

Editorial Aspects of Reporting into Structured Narratives

David Caswell
Reynolds Journalism Institute
Columbia, MO
david@
structuredstories.com

Frank Russell
University of Missouri
Columbia, MO
frank.russell@
mail.missouri.edu

Bill Adair
Duke University
Durham, NC
bill.adair@duke.edu

ABSTRACT

An experiment is described in which reporters captured real-world news stories as semantically structured representations of journalistic events and narratives, and preliminary results from the experiment are discussed.

The experiment was conducted using the Structured Stories prototype platform, which provides a library of abstractions of journalistic events grounded in frame semantics. Real-world journalistic events are then instantiated from those abstractions by filling semantic roles with identifiers to various Linked Open Data. Various media elements are attached to each instantiated event to facilitate presentation within a structured discourse. The resulting structured events are then assembled into larger narrative structures with recursive, fractal and network characteristics that seek to correspond with the cognitive experience of narrative in humans.

The goal of the experiment was to evaluate the editorial feasibility of using this approach to capture and represent real-world news stories within a realistic journalistic workflow.

Categories and Subject Descriptors

H.1.2 [User/Machine Systems]: Human factors; H.5.4 [Hypertext/Hypermedia]: User issues; I.2.4 [Knowledge Representation Formalisms and Methods]: Frames and scripts

General Terms

Experimentation, Human Factors

Keywords

Computational Narrative, Structured Journalism, Event Representation, Narrative Representation, Structured Events, Structured Narratives

1. INTRODUCTION

Structured Journalism is a new form of journalism based on reporting news as structured components into a database, and subsequent retrieval of those structured components to generate news products. The approach is still nascent but it directly addresses several systemic problems facing news producers and news consumers in the digital media ecosystem, and it may potentially facilitate the rebundling of journalism as networks and the creation of consumer-controlled news products with context that extends beyond the article. This potential has caused growing interest in Structured Journalism, including experiments that are currently ongoing at the BBC, the New York Times, and other media companies [8] [3].

Pioneering examples of Structured Journalism, such as PoliFact and D.C. Homicide Watch, have been domain-specific and have used fixed database schemas to represent the news components of their chosen domains, often as small text artifacts. Other early examples, such as Google's Living Stories and more recently the Guardian's Live Pages, have gathered various news components and meta-data around text articles, using both automated tagging and manual creation of components by reporters or editors. In contrast to both of these earlier approaches an alternative approach, called Structured Stories, attempts to provide a method of representing ad hoc news events and stories entirely as structure, without reliance on text artifacts except as adjunct and optional descriptors attached to structured events.

The basic units of Structured Stories are event frames (Figure 1), which are abstractions of discrete journalistic events. Each event frame describes the general activity of a 'type' of event, centered on a verb, and includes typed semantic roles that each play a specific part in that activity. The definition of each event frame is grounded in references to FrameNet - a formal library of meaningful situations encoded as 'semantic frames' and developed by the International Computer Science Institute at U.C. Berkeley [1]. FrameNet provides the formal semantic basis for the predicate portion of event frames within Structured Stories, however each event frame also contains a small degree of editorial flexibility that enables different event frames to express slightly different 'types' of events while referencing the same FrameNet frame and the same semantic roles.

The representation of a real-world journalistic event within Structured Stories is achieved by instantiating the event




Figure 1: A simplified event frame, referenced to the ‘Publishing’ FrameNet frame.

frame as a specific event record (Figure 2) and by populating the semantic roles within the event frame with unique identifiers specifying the noun participants (or ‘arguments’) in the represented event. These noun participants are represented by Linked Open Data, typically nodes in various knowledge graphs, including Freebase, WikiData, GeoNames and Facebook, and are organized into seven top-level types - characters, entities/concepts, locations, information artifacts, other structured events, other structured stories and constants. Each instantiated event is assigned a location at an appropriate scale, using a GeoNames location identifier, and a package of temporal information that includes the temporal granularity of the event ranging from minutes to years. The journalistic event is therefore unambiguously structured such that its core meaning can be recovered, and the event can therefore be uniquely identified.

This approach to the representation of journalistic events is based on the view that such events are discrete, quantifiable and uniquely identifiable ‘things in the world’ that originate as instantiations of abstractions of forms of action. This view is broadly consistent with research in the cognitive science of event perception [10] and also with a long history of work in the semantics of events [7] [9].

A simple database of news events, whether represented as structured records or otherwise, would not enable any method of contextually organizing those events other than as temporally ordered lists. This is a significant problem, because such a database could therefore not represent sprawling and interrelated collections of diverse events in a coherent manner that supports context, provides efficient access to detail, communicates values and builds understanding. The concept of ‘story’, or ‘narrative’, is a representation scheme that organizes events in a way that provides such coherence, and is a central cognitive function in humans [11] [4].

