) and M. Rohangis Mohseni (TU Ilmenau, Germany).

© 2020 by the authors; licensee Cogitatio (Lisbon, Portugal). This article is licensed under a Creative Commons Attribution 4.0 International License (CC BY).

1. Introduction

Art is fire plus algebra. (Jorge Luis Borges)

Humans invest a large amount of their time and money engaging with fictional narratives. These narratives may be described as “cohesive and coherent stor[ies] with an identifiable beginning, middle, and end that provide information about scene, characters, and conflict; raise unanswered questions or unresolved conflict; and provide resolution” (Hinyard & Kreuter, 2007, p. 778). Fictional narratives on TV, like *Game of Thrones*, attract more than 44 million views per episode (Everett & Crockett, 2019), and screenwriters create up to 15,000 movie scripts per year (Eliashberg, Elberse, & Leenders, 2006). Successful movies can rake in billions of dollars.

Avengers: Endgame achieved a record-breaking, world-wide gross of nearly three billion dollars (Box Office Mojo, 2019). Though, not all narratives are created equal. A majority of movies produced for entertainment underperforms, leaving producers and studios with a loss, and one of the main reasons for underperformance is a bad story (Eliashberg et al., 2006). Clearly, decision makers in the film industry face a critical decision when “choos[ing] among thousands of scripts to decide which ones to turn into movies” (Eliashberg, Hui, & John Zhang, 2007, p. 881).

But what exactly makes a ‘good’ story? Attempts to answer this question date back to the beginnings of storytelling itself (László, 2008), and communication and media psychology research have since formulated many different theories to explain narrative enjoyment (e.g.,

Oliver, 1993; Tamborini et al., 2013; Vorderer, Klimmt, & Ritterfeld, 2004; Zillman & Cantor, 1977). Some assert that a good narrative is all about the appeal and complexity of the characters (McKee, 2005), while others argue that a good story is mainly about a sequence of events (Franzosi, 2010). Recent evidence suggests that appealing stories emphasize interactions between the traits of characters, the events that befall them, and the specific sequence in which character traits and events are intertwined (Cron, 2012; Tamborini & Weber, 2020). In evolutionary media psychology, the fictional creation of narrative character and event chains has been hypothesized as an evolved principle to organize humans' cognitive machinery. According to this view, good stories are "attended to, valued, preserved, and transmitted because the mind detects that such bundles of representations have a powerfully organizing effect on our neurocognitive adaptations, even though the representations are not literally true" (Tooby & Cosmides, 2001, p. 21).

In this article, we assume that 'conflict,' and in particular 'moral conflict,' are crucial organizing principles permeating universally appealing narratives. In narrative psychology, conflict patterns are argued to undergird all fictional stories (László, 2008), elevating "the strength of any exciting character and story...without [conflict], characters don't have drive, desire, or desperation...there is no story, just words" (Ballon, 2014, pp. 49–51). Narratology further distinguishes external and internal conflict, with the former describing a character's goals and resistance from the environment (including interactions with other characters) and the latter occurring within a character and involving their internal needs (for an overview, see Szilas, Estupiñán, & Richle, 2018). Moral conflict—violating moral norms in order to uphold others (Tamborini, 2011, 2013)—has been related to the moral values of characters and to the ethical dimension of the story as a whole (Altman, 2008; McKee, 2005; Truby, 2007). In addition, moral conflict is regarded an especially salient component of cross-culturally appealing narratives (Knop-Huels, Rieger, & Schneider, 2019; Lewis, Grizzard, Choi, & Wang, 2019; Lewis, Tamborini, & Weber, 2014; Tamborini, 2013; Weber, Popova, & Mangus, 2013). Yet, empirical assessments of the kinds of moral conflict that are present in existing, fictional narratives are absent, with only a few studies simulating agent-based moral conflicts in digital interactive storytelling (e.g., Battagliano & Damiano, 2014). This lack of knowledge hinders media and communication scholars to develop an inventory of moral conflict in narratives and stymies further research relating moral conflict and narrative appeal. As a first step toward filling this gap, we herein address two interlocking research questions:

RQ1: How can moral conflict be conceptualized and mathematized?

RQ2: How can complex narratives be abstracted to guide the detection of moral conflict?

To address these questions, we propose a scalable, computational approach to identify moral conflict in story lines of movie scripts. This approach considers narratives as an evolving network of relationships between characters, leveraging the community structure of this network to reveal important points in the narrative's arc. Using a total of 894 movie scripts encompassing 82,195 scenes, we show that scenes in which characters from different network modules interact with one another often contain morally-charged conflict that pushes the narrative forward.

2. Moral Conflict and Narrative Appeal

Moral conflict permeates human culture and history, manifested in early philosophical discussions between Plato and Socrates on the relative priority of repaying one's debts over protecting others from harm; in ancient literature, such as Aeschylus's *Agamemnon*, where the protagonist has to decide between saving his daughter or leading Greek troops to Troy; in modern philosophy, including Jean-Paul Sartre's tale of a student who is torn between his personal devotion to his mother and the attempt to contribute to the defeat of an unjust aggressor; and in recent debates on algorithmic judgment during moral dilemmas (Awad et al., 2018). Traditionally, moral conflict has been defined as conflict between moral requirements: Moral reasons for adopting an alternative such that "it would be morally wrong not to adopt that alternative if there were no moral justification for not adopting it" (Sinnott-Armstrong, 1988, p. 12). When these requirements conflict, a moral dilemma arises if neither requirement overrides the other (i.e., is "stronger overall in some morally relevant way," Sinnott-Armstrong, 1988, p. 20). According to this view, moral conflict can only be resolved if one moral requirement overrides the other, otherwise there exists a moral dilemma.

Mounting evidence suggests that moral conflict motivates a variety of actions (Weber & Hopp, in press) and is essential for the construction, processing, and evaluation of fictional narratives (Altman, 2008; Eden, Daalmans, & Johnson, 2017; Lewis et al., 2014). Stories featuring moral conflict are processed and appraised in a slower, more rational fashion, compared to stories that do not feature conflict, and are therefore processed fast and intuitively (Lewis et al., 2014). In the parlance of entertainment research, exposure to morally conflicting stories elicits higher levels of appreciation, an entertainment experience characterized by thought-provoking, deeper, and meaningful insights into life (Knop-Huels et al., 2019; Lewis et al., 2014, 2019; Oliver & Bartsch, 2011). But why are audiences often drawn toward stories that prominently feature moral conflict, despite the cognitively effortful process of appraising and resolving these dilemmas? The Model of Intuitive Morality and Exemplars (MIME; Tamborini, 2011, 2013; for the latest update of the model see Tamborini & Weber, 2020) provides an answer to this question.

2.1. The Model of Intuitive Morality and Exemplars

The MIME describes short- and long-term reciprocal relationships between individuals' moral intuitions and their mediated and non-mediated environments (Tamborini, 2011, 2013). The MIME draws on Moral Foundations Theory (Graham et al., 2013) to conceptualize these moral intuitions. Moral Foundations Theory postulates the existence of five universal, innate moral sensibilities: care/harm, fairness/cheating, loyalty/betrayal, authority/subversion, and sanctity/degradation. The intuitive nature of these moral intuitions is described as "the sudden conscious appearance of a moral judgment, including an affective valence (good–bad, like–dislike) without any conscious awareness of having gone through steps of weighing evidence of inferring a conclusion" (Haidt, 2001, p. 818). Accordingly, the MIME suggests that media users experience an automatic, pleasurable response when their moral intuitions are reinforced, and an aversive reaction when these intuitions are violated, a prediction that has accumulated much evidence (e.g., Grizzard, Lewis, Lee, & Eden, 2011; Lewis et al., 2014; Tamborini et al., 2013; Weber, Tamborini, Lee, & Stipp, 2008).

The MIME asserts that individuals only exert these fast, pre-conscious evaluations in scenarios wherein a foundation is either upheld (e.g., care) or violated (e.g., harm) by character actions in a dominantly salient fashion (i.e., a moral foundation is distinctively upheld or violated). Conversely, within-foundation conflict (WFC) would arise if a foundation is both upheld and violated by a character's actions. For example, a character may need to enact (physical or emotional) harm in order to care for the (physical or emotional) well-being of another character. In contrast, between-foundation conflict (BFC) arises in situations in which different moral foundations are placed into conflict (Tamborini, 2011, 2013). For example, a character may physically harm innocent people (violating the care foundation) in order to be loyal to his/her group or country (upholding the loyalty foundation). To derive an affectively valenced response from these moral conflicts, individuals must decide whether the violation of one (or multiple) foundations as means of upholding another is justifiable in the scope of their individual moral preferences (also called moral intuition salience). While moral intuitions can exert chronic influence on moral judgments, stories may contain particular characters (e.g., highly moralized entities, such as Jesus or Ghandi) that may bias audiences' moral judgments in view of the character's general exemplification of particular foundations. For instance, a police officer may exemplify the fairness and care foundations in a story and at the same time may violate the authority foundation (subversion) by breaking the law (a 'bad cop' story). The initial exemplification may lead to a biased judgment of the police officer's motivation for violating another foundation.

