Automated Journalism 2.0: Event-Driven Narratives

From simple descriptions to real stories

A revised version of this manuscript is published in *Journalism Practice*:

David Caswell & Konstantin Dörr (2017): Automated Journalism 2.0: Event-driven narratives, Journalism Practice, DOI: 10.1080/17512786.2017.1320773

To link to this article: http://dx.doi.org/10.1080/17512786.2017.1320773

AUTHORS

David Caswell and Konstantin Dörr

ADDRESS

David Caswell, Donald W. Reynolds Journalism Institute at the University of Missouri School of Journalism. E-mail: david@structuredstories.com, @StructStories

Konstantin Dörr, Institute of Mass Communication and Media Research - IPMZ, Media Change & Innovation Division, University of Zurich, Switzerland. E-mail: k.doerr@ipmz.uzh.ch, @kndoerr

ACKNOWLEDGEMENTS

The authors would like to thank Automated Insights, Inc. (in particular James Kotecki and Locky Stewart), NBC SoCal (in particular Sara Catania, Jason Kandel and Ryan Smythe), the Reynolds Journalism Institute at the University of Missouri, and especially one reviewer for helpful and constructive comments.

DISCLOSURE STATEMENT

No potential conflict of interest was reported by the authors.

FUNDING

The research was supported by the Donald W. Reynolds Journalism Institute at the University of Missouri School of Journalism, The Reporter's Lab at Duke University, and the Online News Association's Challenge Fund.

Abstract

This article introduces an exploratory computational approach to extending the realm of automated journalism from simple descriptions to richer and more complex event-driven narratives, based on original applied research in structured journalism. The practice of automated journalism is reviewed and a major constraint on the potential to automate journalistic writing is identified, namely the absence of data models sufficient to encode the journalistic knowledge necessary for automatically writing event-driven narratives. A detailed proposal addressing this constraint is presented, based on the representation of journalistic knowledge as structured event and structured narrative data. We describe a prototyped database of structured events and narratives, and introduce two methods of using event and narrative data from the prototyped database to provide journalistic knowledge to a commercial automated writing platform. Detailed examples of the use of each method are provided, including a successful application of the approach to stories about car chases, from initial data reporting through to automatically generated text. A framework for evaluating automatically generated event-driven narratives is proposed, several technical and editorial challenges to applying the approach in practice are discussed, and several high-level conclusions about the importance of data structures in automated journalism workflows are provided.

Keywords

Automated Journalism, Computational Journalism, Structured Journalism, Structured Events, Structured Narratives, Natural Language Generation, Computational Narrative, Data Journalism

Introduction

Natural language generation (NLG) technology – the automatic creation of text from digital structure data – has been productized and commercialized over the past decade by several companies, such as Arria, Applied Semantics, Automated Insights, and Narrative Science. This technology first appeared in the 1950's as part of machine translation efforts (McDonald 2010; Reiter 2010), but its use has expanded substantially in recent years. Modern NLG products are primarily marketed as enterprise tools for making insights from 'big data' accessible to people without specialized expertise, however these technologies are also increasingly used to automate certain routine writing tasks within news organizations – a practice known as 'automated journalism' (for an overview see Dörr 2016).

Previous research into automated journalism classified it as a form of quantitative journalism, sometimes also referred to as 'algorithmic journalism' or 'robot journalism' (Anderson 2012; Coddington 2015). This research has focused on its professional use by journalists (Thurman, Dörr, and Kunert 2017), on audience perceptions of automated content (Clerwall 2014; Graefe et al. 2016), on societal or economic implications (van Dalen 2012; Carlson 2014; Young and Hermida 2015; Latar 2015; Dörr 2016), on transparency (Diakopoulos 2015; Diakopoulos and Koliska 2016), on attribution (Montal and Reich 2016),

on professional, organizational, and societal levels of responsibility (Dörr and Hollnbuchner 2016), and on a broad range of technical, cultural, political and ethical questions (Boyd and Crawford 2012; Lewis and Westlund 2015a,b; Hammond 2015; Saurwein et al. 2015). Despite various concerns revealed by this academic attention, the use of automated journalism in newsrooms has steadily increased, motivated by an urgent need for sustainable business models for news organizations and by intense economic pressure to increase productivity in producing journalistic content (Anderson et al. 2012; Örnebring and Ferrer Conill 2016).

Early examples of the use of NLG technology to automate journalism are mostly confined to relatively short texts in limited domains, but are nonetheless impressive in terms of both quality and quantity. The text produced is generally indistinguishable from text written by human writers and the number of text documents generated substantially exceeds what is possible from manual editorial processes (Clerwall 2014; Graefe et al. 2016). Automated writing tools are therefore increasingly being integrated into publishing workflows, including at the Associated Press and at the Los Angeles Times. Despite this success there is a widespread consensus that the potential for NLG technology to automate journalistic writing is limited to simple descriptions that are already commonly represented as data, the most obvious examples being the writing of routine sports and financial news. This descriptive reporting of routine news is an important part of the journalistic portfolio of most newsrooms, but NLG technology has not yet been shown to be suited to the production of longer-form news articles and it is unable to incorporate the storytelling value of more sophisticated, and more valuable, journalism (Dörr 2016).

Extending the use of automated journalism beyond simple descriptive reports is not prohibited by deficiencies in NLG software, but by the absence of appropriate data needed to automate more complex journalistic writing, such as event-driven storytelling. This 'absence of data' does not merely refer to the idea that certain data are not routinely collected or cannot easily be procured, but to the much more challenging idea that the 'kind' of data required for the application of automation to complex journalistic storytelling does not exist. In other words, the major barrier to the automation of more sophisticated journalism is not just the absence of data but also the absence of methods of encoding information ('data models') that are appropriate for capturing more sophisticated journalistic knowledge as data. Events are the central semantic feature of news, and therefore the lack of data models suitable for encoding journalistic events and stories as data prevents the adoption of automated journalism beyond routine descriptive reporting.

This article focuses on an exploratory computational approach to overcoming this 'absence of data' barrier and thereby extending the realm of automated journalism from simple descriptions to richer and more complex event-driven narratives. We will briefly review the research project upon which this work is based, and then will describe a generic NLG platform, illustrating why the first ('version 1.0') generation of automated journalism practices have been restricted to simple description. We will then present a computational view of 'story', showing how simple event-driven narratives can be represented as data. This is followed by a proposal for using NLG and computational story representations together to generate complex journalistic narratives, with a description of two methods to doing so. We provide detailed examples of each method, including a successful demonstration of Automated Insights´ WordSmith platform to automatically generate event-driven journalistic narratives of car chases ('pursuit stories') using data from a prototyped database of fully structured journalistic events. These approaches

show how automated journalism might facilitate more complex digital storytelling in editorial environments, enabling a new ('version 2.0') generation of automated journalism practices. Finally a discussion of the challenges associated with more complex automated journalism is presented, including the challenge of evaluating automated texts and the challenge of extending journalistic education and skill sets to encompass 'computational thinking'.

