Computational Humanities Research

www.cambridge.org/chr

McK: Ontology for mapping semantic conflict in plot summary

Nicolas Chiappucci

Department of Philosophy, Alma Mater Studiorum University of Bologna, Italy

Research Article

Cite this article: Chiappucci Nicolas. 2026. "McK: Ontology for mapping semantic conflict in plot summary" *Computational Humanities Research*, **2**:e8, https://doi.org/10.1017/chr.2026.10028

Received: 23 September 2025
Revised: 19 December 2025
Accepted: 6 March 2026

Keywords:
conflict; ontology; screenwriting

Corresponding author:
Nicolas Chiappucci;
Email: nicolas.chiappucci2@unibo.it

Abstract

This article investigates the use of a neurosymbolic approach in the analysis of conflict within narrative discourse. Due to the difficulty of providing a precise conceptual definition, abstract elements such as conflict prove to be hard to analyze through computational systems. The retrieval of a neurosymbolic approach, through the combination of knowledge graphs and LLMs, can open new perspectives for the analysis of abstract elements typical of narrative discourse at the level of plot summary. Starting from the schematization of conflict elaborated by Robert McKee, an ontology was constructed and subsequently populated with the entities present in the screenplay of the film *A Clockwork Orange*. Afterward, through experimental validation, the approach was tested by means of a prompt injection of the ontology into the request. Through a comparative-qualitative approach, the experimentation considered the analysis of the narrative discourse first by combining synopsis (The use of the plot summary, instead of the screenplay, is connected to specific provisions concerning copyright law) and knowledge graph, and then relying only on the synopsis. In conclusion, the potentialities of the neurosymbolic approach are presented with regard to screenplay analysis, opening up the possibilities of the approach to a larger portion of text, such as the screenplay.

Plain language summary

This study investigates a possibility for analyzing conflict within narrative discourse. Due to the issues identified in the field of computational narratology and in the analysis of screenplays, computational systems encounter difficulties in the investigation of abstract elements such as conflict, which even humans struggle to define. Starting from a neurosymbolic approach, the intention is to enhance the process of conflict analysis within narrative discourse by means of a schematization of it. Indeed, in order to allow a better grasp of the screenplay, it is necessary to define detailed and dynamic models suitable for computational processing, so as to provide a schematization of conflict useful as an intermediate representation in the process.

The theoretical model, necessary for the schematization of conflict, is inspired by Robert McKee and his definition of conflict formulated in 1997, in the context of screenwriting. Consequently, the research offers a concrete possibility for analyzing narrative themes within film screenplays, enabling a better positioning already from the earliest stages of the cinematic production chain.

Introduction

In this context, the topic of interest is conflict. According to McKee (1997), during the act of writing the screenplay, narration originates from conflict. A story arises when a character does something expecting a certain reaction, but the world responds in an unexpected way. At the same time, for Snyder (2005), it is the catalyst that represents the moment in which, within the story, something happens for the first time, where the action of a character triggers an unforeseen reaction. Looking at the field of narratology, in Herman, Jahn, and Ryan (2005), it is argued that a minimal condition for narrative corresponds to the failure of intended actions caused by emergent events, which may originate from other characters' intentions but are not necessarily planned as such. Indeed, before contemporary approaches, Todorov argued that an ideal narrative begins with a stable situation that some force disrupts, resulting in a state of disequilibrium (1968). This formulation does not presuppose conflict as an organizing principle, but rather identifies a minimal structural condition for narrative emergence: the interruption of an expected state.

Even narrative traditions that do not privilege antagonism or goal opposition rely on these moments of deviation or reconfiguration that disrupt an expected progression. While in these cases, narrative development is not organized around a conflict in a dramatic sense, it nonetheless involves the structural introduction of disequilibrium. With the advent of computational narratology, deeply influenced by the functionalist frameworks developed by scholars,

This article was awarded Open Materials badge for transparent practices. See the Data availability statement for details.

© The Author(s), 2026. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

such as Propp, Bremond and Genette, conflict has effectively become a category of inquiry within textual analysis. Ware and Michael Young (2010), for example, present a computational model of narrative conflict based on AI planning, extending POCL planning to allow for the representation of thwarted or failed plans. Most recently, Yavuz et al. (2025) built an approach to analyze character relationships in Shakespeare's *Julius Caesar*, in a process integrating semantic, emotional and relational information derived from character dialogue. While these approaches offer rich and theoretically grounded representations of narrative conflict, their complexity and modeling requirements limit their practical adoption in applied contexts – such as workflows in the creative industries (Osborne et al. 2024) – particularly when dealing with large-scale corpora of narrative texts like novels or screenplays.

The contribution of this essay is threefold: 1) to offer a computational framework in support of a neurosymbolic approach to narratological understanding, exemplified through the construction of an ontology of conflict; 2) provide an experimental evaluation of the use of a knowledge graph, in interaction with an LLM, within the process of narratological analysis of conflict in narrative discourse; and 3) discuss the possible implementation and practical approaches of the same.

Theoretical motivation

A first point concerns the practical dimension. The automated and collaborative annotation of narrative texts, enriched with narratological metadata regarding aspects, such as temporality, spatiality, enunciative perspectives (Piper and Bagga 2024), plot structure (Lin et al. 2025), characterization (Stammbach, Antoniak, and Ash 2022) and emotional dynamics, represents a methodological support tool of primary importance. While many of the cited works refer to other forms of narrative discourse, the aim here is to transpose such approaches to the domain of film screenplays. Conceived as a collective enterprise, it not only makes it possible to produce narratological repertoires of unprecedented breadth and depth, but also renders the analysis of dramaturgical structures more transparent and systematic. In this perspective, the "crowdsourcing" (Meister 2012) of annotation is not to be understood as an end in itself, but rather as a twofold support device: on the one hand, it provides researchers and film scholars with more articulated and coherent analytical descriptions of screenplays; on the other hand, it enables advanced computational access to their narrative patterns, facilitating the identification of regularities, variations and formal innovations.

A second point regards the theoretical implications of relying on LLMs for narratological analysis. Currently, several studies (Piper and Bagga 2022; 2024; Jenner et al. 2025) highlight the potential of LLMs in the analysis of narrative discourse, often by redefining existing analytical standards and challenging previously established methodological frameworks. However, as shown by Murthy, Ullman, and Hu (2025), aligned models exhibit persistent difficulties in generating responses that display a level of conceptual diversity comparable to human reasoning. This limitation is particularly evident in widely adopted LLMs, whose alignment objectives necessarily favor predictable and socially recognizable outputs. At the same time, the autoregressive nature of these systems tends to amplify output inconsistencies (Vsevolodovna and Monti 2025). Consequently, simply increasing the amount of contextual information in the prompt or refining the user's request does not necessarily lead to more reliable or conceptually rich results, and may fail to resolve the underlying limitations of the task.