Structured Stories provides a narrative representation scheme (Figure 3) for organizing structured events that attempts to replicate three aspects of cognitive narrative function - semantic zoom, differential value and narrative network. Semantic zoom is an ability to efficiently ‘drill down’ into more detail for a specific event, and is provided in Structured Stories using a recursive representation in which events within stories can be linked to another story containing a detailed narrative of that event. Differential value is the ability to assign different values to the representation of different events within a particular story, as opposed to assigning absolute values to events across all stories. Multiple values can be assigned to events within stories, enabling the same story to be simultaneously represented using different value schemes. The narrative network function arises from the existence of multiple connections between events, including

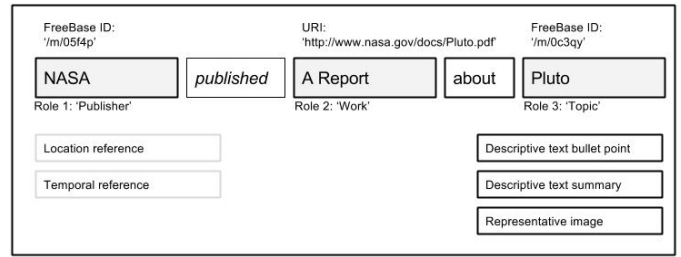


Figure 2: A simplified instantiated event, using the event frame of Figure 1.

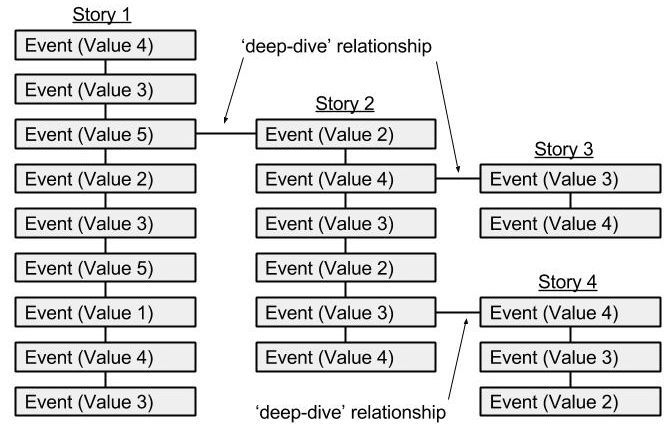


Figure 3: A simplified narrative structure with ‘deep-dive’ stories and values, illustrating references to instantiated events similar to Figure 2.

story order, temporal order, common events, common event frames, common characters, common entities, common locations and, in an earlier research prototype, assigned cause and effect relationships. These connections enable traversal between stories (and within stories) from any structural element. A more detailed technical description of the Structured Stories prototype as a computational model of narrative applied to journalism is provided in Caswell 2015 [5].

The technical feasibility of Structured Stories for representing simple journalistic events and narratives was demonstrated in late 2014 using a publicly accessible prototype populated with reporting on Los Angeles local government stories. Significant editorial questions remained unanswered, however, including questions about the sufficiency of FrameNet for semantically representing journalistic events, the sufficiency of knowledge graphs for providing references to characters, entities/concepts and locations, the feasibility of creating and maintaining event frames sufficient to capture journalistic events, the practicality of arranging structured events into narrative structures, and the ability of journalists without a technical background to identify and capture discrete journalistic events under realistic conditions.

An experiment intended to explore these unanswered editorial questions was therefore conducted. A team of three students from Duke University’s Reporter’s Lab performed full-time reporting in New York City during June and July

of 2015. During that time the Structured Stories New York City experiment generated about 60 structured stories covering housing, law and order, transport, and other topics. The reporters originated and pursued stories and recorded their reporting as structured events and narratives in the Structured Stories prototype platform, using a sequential and semi-automated browser-based user interface. Reporters were unrestricted as to which stories they could pursue (within broad topic areas) and were unrestricted as to which events they chose to structure. Events were structured using existing event frames, or using event frames created on the first encounter with a new type of event. The reporters received several days of introduction to the Structured Stories concept and process, and were provided with remote support and coaching throughout the experiment. Reporters worked from an office in New York City, and also within a ‘virtual newsroom’ environment using the Slack messaging system and the Skype video channel. All structured events and narratives were published as they were reported and structured.

