VeriDash: an AI-Driven, User-Centric Open Source Dashboard for Enhancing Multimedia Verification

Johannes Skivdal, Laurence Dierickx and Duc-Tien Dang-Nguyen

November 12, 2024

VeriDash: An AI-Driven, User-Centric Open Source Dashboard for Enhancing Multimedia Verification

Johannes Skivdal[1], Laurence Dierickx[1], and Duc-Tien Dang-Nguyen[1]

University of Bergen, Norway
Johannes.Skivdal@student.uib.no, laurence.dierickx@uib.no, ductien.dangnguyen@uib.no

Abstract. This paper presents VeriDash, an open source dashboard that integrates AI-based technologies to streamline the multimedia verification process for fact-checkers. VeriDash offers advanced features such as automated transcription, geolocation, and an intuitive interface that supports a human-driven fact-checking process while ensuring ease of use. By incorporating a human-in-the-loop approach, VeriDash balances technological efficiency with human expertise, promoting trusted and responsible AI technology to support and enhance the fact-checking process.

Keywords: fact-checking · multimedia verification · human-in-the-loop

1 Introduction

The rise of fact-checking as a distinct journalistic sub-genre over the past decade can be attributed to a convergence of internal and external factors within the news media ecosystem [26]. The proliferation of digital media and social platforms has led to an unprecedented spread of mis- and disinformation [28]. This environment has increased the need for rigorous verification processes, underscoring the importance of fact-checking as an effective strategy to combat information disorder [1,3]. At the same time, the increasing use of political polarisation for spreading disinformation, coupled with a growing demand for accountability, has increased the need to provide accurate, balanced and reliable counter-information [21].

Digital tools and AI-based technologies are an integral part of the fact-checking apparatus, utilising a range of resources that are often not initially designed for this purpose but adapted to increase the efficiency and scope of fact-checking efforts [16]. However, fact-checkers often need more time to explore the wide range of tools available, some of which require acquiring specific skills, and they tend to use the same tools for the same purposes [7]. In addition, fragmented practices tend to increase workloads [16], whereas automated solutions are valued for streamlining time-consuming tasks, which depend largely on the type of claim being verified [7].

As digital media, including images, videos and audio recordings, increasingly populate online platforms, the ability to verify multimedia content has become a critical part of fact-checkers' work. Multimedia verification is further complicated as image and video manipulation is not necessarily about altered, cropped or fake content, and it often relate to authentic but decontextualised content [29]. Visual content verification, therefore, requires a combination of techniques, including reverse image search, metadata analysis and cross-referencing with reliable sources, to ensure that the content is accurate and contextually appropriate. In addition, the rise of deepfakes and cheapfakes (more results on fighting with these challenges can be found in [13,4]), aided by sophisticated generative AI technology, makes it increasingly important to distinguish between genuine and deceptive content [14].

In response to these challenges, developing a comprehensive tool that integrates the full range of multimedia verification capabilities into a single, easy-to-use platform is critical to supporting fact-checkers. A prominent example of this integrated approach is the InVID plug-in, which offers a range of functionalities, including image, video and metadata analysis, as well as the ability to benchmark different reverse image search engines [27]. These features were developed following a user-centred approach and have made the tool one of the most cited and used by fact-checkers [7]. Nevertheless, fact-checkers have expressed a need for new tools to address the increasing complexity of media verification, which comes with the increasing complexity of misleading media content. To address these specific needs, fact-checkers have emphasised the importance of a human-in-the-loop approach that allows them to exercise essential human skills involving critical thinking and contextual analysis [7].

Aiming to answer these needs, this paper introduces VeriDash, an AI-based and open-source dashboard designed to combine multimedia verification tools into one cohesive platform to support additional tasks not taken in charge by other existing solutions. With features such as automated transcription, geolocation and a user-friendly interface, the tool streamlines the entire multimedia verification process while prioritising the user experience. The strength of VeriDash lies in its ability to enhance human agency, recognising that effective fact-checking requires a balance between technology and human expertise [9]. Designed to be adaptable, VeriDash allows for easy updates and integration of new libraries and tools as they become available, emphasising collaboration between researchers and fact-checkers. This flexibility ensures that the dashboard remains a relevant and effective resource in an ever-evolving landscape.

