

December 2024

MaskAnyone Toolkit : Offering Strategies for Minimizing Privacy Risks and Maximizing Utility in Audio-Visual Data Archiving

Babajide Alamu Owoyele
Hasso Plattner Institute, babajide.owoyele@gmail.com

Martin Schilling
University of Potsdam, martin.schilling@student.hpi.uni-potsdam.de

Rohan Sawahn
University of Potsdam, rohan.sawahn@student.hpi.uni-potsdam.de

Niklas Kaemer
University of Potsdam, niklas.kaemmer@student.hpi.uni-potsdam.de

Pavel Zherebenkov
University of Potsdam, pavel.zherebenkov@student.hpi.uni-potsdam.de

See next page for additional authors

Follow this and additional works at: <https://aisel.aisnet.org/icis2024>

Recommended Citation

Owoyele, Babajide Alamu; Schilling, Martin; Sawahn, Rohan; Kaemer, Niklas; Zherebenkov, Pavel; Verma, Bhuvanesh; Pouw, Wim; and de Melo, Gerard, "MaskAnyone Toolkit : Offering Strategies for Minimizing Privacy Risks and Maximizing Utility in Audio-Visual Data Archiving" (2024). *ICIS 2024 Proceedings*. 1. https://aisel.aisnet.org/icis2024/adv_theory/adv_theory/1

This material is brought to you by the International Conference on Information Systems (ICIS) at AIS Electronic Library (AISeL). It has been accepted for inclusion in ICIS 2024 Proceedings by an authorized administrator of AIS Electronic Library (AISeL). For more information, please contact elibrary@aisnet.org.

Presenter Information

Babajide Alamu Owoyele, Martin Schilling, Rohan Sawahn, Niklas Kaemer, Pavel Zherebenkov, Bhuvanesh Verma, Wim Pouw, and Gerard de Melo

MaskAnyone Toolkit: Offering Strategies for Minimizing Privacy Risks and Maximizing Utility in Audio-Visual Data Archiving

Completed Research Paper

Babajide Alamu Owoyele

Hasso Plattner Institute,
Prof.-Dr.-Helmert-Straße 2-3, 14482
Potsdam, Germany
babajide.owoyele@hpi.de

Martin Schilling

University of Potsdam
Prof.-Dr.-Helmert-Straße 2-3, 14482
Potsdam, Germany
martin.schilling@student.hpi.uni-
potsdam.de

Rohan Sawahn

University of Potsdam
Prof.-Dr.-Helmert-Straße 2-3, 14482
Potsdam, Germany
rohan.sawahn@student.hpi.uni-
potsdam.de

Niklas Kaemer

University of Potsdam
Prof.-Dr.-Helmert-Straße 2-3, 14482
Potsdam, Germany
niklas.kaemmer@student.hpi.uni-
potsdam.de

Pavel Zherebenkov

University of Potsdam
Prof.-Dr.-Helmert-Straße 2-3, 14482
Potsdam, Germany
pavel.zherebenkov@student.hpi.uni-
potsdam.de

Bhuvanesh Verma

University of Potsdam
Am Neuen Palais 10, 14469
Potsdam, Germany
bhuvanesh.verma@uni-potsdam.de

Wim Pouw

Donders Centre for Cognition,
Radboud University, Nijmegen,
Thomas Van Aquinostraat 4
6525GD Nijmegen, The Netherlands
wim.pouw@donders.ru.nl

Gerard de Melo

Hasso Plattner Institute, University of
Potsdam
Prof.-Dr.-Helmert-Straße 2-3, 14482
Potsdam, Germany
gerard.demelo@hpi.de

Abstract

This paper introduces MaskAnyone, a novel toolkit designed to navigate some privacy and ethical concerns of sharing audio-visual data in research. MaskAnyone offers a scalable, user-friendly solution for de-identifying individuals in video and audio content through face-swapping and voice alteration, supporting multi-person masking and real-time bulk processing. By integrating this tool within research practices, we aim to enhance data reproducibility and utility in social science research. Our approach draws on Design Science Research, proposing that MaskAnyone can facilitate safer data

sharing and potentially reduce the storage of fully identifiable data. We discuss the development and capabilities of MaskAnyone, explore its integration into ethical research practices, and consider the broader implications of audio-visual data masking, including issues of consent and the risk of misuse. The paper concludes with a preliminary evaluation framework for assessing the effectiveness and ethical integration of masking tools in such research settings.

