Fact-Checked Claim Detection in Videos Using a Multimodal Approach

Frédéric Rayar[a]
LIFAT, University of Tours, Tours, France

Keywords: Multimodal Multimedia Analysis, Computational Journalism, Fact-Checking, Verified Claims Retrieval.

Abstract: The recent advances in technology and social networks have led to a phenomenon that is so called, information disorder. Our information environment is polluted by misinformation, disinformation and fake news and at a global scale. Hence, assessing the veracity of information becomes mandatory nowadays. So far, most of the efforts in the last decades have focused on analysing textual content, scraped from blogs and social media, and trying to predict the veracity of information. However, such false information also appears in multimedia content and have a peculiar life cycle throughout time that is worth leveraging. In this paper, we present a multimodal approach for detecting claims that have already been fact-checked, from videos input. Focusing on political discourse in French language, we demonstrate the feasibility of a complete system, offline and explainable.

1 INTRODUCTION

The advent of internet and social networks in the last decades has emphasized the proliferation of misinformation and disinformation. The expression "fake news", popularised during the 2016 United States presidential election, has nowadays became part of our daily lives. Appearing in several domains, such as politics, health or finance, it has become a major and urgent concern, due to its impact on the society (acceptance of false beliefs, impact on brand and organisations, etc.). This concern is emphasized when it touches the political sphere and has to be addressed to improve the democratic accountability, restore confidence in political institutions and the political discourse.

To do so, several fact-checking initiatives and organisations have emerged to validate the veracity of claims that can be found online or in traditional media (newspaper, radio, television). Fact-checkers play an essential role in today's information environment, however coping with the deluge of information that has to be verified appears to be impossible. Furthermore, recent advances in Artificial Intelligence have lead to the curation of more suspicious content that can be disseminated online, such as *DeepFakes* (Westerlund, 2019) or *ChatGPT* outcomes [1]. Hence, finding effective and large-scale solutions to assist human fact-checkers becomes a necessity. This has shed light on the potential of *automated fact-checking* technologies that can assist human fact-checkers. Automated fact-checking is an umbrella expression in academic research that encompasses several sub-tasks such as identifying check-worthy claims, verifying claims or identifying rumors and satire, among others (Kotonya and Toni, 2020). Thorough surveys on current automated fact-checking technologies can be found in (Kotonya and Toni, 2020)[2], (Nakov et al., 2021a), or (Guo et al., 2022).

Figure 1 illustrates a classic fact-checking pipeline: first, one has to find claims that are worth fact-checking from various sources. Second, it appears relevant to verify if the selected claim has already been fact-checked previously. Third, one has to gather relevant evidence to help understanding the context and the veracity of the claim. Finally, one has to decide either the claim is wrong, mostly wrong, imprecise, mostly true or true. Note that several graduation scales can be considered in the latter step.

Most of the efforts on asserting the veracity of information in the last decade have focused on textual content, whether scraped from blogs and social media or obtained with video/radio transcripts, thanks to techniques from the Natural Language Processing (NLP) field. However, political information mainly appears in television broadcasts, such TV shows (in-

[a] https://orcid.org/0000-0003-1927-8400
[1] https://openai.com/blog/chatgpt

[2] https://github.com/neemakot/Fact-Checking-Survey

Rayar, F.
Fact-Checked Claim Detection in Videos Using a Multimodal Approach.
DOI: 10.5220/0012545500003660

In *Proceedings of the 19th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP 2024) - Volume 4: VISAPP*, pages 614-620
ISBN: 978-989-758-679-8; ISSN: 2184-4321

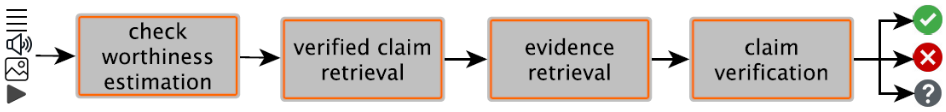

Figure 1: Classic fact-checking pipeline (illustration from (Nakov et al., 2021a)).

terviews or debates) or breaking news. This fact leads to the need to leverage methods that can analyse multimedia content (image, audio and video) in order to perform *multimedia (or multimodal) automated fact-checking* (Akhtar et al., 2023). Furthermore, as manual fact-checking is very time-consuming, it is worth saving these efforts and leverage them to help others fact-checkers, avoiding them to waste time on claims that have already been fact-checked. Indeed, we rely on the anterior logic of false news, that are often slightly modified (or not) and repeated over different time scales.