A third point regards conflict specifically. Thematically, engaging with conflict means dealing with a wide set of properties tied to its own dimension and structure, generally irreducible to a mere series or to a pure categorization. As Deutsch, Coleman, and Marcus (2006) observe, discussing conflict entails addressing a multiplicity of elements, including cooperation–competition, social justice, motivation, trust, communication, language and attribution processes.

In connection with the third point, difficulties arise in the conceptual definition of conflict. Although, starting from anthropological and strategic studies, it is possible to identify a sort of structural matrix of conflictual forms (Montanari 2004), it remains difficult for human annotators to account for what conflict actually is, as demonstrated by Piper and Bagga (2024). Consequently, since LLMs operate through the recognition and generalization of recurring patterns within training datasets, they tend to reproduce this very indeterminacy: for LLMs, there exists no stable and shared notion of "conflict" to rely on, but only variously plausible statistical configurations. As a result, it is unlikely to obtain perfectly concordant answers, precisely because the definition of conflict remains structurally polysemic and contextual. In what follows, we address: 1) previous work concerning narrative understanding and the use of neurosymbolic approaches in combination with knowledge graphs; 2) the theoretical framework, inspired by Robert McKee's theory of conflict, for constructing an ontology of conflict; and 3) testing how, and whether, the quality and coherence of insights on conflict in *A Clockwork Orange* change when an LLM operates:

- with a schematization of conflict (KG + plot summary);
- without schematization (plot summary only).

Prior work

The scientific literature specifically related to the thematic analysis of screenplays, with a view to fostering narrative depth during the early stages of the cinematic production chain, has focused on properties related to box office success derived from a series of textual features (Eliashberg, Hui, and John Zhang 2014; Hunter, Smith, and Singh 2016; Agah et al. 2025). Other works have addressed the possibility of identifying the local degree of narrativity in texts, focusing on specific textual passages (Piper and Bagga 2022; 2024; Steg, Slot, and Pianzola 2022). However, the challenges of the computational approach to narratology have revolved around a series of persistent issues: among these, the processes of cognitive economy in the recognition of narrative structures (Pronin 2013); prototypical elements of post-classical narratology that are difficult to translate into features, such as the Most Reportable Events (Labov and Waletzky 1967; Labov 1997), where an MRE is an unusual event that has the greatest emotional impact on the narrator and the audience and non-events (Passalacqua and Pianzola 2011). Indeed, Steg, Slot, and Pianzola (2022) analyze to what extent suspense, surprise and curiosity – considered as perceptual modes of the reader – can serve as indicators of narrativity, and they relate these forms of response to textual characteristics that narrative universals associate with storytelling. Consequently, the attempt to axiomatize conflict, from a computational perspective through a neurosymbolic approach in interaction with LLMs, would allow the semantic and multimodal capital of the text to be better valued. As shown by Baek, Aji, and Saffari (2023), the use of knowledge derived from a knowledege graph within the prompt makes it possible to mitigate LLM hallucinations, through a process of knowledge injection whereby meaningful triples are inserted into the prompt to support the

https://doi.org/10.1017/chr.2026.10028 Published online by Cambridge University Press

model in generating its response. However, while in this case, the task focuses on detecting the correct answer in response to a given query, KGs can also provide LLMs with greater interpretability, offering insights into their reasoning processes (Kau et al. 2024; Meyer et al. 2024; Vsevolodovna and Monti 2025). The creation of a knowledge graph can be understood as the construction of a directed, labeled graph in which nodes represent entities or concepts, and edges represent the relationships between them (Kau et al. 2024). Its purpose is to make explicit the assumptions about a given domain of knowledge (Noy and McGuinness 2001) through its schematization. The use of a neurosymbolic approach, in combination with knowledge graphs, aims to mitigate the well-recognized limitations of LLMs, particularly with respect to the lack of interpretability of the model. This has led, for example, Baek, Aji, and Saffari (2023) to introduce KAPING (Knowledge-Augmented PromptING), a method that enriches prompts for LLMs by drawing factual information from a knowledge graph. As noted by Kau et al. (2024), the present neurosymbolic approach seeks to balance the potential interpretability of the model by offering it an enhanced capacity for semantic understanding of narrative discourse. In the case of conflict, the construction of a knowledge graph would make it possible to overcome the problems of semantic understanding of the theme observed, for instance, in the experiments of Piper and Bagga (2024). A specific strength of LLMs lies in their linguistic specificity and their ability to capture patterns, relations and correlations. However, within the scope of computational narratology, such systems are hardly able to capture the semiotic nature of conflict as it pertains to collective, relational, perceptual or aesthetic dimensions. For an LLM to comprehend a conflictual structure, it must, in the act of analysis, assume a radically encyclopedic (Eco 1984) rather than merely vocabularistic dimension. Here, the encyclopedic dimension refers to a model of meaning in which signs are understood not through fixed and isolated definitions but through an open network of interpretants, where meaning emerges from continuous shifts and contextual relations. From this perspective, the use of ontologies provides a computational analog: not as static vocabularies, but as relational frameworks that organize meaning through interconnected concepts rather than rigid lexical correspondences, once placed in interaction with LLMs. In this sense, ontologies structure the space within which interpretation can be computationally guided (Vsevolodovna and Monti 2025). As shown by Umberto Eco, an encyclopedic model of meaning is inherently rhizomatic, characterized by relational connectivity, multiple access paths and the absence of a central organizing principle (Eco 1984). The construction of an ontology exhibits strong structural affinities with this model: beyond relationality and plurality of access, ontological schemas are irreducible to dictionary-based representations, as meaning emerges from the network of relations rather than from isolated definitions. Rather than fixing meaning, ontologies function as interpretative scaffolds, structurally compatible with the encyclopedic model as they organize access to meaning without collapsing into a closed set of semantic assignments. This holds provided that contextual modulation is introduced at the level of their interaction with LLMs, rather than being assumed as an intrinsic property of ontological formalisms. At this point, the question shifts from the structure of meaning to the conditions under which such heterogeneous relations can be integrated and rendered operational. If language functions as a unifying mode that transcends individual media, by virtue of its supramodal nature (Gangemi and Nuzzolese 2025), serving as a bridge between sensory data and symbolic representation, then the operation of LLMs, as shown elsewhere (Agah et al. 2025; Qin et al. 2024), appears optimal for integrating modal signs into a coherent and human-readable format. Yet, the semantic theory of the system remains bound to vector space and to its own distributional hypothesis. These agents exploit vast corpora and the supervised pre-training of language models to infer patterns and relations from raw text, enabling them to manage ambiguity, while lacking true semantic structuring – hence their tendency to be used as supportive tools. Nevertheless, the process of machinic perception remains akin to an informal encapsulation (Fodor 1983), lacking an attentional process capable of linking perceptual and cognitive processing at the computational level. This is where the intention arises to integrate a neurosymbolic approach – namely, techniques designed to enhance the analytical process of AI systems through a knowledge-based symbolic approach in relation to LLMs – thus enabling the neural network to reason more effectively about concepts and to produce explainable results (Sheth, Roy, and Gaur 2023).