The Structured Stories Project

The project described in this article is part of a larger research and development effort focused on exploring the application of structured events and structured narratives to journalism. This effort is based on an experimental platform called Structured Stories, which enables journalists to enter journalistic events and narratives into a 'story database' that is not based on written text, but instead uses the semantics, or 'meaning', of journalistic events to represent news stories entirely as structured data (Caswell 2015). Story databases like Structured Stories demonstrate a radically new approach to working with journalistic information and enable many novel ways to collect, distribute and consume news. The platform was evaluated during 2015 and early 2016 by reporters from Duke University's Reporter's Lab, from the University of Missouri School of Journalism and from a major Los Angeles-based newsroom (Caswell et al. 2015).

For this project the structured event and structured narrative information stored in the Structured Stories database were used to enable the automatic generation of event-driven narratives using WordSmith, an automated writing tool developed by Automated Insights, one of the leading providers of commercial NLG technologies. A series of trial-and-error experiments were conducted during early 2016, seeking to format data from the Structured Stories database in a way that could be used with WordSmith to generate event-driven narrative texts. Results from these experiments were shared with Automated Insights in April 2016 and the approaches are now being applied in ongoing development of the Structured Stories platform.

Natural Language Generation

Commercialized automated writing platforms suitable for application in journalism (such as Automated Insights' WordSmith platform or Narrative Science's Quill platform) are usually based on an approach to NLG that is sometimes referred to as 'dynamic templating' (Deemter et al. 2005). These platforms enable the construction of language generation templates – complex arrangements of logic and associated text fragments. Each template is designed to operate with a particular data model, and can then be used by the platform to assemble blocks of written text from its text fragments based on any dataset with fields formatted in accordance with that data model. These templates are essentially discrete hierarchical trees of conditional (i.e. 'if-then-else') logic that provide a text-assembly procedure for all eventualities that may occur in any dataset formatted according to the designated data model. For example: If data field 1 is X then write "this", else if data field 1 is Y then write "that". Although this technique appears to be relatively simple, implementing it in a way that can be practically used in a production editorial environment is not trivial, because the complexity of the template increases very quickly for most useful applications.

The central challenge of a template-based NLG platform is in enabling template design in a way that facilitates the interplay between data-determined logical choices and the authoring of text fragments that ensure grammatical coherence in the assembled block of text. The WordSmith platform, for example, provides an array of tools that enable template authors to manage trigger data, to detect and use features and trends in data, to use nesting of conditional logic branches, to apply comparative logic, to randomly use lexical and phrase synonyms, to format numerical text, etc., as well as providing a simple authoring environment and testing process. Authoring NLG templates for automated journalism requires continually trading off between the multiple criteria that contribute to the quality of the possible text outputs created using the template. These criteria include the target data model, the type and range of values possible in each field of the data model, the quality and completeness of data, the grammatical coherence of the text output under all data values, the editorial voice of the text output, and, of course, the author's structural and stylistic choices (or 'angle') to be expressed in the text output. The ability to work with these criteria editorially, as a writer, rather than technically, as a programmer, is a key distinction between modern commercial NLG systems and their predecessors, which required the building of complex text plans as 'discourse structure trees' and sentence plan scripts written in 'Sentence Planning Language' (Reiter and Dale 1997). Thankfully, template authoring for journalism is now largely an editorial task, however it still requires 'computational thinking' – an ability to think abstractly about the use of language while simultaneously applying logical rules and practical language skills (Wing 2006; Creech and Mendelson 2015). This is a skillset that is not yet a common component of journalism education.

This description of NLG and its use in automated journalism reveals several major limitations relating to data. The obvious limitation is that the entire process is completely determined by the available data. A less obvious limitation is that the use of available data is also restricted by the data model by which the data is organized. It doesn't matter, for example, that a certain type of data is available if it is blended with other data within a single field, or if it is present in an arbitrary manner across multiple fields, or if its interpretation is ambiguous either within individual fields or across multiple fields. An additional limitation is imposed by the logical complexity of templates, which can quickly become unmanageable if the available data fields are not directly related to their intended use within the template. A substantial intellectual burden is imposed on the template author if he/she is forced to use conditional logic to isolate aspects of the data needed for text description, rather than just access the necessary information directly as a data field.

These data-related limitations combine to restrict most uses of NLG to description – the static portrayal of a situation or the status of some multi-variant system or scenario, such as an earnings report containing multiple financial variables. Even the use of NLG platforms to describe gameplay, to compare descriptions of periodic snapshots, or to describe time series information, is typically done in a primarily descriptive manner that recites aspects of the underlying data as a 'state', or as a static picture rendered in words. In practice this often results in output text that contains a relatively narrow range of verbs, such as 'was', 'is', 'increased' and 'decreased', or their synonyms. Automated journalism using existing data and existing data models might therefore often be more accurately referred to as 'automated description'.

Figure 1 shows an example of a descriptive text written with NLG software. This text was pro-

duced and published in 2016 by the Associated Press using Automated Insights' WordSmith NLG platform driven by quarterly financial data provided by Zacks Investment Research. This text can be automatically generated using only a handful of data variables, such as share price at and after the close of the market, posted revenue and earnings, and averaged analyst expectations of earnings.

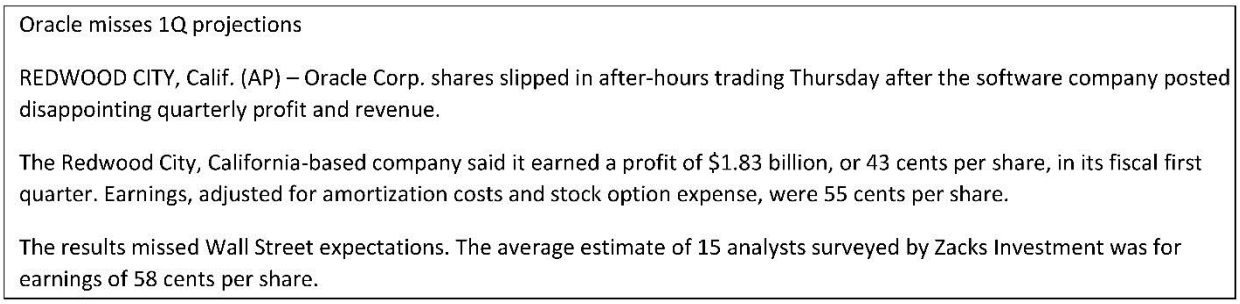

Oracle misses 1Q projections

REDWOOD CITY, Calif. (AP) – Oracle Corp. shares slipped in after-hours trading Thursday after the software company posted disappointing quarterly profit and revenue.

The Redwood City, California-based company said it earned a profit of $1.83 billion, or 43 cents per share, in its fiscal first quarter. Earnings, adjusted for amortization costs and stock option expense, were 55 cents per share.

The results missed Wall Street expectations. The average estimate of 15 analysts surveyed by Zacks Investment was for earnings of 58 cents per share.

