Assisting News Media Editors with Cohesive Visual Storylines

Gonçalo Marcelino
Universidade NOVA de Lisboa
NOVA LINC3, Portugal
goncalo.bfm@gmail.com

David Semedo
Universidade NOVA de Lisboa
NOVA LINC3, Portugal
df.semedo@fct.unl.pt

André Mourão
Universidade NOVA de Lisboa
NOVA LINC3, Portugal
a.mourao@campus.fct.unl.pt

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

João Magalhães
Universidade NOVA de Lisboa
NOVA LINC3, Portugal
jmag@fct.unl.pt

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

ABSTRACT

Creating a cohesive, high-quality, relevant, media story is a challenge that news media editors face on a daily basis. This challenge is aggravated by the flood of highly-relevant information that is constantly pouring onto the newsroom. To assist news media editors in this daunting task, this paper proposes a framework to organize news content into cohesive, high-quality, relevant visual storylines. First, we formalize, in a nonsubjective manner, the concept of visual story transition. Leveraging it, we propose four graph based methods of storyline creation, aiming for global story cohesiveness. These were created and implemented to take full advantage of existing graph algorithms, ensuring their correctness and good computational performance. They leverage a strong ensemble-based estimator which was trained to predict story transition quality based on both the semantic and visual features present in the pair of images under scrutiny. A user study covered a total of 28 curated stories about sports and cultural events. Experiments showed that (i) visual transitions in storylines can be learned with a quality above 90%, and (ii) the proposed graph methods can produce cohesive storylines with a quality in the range of 88% to 96%.

CCS CONCEPTS

• Computing methodologies → Visual content-based indexing and retrieval;

KEYWORDS

News Media Editors, Visual Storylines, Graph-based News Illustration

ACM Reference Format:

Gonçalo Marcelino, David Semedo, André Mourão, Saverio Blasi, João Magalhães, and Marta Mrak. 2021. Assisting News Media Editors with Cohesive Visual Storylines. In *Proceedings of the 29th ACM International Conference on Multimedia (MM '21)*, October 20–24, 2021, Virtual Event, China. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3474085.3475476>

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.

MM '21, October 20–24, 2021, Virtual Event, China

© 2021 Copyright held by the owner/author(s). Publication rights licensed to the Association for Computing Machinery.

ACM ISBN 978-1-4503-8651-7/21/10...\$15.00

<https://doi.org/10.1145/3474085.3475476>

1 INTRODUCTION

Media editors are constantly judging the quality of news material to decide if it deserves being published. The task is highly skilful 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 complex. The motivation for why some content may be used to illustrate specific segments can derive from a variety of factors. 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 [5, 9, 24, 31]. Unfortunately though, exploiting this information is not easy, in that media professionals are faced with the daunting task of needing to manually search over a massive amount of data [4], which can be resource expensive, time consuming and lead to unsatisfactory results. Finding ways to support or automate this process would, therefore, be highly beneficial. This is not an easy task because such systems would need to (1) find appropriate information to illustrate a news timeline using camera crew and newsroom specific social media tools, and (2) make sense (from an editorial perspective) of the timeline of events across different textual and/or visual documents.

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

These timelines, illustrated in Figure 1, refer to a story topic and are structured into *story segments* that should describe narrow occurrences over the course of the event. In practice, **the task of the media editor is to illustrate each story segment by selecting the most cohesive set images/videos from a set of candidate**

images/videos. In this setting, we propose a graph-based framework to organize content into a *cohesive visual storyline*. The first contribution tackles a non-trivial task: mimicking the human perception of *visual story transition quality*. To do so, we propose a novel formal definition of transition between two media documents. Leveraging this definition we are then able to study the impact of a large set of semantic and visual criteria in the quality of image transitions. This proposal is utterly important as it allows predicting the perceived transition quality of pairs of images by analyzing their visual and semantic characteristics. The second contribution moves beyond individual story transitions and looks at the storyline as a whole. We propose a set of multimodal graph based methods that, given a story with N segments, and a list of sets of candidate multimodal documents to illustrate each segment, produce visual storylines composed of images.

In summary, the proposed framework targets the media professional who needs to find content to illustrate a news story. Our framework works as follow: (1) the media editor submits a story topic and its story segments as a query; (2) the framework computes a sequence of story illustrations in a cohesive way; and (3) this rank of story illustrations is then presented to the media editor for selection and post-editing. Hence, the gain is obvious: instead of struggling with a wide range of content permutations to illustrate story segments, the media editor can now browse different sequences of story illustrations and work from there.

2 RELATED WORK

From our discussions with media professionals, our goal was to design algorithms that could assist them in illustrating event plots. In this context, one of the few works that provide insight into the task is [2]. In it, the authors observe that presenting repeated images/videos in the context of news related storylines, makes viewers perceive said storylines as having less quality, even if relevant information is being shown. Strong transitions [11, 25] to cause a laugh, surprise, fear, etc., can be introduced by editors, who always have the final decision, by modifying candidate illustrations.