2 AI in Multimedia Verification

Fact-checking is a multifaceted process that involves not only verifying claims after they have been publicly disseminated but also a comprehensive approach that includes monitoring social media, identifying and selecting claims, verifying their content, gathering evidence, making a judgement on the factuality of the claim, and writing and disseminating a story [8,17]. Each stage of this process is characterised by its time-consuming nature, reflecting the investigative rigour required to ensure the accuracy and reliability of the published fact-check [10]. In this process, the use of AI technologies is becoming increasingly important to support fact-checking routines and workflows. Natural language processing (NLP) combined with machine learning holds significant potential for verifying facts against reputable sources and identifying patterns indicative of falsehoods [5]. Although AI is also proving valuable for image verification and video authentication, from an end-user perspective, fact-checkers face challenges in adopting AI tools, mainly due to a lack of awareness and trust. Furthermore, while the tools developed by computer scientists propose advanced solutions, they are often under-utilised because fact-checkers are unaware of their existence or uncertain about their reliability [7,17]. In addition, despite the availability of robust image and video verification tools, their use is practically limited because of a lack of alignment with the needs, practices and professional values of fact-checkers [6]. Furthermore, many existing tools have a narrow focus and were developed outside of journalism, making them difficult to adapt for tailored solutions [17].

Professional fact-checkers consider current AI tools to be fast and cost-effective, but also less accurate and explainable than humans. In many cases, they found that AI technologies cannot fully replicate human skills, which limits their effectiveness in automating fact-checking in practice [5]. As a result, fact-checkers have expressed a need for tools that enhance their soft skills, such as critical thinking and judging newsworthiness - human know-how that cannot be automated [7]. Enabling a human-in-the-loop approach not only addresses these needs, but also serves as a valuable strategy for promoting trustworthy AI, especially given the critical role of trust in the adoption of AI technologies [11,20]. Moreover, fact-checkers often perceive algorithmic accuracy as inferior to human accuracy, making human oversight essential to increase the overall reliability and effectiveness of the fact-checking process [15].

Collaboration between researchers and practitioners is, therefore, critical to developing tools that effectively address the needs of fact-checkers and enhance the integration of AI in the verification process [17]. This partnership is particularly important as fact-checkers often use the same tools for multimedia verification, which presents unique challenges [7]. Evaluating different types of content - such as images, video and audio - requires different verification techniques and contextual analysis to accurately assess authenticity. It also relies heavily on the expertise of fact-checkers, as technology should primarily support human decision-making rather than replace it [18]. Effective collaboration between researchers and fact-checkers enhances the overall verification process and helps avoid the pitfalls of tools whose complex behaviour can undermine trust - an essential factor in technology adoption [23].

AI techniques such as micro-targeting, user profiling and deepfakes are increasingly being used to manipulate public opinion, undermine trust in institutions or influence elections [12]. This rapid evolution of multimedia content, exacerbated by the advanced capabilities of generative AI, further complicates the verification work of fact-checkers. Deepfakes in particular pose a significant challenge due to their advanced technology, which can easily distort public perception and create false evidence. Though less common than simpler disinformation tactics, deepfakes are especially deceptive, undermining public trust in news media [29]. For fact-checkers, decontextualisation is the most intricate issue, whether in deepfakes or cheapfakes – which involves minimum editing of the video [29]. To tackle all these challenges, fact-checkers need platforms that combine existing functionalities and offer greater flexibility, adaptability and ease of use, especially since multimedia contents are becoming increasingly complex. Digital forensics plays a crucial role in this context, focusing on detecting traces of artificial intelligence in these images. While inconsistencies in scenes and objects can sometimes reveal synthetic origins, advances in synthesis techniques may render these inconsistencies invisible [2].