Keywords: Audio-visual data, open science, de-identification strategies, design science research, data sharing

Introduction

Audio-visual data with human subjects is crucial for behavioral sciences and linguistics and provides insights into human behavior and communication (Abney et al. 2018; Cienki 2016; D’Errico et al. 2015; Gregori et al. 2023). However, including identifiable human data raises ethical and privacy concerns (Abay et al. 2019; Benson et al. 2020; Bishop 2009; Jarolimkova and Drobikova 2019; Jeng et al. 2016; Johannesson and Perjons 2021). GDPR outlines frameworks for handling such data responsibly (Nautsch et al. 2019). In the spirit of open science, more platforms, artifacts, and tools are needed to balance privacy and data sharing, especially in social sciences and humanities (Hunyadi et al., 2016; Qian et al., 2018). Using audio-visual data in social and behavioral sciences requires a deep understanding of ethical, legal, and methodological aspects. However, using such data is necessary to ensure research transparency and reproducibility, with recent work arguing that retaining interview data should be the default (Resnik et al. 2024). Incorporating audio-visual data in research introduces a host of complexities, from data collection and analysis to interpretation. This can impact the reliability and validity of findings, raising concerns about potential biases, misinterpretations, and inadvertent capture of unrelated audiovisual data. Aligning and harmonizing audio-visual inputs is particularly challenging, especially with recent trends in generating audio-visual data. Considering these challenges, the proposed toolkit, MaskAnyone, becomes a necessity. Researchers must take the lead in considering the unique challenges associated with audio-visual data to safeguard their research’s integrity, validity, and ethical conduct.

Building on existing work (Khasbage et al. 2022; Owoyele et al. 2022), we propose MaskAnyone, a toolkit for de-identifying individuals using audio-visual data. MaskAnyone offers multi-person masking, real-time bulk processing, and a user-friendly interface - essential for handling large datasets efficiently while reducing privacy risks. With accessibility in mind, social scientists can use MaskAnyone techniques to mask identifiable information, such as face-swapping and auditory elements. This modular approach allows researchers who do not write code to customize the anonymity level based on the sensitivity of the data and available computing resources. The toolkit’s flexibility is essential as it enables researchers to use it on personal machines for smaller projects or scale it up to server-based environments for institutional research involving larger datasets. This adaptability allows researchers to tailor the toolkit to their needs and resources. Considering the challenges and related issues above, as well as the opportunities to leverage developments in information systems and computer science, our paper is guided by the following research questions and objectives:

1. How can researchers effectively navigate/balance subjects’ privacy with the utility of audio-visual data?
2. What techniques can be employed to ensure the ethical use of audio-visual data in compliance with stringent regulatory frameworks?
3. What implications does masking audio/visual data have for generating and analyzing more synthetic data, and what evaluation-related challenges remain in designing and iterating such tools?

Drawing explicitly on the design science research framework (Hevner et al. 2004), we iteratively developed *a scalable toolkit to de-identify audiovisual data. Co-developed with researchers and data stewards, the tool can be integrated into current research practices to promote ethical data sharing.* MaskAnyone aims to navigate the privacy risks associated with audio-visual data sharing. Distinct from existing solutions (Khasbage et al. 2022; Owoyele et al. 2022), the toolkit also supports exporting body and face pose data as JSON and Csvs along with multi-person masking and real-time bulk processing while offering a modular,

user-friendly interface. It incorporates various masking techniques for visual (e.g., face-swapping) and auditory elements of identifiable information and is designed to be scalable for handling large datasets. By providing customizable options, MaskAnyone allows researchers to balance privacy and data utility, and it is specifically designed for scalability. Tailoring to data sensitivity requirements and computing needs, MaskAnyone can run on personal machines or a (secured) server with a large amount of computing power to accommodate multiple researchers with large datasets. Our paper proceeds as follows: First, we delve into recent literature on the masking of audiovisual data and discuss the design methodology used. Next, we detail the techniques and algorithms we employed for de-identification in video and audio domains. We present our preliminary evaluations, describing the system, sample datasets, and proposed evaluation metrics. The discussion section offers a broader outlook on how we want to evaluate toolkits like MaskAnyone, ending with broader implications, opportunities, and issues of masking practices.