Collaborative Timelines. A highly appealing characteristic of social media platforms is that they make available large amounts of real-time data (news updates, opinions and first-hand reports) about specific events [20] in a streaming fashion, including important temporal clues about shared content [28, 29]. Hence, the synchronization and organisation of collaborative media along the temporal dimension leads to the generation of event timelines. TwitInfo [18] is a system that displays information about an event, sub-divided into several sub-events identified by automatically-detected peaks. Peak detection is also employed by Nichols et al. [23] to detect and summarise important updates to an event, and by Lin et al. [14] to build summaries in the form of storylines. Wang et al. [33] propose to generate summaries and timelines for continuous social media streams with an online clustering algorithm, maintaining distilled statistics in a data structure later used for summarization. Semedo and Magalhães [28] learn a temporal cross-modal embedding space, enabling the creation of multimodal summaries by navigating over that space. Freddy Chong, Tat Chua [6] propose a topic model that explores the correlation among posts published around the same time, extracting topics that identify different episodes of the same

event. The event is summarised by displaying the most representative post for each episode.

To enrich summaries with visual information, media needs to be processed and organized based on its content. Graph-based approaches [12, 13, 26, 35] are among the most popular methods for summarising collaborative visual streams. Yang et al. [35] focused on the alignment of photo streams in a common timeline using a bipartite kernel sparse representation graph. A master stream, corresponding to the summary, is obtained by removing redundant photos. The construction of large-scale storyline graphs is the central challenge addressed by Kim and Xing [13]. The optimisation framework infers a time-varying directed storyline graph that supports multiple storyline branches. Kim et al. [12] propose to use videos to model the sequence of actions that are more likely to occur on a given storyline. These sequences are captured on a sparse graph to later help the temporal grounding of images and the reconstruction of the timeline. Schinas et al. [26] propose to use a graph-based approach to detect topics across a dataset of microblog messages organised in a multi-graph, followed by ranking the images embedded in those messages according to relevance and diversity, to produce a visual summary.

Qualitative attributes in photos and films. Solving the problem of information relevance is hardly enough, in the creation of visual summaries that matches user expectations. Past work in visual quality [15], and more specifically in visual aesthetics [3] and interestingness [8], have proven that subjective visual attributes can computationally modelled and be applied to improve user experience. In illustrating visual storylines, visual memorability [10] plays an important role in the narrative of the story. For example, it has been observed [2] that repeating images/videos in a story is not appreciated, even if relevant information is being shown.

In [23], Nichols et. al. propose a methodology to perform automated text summarization of sporting events using data collected from Twitter. In this case, the authors find the most important moments of an event by identifying spikes in the volume of tweets over time, an approach also applied in TwitInfo [19]. After acquiring these tweets and applying spam removal techniques to the set, a phrase graph is created from the text present in the tweets. This graph details the chance of a word appearing next to a previously established sequence of words in the context of a phrase. The phrase graph is used to generate possible sentences that summarize the event. These sentences are also scored through the phrase graph and the best scoring ones are outputted. As a possible alternative to this graph based algorithm, the authors propose Sharifi's modified TF.IDF [30] as a method for generating a text summary of the event. Although not directly related to visual storyline generation, these works highlight the importance of considering the amount of content being posted to social media, at any particular time, as metric for finding interesting and important events.

Three years after Nichols et. al. published [23], the authors of [26], Schinas et. al., tackled a similar the problem, by rooting their research in a similar approach. In this work the task of visual event summarization using social media content, again from Twitter, is tackled through the use of topic modelling and graph based algorithms. The authors start by filtering a stream of tweets from a specific event in order to obtain only the most informative ones. This filtering process takes into account image size, image type,

text size, text morphology through part-of-speech tagging, among other criteria, leading to the creation of a multigraph. Re-posts are then discarded and tweets with duplicated images are clustered with the help of the graph. The tweets in these clusters are removed from the graph and are replaced by a single node that encompasses the information of the removed tweets. The authors then intent on discovering which tweets are part of the targeted event. In order to do this, they apply the SCAN [34] algorithm to the graph, which returns a set of subgraphs, each corresponding to a different topic. Finally, for each subgraph, the authors extract the respective images and rank them according to their popularity, relevance to the topic and the amount of information they provide.

In [21] the authors also tackle the problem of visual summarization, although in this case the summarization is done with images not only from Twitter but also from other websites found in the URLs of tweets. After filtering the obtained images in a manner similar to the one proposed in [27], the authors rank the remaining images using social signals. Finally, besides taking into account the popularity of an image the authors also factor into its rank the diversity of said image when compared to the remaining available ones. Of note is that both authors of [27] and [21] tackle the problem of filtering duplicated content, although through different methodologies. In order to reconcile both approaches one could use pHash as proposed by [21] and cluster the results through a graph based approach in a way similar to what is described in [27].

3 PRELIMINARIES

A visual storyline must be succinct and cohesive but also pleasing to the viewer. Overall it must present a sequence of images/videos as an interesting and relevant narrative. Before building the story graph, we first need to retrieve content that is relevant to the different story segments. In practice, a media editor defines a story composed of N segments, $Story_N = (u_1, u_2, \dots, u_N)$, that need to be illustrated (u_i). The segments ($u_1, u_2, \dots, u_N$) are then used to query the set of media documents, D , and the BM25 retrieval method produces a list of sets of candidate images/videos ($C_1, C_2, \dots, C_N$), where $C_i = \{a_1, \dots, a_k\}$ contains candidate images a_j to illustrate the story segment u_i .

