

Video Verification in the Fake News Era

Vasileios Mezaris · Lyndon Nixon ·
Symeon Papadopoulos · Denis Teyssou
Editors

Video Verification in the Fake News Era

Editors

Vasileios Mezaris
Centre for Research and Technology Hellas
Information Technologies Institute
Thermi, Thessaloniki, Greece

Lyndon Nixon
MODUL Technology GmbH
MODUL University Vienna
Vienna, Austria

Symeon Papadopoulos
Centre for Research and Technology Hellas
Information Technologies Institute
Thermi, Thessaloniki, Greece

Denis Teyssou
Agence France-Presse
Paris, France

ISBN 978-3-030-26751-3

ISBN 978-3-030-26752-0 (eBook)

<https://doi.org/10.1007/978-3-030-26752-0>

© Springer Nature Switzerland AG 2019

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, expressed or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Preface

The digital media revolution is bringing breaking news to online video platforms, and news organizations often rely on user-generated recordings of breaking and developing events shared in social media to illustrate the story. However, in video there is also deception. In today's 'fake news' era, access to increasingly sophisticated editing and content management tools and the ease in which fake information spreads in electronic networks require the entire news and media industries to carefully verify third-party content before publishing it. This book presents the latest technological advances and practical tools for discovering, verifying and visualizing social media video content, and managing related rights. These are expected to be of interest to computer scientists and researchers, news and media professionals, and even policymakers and data-savvy media consumers.

The book is organized in four main parts. Part I presents the necessary Problem Statement, Part II covers the various Technologies that can contribute to video verification, Part III introduces three complete Applications that integrate several verification technologies and Part IV presents some Concluding Remarks.

Part I Problem Statement

The first step in addressing the problem of 'fake news', or disinformation, is to understand the problem. Chapter 1, 'Video Verification: Motivation and Requirements', attempts to introduce us to the peculiarities of the video verification problem by initially presenting the motivations of those involved in video verification, showcasing the respective requirements and highlighting the importance and relevance of tackling disinformation on social networks. Then, this chapter provides an overview of the state of the art of techniques and technologies for video verification. It also highlights the emergence of new threats, such as the so-called 'deep fakes'. Finally, the chapter concludes by formulating an empirical typology of false videos spreading online.

Part II Technologies

In this part of the book, Chaps. 2 through 8 present in-depth analyses of different technologies that contribute to video verification. Chapter 2, ‘Real-Time Story Detection and Video Retrieval from Social Media Streams’, starts one step before coming to verifying a specific video: it discusses how a journalist can detect emerging news stories online and find videos around that story, which may then require verification. The chapter starts by reviewing the prior research in the area of topic detection, and then presents a keyword-graph-based method for news story discovery out of Twitter streams. Subsequently, it presents a technique for the selection of online videos that are candidates for news stories, by using the detected stories to form a query against social networks. This enables relevant information retrieval at web scale for news-story-associated videos. These techniques are evaluated by observation of the detected stories and of the news videos that are presented for those stories, demonstrating how journalists can quickly identify videos for verification and reuse.

Chapter 3 focuses on ‘Video Fragmentation and Reverse Search on the Web’. Such search is a first and simple, yet often very valuable, means for checking if a video under examination, or a slightly modified version of it, has appeared in previous times in the web and social sphere. Video reuse is in fact the ‘easy fake’: it does not take elaborate editing tools and effort to fake an event in this way; it suffices to fetch some older footage of, e.g. a terrorist attack or a plane crash from the web, and repost it claiming that this is happening right now, right before your eyes. Chapter 3 presents technologies for the fragmentation of a video into visually and temporally coherent parts and the extraction of a representative keyframe for each defined fragment that enables the provision of a complete and concise keyframe-based summary of the video; these keyframes can then be used for performing a fragment-level search for the video on the web. Following a literature survey on the topic, the chapter describes two state-of-the-art methods for video subshot fragmentation—one relying on the assessment of the visual coherence over sequences of frames, and another one that is based on the identification of camera activity during the video recording. It then goes on to present a web application that enables the fine-grained (at the fragment-level) reverse search for near-duplicates of a given video on the web, and evaluation results and conclusions about the effectiveness of these technologies as well as some thoughts on future developments.

