

Analysing the Memorability of a Procedural Crime-Drama TV Series, CSI

Seán Cummins, Lorin Sweeney and Alan F. Smeaton

Alan.Smeaton@DCU.ie

Insight SFI Research Centre for Data Analytics,

Dublin City University, Glasnevin, Dublin 9

Ireland

ABSTRACT

We investigate the memorability of a 5-season span of a popular crime-drama TV series, CSI, through the application of a vision transformer fine-tuned on the task of predicting video memorability. By investigating the popular genre of crime-drama TV through the use of a detailed annotated corpus combined with video memorability scores, we show how to extrapolate meaning from the memorability scores generated on video shots. We perform a quantitative analysis to relate video shot memorability to a variety of aspects of the show. The insights we present in this paper illustrate the importance of video memorability in applications which use multimedia in areas like education, marketing, indexing, as well as in the case here namely TV and film production.

CCS CONCEPTS

• **Information systems** → **Video search; Multimedia information systems.**

KEYWORDS

Video memorability, vision transformers, CSI TV series

ACM Reference Format:

Seán Cummins, Lorin Sweeney and Alan F. Smeaton. 2022. Analysing the Memorability of a Procedural Crime-Drama TV Series, CSI. In *International Conference on Content-based Multimedia Indexing (CBMI 2022)*, September 14–16, 2022, Graz, Austria. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3549555.3549592>

1 INTRODUCTION

Memorability is a subconscious measure of importance imposed on external stimuli by our internal cognition which ultimately decides the experiences we either carry with us into our long-term memory, or we discard. Memory is an unequivocally unreliable phenomenon for any person and despite its importance, it is challenging for us to influence what we will ultimately remember or forget. This absence of any unified meta-cognitive insight brings meaning to the study of memory, or more specifically into the study of memorability, more generally known as the likelihood that something will be remembered or forgotten [46].

We have evolved to store the important moments in life while forgetting uninteresting or redundant information. Thus, a memorability-based model of our experiences could allow us to filter and sort information we encounter using a human-memorability criterion [29]. Researching how our memories take shape could have a profound impact on a plethora of research areas such as human health or the curation of educational content [12, 24].

The study of memorability may also have a commercial impact in fields such as marketing, film or TV production, or content retrieval [49]. The latter of these use-cases serves as a particularly important application of memorability as the volume of digital media content available grows exponentially over time via platforms such as social networks, search engines, and recommender systems [29, 40, 43].

In this research, crime-drama exemplified in television (TV) programs such as *CSI: Crime Scene Investigation* (referred to as CSI) acts as an experimental test bed in which computational memorability as a surrogate for cognitive processing, can be analysed. CSI is a procedural crime-drama TV series with each episode following the same format. Viewers are introduced to a crime at the beginning of an episode. As the episode progresses, various clues are revealed, concluding by revealing the true perpetrator of the crime. The success of CSI can be attributed partly to its aesthetic high production values which many media critics comment on, and we argue its success is also due to its ability to promote critical thinking and memory in its viewers.

Previous studies have already proven crime-drama TV series to be useful as test cases in multiple research areas including NLP, summarisation and multi-modal inference [17, 30, 31]. In particular, the work in [37] is a large-scale anthology of Barry Salt's essays on statistical analysis of film over many years of investigation. Salt has single-handedly established statistical style analysis as a research paradigm in film studies. He argues that "many directors have sharply different styles that are easily recognized" [37, p. 13]. Salt's more recent work includes the application of film theory to the creation of TV series and in [36] he presents a stylistic analysis of episodes in twenty TV drama series, finding that the structure of TV style is uniform, a similar finding to movies.

The work by Salt and others uses a variety of measures such as shot duration and shot distribution and since then several further studies have attempted to perform statistical analysis of TV content [3, 7, 23, 35, 38]. Yet none of this, or any other, forms of analysis considers the memorability of the visual content in its analysis.

We investigate the distribution of computational memorability scores of CSI episodes over 5 seasons by leveraging modern vision transformer architectures [4, 13, 34] fine-tuned on the task of media memorability prediction. We interleave our predicted memorability

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).
CBMI 2022, September 14–16, 2022, Graz, Austria
© 2022 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-9720-9/22/09.
<https://doi.org/10.1145/3549555.3549592>

scores for given episodes of CSI with a detailed annotated corpus to investigate correlation between memorability and aspects of the show including scene and character significance.