In this paper, we propose to address this issues by proving the feasibility of a multimodal detection tool for fact-checked claims in videos. Our contributions can be summarized as follows:

- To the best of our knowledge, this study is the first one that searches French fact-checked claims from videos input.
- We demonstrate the feasibility of a complete system, with a simple, yet promising, offline and explainable approach.
- We experiment on a video dataset and demonstrate the effectiveness of our approach. Future improvements are also discussed.

The rest of the paper is organized as follows: Section 2 provides a brief review of related works. Section 3 presents the proposed multimodal approach. Section 4 details the methodology and results of the system's evaluation on a toy dataset. Finally, Section 5 states conclusions and perspectives of this work.

2 RELATED WORKS

To the best of our knowledge, the earliest work on detecting whether a claim has already been fact-checked is rather recent (2020). The authors (Shaar et al., 2020) leverage BERT models to do so and present their results on the PolitiFact dataset [3] and Snopes dataset [4], that are only in English. Following this first study, the authors have built a specific lab, namely *CheckThat!*, in the Conference and Labs of the Evaluation Forum (CLEF). This conference contributes to the evaluation of information access systems, primarily through experimentation on shared tasks (i.e. challenges). A dedicated task on the claim retrieval has been running since 2020 (Barrón-Cedeño et al., 2020), (Nakov et al., 2021b) and (Nakov et al., 2022). As defined in their website [5], the task consists in *"given a check-worthy claim, and a set of previously-checked claims, determine whether the claim has been previously fact-checked with respect to a collection of fact-checked claims"*. If we focus on the latest edition (Nakov et al., 2022), a subtask is more related to our work: *"detecting previously fact-checked claims in political debates/speeches [...] This is a ranking task, where systems are asked to produce a list of top-n candidates"*. As the earliest work (Shaar et al., 2020), the subtask is only given in English and uses the PolitiFact dataset [6].

Hence, contrary to our work, these studies only focus on an unimodal detection, based on NLP techniques, and do not take into account French language.

3 MULTIMODAL DETECTION OF FACT-CHECKED CLAIMS FROM VIDEOS

Figure 2 illustrates the workflow that we propose in this paper to perform multimodal detection of fact-checked claims from videos. The video is processed as follows : first, facial regions of interest are detected and we verify if the face is recognized with regard to a set of persons of interest that have been learnt. Second, the audio of the video is extracted and used to retrieve the transcription using an Automatic Speech Recognition system (ASR), along with timestamps of words. Third, we perform a search in an existing database of fact-checked claims, described by their authors and a list of keywords. To de so, we use both visual (face recognized) and textual (keywords from the transcription) features. If a relevant fact-checked claim is found in the database, we link it at the specific time of the video where the claim is being said and can notify the viewer using, for instance, a pop-up. These steps are detailed below.

[3] https://www.politifact.com/

[4] https://www.snopes.com/fact-check/

[5] https://sites.google.com/view/clef2022-checkthat/tasks/task-2-detecting-previously-fact-checked-claims

[6] https://www.politifact.com/

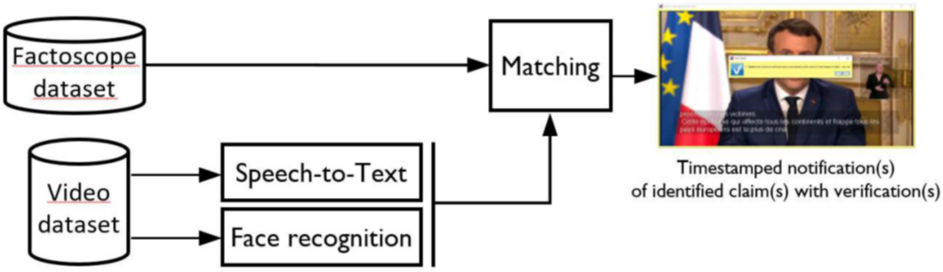

Figure 2: Proposed workflow to perform multimodal detection of fact-checked claims from videos.

3.1 Face Detection

The face detection step is performed using Google BlazeFace (Bazarevsky et al., 2019). Adapted from the Single Shot Multibox Detector (SSD) framework (Liu et al., 2016) and MobileNetV2 (Sandler et al., 2018), BlazeFace has been optimized for inference on mobile GPUs. Therefore it fits perfectly our real time video processing consideration, since it runs at a speed of 200–1000+ FPS on flagship smartphones. Furthermore, it can provide an accurate facial region of interest or a task-specific model of the face. In our experimentations, we have used the Python implementation of BlazeFace provided by the Face Library [7] package. Since this implementation doesn't use heavy frameworks (such as TensorFlow, Keras or PyTorch), it can be easily embed for production.