For computer scientists, the key struggle is to deal with multiple modalities - such as text, audio, video and images - as well as the complexities of interpreting context, sarcasm and irony [22]. At the same time, several technological issues still need to be addressed in terms of accuracy, accessibility and interoperability [29]. Therefore, moving towards a human-centred approach remains a challenge for building trust, which can be supported not only by improving the usability of tools, but also by providing more explainable systems to enhance human understanding [20]. It is also about creating a balance between AI-augmented human tasks and AI-automated routines tasks [20]. For example, Polzehl et al. identified five significant functions for human-machine collaboration: collecting data, ensuring data quality, monitoring system performance, providing context for AI results, and providing feedback to improve AI systems. The authors highlighted the importance of considering explainability, bias and usability in a human-centred approach [22].

Lessons learned from this literature review suggest that while AI technologies have potential to improve fact-checking, particularly in multimedia verification, challenges remain. These include the limitations of AI in replicating human skills, a lack of awareness and trust among fact-checkers of the tools available, and the need for support in human soft skills such as critical thinking. Therefore, effective collaboration between researchers and practitioners is essential to develop tailored solutions that consider a human-in-the-loop approach for building trustworthy AI. All these insights have informed the design of VeriDash, which supports the verification process by complementing existing tools and enhancing human soft skills alongside technical capabilities.

3 The Dashboard Design

VeriDash is a dashboard designed to help fact-checkers verify video content by providing them with advanced AI capabilities that incorporate a human-in-the-loop approach. This ensures that human expertise is integrated with AI tools for effective decision-making. The dashboard provides fact-checkers with a set of complementary tools organised into a cohesive interface for verifying the authenticity of video content. This perspective also aims to avoid fragmented usage by allowing fact-checkers to seamlessly navigate different tasks without switching between multiple tools, thereby streamlining their workflow. Key functionalities of VeriDash include:

- Geolocation: This tool helps fact-checkers verify where a video was filmed by comparing visual cues in the footage with known geographic data. The location of the video is displayed on an interactive map.
- Automated audio transcription: The dashboard includes an automated speech-to-text system that converts audio into text, helping to verify content quickly and efficiently.
- Machine translation: A built-in translation tool allows the transcription to be translated into multiple languages, enabling fact-checkers to work across different language contexts.
- Information retrieval: The tool includes a search engine that allow for contextualisation.
- Frame stitching: The dashboard has a feature for stitching frames of a video together to create a bigger picture, creating a better overview of a chaotic situation.
- Object detection: The tool scans the video looking for objects of interest. The fact-checker can then see where these are located in the video, or use the images created from this to perform reverse image searches.

VeriDash is designed as a dashboard that integrates various widgets for an interactive user experience. Shown in Figure 1 is a screenshot of it. On the top row, the leftmost widget allows users to upload and play an input video, with the option to view it in full-screen mode in a new tab. The next widget detects the video's location. Currently, geolocation is determined based on metadata, but it can be extended using content-based methods. Users can navigate the map and switch between different map services (for example, Google, Bing, or Yandex) for a customised view. In the bottom row, the leftmost widget displays extracted keyframes from the video, while the adjacent widget provides their corresponding metadata. A built-in translation tool in the third widget allows transcription to be translated into multiple languages, facilitating cross-language fact-checking. Finally, in the rightmost column, widgets for object detection are available, and users can explore further by performing reverse image searches, by combining with external services, for deeper analysis.

The tool is specifically designed to meet the needs of fact-checkers by prioritising user-centred principles. As a result, it has a simple and intuitive interface that enhances usability and ease of use. By following responsible design principles, the tool ensures that users can easily access the functionality they need without unnecessary complexity. This approach aims not only to improve the efficiency of the fact-checking process, but also to build trust in the technology. Another major advantage of a user-centred approach is alleviating concerns of AI inaccuracy. VeriDash only employs AI tools in ways to improve productivity of the fact-checker, but never to draw direct conclusions about the situation.