Developing an ontology of conflict

To construct an ontology suitable for enhancing the interpretability of an LLM in the analysis of conflict, it was first necessary to identify a schematization of conflict consistent with the use of an ontology. As Lakoff and Narayanan (2010) note, an ontology capable of aligning with a cognitive model of narrative requires precise components: 1) they must be granular and detailed, so as to represent the full spectrum of possible events and interactions; 2) context-sensitive and supported by evidence, so as to adapt to dynamic and uncertain scenarios; 3) inspired by human cognitive processes, in order to make responses easily queryable and understandable; and 4) open to processing and extendable, allowing new domain models to specialize existing representations without altering their basic principles.

In his major work *Story: Substance, Structure, Style and the Principles of Screenwriting*, Robert McKee develops a schema for the formalization of conflict. According to the author, once a screenplay is written, the world of a character can be imagined as a series of concentric circles surrounding a core of identity – circles that mark the levels of conflict within a character. For McKee, the identity of a character can be conceived as a tripartition of concentric circles, with the innermost core defined as body, mind and emotions. From there, the concepts defining the first circle expand outward through increasingly external conflicts to the character's personal sphere, colliding with other circles. Starting from the center, the three forms of conflict that affect the character are defined as inner conflict, personal conflict and extra-personal conflict.

1) The first level of conflict is the one that arises within the character, in their interiority. It is a struggle involving mind, body and emotions and manifests itself as a tension between desires, fears, beliefs and personal limits. Here, the character is forced to confront their own contradictions, moral dilemmas and difficult choices that call their identity into question. It is the deepest and most intimate level, the one that allows the viewer to grasp the psychological complexity of the protagonist. Without inner conflict, characters risk appearing flat, since it is precisely the fracture within them that makes them alive and credible.

2) The second level refers to the direct relationships of the character with those closest to them: family, friends and romantic partners. At this level emotional tensions, rivalries, misunderstandings,

https://doi.org/10.1017/chr.2026.10028 Published online by Cambridge University Press

jealousies and betrayals are developed. Personal conflict represents the terrain of human relationships, the place where intimacy and affection can turn into sources of pain and confrontation. It is at this level that the most recognizable dynamics emerge for the viewer, since everyone can relate to a family quarrel or a couple's crisis. This conflict is more external than the inner one, yet it remains strongly tied to the private and emotional sphere of the character.

3) The third level of conflict concerns the character's relationship with broader and more impersonal external forces: society, institutions, the physical environment and historical or natural circumstances. It is the level at which the protagonist clashes with oppressive laws, social structures, injustices or events beyond their control, such as disasters or extreme environmental conditions. Here, the struggle is no longer only internal or relational, but becomes a battle with the external world and its constraints. This level of conflict often confers an epic dimension to the story, as the character must confront forces greater than themselves, putting their courage and resilience to the test.

The use of Robert McKee's conflict framework is motivated by its structurally explicit organization, which lends itself to formal modeling. As demonstrated in both semiotic and narratological theory, general narratological models (Federman 2016) are frequently employed even in the analysis of specific narrative forms, relying on the versatility of semio-narrative categories that can be readily translated into diagrammatic representations. Although McKee's framework was not originally developed as a theoretical model, its schematization of conflict dynamics proves particularly suitable for direct transposition into an ontological structure.

McK ontology structure

The McK ontology provides an event-centric formalization of dramatic conflict according to three levels (inner, personal and extra-personal), where a conflict is modeled as an event/situation entity involving one or more subjects and potentially originating in the internal components of the agent (body, emotions and mind) while opposing external forces. The ontology is expressed in OWL 2 DL.

Classes like "Body," "PhysicalEnvironment" and "SocialInstitutions" represent different groups of individual instances relevant to one or more conflictual events.

`mck:Conflict`: Root class for events/occasions of narrative opposition. It constitutes the natural domain of participation/impact relations and of internal etiological relations.

`mck:InnerConflict` $\sqsubseteq$ `mck:Conflict`: Specialization for intrapersonal conflicts; in the ontology, it is constrained by a minimum cardinality on an internal etiological link (see the Axioms section).

`mck:PersonalConflict` $\sqsubseteq$ `mck:Conflict`: Specialization for interpersonal conflicts (between characters or groups).

`mck:ExtraPersonalConflict` $\sqsubseteq$ `mck:Conflict`: Specialization for conflicts against external forces. The characteristic property is the link to the external object of opposition (see `mck:isAgainst`).

`mck:ExternalForce` $\equiv$ `mck:PhysicalEnvironment` $\sqcup$ `mck:SocialEnvironment` $\sqcup$ `mck:SocialInstitutions`: Abstract category that groups together forms of pressure or opposition coming from the outside, whose three subclasses are:

`mck:PhysicalEnvironment`
`mck:SocialEnvironment`
`mck:SocialInstitutions`

`mck:Origin` $\equiv$ `mck:Body` $\equiv$ `mck:Emotions` $\equiv$ `mck:Mind`: Abstraction of the internal sources of conflict, whose subclasses are:

`mck:Mind`
`mck:Body`
`mck:Emotions`

The *Object properties* of the McK ontology are intended to connect conflicts to the relevant elements that define and contextualize them.

`mck:affects`(inverse: `mck:isAffectedBy`): General relation of participation/impact: links an instance of `mck:Conflict` to the entities involved (characters, groups, institutions, etc.). The inverse `mck:isAffectedBy` is defined and is not typed with further domains beyond the inversion.

`mck:hasOriginIn`: Internal etiological relation: connects an individual of the class `mck:Conflict` to an individual belonging to `mck:Origin`. Three sub-properties are provided and typed according to level of detail:

`mck:hasOriginInBody` $\sqsubseteq$ `mck:hasOriginIn` (range: `mck:Body`)
`mck:hasOriginInEmotions` $\sqsubseteq$ `mck:hasOriginIn` (range: `mck:Emotions`)
`mck:hasOriginInMind` $\sqsubseteq$ `mck:hasOriginIn` (range: `mck:mind`)

`mck:isAgainst`: Characteristic relation of extra-personal conflicts: links an instance of `mck:ExtraPersonalConflict` to the `mck:ExternalForce` against which the opposition is configured.