3. Detecting Moral Conflict in Narratives

Attempts to computationally extract moral conflict from textual corpora are sparse. This is partially due to the latent, context-dependent nature of moral language that complicates moral information extraction (Weber et al., 2018) and missing operationalizations for moral conflict measurements. Recently, progress has been made toward the automated extraction of moral information from text (for an overview, see Hopp, Fisher, Cornell, Huskey, & Weber, in press). For example, Sagi and Dehghani (2014) combined dictionary-based content analysis with latent semantic analysis to identify moral rhetoric in news articles discussing the World Trade Center before and after the terror attacks on September 11th, blog posts and comments surrounding the Ground Zero Mosque, and speeches during the abortion debate in the US senate. Building on this work, Hopp et al. (in press) recently introduced the extended Moral Foundations Dictionary (eMFD) to computationally identify morally charged words in text. Compellingly, the eMFD is based on a large, crowd-sourced content-analysis study (Weber et al., 2018) and has been shown to outperform previous, word-count based moral extraction procedures.

We herein seek to build on this work to detect and mathematize moral conflict in narratives. To accomplish this aim, we theorize moral conflict as a set of hypotheses about the moral content and social network structure of narratives. We start with the simple assumption that moral conflict is likely to arise when characters from different groups, communities, or factions collide. Screenwriters frequently describe different communities of characters alternately, thereby constructing simultaneously evolving storylines (Weng, Chu, & Wu, 2007). Thus, as a story progresses, communities of characters come into focus as particular characters more frequently interact with each other. More specifically, the underlying cognitive process through which narrative consumers assign interacting characters to communities likely follows a stochastic, graph-learning process (Lynn & Bassett, 2019): The continuous registration of characters' spatiotemporal dependency leads viewers to develop expectations about characters' social relations, including their group affiliation, and hence expect particular characters to be more likely to 'flock together.' Based on this rationale and in line with narrative theory (Altman, 2008; Booker, 2004), we predict that narrative events in which characters from different communities collide highlight discrepancies across characters' group-based, moral motivations and thereby presage moral conflict (Tamborini & Weber, 2020).

To test these predictions, a computational translation is required that abstracts narratives into social network representations. In what follows, we make the bold claim that reducing a story to its main characters and their dialogue achieves a level of abstraction that enables computational learning of moral conflict.

3.1. Narratives as Social Networks of Characters in Movie Scripts

To examine characters and their interactions, mounting research draws on network science to conceptualize a narrative as an evolving social network in which nodes correspond to characters and edges between characters denote some dimension of their interaction (Ding & Yilmaz, 2010; Gleiser, 2007; Mac Carron & Kenna, 2012; Skowron, Trapp, Payr, & Trappl, 2016; Tran & Jung, 2015; Weng et al., 2007). This network representation of narratives has been fruitful in a variety of areas, from understanding character relations (Ding & Yilmaz, 2010; Park, Oh, & Jo, 2012) to identifying character types (Gleiser, 2007; Skowron et al., 2016) as well as learning leading roles, hidden communities, and storylines (Weng et al., 2007).

Due to their pre-structured format, movie scripts have been especially prominent for computationally constructing social networks of characters. Properly formatted movie scripts allow the distinction between several structural screenplay features such as scenes and stage directions, action descriptions, characters, and dialogue (Figure 1). Structural elements can subsequently be harnessed when computationally parsing a script, allowing one to extract, for example, which characters appear in the same scene or appear next to each other in dialogue. We utilize these structural features for computationally learning moral conflict patterns that permeate movie scripts.

4. Method

All data, code, analyses, and supplemental materials (SM) are made available at https://osf.io/rbdws/?view_only=24b117f22708457ca91f43ea2a4a6803

4.1. Movie Script Collection

We acquired scripts from the *ScriptBase J* dataset (Gorinski & Lapata, 2018), which contains 917 movie scripts. Each movie script is formatted in extended

markup language (XML), with tags corresponding to particular script elements (e.g., speech, speaker, scene description, etc.). We wrote a custom Python script to query the XML tree for each script to extract the scene number, stage direction, speaker, and dialogue for each script. 63 movie scripts had to be discarded due to formatting errors, resulting in a total of 894 movie scripts spanning 82,195 scenes.

4.2. Character Dialogue Networks

To create the character dialogue network for each movie (Figure 3a), we divided the corresponding script into individual scenes. We created a node in the network for each unique character that has at least one line of dialogue in the script. We incremented the weight of the edge between two characters by one each time the two characters had lines of dialogue immediately adjacent to one another in the scene. For example, consider a scene containing the set of characters $\{i, j, k\}$, and the sequence of lines $[i, j, k, i, k, i]$. The network for this scene would contain one node for each character, and would contain three edges: an edge of weight one between i and j ; an edge of weight one between j and k ; and an edge of weight three between i and k . The final edge weight between two characters thus reflects how often the two characters appeared next to each other across the entire movie script.

Modularity is commonly used to detect sub-units or communities within a network. The extraction of these communities is based on computations that partition networks into classes that maximize the density of links between nodes inside that class compared to links between classes. Applied to our character dialogue networks (Figure 3b), we expect modularity to reveal communities among characters such that characters assigned to the same module are more likely to interact with each other in contrast to characters from different modules. To extract topological information about the characters and their relationships, we calculated network modularity using the Louvain modularity maximization algorithm (Blondel, Guillaume, Lambiotte, & Lefebvre, 2008).

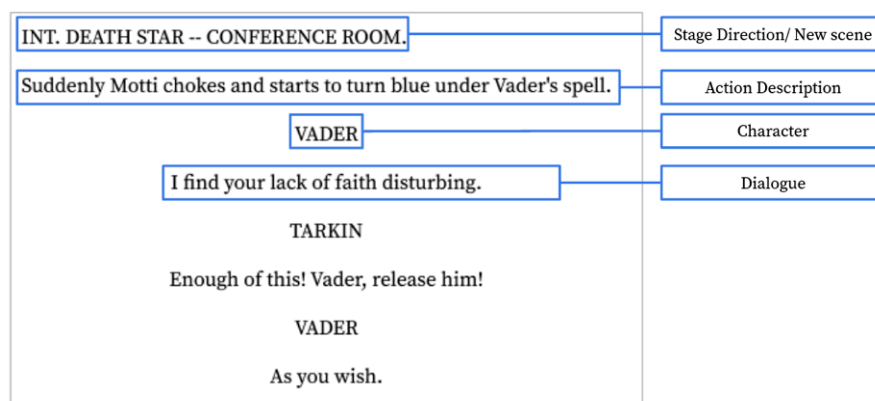


Figure 1. Elements and syntactic structure of a movie script excerpt from *Star Wars Episode IV: A New Hope* (Lucas, 1977).

4.3. Moral Conflict Scoring

4.3.1. Detection of Cross-Module and Within-Module Scenes

According to our network rationale for detecting conflict, we assume that the majority of conflicts—both moral and non-moral—occur between characters that belong to *different* communities as indexed by modularity. Hence, for each movie script, we label each scene as either a ‘cross-module’ scene if this scene contains characters from different modules or ‘within-module’ scene if this scene solely features characters from the same module (Figure 3c). Furthermore, we assert that the dialogue across character modules contains language cues that reveal: (a) whether the conflict has a moral or non-moral basis; and (b) how and which moral foundation(s) are conflicted.

4.3.2. Moral Content Extraction

To extract moral content from characters’ dialogue, we utilized the extended eMFD (Hopp et al., in press; Weber et al., 2018). The eMFD contains word lists in which each word is assigned a vector of five probabilities, denoting the likelihood that this word belongs to any one of the five moral foundations. These probabilities were derived from crowd-sourced content annotations and have been shown to improve moral signal detection compared to extant moral content classification procedures.

In turn, we tokenized each character’s dialogue and assessed which words in the dialogue are contained in the eMFD (Figure 3d). We then obtained the moral foundation probabilities for each detected word along with whether the given word denotes that the given founda-

tion was ‘upheld’ (positive valence, e.g., care) or whether it was ‘violated’ (negative valence, e.g., harm). Finally, all probabilities for words in the upheld and violated categories were averaged, returning a vector of ten scores per scene—two (upheld versus violated) per foundation.

4.3.3. Moral Conflict Measurement

We operationalized moral conflict as a situation in which one foundation (or perhaps more than one) is violated in order to uphold the same foundation or other foundations (Tamborini, 2011, 2013). Thus, moral conflict can occur in two ways: WFC describes situations in which the same foundation is simultaneously upheld and violated, whereas BFC can occur in contexts where some moral foundations are upheld while others are violated.