The user interface is built in a highly modular way to enable tailoring of features to different preferences in the future. Similarly, the specific implementation

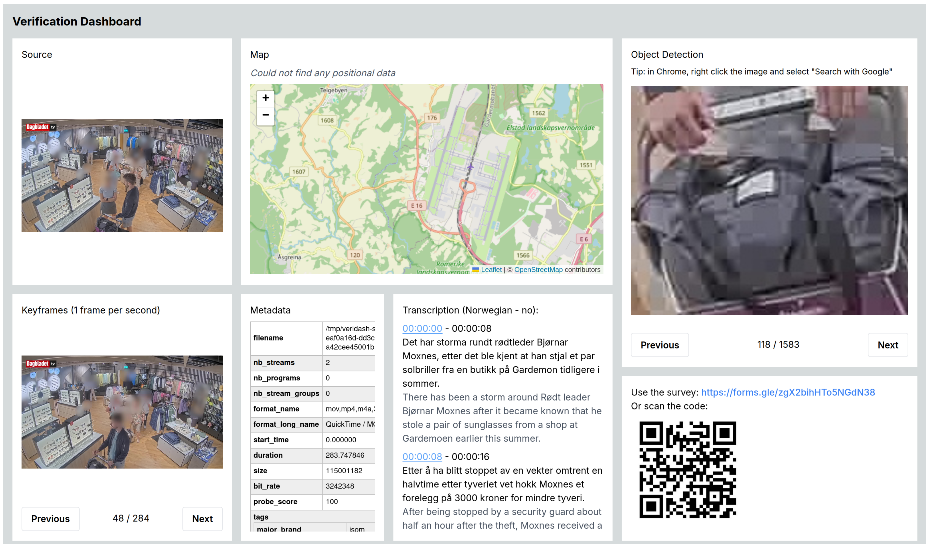

Fig. 1. The dashboard features widgets for video upload and playback, location and individual detection, map navigation, keyframe extraction, metadata display, translation, and object detection with reverse image search for deeper analysis.

of the software is also highly modularized. This allows the addition of new techniques as the field develops without requiring a significant re-engineering effort.

3.1 Implementation details

The source code for VeriDash is availiable at github.com/skivdal/veridash. It is implemented as a React application, communicating with a server over a WebSocket. This model allows us to dynamically update the content that is displayed to the user as the server works through tasks. On the backend, the server is implemented as two discrete Python services, one handling interaction with the frontend (webserver), the other handling the running of analysis tasks (worker). This split in the backend is done for reasons of concurrency and scalability.

We also depend on a few open-source services, PostgreSQL as our main database, Valkey for locks and communication between the webserver and worker services, and MinIO for file storage and upload/download handling.

Concretely, we implement each module (Source, Map, Object Detection, Keyframes, Metadata, and Transcription) as a distinct React component. All of these components receives a VideoId when a video is successfully uploaded, and uses the same custom React hook to request its work to be done over the same WebSocket connection. Specializations to this hook has been made to facilitate more complex requests, such as video upload. This system makes future development easy, as you would just have to make a new component using that

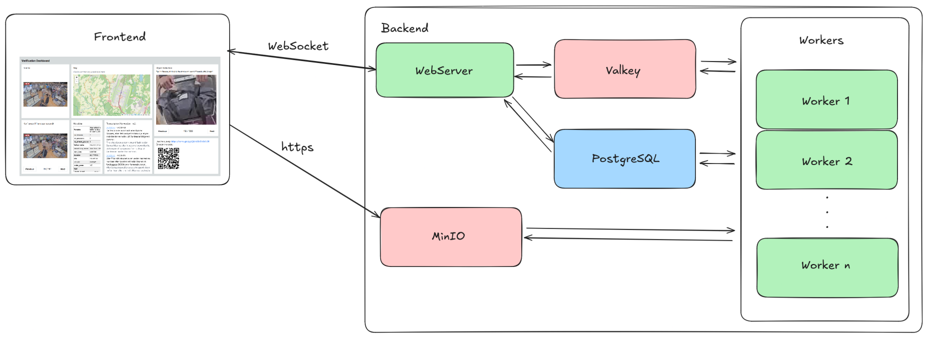

Fig. 2. High level overview of the VeriDash system architecture.

same hook for backend interaction. The hook also allows us to have live loading indicators in the future, although this has not been implemented yet.