The **axioms** of the McK ontology are intended to define the logical constraints that govern the relationships between classes and properties:

Union equivalences:

`mck:ExternalForce` $\equiv$ `mck:PhysicalEnvironment` $\sqcup$ `mck:SocialEnvironment` $\sqcup$ `mck:SocialInstitutions`
`mck:Origin` $\equiv$ `mck:Body` $\sqcup$ `mck:Emotions` $\sqcup$ `mck:Mind`

Structural disjointness:

`mck:ExtraPersonalConflict` $\bot$ `mck:InnerConflict`
`mck:ExtraPersonalConflict` $\bot$ `mck:PersonalConflict`
`mck:InnerConflict` $\bot$ `mck:PersonalConflict`

Cardinality constraints:

`mck:InnerConflict` $\sqsubseteq$ `mck:Conflict` $\sqcap$ (≥ 1 `mck:hasOriginIn`. (`mck:Body` $\sqcup$ `mck:Emotions` $\sqcup$ `mck:Mind`))

Inverse properties:

`mck:affects` $\equiv$ (`mck:isAffectedBy`)$^{-1}$

After defining the conceptual structure of McK, we proceeded with its population by employing the screenplay of the film *A Clockwork Orange* as the observation domain. The objective was to verify that the classes (`mck:Conflict` and its three specializations; `mck:Origin`; `mck:ExternalForce`) and the object properties (`mck:affects`/`mck:isAffectedBy`, `mck:hasOriginIn`$*$, `mck:isAgainst`) naturally supported the annotation of narrative conflicts in a concrete plot summary. In this phase, the modeling remained event-centric: each conflict instance denotes a situated occurrence, connected to the entities involved and, where relevant, to its internal origins or to the external opposing forces. The choice of a plot such as *A Clockwork Orange* was motivated by the depth (Eder 2008) of the conflictual structures embedded in the text, which preserves a fundamental

https://doi.org/10.1017/chr.2026.10028 Published online by Cambridge University Press

```
digraph g {
      "owl:Thing" -> "Mind" [label="has subclass"]
      "owl:Thing" -> "SocialInstitutions" [label="has subclass"]
      "Conflict" -> "InnerConflict" [label="has subclass"]
      "Conflict" -> "ExtraPersonalConflict" [label="has
subclass"]
      "Conflict" -> "PersonalConflict" [label="has subclass"]
      "owl:Thing" -> "Origin" [label="has subclass"]
      "owl:Thing" -> "PhysicalEnvironment" [label="has subclass"]
      "owl:Thing" -> "Emotions" [label="has subclass"]
      "owl:Thing" -> "Conflict" [label="has subclass"]
      "owl:Thing" -> "Body" [label="has subclass"]
      "Conflict" -> "Origin" [label="hasOriginIn (Domain>Range)"]
      "owl:Thing" -> "ExternalForce" [label="has subclass"]
      "owl:Thing" -> "SocialEnvironment" [label="has subclass"]
      "ExtraPersonalConflict" -> "ExternalForce"
[label="isAgainst (Domain>Range)"]
}
```

Figure 1. DOT specification of the OWL-based conflict ontology McK.

thematic irreducibility due to the presence of multiple narrative tension nodes throughout the work.[1]

Representative examples. Three instances illustrate the coverage of the schema; in addition, a concrete example of ontology population is provided to demonstrate how the schema is effectively instantiated in the knowledge graph (for the ontology DOT specification see Figure 1).

(i) An *Inner Conflict* such as the conditioned aversion to violence or to Beethoven's Ninth links the instance to `mck:Body`, `mck:Emotions` or `mck:Mind` through the sub-properties of `mck:hasOriginIn`, making the nature of the internal cause (somatic, emotional and cognitive) subject to querying.

```
mck:AlexsAversionToBeethoven
    rdf:type mck:InnerConflict ;
    mck:affects mck:Alex ;
    mck:hasOriginInBody mck:AlexsBody ;
    mck:hasOriginInEmotions
    mck:AlexsEmotions ;
    rdfs:label "Alex's love for Beethoven's
    9th versus his conditioned sickness" .
```

(ii) A *Personal Conflict*, such as the famous underpass assault, links the event to the main participants through `mck:affects`, without imposing directional roles: this choice allows querying both from the protagonist's perspective and from that of the opponent.

```
mck:AlexVsTramp_Underpass
    rdf:type mck:PersonalConflict ;
    mck:affects mck:Alex ,
                mck:Tramp ;
    rdfs:label "Alex and his droogs assault
    the tramp in the underpass" .
```

(iii) An *Extra-Personal Conflict* such as the "Ludovico Technique" procedure makes explicit Alex's opposition to an institutional `mck:ExternalForce` (`mck:TheState`) via `mck:isAgainst`, enabling local inferences consistent with domain and range.

```
mck:AlexVsState_LudovicoTechnique
    rdf:type mck:ExtraPersonalConflict ;
    mck:affects mck:Alex ;
    mck:isAgainst mck:TheState ;
    rdfs:label "Alex undergoes the Ludovico
    Technique" .
```

Evaluation

Formal evaluation

We verified the Mck Ontology with the HermiT 1.4.3.456 reasoner in Protégé 5.6.4. HermiT precomputed all inferences (class and property hierarchies, assertions and same individuals) in 300 ms, showing that the model is tractable for interactive work.

Experimental evaluation

We evaluated the use of McK ontology for conflict extraction starting from the synopsis of the film *A Clockwork Orange*. The experiment was structured in two conditions:

1. **Synopsis** + **KG**: In the first phase, we use the synopsis of the screenplay (available on Wikipedia.org) in combination with the knowledge graph, which is introduced through prompt injection.
2. **Synopsis only**: In the second phase, the input consists solely of the screenplay synopsis.

Our prompting framework consists of the following elements: role prompt, framing question, narrative feature and individual passage (synopsis). Here is the first implementation:

You are given an ontology of conflict populated with characters, events, and relationships, together with a synopsis of the film from which the ontology is derived. Your task is to analyze the deeper layers of meaning embedded in the narrative using the ontology as the guiding interpretive framework. The ontology defines the structure

[1] We acknowledge the potential limitations in the generalizability of the model due to the choice of a single case study. Future work will aim to extend this approach to corpora of previously unexamined plots and screenplays.

https://doi.org/10.1017/chr.2026.10028 Published online by Cambridge University Press

of entities, roles, and relations; the synopsis serves only as narrative evidence to be interpreted through this structure. Treat the ontology as the primary lens: all reasoning, categories, and analytical moves must be grounded in its relational framework. Use the synopsis only to illustrate, support, or instantiate the entities and conflicts already modeled in the ontology. If parts of the synopsis fall outside the ontology, flag them as out of scope rather than introducing external categories or knowledge. Your goal is to uncover what the ontology, when applied to the synopsis, implies but does not directly state: latent dynamics, hidden structures, and suggested meanings. Formulate inferences and hypotheses only if they can be expressed through the ontology's internal logic and supported by the synopsis. Interpret character development, symbolic gestures, and thematic elements by mapping them back to ontology entities and relations, focusing on what the interplay between ontology and synopsis structurally and narratively suggests.

Here is the synopsis: . . .

Here is the content of the ontology in Turtle format . . .

Here, the second implementation without the content of the knowledge graph:

You are given a summary of a film that includes characters, events, and relationships. Your task is to analyze the deeper layers of meaning embedded within the film's narrative structure. Examine the conflict dynamics between characters and factions. Use the internal logic of the story—its progression of events, character arcs, and symbolic choices—as a framework to infer hidden mechanisms at play. Your goal is to uncover what the summary implies but does not directly state: patterns of influence, unspoken hierarchies, ethical contradictions, cultural subtexts, and systemic tensions that define the world depicted in the film.

Here is the synopsis: . . .

The evaluation favors a qualitative-comparative approach, based on the comparison between the types of insights produced with and without the knowledge graph.