Our research investigates whether there is a relationship between the memorability of a shot/scene and the characters present in it, how the memorability associated with particular characters develops over multiple episodes and seasons, and whether the memorability of a shot/scene is correlated with the significance of the scene. The significance of this and the reason we study the CSI TV series is precisely because of its importance in popular culture. Since its first broadcast in the early 2000s it has led to the notion of the “CSI Effect” which has altered public understanding of forensic science [27, 39] and thus has had a societal impact. CSI has also been the subject of much investigation in media studies because of its importance.

The remainder of the paper is structured as follows: In Section 2, we discuss some of the seminal related work and in Section 3, we discuss the CSI dataset used throughout this study and describe the data manipulation and augmentation techniques we performed. Section 4 describes our experiments and discusses the results of these experiments. Finally, in Section 5, we make some closing remarks and discuss future directions.

It is important for the reader to note that this work does not include any direct comparison to the work of others because there is nothing for us to compare against. Thus a direct evaluation of this work is not possible because of the nature of the topic. Instead, our novel findings present the kinds of insights which can be gleaned from computational memorability analysis of visual content and this points to the promise that such analysis can have in areas like educational content, marketing, film and TV production, and content retrieval.

2 RELATED WORK

2.1 Media Memorability

Video Memorability (VM) is a natural progression from a related research discipline: *image memorability* (IM). The work of [21] is regarded as a seminal paper on IM in which it is stated that memorability is an intrinsic property of an image, an abstract concept similar to image aesthetics which can be computed automatically using modern image analysis techniques [20]. The use of deep-learning frameworks has led to results at near human levels of consistency in the task of IM prediction [2, 6, 15] partly attributed to the development of image datasets with annotated memorability scores [22]. In contrast, VM is still in its early stages and only recently has there been available data sets [8, 29] on which to train and test VM models.

The Mediaeval Predicting Media Memorability task first took place in 2018, has run annually since then [9, 11, 12, 24, 25] and has been central to the development of VM. The 2021 edition saw the use of vision transformers [4, 13, 34] across all participant submissions. The focus of work in the area is predominantly on developing memorability prediction techniques and to date the most applicable use of VM has been in the analysis of TV advertisements [28, 41].

2.2 Affective Video Content Analysis

Affective Video Content Analysis (AVCA) is a research area closely related to VM which investigates the emotions elicited by video among its viewers [5]. Like VM, AVCA-based research aims to generate metacognitive insights capable of improving video content indexing and retrieval [19, 50]. Besides this, AVCA research could, for example, protect children from emotionally harmful media [48], as well as create mood-dependent video recommendation [18]. Past research in this domain has shown the modelling of some signatures of film such as tempo [1], violence [14, 32], and emotion [45].

2.3 Macabre Fascination in CSI: The Rise of the Corpse

The impact of crime-drama TV series such as CSI on modern pop-culture stretches far beyond academic research, even resulting in the coining of a new term: ‘The CSI Effect’ [27, 39], mentioned earlier. Forensic autopsy portrayed in CSI acts as a softening lens in which exploration of the dead who are usually the victims of the crime which is the focus of the episode, is socially-palatable entertainment [33]. As CSI viewership increased from the first series in 2001, so did the number of forensic science courses available across educational bodies¹. Early 2000s crime-drama TV such as CSI sexualises the cadavers present in their series through the ever-present “beautiful female victim” motif. The sexualised corpse is a stock character that makes regular appearances. Our eyes are invited to linger on the passive, beautiful, dead bodies that lay in the morgue. This phenomenon captured the public’s attention and increased the abundance of sexualised cadavers throughout subsequent seasons of CSI while other series such as *Law & Order* began adding autopsies to their shows [16].

3 THE CSI DATASET

3.1 Source Data

We use 39 annotated episodes of CSI in our analysis, curated by Frermann *et al.* [17] and available on the Edinburgh NLP GitHub². The original source data consists of two corpora, the first of which is *Perpetrator Identification* (PI). PI corpus’ files are annotated at word-level for each episode. A word can be uttered by a speaking character or is part of the screenplay. Each word is accompanied by *case ID*, *sentence ID*, *speaker*, *sentence start time*, *sentence end time*. We also have binary indicators such as *killer gold*, *suspect gold*, and *other gold* which indicate whether the word uttered mentions the killer, suspect, or otherwise.