2. STRUCTURED EDITORIAL ASPECTS

The reporting phase of the Structured Stories New York City experiment took place during June and July of 2015. The results have not yet been fully analyzed, however there are several important observations available from the operational experience of the experiment and from a general analysis of the structured reporting that it produced. These observations are of the editorial aspects of reporting into structure - i.e. those aspects that concern the application of human editorial judgement at the event frame, structured event and structured story levels.

The experiment succeeded in capturing events and narratives as structure. Approximately 60 stories in 5 topic areas were recorded, containing approximately 615 discrete journalistic events, based on more than 220 event frames. This demonstrates that it is possible for reporters without a technical background to internalize this approach to structuring journalism, and to routinely carry out the series of specific actions required to apply it. These required actions were:

1. Identify discrete journalistic events from reporting.
2. Identify the primary activity of the event, sufficient to find (or request) an event frame that described it.
3. Identify ‘participants’ in the event, including characters, entities/concepts, locations and information artifacts.
4. Specifically locate the event temporally and spatially.
5. Create discourse elements (text bullet points and summaries) relating to the event.
6. Assemble events into a coherent structured story representing a journalistic narrative.
7. Assign differential importance values to events within a structured story.
8. Assign ‘deep-dive’ references to some events within a structured story, linking to stories with further detail.
9. Identify edits and refinements of events and stories after they had been entered and published.

For example, in reporting a story about the relationship between Mayor Bill de Blasio and the NYPD, the reporters first identify discrete journalistic events within the story,

such as Mayor de Blasio’s commitment to review the use of stop-and-frisk tactics by the NYPD. They then identified the primary activity of the event (commitment) and select an appropriate event frame to capture the event - in this case ‘[A Character] Committed To Review The Use Of [A Method]’ based on the FrameNet ‘Commitment’ frame. The reporters then identify the specific occupants of the ‘character’ and ‘method’ roles in this event frame and, using the API-driven user interface, select ‘Bill de Blasio’ (Freebase unique ID ‘/m/0gjsd3’) and ‘Stop and Frisk’ (Freebase unique ID ‘/m/0j7mpr7’) to fill those roles. The specific time and location (as a GeoNames unique ID) of the event are then added, and an image, a text bullet point and a short text summary are attached as discourse elements. Other events in the story are added in a similar way, and each is assigned a value indicating its importance within this specific story. Some events in the story are expanded by associating them with other ‘deep-dive’ stories, thereby enabling a hierarchical recursion of stories that encapsulate detail.

All three reporters successfully and routinely performed all of these actions, however the performance of actions 1, 7 and 8 displayed a significant variance in editorial decision-making regarding structure. The identification of discrete journalistic events from secondary reporting (i.e. reporting from source documents, archives, etc.), sometimes revealed a tendency to attempt structuring of event sentences from text, rather than of the underlying journalistic event. This tendency was explicitly identified as an error and diminished over the course of the experiment in all reporters. Variance in assigning values to events (action 7) and in structuring hierarchies of ‘deep-dive’ stories (action 8) was observed, reflecting personal preferences of the reporters. This variance was not explicitly identified as an error during the experiment, but may be usefully subjected to editorial direction via guidelines.

The semantic frames provided by FrameNet, and the semantic roles within those frames, appeared to be sufficient to semantically describe the great majority (approximately 95 percent) of discrete journalistic events that were encountered. A significant portion of events (approximately 20 percent), however, required two FrameNet frames to be fully described - for example the beginning of a legal trial required both the ‘Process Start’ frame and the ‘Trial’ frame. A few events may require three FrameNet frames for semantic description, although these are still being investigated. The FrameNet dataset appears to be suitable as a basis for structuring most journalistic events.

The event frames that were created and used during the experiment appeared to be sufficient to represent the discrete journalistic events that were encountered from reporting. This was not particularly surprising because of the editorial flexibility within the definitions of event frames, however it did suggest that several key assumptions about the components of event frames are valid. The assumption that only 7 top-level types of semantic roles (character, entity/concept, location, information artifact, event, story and constant) are sufficient to capture all forms of noun participants in events appeared to be correct, as no situation was encountered in which an event participant was not a member of one of those types. The assumption that a single verb

lexical unit was sufficient as the basis of each event frame appeared to be correct. A few reported events did contain multiple verbs, however these events were easily reduced to a single verb without loss of the core meaning of the event. Several insufficiencies in the handling of event frames within the prototype platform were revealed, including the need for a taxonomy of event frames to improve the findability of suitable event frames by reporters, the need to manage small lists of identified characters as a single character, and the need to distinguish negative events from otherwise identical positive events (e.g. ‘attended’ or ‘did not attend’). The requirement for semantically describing a significant portion of events using references to two FrameNet frames (described above) will also require changes to event frames, although this deficiency did not influence this experiment.