Implementation details

The experiments were carried out in Google Colab (Python 3 runtime, GPU accelerator enabled), with a Large Language Model (Claude-Sonnet-4) guided by structured prompts.

Replication

For each condition, five total runs were performed (with identical parameters).

Results

Condition 1: Synopsis + knowledge graph

In the analysis supported by the knowledge graph, for each execution, it can be observed that: 1a) Every portion of the conflict is typed, integrated with origins, roles and relations. Each episode is classified within McKee's triad (Personal, Extra-personal and Inner), while the story is conceived as passages and transformations across categories. Consequently, the knowledge graph makes it possible to capture the logical transitions of narrative transformation, starting from a grammar of conflict that conditions the interpretation of the story. Here, a series of responses from the model:

The Transformation Cascade
The ontology maps a clear progression: Alex begins dominating *Personal Conflicts*(`AlexVsDroogs_LeadershipChallenge`, `AlexVsGeorgieAndDim_Marina`), transitions to subjugation under **Extra-Personal Conflicts** (`AlexVsState_LudovicoTechnique`), and ultimately experiences resolution through *Inner Conflicts* that restore his original nature (`AlexsRecovery`). This structural pattern reveals that the narrative's central concern is not violence per se, but the *source of volition* – whether actions originate from authentic self determination or external manipulation.

The Stratification of Origins in Inner Conflict
The ontology's `hasOriginIn` properties reveal a crucial progression in Alex's Inner Conflicts: 1. `AlexsAversionToViolence`: originates in both Body and Mind (dual conditioning) 2. `AlexsAversionToBeethoven`: originates in Body and Emotions (aesthetic violation) 3. `AlexsSuicidalImpulse`: originates in Mind and Body (existential, physical crisis) 4. `AlexsRecovery`: originates only in Mind (pure psychological restoration) This progression suggests that **authentic selfhood**, as modeled by the ontology, requires the reduction of inner conflict to a single origin (`Mind`). The Ludovico technique fails precisely because it creates conflicts with multiple, contradictory origins that cannot be sustained.

It is possible to observe the model's tendency to read the story – and thus the individuals inhabiting our screenplay – through the lens of the synopsis, consistently aligning with McKee's triad. As a result, the progression is interpreted as a movement from personal conflicts to extra-personal ones, then to internal conflicts, accounting for the underlying dynamism. Similarly, in the second case, the system interprets the conflicts as belonging to different dimensions, or classes, that define Alex's character – namely, body, mind and emotions – yet its interpretation remains semi-typified by the classes established within the ontology.

2a) The conflict is managed in a systematic way, through the observation of cyclical reconfigurations. The conflict itself is read as a recurring pattern. The nodes are treated as transit points: if they have outputs toward new conflicts, the analysis tends to recognize a loop within the story. Consequently, even if some theses are not stated in the synopsis, they are deduced from the topology. Moreover, certain elements are inferred directly from the ontology: by virtue of class definitions, property restrictions and logical axioms, the system can derive implicit structures (e.g., conflict categorization, origin typology and role equivalences) that go beyond what is explicitly present in the text. However, the synopsis remains the reference point: the system's inferences are not detached from the narrative content but emerge as structural extensions of the synopsis itself, made possible by the ontological framework.

Here, a series of responses from the model:

The Origin Inversion Pattern
The ontology's `mck:hasOriginIn` properties reveal a crucial structural irony. Alex's **Inner Conflicts** originate in Body, Mind, and Emotions – but these origins are artificially imposed by external forces rather than naturally arising from his psychology. The `AlexsAversionToViolence` and `AlexsAversionToBeethoven` conflicts have origins in `AlexsBody` and `AlexsMind`, but the synopsis shows these bodily and mental responses were technologically implanted.
Ontological Implication: The framework exposes how institutional control operates by colonizing the very categories (Body,

https://doi.org/10.1017/chr.2026.10028 Published online by Cambridge University Press

Mind, Emotions) that should generate authentic inner conflict, creating a simulated interiority.

The Recursive Recognition Pattern
Multiple conflicts involve recognition scenarios: the tramp recognizes Alex (`AlexVsTramp_Revenge`), Mr. Alexander recognizes Alex through "Singin' in the Rain" (`AlexVsMrAlexander_Revenge`). The ontology's affect relations show these recognition moments convert past Personal Conflicts into present Extra-Personal ones, as victims align with broader institutional forces.
Ontological Revelation: Recognition functions as the mechanism by which Personal Conflicts retrospectively become Extra-Personal ones, collapsing the temporal boundaries between conflict types.

In the first case, the model uncovers what it designates as an Origin Inversion Pattern: whereas in a conventional narrative a character's inner conflicts emerge from within – from body, mind and emotions – here the narrative develops differently: conflicts first arise at the personal level, then expand into extra-personal dimensions, and ultimately return as internal conflicts. Indeed, conflicts are artificially induced by technology and conditioning. The aversion to violence and to Beethoven's music, for instance, are not spontaneous but implanted through an institutional process. Thus, the fundamental categories of body, mind and emotions, instead of serving as natural sources of conflict, are transformed into instruments of control, constructing what the system describes as a simulated interiority – whether one agrees with this reading or not. In the second case, a Recursive Recognition Pattern appears, meaning that the ontology reveals how certain past personal conflicts resurface later in the story in another form. In the case of the tramp's assault, personal conflicts are later "recycled" as extra-personal ones, showing how the ontology frames their evolution beyond the initial level.

Condition 2: Synopsis only

Looking at the analysis carried out without the use of the knowledge graph, different processes are highlighted: 1b) There is a strong interpretative hypothesis in the analysis process (power, authoritarianism and consensus), to which scenes are linked as "evidence," using a consistent theoretical lexicon (e.g., totalitarianism, technocracy and consensus):

The Paradox of State Violence
The film exposes a fundamental contradiction: the state condemns Alex's violence while systematically employing far more sophisticated brutality. The Ludovico technique represents institutional sadism—torture disguised as rehabilitation. The government doesn't eliminate violence; it monopolizes it, creating a sanitized brutality that serves political ends. Alex's crude "ultra-violence" is merely the unrefined version of what the state practices with scientific precision.

The Cycle of Retribution as Social Order
The narrative suggests that this society has replaced justice with an endless cycle of revenge. The vagrant's attack on Alex, Dim and Georgie's police brutality, and Alexander's manipulation all represent not justice but the perpetuation of violence through different institutional channels. The system doesn't break cycles of violence—it legitimizes and systematizes them.

2b) The analysis is directed toward the refunctionalization of events (Beethoven, "beauty"), with a tendency toward the "armed" structuring of the event starting from the discourse of power.

The State as Ultimate Predator
The Ludovico technique represents the state's evolution from punishing crime to manufacturing compliant citizens. The Minister's casual mention of needing "more space for political prisoners" unveils the true hierarchy of threats: common criminals like Alex are merely nuisances, while political dissidents represent existential challenges to power. The technique doesn't eliminate violence—it centralizes it within state institutions.