3.2 Face Recognition

Several state-of-the-art face recognition models can be found in the literature, such as VGG-Face (Parkhi et al., 2015) (Oxford University), FaceNet (Google Research) (Schroff et al., 2015), OpenFace (Baltrusaitis et al., 2018) (Carnegie Mellon University) or DeepFace (Taigman et al., 2014) (Facebook Research). According to a recent comparative study (Serengil and Ozpinar, 2020), FaceNet outperforms the other three aforementioned algorithms and achieves the best performance, in terms of F1-score (98.55%). Therefore, we have selected FaceNet in the rest of our experimentations. Regarding its implementation, we have used the DeepFace Python framework [8] (Serengil and Ozpinar, 2020). This framework is not only appropriate for real-time face recognition, but could also be leveraged for taking into consideration several facial attributes in future works (e.g. age, emotion and gender). Furthermore, it allows to easily fine tune the model to consider new persons of interest.

[7]https://github.com/a-akram-98/face_lib

[8]https://github.com/serengil/deepface

3.3 Transcription Generation

In order to leverage textual features in the proposed multimodal approach, we have used the videos' transcriptions. To generate these transcriptions, an automatic speech recognition, also known as Speech-to-Text (STT), algorithm is needed. In our experimentations, we have selected a solution that could be used in an offline mode, namely Vosk [9]. Vosk is an offline open source speech recognition toolkit. It enables speech recognition for more than 20 languages and dialects, including French, which is the language of interest in this study. Two models are provided for French : a big and accurate model (1.4GB), that is relevant for servers and a smaller one (41MB) more adapted for desktop and mobile applications. The latest model have be used in our experimentations. Another plus of Vosk is that it provides word-level time offset (timestamp) values, allowing us to align a detected fact-checked claim to the transcription, hence the video.

3.4 Multimodal Retrieval

We assume that the fact-checked claims are stored in a dataset and are described by, at least, their authors and a set of keywords (let us name this set *claim_keywords*). The retrieval of fact-checked claims in the videos is performed in a multimodal scheme, using both visual and textual features. First, we select the fact-checked claims in the dataset that have been said by people that have been recognized in the current video. Let us name this set *candidates_claims*. Second, from the transcription, a part-of-speech tagging is performed and only nouns, pronouns, verbs and adjectives (tagged as 'ADJ', 'NOUN', 'PROPN', 'ADJ' and 'VERB' respectively), that are unique, are kept as keywords. Let us name this set of keywords *transcription_keywords*. Third, we compute the similarity between *transcription_keywords* and

[9]https://alphacephei.com/vosk/

each set *claim_keywords* of a claim that belong to *candidates_claims*. This similarity is computed by comparing word vectors (or word embeddings) that have been generated by the classical Word2vec algorithm (Mikolov et al., 2013). The Word2vec has been outperformed by transformer models in the recent years, however we chose this simple method in our first experimentations to validate the feasibility of our multimodal approach in real-time. Finally, if at least or more than N keywords match ($similarity > sim\_thresh$), then we consider that a sentence correspond to the fact-checked claim. For the experimentations presented in the next section, we have empirically set $N = 2$ and $sim\_thresh = 0.8$. Hence, by construction, the proposed multimodal approach to retrieve claims that have been fact-checked in videos , is explainable. The Python free open-source library spaCy [10] has been used for the implementation of NLP algorithms.