The backend webserver service asynchronously receives messages, and serves results for each module from cache if they've been done before, or alternatively passes the request on to the worker service. It then asynchronously waits for results from running worker tasks, caching them and passing them on to the frontend. Some tasks are dependent on other tasks, in this case, receiving a response from a task causes the webserver to request the worker to run dependent tasks. The webserver service is implementing using the FastAPI library, and the worker service is implemented using the Celery library, with task queuing and responses handled with Valkey.

The Source task allows for uploading a video for further analysis. The frontend will compute a sha256 hash of the video, and the backend will add this and provision a filename in the database. It will then create URLs for upload and download from MinIO, passing these back to the frontend. The frontend will upload the video, and distribute the provisioned video ID to the other components, causing them to request other tasks to be ran against the video.

The task cache is stored in PostgreSQL, and is used for tasks other than Source, Keyframes, and Object Detection. Cache hits are detected by matching video sha256 hashes. The reason for the mentioned tasks not being cached is that they produce files in MinIO, indexed in the database. Responses are presigned MinIO download URLs, which are inherently temporary.

All tasks require the video to be downloaded to the Worker's filesystem. Often, many tasks run at the same time, so we make atomic locks in Valkey to ensure only one task downloads the file at a time, allowing all tasks to use the same downloaded file. Similarly, GPU VRAM is a limited resource for the tasks that leverage the GPU (Transcription and Object Detection), and we also use a Valkey-based lock for limiting GPU access to one task at a time.

The Keyframes task requires the video to be downloaded. It uses the ffmpeg library to extract one JPEG-frame per second of the video, provisioning filenames connected to the Source video for them in PostgreSQL, then uploading them to

MinIO. Finally, we return presigned MinIO download URLs for the frontend to display to the user. As previously mentioned, these URLs are not stored in cache.

The Map task requires the video to be downloaded. It uses the ffmpeg library to extract metadata from the file, and looks for some common tags for location data. This is returned to the webserver service, added to the cache, and returned to the client. The metadata task works in a very similar way, using ffmpeg to extract metadata without doing any kind of selection, returning to the webserver, caching, and sending to the client.

The Transcription task requires the video to be downloaded. It requires using the GPU, and thus waits for that resource to be free before starting. Transcription is done using the Whisper medium model from OpenAI, supporting a wide variety of languages. This is ran locally on the GPU. If the transcript is any other language than English, it will be ran through GPT using the API for translation to English. This requires OpenAI API credentials to be provided. Transcription and translation texts are returned to the webserver, cached, and sent to the client.

The Object Detection task requires the keyframes extracted from the video to be downloaded, using a similar mechanism to the deduplicated video download. It also requires using the GPU, and needs to wait for that to be free to start. Object detection is ran on every keyframe, with potentially multiple objects being detected per frame. We leverage the YOLO-v8 world model, and create new images from the content of each detected object. These object images then get filenames provisioned by PostgreSQL, uploaded to MinIO, and the task result is returned to the webserver, not cached, and pushed to the client.

This task-based architecture for workers with deduplicated downloads, caching and dependencies handled by the webserver, and messages passed along through Valkey creates lots of common structure. It allows for a new type of task to be added as simply as creating a function in the worker application, modelling its dependency relations and caching behaviour, and adding a response type to the frontend React hook. Having the ability to do simple scalable services has been a main priority in this project, as this allows fast iteration in the future to keep up integrating new verification methods to the dashboard.

4 Evaluation

A preliminary evaluation was conducted to assess the functional features of the dashboard, its compliance with Nielsen's heuristic evaluation [19], and user feedback. This evaluation included gathering insights from professional fact-checkers, testing the function, and following Nielsen's heuristics. Nielsen's principles are particularly prioritized, as they align closely with the core purpose of the dashboard: to facilitate the verification process. These evaluations, commonly used to build an effective user experience, support iterative improvements based on user needs and practices [24].

4.1 Functional Evaluation

The functional evaluation of the dashboard highlights its effectiveness in multimedia verification through a range of powerful tools, most of which are off-the-shelf services/solutions. Exploring such tools ensures the system remains dynamic, with improvements continually introduced as these tools evolve. This combination of tools ensures the dashboard remains a robust and adaptable solution for multimedia verification.