The Cyclical Nature of Institutional Power
The narrative reveals a closed system where former criminals become enforcers (Dim and Georgie as police), victims become perpetrators (the vagrant's revenge), and reformers become oppressors (Alexander's manipulation). This suggests that the society's fundamental problem isn't individual pathology but systemic corruption that transforms everyone who participates in it.

Discussion

According to the classical theory of narration, a dramatic situation is defined by systems of forces endured or enacted by the characters inhabiting the situation itself (Souriau 1950). At the same time, for classical semiotic theory, in order to define the narrative value of discourse, it is necessary to reduce the relevance of syntactic-semantic structures within a scene in favor of the definition of semio-narrative structures defined by actants and action modules. This is the case with the construction of the ontology which, although presenting issues of univocity (that will be discussed shortly), allows for a medial representation of conflict by anchoring it to a form of regularity. If the semio-narrative level must be made relevant as a collective phenomenon (Marrone 2001), then the univocal construction of a conflict ontology can be considered valid, aiming at a structured representation of conflict and stability in its interpretation. From here derives the idea of combining structured symbolic knowledge with neural pattern recognition systems (Sheth, Roy, and Gaur 2023). The results obtained allow us to reflect on the value of narration and the schematization of the narrative path: the first experiment, as is common practice, tends to search for patterns and repetitions within the synopsis starting from the associative constraints present in the componential model, using the synopsis itself as a proxy prior to reasoning. What has been privileged is the conceptual organization, leading the responses to reflect a better cognitive dimension (Violi 2001) of how meaning is conceived. At the same time, the cognitive dimension is reflected in the model's ability to make inferences by using apparently different elements, identifying homologies between different phenomena; ontological properties (like `hasOriginIn`) make it possible to treat phenomena as sets of related components rather than as unitary blocks. This allows the model to recognize when an entity or event is the result of multiple simultaneous factors. The knowledge graph functions as a system of explicit formal semantics, reducing ambiguities and guiding the model's inference.

In Condition 2, attention is distributed across two predominant themes – violence and politics – starting from the constant identification of violence as the foundation of power and at the root of the corruption of institutions. The system tends to draw on value categories typically associated with the concept of conflict, such as

https://doi.org/10.1017/chr.2026.10028 Published online by Cambridge University Press

free will or the legitimacy of power. This recurrent association of conflict with stereotypical markers, such as violence and power, is related to the model's pretraining regime. As a result, the model's interpretive scope becomes restricted, with conflict recurrently instantiated through entrenched frames of violence and authority. By contrast, Condition 1 does not seem to privilege a specific form of conflict, but rather distributes attention across the interaction between individuals and the metamorphosis of conflicts. Here, the knowledge graph drives the system to see forms of conflict as co-dependent and convertible. In this sense, it pushes toward concepts made operational by the grid itself: classes, individuals and properties interpolate the contents generated by the LLM into a semi-typed system, starting from explicit constraints that enable stable inferences. As noted by Gruber (1993), an ontology is nothing more than the explicit specification of a conceptualization subject to formal rigor. Since, in the field of computer science and knowledge-based systems, what "exists" effectively corresponds to what can be represented, the process of ontological design does not merely describe a domain. Rather, it establishes constraints on what phenomena can be rendered computationally intelligible, thereby, in this context, guiding inference and reducing perspectival ambiguity. In the specific case of adopting McKee's perspective, the structural solidity of the model should not be interpreted as the formalization of a narrative theory. McKee's approach is grounded in the aggregation of a set of recurring macro-frames within a schematic space (as Friends, Individuals in Society and Social Institutions) which have been made explicit at the ontological level. Accordingly, the model's structural robustness does not stem from a fully articulated narrative theory, but from the operationalization of a series of macro-frames that enable the system to identify patterns by referring to configurations of states of affairs within the plot. This process does not rely on positional structures grounded in a relativistic logic, as would be the case with the adoption of a formal narrative theory such as Greimas's canonical generative trajectory (1970).

More broadly, these systemic dynamics should be understood as effects of the theoretical framework encoded in the ontology. The forms of conflict, inspired by Mckee's configuration, are not merely discovered by the system, but emerge from the ontological distinctions, constrains and relations that structure the ontology. The validity and interpretability of the analysis are intrinsically dependent on the validity of the theory, here formalized at the ontological level.

The conflict system inspired by McKee's schema, when formalized within an ontology, does not, however, explicitly account for the participating components of the nature of conflict, tending to neglect details and privileging the detection of macro-structural patterns and systemic dynamics. This derives both from the nature of the knowledge graph, which represents entities and relations on an eventual scale, and from the type of query that was formulated, oriented toward capturing general configurations or structures rather than micro-details. The modeling of more fine-grained dynamics and the inclusion of a specific conflictual scene could make it possible to perform the opposite process with respect to what was tested. Moreover, it is possible that conflict, by definition, resists precise definability due to its semantic nature. In the structuring of a narrative text such as a screenplay, however, conflict tends to be anchored to well-defined structures. At the same time, axiomatizing conflict risks producing an overly rigid theory, particularly in cases where screenplays or filmic works do not adhere to well-established canons. In conclusion, our findings should be interpreted as conditional on a specific model-prompt configuration. Accordingly, we do not claim that the produced annotations are model-invariant. The contribution of this article lies in the proposed ontological scaffolding and the possible integration in a workflow, whose transferability should be tested across models and more diverse corpora in future.

Prospects for future research

The use of such a neurosymbolic approach, although at this stage limited to the analysis of a short narrative discourse such as the synopsis, proves to be scalable and applicable to a more complex and larger discourse, namely, the screenplay. Indeed, the human reader is required to grasp subtexts and multimodal representations in the act of screenplay coverage (Maras 2009), exercising a refined interpretative competence which, however, necessarily remains partial, since the screenplay at this stage carries with it an inevitably unfinished character. The use of LLMs equipped with a context length sufficient to encompass the entire screenplay opens up the possibility of addressing this very incompleteness. Furthermore, the possibility of employing, as in this case, precise thematic knowledge graphs would enable the model to capture complex articulations starting from the desired narrative feature used to construct its ontology. A neurosymbolic approach of this kind could also prove particularly effective in capturing cultural and formal differences: the ability to anchor semantic representation to knowledge graphs would indeed allow the modelling of the interpretative process of a given theme within the screenplay, while respecting narrative sensibilities other than the Western ones. In addition, such an approach could be integrated as a supportive and easily implementable tool in situations where development processes are lacking or inadequate – often due to budget constraints or tight deadlines – by intervening in the act of coverage and ensuring a minimum level of structured analysis and interpretative coherence even under limiting production conditions. The adoption of an LLM-based infrastructures can reduce costs compared to traditional forms of script reading, coverage and manual analysis, as it enables the rapid scaling of repetitive tasks – such as pattern extraction, summarization and annotation – while maintaining a controllable marginal cost per text through token constraints and hybrid pipelines. In this sense, data-driven innovation can support more sustainable and accessible practices, in line with Osborne et al. (2024) on changing patterns of consumption and ownership, tools for fostering sustainability awareness and platforms that empower circular economy communities. Further exploration of these models and their potential is therefore necessary in order to demonstrate and refine their effective integration within such workflows.