Chapter 4, ‘Finding Near-Duplicate Videos in Large-Scale Collections’, sticks to the topic of detecting video reuse for combating the ‘easy fakes’ that we started to deal with in Chap. 3, but views this from a different—and complementary—perspective. In Chap. 3, we discussed web-scale search, which inevitably relies on the reverse image search functionalities that are offered by popular web search engines. The latter provide excellent coverage of the whole web, but on the other hand only allow us to deal with reverse video search in a ‘quick and dirty’ way: by searching for and matching just isolated keyframes. In Chap. 4, we deal with finding duplicate

or near-duplicate videos (Near-Duplicate Video Retrieval—NDVR) in our own, closed collections of videos. This means that the coverage of the search is more limited (since we cannot index the whole web in the way that major web search engines do), but on the other hand we can do a much more elaborate and accurate search, because we have full control of the indexing and searching process. Thus, having indexed, for instance, a large number of web videos that show previous terrorist attacks and related content, if we want to check the possible prior use of an (allegedly) new terrorist attack video we can complement the web-scale search of Chap. 3 with a more accurate search in our own collection of such videos. As the main objective of a typical NDVR approach is, given a query video, to retrieve all near-duplicate videos in a video repository and rank them based on their similarity to the query, the chapter starts by reviewing the literature on this topic, and then goes on to present two methods for video-level matching. Extensive evaluation on publicly available benchmark datasets documents the merits of these approaches, and their complementarity to keyframe-based web-scale search.

Chapter 5, ‘Finding Semantically-Related Videos in Closed Collections’, takes the search for similar video content one step further. When trying to verify a video, and an associated news story, important cues can come from looking at the greater picture: what other videos out there (and thus also in our closed collection of videos, as long as we keep collecting videos related to specific, potentially newsworthy events, such as terrorist attacks) can support (or disprove) the claims made with the help of the specific video in question? For this, besides any near-duplicate videos (as discussed in Chaps. 3 and 4), we would like to detect semantically similar videos. That is, videos showing the same event/actors/activities from a different viewpoint or videos coming from the same source (‘channel’—in the broad sense). In other words, we need to be able to organize any content that we collect from the web and social media sources. For this, we discuss two classes of techniques in this chapter: the detection of semantic concepts in video (a.k.a. the annotation of the video with semantic labels) and the detection of logos that are visible in videos and can help us to identify their provenance. Both classes of techniques rely on deep learning (deep neural networks), which is a learning paradigm that is considered to be a key element of Artificial Intelligence (AI). The chapter discusses the state of the art in these two sub-problems of video understanding and presents two techniques developed by the authors of the chapter and their experimental results.

Chapter 6, ‘Detecting Manipulations in Video’, discusses another fundamental problem related to video verification: can we trust what we see? If an event really unfolded before our eyes, the answer would be yes. But if it is shown on video, how can we assess if the video is an accurate depiction of (some) reality or an alternate ‘reality’ whose capture in video was only made possible with the help of digital video editing tools? To answer this question, this chapter presents the techniques researched and developed within InVID for the forensic analysis of videos, and the detection and localization of forgeries. Following an overview of state-of-the-art video tampering detection techniques, the chapter documents that the bulk of current research is mainly dedicated to frame-based tampering analysis or

encoding-based inconsistency characterization. The authors built upon this existing research, by designing forensics filters aimed to highlight any traces left behind by video tampering, with a focus on identifying disruptions in the temporal aspects of a video. Subsequently, they proceeded to develop a deep learning approach aimed to analyse the outputs of these forensics filters and automatically detect tampered videos. Experimental results on benchmark and real-world data, and analyses of the results, show that the proposed deep-learning-based method yields promising results compared to the state of the art, especially with respect to the algorithm's ability to generalize to unknown data taken from the real world. On the other hand, the same analyses also show that this problem is far from being resolved, and further research on it is in order.

Chapter 7, 'Verification of Web Videos Through Analysis of Their Online Context', continues in the direction of previous chapters, most notably Chap. 5, of looking at the greater picture for verifying a specific video. Contrary (and complementarily) to Chap. 5, though, we are not examining here other related videos that can help debunk the video under examination; instead, we are looking at the online 'context' of this video. The goal is to extract clues that can help us with the video verification process. As video context, we refer to information surrounding the video in the web and/or the social media platforms where it resides, i.e. information about the video itself, user comments below the video, information about the video publisher and any dissemination of the same video through other video platforms or social media. As a starting point, the authors present the Fake Video Corpus, a dataset of debunked and verified UGVs that aim at serving as reference for qualitative and quantitative analysis and evaluation. Next, they present a web-based service, called Context Aggregation and Analysis, which supports the collection, filtering and mining of contextual pieces of information that can serve as verification signals.

