

A Benchmark of Visual Storytelling in Social Media

Gonçalo Marcelino
NOVALINCS
Uni. NOVA de Lisboa, Portugal
goncalo.bfm@gmail.com

David Semedo
NOVALINCS
Uni. NOVA de Lisboa, Portugal
df.semedo@campus.fct.unl.pt

André Mourão
NOVALINCS
Uni. NOVA de Lisboa, Portugal
a.mourao@campus.fct.unl.pt

Saverio Blasi
BBC Research and Development
London, UK
saverio.blasi@bbc.co.uk

Marta Mrak
BBC Research and Development
London, UK
marta.mrak@bbc.co.uk

João Magalhães
NOVALINCS
Uni. NOVA de Lisboa, Portugal
jm.magalhaes@fct.unl.pt

ABSTRACT

Media editors in the newsroom are constantly pressed to provide a "like-being there" coverage of live events. Social media provides a disorganised collection of images and videos that media professionals need to grasp before publishing their latest news updated. Automated news visual storyline editing with social media content can be very challenging, as it not only entails the task of finding the right content but also making sure that news content evolves coherently over time. To tackle these issues, this paper proposes a benchmark for assessing social media visual storylines. The SocialStories benchmark, comprised by total of 40 curated stories covering sports and cultural events, provides the experimental setup and introduces novel quantitative metrics to perform a rigorous evaluation of visual storytelling with social media data.

KEYWORDS

Storytelling, social media, benchmark

ACM Reference Format:

Gonçalo Marcelino, David Semedo, André Mourão, Saverio Blasi, Marta Mrak, and João Magalhães. 2019. A Benchmark of Visual Storytelling in Social Media. In *International Conference on Multimedia Retrieval (ICMR '19)*, June 10–13, 2019, Ottawa, ON, Canada. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3323873.3325047>

1 INTRODUCTION

Editorial coverage of events is often a challenging task, in that media professionals need to identify interesting stories, summarise each story, and illustrate the story episodes, in order to inform the public about how an *event* unfolded over time. Thanks to its widespread adoption, social media services offer a vast amount of available content, both textual and visual, and is therefore ideal to support the creation and illustration of these event stories [4–6, 12, 16].

The timeline of an event, e.g. a music festival, a sport tournament or a natural disaster [13], contains visual and textual pieces

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ICMR '19, June 10–13, 2019, Ottawa, ON, Canada

© 2019 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery.

ACM ISBN 978-1-4503-6765-3/19/06...\$15.00

<https://doi.org/10.1145/3323873.3325047>

of information that are strongly correlated. There are several ways of presenting the same event, by covering specific *storylines*, each offering different perspectives. These storylines, illustrated in Figure 1, refer to a story topic and related subtopics, and are structured into *story segments* that should describe narrow occurrences over the course of the event. More formally, we define a *Visual Storyline* as a sequence of segments, referring to an event topic, with each segment being defined by a textual description and comprising an image or a video.

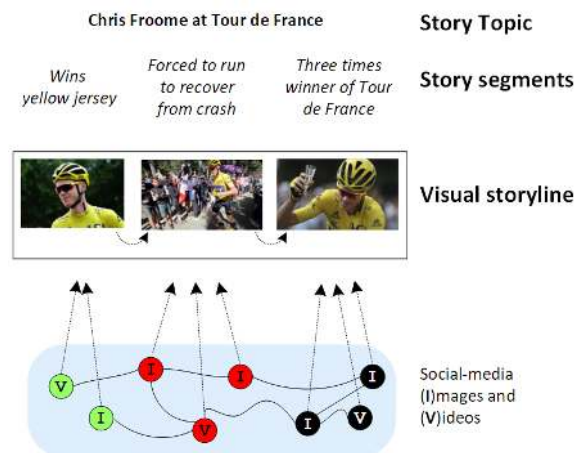


Figure 1: Visual storyline editing task: a news *story topic* and *story segments* can be illustrated by social media content.