4 MULTIMODAL STORY GRAPHS

We now tackle the task of automatic visual storyline illustration. This method takes as input, sets of candidate images C_i to illustrate each segment u_i of a story and is tasked with outputting visually and semantically cohesive storylines composed of images in the said sets. We propose two different multimodal graph based approaches to tackle this problem, each with two variants.

4.1 Scoring Multimodal Story Transitions

Rating a transition between a pair of images according to its quality is a non-linear process that results from the interpretation of the features of the individual images and of the manner in which they interact. To tackle the task of optimizing the transition quality between two pairs of images we first need a computationally valid approach to describe the concept of transition. From a non-computational, professional perspective, literature characterizes transitions based on the relations between semantic and visual

Feature Name (p)	$distance_p(f_1, f_2)$	$feature_p(a)$
Luminance	$abs(f_1 - f_2)$	A positive real value representing the luminance.
Color histogram	$\sum abs(f_1 - f_2)$	A 3D color histogram with 16 bins per RGB channel converted to CIELAB color space.
Color moment	$euclidean(f_1, f_2)$	A vector representing the first color moment of the image in CIELAB color space.
Color correlogram	$\sum abs(f_1 - f_2)$	A 16 bins 3D color correlogram in CIELAB color space.
Entropy	$abs(f_1 - f_2)$	A positive real value representing the entropy of the image.
#Edges	$\sum abs(f_1 - f_2)$	A vector containing the number of horizontal, vertical and diagonal edges.
pHash	$hamming(f_1, f_2)$	A pHash vector.
Concepts	$\#(f_1 \cap f_2)$	A set of image concepts extracted using VGG16.
CNN Dense	$euclidean(f_1, f_2)$	The embeddings extracted from the last layer of the ResNet CNN.
Environment	$f_1 = f_2$	Either "outdoors" or "indoors".
Scene category	$\#(f_1 \cap f_2)$	The location depicted in an image described through labels (e.g.: "bridge", "forest path", "skyscraper", etc.).
Scene attributes	$\#(f_1 \cap f_2)$	The attributes of the location depicted in an image described through labels (e.g.: "man-made", "open area", "natural light", etc.).

Table 1: The multimodal set of features I_f , corresponding distance functions and descriptions.

characteristics of the images that compose them and by the ways in which these images interact. We emulate this approach proposing a novel formalization of transition based on the concept of a distance ensemble.

To tackle the automation of this process we resort to Gradient Boosted Regression Trees [7, 22], and define the problem as one of predicting a rating given the set of *transition distances* between a given pair. Aiming to build a robust model, we propose the use of large set of features to compose the transitions' distances between image pairs. More specifically, we define the *story transition quality estimator* as the Gradient Boosted Regressor (GBR)

$$transQuality(a, b) \in [0, 1], \quad (1)$$

for a pair of any two images a and b . In practice the GBR needs to estimate a transition quality score between images a and b , from a set of *transition distances* $D = \{\forall c \in I_f, dist_c(feac_c(a), feac_c(b))\}$, where I_f is the set of features under consideration. Hence, a transition between two images is formalized as an ensemble of regression trees over a set of distances between the features of said images. These features I_f are presented in Table 1. As training data we considered the transition quality ground truth provided by human annotators.

4.2 MaxTransitions: Bi-partite graph

This approach optimizes for storylines with the best possible sequential transitions. In this context, the transition quality is only measured between the candidate images of consecutive segments.

We follow a graph based approach to tackle this problem. Specifically, we define $G = (V, E)$, as a *sequence of bipartite weighted directed graphs*, where given a story $Story_N = (u_1, u_2, \dots, u_N)$ of N segments, the graph G is constructed as follows:

- (1) **Vertices:** the graph's vertices V correspond to all the candidate images in the sets $(C_1, C_2, \dots, C_N)$, of k images each. Each set $C_i = \{a_1, \dots, a_k\}$, corresponds to a story segment u_i . Hence, each candidate image a_* of candidate set C_i becomes a vertex v_i in the graph;
- (2) **Edges:** the graph's edges E , associate all the candidate images from neighbouring candidate sets. In other words, all vertices in set C_i are fully connected and directed to vertices in set C_{i+1} . Hence, the bipartite property of graph;
- (3) **Edges-weight:** the weight associated with each edge $e \in E$, connecting two vertices v_1 and v_2 is given by a function $pairCost(v_1, v_2)$

Leveraging this graph structure, exemplified in Figure 2a, for a $N = 4$ segment story, we propose two different story transition methods to extract the visual storyline.

4.2.1 Shortest-Path Solution. The first approach, regarded as *Shortest Path without relevance (MaxTransition)*, is optimized for creating storylines with the best possible sequential transitions, regardless of the relevance of the candidate images to the segments they attempt to illustrate. It is designed to present added value in situations where most or all candidate images are already highly relevant to their respective segments or where content relevance is not the highest priority in the context of the story being illustrated.

Formally, this first proposed method computes the shortest path of size N in the graph (i.e. the path will contain a number of vertices equal to the number of segments in the story being illustrated). Hence, we aim to minimize the following expression:

$$\min_{v_1 \in C_1, v_2 \in C_2, \dots, v_N \in C_N} \sum_{i=1}^{N-1} pairCost(v_i, v_{i+1}) \quad (2)$$

where $pairCost(v_x, v_y) = transQ(v_x, v_y)$ and

$$transQ(v_x, v_y) = 1 - transQuality(v_x, v_y) \quad (3)$$

the function $transQuality$ in equation 1.