Furthermore, we assert that the degree of moral conflict is a function of: (1) the magnitude to which a foundation is upheld and violated; and (2) the polarity (i.e., difference) of simultaneously upholding and violating a foundation. Magnitude signals the general moral relevance of a particular event (e.g., killing might have a stronger moral relevance than hitting), whereas polarity denotes the degree to which conflicting moral requirements are incomparable, such that one (or multiple) moral requirement cannot override each other as “neither is stronger than, weaker than, or equal in strength to the other” (Sinnott-Armstrong, 1988, p. 58; see also Tamborini, 2013). To better capture this relationship, we can visualize the above model on a two-dimensional grid, where one axis reflects the degree to which a foundation is upheld, with the other axis capturing the degree to which the same (or another) foundation is violated (see Figure 2). Accordingly, the degree of moral conflict on this grid is a function of: (a) the magnitude of the up-




Figure 2. Computational model of moral conflict. Source: Authors.

held and violated foundation, where a higher magnitude results in moving further to the top right of the grid and increasing the degree of moral conflict; and (b) the polarity with which the moral requirements diverge from each other, with higher polarity indicating a greater distance to the diagonal of the grid. Mathematically, this relationship can be expressed as the Equation 1 shows below:

$$\text{Moral Conflict} = \sqrt{\frac{f_{i,u} + f_{i,v}}{2}} \times (1 - |f_{i,u} - f_{i,v}|)^2$$

Here, $f_{i,u}$ reflects the degree to which foundation i has been upheld and $f_{i,v}$ reflects the degree to which the same foundation i (in case of a WFC) has been violated. The first term expresses the average magnitude of upholding and violating a foundation: There will be no WFC if $f_{i,u} = 0$ or $f_{i,v} = 0$ or both $f_{i,u}$ and $f_{i,v} = 0$ (i.e., either only upholding or only violating a foundation, or neither upholding nor violating a foundation). Likewise, there will be maximum WFC if $f_{i,u} = 1$ and $f_{i,v} = 1$ (i.e., simultaneous, maximum upholding and violating a foundation). The second term integrates the polarity between virtue and vice: When a foundation is both maximally upheld and maximally violated ($f_{i,u} = 1$ and $f_{i,v} = 1$) the polarity between upheld and violated is maximized. In this scenario, there should not be any regularization of the moral conflict score and hence $f_{i,u} - f_{i,v} = 0$. However, if a foundation is relatively more upheld or violated, this should cause a decrease in moral conflict as either the upholding or the violation of the foundation becomes overwhelmingly salient (Sinnott-Armstrong, 1988; Tamborini, 2011, 2013).

Building on this formula, we can express within- and between-foundation conflict in a straightforward manner. For WFC, the degree to which the same foundation is both upheld and violated can simply be computed by entering the virtue ($f_{i,u}$) and vice ($f_{i,v}$) score of the same foundation (e.g., care and harm; fairness and cheating; etc.) into Equation 1. Hence, five WFC scores can be computed that reflect the degree to which any one or all foundations are internally conflicted.

Analogously, we can compute the degree of BFC by entering the virtue ($f_{i,u}$) score of one foundation (e.g., care) and the vice ($f_{j,v}$) score of a different foundation (e.g., cheating). To obtain the average degree to which a particular foundation (e.g., care) is conflicted with all other foundations, the following Equation 2 can be applied:

$$\text{BFC} = \frac{\sum_{1 \leq j \leq 5, j \neq i} \sqrt{\frac{f_{i,u} + f_{j,v}}{2}} \times (1 - |f_{i,u} - f_{j,v}|)^2}{4}$$

In this equation, $f_{i,u}$ reflects the degree to which foundation i has been upheld (e.g., care) and $f_{j,v}$ reflects the degree to which a different foundation j has been violated (e.g., cheating). Accordingly, for every foundation, four conflict scores are calculated to measure the degree to which a single foundation has been upheld while violating any of the remaining four foundations. Averaging across these sums captures the mean degree to which a particular foundation has been upheld while violating all other foundations. WFC and BFC can also be expressed in terms of a conflict matrix C (see Table 1), where each cell reflects a computed moral conflict score. The diagonal of this matrix denotes the previously introduced WFC, whereas values on the off diagonal capture BFC. Summing the diagonal of this matrix produces the total WFC, whereas summing the upper and lower triangular elements of C produces the total BFC.

5. Results

We first provide a detailed, small-scale evaluation of our moral conflict detection algorithm across a set of three diverse movie scripts. Thereafter, we present results of a large-scale validation of our algorithm, scaling it up to a total of 894 screenplays.

5.1. Character Networks and Modularity

We first evaluated the computational construction of character networks and assignment of modularity

Table 1. Moral conflict matrix representing WFC and BFC.

| i, j | Harm | Cheating | Betrayal | Subversion | Degradation |
|-----------|---|---|----------|------------|-------------|
| Care | $\sqrt{\frac{f_{i,u} + f_{i,v}}{2}} \times (1 - f_{i,u} - f_{i,v})^2$ * | $\sqrt{\frac{f_{i,u} + f_{j,v}}{2}} \times (1 - f_{i,u} - f_{j,v})^2$ | . | . | . |
| Fairness | $\sqrt{\frac{f_{i,u} + f_{j,v}}{2}} \times (1 - f_{i,u} - f_{j,v})^2$ | .* | . | . | . |
| Loyalty | . | . | .* | . | . |
| Authority | . | . | . | .* | . |
| Sanctity | . | . | . | . | .* |

Note: Cells with * represent WFC, and the ones without represent BFC.