In Figure 1 we illustrate the newsroom workflow tackled by this paper: once the *story topic* and *story segments* are created, the media editor selects images/videos from social media platforms and organise the retrieved content according to a coherent narrative, i.e. the *visual storyline*. Many social media platforms, including [Twitter](#), [Flickr](#) or [YouTube](#) provide a stream of multimodal social media content, naturally yielding an unfiltered event timeline. These timelines can be listened to, mined [2, 3, 6], and exploited to gather visual content, specifically image and video [11, 14].

The primary contribution of this paper is the introduction of a quality metric to assess visual storylines. This metric is designed to evaluate the quality of an automatically illustrated storyline, based on computational aspects that attempt to mimic the human-driven editorial perspective. The quality metric focuses on the *relevance*

Table 1: Dataset statistics for each event, including both terms and hashtags. The event and crawling time spans are shown.

| Event | Stories | Docs | Docs w/images | Docs w/videos | Crawling span | Crawling seeds |
|--------|---------|---------|--------------------------------|-------------------------------|---------------------------------------|---|
| EdFest | 20 | 82,348 | Twitter: 15439 Flickr: 5908 | Twitter: 3690 Youtube: 293 | From: 2016-07-01 Until: 2017-01-01 | Terms Edinburgh Festival, Edfest, Edinburgh Festival 2016, Edfest 2016 Hashtags #edfest, #edfringe, #EdinburghFestival, #edinburghfest |
| TDF | 20 | 325,074 | Twitter: 34865 Flickr: 6442 | Twitter: 8677 Youtube: 983 | From: 2016-06-01 Until: 2017-01-01 | Terms le tour de france, le tour de france 2016, tour de france Hashtags #TDF2016, #TDF |

of each individual segment's illustration and on the general flow of the storyline, i.e. the transitions from segment's illustration to segment's illustration. The second contribution of this paper is a social media dataset with news storylines to allow the research of visual storytelling with social media data. The benchmark, developed for TRECVID2018 [1], provides a rigorous setting to research the underpinnings of visual social-storytelling. The most relevant aspect of this benchmark is the realistic nature of the storylines, that mimic the newsroom media editorial process: some stories were manually investigated and inferred from social media and other stories were constructed from existing news articles.

2 SOCIAL STORIES BENCHMARK¹

Assessing the success of news visual storyline creation is a complex task. In this section, we address this task and propose the SocialStories benchmark. Visual storytelling datasets like [7] and [8] contain sequences of image-caption pairs, that capture a specific activity, e.g., "playing frisbee with a dog". A characteristic of these stories is that the sequence of visual elements is very coherent (visually and textually), which is highly unlikely to occur in social media. Hence, these stories do not match the ones a journalist or a media professional needs to create and illustrate on a daily basis. This highlights the importance of a suitable experimental test-bed for the task at hand.

The SocialStories benchmark provides the experimental setup and metrics to perform a rigorous evaluation of the task of creating visual storylines from social media data. The core aspects of the SocialStories benchmark are:

- **Storylines:** We created three types of storylines: news article, investigative topics, and review topics. These were either obtained from newswire articles, or created manually through data inspection.
- **Assessing Visual Storylines Quality:** Assessing the quality of a sequence of information is a novel and challenging task. We propose a new metric that assesses the quality of a visual storyline in terms of its relevance and transition between segment illustrations.

The following sections detail the types of storylines considered in the benchmark and proposes a story quality assessment metric.

2.1 SocialStories: Event Data and Storylines

To enable social media visual storyline illustration, a data collection strategy was designed to create a suitable corpora, limiting the number of retrieved documents to those posted during the span of the event. Events adequate for storytelling were selected, namely those with strong social-dynamics in terms of temporal variations

with respect to their semantics (textual vocabulary and visual content). In other words, the unfolding of the event stories is encoded in each collection. Events that span over multiple days like music festivals, sports competitions, etc., are examples of good storyline candidates. Taking the aforementioned aspects into account, the data for the following events was crawled (Table 1):

The Edinburgh Festival (EdFest) consists of a celebration of the performing arts, gathering dance, opera, music and theatre performers from all over the world. The event takes place in Edinburgh, Scotland and has a duration of 3 weeks in August.