The resulting storyline from this approach is the one composed by the images represented by the vertices that minimize expression 2. In practice we resort to a variation of Dijkstra's minimum cost path algorithm to solve this problem.

4.2.2 Weighted Shortest-Path. The previous alternative considers only transition quality when generating visual storylines. For situations where relevant content is scarce, we propose a second approach, *Sequential with relevance (MaxTransitionsRel)*. In such situations, some of the available candidate images might not be relevant to the segment they attempt to illustrate. Hence, we leverage both transition quality and relevance of the candidate images, encouraging the creation of storylines with the most relevant candidate images that also present quality transitions.

Figure 2: Media is represented by the vertices of the graph, each vertex belonging to a candidate set C_i . The cost associated with an edge directed from vertices v_x to v_y is given by the $pairCost(v_x, v_y)$ function. The bi-partite graph connects all media in one segment to all media in the next segment – the shortest path finds the solution with maximal similarity between neighboured segments. The Multipartite graph connects all media in one segment to all media in the other segments – the maximal clique finds the solution with maximal similarity between all documents.

Thus, for this approach, we again aim to find a path composed of N vertices in the graph. However, this time, the expression to minimize weighs both the importance of the first segment in the storyline being relevant, as well as the importance of transitions vs. relevance, to score overall storyline quality. We do so by basing ourselves in the transition quality estimator proposed in section 4.1. In practice, we attain this path by minimizing the following expression:

$$\min_{v_1 \in C_1, v_2 \in C_2, \dots, v_N \in C_N} \beta \cdot \frac{1}{2(N-1)} \cdot \sum_{i=1}^{N-1} pairCost(v_i, v_{i+1}) + (1 - \beta) \cdot relC(v_1) \quad (4)$$

where β is used to control the importance of the first segment, and the function $pairCost(v_x, v_y)$, that ranges from 0 to 2, is defined as:

$$pairCost(v_x, v_y) = \underbrace{(1 - \gamma) \cdot (relC(v_x) + relC(v_y))}_{\text{segments illustration}} + \underbrace{\gamma \cdot (relC(v_x) \cdot relC(v_y) + transQ(v_x, v_y))}_{\text{transition}} \quad (5)$$

capturing segment illustration relevance in the first term, and accounting for relevance and quality of the transition, in the second term. γ controls this trade-off. Here, $relC(v) = 1 - rel(c)$. In turn, $rel(c)$ is the normalized relevance of the candidate image represented by vertex v to the text segment u_i to be illustrated, as calculated through the BM25 retrieval method.

The resulting storyline is the one composed by the images represented by the vertices that minimize expression 4.

4.3 MaxCohesion: Multipartite graph

The *MaxTransition* approach aims to produce storylines with high transition quality for sequential pairs of images. However, a visual storyline is consumed as a whole by its viewers, not as a disconnected set of pairs. Consequently we posit a second approach, *Fully connected*, designed to ensure maximal cohesion between all elements of the generated visual storylines, leveraging the possibility that individual transition quality is affected by the remaining elements of the storyline they are part of. Hence, we need to compute the transitions between all elements of the storyline to ensure maximal cohesion.

Again, we follow a graph based approach to tackle this problem. We define $G = (V, E)$, a *N-partite weighted graph*, where given a story $Story_N = (u_1, u_2, \dots, u_N)$ of N segments, the graph G is constructed as follow:

- (1) **Vertices:** the graph's vertices V correspond to all the candidate images in the sets $(C_1, C_2, \dots, C_N)$, of k images each. Each set $C_i = \{a_1, \dots, a_k\}$, corresponds to a story segment u_i . Hence, each candidate image a_* of candidate set C_i becomes a vertex v_i in the graph;
- (2) **Edges:** the graph's edges E , associate all the candidate images from *all other candidate sets*. In other words, all vertices in set C_i are connected to vertices in all the candidate sets except C_i . Hence, the multipartite property of the graph;
- (3) **Edges-weight:** the weight associated with each edge $e \in E$, connecting two vertices v_1 and v_2 is given by a function $pairCost(v_1, v_2)$

This graph construct, exemplified in Figure 2b for a $N = 4$ segment story, is well suited to extract storylines with maximal cohesion as we are going to see next.

4.3.1 Maximal-Clique Solution. First, we optimize for maximal cohesion, i.e. for maximal similarity among all story illustrations. Hence, for this first *MaxCohesion* approach, we compute the maximal weighted clique containing N vertices of graph G . (i.e. the clique will contain a number of vertices equal to the number of segments in the story being illustrated). Additionally, we pose the following restriction to the clique: it can only contain one vertex per candidate set. The bottom graph of Figure 2b provides an example of such a clique.

We attain this clique by minimizing the following expression:

$$\min_{v_1 \in C_1, v_2 \in C_2, \dots, v_N \in C_N} \sum_{i=1}^{N-1} \sum_{k=i+1}^N pairCost(v_i, v_k) \quad (6)$$

The resulting storyline from this approach is the one composed by the images represented by the vertices that minimize expression 6. We refer to this approach as *maximal clique without relevance* (*MaxCohesion*).