FIGURE 1
Example of a descriptive quarterly earnings report generated automatically for the Associated Press using WordSmith and data provided by Zacks Investment Research

Computational Stories

Journalism is not fundamentally based on description, but on event-driven story. Journalists are highly skilled at recognizing, creating, organizing and communicating stories in natural language, and the centrality of story to journalism is almost unanimously acknowledged by journalists (Roeh 1989; Pavlik 2000; Deuze 2005). The concept of 'story' is sometimes difficult for human beings to clearly define, because consuming stories is such an effortless and ever-present experience for humans that defining the mechanism seems unnecessary (Boyd 2009). But story is a mechanism. More specifically it is a knowledge representation mechanism, possibly of biological origin, for organizing temporal events in a manner optimized for cognitive access and use. A rapidly growing body of literature in a wide variety of scientific disciplines, from cognitive psychology to linguistics to artificial intelligence, is identifying the components and applications of story (Radvansky and Zacks 2014; Baldassano et al. 2016; Mani 2013), and an array of technologies have recently appeared that enable working with story computationally to become practical for the first time (Rospocher et al. 2016; Caswell 2015).

Stories are arrangements of events and therefore representing news stories computationally (i.e. as data) requires first representing discrete news events computationally (Abbott 2008). Event representation requires constructing libraries of 'event abstractions' – generalized forms of fundamental types of events that define the action of the event and the semantic roles that various participants, entities and locations play in that action. For example, a simplified event abstraction might take the form of '[X] appointed [Y] to the position of [Z]', in which the activity of the event is narrowly specified (the appointment of a person to a position by another person or organization), and the specific semantic roles (the appointer, the appointee, the position) are variables. These event abstractions can then be used to capture particular news events as structured data, a process known as 'instantiation'. Particular news events are instantiated by 'filling in' the semantic role variables with references to the participants, entities, locations, etc. to create a unique record of that event. This approach, while still experimental, has become

much more practical in recent years because of the new availability of libraries of event abstractions such as FrameNet (Baker 2008), the Events and Situations Ontology (Segers et al. 2015) or Schema.org 'actions' (Schema.org), and also because of the availability of common digital references for participants, entities, locations, etc. (Knowledge graphs and other 'semantic web' infrastructure).

The arrangement of structured events into coherent story representations requires more than merely assembling a temporally-ordered list of events. Stories have at least three properties that should ideally be represented in the arrangement of structured events, all of which are relatively simple:

1. Events should exist independently of the stories in which they are used, thereby allowing any specific event to be present in multiple stories and facilitating the development of 'narrative networks'.
2. Events should be assigned different values within a particular story, indicating their relative importance within that story based on one or more value schemes.
3. Events should be arranged in a way that facilitates 'semantic zoom' - the ability to recursively nest story detail such that it can be directly accessed from particular events.

The combination of a representation scheme for capturing news events as structure and a representation scheme for assembling structured events into stories together provide a possible data model for general news stories.

Projects that implement news story representations in a manner similar to that described above are increasingly common. An early example is the BBC's Storyline ontology (Rissen et al. 2013), and a more recent example is the recently completed NewsReader project (Rospocher et al. 2016), which is centered on what its developers call 'an event-centric knowledge graph'. An academic workshop on "computing news storylines" is held annually, and several similar projects are underway in the government and intelligence communities. Simultaneously, a new form of journalism, known as 'structured journalism' has recently emerged, which focuses on technologies and workflows for capturing and publishing journalistic knowledge as structured data. Although practical datasets based on taxonomies of activity are still rare, there are several examples of practical datasets that are very close to an event-based format, particularly in the government domain (e.g. GovTrack.us, TrackBill.com, etc.).

The news database used for the demonstration of NLG in this article was Structured Stories (Caswell 2015), a prototyped news database that uses FrameNet (Baker 2008), a library of semantic frames, as the formal basis of event structure. Unlike some other event and story representations, Structured Stories relies on the entry of structured events and stories manually, by journalists, and it is therefore particularly well suited as a basis for experiments in automated journalism due to the high quality of its data. The structured event representation used by Structured Stories is based on the concept of frame semantics – a view that meaning in language exists only within specific contexts, known as a 'semantic frames'. A simple example is the 'commerce' semantic frame, which implies certain roles (buyer, seller, goods, payment, etc.) and certain actions (buying, selling, paying, etc.). Structured Stories uses FrameNet to define detailed 'event frames' that provide abstractions of specific journalistic events, which can then be instantiated as records of actual events – for example "Microsoft Inc. bought Skype Communications from Silver Lake L.L.C. for $8.5 billion USD". These structured records of actual events

are then assembled into structured stories based on the story characteristics described above.

Event-driven Natural Language Generation

The availability of structured databases of news events and news stories opens up the possibility of using template-based natural language generation as a mechanism for automating the production of journalistic text articles based on event-driven stories rather than just static description. While some forms of knowledge representation existed in earlier NLG systems (Reiter and Dale 1997), these were typically entities, concepts and relations rather than events and stories. In contrast, the availability of structured events and structured narratives offers new potential to represent news events and journalistic stories directly, for use in automated writing. Realizing that potential is not straightforward, however, because the contents of the structured news database must still be translated into a standardized form suitable for use with specific NLG templates. This is difficult because news stories (whether structured or not) are inherently unpredictable, whereas NLG templates essentially require data that is perfectly predictable.

One approach to converting a structured story into a standardized format suitable for use with an NLG template is to develop abstractions of stories – a standardized 'kind' of story, within a particular domain, that can be loosely expected to contain particular 'kinds' of events in a particular order. A familiar example of a story abstraction is the 'road trip story' – each one is different, but each also can be expected to contain particular kinds of events in a particular order (a motivating event, setting out, various geographic milestones, hotel stays, meals, various adventures, the arrival or return, and possibly some consequences). If the domain is constrained enough, and if the variation from story to story is small enough, then events from the structured story database can be mapped to 'event kinds' in the story abstraction to produce a standardization of each story 'kind'. This approach can be relatively easy to implement and was used for the project described below, demonstrating automated generation of car chase (or 'pursuit') stories from structured events (example 1). It would likely work equally well with a wide variety of relatively predictable stories in journalistic domains such as crime, business, politics, disaster reporting, etc., but would be unsuited for less formulaic kinds of stories.

A second approach is therefore needed in order to convert relatively unpredictable structured stories into a standardized format. One possibility is to break up the story into a series of 'story fragments', each of which is then mapped to a standardized format matched with a particular NLG template. These 'story fragments' are converted into natural language independently and the resulting text blocks are assembled into a complete text document. This is a much more challenging approach because it requires that each 'story fragment' be correctly recognized within the story database and correctly mapped into the correct standardized format. It also requires that the text within each of the NLG templates be authored such that the final text document assembled from the text blocks is grammatically and logically coherent. This is the approach taken for an ongoing project attempting to automatically generate local and state government stories from structured events (see example 2 below). Although more challenging, this approach also offers the possibility of generating much more flexible stories that may not need to be constrained to narrow domains.

In terms of the specific data inputs to an natural language generation platform, the availability of

structured event data enables an important and possibly novel characteristic in how events are represented in the standardized data used by the platform: it enables events and their semantic roles to be handled as data values instead of as data fields. This emphasis on news events as variables enables the automation of text that approaches true storytelling, with a much wider variety of verbs communicating a much richer range of activity. Direct access to structured news events might also eventually permit the encoding of causality as a structured relationship between discrete events, thereby enabling automated journalism to communicate causation and explanation.