Related Work/Problem Space

Masking has become even more relevant due to the development of MaskAnyone, with several implications that need to be considered. One of the issues that need to be addressed is consent. Regulatory frameworks like GDPR require that human subjects be informed about de-identification techniques. Therefore, masking previously collected data may require additional consent. MaskAnyone is designed to safeguard privacy, but there is also potential for misuse by enabling the creation of deepfakes. However, it might also generate data to detect deepfakes. There are also questions about what degree of masking is sufficient and how to determine it. For instance, sensitive contexts like healthcare may require complete anonymization and additional safeguards to enable audio-visual data sharing. As a field, we need to be able to judge and validate what degree of masking is considered a reduction of identifiable information, de-identification, pseudo-anonymization, or complete irreversible anonymization. This question is not just an abstract theoretical question, as failure to comply with regulatory privacy frameworks can have direct legal repercussions. These considerations form the challenging context within which audio-visual masking tools must operate. Nonetheless, when integrated into ethical procedures already upheld by research institutes, toolkits like MaskAnyone promise to support the safe sharing of multimodal language data.

Person de-identification in Video

For video de-identification (Gafni et al. 2019; Kasturi and Ekambaram 2014), we distinguish two principal strategies: hiding and masking. The hiding strategy focuses on obscuring or eliminating video segments containing personally identifiable information, thereby enhancing privacy. In contrast, the masking strategy substitutes the individual with an alternate representation, preserving essential attributes for the video's utility while masking the individual's identity. Recently, tools such as Masked-Piper (Owoyele et al. 2022) and Red Hen Anonymizer (Khasbage et al. 2022) have been developed to use hiding and masking strategies; however, their usability remains challenging, and they lack modularity, adding new modules for privacy masking and multimodal analytics. Regardless of the strategy employed, a preliminary step of person detection is needed. This can be achieved through generic object detection models like 'You Only Look Once' (YOLO) (Redmon et al. 2016) or more specialized models like MediaPipe's BlazePose (Bazarevsky et al. 2020). Both categories of models offer robust capabilities, catering to a range of requirements and applications.

Hiding

Upon detecting a person in a video, the hiding strategy offers multiple options for concealment. One common approach is to obfuscate the identified area using techniques such as Gaussian blurring, pixelation, or Laplacian edge detection, which can be implemented using image processing libraries such as OpenCV (Bradski 2000; OpenCV 2023). While machine learning-based methods such as PiDiNet (Su et al. 2021) and UAED (Zhou2023) may offer enhanced quality and privacy, they do not necessarily guarantee robust anonymity, particularly for videos featuring well-known individuals (Hasan et al. 2018; Lander et al. 2001). Another tactic involves overlaying the detected area with a non-translucent color, effectively hiding person-specific attributes except for potentially identifiable characteristics like height or clothing. Alternatively, one could employ inpainting techniques to estimate and replace the background behind the detected individuals, as demonstrated by projects such as Inpaint Anything (Yu et al. 2023), STTN (Zeng et al. 2020),

and E2FGVI (Dang and Buschek 2021; Ripperda et al. 2020)(Li2022). While earlier methods like STTN had limitations such as resolution constraints, recent developments like E2FGVI offer more flexibility and improved performance. The field continues to evolve, mainly focusing on responsible data sharing (Morehouse 2023), with emerging approaches like DMT (Yu 2023-2) promising even more effective results.