4.3.2 Weighted Maximum-Clique Solution. This final alternative *MaxCohesionRel* builds on the previous one, ensuring high transition quality between all pairs of images in the resulting visual storylines, not just sequential pairs, while optimizing for relevance as with the *MaxTransitionsRel* approach.

To do so, we compute a weighted clique containing N vertices of graph G , again with the restriction that it can only contain one vertex per set of candidate images. To find this clique we minimize

the following expression.

$$\min_{v_1 \in C_1, v_2 \in C_2, \dots, v_N \in C_N} \beta \cdot \frac{1}{N(N-1)} \cdot \sum_{i=1}^{N-1} \sum_{k=i+1}^N pairCost(v_i, v_k) + (1 - \beta) \cdot relC(v_1) \quad (7)$$

where the function $pairCost(v_x, v_y)$, a function that ranges from 0 to 2, is defined in Eq 5. The resulting storyline from this approach is the one composed by the images represented by the vertices that minimize equation 7. We refer to this final method as *Fully connected with relevance* (*MaxCohesionRel*).

5 EVALUATION

5.1 Datasets

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). News story topics and their segments were defined collaboratively by media editors and the authors of this paper. These stories were related to events usually covered in the newsroom. News and transitions were evaluated by newsreaders. As annotators, we used 5 newsreaders with ages in the range of 23 to 29 years old [17].

5.1.1 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.

5.1.2 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).

5.1.3 Data Filtering and Curation. Having gathered all tweets relevant to the topic and posted within an event's time frame, we proceed to filter documents unsuitable for summarization, according to several quality criteria [16]:

Retweets. Since retweets consist of a copy of an original tweet, they are redundant, thus they are excluded.

Language filtering. The Twitter language filter was already used, i.e., only tweets with text that Twitter deemed as English were collected. A second language filter is applied to exclude Non-English tweets using the LangDetect library¹.

SPAM In social-media platforms like Twitter, spam is abundant and a serious issue. Thus, we apply a set of filtering methods in an attempt to clean the collected data. We follow the approach in [20] and excluded tweets containing more than 3 hashtags, more than 3 mentions or more than 2 URLs. We adopt these empirical values which were chosen by the authors due to [1].

Visual SPAM. To detect synthetic and captioned visual elements, we train a logistic regression classifier (with L1 penalty and $\lambda = 1.0$) that predicts if an image is a photograph or not. Features include luminance and edge histograms of each image (with edges grouped into vertical, horizontal, 45 and 135 degrees orientations). We employ 5-fold cross-validation

¹<https://pypi.python.org/pypi/langdetect>

Figure 3: Story transition confusion matrices for different of features. On each matrix, the y-axis regards the score provided by the individual annotators and the x-axis regards the score attained through majority vote. The diagonal is clearly visible for all features, hence, showing a good user agreement.

and train the model with the dataset from [32] plus extra data identified by us. Additionally, we perform OCR to detect and exclude visual elements containing text using the PyTesseract² library. See [16] for details.

5.1.4 Story Transitions ground truth. 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. In practice, the 320 distinct visual (232 composed of 4 segments and the remaining 88 composed of 3 segments) were presented to annotators and were asked to rate (i) each segment illustration as relevant or non-relevant, (ii) the transitions between each of the segments in a Likert-scale of 1 to 5, and finally, (iii) the overall story quality, also in Likert scale. Stories were visualized and assessed in a specifically designed prototype interface. It presents media in a sequential manner to create the right *story* mindset to the user. By taking advantage of the annotations made to the 320 storylines we attained ground truth for a total of 872 pairs of images regarding transition quality, by averaging the scores (0 or 1) provided by the annotators to each pair.

5.2 User Study: Multimodal Story Transitions

To understand which individual features have more impact in the story transition quality, we conducted a user study. The protocol is as follows. To ensure that no irrelevant content would interfere with the assessment of story transitions, we looked at the ground truth and selected a total of 1572 relevant images to illustrate the segments of the 40 stories related to the EdFest and TDF datasets and then sorted them by their BM-25 score. This corresponds to an average of 10 relevant candidates for each story segment. Afterwards, each story was illustrated by considering the 10 relevant images per segment, and selecting the segment's illustration sequence that minimizes the sum of the pairwise *transition distances*,

²<https://github.com/madmaze/pytesseract>

Figure 4: Storyline transitions for two different stories. On the top story example, the transitions of the storyline created with the *Color histogram* baseline obtained an average score of 1 while the ones in the storyline created by the *Entropy* baseline obtained an average score of 0.6. On the bottom story example, the transitions of the storyline created with the *CNN Dense* baseline obtained an average score of 0.93 while the ones in the storyline created by the *Luminance* baseline obtained an average score of 0.87.

composed only of a single feature, between sequential images. Each baseline focuses on creating storylines with the following features individually: *Luminance*, *Color histogram*, *Color moment*, *Entropy*, *#Edges*, *pHash*, *Concepts*, *CNN Dense*, see Figure 4.

Figure 3 and Figure 4 show the performance of the proposed transitions baselines on the task of illustrating the EdFest and TDF storylines. The performance was measured by averaging the sum of the scores given by the annotators to each pairwise transition. All baselines use the same pool of manually selected relevant visual content, for each segment.