The second corpus in the original annotated data set is *Screenplay Summarisation* (SS) presented at scene-level (a section of the story with its own unique combination of setting, character, and dialogue). For a given episode, the dataset consists of each scene annotated with *Scene ID*, *screenplay*, as well as the scene *aspect*. Aspect is a categorical feature indicating the scene’s significance and can be

¹Retrieved from: <http://www.telegraph.co.uk/education/universityeducation/3243086/CSI-leads-to-increase-in-forensic-science-courses.html>

²The repository hosting the original CSI annotation data is available at <https://github.com/EdinburghNLP/csi-corpus> and the memorability scores we computed on this corpus are available at https://github.com/scummins00/CSI-Memorability/blob/main/data/shot_memorability.csv

one or many of the following: *Crime Scene*, *Victim*, *Death Cause*, *Evidence*, *Perpetrator*, *Motive*, or *None*.

Each of the 39 episodes was also available to us in video format by scraping directly from the DVDs purchased from Amazon, as well as an extra 3 episode video files.

3.2 Data Preprocessing

3.2.1 Data Manipulation: We aggregated each of our PI episode datasets from word-level to sentence-level, combining the *killer gold*, *suspect gold*, and *other gold* binary indicators previously mentioned into a single categorical column. We then disaggregated our SS corpora from scene-level to sentence-level. With equivalent datasets of the same configuration, we then merged them and Table 1 shows a section of one of our augmented and merged datasets for S1.E8.

3.2.2 Data Augmentation: Features available from the annotation are rich and allow us to analyse data in a variety of ways. However, we noted two shortcomings: (i) the data provides little context regarding suspects throughout an episode other than the *type mentioned* feature and (ii) the *aspect* feature does not allow us to identify scenes involving suspects. Also, a portion of scenes involving the killer are not labelled with the ‘Perpetrator’ aspect. We extended the *aspect* feature of any sentence in a scene in which the perpetrator or a suspect speaks so as to have the values of ‘Perpetrator’ or ‘Suspect’, respectively.

3.3 Pre-processing of Video Files

We subdivided our clips into shots which allows for more specific and fine-grained indexing, in turn reducing the amount of video we needed to process per analysis. We used a neural shot-boundary-detection (SBD) framework, *TransNet V2*,³ described in [44], to segment episodes into individual shots.

3.4 Data Generation

The use of vision transformers in predicting memorability scores is shown in [10, 26]. Our vision transformer architecture is the CLIP [34] pre-trained image encoder. We use this encoder to train a Bayesian Ridge Regressor (BRR) on the Memento10k dataset [29]. BRR was used previously by [46] for computing memorability scores based on extracted features. For each training sample, the encoder extracts features from frames sampled at a rate of 3 frames per second. The BRR is then trained on the features and memorability score pairs. To compute the memorability of unseen CSI clips, we first extract representative frames using a temporal-based frame extraction mechanism. These are used as input to the vision transformer architecture to extract features. Our BRR model then produces a memorability score based on these features. This framework is shown in Figure 1.

4 METHODOLOGY AND EXPERIMENTAL RESULTS

We now present the distribution of memorability scores across each episode, relating it to its metadata from annotations taken from [17].

For each episode, we plotted the memorability scores associated with shots in chronological order to create a time-series signal representing the intrinsic memorability of the episode. We then annotated each episode’s signal according to various metadata including (i) the scene aspect, and (ii) the character speaking (including None). We used Laplacian smoothing described in Equation 1 to remove the jaggedness present in raw memorability scores

$$\hat{x}_i = \frac{1}{N} \sum_{j=1}^N \hat{x}_j \quad (1)$$

where N is the size of the smoothing window. We generated signals for each episode using a range of smoothing window sizes from 15 to 305 shots.

To compliment the shot memorability score distributions we gathered memorability scores across the distribution of characters, and of aspects. By aggregating the overall speaking time of the main characters, we identify the lead as well as the secondary characters. We performed a similar analysis to understand the scene count per aspect value. Our results are shown in Figure 2.

We investigated the memorability distribution associated with the series characters and aspects and this is presented in Figure 3 parts (a) and (b) in each individual season as well as across all 5 seasons. We investigated: i) the main character speaking, and (ii) the scene aspect. The lack of variation across Figure 3 (a) and (b) and the relatively high memorability scores is an artefact of the Memento10k memorability dataset [29] used to train the memorability model we used where we rarely see a video segment given a memorability score below 0.7. Instead we are interested in repeating patterns e.g. which parts of shows consistently result in higher or lower memorability scores and what these patterns might reveal to us about the series.