Le Tour de France (TDF) is one of the main road cycling race competitions. The event takes place in France (16 days), Spain (1 day), Andorra (3 days) and Switzerland (3 days).

2.1.1 Crawling Strategy. Our *keyword-based* approach, consists of querying the social media APIs with a set of keyword terms. Thus, a curated list of keywords was manually selected for each event. Furthermore, hashtags in social media play the essential role of grouping similar content (e.g. content belonging to the same event) [9]. Therefore a set of relevant hashtags grouping content of the same topic was also manually defined. The data collected is detailed in Table 1. With no loss of generality, Twitter data was used for the experiments reported in the following sections.

2.1.2 Story Segments. For each event, newsworthy, informative and interesting topics are considered, containing diverse visual material either in terms of low-level visual aspects (colour, backgrounds, shapes, etc.) and/or semantic visual aspects. Each storyline contains 3 to 4 story segments.

2.2 Visual Storyline Quality Metric

Media editors are constantly judging the quality of news material to decide if it deserves being published. The task is highly skillful and deriving a methodology from such process is not straightforward. The task of identifying visual material suitable to describe each story segment is, from the perspective of media professionals, highly subjective. The motivation for why some content may be used to illustrate specific segments can derive from a variety of factors. While subjective preference obviously plays a part in this process (which cannot be replicated by an automated process), other factors are also important which come from common practice and general guidelines, and which can be mimicked by objective quality assessment metrics.

The first step towards the quantification of visual storyline quality concerns the human-judgement of these different dimensions. This is achieved in a sound manner by judging specific objective characteristics of the story – Figure 2 illustrates the visual storyline quality assessment framework. In particular, storyline illustrations

¹[HTTPS://NOVASEARCH.ORG/DATASETS/](https://novasearch.org/datasets/).

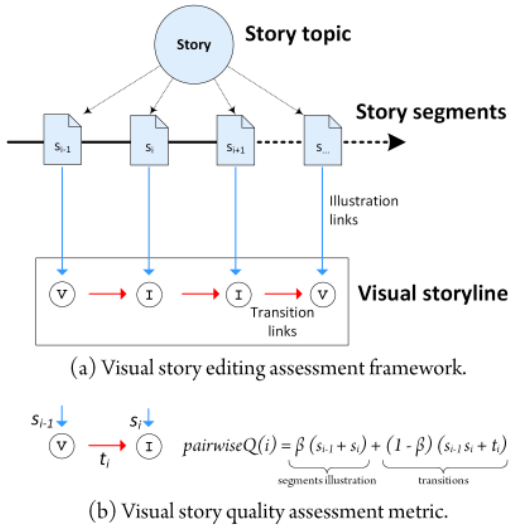


Figure 2: Benchmarking visual storytelling creation.

are assessed in terms of *relevance of illustrations* (blue links in Figure 2) and *coherence of transitions* (red links in Figure 2). Once a visual storyline is generated, annotators will judge the relevance of the *story segment illustration* as:

- $s_i=0$: the image/video is not relevant to the story segment;
- $s_i=1$: the image/video is relevant to the story segment;
- $s_i=2$: the image/video is highly relevant to the story segment.

Similarly with respect to the *coherence* of a visual storyline, each *story transition* is judged by annotators as the degree of affinity between pairs of story segment illustrations:

- $t_i=0$: there is no relation between the segment illustrations;
- $t_i=1$: there is a relation between the two segments;
- $t_i=2$: there is an appealing semantic and visual coherence between the two segment illustrations.

These two dimensions can be used to obtain an overall expression of the "quality" of a given illustration for a story of N segments. This is formalised by the expression:

$$Quality = \alpha \cdot s_1 + \frac{(1 - \alpha)}{2(N - 1)} \sum_{i=2}^N pairwiseQ(i) \quad (1)$$

The function $pairwiseQ(i)$ defines quantitatively the perceived quality of two neighbouring segment illustrations based on their relevance and transition:

$$pairwiseQ(i) = \underbrace{\beta \cdot (s_i + s_{i-1})}_{\text{segments illustration}} + \underbrace{(1 - \beta) \cdot (s_{i-1} + t_{i-1})}_{\text{transition}} \quad (2)$$

where α weights the importance of the first segment, and β weights the trade-off between *relevance of segment illustrations* and *coherence of transitions* towards the overall quality of the story.