4 EXPERIMENTATIONS

4.1 Data

4.1.1 Fact-Checked Dataset

Factoscope is an initiative that has begun at the end of 2016, to address the political discourse claim fact-checking during the debates of the 2017 French presidential elections. The fact-checked claims are presented in a website powered by a content management system software, namely WordPress. Each claim corresponds to a given article in a web page. Each article presents several features of the claim: its author and a picture of the author's face, the claim itself, the level of veracity of the claim, a contextual description of the claim, and links that have been used to perform the fact-checking. We have led a web scraping campaign to gather the fact-checked claims in a structured scheme. The scraping of the Factoscope website led us to create a dataset of more than $1,300$ fact-checked claims. These claims deal with 12 various topics (*e.g.* education, environment, justice, health, security, etc.) that have been stated by more than 200 political figures between 2017 and 2022. Three levels of veracity are considered in the Factoscope fact-checking service: the claims that have been fact-checked are composed of 55% false claims, 25% of imprecise claims and 20% true claims. Sources to check the veracity of the claims are either original or extracted from well-known national fact-checkers, such as *Factuel* (l'AFP - Agence France Presse)[11], *Les Decodeurs* (le Monde)[12] or *Le Vrai du faux* (France Info)[13] . This dataset is publicly available for the community [14] (upon signing an agreement) and described in (Rayar et al., 2022). Early NLP analysis have been conducted on this dataset using Unitex/Gramlab (Paumier et al., 2009), an open source [15], cross-platform, multilingual, lexicon- and grammar-based corpus processing suite. Focusing on the description attribute of each claim, we consider a corpus of $241,202$ words (with $17,998$ unique words), that contains $23,886$ named entities (mostly persons, organisations, roles, places and dates).

4.1.2 Video Dataset

From the Factoscope dataset described above, we have selected the top-10 political figures that have fact-checked claims. First, we have trained the face recognition model to take into account these new 10 person of interest, with an average of only 80 training images per person (crawled online from image search results). Then, we have built a toy video dataset where at least one of the selected political figure is present to validate the feasibility of our multimodal approach. This dataset consists in 15 short videos featuring political discourse (mostly interviews) in French. They have been cherry-picked from online streaming platforms such as they contain at least one claim that appears in the Factoscope dataset. Table 1 describes the dataset, that is also available online [16]. The original resolution of these videos are either $720p$ or $1080p$, but we have also conducted experimentations with a reduced video resolution of $240p$. Indeed, this $240p$ resolution corresponds to the resolution of the videos that are in the STVD-FC dataset, a larger video dataset that has been curated for future experimentations, composed of $6,730$ TV programs, that represent a total duration of $6,540$ hours and a memory storage of 2TB. It is currently publicly available [17] and described in (Rayar et al., 2022).

4.2 Results and Discussion

Table 2 presents the results that are obtained by our approach on the video toy dataset. Each video has

[10] https://spacy.io/

[11] https://factuel.afp.com/

[12] https://www.lemonde.fr/les-decodeurs/

[13] https://www.francetvinfo.fr/replay-magazine/franceinfo/vrai-ou-fake-1-emission/

[14] https://dataset-stvd.univ-tours.fr/fc/

[15] https://unitexgramlab.org/

[16] http://frederic.rayar.free.fr/fact-checking/videos.html

[17] https://dataset-stvd.univ-tours.fr/fc/

Table 1: Video toy dataset description.

id	Political figure	Video name	Veracity	Timestamp
1	Jordan Bardella	Bardella-Travailleurs-Detaches	False	00:17
2	Jordan Bardella	Bardella_Soins_Gratuits	False	00:21
3	Nicolas Bay	Bay_France_Raciste	False	00:24
4	Nicolas Bay	bay_kamikaze_manchester	False	00:27
5	Christophe Castaner	Castaner-Attaque-Hopital	False	00:03
6	Nicolas Dupont-Aignan	Dupont_Aignan_Constitution_Laïcité	False	00:32
7	Nicolas Dupont-Aignan	dupont_aignan_seul_parti_progresse	False	00:14
8	Nicolas Dupont-Aignan	Dupont-Aignan-Taire	False	00:09
9	Emmanuel Macron	macron_croissance_six_pourcent	True	00:18
10	Emmanuel Macron	macron_plus_grave_crise	True	00:20
11	Jean-Luc Mélenchon	melenchon_vaccins_russes	False	00:23
12	Marine Le Pen	MLP_5000_expulsés	False	00:36
13	Marine Le Pen	MLP_Nationalité_Coulibaly	False	00:24
14	Marine Le Pen	MLP-5-Ans-Nationalite	False	00:12
15	Laurent Wauquiez	wauquiez_jamais_gilet_jaune	False	00:22

Table 2: Face recognition and fact-checked claim detection results.