Masking

Hiding leads to a high loss of information, such as facial expressions or other expressive movements, that may be important for linguistic and behavioral research. Masking aims to maintain a representation of this information without retaining personally identifiable information. **Landmark Detection** for humans in monocular video data refers to identifying and localizing specific points on the human body, such as joints or facial features. The task can be further divided into 2D and 3D Landmark detection. Several models offer advanced features for landmark detection but have specific limitations -- be it computational speed, accuracy, or full-body image handling. For example, OpenPose is noted for its accuracy but is computationally slower (Mroz et al. 2021). AlphaPose, an open-source model, has gained attention for innovative approaches like Symmetric Integral Keypoints Regression (Fang2022). ViTPose offers a scalable architecture for human pose estimation based on a Vision Transformer that can scale up to 1B parameters (Xu et al. 2022). MediaPipe's Holistic model was developed to mitigate limitations in full-body detection and offers an approximation of 3D positions. We, therefore, settled on using this approach for the current version of MaskAnyone. The BlazePose model underlying MediaPipe uses a two-stage approach, in which a single-shot-detection (SSD) based detector first locates the bounding boxes of people (Grishchenko and Bazarevsky 2020). Subsequently, an estimation model applies a regression approach supervised by a combined heat map/offset prediction of all key points.

MediaPipe's holistic model computes fine-grained landmarks for the complete body (including the pose, detailed hand landmarks, and face mesh). To do so, pose detection is performed first. Subsequently, a region-of-interest (ROI) detection is performed to locate the hands and face based on the detected pose. A unique transformer re-crop model is then used to improve the ROI at only 10% of the corresponding model's inference time (Grishchenko and Bazarevsky 2020). Finally, these crops are passed to the detailed hand landmarks or face mesh models. **Face swapping**, widely used in entertainment and social media, replaces one person's face with another in images or videos. Despite its potential misuse in creating malicious deep fakes, it can serve as a de-identification tool by swapping faces to obscure identities while retaining expressions. This approach is implemented in the Red Hen Anonymizer using an FSGAN-based method (Nirkin et al. 2019). DeepFakes and DeepFaceLab offer pipelines for face replacement but often require post-editing for natural results and are limited by their need for retraining for each new face pair (Korshunov and Marcel 2018; Liu et al. 2023). InsightFace could face-swap without retraining but turned commercial and withdrew public models (InsightFace 2023). The Roop project, based on InsightFace, was discontinued due to ethical concerns. Other methods, such as Thin-Plate Spline and XFace, are limited in preserving facial expressions or are unsuitable for large-scale de-identification (Balci 2005; Zhao and Zhang 2022). Facebook's de-identification approach uses a feed-forward network, maintains facial features, and does not require retraining, but it is not open-source (Gafni et al. 2019). **Avatar creation** focuses on modeling the human body in 3D, including texture generation. Open-source tools like Blender offer basic capabilities for avatar creation (cgtinker 2021). On the proprietary side, solutions such as Meshcapade, Rokoko, and UnrealEngine MetaHuman Animator provide high-quality results. Still, they are not open-source and can be costly for research (Meshcapade 2023; Rokoko 2023). Our experiments with Meshcapade revealed a discrepancy between the advertised and actual quality.

Person de-identification in Audio

Voice data inherently contains personally identifiable information through spoken content, vocal attributes, or linguistic style (Tomashenko et al. 2020). Our work zeroes in on techniques for obscuring identifiable information while retaining linguistic and prosodic elements of speech. Broadly, we explore Spectral Modification, Pitch Shifting, and Voice Conversion as standard methodologies. Spectral Modification alters speech signals at the spectral level, targeting formant frequencies or the spectral envelope. Pitch Shifting directly modifies pitch but can distort speech naturalness and paralinguistic elements. On the other hand, Voice Conversion adapts one speaker's vocal characteristics to resemble another's, offering a nuanced approach to de-identification. The Voice Privacy Challenge (Tomashenko et al. 2020) aims to advance voice