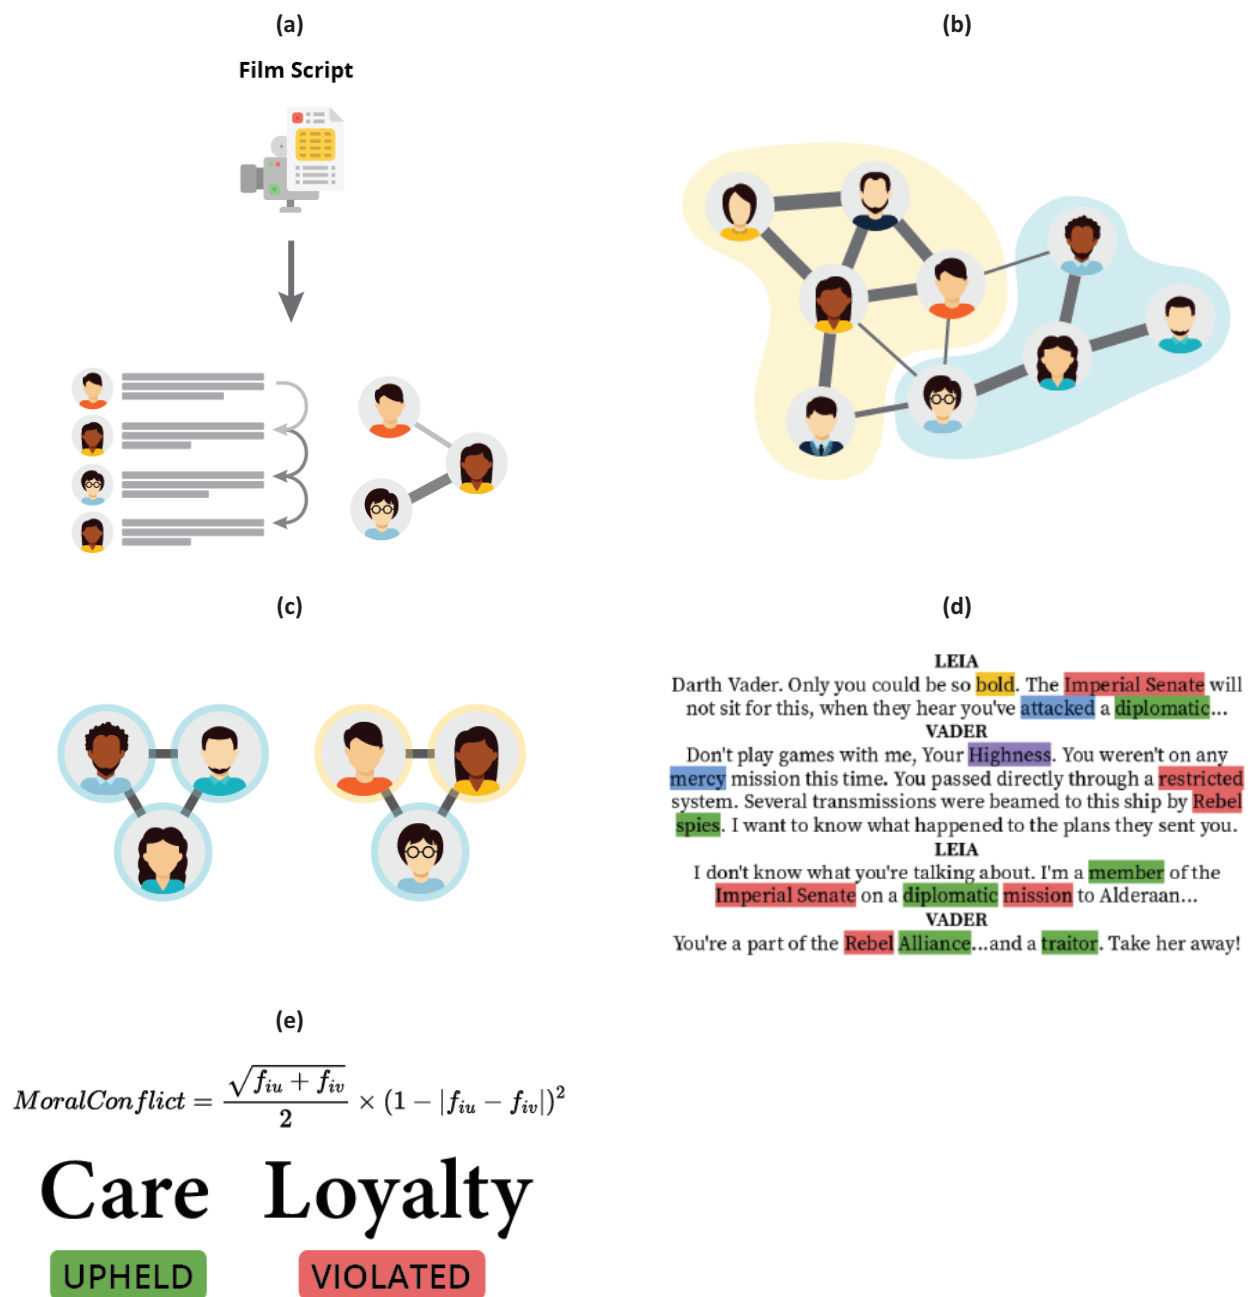


Figure 3. Flowchart of the Moral Conflict Detection Pipeline: (a) Network construction; (b) Modularity maximization; (c) Extracting within-module (left) and cross-module (right) scenes; (d) Coding moral content; (e) Calculating moral conflict. Source: Authors.

classes across three popular movie scripts. Figure 4 demonstrates that character networks provide a high-level summary of each movie, containing main characters and their relative interactions. Corroborating previous findings (Gleiser, 2007; Weng et al., 2007), modularity maximization performs well as a means of partitioning the character networks into different modules, such that characters belonging to the same module more frequently interact with each other compared to characters from a different module. For example, in *Star Wars Episode IV*, the largest module consists of the main hero

(Luke) and his closest allies (Ben, Han, Leia, and Threepio), whereas the second largest module contains the main antagonist (Vader) along with his imperial allies (e.g., Tarkin, Motti, Tagge; Figure 4a). Similar 'good versus evil' communities are detected in *The Matrix* (Figure 4b). However, as the character network for *The Lord of the Rings: The Fellowship of the Ring* (Figure 4c) reveals, modules do not always reflect the social networks around heroes and villains, but rather capture which characters are co-present throughout alternately evolving storylines (Weng et al., 2007).



Figure 4. Character dialogue networks: **(a)** *Star Wars: Episode IV: A New Hope*; **(b)** *The Matrix*; **(c)** *The Lord of the Rings: The Fellowship of the Ring*. Notes: Node size reflects degree centrality, edge weight captures total co-occurrence in dialogue, node color illustrates modularity class. Networks are laid out using the force-directed Fruchterman-Reingold algorithm. See SM Sections 1 and 2 for scree plots and character dialogue adjacency matrices. Source: Authors.

5.2. Case-Study Evaluation of Moral Conflict

5.2.1. Star Wars Episode IV

A total of 20 cross-module scenes (Figure 5b) were detected for which we extracted the moral content across dialogues and calculated the degree of moral conflict (Figure 5b). Overall, the results are promising in that cross-module scenes indeed frequently contain moral conflict: For example, in scene 16, *Vader* and his allies are accusing *Leia* of passing through a restricted system and being part of the *Rebel Alliance*, which both are indicators of violating authority. Yet, *Leia* stresses the authority of the *Imperial Senate* to punish *Vader* for interfering in a diplomatic mission. Accordingly, as indexed by

the scene's corresponding conflict matrix, there is strong WFC between authority and subversion. Moreover, in scene 89, *Leia* faces the BFC of whether to reveal the location of the hidden rebel base (betrayal, but following the authoritative order of *Tarkin*), or keeping loyal to her alliance at the cost of being partially responsible for the killing (harm) of innocent people should she not cooperate.

5.2.2. The Matrix

In *The Matrix*, a total of 13 cross-module scenes were identified (Figure 6A). In scene 19, for example, the highest WFC occurs within the loyalty foundation, denoting the emphasis on *Neo's* double-life as a program writer

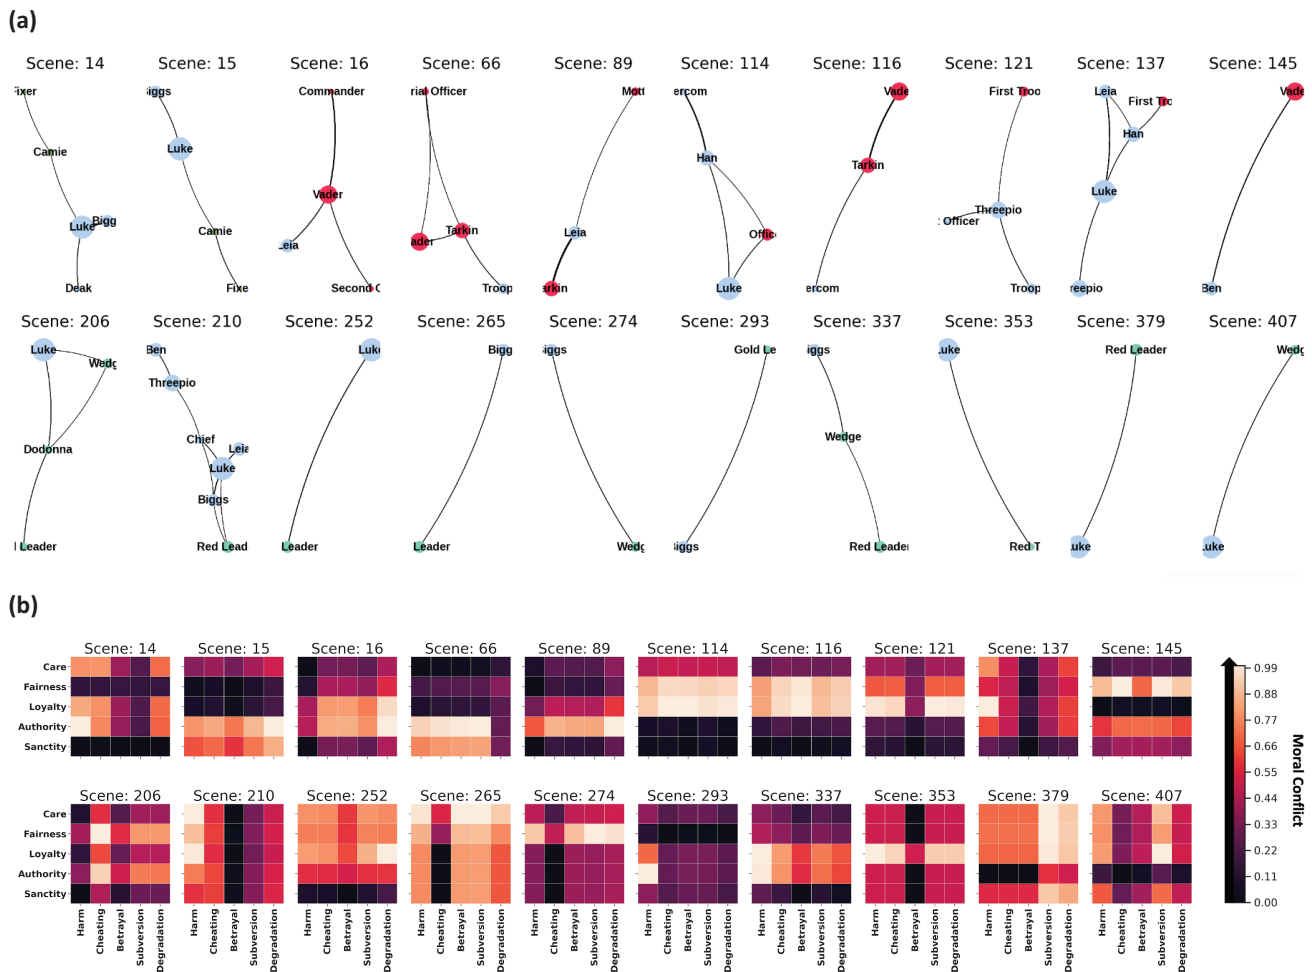


Figure 5. Moral conflict in *Star Wars Episode IV*: **(a)** Cross-module scenes and character interaction networks; **(b)** Moral conflict matrices. Source: Authors.

and a computer hacker who is “guilty of virtually every computer crime we have a law for” (Wachowski & Wachowski, 1999). Likewise, *Agent Smith* stresses that he “believe[s] [Neo] want[s] to do the right thing” (Wachowski & Wachowski, 1999) yet *Neo* emphasizes that he is denied his rights for not having the opportunity to make a phone call.