id	Recognized faces' name	Correct claim found	Wrong claim found
1	Bardella	Yes	No
1 (240p)	Bardella	Yes	No
2	Bardella	Yes	Yes
2 (240p)	Bardella, Wauquiez	Yes	Yes
3	Bay	No	No
3 (240p)	Bay, Wauquiez	No	No
4	Bay	Yes	No
4 (240p)	Bay, Blanquer	Yes	No
5	N/A	N/A	N/A
5 (240p)	N/A	N/A	N/A
6	Bay, Macron, Dupont-Aignan	Yes	Yes
6 (240p)	Blanquer, Dupont-Aignan	Yes	No
7	Dupont-Aignan	Yes	No
7 (240p)	Dupont-Aignan	Yes	No
8	Dupont-Aignan	No	No
8 (240p)	Dupont-Aignan	No	No
9	Castaner, Blanquer	N/A	N/A
9 (240p)	Castaner, Blanquer, Le Pen, Dupont-Aignan	N/A	N/A
10	Macron	Yes	No
10 (240p)	Macron	Yes	No
11	Mélenchon	Yes	Yes
11 (240p)	Mélenchon	Yes	Yes
12	Le Pen	Yes	Yes
12 (240p)	Le Pen, Wauquiez, Dupon- Aignan	Yes	Yes
13	Le Pen, Dupont-Aignan	Yes	No
13 (240p)	Le Pen	Yes	No
14	Le Pen	Yes	Yes
14 (240p)	Le Pen	Yes	Yes
15	Le Pen, Wauquiez	Yes	No
15 (240p)	Castaner, Wauquiez	Yes	Yes

been considered both in their original and reduced resolution. For each video, the table describes if the correct claim has been found (true positive) and if additional wrong claims have also been found (false positive).

When considering the original resolutions, we observe that our approach successfully identify the correct claim in the fact-checked dataset 11 times out of

15 , corresponding to a recall of 73.3%. Failing at finding the correct claim can be explained by two explanations:

1. The claim author's face is not correctly recognized (for instance, in videos 5 or 9, where no faces or wrong faces are recognized, respectively). This could be addressed by using more training images when fine tuning the face recognition model.
2. The quality of the transcription did not allowed a correct matching (for instance, in videos 3 or 8, where the correct claim author's face is recognized). This could be addressed by either determining optimal values of the parameters N and $thresh\_sim$ or considering more recent matching algorithm (for instance, transformer models).

Since wrong claims are also found in 5 cases, the proposed method has a precision of 68.75%. Similar explanations (solutions) could explain (improve) this score:

1. The face recognition module find other faces than the correct claim author's face, either by mistake or because they do actually appear in the video, leading our system to consider more elements in *candidates_claims* (i.e. more claims in the dataset), and thus having the possibility to match a wrong claim.
2. The overlap between *claim_keywords* sets are important, meaning that some claims of the dataset share several identical keywords with some other claims. This could be addressed by refining the quality or the diversity of the keywords describing fact-checked claims (either manually curated or automatically extracted/inferred).

Regarding our experimentations with video resolution reduced to 240p, it appears that the impact is not important. Indeed, for all the videos considered in the toy dataset, when the correct claim was found on the original resolution, it was also found on the 240p version. The influence of this reduced resolution can however be observed at the facial recognition step: for most of the videos (7 out of 15, nearly 50%), the face recognition module find several wrong persons. This argue for the fact that our multimodal approach, that considers both visual and textual features, is relevant when dealing with reduced video resolution.

5 CONCLUSION AND PERSPECTIVES

In this paper, we have introduced a multimodal approach for detecting claims that have already been fact-checked in videos input. Due to the recurring aspect of false information that is propagated throughout different media and the time-consuming task to assess the veracity of information for fact-checkers, we believe that such a system could be provided as an asset to experts such as journalists, but also the general public. Focusing on political discourse in French language, we demonstrate the feasibility of a complete system, offline and explainable. The results that have been obtained are promising towards future real-time applications, and its robustness could be easily improved using more recent and performing state-of-the-art methods such as transformers models. In future works, we also plan to stress out our workflow with a larger fact-checked dataset that is currently being curated and the larger STVD-FC video dataset.

REFERENCES

Akhtar, M., Schlichtkrull, M., Guo, Z., Cocarascu, O., Simperl, E., and Vlachos, A. (2023). Multimodal automated fact-checking: A survey.

Baltrusaitis, T., Zadeh, A., Lim, Y. C., and Morency, L.-P. (2018). Openface 2.0: Facial behavior analysis toolkit. In *2018 13th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2018)*, pages 59–66.

Barrón-Cedeño, A., Elsayed, T., Nakov, P., Martino, G. D. S., Hasanain, M., Suwaileh, R., Haouari, F., Babulkov, N., Hamdan, B., Nikolov, A., Shaar, S., and Ali, Z. S. (2020). Overview of checkthat 2020: Automatic identification and verification of claims in social media. *CoRR*, abs/2007.07997.