anonymization techniques. The 2020 iteration offered two baselines: one using x-vectors (Fang et al. 2019) and neural speech synthesis and another altering voice signal through a McAdams Coefficient method (Patino et al. 2020). The 2022 round (Tomashenko et al. 2022), introduced improved baselines and evaluation metrics, employing the equal error rate (EER) for privacy and word error rate for utility. A noteworthy trade-off between privacy and utility was observed. For instance, submission T04 achieved a high EER of 47.60% but had a low pitch correlation of 37%, limiting its utility. Conversely, submission T18 balanced this trade-off well, boasting an 82% pitch correlation and an EER of 20.8%. Voice changer systems can be organized into four primary taxonomies: Phonetic Models employ phonemes for voice conversion, typically initiating the process by extracting phonemes from the input speech. Statistical Models, such as Gaussian Mixture Models (Reynolds 2015) or Hidden Markov Models (Kong et al. 2020; Wang et al. 2020), can be used to represent vocal features statistically and then map them to a target speaker. Deep Learning Models utilize architectures like CNNs, RNNs, and GANs, including notable examples such as HifiGAN and StarGAN, which generate high-quality converted speech using GAN architectures (Kong2020, Wang2020). Lastly, Retrieval-based Models search and concatenate similar speech segments from a target speaker's database. HifiGAN and StarGAN use GAN-based methods to synthesize speech from mel-spectrogram features. StarGAN also offers the potential for one-shot voice conversion (Pavlakos et al. 2019). VITS incorporates a Conditional Variational Autoencoder coupled with adversarial learning for end-to-end speech synthesis (Kim et al. 2021). The Retrieval-based Voice Conversion system employs feature extractors like Crepe for F0 features and HuBERT for representing the input speech (Hsu et al. 2021; Kim et al. 2018). These are then synthesized into the final vocal output through models such as HiFiGan (Kong et al. 2020). Notably, the literature still lacks explicit benchmarks that evaluate these models regarding utility preservation and privacy assurance. This is understandable as these individual tools were not explicitly developed for our context and problem space – to help social and behavioral science researchers navigate audio-visual data-sharing practices and the risk inherent in such endeavors.

Methodology - Design Science Research

We adopted the design science methodology for designing MaskAnyone. Here, we elaborate on the design requirements guiding the toolkit construction by integrating the design science approach and insights from the literature on audio-visual data sharing in social and behavioral sciences. This ensures that MaskAnyone augments ethical, robust, and effective data management and sharing practices. We have also developed MaskAnyone to balance open science and privacy in audio-visual research based on the call for more FAIR-friendly research infrastructure and tools to support the development of thematic digital competence centers (“Roadmaps from the Three Thematic DCCs – Digital Competence Centres | NWO” n.d.; “Thematic Digital Competence Centers | NWO” 2024). In terms of an artifact and as a DSR instantiation, the toolkit is modular and extendable and enhances the capabilities of existing tools by supporting advanced features like 3D tracking and real-time processing. It offers various masking methods for sharing and secure storage scenarios, thus protecting against unauthorized access risks like data breaches. MaskAnyone meets diverse research needs while ensuring data integrity and privacy by providing an easy-to-use interface and versatile masking options.

Design Requirements

Our toolkit, MaskAnyone, is developed using the Design Science Research framework as an instantiation. (Hevner et al. 2004). We enumerate the following design requirements based on user feedback from two live demo workshops with researchers at a Dutch University(n=16), at a German University (n=10), and data stewards(n=5). We have also used insights from existing literature on multimodal analysis to justify the relevance of these requirements, although we know they are non-exhaustive depending on future scenarios and evolving stakeholder needs (in multimodal behavior research)

Requirement	Description	Source/Literature
R1 - Multimodality	Enable the masking of primary actors, backgrounds, and voice data to ensure comprehensive privacy.	Workshop/Survey

R2-Usability	Toolkit usability: intuitive enough for researchers with limited technical skills for quick deployment.	Workshop/Survey
R3-Flexibility	Flexibility in masking methods and settings to handle various video types and privacy constraints.	Workshop
R4-Multi-level	Offer different levels of de-identification to match specific use-cases and privacy needs.	Workshop/Survey
R5-Efficiency	Ensure performance efficiency on local machines, respecting hardware limitations.	Survey
R6-Sensitivity	Optimize performance for server environments, considering data sensitivities.	Survey
R7-Scalability	Scalability to support distributed video processing for large datasets.	Workshop
R8-Modularity	Extensibility to integrate new techniques seamlessly as they emerge in video processing.	Workshop/Survey