Several problems were revealed regarding the sufficiency of knowledge graphs for describing the noun participants of events. A small but significant proportion of characters and entities/concepts encountered in the reported events did not have unique identifiers in knowledge graphs and therefore required the creation of records that were not ‘grounded’ externally from the Structured Stories database. Furthermore, some characters, particularly in law and order stories, were unavoidably anonymous and challenges were encountered in distinguishing between different anonymous characters. Constraints in the user interface prevented the use of location nouns more granular than neighborhoods, although most required locations were in fact present within the GeoNames knowledge graph and were added later in the editing process. It is expected that the coverage of characters and entities/concepts as external identifiers will improve with an anticipated migration from Freebase to Google Knowledge Graph and WikiData, however it is clear that more focus on handling semantically ‘ungrounded’ characters and entities/concepts will be required.

The assembly of structured events into narratives that were editorially structured using deep-dive stories and assigned importance values was undertaken with apparent ease by the reporters, however the editorial inclusion of the same events in multiple narratives, part of the ‘narrative network’ aspect of structured narratives, was rarely used. This was probably partly because of deficiencies in the user interface, which required the assignment to be made in editing, but appeared to be mostly because very few reported events were relevant to more than one narrative. This observation may not be relevant to other types of stories or reporting scenarios, for example in scenarios where a single event initiates multiple stories, but was clearly the case here. Other aspects of the ‘narrative network’ effect, such as connections between narratives based on common characters or common locations, were frequently observed. Cause and effect relationships between events were not captured during this experiment.

3. MACRO EDITORIAL ASPECTS

In addition to the structured editorial observations discussed above, several phenomena were observed relating to the interaction between the use of structure and the expectations and norms associated with traditional journalism.

A general observation made during the experiment was that, while the structured format is clearly very different from tra-

ditional journalism formats, the essential editorial decisions made by the reporters were similar. Decisions about which stories to cover and which events to report within those stories were obviously similar. Decisions about the differential assignment of value to specific events within the story were similar. Decisions about the use of ‘deep-dive’ stories to encapsulate detail about events had analogs in the use of sidebars in traditional journalism. An approximate analog seems to exist between structured events and paragraphs in the body of text articles. These similarities are perhaps not surprising, as the structured format and the traditional format are both seeking to create coherent narratives that are comprehensible by humans, one as a designed approach and the other as a craft approach that has been highly optimized over many years.

Observations were also made about differences between the structured approach and a more traditional approach to journalism. The specificity required by structuring events and stories, and the limited capacity for description afforded by the format, sometimes forced the reporters to explicitly make journalistic decisions that may not have clear analogs in traditional reporting. The necessity to decide which specific event frame to use to represent each event forced specificity about the action involved, possibly to a greater degree than would have been necessary if the flexibility of writing in natural language had been available - for example debate sometimes occurred about the difference between a speech act or a publishing act and the activity described by the speech or the publication. This need for specificity also appeared to influence the perceived importance of peripheral aspects of each event, providing nuance or ‘color’, which were increasingly seen as superfluous to core event structure and were therefore either constrained to the text discourse elements attached to each structured event or abandoned altogether. Some differences were also observed at the narrative level, where constraint to an explicit story order possibly reduced the addition of anecdotal events or example events to a greater degree than would perhaps have been necessary if the flexibility of writing in natural language had been available.

A common observation was the surprisingly small number of core events in stories, or in portions of stories. This realization appeared to reduce the tendency of the reporters to conduct speculative field reporting, because the investment of time often did not result in commensurate numbers of ‘structurable’ events. That, combined with the need for historical ‘backstories’ to be reported from archives and structured, appeared to lead to a bias toward secondary reporting rather than primary reporting. The journalistic value attributed by the team to sense-making from secondary reporting appeared to be lower than the journalistic value attributed to primary reporting, a observation that has been seen in other structured journalism newsrooms [6]. Finally, there were indications that the reporters may have been somewhat bored by encoding events as structure, and may not have drawn as much intrinsic satisfaction from the sense-making aspect of the approach as was lost by the reduced need for writing. These indications showed some variance by reporter and may be related to individual skill and interest in abstraction and sense-making, however the reporters also noted that the Structured Stories approach opened the

possibility that individuals who were drawn to the reporting and research but not the writing functions of journalism could more readily participate in the production of news.

Of all structured events reported during the experiment it is apparent that a significant portion, perhaps as much as 50 percent, are events that are communication acts or speech acts of some kind. This was observed across all topic areas, and confirms reports from other media observers [2]. This phenomena resulted in substantial use of the FrameNet ‘Statement’ frame, ‘Communication’ frame and similar frames, and produced a large number of related event frames with relatively subtle distinctions in meaning. This will probably require special handling of related verbs that are used to locate these frames, a taxonomy of related event frames and associated specialized user interface elements.