In this experiment the *CNN Dense* baseline was one of the three best performing baselines, highlighting the importance of taking

into account semantics when optimizing the quality of transitions. The *CNN Dense* baseline, minimises distance between representations extracted from the penultimate layer of the visual concept detector. Thus, it is interesting to note that using single concepts as is the case for the *Visual concepts* baseline provides very poor results, stressing the importance of considering other criteria.

Regarding visual aesthetic features, the best performing baselines where the ones that focus on minimizing the color difference between sequential images in a storyline: *Color histograms* and *Color moment*. This supports the assumption that illustrating storylines using content with similar color palettes is a solid way to optimize the quality of visual storylines. Now turning to the *Luminance* and *pHash* baselines, these presented varying results, not always being able to ensure high quality transitions. Regarding luminance, its important to note that two images can be very distinct while still presenting the same overall luminance value. Conversely, illustrating storylines by selecting images with similar entropy and number of edges, using the *Entropy* and *#Edges* baselines, provided worst results. This happens because the aesthetic similarities between the sequential images presented in these storylines are, most of the times, not easily perceptible to the naked eye.

5.3 Learning Multimodal Story Transitions

The user study showed that there is not a single feature that can produce the best quality transitions in all situations. Hence, we trained the Gradient Boosted regressor, section 4.1 to, given a *transition distance* of a pair of images, predict its quality score, according to the ground truth. The *transition distances* for the aforementioned pairs were calculated, using the features defined in Table 1. The resulting values were then standardized. Dividing the resulting dataset into train and test sets, we first trained the Gradient Boosted Trees model with 70% of the data available.

5.4 User Evaluation: Cohesive Storylines

In this section we evaluated the full framework in a realistic setting with a corpus of 35k documents and a standard retrieval method (included in our full framework) - hence, there may be non-relevant content in the final visual storyline. The media editor provides the story segments and the framework computes the optimal illustrations under a transition function trained as described in section 5.3. To do so, we considered a completely new set of EdFest and TDF images and stories. The parameters β and γ were set to 0.9 and 0.4 respectively.

To illustrate each segment of each story, the top 10 candidate tweets were selected with the standard text retrieval method BM25. In total, 953 distinct tweets were retrieved, resulting in an average of 10 candidate images to illustrate each segment. Afterwards the 4 baselines were applied, resulting in a total of 112 storylines. We proceeded to assess the quality of each visual storyline related to each story topic. Hence, the 112 distinct visual storylines were presented to 3 annotators. For each visual storyline, the annotators were again asked to rate the relevance of the images to the segment they illustrate and the transition quality between each sequential pair of images with a score of 0 ("bad") or 1 ("good").

Table 2 shows the performance of the graph based storyline generation methods in terms of the average relevance and transition

Baseline	EdFest			TDF		
	Rel.	Trans.	Qual.	Rel.	Trans.	Qual.
<i>MaxTransitions</i>	0.49	0.72	0.51	0.56	0.81	0.56
<i>MaxTransitRel</i>	0.48	0.71	0.50	0.55	0.78	0.54
<i>MaxCohesion</i>	0.47	0.77	0.52	0.62	0.91	0.64
<i>MaxCohesionRel</i>	0.42	0.61	0.42	0.59	0.72	0.57

Table 2: Average performance of the graph based storyline generation methods on the task of illustrating the Edinburgh Festival and Tour de France stories.

Figure 5: Evaluation of individual storylines for the baselines: ■ MaxCohesion, ■ MaxCohesionRel, ■ MaxTransitions and ■ MaxTransitionsRel.

scores given by the annotators, as well as through the quality assessed by users. By analyzing them we can verify that, the *maximal-cohesion without relevance* (*MaxCohesion*) approach was the best performing one in terms of average storyline quality, as assessed by the users, but also in terms of transition quality. Specifically, the storylines created by this approach were rated as having 91% high quality transitions for the TDF stories, and 77% high quality transitions for the EdFest stories. This shows that, in a storyline, the transition quality between a pair of images is not just affected by the images of said pair but also by the remaining images in the storyline. Not only that, but this approach was also the best performing one according to the user assessments, attaining a score of 0.52 for the EdFest stories and a score of 0.64 for the TDF stories.

Furthermore, by inspecting results we verify that the storylines generated by the approaches that leveraged content relevance, *MaxTransitionsRel* and *MaxCohesionRel*, were not scored as containing more relevant images when compared with the storylines generated by the two remaining approaches. By analyzing the storylines individually we see that, the approaches that leverage relevance are, in certain cases, able to pick relevant images to illustrate story segments while the remaining approaches fail to do so. However, in turn, seeking higher story cohesion seems to also result in an improved sense of relevance in particular situations. This happens because, higher transition quality means the images in a storyline tend to present similar semantics. Since stories are related

to a particular topic, semantically similar images to the ones already relevant to a story have a high chance of also being relevant to that same story. Consequently, the approaches that leverage content relevance were not able to outperform the remaining ones in terms of finding relevant content.