Example 1: Pursuit stories

The pursuit story project originated with a small set of car chase stories reported into the Structured Stories database, using the platform's user interface, by reporters from the NBC SoCal newsroom in Los Angeles during the autumn of 2015. The intent of the project was merely to explore the potential of structured journalism to accumulate pursuit story content, possibly for interactive display, and there was no initial intention to use the data for automated journalism. The events of each story were reported directly as structured data, using information obtained during the newsroom's normal reporting activities.

As the stories accumulated it became apparent that there were significant similarities from story to story - not in the literal activity of the story, which varied wildly, but in the 'part' that each event played within the story. For example, each story contained an 'initiating event', ranging from speeding to murder. Each story contained a 'concluding event', ranging from the suspect simply giving up to the suspect being shot dead by police. Most stories contained one or more 'consequence events', such as jail time or hospitalization. Many stories contained some kind of distinct 'unusual event', some remarkable occurrence but with specific activity that varied dramatically from story to story. The observed similarity was therefore in the form of these structured stories, rather than in the details of the activity within them. The stories were quite literally formulaic.

This observation, combined with access to Automated Insight's WordSmith platform, led to a small development project aimed at transforming the structured event data from these structured pursuit stories into a standardized data format suitable for use with a WordSmith language generation template. The steps taken to accomplish this data transformation task were as follows:

1. Definition of a standardized data model based on data fields representing the 'parts' that events played within a generic pursuit story (i.e. not based on the events themselves).
2. Mapping the structured events from each story to the standardized data model. This was first done by hand, and then semi-automatically using signals from the frequency of use of 'kinds' of events.
3. Authoring of text fragments for each of the 'event abstractions' used to define the structured events in the stories. To be clear, these text fragments are not authored for the many instantiated events, but for the relatively few event abstractions upon which the instantiated events were based.
4. Development of a WordSmith template for the standardized data model, which could then be tested against data formatted with that data model.

5. Iteration of the process until the standardized data model and the WordSmith template stabilized.

The final standardized data model that resulted from this process contained 48 data fields, listed in figure 2. Twenty-one of these data fields were based on structured events from the story database, and these are highlighted. Some structured events generated multiple data fields, and not every story generated data values for each data field.

Story_ID	Handover_location	Route_2	Unusual_event_2
Headline	Suspect_1st_name	Location_3	Quote_char_name
Date	Suspect_2nd_name	Route_3	Quote_char_title
Year	Suspect_description	Location_3_area	Quote_text
Time	Passenger_num	Concluding_event	Link
Time_colloquial	Vehicle_type	Conseq_event_1A	Link_desc
Time_type	Vehicle_description	Conseq_event_1B	DB_num
Dept1_long	Originating_event	Conseq_event_pax	DB_num_fatal
Dept1_short	Unusual_event_1	Conseq_event_2	DB_num_dept
Dept2_long	Location_1	3rd_party_name	DB_num_prev
Dept2_short	Route_1	Conseq_event1_3rd_party	DB_num_fatal_prev
Handover_type	Location_2	Conseq_event2_3rd_party	DB_num_dept_prev

FIGURE 2
Data fields filled from structured event data and used to generate story text

The data values required for each of the event-based data fields were fragments of descriptive text that could be used by WordSmith to assemble the output text using the template. This required associating a fragment of descriptive text with each of the event frames used to instantiate the structured events of the stories, enabling that text to be used for each event instantiated from that event frame. These text fragments were relatively simple.

Some text fragments were single words describing an outcome of an activity, for example:

"injury"
"death"

Some text fragments were short statements based on the police designation upon which the event was based as a definite or indefinite article, for example:

"a possible DUI [driving under the influence]"
"a report of shots fired"
"the identification of a wanted person"
"the observation of suspicious activity"

Some text fragments referred implicitly to the main character of the event in describing an outcome of the event, for example:

"treated for injuries"
"pronounced dead"
"taken into custody"
"captured following a search"

Finally, some text fragments were more complex descriptions of the structured event. Unlike the text fragments described above, these complex text fragments cannot be constructed solely from the information contained in the event frame, because they use information from the semantic roles of the instantiated structured event. Examples of these complex text fragments are:

"the suspect deliberately crashed into a police vehicle"
"the driver hit speeds of 100 mph"

The generation of standardized data for each structured pursuit story already directly uses information from the semantic roles of the instantiated structured events (for example character names, locations and routes). The technical aspect of constructing these complex text fragments is therefore not particularly difficult, although the editorial aspect may be more challenging.

Each structured pursuit story was therefore mapped into the single standardized data format created using the techniques described above. Although some of this mapping was performed manually for this experiment, manual intervention was only permitted via clear procedures that could be easily automated. In practice this was restricted to aspects of the selection process by which specific structured events were chosen to fill specific event-based data fields in the standardized data model.

An example of an automatically generated pursuit story article generated by WordSmith from the standardized data is shown in Figure 3:

Driver Drives Off Of 300 Ft Cliff During Pursuit

July 13TH, 2015 (Point Fermin Park, San Pedro, California) – A vehicle pursuit that began in Wilmington during the late evening of July 13th later ended with the crash of the suspect in Point Fermin Park, San Pedro. The incident began at about 11:00 PM when an unidentified driver, driving a Toyota Prius, fled from officers of the Los Angeles Port Police following a traffic stop on Pacific Coast Highway in Wilmington. The suspect was then pursued by the LAPP along Alameda Street to San Pedro, and then further along Alameda Street to Point Fermin Park in San Pedro. The incident concluded with the crash of the suspect in Point Fermin Park, where the suspect drove over a cliff. The unidentified suspect was injured. Referring to the incident, witness Manuel Castro said "We peeked our heads and it was just a gray Prius and we saw the wreckage and the cops over here". The suspect was treated for injuries at the scene and was expected to be arrested.

This was the 0th vehicle pursuit in the greater Los Angeles region during 2015, and the 0th pursuit of the year conducted by the Los Angeles Port Police. In the previous year there were a total of 0 vehicle pursuits in the Los Angeles region, of which just 0 were conducted by the LAPP.

FIGURE 3
Example of a dynamic "pursuit story" generated automatically from event data from a structured story. Note: The last paragraph in this article is intended to be driven by accumulated data within a structured news database. This project involved only a small number of stories and therefore the values were left at zero to avoid confusion that might result from the use of test data

Although not longer than many text documents already generated automatically, and not based on more data fields than many text documents already generated automatically, these 'pursuit stories' stories are unlike other output of automated journalism using NLG, because they are clearly reciting an event-by-event narrative. The authors are not aware of any examples of automated journalism that exhibit a similar event-driven narrative structure. The closest such examples are automatically produced articles that compare temporally sequential values, such as quarterly earnings, or that describe generalized gameplay, such as inning-by-inning descriptions of baseball games. The apparent novelty of the event-driven narratives over other template-driven outputs appears to result from at least two factors. First, variability in kinds of events appears to be far more unusual than variability in numbers or names, and second, the combination of variability in events with variability in the characters, entities and locations that participate in those events appears to provide an additional layer of complexity, beyond that typically achievable from a one-dimensional data record. The authors therefore hypothesize that readers may at least partly respond to the fact that these articles are event-driven narratives rather than merely descriptions. A formal assessment of the qualitative differences in reader experience between static descriptions (figure 1) and event-driven narratives (figure 2) is planned as future work, using comparative surveys that remove the possible difference in interest level between article subject matter.