In our evaluation, we explore mainly real cases from the two on-going wars in Gaza and Ukraine. At this stage, only a small subset of the dataset has been used in our evaluation. We report here the results on the short videos (less than 5 minutes each) in our evaluation.

- Geolocation: The tool effectively verifies a video's location using metadata and visual cues, displayed on an interactive map.
- Automated audio transcription: The system quickly converts audio to text, utilising various transcription tools (in the current version, OpenAI Whisper [25] is used), though shorter clips (audio less than 2 seconds) may affect accuracy. Ongoing updates ensure improvements in speed and precision.
- Machine translation: The tool supports cross-language verification by incorporating multiple translation technologies though accuracy may vary with brief audio samples.
- Object detection and reverse image search: The integration of object detection with reverse image search provides a flexible way to analyse video content, allowing fact-checkers to identify objects and perform further investigations efficiently.

4.2 Nielsen's Heuristic Evaluation

To analyse VeriDash in a more uniform manner, a heuristic evaluation with the goal of generic usability testing - Nielsen's heuristic evaluation [19] - is used. We are using the general principles of Nielsen heuristics to define the product's usability. We use a set of criteria to assess where our product is at in terms of usability, on a scale of 1 (lowest score) to 5 (highest score), as indicated in Table 1.

4.3 User feedback

The preliminary feedback received from fact-checkers and journalists in Norway, Sweden, Finland, Denmark, and the BBC in the UK confirms that exploring AI-based tools with a human-in-the-loop approach is the right direction for advancing multimedia verification. This feedback aligns with current research underscoring the need for AI tools that not only enhance verification capabilities but also complement human expertise. Additionally, it reflects key principles for building responsible AI tools in journalism by considering user needs, values, and practices.

Table 1. Own usability inspection based on Nielsen's heuristic evaluation.

No.	Principle	1	2	3	4	5
1	Visibility of system status			X		
2	Match between the system and the real world					X
3	User control and freedom					X
4	Consistency				X	
5	Error prevention				X	
6	Recognition rather than recall				X	
7	Flexibility and efficiency of use			X		
8	Aesthetic and minimalism in design			X		
9	Recognition, diagnosis, and recovery from errors				X	
10	Help and Documentation			X		

Users particularly appreciated VeriDash's speed and keyframe extraction capabilities, especially when integrated with the InVid plugin[27]. As one user expressed, the keyframe extraction feature is a *"vital part of our workflow,"* underscoring its importance in streamlining verification tasks. The metadata and transcription functionalities were also highly regarded, with these tools proving particularly effective for geolocation and video verification. Moreover, the image stitching feature was described as a unique asset, with one user noting, *"I don't know of any services worldwide that does this."* Overall, the combination of speed, essential functionality, and compatibility with tools like InVid positions VeriDash as a valuable resource in verification, closely aligned with the workflows of fact-checking professionals.

However, several areas for enhancement were identified. Users raised concerns regarding metadata accuracy, specifically the "creation time" field, which currently does not reflect the actual video creation date and affects the tool's reliability in fact-checking tasks. Technical issues, such as inconsistent performance in the map and object detection tools and occasional transcription errors, were also noted. Fixing these bugs will significantly enhance reliability, with one user remarking, *"With these few changes, it will become immensely useful to verify videos."* Some feedback also touched on the interface design, suggesting a more modern layout for improved usability. However, as this current version is primarily a functional test prototype, these design suggestions will be noted for future updates.

In summary, VeriDash has built a strong foundation with its speed, keyframe extraction, metadata, and transcription and translation features, which are highly valued by users. Unique capabilities, like image stitching, position it as a valuable tool in verification. Addressing technical issues and metadata accuracy, along with incorporating user-suggested improvements, will be essential for reliability and functionality. As the tool moves beyond functional testing, a more user-friendly interface and expanded capabilities, such as video URL inputs and metadata verification integrations, will further enhance its appeal. With these improvements, VeriDash can evolve into a robust, competitive tool, solidifying its role in fact-checking workflows.