Figure 5 presents the performance of the graph approaches at illustrating individual stories. By analyzing these results we can see that the performance is sensitive to the story being illustrated. Although the best average performing approach was *MaxCohesion*, in a significant number of cases, the storylines created by the other approaches were scored higher. This proves the importance of having different methods of storyline creation, as they provide different illustration alternatives that news editors can work with and build upon. Taking this into account, we can state that in 10 out of 15 TDF stories, at least one of the approaches was able to provide a storyline that was rated with a score near or above 0.7, according to the quality metric. Additionally, in 8 out of 13 EdFest stories, the same can be stated for a score near or above 0.6.

Observing the examples in Figure 6, one can get a deep insight of how the different storylines created by each approach illustrated the “*What is Edinburgh Festival*” story. This compares favourably to Figure 4 examples of stories illustrated by the *Color histogram*, *Entropy*, *CCN Dense* and *#Edges* baselines. More storylines created by the graph approaches to illustrate the EdFest and TDF stories are available for inspection.³

6 CONCLUSION

In this paper we proposed and formalized a novel computational approach to assist media editors in producing visually cohesive news storylines. Through it, it is possible to express the visual and semantic relationships present in said pairs, in a non-subjective manner. Consequently, this is an important step towards the research of visual news storyline co-creation from a computational perspective. In this context, the key contributions are:

- **Story graphs:** We propose four distinct graph based approaches to visual storyline creation, Figure 2. These approaches proved to be successful at creating high quality visual news storylines, while also providing a solid baseline for future research related to this novel task.
- **Learning story transitions:** Leveraging this novel definition, we study the impact of semantic and visual characteristics in transition quality and propose and test a method for predicting the quality of transitions using a Gradient Boosted Trees regressor. This approach presented a good performance, proving that, although transition quality is a subjective topic, it is possible to systematically and accurately predict transition quality in an automated manner. This could only be achieved by providing the model with the large set of the carefully picked, low and high level features.
- **Cohesive storyline illustration:** The proposed methods provided novel insights into what impacts the perception of news storyline illustrations. We conclude that: i) the **best perceived quality** of a storyline illustration is achieved

Figure 6: Illustrations of the “What is Edinburgh Festival”. It is noticeable that shortest path solutions maximize local similarities (the bottom two stories have very strong relations for image transitions (1, 2) and (3, 4), and low relation for the image pair (2, 3). In contrast, in the top two story illustrations, all images are more homogenous and centered in its semantics (people are always present).

when targeting maximum cohesion among all story illustrations; and ii) maximum cohesion also leads to a **better sense of relevance** of the news story illustration.

Finally, although the four approaches to storyline generation presented different performances during evaluation, in the context of news illustration they can all be considered together. They present four alternative illustrated storylines, that can later be reviewed and polished by a news media editor.

Acknowledgements. This work has been partially funded by the CMU Portugal research project GoLocal Ref. CMUP-ERI/TIC/0033/2014, by the H2020 ICT project COGNITUS with the grant agreement No 687605 and by the project NOVA LINC Ref. UIDP/04516/2020. We also gratefully acknowledge the support of NVIDIA Corporation with the donation of the GPUs used for this research.