Chapter 8, 'Copyright Management of User Generated Video for Journalistic Reuse', concludes this part of the book on technologies, by considering what comes after a newsworthy piece of user-generated video is verified: how can the journalist use it in a legal way? For this, i.e. for reviewing the copyright scope of reuse of user-generated videos usually found in social media, for journalistic purposes, the starting point of this chapter is the analysis of current practices in the news industry. Based on this analysis, the authors provide a set of recommendations for social media reuse under copyright law and social networks terms of use. Moreover, they describe how these recommendations have been used to guide the development of the InVID Rights Management module, focusing on EU copyright law given the context of the InVID EU project.

Part III Applications

Chapter 9, ‘Applying Design Thinking Methodology: The InVID Verification Plugin’, kick-starts the presentation of integrated, complete tools for journalists who want to verify user-generated videos. It describes the methodology used to develop and release a browser extension which has become one of the major tools to debunk disinformation and verify videos and images, in a period of less than 18 months. This is a tool that combines several of the technologies discussed in Chaps. 2 through 8 in a free, easy-to-use package, which has attracted more than 12,000 users worldwide from media newsrooms, fact-checkers, the media literacy community, human rights defenders and emergency response workers dealing with false rumours and content.

Chapter 10, ‘Multimodal Analytics Dashboard for Story Detection and Visualization’, is the second tool presented in this part of the book. The InVID Multimodal Analytics Dashboard is a visual content exploration and retrieval system to analyse user-generated video content from social media platforms including YouTube, Twitter, Facebook, Reddit, Vimeo and Dailymotion. That is, it is not a tool for video verification, but rather a tool for discovering emerging newsworthy stories and related video content, which then may be verified (either using the InVID Verification plugin, presented in the previous chapter; or by directly transferring the video in question, with a click of a button, to the InVID Verification Application that will be discussed in the following chapter). The InVID Multimodal Analytics Dashboard uses automated knowledge extraction methods to analyse each of the collected postings and stores the extracted metadata for later analyses. The real-time synchronization mechanisms of the dashboard help to track information flows within the resulting information space. Cluster analysis is used to group related postings and detect evolving stories, which can be analysed along multiple semantic dimensions—e.g. sentiment, geographic location, opinion leaders (persons or organizations) as well as the relations among these opinion leaders. The result can be used by data journalists to analyse and visualize online developments within and across news stories.

Chapter 11, ‘Video Verification in the Newsroom’, comes as a natural extension of both Chap. 9 (which presented a first tool for the exact same problem: video verification) and Chap. 10, whose Multimodal Analytics Dashboard provides a direct, one-click link for importing newsworthy videos detected with the latter tool into the newsroom’s video verification pipeline. The chapter starts by describing the integration of a video verification process into newsrooms of TV broadcasters or news agencies. The authors discuss the organizational integration concerning the workflow, responsibility and preparations as well as the inclusion of innovative verification tools and services into an existing IT environment. Then the authors present the InVID Video Verification Application or Verification App for short. This can be considered to be an ‘InVID Verification plugin on steroids’, i.e. a more complete and professional application for video verification, which can serve as a blueprint for introducing video verification processes in professional newsroom systems. This verification application, similarly to the InVID Verification plugin, combines several of the technologies discussed in Chaps. 2 through 8.

Part IV Concluding Remarks

The book concludes with Chap. 12, ‘Disinformation: the Force of Falsity’, which departs a bit from the primarily technology-oriented presentation in previous chapters, to engage in a more forward-looking discussion on how can we avoid the proliferation of fake videos, and stop them from spreading over and over again. This final chapter borrows the concept of force of falsity from the famous Italian semiotician and novelist Umberto Eco, to describe how manipulated information remains visible and accessible despite efforts to debunk it. It illustrates, with the help of real-life examples, how search engine indexes are getting confused by disinformation and they too often fail to retrieve the authentic pieces of content, the ones which are neither manipulated nor decontextualized. The chapter concludes with some further thoughts on how to address this problem.

Thessaloniki, Greece
Vienna, Austria
Thessaloniki, Greece
Paris, France
May 2019

Vasileios Mezaris
Lyndon Nixon
Symeon Papadopoulos
Denis Teyssou

Acknowledgements Most of the work reported throughout this book was supported by the European Unions Horizon 2020 research and innovation programme under grant agreement No 687786 ‘InVID: In Video Veritas—Verification of Social Media Video Content for the News Industry’, 2016–2018.