In Figure 3 (a), we see that ordering the cast in terms of memorability over all 5 seasons is almost identical to that seen in Figure 2 (a). Catherine has less speaking time than Grissom, but is consistently considered more memorable. With the exception of Season 3, Nick is simultaneously the least-memorable character and the character with the least time spent speaking. We observe a somewhat positive linear relationship between a character’s importance (in terms of screen-time), and their memorability.



In Figure 2 (b), the proportion of scenes in which each aspect value appears correlates with the purpose in which that aspect serves. For example, ‘Motive’ appears least across all seasons as only a small portion near the end of each episode is dedicated to revealing a perpetrator’s motive. Similarly, ‘Crime scene’ and ‘Victim’ scenes usually occur near the beginning of an episode. Despite this, as shown in Figure 3 (b), ‘Motive’ scenes are considered the most memorable across all 5 seasons. In contrast, both ‘Crime scene’ and ‘Victim’ scenes are considered the least memorable.

Scenes described with the ‘Death Cause’ aspect value are associated with autopsy scenes in which viewers are shown flashbacks to the crime committed, as well as the cadaver of a usually youthful, attractive victim. We discussed the importance of these scenes and their pivotal role in the development of TV crime series in Section 2.3. In Figure 3 (b), the intrinsic memorability of this scene type increases as the seasons progress. These pedagogic, grotesque, scenes capture the audience’s fascination because of their brash

³TransNet V2 is an SBD framework: <https://github.com/soCzech/TransNetV2>.

Table 1: Sample augmented dataset from Season 1, Ep. 8. Original data from [17].

caseID	sentID	speaker	type mentioned	start	end	sentence	aspect
1	6	Grissom	other	00:36.5	00:41.5	where’s the girl?	Victim
1	7	Officer	other	00:41.5	00:46.6	she’s down this hall	Victim
1	8	None	None	00:46.6	00:51.7	Grissom turns the corner	Crime scene
1	9	None	None	00:51.7	00:56.8	the officer signals Grissom inside	Crime scene

**Figure 1: Our memorability prediction framework, extracting representative frames and computing a memorability score for each shot.****Figure 2: (a) The time (mins) spent speaking by each main cast member. (b) The number of scenes in which each aspect value is present.**

autoptic vision [47] and are a staple of the crime-drama TV genre and contribute to its popular success. The visual memorability of these scenes correlates with the success of this scene-type.

5 CONCLUSIONS AND FUTURE WORK

Video memorability is a recent computational tool for analysis of visual media which is more abstract and higher-level than raw content analysis like object detection or action classification. In this paper we have used the computation of video memorability at video shot level to analyse a popular procedural crime-drama TV series revealing insights not previously visible.

Our investigation creates a mélange of intertwined arguments ranging from more statistically-driven interpretations of film such as cinematics, to artistic studies on the relationship between a series’ popularity and its most memorable moments. We hypothesise that scene and character importance is related to the memorability of a scene, and show that correlation exists between factors of the show not previously visible. We show that these correlations are not random, but instead have complex interpretation. The memorability scores we generated uncovered patterns within the early seasons of CSI, highlighting the significance and memorability of an episode’s finale in comparison to other scenes. Our memorability