Bazarevsky, V., Kartynnik, Y., Vakunov, A., Raveendran, K., and Grundmann, M. (2019). Blazeface: Sub-millisecond neural face detection on mobile gpus. *CoRR*, abs/1907.05047.

Guo, Z., Schlichtkrull, M., and Vlachos, A. (2022). A survey on automated fact-checking. *Transactions of the Association for Computational Linguistics*, 10:178–206.

Kotonya, N. and Toni, F. (2020). Explainable automated fact-checking: A survey. *CoRR*, abs/2011.03870.

Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.-Y., and Berg, A. C. (2016). Ssd: Single shot multibox detector. In *Proceedings of the 14th European Conference on Computer Vision*, pages 21–37.

Mikolov, T., Sutskever, I., Chen, K., Corrado, G., and Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2*, page 3111–3119.

Nakov, P., Barrón-Cedeño, A., da San Martino, G., Alam, F., Struß, J. M., Mandl, T., Míguez, R., Caselli, T., Kutlu, M., Zaghouani, W., Li, C., Shaar, S., Shahi, G. K., Mubarak, H., Nikolov, A., Babulkov, N., Kartal, Y. S., Wiegand, M., Siegel, M., and Köhler, J.

(2022). Overview of the clef–2022 checkthat! lab on fighting the covid-19 infodemic and fake news detection. In Barrón-Cedeño, A., Da San Martino, G., Degli Esposti, M., Sebastiani, F., Macdonald, C., Pasi, G., Hanbury, A., Potthast, M., Faggioli, G., and Ferro, N., editors, *Experimental IR Meets Multilinguality, Multimodality, and Interaction*, pages 495–520.

Nakov, P., Corney, D. P. A., Hasanain, M., Alam, F., Elsayed, T., Barrón-Cedeño, A., Papotti, P., Shaar, S., and Martino, G. D. S. (2021a). Automated fact-checking for assisting human fact-checkers. *CoRR*, abs/2103.07769.

Nakov, P., Da San Martino, G., Elsayed, T., Barrón-Cedeño, A., Míguez, R., Shaar, S., Alam, F., Haouari, F., Hasanain, M., Mansour, W., Hamdan, B., Ali, Z. S., Babulkov, N., Nikolov, A., Shahi, G. K., Struß, J. M., Mandl, T., Kutlu, M., and Kartal, Y. S. (2021b). Overview of the clef–2021 checkthat! lab on detecting check-worthy claims, previously fact-checked claims, and fake news. In Candan, K. S., Ionescu, B., Goeuriot, L., Larsen, B., Müller, H., Joly, A., Maistro, M., Piroi, F., Faggioli, G., and Ferro, N., editors, *Experimental IR Meets Multilinguality, Multimodality, and Interaction*, pages 264–291.

Parkhi, O. M., Vedaldi, A., and Zisserman, A. (2015). Deep face recognition. In *Proceedings of the British Machine Vision Conference (BMVC)*, pages 41.1–41.12.

Paumier, S., Nakamura, T., and Voyatzi., S. (2009). Unitex, a corpus processing system with multi-lingual linguistic resources. In *eLexicography in the 21st century: new challenges, new applications (eLEX)*, pages 173–175.

Rayar, F., Delalandre, M., and Le, V.-H. (2022). A large-scale tv video and metadata database for french political content analysis and fact-checking. In *Proceedings of the 19th International Conference on Content-Based Multimedia Indexing*, CBMI '22, page 181–185. Association for Computing Machinery.

Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., and Chen, L. (2018). Mobilenetv2: Inverted residuals and linear bottlenecks. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4510–4520.

Schroff, F., Kalenichenko, D., and Philbin, J. (2015). Facenet: A unified embedding for face recognition and clustering. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 815–823.

Serengil, S. I. and Ozpinar, A. (2020). Lightface: A hybrid deep face recognition framework. In *2020 Innovations in Intelligent Systems and Applications Conference (ASYU)*, pages 23–27.

Shaar, S., Babulkov, N., Da San Martino, G., and Nakov, P. (2020). That is a known lie: Detecting previously fact-checked claims. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3607–3618.

Taigman, Y., Yang, M., Ranzato, M., and Wolf, L. (2014). Deepface: Closing the gap to human-level performance in face verification. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1701–1708.

Westerlund, M. (2019). The emergence of deepfake technology: A review. *Technology Innovation Management Review*, 9:40–53.