The approach described here is not unique to pursuit stories, but is illustrative of an entire class of news stories that exhibit relatively stable patterns of events even if the specific activity of those events varies from story to story. Other domains that tend to exhibit this 'formulaic' characteristic include many crime stories, celebrity stories, stories about natural disasters, stories about business activity, many political stories, legal stories and probably stories in many other domains. The data necessary to support the automated writing of such stories could be gathered from a combination of automation and manual reporting, perhaps managed by news organizations for use in generating multiple forms of text output, interactive content and other data-centric news products.

Example 2: Local and State Government stories

The relative ease with which structured pursuit stories can be expressed as written text using WordSmith suggests that applying the technique to less predictable structured stories might also be possible. The domain of local and state government stories was selected for a second experiment, because the Structured Stories database already contained a selection of such stories, and because automating these stories would be particularly valuable.

Journalism about the day-to-day activities of local and state governments is in crisis. Newsrooms no longer have the resources to cover the routine, cumulative activities in legislatures and city councils, and the long-term and periodic nature of these stories makes them difficult for citizens to efficiently follow and understand. This crisis is occurring as political power is increasingly being exercised at the state and local level, and therefore more efficient approaches to state and local journalism are needed. Many of the procedural events for these stories is already available as structured data from sources such as TrackBill.com, and therefore the domain is particularly well-suited to early automation of complex stories.

Like the pursuit stories reported by NBC SoCal, local and state government stories have a particular cadence. Similar kinds of stories tend to happen again and again, but with different specific activity, different characters, at different locations and often in a different order. Unlike the pursuit stories, local and state government stories tend to be longer and more structurally complex. They have very distinct phases, and often branch and intersect in ways that pursuit stories never do. They also tend to vary much more from story to story than pursuit stories do, and they clearly would require WordSmith templates that were much too complex to be feasible for use in a realistic editorial environment. In examining a selection of structured local and state government stories it therefore quickly became apparent that the technique used on the pursuit stories would not be sufficient.

A series of trial-and-error sessions involving the manual manipulation of versions of structured local and state government stories resulted in two candidate approaches to handling the increased complexity. The first of these involved reducing the unit of language generation from the structured story to the structured event – i.e. developing separate WordSmith templates for each individual structured event within a structured local/state government story, based on the event frame used to instantiate the structured event. The idea was that the resulting text could then be assembled into a finished text article from those multiple outputs, based on the story-level structure from the database. This approach quickly failed, because although it was easy to come up with ways to generate coherent text event by event, it was extremely difficult to do so in a way that made the assembled text coherent and comprehensible. The barrier to this approach is therefore editorial, and it is difficult to see how text fragments could be authored within event-based 'micro templates' to enable grammatical coherence in the assembled text.

The second candidate approach was to break down each structured story into multiple 'story segments' based on the activity of portions of the story. The idea here was that each segment could be handled as a separate story in the same manner as the pursuit stories described above. Each would be matched to a specific WordSmith template to generate output text for the segment and it would then be relatively easy to assemble the finished complete text document from the relatively small number of segment text blocks. The diagram in Figure 4 illustrates this candidate approach applied to a structured

story containing three segments relating to the origin, passage and consequences of a piece of proposed legislation becoming law.

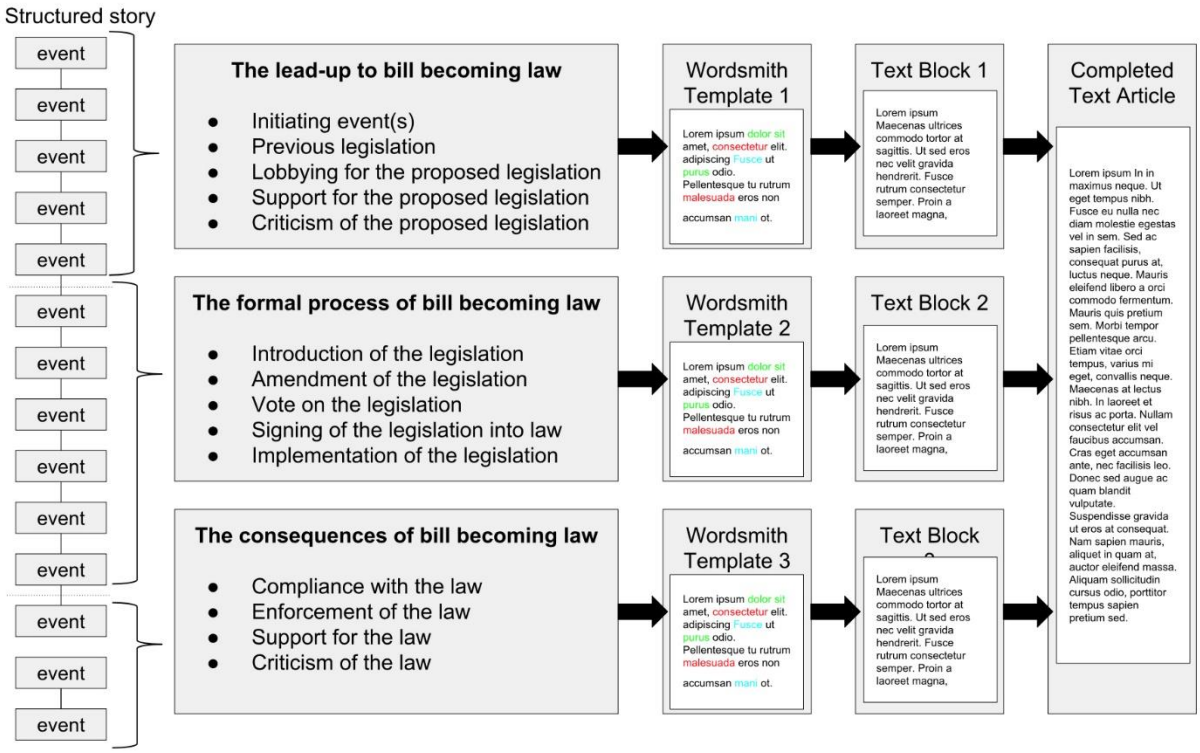

FIGURE 4
Example of a multi-segment approach to generating NLG text from relatively unpredictable structured stories

Again, each of these segments follows a process that is identical to that used for the pursuit stories and the only new challenges that arise here are finding a method to accurately split up the structured story into segments that match appropriate WordSmith templates, and addressing the editorial challenge of transitioning between segments in assembling the final text article.

This exploration of multi-segment processing of structured stories is still at an early stage. Although the process works for a narrow range of local/state government stories, the procedure for recognizing segments and matching them to templates is not yet automatable. More general use of the approach may require machine learning to achieve reliable performance, which would in turn require the collection of more training data. Nonetheless the approach appears to be possible, and the remaining challenges are well-defined. If it can be made routine then the implications could be profound. Such a process could, for example, enable automation to extend substantially beyond the formulaic stories described above in example 1, enabling the automated writing of longer, one-off stories that are more complex and less predictable. Such a process might also even eventually be applied to the automated expression of news stories as video, using techniques similar to those currently employed by the Wibbitz.com automated video production technology, which currently creates video based on text articles.