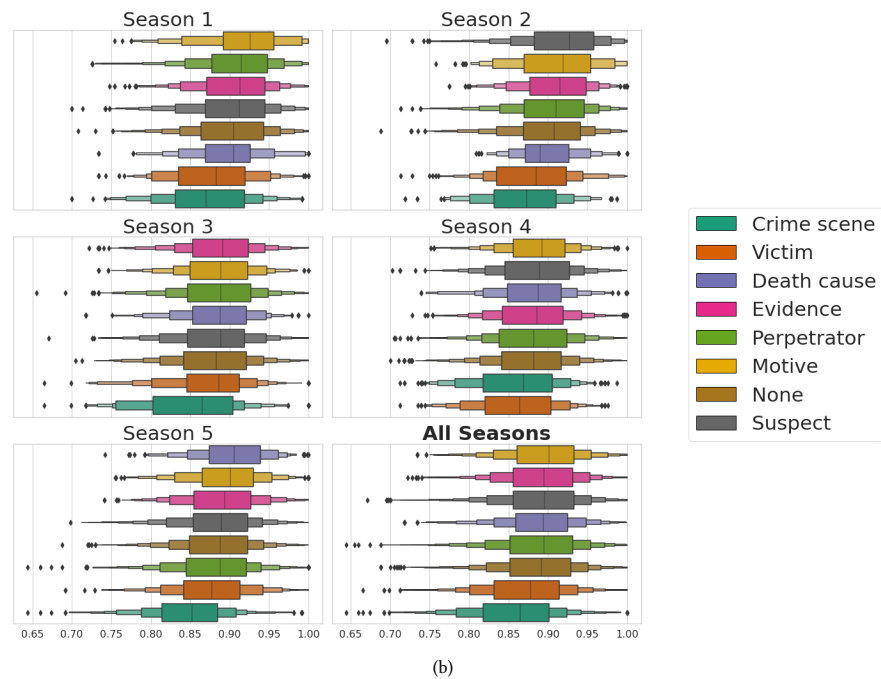



Figure 3: Graphs display results from Seasons 1 to 5 and then results for all seasons combined. (a) displays computed memorability scores for shots with the main characters, while (b) displays computed memorability scores for shots with associated with aspect values.

scores for lead cast members heavily correlate with the characters' importance, considering on-screen time as importance. The work

here is one of the first efforts in the application of memorability scores beyond benchmarking environments.

Our first area for future work is to increase the data size by including more episodes of CSI, as well as episodes from other shows such as *CSI: Miami* or other ‘Whodunnit-styled’ series such as *Law & Order*. In our study, we used a temporal-based frame extraction method to extract representative frames from shots. This could be enhanced via Representative Frame Extraction techniques [42], which are capable of automatically selecting the most representative frames from video content.

It is clear that we have developed an importance-proxy by analysing various metadata from CSI. Ideally, we would define a user task in which participants’ genuine memorability of aspects of a crime-drama series can be examined. This gives a real measure of importance with respect to the show, rather than an importance proxy.

The novel findings we have presented here reveal insights from computational memorability scores for video shots of a TV crime series. As mentioned earlier, this is useful and interesting from a film or TV studies perspective but it also shows the potential that computational memorability has in other areas. We could now analyse or even edit videos to be used in online education so as to maximise the memorability of key moments. We could structure video content used in marketing and advertising, or as we have seen here in film and TV production, so that key messages are delivered with maximised memorability. Future work in these and other areas may indeed exploit computational memorability as a criterion for composing video content.

ACKNOWLEDGMENTS

This publication has emanated from research partly supported by Science Foundation Ireland (SFI) under Grant Number SFI/12/RC/2289_P2 (Insight SFI Research Centre for Data Analytics).