³ <https://run.unl.pt/handle/10362/66267>

REFERENCES

- [1] Fabricio Benevenuto, Gabriel Magno, Tiago Rodrigues, and Virgilio Almeida. 2010. Detecting spammers on twitter. In *Collaboration, electronic messaging, anti-abuse and spam conference (CEAS)*, Vol. 6.
- [2] Diogo Delgado, Joao Magalhaes, and Nuno Correia. 2010. Assisted news reading with automated illustration. In *Proceedings of the 18th ACM international conference on Multimedia*. ACM, 1647–1650.
- [3] Sagnik Dhar, Vicente Ordóñez, and Tamara L Berg. 2011. High level describable attributes for predicting aesthetics and interestingness. In *IEEE CVPR*.
- [4] Nicholas Diakopoulos, Munmun De Choudhury, and Mor Naaman. 2012. Finding and assessing social media information sources in the context of journalism. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2451–2460.
- [5] Erika Doggett and Alejandro Cantarero. 2016. Identifying Eyewitness News-worthy Events on Twitter. In *SocialNLP@EMNLP*.
- [6] Sitaram Asur Freddy Chong, Tat Chua. 2013. Automatic Summarization of Events From Social Media. *Int'l AAAI Conf. on Weblogs and Social Media* (2013).
- [7] Jerome H Friedman. 2001. Greedy function approximation: a gradient boosting machine. *Annals of statistics* (2001), 1189–1232.
- [8] Michael Gygli, Helmut Grabner, Hayko Riemenschneider, Fabian Nater, and Luc Van Gool. 2013. The interestingness of images. In *ICCV*. IEEE.
- [9] 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)*.
- [10] Phillip Isola, Devi Parikh, Antonio Torralba, and Aude Oliva. 2011. Understanding the intrinsic memorability of images. In *Advances in Neural Information Processing Systems*.
- [11] Mohit Iyyer, Varun Manjunatha, Anupam Guha, Yogarshi Vyas, Jordan Boyd-Graber, Hal Daume, III, and Larry S. Davis. 2017. The Amazing Mysteries of the Gutter: Drawing Inferences Between Panels in Comic Book Narratives. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [12] Gunhee Kim, Leonid Sigal, and Eric P. Xing. 2014. Joint Summarization of Large-Scale Collections of Web Images and Videos for Storyline Reconstruction. In *Proceedings of the 2014 IEEE Conference on Computer Vision and Pattern Recognition (CVPR '14)*. IEEE Computer Society, Washington, DC, USA, 8.
- [13] Gunhee Kim and Eric P Xing. 2014. Reconstructing storyline graphs for image recommendation from web community photos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- [14] C Lin, J Li, D Wang, Y Chen, and T Li. 2012. Generating event storylines from microblogs. *Proceedings of ACM CIKM* (2012).
- [15] Yiwen Luo and Xiaoou Tang. 2008. Photo and video quality evaluation: Focusing on the subject. In *European Conference on Computer Vision*. Springer.
- [16] Gonçalo Marcelino, Ricardo Pinto, and João Magalhães. 2018. Ranking News-Quality Multimedia. In *Proceedings of the 2018 ACM on International Conference on Multimedia Retrieval (ICMR '18)*. Association for Computing Machinery, New York, NY, USA, 10–18. <https://doi.org/10.1145/3206025.3206053>
- [17] Gonçalo Marcelino, David Semedo, André Mourão, Saverio Blasi, Marta Mrak, and Joao Magalhaes. 2019. A Benchmark of Visual Storytelling in Social Media. In *Proceedings of the 2019 on International Conference on Multimedia Retrieval (ICMR '19)*. Association for Computing Machinery, New York, NY, USA, 324–328. <https://doi.org/10.1145/3323873.3325047>
- [18] Adam Marcus, Michael S. Bernstein, Osama Badar, David R. Karger, Samuel Madden, and Robert C. Miller. 2011. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*.
- [19] Adam Marcus, Michael S Bernstein, Osama Badar, David R Karger, Samuel Madden, and Robert C Miller. 2011. Twitinfo: aggregating and visualizing microblogs for event exploration. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 227–236.
- [20] Andrew J. McMin, Yashar Moshfeghi, and Joemon M. Jose. 2013. Building a Large-scale Corpus for Evaluating Event Detection on Twitter. In *Proceedings of the 22nd ACM International Conference on Information & Knowledge Management (CIKM '13)*. ACM, New York, NY, USA, 409–418.
- [21] Philip J. McParlane, Andrew James McMin, and Joemon M. Jose. 2014. "Picture the Scene...": Visually Summarising Social Media Events. In *ACM CIKM*.
- [22] Ananth Mohan, Zheng Chen, and Kilian Weinberger. 2011. Web-search ranking with initialized gradient boosted regression trees. In *Proceedings of the learning to rank challenge*. PMLR, 77–89.
- [23] Jeffrey Nichols, Jalal Mahmud, and Clemens Drews. 2012. Summarizing sporting events using twitter. In *Proceedings of the 2012 ACM international conference on Intelligent User Interfaces - IUI '12*. ACM Press, New York, New York, USA, 189.
- [24] 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).
- [25] Alberto Piacenza, Fabrizio Guerrini, Nicola Adami, Riccardo Leonardi, Julie Porteous, Jonathan Teutenberg, and Marc Cavazza. 2011. Generating Story Variants with Constrained Video Recombination. In *Proceedings of the 19th ACM International Conference on Multimedia (MM '11)*. Association for Computing Machinery, New York, NY, USA, 223–232. <https://doi.org/10.1145/2072298.2072329>
- [26] 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*.
- [27] 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*. ACM Press, New York, New York, USA, 203–210. <https://doi.org/10.1145/2671188.2749407>
- [28] David Semedo and João Magalhães. 2020. Adaptive Temporal Triplet-Loss for Cross-Modal Embedding Learning. In *Proceedings of the 28th ACM International Conference on Multimedia (MM '20)*. Association for Computing Machinery, New York, NY, USA, 1152–1161. <https://doi.org/10.1145/3394171.3413540>
- [29] David Semedo and Joao Magalhaes. 2019. Diachronic Cross-Modal Embeddings. In *Proceedings of the 27th ACM International Conference on Multimedia (MM '19)*. Association for Computing Machinery, New York, NY, USA, 2061–2069. <https://doi.org/10.1145/3343031.3351036>
- [30] Beaux Sharifi, Mark-Anthony Hutton, and Jugal K Kalita. 2010. Experiments in microblog summarization. In *Social Computing (SocialCom), 2010 IEEE Second International Conference on*. IEEE, 49–56.
- [31] 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*.
- [32] Fei Wang and Min-Yen Kan. 2006. NPIC: Hierarchical synthetic image classification using image search and generic features. In *International Conference on Image and Video Retrieval*. Springer.
- [33] Zhenhua Wang, Lidian Shou, Ke Chen, Gang Chen, and Sharad Mehrotra. 2014. On Summarization and Timeline Generation for Evolutionary Tweet Streams. *IEEE Transactions on Knowledge and Data Engineering* 27, 5 (2014), 1301–1315.
- [34] Xiaowei Xu, Nurcan Yuruk, Zhidan Feng, and Thomas AJ Schweiger. 2007. Scan: a structural clustering algorithm for networks. In *Proceedings of the 13th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 824–833.
- [35] Jianchao Yang, Jiebo Luo, Jie Yu, and Thomas S Huang. 2012. Photo stream alignment and summarization for collaborative photo collection and sharing. *IEEE Transactions on Multimedia* 14, 6 (2012).