When evaluating the representation of moral conflict in scene 56, it becomes apparent that harm and degradation are most salient and conflicted: *Agent Smith* multiple times emphasizes the deal between him and *Cypher*, “I’ll get you what you want” (Wachowski & Wachowski, 1999). While not explicitly mentioned in the dialogue, this scene clearly contains moral conflict: Deciding whether *Cypher* should betray his allies in return to being loyal to upholding the deal with *Agent Smith*.

5.2.3. The Lord of the Rings: The Fellowship of the Ring

In *The Lord of the Rings: The Fellowship of the Ring*, a total of 35 cross-module scenes were identified (Figure 7). Remarkably, the majority of cross-module scenes were

high in (moral) conflict: The secret council meeting in *Rivendell* (scene 94), *Boromir*’s attempt to take the ring from *Frodo* (scene 133), the breaking of the fellowship and the death of *Boromir* (scene 136), and the appearance of the *Crebain* in the *Eregion Hills* (scene 98). For instance, in scene 94, the main conflict revolves around the appropriate manner in dealing with the one ring, which induced the highest conflict, both within and across fairness (e.g., who should be given the ring), loyalty, (e.g., who can be trusted with the ring), and authority (e.g., who decides what happens with the ring). Likewise, in scene 133, moral conflict emerges from *Boromir*’s attempt to steal the ring from *Frodo* (cheating), although assuring that he is no thief (fairness). Furthermore, *Boromir* accuses *Frodo* of betrayal, which explains the high degree of conflict surrounding the loyalty foundation.

5.3. Large-Scale Validation of Moral Conflict Detection Algorithm

The previous section demonstrated convergent validity of our moral conflict algorithm across three selected

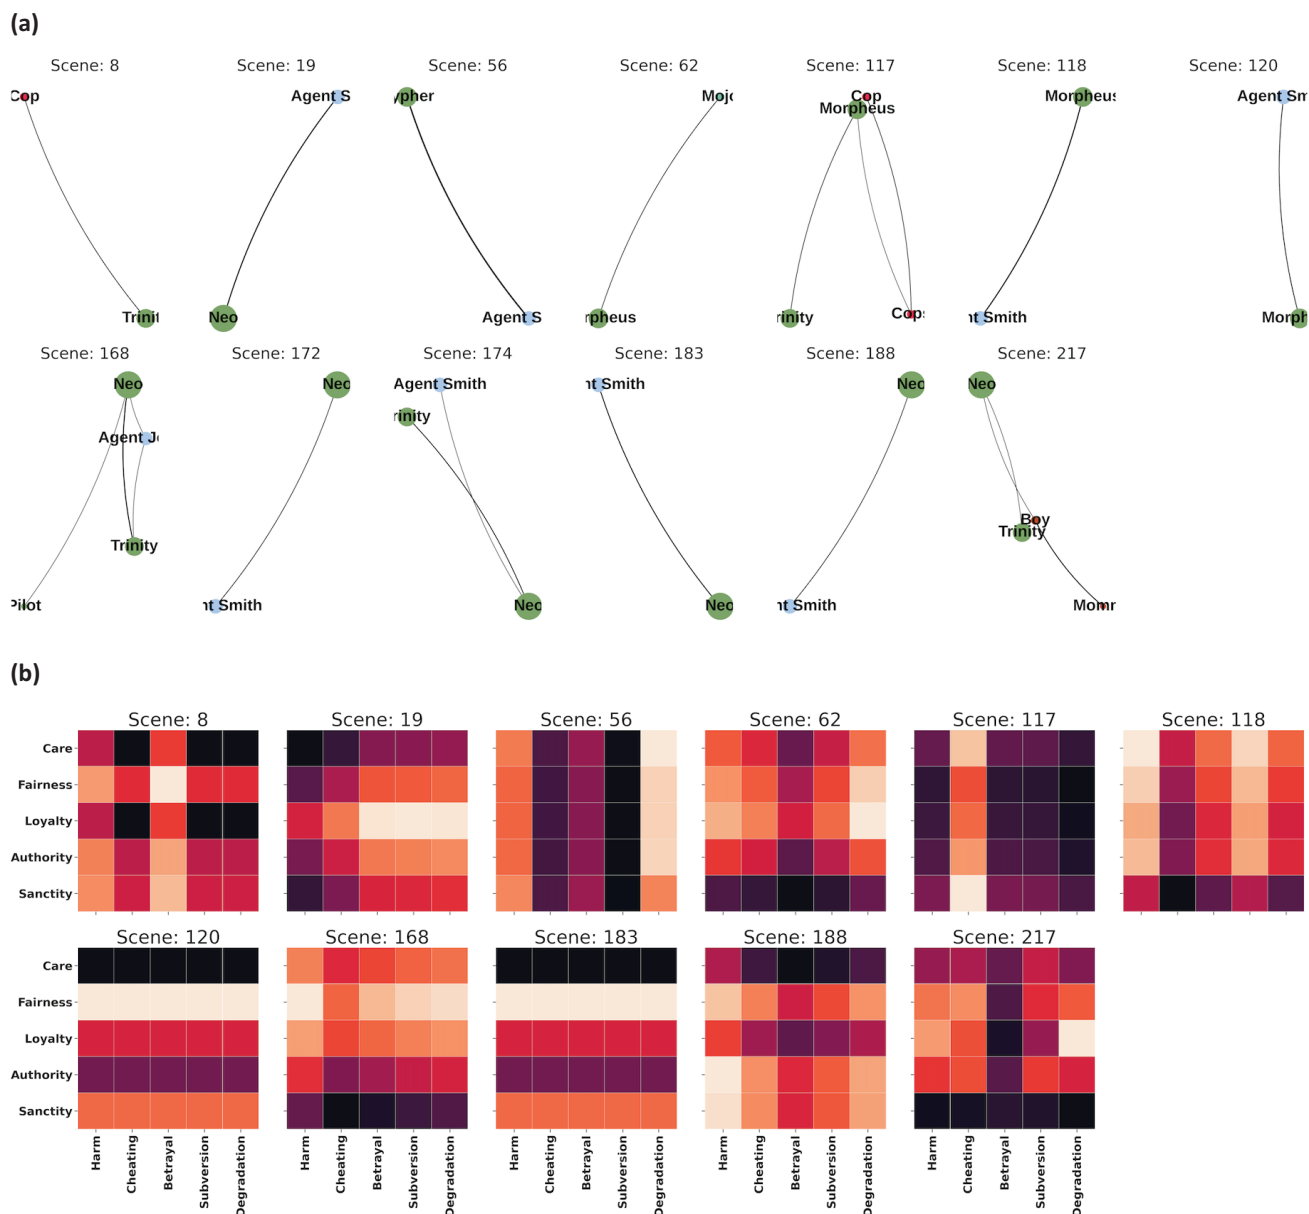


Figure 6. Moral conflict in *The Matrix*: (a) Cross-module scenes and character interaction networks; (b) Moral conflict matrices. Notes: Two cross-module scenes (172 and 174) were discarded for moral conflict computations as no moral content was identified. Source: Authors.

movies. To corroborate these findings across a larger set of screenplays, we scaled our algorithm to 894 movie scripts, containing a total of 82,195 scenes. First, we found a much higher frequency of within-module scenes ($N = 69,042$) compared to cross-module scenes ($N = 13,153$), although the ratio of cross-module to within-module scenes varies across movies (see Figure 3 in the SM). Furthermore, this result can be expected for two reasons: First, the modularity maximization algorithm assigns characters to modules that do not interact with each other frequently. As such, scenes in which characters from different modules are co-present will be rare. Second, screenplay writers frequently describe different communities of characters alternately, such that there will be a higher representation of scenes featuring

characters from the same module compared to scenes in which characters from different modules are present.