REFERENCES

- [1] B. Adams, C. Dorai, and S. Venkatesh. 2000. Novel approach to determining tempo and dramatic story sections in motion pictures. In *Proceedings 2000 International Conference on Image Processing (Cat. No.00CH37101)*, Vol. 2. IEEE, Vancouver, BC, Canada, 283–286 vol.2. <https://doi.org/10.1109/ICIP.2000.899358>
- [2] Erdem Akagunduz, Adrian G. Bors, and Karla K. Evans. 2020. Defining Image Memorability Using the Visual Memory Schema. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 42, 9 (Sep 2020), 2165–2178. <https://doi.org/10.1109/TPAMI.2019.2914392>
- [3] Taylor Arnold, Lauren Tilton, and Annie Berke. 2019. Visual Style in Two Network Era Sitcoms. *Journal of Cultural Analytics* 4, 2 (Jul 2019), 11045. <https://doi.org/10.22148/16.043>
- [4] Hangbo Bao, Li Dong, and Furu Wei. 2021. BEiT: BERT Pre-Training of Image Transformers. *arXiv:2106.08254 [cs]* (Jun 2021). <http://arxiv.org/abs/2106.08254> arXiv: 2106.08254.
- [5] Yoann Baveye, Christel Chamaret, Emmanuel Dellandréa, and Liming Chen. 2018. Affective Video Content Analysis: A Multidisciplinary Insight. *IEEE Transactions on Affective Computing* 9, 4 (Oct 2018), 396–409. <https://doi.org/10.1109/TAFFC.2017.2661284>
- [6] Yoann Baveye, Romain Cohendet, Matthieu Perreira Da Silva, and Patrick Le Callet. 2016. Deep Learning for Image Memorability Prediction: the Emotional Bias. In *Proceedings of the 24th ACM international conference on Multimedia (MM '16)*. Association for Computing Machinery, New York, NY, USA, 491–495. <https://doi.org/10.1145/2964284.2967269>
- [7] Jeremy Butler. 2014. Statistical Analysis of Television Style: What Can Numbers Tell Us about TV Editing? *Cinema Journal* 54, 1 (2014), 25–44.
- [8] Romain Cohendet, Claire-Hélène Demarty, Ngoc Duong, and Martin Engilberge. 2019. VideoMem: Constructing, Analyzing, Predicting Short-Term and Long-Term Video Memorability. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*. IEEE, Seoul, Korea (South), 2531–2540. <https://doi.org/10.1109/ICCV.2019.00262>
- [9] Romain Cohendet, Claire-Hélène Demarty, Ngoc Duong, Mats Sjöberg, Bogdan Ionescu, Thanh-Toan Do, and France Rennes. 2018. MediaEval 2018: Predicting Media Memorability Task. *arXiv:1807.01052 [cs]* (Jul 2018). <http://arxiv.org/abs/1807.01052> arXiv: 1807.01052.
- [10] Mihai Gabriel Constantin and Bogdan Ionescu. 2021. Using Vision Transformers and Memorable Moments for the Prediction of Video Memorability. (2021), 3.
- [11] Mihai Gabriel Constantin, Bogdan Ionescu, Claire-Hélène Demarty, Ngoc Q K Duong, Xavier Alameda-Pineda, and Mats Sjöberg. 2019. The Predicting Media Memorability Task at MediaEval 2019. (Dec 2019), 3.
- [12] Alba García Seco De Herrera, Rukiye Savran Kiziltepe, Jon Chamberlain, Mihai Gabriel Constantin, Claire-Hélène Demarty, Faiyaz Doctor, Bogdan Ionescu, and Alan F. Smeaton. 2020. Overview of MediaEval 2020 Predicting Media Memorability Task: What Makes a Video Memorable? *arXiv:2012.15650 [cs]* (Dec 2020). <http://arxiv.org/abs/2012.15650> arXiv: 2012.15650.
- [13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. *arXiv:2010.11929 [cs]* (Jun 2021). <http://arxiv.org/abs/2010.11929> arXiv: 2010.11929.
- [14] Florian Eyben, Felix Weninger, Nicolas Lehment, Björn Schuller, and Gerhard Rigoll. 2013. Affective Video Retrieval: Violence Detection in Hollywood Movies by Large-Scale Segmental Feature Extraction. *PLOS ONE* 8, 12 (Dec 2013), e78506. <https://doi.org/10.1371/journal.pone.0078506>
- [15] Jiri Fajtl, Vasileios Argyriou, Dorothy Monekosso, and Paolo Remagnino. 2018. AMNet: Memorability Estimation with Attention. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, Salt Lake City, UT, USA, 6363–6372. <https://doi.org/10.1109/CVPR.2018.00666>