Evaluating Event-driven Narratives

Evaluating the quality of automated texts from NLG systems has been the subject of ongoing research since at least the late 1990's (Dale and Mellish 1998; Krahmer and Theune 2010). Initial attempts at evaluation focused on assessing the output of each of the technical sub-systems contributing to the completed text, e.g. the text planning, lexical selection and sentence planning sub-systems (Barr 2003). This approach was appropriate for evaluating these developer-operated early NLG systems, but is less useful for assessing modern commercial NLG systems in which templates are created editorially by writers. Evaluating the quality of texts produced by these newer systems is challenging because it is not clear what is being evaluated: the underlying data model, the quality of the underlying data, the template author's writing ability, the template author's logic or the consumers perception of the output text or of its origin.

Objective evaluation of the output of automated journalism by academic researchers has typically been performed using a 'blind taste test' methodology, in which consumers are shown various examples of automatically-generated and human-written texts and asked to rate each example according to various criteria, such as credibility and readability (Clerwall 2014; van der Kaa and Krahmer 2014; Graefe et al. 2016). As automated journalism becomes more common, however, human-versus-machine comparisons may become insufficient, because the knowledge that automated and human writing can be broadly comparable has now already been established. The development of more advanced forms of automated journalism – such as event-driven narratives – may therefore require a new evaluation framework that directly considers the structural and semantic complexity of output texts. Such a framework should provide feedback that facilitates technical improvements in journalistic data models and NLG systems, and that thereby helps to advance the sophistication of texts available from automated journalism.

A framework for evaluating the sophistication of machine-generated texts would likely require multiple components that together enable comparisons between output from different data models, different templates and different NLG systems. Some candidate components include:

- *Discourse categories*, identifying the type of discourse communicated in the text (for example description, narrative, explanation, exposition, etc.). Related work on discourse analysis in news includes Bell (1991) and Van Dijk (1988).
- *A story complexity metric*, measuring the degree to which the text exhibits the characteristics of narrative, such as episodic structure, setting, the presence of characters, etc. Related work on assessing narrative structure as a diagnostic tool has been done by Gillam et al. (1999), and similar metrics have been proposed for other forms of discourse, such as exposition.
- *A language complexity metric*, measuring the degree to which aspects of the language of the text exhibits structure and variety. Several methods of measuring language complexity exist, including lexical density (Gregori-Signesa and Clavel-Arroitiab 2015), idea density (Covington 2009) and various metrics used to assess student performance in writing and reading (Nelson et al. 2012).
- *A data complexity metric*, measuring the complexity of the underlying data model used to generate the output texts. Various formal measurement of the complexity of data models have been

proposed (Calero et al. 2001), and heuristic metrics, such as the dimensionality of data models, are in common use.

- *A consumer perception metric*, measuring consumer perception of the sophistication of communication exhibited in the text. This metric would be analogous to the descriptors of credibility and quality mentioned by Clerwall (2014).

Collectively, the availability of these, or similar, metrics would enable text-to-text comparison of the complexity and sophistication of machine-generated news. Individually, their availability would enable testing of various hypotheses about the differences between automated texts output from descriptive data (e.g. figure 1) and texts output from event-driven narrative data (e.g. figure 2). Such hypothesis tests might include:

- Whether the latter are of a different class of discourse (narrative) than the former (static description);
- Whether the latter exhibit greater story complexity than the former;
- Whether the latter exhibit greater language complexity than the former, specifically regarding the variety of verbs and the role of nouns within the activity of verbs;
- Whether the latter are based on underlying data of greater dimensionality than the former;
- Whether consumers of these news texts will perceive the latter to be more complex or sophisticated than the former.

By enabling comparisons between automated texts on the basis of sophistication, such a framework may provide incentives that impel the practice of automated journalism within newsrooms towards greater sophistication. Developing and demonstrating this proposed evaluation framework, using metrics with firm criteria based on accepted principles, is, however, a significant challenge and remains as future work.

Challenges and Limitations

It is clear that template-based NLG platforms such as WordSmith are already capable of supporting the automatic generation of event-based text narratives at scale. But it is also clear that there are numerous significant challenges to overcome before event-based approaches could be used in production editorial environments to automate journalism. These challenges can be loosely grouped into two categories: those relating to the availability of structured event and story data, and those relating to the editorial skills required to author the logic and text fragments of templates, and the text descriptions of activity associated with structured event frames.

From a technical perspective the ability to capture and store structured event and story data for journalism is not yet a completely solved problem, although it is close to solution. It is necessary, for example, to develop taxonomies of abstractions of journalistic events (such as the event frames used by Structured Stories) at granularities much finer than currently provided by FrameNet, the Events and Situations Ontology and schema.org. It is also necessary to find ways to simplify the acquisition of structured

event and story data, likely using a combination of ingestion of already-structured data (such as Track-Bill), the automated extraction of structured event data from text (as demonstrated by the NewsReader project), the use of sequentially guided user interfaces like that used by Structured Stories, and possibly new techniques such as controlled natural language interfaces.

From a usage perspective the development of technology, infrastructure and processes that ensure the availability of structured event and story data is a 'bootstrap' problem, where infrastructure and data must be accumulated in order to demonstrate value, attract investment and therefore develop further infrastructure and data. It is important to appreciate that automated journalism is just one use case of structured event and story platforms, and that other uses might drive their development - for example chat interfaces, smart content and business or government intelligence applications.

The editorial challenges associated with the automation of complex journalism may perhaps be more difficult to overcome than the technical challenges (see also Flew et al. 2012). For example it is not easy to author the text fragments and logic of NLG templates such that any variation of the input data produces quality text, and this difficulty is compounded when using event-based data values that are also text fragments, authored separately. Additional editorial challenges arise from authoring the transitions between blocks of text generated separately for assembly into a single final text document, and indeed from the logic by which text blocks are assembled into larger documents. Challenges relating to expressing narrative structure (i.e. the narrative network effect, the differential value of events and the semantic zoom effect) in a grammatically and structurally coherent NLG template are also not trivial. These, and other, editorial challenges suggest a need to develop new journalistic styles and grammatical vocabularies appropriate to automated journalism, and a need to develop editorial guidelines and style guides for text fragments, template authoring and document assembly.

Underlying these editorial challenges is the unavoidable requirement that journalists develop the 'computational thinking' skills necessary to work with journalistic information and with narrative as abstractions. Building and operating an automated journalism workflow based on structured events and computational narratives is an intrinsically abstract endeavor, requiring journalists to move back and forth between specific events and stories and the generalized patterns that those events and stories exhibit. Such skills are not primarily technical or technological. Given appropriate platforms and user interfaces, such as the WordSmith user interface, journalists should be able to work with structured events, narratives and narrative abstractions without any coding skills, database skills, mathematical skills, or other technical expertise. Instead, they need only develop and apply a familiarity with the patterns of journalistic information, as distinct from the particulars (Anderson 2014). This is not an easy skill to master, and would require substantial changes in the goals and curricula of journalistic education and in the incentives and expectations of journalists within newsrooms. Ironically, a primarily reason for journalists to learn to code may not be so that they can code, but so that they can learn to think about journalism in the same way that software developers think about code.