Conclusion

In this essay, we presented a neurosymbolic approach to the analysis of narrative discourse. Starting from the framework provided by McKee (1997), an ontology was constructed to capture the narrative levels of conflict, based on the components identified by Lakoff and Narayanan (2010) for the development of a computational-cognitive model of narration. The ontology was then populated with individuals and events drawn from the screenplay of the film *A Clockwork Orange*. Subsequently, we tested a neurosymbolic approach by combining the use of a knowledge graph, introduced through prompt injection, with an LLM in order to assess the potential increase in the interpretability of the model, starting from

https://doi.org/10.1017/chr.2026.10028 Published online by Cambridge University Press

the analysis of an abstract and theoretically elusive topic such as conflict, which remains difficult to identify and operationalize even within established narratological frameworks. The results showed that the use of an ontology enables the modeling of the inferential process, allowing the model to stabilize interpretation on the basis of the ontology's categories. Future work will focus on a series of identified issues: using a longer narrative text, such as a full screenplay, while taking into account the copyright concerns encountered in this experiment, thereby testing the approach on a more substantial textual structure. Furthermore, we plan to introduce a phase of validation through human annotation, with the aim of constructing a reference set against which to carry out a systematic comparative evaluation, based on standard quantitative metrics.

Acknowledgements. I would like to thank Anna Sofia Lippolis, who played an important role in the development of this project, Prof. Cristina Demaria for her unconditional trust, Prof. Aldo Gangemi for his valuable advice and for inspiring this work and Prof. Paolo Martinelli for the insightful discussions. Many thanks to I Wonder Pictures for funding this work.

Data availability statement. Ontology, knowledge graph and transcription of the executions are available at https://github.com/ChiappucciN/MCK.

Disclosure of use of AI tools. Artificial intelligence tools were used as part of the methodological framework for exploratory analysis. The final interpretation, design of the study and conclusions remain entirely the author's responsibility.

Author contributions. Conceptualization: N.C.; Data curation: N.C.; Data visualization: N.C.; Methodology: N.C.; Writing – original draft: N.C. The author approved the final submitted draft.

Funding statement. My research work was supported by Top Contents, the research department of the Italian film distribution company I Wonder Pictures.

Competing interests. The author declares none.

Ethical standards. The research meets all ethical guidelines, including adherence to the legal requirements of the study country.

References

Agah, Shaghayegh, Yejin Kim, Neeraj Sharma, Mayur Nankani, Kevin Foley, H. Howie Huang, and Sardar Hamidian. 2025. "Predicting Movie Hits Before They Happen with LLMs." In *UMAP '25: Proceedings of the 33rd ACM Conference on User Modeling, Adaptation and Personalization*, 370–373. New York, NY: Association for Computing Machinery. https://doi.org/10.1145/3699682.3730972.

Baek, Jinheon, Alham Fikri Aji, and Amir Saffari. 2023. "Knowledge-Augmented Language Model Prompting for Zero-Shot Knowledge Graph Question Answering." In *Proceedings of the First Workshop on Matching From Unstructured and Structured Data (MATCHING 2023), 70–98*. Toronto, ON: Association for Computational Linguistics. https://doi.org/10.18653/v1/2023.matching-1.7.

Deutsch, Morton, Peter T. Coleman, and Eric C. Marcus. 2006. *The Handbook of Conflict Resolution: Theory and Practice*, 2nd ed. San Francisco, CA: Wiley.

Eco, Umberto. 1984. *Semiotics and the Philosophy of Language*. Bloomington, IN: Indiana University Press.

Eder, Jens. 2008. "Feelings in Conflict: "A Clockwork Orange" and the Explanation of Audiovisual Emotions." *Projections: The Journal for Movies and Mind* 2, no. 2: 66–84.

Eliashberg, Jehoshua, Sam K. Hui, and Z. John Zhang. 2014. "Assessing Box Office Performance Using Movie Scripts: A Kernel-Based Approach." *IEEE Transactions on Knowledge and Data Engineering* 26, no. 11: 2639–48. https://doi.org/10.1109/TKDE.2013.131.

Federman, Sarah. 2016. "Narrative Approaches to Understanding and Responding to Conflict." *International Journal of Conflict Engagement and Resolution* 4, no. 2: 154–71.

Fodor, Jerry A. 1983. *The Modularity of Mind*. Cambridge, MA: MIT Press.

Gangemi, Aldo, and Andrea Giovanni Nuzzolese. 2025. "Logic Augmented Generation." *Journal of Web Semantics* 85. https://doi.org/10.1016/j.websem.2025.100859.

Greimas, Algirdas Julien. 1970. *Du Sens I: Essais sémiotiques*. Paris: Seuil.

Gruber, Thomas R. 1993. "Toward Principles for the Design of Ontologies Used for Knowledge Sharing." In *International Journal of Human-Computer Studies* 43, edited by **B. R. Addis**, 907–928. Revised version of paper presented at the International Workshop on Formal Ontology, March 1993, Padova, Italy. Technical Report KSL 93-04, Stanford Knowledge Systems Laboratory.

Jenner, S., D. Raidos, E. Anderson, S. Fleetwood, B. Ainsworth, K. Fox, J. Kreppner, and M. Barker. 2025. "Using Large Language Models for Narrative Analysis: A Novel Application of Generative AI." *Methods in Psychology* 12: 100183. https://doi.org/10.1016/j.metip.2025.100183.

Herman, David, Manfred Jahn, and Marie-Laure Ryan. 2005. *Narrative Theory: Core Concepts and Critical Debates*. London: Routledge.

Hunter, Starling, Susan Smith, and Saba Singh. 2016. "Predicting Box Office from the Screenplay: A Text Analytical Approach." *Journal of Screenwriting* 7, no. 2: 135–54. https://doi.org/10.1386/josc.7.2.135_1.

Kau, Amanda, Xuzeng He, Aishwarya Nambissan, Aland Astudillo, Hui Yin, and Amir Aryani. 2024. "Combining Knowledge Graphs and Large Language Models." Preprint, arXiv:2407.06564. https://arxiv.org/abs/2407.06564.

Labov, William. 1997. "Some Further Steps in Narrative Analysis." *Journal of Narrative and Life History* 7: 395–415.

Labov, William, and Joshua Waletzky. 1967. "Narrative Analysis: Oral Versions of Personal Experience." In *Essays on the Verbal and Visual Arts*, edited by **June Helm**, 12–44. Seattle, WA: University of Washington Press.

Lakoff, George, and Srini Narayanan. 2010. "Toward a Computational Model of Narrative." In *Computational Models of Narrative: Papers from the AAAI Fall Symposium (FS-10-04), 21–28*. Arlington, VA: Association for the Advancement of Artificial Intelligence.