Given the underlying subjectivity of the task, the values of α or β that optimally represents the human perception of visual stories, are in fact average values. Nevertheless, we posit the following two reasonable criteria: (i) illustrating with non-relevant elements ($s_i = 0$) completely breaks the story perception and should be

penalised. Thus, we consider values of $\beta > 0.5$; and (ii) the first image/video perceived is assumed to be more important, as it should grab the attention towards consuming the rest of the story. Thus, α is a boost to the first story segment s_1 . It was empirically found that $\alpha = 0.1$ and $\beta = 0.6$ adequately represent human perception of visual stories editing.

3 EVALUATION

3.1 Protocol and Ground-truth

Protocol. The goal of this experiment is demonstrate the robustness of the proposed benchmark. Target storylines and segments were obtained using several methods, resulting in a total of 40 generated storylines (20 for each event), each comprising 3 to 4 segments. Ground truth for both relevant segment illustrations, transitions and global story quality were obtained as described in the following section.

Ground-truth. Three annotators were presented with each story title, and asked to rate (i) each segment illustration as relevant or non-relevant, (ii) the transitions between each of the segments, and finally, (iii) the overall story quality. Stories were visualised and assessed in a specifically designed prototype interface. It presents media in a sequential manner to create the right *story* mindset to the user. Using the subjective assessment of the annotators, the score proposed in Section 2.2 was calculated for each story.

3.2 Quality Metric vs Human Judgement

In order to test the stability of the metric proposed to emulate the human's perception of visual stories quality, we resorted to crowd-sourcing. To do so, we computed the metric based on the relevance of segments and transitions between segments, and related it to the overall story rating assigned by annotators. Figure 3 compares the annotator rating to the quality metric. These values show that linear increments in the ratings provided by the annotators were matched by the metric. As can be seen, the relation is strong and relatively stable, which is a good indicator of the metric stability. Thus, these results show that the metric *Quality* effectively emulates the human perception of visual storyline quality.

3.3 Automatic Visual Storytelling

3.3.1 Results: Illustrations. Figure 4 (a) presents the influence of illustrations in the story *Quality* metric introduced in Section 2.2.

A Text Retrieval baseline (*BM25*) was the best performing baseline for both events. This shows the importance of considering the text of social media documents when choosing the images to illustrate the segments. Approaches based on social signals - (*#Retweets* and *#Duplicates*) - also attained good results. Particularly, when inspecting the storylines that result from using the *#Duplicates* baseline (Figure 4), an increase in aesthetic quality of the images selected for illustration can be noticed. This shows that social signals are a powerful indicator of the quality of social media content. However, in scenarios where relevant content is scarce, the approach is hindered by noise. This is especially problematic in cases where there are no strong social signals associated with the available content.

A visual concept detector baseline, based on a pre-trained VGG-16 [15] CNN, did not perform as expected. A *Concept Pool* method

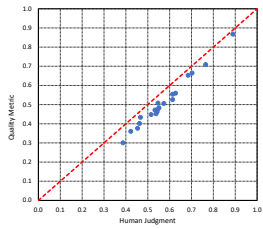


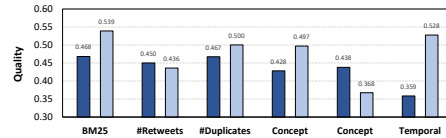
Figure 3: Correlation between human judgement and quality metric.

selects the image with the 10 most popular visual concepts. It failed often in correctly attributing concepts to the visual content of both datasets. As an example, in TDF, for segments featuring cyclists, unrelated concepts such as "unicycle", "bathing_cap" or "ballplayer" appeared very frequently. However, even though the extracted concepts lack precision, the concepts are extracted consistently among different pieces of content. This improves the performance of the method, since concepts are used to assess similarities in the content. The *Concept Query* is a pseudo-relevance feedback-based method that performed even worse. Hence, the performance of these baselines was lower than that of Text Retrieval. Nevertheless, they were able to effectively retrieve relevant content in situations where relevant content that could be found through text retrieval was scarce or non-existing.