It is important to emphasize that the proportion of news stories that are susceptible to automation using these techniques, while far greater than is currently possible, is still quite limited. Only news events for which event frames can be defined can be captured as data, and only stories that exhibit some formulaic structure can be processed into a form useable by NLG systems. Also, the mere availability of data structures for news events and stories is insufficient for automation unless workflows for actually

reporting news as structured data are in place. A simple economic limitation also arises from the need for a minimum quantity of any one 'kind' of story in order to justify the investment in automating its production. These limitations, and others, ensure that manual writing will remain the only viable method for producing the most complex, impactful and valuable journalism for the foreseeable future. Nonetheless, as demonstrated in this article, it is indeed possible to encode many journalistic events and stories as data, and thereby automate the writing of news that is more complex than routine sports and finance reporting.

Conclusion

Much of the early discussion about automated journalism has focused on 'algorithms' – the step-by-step or statistical procedures by which information is translated from data into natural language text. This focus may be insufficient, however, because data models likely have much more influence than 'algorithms' in determining how much automation is possible, or desirable, in journalism. The full potential of automation in journalism should therefore be considered in terms of journalism-specific data models that can be intentionally designed and used to drive NLG text output, rather than merely in terms of what data models (and data) already happen to be available. Future research on automated journalism should not merely address questions about the quality of automated texts relative to human-written texts – NLG software is already of comparable quality (see Graefe et al. 2016) – but should also examine how journalists and media organizations can deliberately model the structured data used by NLG software to advance the sophistication of automated journalism. Similarly, ethical questions currently being asked about aspects of data in journalism, such as data origin, selection, and usage (Zion and Craig 2014), bias and objectivity (Gillespie 2014) and algorithmic transparency and accountability (Diakopoulos 2015) may need to be expanded to include questions about the data models used to capture that journalistic data.

The intentional design and use of novel data models for automated journalism, as described in this paper, requires precise definition of the nature of journalistic information. In particular, recognizing the distinction between static 'description' and dynamic event-driven 'storytelling' enables the conception of data models that can automate more sophisticated and engaging narration. These story-like data models, sometimes called 'event-centric knowledge graphs' or 'narrative networks', are already a reality. It is possible that similarly precise definition of other forms of journalistic information may enable further progress in automating complex journalistic writing, for example in the use computational representations of causality to automate 'explanation' journalism.

The experience of this project suggests that the editorial challenges of automating more sophisticated journalism will likely be greater than the technical challenges. It is important to recognize that it is unlikely that formal scientific models of the entirety of natural language will be available in the foreseeable future, and therefore that automated journalism will probably develop in the absence of computationally-accessible grammars that might permit the automated generation of journalistic text without human-generated templates. Automating journalism, from designing data models to authoring NLG templates, will therefore remain an almost entirely editorial activity, albeit at a much more abstract level than is seen in traditional editorial work. A new kind of 'meta-editorial' skill that combines the ability to

abstract the use of language with the ability to apply that abstraction to crafting language-based editorial products will be required.

Given the economic realities facing journalism, increased automation is probably essential to fulfilling journalism's societal mission and will therefore probably become a more common aspect of the practice of journalism. If that automation is to be comprehensive and engaging enough to be useful then it will likely need to be in story-like form, based on story-like data. Creating and formalizing the data models, workflows and editorial skills necessary to support such automation is therefore a worthy investment in economically sustainable journalism, and an exciting opportunity.

References

Abbot, Porter H. 2008. The Cambridge Introduction to Narrative. New York: Cambridge University Press.

Anderson, Chris W. 2012. "Towards a Sociology of Computational and Algorithmic Journalism." New Media and Society 7 (15): 1005–1021.

Anderson, Chris W. 2015. "Between the unique and the pattern: Historical tensions in our understanding of quantitative journalism." Digital Journalism 3 (3): 349–363.

Anderson, Chris W., Emily Bell, and Clay Shirky. 2012. Post-industrial Journalism: Adapting to the Present. New York, NY: Tow Center for Digital Journalism Columbia University.

Baker, Colin. 2008. "FrameNet, Present and Future." Proceedings of First International Conference on Global Interoperability for Language Resources.

Baldassano, Christopher, Janice Chen, Asieh Zadbood, Jonathan W. Pillow, Uri Hasson, and Kenneth A. Norman. 2016. "Discovering event structure in continuous narrative perception and memory." bioRxiv Website. Accessed January 2017, http://biorxiv.org/content/early/ 2016/10/14/081018.

Barr, Valerie. 2003. "A Proposed Model for Effective Verification of Natural Language Generation Systems." Proceedings of the Sixteenth International Florida Artificial Intelligence Research Society Conference, May 11-15, St. Augustine, USA.

Bell, Allan. 1991. The Language of News Media. Cambridge, MA: Wiley-Blackwell.

Boyd, Brian. 2009. On the Origin of Stories: Evolution, Cognition, and Fiction. Cambridge, Massachusetts: The Belknap Press of Harvard University Press.

Boyd, Danah, and Kate Crawford. 2012. "Critical Questions for Big Data: Provocations for a Cul- tural, Technological, and Scholarly Phenomenon." Information, Communication & Society 15 (5): 662–679.

Calero, Coral, Mario Piattini, and Marcela Genero. 2001. "Database Complexity Metrics." 4th Inter- national Conference on the Quality of Information and Communications Technology, March 12-14, Lisbon, Portugal.

Carlson, Matt. 2015. "The Robotic Reporter. Automated Journalism and the Redefinition of Labor, Compositional Forms, and Journalistic Authority." Digital Journalism 3 (3): 416–431.

Caswell, David 2015. "Structured Narratives as a Framework for Journalism." Proceedings of the Computational Models of Narrative workshop. http://narrative.csail.mit.edu/cmn15/ paper20.pdf.

Caswell, David, Frank Russell, and Bill Adair. 2015. "Editorial Aspects of Reporting into Structured Narratives." Computation + Journalism Symposium. http://cj2015.brown.columbia.edu/ papers/editorial.pdf.

Clerwall, Christer. 2014. "Enter the Robot Journalist. Users' Perceptions of Automated Content." Journalism Practice 8 (5): 519–531.

Coddington, Mark. 2015. "Clarifying Journalism's Quantitative Turn: A Typology for Evaluating Data Journalism, Computational Journalism, and Computer-Assisted Reporting." Digital Journalism 3 (3): 331–348.

Covington, Michael A. 2009. "Idea density – A potentially informative Characteristic of retrieved Documents." IEEE Southeastcon 2009, 5-8 March, Atlanta, USA.

Creech, Brian, and Andrew L. Mendelson. 2015. "Imagining the Journalist of the Future: Technological Visions of Journalism Education and Newswork." The Communication Review 18 (2): 142–165.

Dale, Rrobert, and Christopher Mellish. 1998. "Towards the Evaluation of Natural Language Gen- eration." Proceedings of the First International Conference on Evaluation of Natural Language Processing Systems, May 28-30, Granada, Spain.

Deemter, Kees Van, Emiel Krahmer, and Mariët Theune. 2005. "Real versus Template-Based Natural Language Generation: A False Opposition?" Computational Linguistics 31 (1): 15–24.