- [16] Jacque Lynn Foltyn. 2008. Dead famous and dead sexy: Popular culture, forensics, and the rise of the corpse. *Mortality* 13, 2 (May 2008), 153–173. <https://doi.org/10.1080/13576270801954468>
- [17] Lea Frermann, Shay B. Cohen, and Mirella Lapata. 2018. Whodunnit? Crime Drama as a Case for Natural Language Understanding. *Transactions of the Association for Computational Linguistics* 6 (Dec 2018), 1–15. https://doi.org/10.1162/tacl_a_00001
- [18] A. Hanjalic. 2006. Extracting moods from pictures and sounds: towards truly personalized TV. *IEEE Signal Processing Magazine* 23, 2 (Mar 2006), 90–100. <https://doi.org/10.1109/MSP.2006.1621452>
- [19] Ching Hau Chan and Gareth J. F. Jones. 2005. Affect-based indexing and retrieval of films. In *Proceedings of the 13th Annual ACM International Conference on Multimedia*. Association for Computing Machinery, Singapore, 427–430. http://portal.acm.org/ft_gateway.cfm?id=1101243&type=pdf&CFID=19879361&CFTOKEN=18605329
- [20] Feiyang Hu and Alan F. Smeaton. 2018. Image Aesthetics and Content in Selecting Memorable Keyframes from Lifelogs. In *Multimedia Modeling*. Klaus Schoeffmann, Thanarat H. Chalidabhongse, Chong Wah Ngo, Supavadee Aramvith, Noel E. O’Connor, Yo-Sung Ho, Moncef Gabbouj, and Ahmed Elgammal (Eds.). Springer International Publishing, Bangkok, Thailand, 608–619.
- [21] Phillip Isola, Devi Parikh, Antonio Torralba, and Aude Oliva. 2011. Understanding the intrinsic memorability of images. In *Proceedings of the 24th International Conference on Neural Information Processing Systems (NIPS’11)*. Curran Associates Inc., Red Hook, NY, USA, 2429–2437.
- [22] Aditya Khosla, Akhil S. Raju, Antonio Torralba, and Aude Oliva. 2015. Understanding and Predicting Image Memorability at a Large Scale. In *2015 IEEE International Conference on Computer Vision (ICCV)*. 2390–2398. <https://doi.org/10.1109/ICCV.2015.275>
- [23] Yunhwan Kim and Sunmi Lee. 2020. Exploration of the Characteristics of Emotion Distribution in Korean TV Series: Common Pattern and Statistical Complexity. *IEEE Access* 8 (2020), 69438–69447. <https://doi.org/10.1109/ACCESS.2020.2985673>
- [24] Rukiye Savran Kiziltepe, Mihai Gabriel Constantin, Claire-Hélène Demarty, Graham Healy, Camilo Fosco, Alba García Seco de Herrera, Sebastian Halder, Bogdan Ionescu, Ana Matran-Fernandez, Alan F. Smeaton, and Lorin Sweeney. 2021. Overview of The MediaEval 2021 Predicting Media Memorability Task. *arXiv:2112.05982 [cs]* (Dec 2021). <http://arxiv.org/abs/2112.05982> arXiv: 2112.05982.
- [25] Rukiye Savran Kiziltepe, Lorin Sweeney, Mihai Gabriel Constantin, Faiyaz Doctor, Alba García Seco de Herrera, Claire-Hélène Demarty, Graham Healy, Bogdan Ionescu, and Alan F. Smeaton. 2021. An annotated video dataset for computing video memorability. *Data in Brief* 39 (2021), 107671.
- [26] Ricardo Kleinlein, Cristina Luna-Jiménez, and Fernando Fernández-Martínez. 2021. THAU-UPM at MediaEval 2021: From Video Semantics To Memorability Using Pretrained Transformers. (2021), 3.
- [27] Evelyn M. Maeder and Richard Corbett. 2015. Beyond frequency: Perceived realism and the CSI effect. *Canadian Journal of Criminology and Criminal Justice* 57, 1 (2015), 83–114.
- [28] Li-Wei Mai and Georgia Schoeller. 2009. Emotions, attitudes and memorability associated with TV commercials. *Journal of Targeting, Measurement and Analysis for Marketing* 17, 1 (Mar 2009), 55–63. <https://doi.org/10.1057/jt.2009.1>
- [29] Anelise Newman, Camilo Fosco, Vincent Casser, Allen Lee, Barry McNamara, and Aude Oliva. 2020. Multimodal Memorability: Modeling Effects of Semantics and Decay on Video Memorability. In *Computer Vision – ECCV 2020 (Lecture*