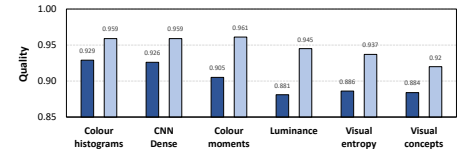
Finally, we tested a temporal smoothing method [10], *Temp. Modeling* to analyse the importance of temporal evidence. It performed the worst on the EdFest test set, while performing second best for the TDF test set. This high variation in performance is justified by the strong impact the type of story being illustrated has on the behaviour of this baseline. For story segments that take place over the course of the whole event, the approach is not well suited as the probabilities attributed to each segment illustration candidate are virtually identical, and therefore the baseline does not differentiate well between visual elements. However, in story segments with large variations on the amount of content posted per day, the approach provides a noticeably good performance.

Overall, as shown by the results presented in Figure 4, generating storylines for Tour de France stories is an easier task than doing so for the Edinburgh Festival stories. As a result 4 of the 6 baselines tested performed better in the task of illustrating Tour de France stories. This happens due to the heterogeneous nature of the media and stories associated with Edinburgh Festival (where very different activities happen) which makes the task of retrieving relevant content more difficult. Additionally, the availability of fewer images and videos for Edinburgh Festival accentuates this problem, as there is clearly less data to exploit when creating the storylines.

3.3.2 Results: Transitions. Figure 4(b) shows the performance of the proposed transitions baselines on the task of illustrating the EdFest and TDF storylines, using the story quality metric introduced in Section 2.2, calculated based on the judgements of the annotators. All baselines use the same pool of manually selected relevant visual content, for each segment.



(a) Illustrations quality



(b) Transitions quality.

Figure 4: Analysis of quality metric in terms of illustration relevance and transition consistency.

The *CNN Dense* baseline, minimises distance between representations extracted from the penultimate layer of the visual concept detector. This baseline provided the best performance, highlighting the importance of taking into account semantics when optimising the quality of transitions. Semantics may not be enough to evaluate the quality of transitions though. In fact, using single concepts as is the case for the *Visual Concepts* baseline provides very poor results, stressing the importance of taking into account multiple aspects.

The second and third best performing baseline focus on minimising the colour difference between images in a storyline: *Colour Histograms* and *Colour Moments*. This supports the assumption that illustrating storylines using content with similar colour palettes is a solid way to optimise the quality of visual storylines. Conversely, illustrating storylines by selecting images with similar degrees of entropy and luminance, using the *Visual Entropy* and *Luminance* baselines, provided worst results. Additionally, and similarly to what was observed while assessing the segment illustration baselines, Figure 4 shows that creating storylines with good transitions is easier for the TDF dataset than for the EdFest dataset.

4 CONCLUSIONS

This paper addressed the problem of automatic visual story editing using social media data that run in TRECVID2018. Media professionals are asked to cover large events and are required to manually process large amounts of social media data to create event plots and select appropriate pieces of content for each segment. We tackle the novel task of automating this process, enabling media professionals to take full advantage of social media content. The main contribution of this paper is a *benchmark to assess the overall quality of a visual story* based on the relevance of individual illustrations and transitions between consecutive segment illustrations. It was shown that the proposed experimental test-bed proved to be effective in the assessment of story editing and composition with social media material.

Acknowledgements. This work has been partially funded by the GoLocal CMU-Portugal project Ref. CMUP-ERI/TIC/0046/2014, by the COGNITUS H2020 ICT project No 687605 and by the project NOVA LINCS Ref. UID/CEC/04516/2013.