> A feedback from a professional fact-checker.
>
> I tested on a few harmless videos for transcription, and it performs really well on geolocation with metadata, transcription and keyframes. Already an improvement for our verifications.
>
> (Henrik B. Vold, Associate Director, Norwegian Institute of Journalism)

5 Conclusions

In this paper, we introduced VeriDash, an open-source dashboard that leverages AI-based technologies to enhance the multimedia verification process for fact-checkers. The initial version of VeriDash was presented alongside preliminary evaluation results, which received highly positive feedback, affirming the effectiveness of the proposed approach. VeriDash offers advanced features such as automated transcription, geolocation, and a user-friendly interface, streamlining the fact-checking process while maintaining ease of use. By incorporating a human-in-the-loop approach, VeriDash combines technological efficiency with human expertise, fostering trusted and responsible AI-driven verification to support and enhance fact-checking efforts. In addition, the dashboard is designed with input from research users to ensure it meets the practical needs of fact-checkers. This emphasis on user-centred design enhances the tool's relevance and usability in real-world scenarios, contributing to more effective fact-checking processes. Further development will incorporate user suggestions to improve the tool, particularly concerning user interface enhancements.

Acknowledgements

This research was supported by NORDIS, European Horizon 2020 grant number 101158604. The authors extend their gratitude to the fact-checkers at Faktisk.no for their valuable discussions and active involvement from the initial stage through to the testing phase of the project.

References

1. Bailey, T.C., Hsieh-Yee, I.: Combating the sharing of false information: History, framework, and literacy strategies. Internet Reference Services Quarterly **24**(1-2), 9–30 (2020)
2. Cardenuto, J.P., Yang, J., Padilha, R., Wan, R., Moreira, D., Li, H., Wang, S., Andaló, F., Marcel, S., Rocha, A., et al.: The age of synthetic realities: Challenges and opportunities. APSIPA Transactions on Signal and Information Processing **12**(1) (2023)

3. Carnahan, D., Bergan, D.E.: Correcting the Misinformed: The Effectiveness of Fact-checking Messages in Changing False Beliefs. Political Communication **39**, 166–183 (2021), `https://api.semanticscholar.org/CorpusID:238730423`
4. Dang-Nguyen, D.T., Khan, S.A., Riegler, M., Halvorsen, P., Tran, A.D., Dao, M.S., Tran, M.T.: Overview of the Grand Challenge on Detecting Cheapfakes at ACM ICMR 2024 . In: Proceedings of the 2024 International Conference on Multimedia Retrieval. pp. 1275–1281 (2024)
5. Das, A., Liu, H., Kovatchev, V., Lease, M.: The state of human-centered NLP technology for fact-checking. Information processing & management **60**(2), 103219 (2023)
6. Dierickx, L., Lindén, C., Opdahl, A.: Automated fact-checking to support professional practices: systematic literature review and meta-analysis. International Journal of Communication (2023)
7. Dierickx, L., Lindén, C.G.: Journalism and Fact-Checking Technologies: Understanding User Needs. Communication+ 1 **10**(1) (2023)
8. Dierickx, L., Van Dalen, A., Opdahl, A.L., Lindén, C.G.: Striking the Balance in Using LLMs for Fact-Checking: A Narrative Literature Review. In: Multidisciplinary International Symposium on Disinformation in Open Online Media. pp. 1–15. Springer (2024)
9. Himma-Kadakas, M., Ojamets, I.: Debunking false information: investigating journalists' fact-checking skills. Digital journalism **10**(5), 866–887 (2022)
10. Juneja, P., Mitra, T.: Human and technological infrastructures of fact-checking. Proceedings of the ACM on Human-Computer Interaction **6**(Cscw2), 1–36 (2022)
11. Kavtaradze, L., Kalsnes, B.: AI-powered fact-checking: Strategic framing of AI use. Strategic Communication–Contemporary Perspectives p. 9177 (2024)
12. Kertysova, K.: Artificial intelligence and disinformation: How AI changes the way disinformation is produced, disseminated, and can be countered. Security and Human Rights **29**(1-4), 55–81 (2018)
13. Khan, S.A., Dang-Nguyen, D.T.: Deepfake Detection: Analysing Model Generalisation Across Architectures, Datasets and Pre-Training Paradigms. IEEE Access (2023)
14. Khan, S.A., Sheikhi, G., Opdahl, A.L., Rabbi, F., Stoppel, S., Trattner, C., Dang-Nguyen, D.T.: Visual user-generated content verification in journalism: An overview. IEEE Access **11**, 6748–6769 (2023)
15. Liu, H., Das, A., Boltz, A., Zhou, D., Pinaroc, D., Lease, M., Lee, M.K.: Human-centered NLP Fact-checking: Co-Designing with Fact-checkers using Matchmaking for AI. arXiv preprint arXiv:2308.07213 (2023)
16. Micallef, N., Armacost, V., Memon, N., Patil, S.: True or false: Studying the work practices of professional fact-checkers. Proceedings of the ACM on Human-Computer Interaction **6**(Cscw1), 1–44 (2022)
17. Nakov, P., Corney, D., Hasanain, M., Alam, F., Elsayed, T., Barrón-Cedeño, A., Papotti, P., Shaar, S., Da San Martino, G., et al.: Automated Fact-Checking for Assisting Human Fact-Checkers. In: Proceedings of the Thirtieth International Joint Conference onArtificial Intelligence,{IJCAI-21}. pp. 4551–4558. International Joint Conferences on Artificial Intelligence Organization (2021)
18. Nguyen, A.T., Kharosekar, A., Krishnan, S., Krishnan, S., Tate, E., Wallace, B.C., Lease, M.: Believe it or not: designing a human-ai partnership for mixed-initiative fact-checking. In: Proceedings of the 31st annual ACM symposium on user interface software and technology. pp. 189–199 (2018)
19. Nielsen, J.: Usability engineering. Morgan Kaufmann (1994)