Next, we tested the hypothesis that cross-module scenes are more moralized than within-module scenes. Compellingly, we indeed found that across moral foundation categories, cross-module scenes contain on average more moralized language as identified by the eMFD compared to within-module scenes (see Figure 8a). In a final step, we tested whether cross-module scenes are higher in moral conflict than within-module scenes. As illustrated in Figure 8b and Figure 8c, we again found support for our prediction that cross-module scenes are consistently higher in moral conflict than within-module scenes. Interestingly, we observed highest BFC for the loyalty foundation, which suggests that characters are




Figure 7. Moral conflict in *The Lord of the Rings: The Fellowship of the Ring*: (a) Cross-module scenes and character interaction networks; (b) Moral conflict matrices. Notes: Two cross-module scenes (85 and 106) were discarded for moral conflict computations as no moral content was identified. Source: Authors.

more concerned with upholding their in-group loyalty at the cost of violating other moral foundations. In contrast, we found that WFC is highest for the care foundation, indicating that characters frequently face situations in which they, for example, have to commit an act of (physical or emotional) harm in one context in order to uphold (physical or emotional) care in another. Likewise, we observed that BFC is higher than WFC across the loyalty and authority foundations, whereas WFC is higher in care and fairness foundations.

6. Discussion

This article introduced a computational approach for extracting moral conflict from movie scripts. While moral conflict is an important factor for story development (Ballon, 2014) and narrative appeal (Eden et al., 2017;

Knop-Huelss et al., 2019; Lewis et al., 2014), no tools exist for the large-scale, standardized extraction of moral conflict from narratives. Our suggested approach provides a first step towards the automated, computational assessment of moral conflict in narrative texts. We found that network structures among characters serves as a useful heuristic for identifying movie scripts' main characters and their interactions. Furthermore, we demonstrated that modularity maximization is a promising method to extract communities of characters within a story. Lastly, we showed that scenes in which characters from different network modules co-occur are more moralized and higher in moral conflict than scenes in which characters from the same network module are co-present. In order to measure the degree and kinds of moral conflict that permeate the dialogue of rival characters within identified conflict scenes, we derived two different metrics:

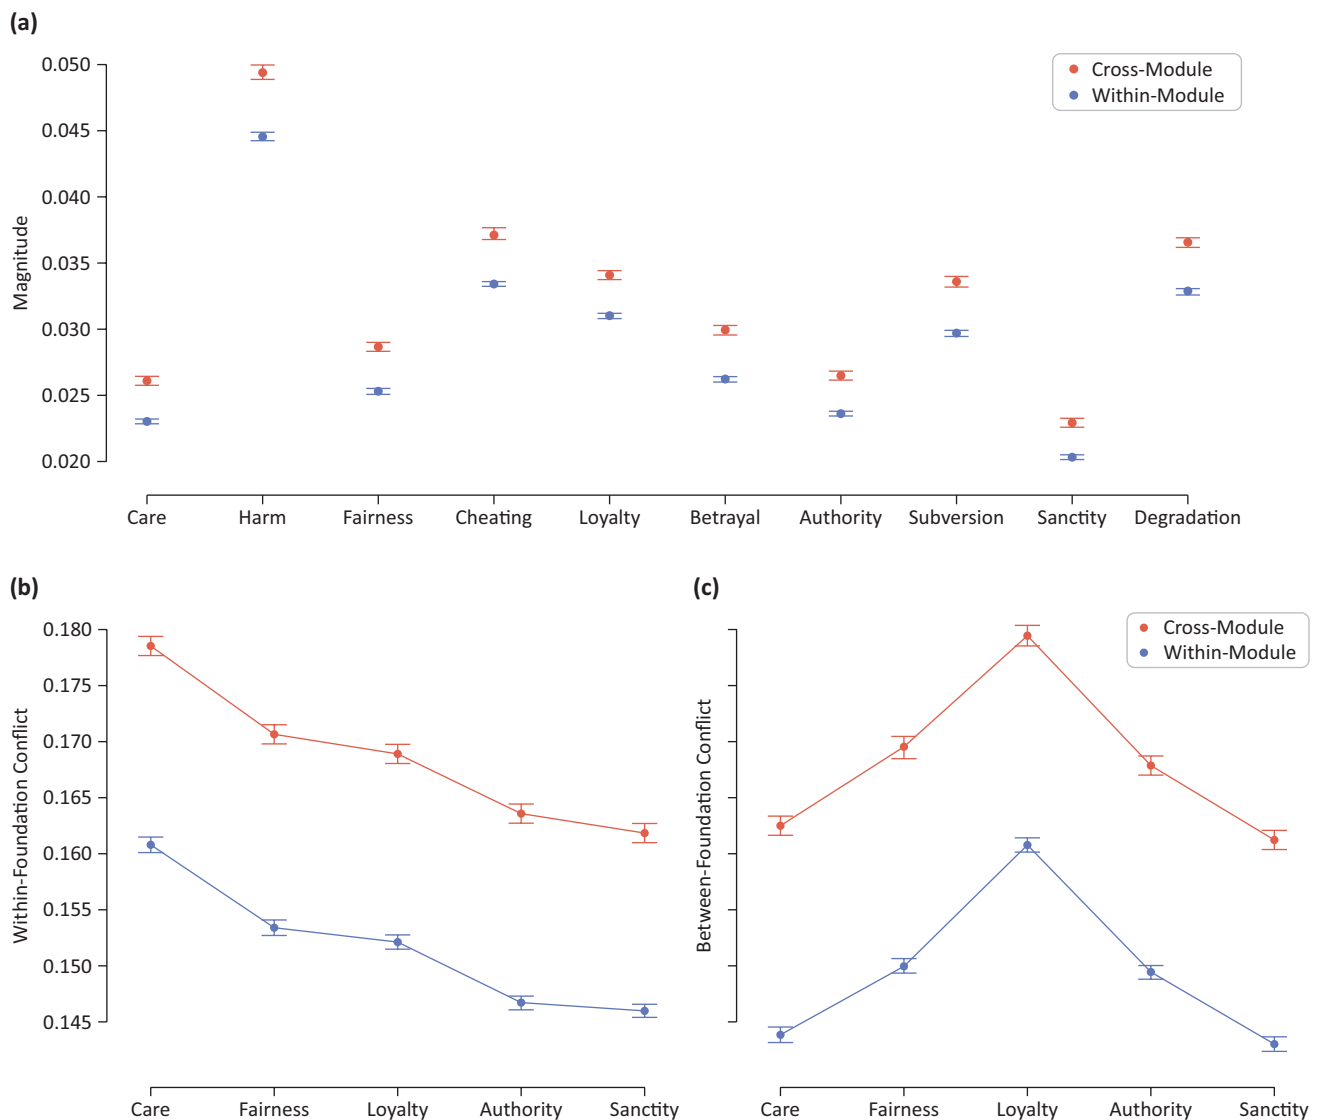


Figure 8. Moral content across within-module and cross-module scenes: **(a)** Representation of moral foundations; **(b)** Representation of within-foundation conflict; **(c)** Representation of between-foundation conflict. Notes: $N = 894$ screenplays. Error bars represent 99% confidence intervals estimated with 1,000 bootstrap iterations. Outliers above 95% were removed. Source: Authors.

First, we operationalized WFC as the degree to which distinct moral foundations are simultaneously upheld and violated; second, we defined BFC as the degree to which a single moral foundation is upheld, while other moral foundations are being violated. Our results indicate that these measures demonstrate acceptable validity as indicated by their alignment with expressed moral sentiment and conflict in characters' dialogue.

6.1. Implications

The herein introduced methodology is inspired by recent research synergizing structural and content features of narratives (e.g., Skowron et al., 2016) to push the envelope of computational entertainment research. Specifically, our finding that cross-module scenes are richer in moral content and moral conflict compared

to within-module scenes is compelling for two reasons: First, it emphasizes the attentional capture of moral cues (Gantman & Van Bavel, 2015) for assisting viewers in directing their attention towards important plot points; second, the collision of characters from different network modules likely highlights characters' group-identity and reinforces motivations for coalition building and group cohesion. Hence, to emphasize distinctiveness between group affiliations, screenwriters may attribute characters in cross-module scenes with more moralized dialogue in order to highlight group differences and evoke moral conflict. This leads to the interesting future prediction whether: a) Character network modules are more homogenous in their moral language use; and b) whether this similarity increases during cross-module scenes where group-identity becomes hyper-salient.

Moreover, our herein developed approach paves the way for testing various central predictions of the MIME at an unprecedented scale. First, if moral conflict is indeed a central predictor of narrative appeal, then our developed conflict metrics should correlate with viewership and story evaluations, including movie performance as indexed by box office sales or film critic ratings (Weber, Hopp, & Fisher, 2020). Second, our approach may illuminate which kinds of moral conflict are more likely to engage particular morality subcultures: communities among audiences with distinct moral sensitivities (Mastro, Enriquez, & Bowman, 2013). Ongoing research demonstrates that subgroups among audiences differ in their moral sensitivities and hence tend to enjoy different moral content patterns (Bowman, Jöckel, & Dogruel, 2012). Accordingly, future work may examine how violating particular moral norms in order to uphold others is enjoyed in some morality subcultures but not in others. Accordingly, we envision that the evaluation of a story's moral conflict pattern may be of assistance to screenwriters during the script creation process and furthermore inform decision-makers in the film industry during content marketing (see e.g., The Moral Narrative Analyzer [MoNA], <https://mona.medianeuroscience.org>; Dramatica, <https://dramatica.com>; StoryFit, <https://www.storyfit.com>).