Lin, Zefeng, Yi Xiao, Zhiqiang Mo, Qifan Zhang, Jie Wang, Jiayang Chen, Jiajing Zhang, Hui Zhang, Zhengyi Liu, Xianyong Fang, and Xiaohua Xu. 2025. "R^2: A LLM Based Novel-to-Screenplay Generation Framework with Causal Plot Graphs." Preprint, arXiv:2503.15655. https://arxiv.org/abs/2503.15655.

Maras, Steven. 2009. *Screenwriting: History, Theory and Practice*. New York: Wallflower.

Marrone, Gianfranco. 2001. *Corpi Sociali. Processi Comunicativi e Semiotica del Testo*. Torino: Einaudi.

McKee, Robert. 1997. *Story: Substance, Structure, Style, and the Principles of Screenwriting*. New York: ReganBooks/HarperCollins.

Meister, Jan Christoph. 2012. "Crowd Sourcing "True Meaning": A Collaborative Markup Approach to Textual Interpretation." In *Collaborative Research in the Digital Humanities*, edited by **Willard McCarty and Marilyn Deegan**, 105–122. Farnham: Ashgate.

Meyer, L.-P., C. Stadler, J. Frey, N. Radtke, K. Junghanns, R. Meissner, G. Dziwis, K. Bulert, and M. Martin. 2024. "LLM-Assisted Knowledge Graph Engineering: Experiments with ChatGPT." In *First Working Conference on Artificial Intelligence Development for a Resilient and Sustainable Tomorrow*, 103–115. Wiesbaden: Springer Fachmedien. https://doi.org/10.1007/978-3-658-43705-3_8.

Montanari, Massimo. 2004. *Linguaggi Della Guerra*. Roma: Meltemi.

Murthy, Sonia K., Tomer Ullman, and Jennifer Hu. 2025. "*One Fish, Two Fish, But Not the Whole Sea: Alignment Reduces Language Models Conceptual Diversity*." In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 11241–11258. Albuquerque, NM: Association for Computational Linguistics. https://doi.org/10.18653/v1/2025.naacl-long.561.

Noy, Natalya F., and Deborah L. McGuinness. 2001. "Ontology Development 101: A Guide to Creating your First Ontology." Technical report KSL-01-05, Stanford Knowledge Systems Laboratory.

Osborne, Nicola, Ingi Helgason, Susan Lechelt, Lucia Michielin, Inge Panneels, Caroline Parkinson, Michael Smyth, Jen Ross, Yasmin Sulaiman,

and Katherine Warren. 2024. "Digital and Data Literacy: Models for Data Training and Upskilling for the Future Creative Industries." In *Data-Driven Innovation in the Creative Industries*, edited by **Melissa Terras, Vikki Jones, Nicola Osborne, and Chris Speed**, 1st ed., 1–30. London: Routledge.

Passalacqua, Franco, and Federico Pianzola. 2011. "Continuity and Break Points: Some Aspects of the Contemporary Debate in Narrative Theory." *Enthymema* no. 4: 9–34.

Piper, Andrew, and Sunyam Bagga. 2022. "Toward a Data-Driven Theory of Narrativity." *New Literary History* 53, no. 4: 879–901.

Piper, Andrew, and Sunyam Bagga. 2024. "Using Large Language Models for Understanding Narrative Discourse." In *Proceedings of the 6th Workshop on Narrative Understanding*, edited by **Yash Kumar Lal, Elizabeth Clark, Mohit Iyyer, Snigdha Chaturvedi, Anneliese Brei, Faeze Brahman, and Khyathi Raghavi Chandu**, 37–46. Mexico City, Mexico: Association for Computational Linguistics. https://doi.org/10.18653/v1/2024.wnu-1.4.

Pronin, Emily. 2013. "When the Mind Races: Effects of Thought Speed on Feeling and Action." *Current Directions in Psychological Science* 22, no. 4: 283–8.

Qin, Zhen, Rolf Jagerman, Kai Hui, Honglei Zhuang, Junru Wu, Le Yan, Jiaming Shen, Tianqi Liu, Jialu Liu, Donald Metzler, Xuanhui Wang, and Michael Bendersky. 2024. "Large Language Models are Effective Text Rankers with Pairwise Ranking Prompting." In *Findings of the Association for Computational Linguistics: NAACL 2024*, 1504–1518. Mexico City, Mexico: Association for Computational Linguistics. https://doi.org/10.18653/v1/2024.findings-naacl.97.

Sheth, Amit, Kaushik Roy, and Manas Gaur. 2023. "Neurosymbolic AI—Why, What, and How." Preprint, arXiv:2305.00813. https://arxiv.org/abs/2305.00813.

Snyder, Blake. 2005. *Save the Cat! The Last Book on Screenwriting You'll Ever Need*. Studio City, CA: Michael Wiese Productions.

Souriau, Étienne. 1950. *Les Deux Cent Mille Situations Dramatiques*. Paris: Flammarion.

Stammbach, Dominik, Maria Antoniak, and Elliott Ash. 2022. "Heroes, Villains, and Victims, and Gpt-3: Automated Extraction of Character Roles without Training Data." In *Proceedings of the 4th Workshop on Narrative Understanding (WNU2022)*, edited by **Elizabeth Clark, Faeze Brahman, Mohit Iyyer**, 47–56. Seattle, WA: Association for Computational Linguistics.

Steg, Max, Karlo Slot, and Federico Pianzola. 2022. "Computational Detection of Narrativity: A Comparison Using Textual Features and Reader Response." In *Proceedings of the 6th Latech-CLFL Workshop*, edited by **Stefania Degaetano, Anna Kazantseva, Nils Reiter, Stan Szpakowicz**, 105–14. Gyeongju, Republic of Korea: International Conference on Computational Linguistics.

Todorov, Tzvetan. 1968. "La Grammaire du Récit." *Langages 3* no. 12: 94–102.

Violi, Patrizia. 2001. *Meaning and Experience*. Bloomington, IN: Indiana University Press.

Vsevolodovna, Ruslan Idelfonso Magaña, and Marco Monti. 2025. "Enhancing Large Language Models Through Neuro-Symbolic Integration and Ontological Reasoning." Preprint, arXiv:2504.07640. https://arxiv.org/abs/2504.07640.

Ware, Stephen G., and R. Michael Young. 2010. "Modeling Narrative Conflict to Generate Interesting Stories." In *Proceedings of the Sixth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 210–215. Palo Alto, CA: Association for the Advancement of Artificial Intelligence.

Yavuz, Mehmet Can, Lucia Cascone, Aylin Özkan, and Irem Ertaş. 2025. "Narrative Conflicts: A Tri-Modal Computational Analysis of Antagonism in Shakespeares Julius Caesar." In *Proceedings of the Eleventh Italian Conference on Computational Linguistics (CLiC-it 2025)*, edited by **Cristina Bosco, Elisabetta Jezek, Marco Polignano, Manuela Sanguinetti**, 849–859. Cagliari, Italy: CEUR Workshop Proceedings.

https://doi.org/10.1017/chr.2026.10028 Published online by Cambridge University Press