20. Opdahl, A.L., Tessem, B., Dang-Nguyen, D.T., Motta, E., Setty, V., Throndsen, E., Tverberg, A., Trattner, C.: Trustworthy journalism through AI. Data & Knowledge Engineering **146**, 102182 (2023)
21. Osmundsen, M., Bor, A., Vahlstrup, P.B., Bechmann, A., Petersen, M.B.: Partisan polarization is the primary psychological motivation behind political fake news sharing on Twitter. American Political Science Review **115**(3), 999–1015 (2021)
22. Polzehl, T., Schmitt, V., Feldhus, N., Meyer, J., Möller, S.: Fighting Disinformation: Overview of Recent AI-Based Collaborative Human-Computer Interaction for Intelligent Decision Support Systems. Visigrapp (2: Hucapp) pp. 267–278 (2023)
23. Procter, R., Arana-Catania, M., He, Y., Liakata, M., Zubiaga, A., Kochkina, E., Zhao, R.: Some Observations on Fact-Checking Work with Implications for Computational Support. arXiv e-prints pp. arXiv–2305 (2023)
24. Quiñones, D., Rusu, C., Rusu, V.: A methodology to develop usability/user experience heuristics. Computer standards & interfaces **59**, 109–129 (2018)
25. Radford, A., Kim, J.W., Xu, T., Brockman, G., McLeavey, C., Sutskever, I.: Robust speech recognition via large-scale weak supervision. In: International conference on machine learning. pp. 28492–28518. Pmlr (2023)
26. Singer, J.B.: Border patrol: The rise and role of fact-checkers and their challenge to journalists' normative boundaries. Journalism **22**(8), 1929–1946 (2021)
27. Teyssou, D., Leung, J.M., Apostolidis, E., Apostolidis, K., Papadopoulos, S., Zampoglou, M., Papadopoulou, O., Mezaris, V.: The InVID plug-in: web video verification on the browser. In: Proceedings of the first international workshop on multimedia verification. pp. 23–30 (2017)
28. Van Raemdonck, N., Meyer, T.: Why disinformation is here to stay. A socio-technical analysis of disinformation as a hybrid threat. In: Addressing Hybrid Threats, pp. 57–83. Edward Elgar Publishing (2024)
29. Weikmann, T., Lecheler, S.: Cutting through the hype: Understanding the implications of deepfakes for the fact-checking actor-network. Digital Journalism pp. 1–18 (2023)