Table 1. MaskAnyone Design Science Approach and Requirements

Solution-Space: MaskAnyone Architecture and User Interaction Layout

Architecture

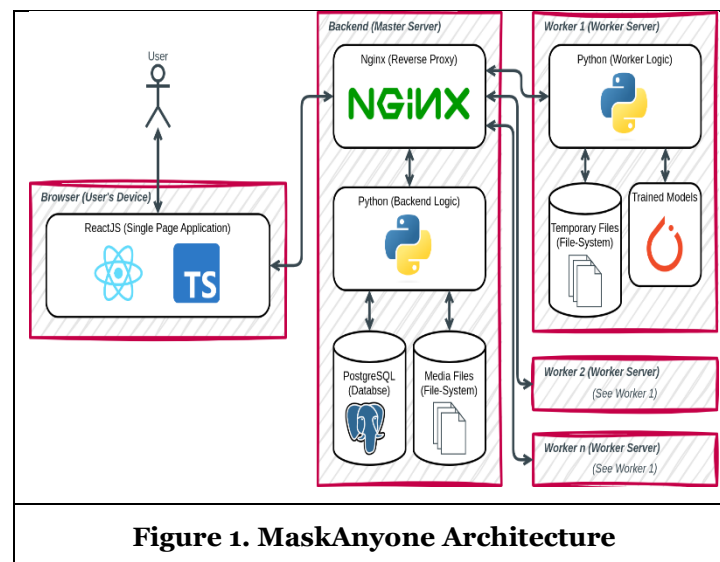


Figure 1. MaskAnyone Architecture

In response to the requirements for server environment support from our survey and workshops with data stewards and early career researchers, scalability, and extensibility (R5, R7, R8), we designed MaskAnyone as a web-based application accessible via standard web browsers. The architecture employs a Manager--Worker pattern to distribute tasks efficiently. The backend serves as the central hub for user interactions and job management, while workers handle the computational heavy-lifting of video masking. This setup allows for scalable and extensible operations without limiting the application's ability to run locally. Figure 1 illustrates this high-level architecture. The front-end uses React and TypeScript to create a single-page

application (SPA) that interacts with the backend via an HTTP API. The backend is accessible through a Nginx reverse proxy. %and is built with Python for consistency with the worker component. Data persistence is managed through PostgreSQL RDBMS, although media files are stored directly within the file system. Worker processes, also implemented in Python, register with the backend and receive masking jobs via an HTTP API. Docker and docker-compose are used for orchestration, making the application easy to set up with minimal prerequisites.

Multimodal Masking Process

The core functionality of MaskAnyone lies in its masking process, as depicted in Figure 2 below. The user initiates this by uploading a target video and selecting it within the interface, prompting the Masking Dialog (Figure -masking-presets). Users are offered three routes: select a predefined preset, choose a custom preset, or manually configure the masking options. If a preset is selected, users can either commence masking directly, satisfying the usability requirement (R2), or refine the preset via the Masking Configuration Dialog.