REFERENCES

- [1] George Awad, Asad Butt, Keith Curtis, Yooyoung Lee, Jonathan Fiscus, Afzal Godil, David Joy, Andrew Delgado, Alan F. Smeaton, Yvette Graham, Wessel Kraaij, Georges Quénot, Joao Magalhaes, David Semedo, and Saverio Blasi. 2018. TRECVID 2018: Benchmarking Video Activity Detection, Video Captioning and Matching, Video Storytelling Linking and Video Search. In *Proceedings of TRECVID 2018*. NIST, USA.
- [2] Deepayan Chakrabarti and Kunal Punera. 2011. Event Summarization Using Tweets. In *International AAAI Conference on Web and Social Media*.
- [3] Freddy Chong Tat Chua and Sitaram Asur. 2013. Automatic Summarization of Events from Social Media. In *ICWSM*, Emre Kiciman, Nicole B. Ellison, Bernie Hogan, Paul Resnick, and Ian Soboroff (Eds.). The AAAI Press.
- [4] Diogo Delgado, Joao Magalhaes, and Nuno Correia. 2010. Automated illustration of news stories. In *2010 IEEE Fourth International Conference on Semantic Computing*. IEEE, 73–78.
- [5] Erika Doggett and Alejandro Cantarero. 2016. Identifying Eyewitness News-worthy Events on Twitter. In *SocialNLP@EMNLP*.
- [6] Mengdie Hu, Shixia Liu, Furu Wei, Yingcai Wu, John Stasko, and Kwan-Liu Ma. 2012. Breaking News on Twitter. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*.
- [7] Ting-Hao K. Huang, Francis Ferraro, Nasrin Mostafazadeh, Ishan Misra, Jacob Devlin, Aishwarya Agrawal, Ross Girshick, Xiaodong He, Pushmeet Kohli, Dhruv Batra, et al. 2016. Visual Storytelling. In *15th Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL 2016)*.
- [8] Gunhee Kim, Seungwhan Moon, and Leonid Sigal. 2015. Ranking and retrieval of image sequences from multiple paragraph queries. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*. IEEE Computer Society, 1993–2001.
- [9] David Laniado and Peter Mika. 2010. Making Sense of Twitter. In *Proceedings of the 9th International Semantic Web Conference on The Semantic Web - Volume Part I (ISWC'10)*. Springer-Verlag, Berlin, Heidelberg, 470–485.
- [10] Flávio Martins, João Magalhães, and Jamie Callan. 2016. Barbara Made the News: Mining the Behavior of Crowds for Time-Aware Learning to Rank. In *Proceedings of the Ninth ACM International Conference on Web Search and Data Mining (WSDM '16)*. ACM, New York, NY, USA, 667–676. <https://doi.org/10.1145/2835776.2835825>
- [11] Philip J. McParlane, Andrew James McMinin, and Joemon M. Jose. 2014. "Picture the Scene...": Visually Summarising Social Media Events. In *ACM CIKM*.
- [12] Steve Paulussen and Pieter Ugille. 2008. User generated content in the news-room: Professional and organisational constraints on participatory journalism. *Westminster Papers in Communication & Culture* 5, 2 (2008).
- [13] Takeshi Sakaki, Makoto Okazaki, and Yutaka Matsuo. 2010. Earthquake Shakes Twitter Users: Real-time Event Detection by Social Sensors. In *Proceedings of the 19th International Conference on World Wide Web (WWW '10)*. 10.
- [14] Manos Schinas, Symeon Papadopoulos, Yiannis Kompatsiaris, and Pericles A. Mitkas. 2015. Visual Event Summarization on Social Media using Topic Modelling and Graph-based Ranking Algorithms. In *Proceedings of the 5th ACM on International Conference on Multimedia Retrieval - ICMR '15*.
- [15] Karen Simonyan and Andrew Zisserman. 2014. Very Deep Convolutional Networks for Large-Scale Image Recognition. *arXiv preprint arXiv:1409.1556* (2014).
- [16] Peter Tolmie, Rob Procter, Dave W. Randall, Mark Rouncefield, Christian Burger, Geraldine Wong Sak Hoi, Arkaitz Zubiaga, and Maria Liakata. 2017. Supporting the Use of User Generated Content in Journalistic Practice. In *CHI*.