The distribution of the use of event frames during the experiment, as indicated by frame requests, seemed to follow a power law - a result that is obviously related to the predominance of speech and communication acts. This result requires further analysis, however it was very clear that the demand by reporters for new event frames dropped precipitously after several weeks of reporting in a topic area, despite increasing numbers of events reported. It also appears that the distribution of characters and locations within individual stories may also follow a power law, which, if verified, is an unanticipated result.

4. CONCLUSIONS AND NEXT STEPS

The Structured Stories New York City experiment was, to our knowledge, the first attempt to report general news into a computational model of narrative based on frame semantics. The prototype platform and the editorial processes used in the experiment were nascent and exploratory, and the experiment was more of a reconnaissance than a definitive survey. Nonetheless, important and useful observations of structured editorial issues and also of more general editorial issues were obtained, and, equally importantly, no editorial challenges were encountered that present obvious conceptual barriers to further refining the approach and exploring its possible utility as a platform for novel news products.

Although the experiment succeeded in representing real-world news as structured data, many serious challenges were identified and will require solutions in order to refine the approach. A large set of technical bugs and user interface deficiencies must be addressed, including the expansion of event frames to incorporate a second FrameNet frame. Various editorial constraints will likely be required, including hard constraints implemented in the software and soft constraints implemented as editorial guidelines, for example guidelines on the choice of deep-dive stories or importance values for representing detail in stories. The training and communication associated with the approach will also need to be improved, using learning from this experiment.

This experiment also did not meaningfully explore the utility of the Structured Stories representation to consumers of news. This utility may be considerable, and theoretically could include opportunities to accumulate journalism as an

integrated network, to convert journalism from a ‘publisher-decides’ to a ‘consumer-decides’ paradigm and to enable computational tools to be applied directly to journalism. These possibilities, and others, have yet to be explored.

A second reporting experiment is being planned for late 2015 at the Reynolds Journalism Institute (RJI), focused on a narrower set of stories associated with state government in Missouri and incorporating many of the lessons learned during the Structured Stories New York City experiment. An initial study of the consumption aspects of Structured Stories, focused on comprehension of large stories, will also be conducted at RJI in late 2015.

The Structured Stories prototype, containing all reporting from the experiment, is available at www.structuredstories.org.

5. ACKNOWLEDGMENTS

Funding for the Structured Stories New York City experiment was provided by the Online News Association Challenge Fund and by the Reporter’s Lab at Duke University. Evaluation of the Structured Stories platform is supported by the Reynolds Journalism Institute. The authors also thank the Structured Stories NYC reporting team of Natalie Ritchie, Ishan Thakore and Rachel Chason.

6. REFERENCES

- [1] C. Baker. Framenet, Present and Future. In *First International Conference on Global Interoperability for Language Resources*, 2008.
- [2] A. Bell. *The Language of News Media*. Wiley-Blackwell, New Jersey, USA, 1991.
- [3] G. Bhatia. How newsrooms are using machine learning to make journalists lives easier. *Poynter.org*, 2015.
- [4] B. Boyd. *On the Origin of Stories: Evolution, Cognition, and Fiction*. Belknap Press, Cambridge, MA, USA, 2010.
- [5] D. Caswell. Structured Narratives as a Framework for Journalism: A Work in Progress. In *Proceedings of the Sixth International Workshop on Computational Models of Narrative*, Atlanta, GA, USA, 2015.
- [6] M. Coddington. Telling secondhand stories: News aggregation and the production of journalistic knowledge. *Doctoral Dissertation, The University of Texas at Austin.*, pages 178–209, 2015.
- [7] D. Davidson. The Logical Form of Action Sentences. In N. Rescher, editor, *The Logic of Decision and Action*. University of Pittsburg Press, 1967.
- [8] C. Gourarie. Structured journalism offers readers a different kind of story experience. *Columbia Journalism Review*, 2015.
- [9] T. Parsons. *Events in the Semantics of English: A Study in Subatomic Semantics*. The MIT Press, Cambridge, MA, USA, 1990.
- [10] G. A. Radvansky and J. M. Zacks. *Event Cognition*. Oxford University Press, Oxford, United Kingdom, 2014.
- [11] P. H. Winston. The Strong Story Hypothesis and the Directed Perception Hypothesis. In *Proceedings of the AAAI Fall Symposium on Advances in Cognitive Systems*. Association for the Advancement of Artificial Intelligence, 2011.