Deuze, Mark. 2005. "Towards professional participatory storytelling in journalism and advertising." FirstMonday 10 (7). doi:10.5210/fm.v10i7.1257.

Diakopoulos, Nicholas. 2015. "Algorithmic Accountability. Journalistic Investigation of Compu- tational Power Structures." Digital Journalism 3 (3): 398–415.

Diakopoulos, Nicholas, and Michael Koliska. 2016. "Algorithmic Transparency in the News Media." Digital Journalism. doi:10.1080/21670811.2016.1208053

Dörr, Konstantin. 2016. "Mapping the Field of Algorithmic Journalism." Digital Journalism 4 (6): 700–722.

Dörr, Konstantin, and Katharina Hollnbuchner. 2016. "Ethical Challenges of Algorithmic Journal- ism." Digital Journalism. doi:10.1080/21670811.2016.1167612

Flew, Terry, Christina Spurgeon, Anna Daniel, and Adam Swift. 2012. "The Promise of Computational Journalism." Journalism Practice 6 (2): 157–171.

Gillam, Ronald B., Elizabeth D. Pena, and Lynda Miller. 1999. "Dynamic Assessment of Narrative and Expository Discourse." Topics in Language Disorders 20 (1): 33–47.

Gillespie, Tarleton. 2014. "The Relevance of Algorithms." In Media Technologies. Paths Forward in Social Research, edited by Tarleton Gillespie, Pablo Boczkowski, and Kirsten Foot, 167–194. London: MIT Press.

Graefe, Andreas, Mario Haim, Bastian Haarmann, and Hans-Bernd Brosius. 2016. "Readers' Perception of Computer-Generated News: Credibility, Expertise, and Readability." Journalism. doi:10.1177/1464884916641269

Gregori-Signesa, Carmen, and Begoña Clavel-Arroitiab. 2015. "Analysing Lexical Density and Lexical Diversity in University Students' Written Discourse." 7th International Conference on Corpus Linguistics, 5-7 March 2015, Valladolid, Spain.

Hammond, Philip. 2015. "From Computer-Assisted to Data-Driven: Journalism and Big Data." Journalism. doi:10.1177/1464884915620205

Krahmer, Emiel, and Mariët Theune. 2010. Empirical Methods in Natural Language Generation: Data-oriented Methods and Empirical Evaluation. Berlin: Springer.

Krahmer, Emiel, and Hille van der Kaa. 2014. "Journalist versus news consumer: The perceived credibility of machine written news." Computation+Journalism Symposium, October 24- 25, Columbia University, New York, USA.

Latar, Noam Lemelshtrich. 2015. "The Robot Journalist in the Age of Social Physics: The End of Human Journalism?" In The Economics of Information, Communication, and Entertainment. The Impacts of Digital Technology in the 21st Century, edited by Gali Einav, 65–80. Wiesba- den: Springer.

Lewis, Seth C., and Oscar Westlund. 2015a. "Big Data and Journalism." Digital Journalism 3 (3): 447–466.

Lewis, Seth C., and Oscar Westlund. 2015b. "Actors, Actants, Audiences, and Activities in Cross- Media News Work." Digital Journalism 3 (1), 19–37.

Mani, Inderjeet. 2013. Computational Modeling of Narrative. San Rafael, CA: Morgan and Claypool.

McDonald, David D. 2010. "Natural Language Generation." In Handbook of Natural Language Pro- cessing, edited by Nitin Indurkhya, and Fred J. Damerau, 121–141. Boca Raton, FL: Chapman and Hall.

Montal, Tal, and Zvi Reich. 2016. "I, Robot. You, Journalist. Who is the Author?" Digital Journalism. doi:10.1080/21670811.2016.1209083

Nelson, Jessica, Charles Perfetti, David Liben, and Meredith Liben. 2012. Measures of Text Difficulty: Testing their Predictive Value for Grade Levels and Student Performance. Technical Report to the Gates Foundation.

Örnebring, Henrik, and Raul Ferrer Conill. 2016. "Outsourcing newswork." In The Sage Handbook of digital journalism, edited by Tamara Witschge, Chris W. Anderson, David Domingo, and Alfred Hermida, 207–221. London: SAGE.

Pavlik, John. 2000. "The Impact of Technology on Journalism." Journalism Studies 1 (2): 229–237.

Radvansky, Gabriel, and Jeffery Zacks. 2014. Event Cognition. New York, NY: Oxford University Press.

Reiter, Ehud. 2010. "Natural Language Generation." In The Handbook of Computational Linguistics and Natural Language Processing, edited by Alexander Clark, Chris Fox and Shalom Lappin, 574–598. Oxford: Wiley-Blackwell.

Reiter, Ehud, and Robert Dale. 1997. "Building applied natural language generation systems." Natural Language Engineering 3 (1): 57–87.

Rissen, Paul, Helen Lippell, Matt Chadburn, Tom Leitch, Dan Brinkley, Michael Smethurst, and Sebastian Cevey. 2013. "News Storyline Ontology, BBC Ontologies Website". Accessed January 2017, http://www.bbc.co.uk/ontologies/storyline.

Roeh, Itzhak. 1989. "Journalism as Storytelling, Coverage as Narrative." American Behavioral Scientist 33 (2): 162–168.

Rospocher, Marco, Marieke van Erp, Piek Vossen, Antske Fokkens, Itziar Aldabe, Gernam Rigau, Aitor Soroa, Thomas Ploeger and Tessel Bogaard. 2016. "Building event-centric Knowledge Graphs from News." Web Semantics: Science, Services and Agents on the World Wide Web 37-38: 132–151.

Saurwein, Florian, Natascha Just, and Michael Latzer. 2015. "Governance of Algorithms: Options and Limitations." Info 17 (6): 35–49.

Schema.org. 2017. Action Schemas Website. Accessed January 2017, https://schema.org/Action.

Segers, Roxane, Piek Vossen, Marco Rospocher, Luciano

Serafini, Egoitz Laparra, and German Rigau. 2015. "ESO: a frame based ontology for events and implied situations." Proceedings of maplex 2015, Yamagata, Japan.

Thurman, Neil, Konstantin Dörr, and Jessica Kunert. 2017. "When Reporters get Hands-on with Robo-Writing: Professionals Consider Automated Journalism's Capabilities and Consequences." Digital Journalism. doi:10.1080/21670811.2017.1289819

Van Dalen, Arjen. 2012. "The Algorithms behind the Headlines. How machine-written News redefines the Core Skills of Human Journalists." Journalism Practice 6 (5–6): 648–658.

Van Dijk, Teun A. 1988. News as Discourse. New Jersey, NJ: Lawrence Erlbaum Associates.

Wing, Jeannette M. 2006. "Computational thinking in journalism." Communications of the ACM 49 (3). http://www.cs.cmu.edu/afs/cs/usr/wing/www/publications/Wing06.pdf.

Workshop on Computing News Storylines, website accessed January 2016. https://sites.google.com/site/newsstorylines2016/.

Young, Mary Lynn, and Alfred Hermida, A. 2015. "From Mr. and Mrs. Outlier to Central Tendencies." Digital Journalism 3 (3): 381–397.

Zion, Lawrie, and David Craig. 2014. Ethics for Digital Journalists. Emerging Best Practices. New York: Routledge.