- Notes in Computer Science*), Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm (Eds.), Springer International Publishing, Cham, 223–240. https://doi.org/10.1007/978-3-030-58517-4_14
- [30] Pinelopi Papalampidi, Frank Keller, Lea Frermann, and Mirella Lapata. 2020. Screenplay Summarization Using Latent Narrative Structure. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Online, 1920–1933. <https://doi.org/10.18653/v1/2020.acl-main.174>
- [31] Nikos Papasaraktopoulos, Lea Frermann, Mirella Lapata, and Shay B. Cohen. 2019. Partners in Crime: Multi-view Sequential Inference for Movie Understanding. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Association for Computational Linguistics, Hong Kong, China, 2057–2067. <https://doi.org/10.18653/v1/D19-1212>
- [32] Cédric Penet, Claire-Hélène Demarty, Guillaume Gravier, and Patrick Gros. 2012. Multimodal information fusion and temporal integration for violence detection in movies. In *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2393–2396. <https://doi.org/10.1109/ICASSP.2012.6288397>
- [33] Ruth Penfold-Mounce. 2016. Corpses, popular culture and forensic science: public obsession with death. *Mortality* 21, 1 (Jan 2016), 19–35. <https://doi.org/10.1080/13576275.2015.1026887>
- [34] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning Transferable Visual Models From Natural Language Supervision. *arXiv:2103.00020 [cs]* (Feb 2021). <http://arxiv.org/abs/2103.00020> arXiv: 2103.00020.
- [35] Nick Redfern. 2014. The Structure of ITV News Bulletins. *International Journal of Communication* 8, 00 (Jun 2014), 22.
- [36] Barry Salt. 2001. Practical Film Theory and its Application to TV Series Dramas. *Journal of Media Practice* 2, 2 (2001), 98–113. <https://doi.org/10.1386/jmpr.2.2.98> arXiv:<https://doi.org/10.1386/jmpr.2.2.98>
- [37] Barry Salt. 2006. *Moving into Pictures: More on film history, style, and analysis*. London: Starword.
- [38] Richard J. Schaefer and Tony J. Martinez. 2009. Trends in Network News Editing Strategies From 1969 Through 2005. *Journal of Broadcasting & Electronic Media* 53, 3 (Aug 2009), 347–364. <https://doi.org/10.1080/08838150903102600>
- [39] N.J. Schweitzer and Michael J. Saks. 2007. The CSI Effect: Popular Fiction About Forensic Science Affects the Publics Expectations About Real Forensic Science. *Jurimetrics* 47, 3 (2007), 357–364.
- [40] Sumit Shekhar, Dhruv Singal, Harvaneet Singh, Manav Kedia, and Akhil Shetty. 2017. Show and Recall: Learning What Makes Videos Memorable. In *2017 IEEE International Conference on Computer Vision Workshops (ICCVW)*. 2730–2739. <https://doi.org/10.1109/ICCVW.2017.321>
- [41] Wangbing Shen, Zongying Liu, Linden J. Ball, Taozhen Huang, Yuan Yuan, Haiping Bai, and Meifeng Hua. 2020. Easy to Remember, Easy to Forget? The Memorability of Creative Advertisements. *Creativity Research Journal* 32, 3 (Jul 2020), 313–322. <https://doi.org/10.1080/10400419.2020.1821568>
- [42] N Shruthi and S Priyamvada. 2017. Dominant frame extraction for video indexing. In *2017 2nd IEEE International Conference on Recent Trends in Electronics, Information Communication Technology (RTEICT)*. 1799–1803. <https://doi.org/10.1109/RTEICT.2017.8256909>
- [43] Aliaksandr Siarohin, Gloria Zen, Cveta Majtanovic, Xavier Alameda-Pineda, Elisa Ricci, and Nicu Sebe. 2019. Increasing Image Memorability with Neural Style Transfer. *ACM Transactions on Multimedia Computing, Communications, and Applications* 15, 2 (Jun 2019), 1–22. <https://doi.org/10.1145/3311781>
- [44] Tomáš Souček and Jakub Lokoč. 2020. TransNet V2: An effective deep network architecture for fast shot transition detection. *arXiv:2008.04838 [cs]* (Aug 2020). <http://arxiv.org/abs/2008.04838> arXiv: 2008.04838.
- [45] Kai Sun, Junqing Yu, Yue Huang, and Xiaoqiang Hu. 2009. An improved valence-arousal emotion space for video affective content representation and recognition. In *2009 IEEE International Conference on Multimedia and Expo*. 566–569. <https://doi.org/10.1109/ICME.2009.5202559>
- [46] Lorin Sweeney, Graham Healy, and Alan F. Smeaton. 2021. The Influence of Audio on Video Memorability with an Audio Gestalt Regulated Video Memorability System. In *2021 International Conference on Content-Based Multimedia Indexing (CBMI)*. 1–6. <https://doi.org/10.1109/CBMI50038.2021.9461903>
- [47] Sue Tait. 2006. Autoptic vision and the necrophilic imaginary in CSI. *International Journal of Cultural Studies* 9, 1 (Mar 2006), 45–62. <https://doi.org/10.1177/1367877906061164>
- [48] Jianchao Wang, Bing Li, Weiming Hu, and Ou Wu. 2011. Horror video scene recognition via Multiple-Instance learning. In *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 1325–1328. <https://doi.org/10.1109/ICASSP.2011.5946656>
- [49] Fumei Yue, Jing Li, and Jiande Sun. 2021. Insights of Feature Fusion for Video Memorability Prediction. In *Digital TV and Wireless Multimedia Communication*, Guangtao Zhai, Jun Zhou, Hua Yang, Ping An, and Xiaokang Yang (Eds.). Springer, Singapore, 239–248. https://doi.org/10.1007/978-981-16-1194-0_21
- [50] Sicheng Zhao, Hongxun Yao, Xiaoshuai Sun, Pengfei Xu, Xianming Liu, and Rongrong Ji. 2011. Video indexing and recommendation based on affective analysis of viewers. In *Proceedings of the 19th ACM international conference on Multimedia (MM '11)*. Association for Computing Machinery, New York, NY, USA, 1473–1476. <https://doi.org/10.1145/2072298.2072043>