Contents

Part I Problem Statement

- 1 Video Verification: Motivation and Requirements 3**
Denis Teyssou and Jochen Spangenberg

Part II Technologies

- 2 Real-Time Story Detection and Video Retrieval from Social Media Streams 17**
Lyndon Nixon, Daniel Fischl and Arno Scharl
- 3 Video Fragmentation and Reverse Search on the Web 53**
Evlampios Apostolidis, Konstantinos Apostolidis, Ioannis Patras and Vasileios Mezaris
- 4 Finding Near-Duplicate Videos in Large-Scale Collections 91**
Giorgos Kordopatis-Zilos, Symeon Papadopoulos, Ioannis Patras and Ioannis Kompatsiaris
- 5 Finding Semantically Related Videos in Closed Collections 127**
Foteini Markatopoulou, Markos Zampoglou, Evlampios Apostolidis, Symeon Papadopoulos, Vasileios Mezaris, Ioannis Patras and Ioannis Kompatsiaris
- 6 Detecting Manipulations in Video 161**
Grégoire Mercier, Foteini Markatopoulou, Roger Cozien, Markos Zampoglou, Evlampios Apostolidis, Alexandros I. Metsai, Symeon Papadopoulos, Vasileios Mezaris, Ioannis Patras and Ioannis Kompatsiaris
- 7 Verification of Web Videos Through Analysis of Their Online Context 191**
Olga Papadopoulou, Markos Zampoglou, Symeon Papadopoulos and Ioannis Kompatsiaris

8	Copyright Management of User-Generated Video for Journalistic Reuse	223
	Roberto García, Maria Teixidor, Paloma de Barrón, Denis Teyssou, Rosa Gil, Albert Berga and Gerard Rovira	
Part III Applications		
9	Applying Design Thinking Methodology: The InVID Verification Plugin	263
	Denis Teyssou	
10	Multimodal Analytics Dashboard for Story Detection and Visualization	281
	Arno Scharl, Alexander Hubmann-Haidvogel, Max Göbel, Tobi Schäfer, Daniel Fischl and Lyndon Nixon	
11	Video Verification in the Newsroom	301
	Rolf Fricke and Jan Thomsen	
Part IV Concluding Remarks		
12	Disinformation: The Force of Falsity	339
	Denis Teyssou	
	Index	349

Contributors

Evlampios Apostolidis Centre for Research and Technology Hellas, Information Technologies Institute, Thessaloniki, Greece;
School of Electronic Engineering and Computer Science, Queen Mary University, London, UK

Konstantinos Apostolidis Centre for Research and Technology Hellas, Information Technologies Institute, Thessaloniki, Greece

Albert Berga Universitat de Lleida, Lleida, Spain

Roger Cozien eXo maKina, Paris, France

Paloma de Barrón Private Law Department, Universitat de Lleida, Lleida, Spain

Daniel Fischl MODUL Technology GmbH, Vienna, Austria

Rolf Fricke Condat AG, Berlin, Germany

Roberto García Computer Science and Engineering Department, Universitat de Lleida, Lleida, Spain

Rosa Gil Computer Science and Engineering Department, Universitat de Lleida, Lleida, Spain

Max Göbel webLyzard technology gmbh, Vienna, Austria

Alexander Hubmann-Haidvogel webLyzard technology gmbh, Vienna, Austria

Ioannis Kompatsiaris Centre for Research and Technology Hellas, Information Technologies Institute, Thessaloniki, Greece

Giorgos Kordopatis-Zilos Centre for Research and Technology Hellas, Information Technologies Institute, Thessaloniki, Greece;
School of Electronic Engineering and Computer Science, Queen Mary University, London, UK

Foteini Markatopoulou Centre for Research and Technology Hellas, Information Technologies Institute, Thessaloniki, Greece

Grégoire Mercier eXo maKina, Paris, France

Alexandros I. Metsai Centre for Research and Technology Hellas, Information Technologies Institute, Thessaloniki, Greece

Vasileios Mezaris Centre for Research and Technology Hellas, Information Technologies Institute, Thessaloniki, Greece

Lyndon Nixon MODUL Technology GmbH, Vienna, Austria

Symeon Papadopoulos Centre for Research and Technology Hellas, Information Technologies Institute, Thessaloniki, Greece

Olga Papadopoulou Centre for Research and Technology Hellas, Information Technologies Institute, Thessaloniki, Greece

Ioannis Patras School of Electronic Engineering and Computer Science, Queen Mary University, London, UK

Gerard Rovira Universitat de Lleida, Lleida, Spain

Tobi Schäfer webLyzard technology gmbh, Vienna, Austria

Arno Scharl webLyzard technology gmbh, Vienna, Austria

Jochen Spangenberg Deutsche Welle, Berlin, Germany

Maria Teixidor Universitat de Lleida, Lleida, Spain

Denis Teyssou Agence France-Presse, Paris, France

Jan Thomsen Condat AG, Berlin, Germany

Markos Zampoglou Centre for Research and Technology Hellas, Information Technologies Institute, Thessaloniki, Greece