6.2. Limitations and Future Directions

Although our findings are promising, they have limitations. First, our approach does currently not allow for the tracing of the dynamic development of moral conflict related to specific characters, as it considers moral conflict as a holistic concept that is expressed across multiple characters within particular scenes. This limitation largely arises from the fact that the edges in our character networks are undirected and hence do not allow for an assessment of 'who is talking to whom, with what message, and with what effect.' Put differently, characters following each other in dialogue do not necessarily interact with each other, although we assume that dialogue co-occurrence is a more precise proxy for interaction compared to the commonly used scene co-occurrence (see e.g., Gorinski & Lapata, 2018; Weng et al., 2007). A consideration of the directed edges between particular characters may enable more in-depth analyses of moral conflicts between particular characters. As such, we are currently collecting human-annotated data to more precisely detect speaker-addressee relationships across movie script dialogue. Second, we herein applied modularity maximization to detect potential conflicts between different communities of characters. Hence, our approach does not account for potential conflicts that may arise *within* a given community. Future work may thus experiment with more fine-grained modularity parametrization to 'modularize within modules.' Relatedly, modularity maximization might not always extract modules that resem-

ble classical 'good versus evil' splits as illustrated in *Star Wars Episode IV* or *The Matrix*. As such, future work may extend the herein introduced approach to detect stereotypical personas of film, such as heroes and villains (Bamman, O'Connor, & Smith, 2013). Analogously, we herein focus on external moral conflicts as expressed in dialogue between characters. Yet, recent advancements in the MIME (Tamborini & Weber, 2020) suggest that moral conflicts may also arise within characters, expressed as the collision of individuals' egoistic motivations and altruistic concerns for others. Moreover, our use-case of structured movie scripts reflects a 'best-case' scenario for computationally extracting moral conflict from text corpora. Yet, we think that many of the herein introduced steps can be mapped onto more unstructured texts to identify moral conflict in other media formats. For example, news articles may be promising to understand how moral conflict evolves in public discourse and dynamically shapes and is shaped by unfolding sociopolitical events (Hopp, Fisher, & Weber, 2020).

Perhaps most importantly, we encourage future researchers to further validate the herein established conflict computations. First, permutation tests in which words of a script are randomly assigned to scenes may corroborate our findings if our moral conflict matrices consistently deviate from null models. Likewise, it would be helpful to randomly permute characters to scenes in order to examine whether detected modules are preserved or destroyed. In addition, we suggest contrasting the computationally derived moral conflict measures against a crowd-sourced, manually annotated dataset. Doing so may reveal, among other things: (a) Whether humans attribute cross-module scenes higher conflict ratings compared to within-module scenes; or (b) whether algorithmically derived moral conflict matrices correlate with humans' judgment of conflicted moral foundations.

We are optimistic that future applications and enhancements of the computational algorithms can serve to help media psychology researchers—as well as industry professionals—decipher and learn the types of moral conflict that permeate stories, and how these conflicts influence how we engage with and respond to narratives in our everyday lives.

Acknowledgments

Contract grant sponsors: John Templeton Foundation (to R.W.), contract grant number: 61292. We extend our thanks to Cole Hawkins and Isaac Mackey for helpful discussions on moral conflict computations.

Conflict of Interests

The authors declare no conflict of interests.

References

- Altman, R. (2008). *A theory of narrative*. New York, NY: Columbia University Press.
- Awad, E., Dsouza, S., Kim, R., Schulz, J., Henrich, J., Shariff, A., . . . Rahwan, I. (2018). The moral machine experiment. *Nature*, 563(7729), 59–64.
- Ballon, R. (2014). *Blueprint for screenwriting: A complete writer's guide to story structure and character development*. Mahwah, NJ: Routledge.
- Bamman, D., O'Connor, B., & Smith, N. A. (2013). Learning latent personas of film characters. In H. Schuetze, P. Fung, & M. Poesio (Eds.), *Proceedings of the 51st annual meeting of the Association for Computational Linguistics* (pp. 352–361). Stroudsburg, PA: Association for Computational Linguistics. Retrieved from <https://www.aclweb.org/anthology/P13-1035.pdf>
- Battagliano, C., & Damiano, R. (2014). A character model with moral emotions: Preliminary evaluation. In M. A. Finlayson, J. C. Meister, & E. G. Bruneau (Eds.), *5th workshop on computational models of narrative* (pp. 24–41). Wadern: Dagstuhl Publishing. <https://doi.org/10.4230/OASlcs.CMN.2014.24>
- Blondel, V. D., Guillaume, J. L., Lambiotte, R., & Lefebvre, E. (2008). Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2008(10), P10008. <https://doi.org/10.1088/1742-5468/2008/10/P10008>
- Booker, C. (2004). *The seven basic plots: Why we tell stories*. London: Bloomsbury Academic.
- Bowman, N. D., Jöckel, S., & Dogruel, L. (2012). A question of morality? The influence of moral salience and nationality on media preferences. *Communications*, 37(4), 345–369.
- Box Office Mojo. (2019). Avengers: Endgame. *Box Office Mojo*. Retrieved from https://www.boxofficemojo.com/title/tt4154796/?ref=bo_se_r_1
- Cron, L. (2012). *Wired for story: The writer's guide to using brain science to hook readers from the very first sentence*. Berkeley, CA: Ten Speed Press.
- Ding, L., & Yilmaz, A. (2010). Learning relations among movie characters: A social network perspective. In K. Daniilidis, P. Maragos, & N. Paragios (Eds.), *Computer vision: ECCV 2010* (pp. 410–423). Heidelberg: Springer. https://doi.org/10.1007/978-3-642-15561-1_30
- Eden, A., Daalmans, S., & Johnson, B. K. (2017). Morality predicts enjoyment but not appreciation of morally ambiguous characters. *Media Psychology*, 20(3), 349–373. <https://doi.org/10.1080/15213269.2016.1182030>
- Eliashberg, J., Elberse, A., & Leenders, M. A. (2006). The motion picture industry: Critical issues in practice, current research, and new research directions. *Marketing Science*, 25(6), 638–661. <https://doi.org/10.1287/mksc.1050.0177>
- Eliashberg, J., Hui, S. K., & Zhang, Z. J. (2007). From story line to box office: A new approach for green-lighting movie scripts. *Management Science*, 53(6), 881–893. <https://doi.org/10.1287/mnsc.1060.0668>
- Everett, J. A. C., & Crockett, M. J. (2019). What Game of Thrones reveals about moral decision-making. *Scientific American*. Retrieved from <https://blogs.scientificamerican.com/observations/what-game-of-thrones-reveals-about-moral-decision-making>
- Franzosi, R. (2010). *Quantitative narrative analysis*. Thousand Oaks, CA: Sage.
- Gantman, A. P., & Van Bavel, J. J. (2015). Moral perception. *Trends in Cognitive Sciences*, 19(11), 631–633. <https://doi.org/10.1016/j.tics.2015.08.004>
- Gleiser, P. M. (2007). How to become a superhero. *Journal of Statistical Mechanics: Theory and Experiment*, 2007(9), P09020. <https://doi.org/10.1088/1742-5468/2007/09/P09020>
- Gorinski, P., & Lapata, M. (2018). What's this movie about? A joint neural network architecture for movie content analysis. In M. Walker, H. Ji, & A. Stent (Eds.), *Proceedings of the 2018 conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, volume 1 (long papers)* (pp. 1770–1781). Stroudsburg, PA: Association for Computational Linguistics.
- Graham, J., Haidt, J., Koleva, S., Motyl, M., Iyer, R., Wojcik, S. P., & Ditto, P. H. (2013). Moral foundations theory: The pragmatic validity of moral pluralism. *Advances in Experimental Social Psychology*, 2013(47), 55–130. <https://doi.org/10.1016/B978-0-12-407236-7.00002-4>
- Grizzard, M., Lewis, R., Lee, A., & Eden, A. L. (2011). Predicting popularity of mass market films using the tenets of disposition theory. *International Journal of Arts and Technology*, 4(1), 48–60. <https://doi.org/10.1504/IJART.2011.037769>
- Haidt, J. (2001). The emotional dog and its rational tail: A social intuitionist approach to moral judgment. *Psychological Review*, 108(4), 814–834.
- Hinyard, L. J., & Kreuter, M. W. (2007). Using narrative communication as a tool for health behavior change: A conceptual, theoretical, and empirical overview. *Health Education & Behavior*, 34(5), 777–792. <https://doi.org/10.1177/1090198106291963>
- Hopp, F. R., Fisher, J., Cornell, D., Huskey, R., & Weber, R. (in press). The extended moral foundations dictionary (eMFD): Development and applications of a crowd-sourced approach to extracting moral intuitions from text. *Behavior Research Methods*. Advance online publication. Retrieved from <https://psyarxiv.com/924gq>
- Hopp, F. R., Fisher, J., & Weber, R. (2020). Dynamic transactions between news frames and sociopolitical events: An integrative, hidden Markov model approach. *Journal of Communication*, 70(3), 335–355. <https://doi.org/10.1093/joc/jqaa015>
- Knop-Huelss, K., Rieger, D., & Schneider, F. M. (2019). Thinking about right and wrong: Examining the ef-

- fect of moral conflict on entertainment experiences, and knowledge. *Media Psychology*. Advance online publication. <https://doi.org/10.1080/15213269.2019.1623697>
- László, J. (2008). *The science of stories: An introduction to narrative psychology*. Abingdon: Routledge.
- Lewis, R. J., Grizzard, M. N., Choi, J. A., & Wang, P. L. (2019). Are enjoyment and appreciation both yardsticks of popularity? *Journal of Media Psychology: Theories, Methods, and Applications*, 31(2), 55–64. <https://doi.org/10.1027/1864-1105/a000219>
- Lewis, R. J., Tamborini, R., & Weber, R. (2014). Testing a dual-process model of media enjoyment and appreciation. *Journal of Communication*, 64(3), 397–416. <https://doi.org/10.1111/jcom.12101>
- Lucas, G. (Director). (1977). *Star Wars: Episode IV: A new hope* [Screenplay]. USA: Lucasfilm and Twentieth Century Fox.
- Lynn, C. W., & Bassett, D. S. (2019). Graph learning: How humans infer and represent networks. *arXiv.org*. Retrieved from <https://arxiv.org/abs/1909.07186>
- Mac Carron, P., & Kenna, R. (2012). Universal properties of mythological networks. *EPL (Europhysics Letters)*, 99(2), 28002. <https://doi.org/10.1209/0295-5075/99/28002>
- Mastro, D., Enriquez, M., & Bowman, N. D. (2013). Morality subcultures and media production: How Hollywood minds the morals of its audience. In R. Tamborini (Ed.), *Media and the moral mind* (pp. 99–116). New York, NY: Routledge.
- McKee, R. (2005). *Story: Substance, structure, style, and the principles of screenwriting*. London: Methuen Publishing.
- Oliver, M. B. (1993). Exploring the paradox of the enjoyment of sad films. *Human Communication Research*, 19(3), 315–342. <https://doi.org/10.1111/j.1468-2958.1993.tb00304.x>
- Oliver, M. B., & Bartsch, A. (2011). Appreciation of entertainment. *Journal of Media Psychology*, 2011(23), 29–33. <https://doi.org/10.1027/1864-1105/a000029>
- Park, S. B., Oh, K. J., & Jo, G. S. (2012). Social network analysis in a movie using character-net. *Multimedia Tools and Applications*, 59(2), 601–627. <https://doi.org/10.1007/s11042-011-0725-1>
- Sagi, E., & Dehghani, M. (2014). Measuring moral rhetoric in text. *Social Science Computer Review*, 32(2), 132–144. <https://doi.org/10.1177/0894439313506837>
- Sinnott-Armstrong, W. (1988). *Moral dilemmas*. Oxford: Basil Blackwell.
- Skowron, M., Trapp, M., Payr, S., & Trappl, R. (2016). Automatic identification of character types from film dialogs. *Applied Artificial Intelligence*, 30(10), 942–973. <https://doi.org/10.1080/08839514.2017.1289311>
- Szilas, N., Estupiñán, S., & Richle, U. (2018). Automatic detection of conflicts in complex narrative structures. In R. Rouse, H. Koenitz, & M. Haahr (Eds.), *11th International Conference on Interactive Digital Storytelling, ICIDS 2018*. Cham: Springer.
- Tamborini, R. (2011). Moral intuition and media entertainment. *Journal of Media Psychology*, 23(1), 39–45.
- Tamborini, R. (2013). Model of intuitive morality and exemplars. In R. Tamborini (Ed.), *Media and the moral mind* (pp. 43–74). London: Routledge.
- Tamborini, R., Eden, A., Bowman, N. D., Grizzard, M., Weber, R., & Lewis, R. J. (2013). Predicting media appeal from instinctive moral values. *Mass Communication and Society*, 16(3), 325–346. <https://doi.org/10.1080/15205436.2012.703285>
- Tamborini, R., & Weber, R. (2020). Advancing the model of intuitive morality and exemplars. In K. Floyd & R. Weber (Eds.), *The handbook of communication science and biology* (pp. 456–469). New York, NY: Routledge.
- Tooby, J., & Cosmides, L. (2001). Does beauty build adapted minds? Toward an evolutionary theory of aesthetics, fiction, and the arts. *SubStance*, 30(1), 6–27. <https://doi.org/10.1353/sub.2001.0017>
- Tran, Q. D., & Jung, J. E. (2015). CoCharNet: Extracting social networks using character co-occurrence in movies. *Journal of Universal Computer Science*, 21(6), 796–815. <https://doi.org/10.3217/jucs-021-06-0796>
- Truby, J. (2007). *Anatomy of story*. New York, NY: Faber and Faber.
- Vorderer, P., Klimmt, C., & Ritterfeld, U. (2004). Enjoyment: At the heart of media entertainment. *Communication Theory*, 14(4), 388–408. <https://doi.org/10.1111/j.1468-2885.2004.tb00321.x>
- Wachowski, L., & Wachowski, L. (Directors). (1999). *The Matrix* [Screenplay]. USA: Warner Bros.
- Weber, R., & Hopp, F. R. (in press). Moral emotions and conflict motivate actions. *Insights: Consumer Neuroscience in Business*.
- Weber, R., Hopp, F. R., & Fisher, J. T. (2020). *Predicting movie performance from latent moral values in movie scripts*. Santa Barbara, CA: University of California.
- Weber, R., Mangus, J. M., Huskey, R., Hopp, F. R., Amir, O., Swanson, R., . . . Tamborini, R. (2018). Extracting latent moral information from text narratives: Relevance, challenges, and solutions. *Communication Methods and Measures*, 12(2/3), 119–139. <https://doi.org/10.1080/19312458.2018.1447656>
- Weber, R., Popova, L., & Mangus, J. M. (2013). Universal morality, mediated narratives, and neural synchrony. In R. Tamborini (Ed.), *Media and the moral mind* (pp. 50–66). New York, NY: Routledge.
- Weber, R., Tamborini, R., Lee, H. E., & Stipp, H. (2008). Soap opera exposure and enjoyment: A longitudinal test of disposition theory. *Media Psychology*, 11(4), 462–487. <https://doi.org/10.1080/15213260802509993>
- Weng, C. Y., Chu, W. T., & Wu, J. L. (2007). Movie analysis based on roles' social network. In X. Zhuang & W. Gao

(Eds.), *2007 IEEE International Conference on Multimedia and Expo* (pp. 1403–1406). Piscataway, NJ: IEEE. <https://doi.org/10.1109/ICME.2007.4284922>

Zillman, D., & Cantor, J. R. (1977). Affective responses

to the emotions of a protagonist. *Journal of Experimental Social Psychology*, 13(2), 155–165. [https://doi.org/10.1016/S0022-1031\(77\)80008-5](https://doi.org/10.1016/S0022-1031(77)80008-5)


About the Authors



Frederic René Hopp is a Doctoral Candidate in the Department of Communication at the University of California, Santa Barbara. His research advances theory and techniques for the algorithmic extraction of latent moral content and moral conflict from narratives, establishes computational models for predicting large-scale sociopolitical events and narrative performance, and illuminates the neural representations that undergird moral message processing. Frederic holds a BA in Media Psychology from the University of Mannheim and a MA in Communication from UCSB.



Jacob Taylor Fisher is a Doctoral Candidate in the Department of Communication at the University of California, Santa Barbara, and a Trainee of the National Science Foundation IGERT in Network Science and Big Data. He researches multimedia processing and media multitasking from a network neuroscience perspective. His current work investigates how certain digital environments can modulate attentional networks in the brain, and how these modulatory effects can be harnessed to develop novel treatments for cognitive processing disorders like ADHD.



René Weber received his PhD in Psychology from the University of Technology in Berlin, Germany, and his MD in Psychiatry and Cognitive Neuroscience from the RWTH University in Aachen, Germany. He is a Professor at UCSB's Department of Communication and director of UCSB's Media Neuroscience Lab (<https://medianeuroscience.org>). He was the first media psychology scholar to regularly use neuroimaging technology to investigate various media effects, from the impact of violence in video games to flow experiences, attention disorders, and the effectiveness of anti-drug PSAs. He has published four books and more than 140 journal articles and book chapters. His research has been supported by grants from national scientific foundations in the US and Germany, as well as through private philanthropies and industry contracts. He is a Fellow of the International Communication Association.