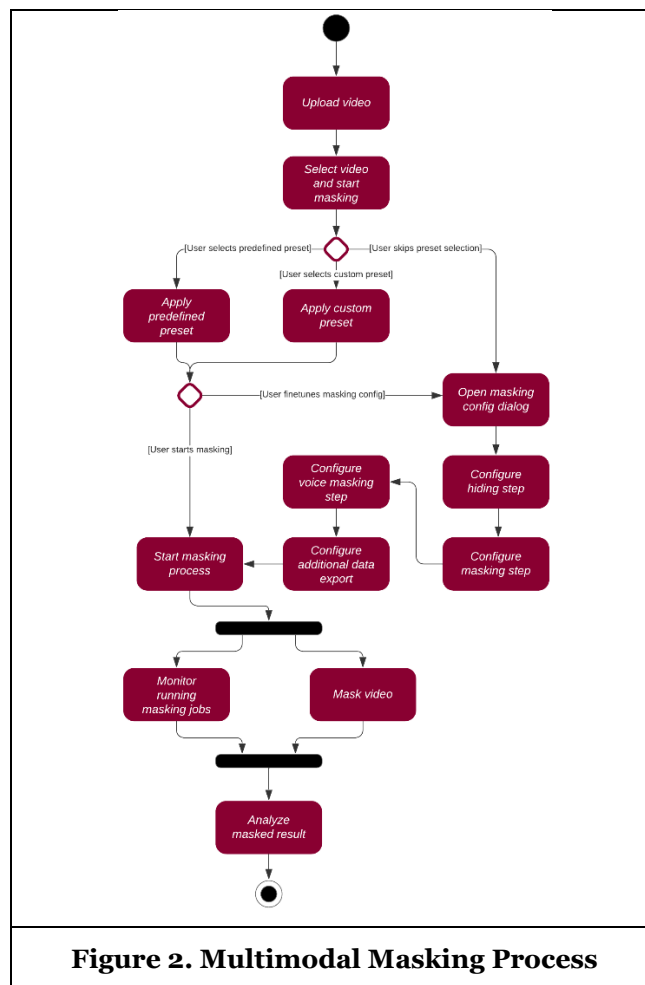


Figure 2. Multimodal Masking Process

The configuration process comprises four steps, offering users granular control over the masking. The first step provides options to control the hiding parameters for detected people and, optionally, the background. The second step involves specifying additional masking techniques to preserve important visual information. In the third step, decisions on audio masking options can be made, such as keeping, removing, or voice-converting the original audio. Finally, the fourth step allows for exporting additional data, such as kinematic information, either for more advanced analytics or external processing, in combination with audio-visual annotation tools (behavioral science) and NLP pipelines specific to the words in such masked

videos. This workflow is designed to meet our articulated usability and flexibility requirements, offering a streamlined yet comprehensive masking solution.

User Interaction, Design Principles.




The user interaction follows two design science principles: extensibility and scalability. The Manager--Worker pattern boosts scalability and extensibility. Users generate masking jobs in the backend; different job types can be added without altering the backend code. A new worker for the specific job type needs to be defined, adhering to established communication protocols. Conflicting dependencies for specific models are managed by creating unique Docker containers for each worker. The front-end integration of new masking methods is simplified using JSON Schema, allowing automated UI generation and job configuration validation. For scaling, MaskAnyone can add more worker processes and split large video files into multiple jobs. The primary bottleneck is the backend, which handles all user and worker interactions. While the backend can handle dozens of concurrent users and workers on robust hardware, further scaling may require partitioning it into multiple services. Significant extensions, like new categories of masking methods, may necessitate codebase modifications.

Audio-Visual Data Masking Strategies


As outlined in our GitHub repository, audio-visual masking in MaskAnyone involves two main aspects: hiding and masking. Hiding focuses on de-identification, while masking aims to preserve as much information in the data as possible. The toolkit offers a range of predefined options (i.e., strategies) for both categories.

Video Masking Strategy: Hiding (S1)

MaskAnyone employs two primary methods for detecting people in videos: YOLOv8 and MediaPipe. YOLO offers multiple pre-trained models with varying levels of accuracy and computational demand, including a specialized face-detection model. MediaPipe provides pose detection and an internal person mask. Users can adjust parameters like confidence thresholds for both methods to fine-tune utility and privacy tradeoffs¹. After successful person detection, MaskAnyone offers several hiding techniques, visualized in Table 2 below:

Strategy	Description	Effectiveness	Illustrative Results
Blackout	Completely blacks out the detected person, removing most identifiable information if the mask is accurate.	High, provided the mask is accurately applied.	
Blur	Applies Gaussian blurring at a user-controllable blur intensity level.	Moderate, may be limited depending on blur intensity.	
Contours	Uses Canny edge detection to preserve essential contours while eliminating finer details. User-controlled detail.	Moderate, balances privacy with utility.	





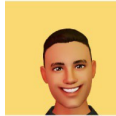
¹ For UI screenshots and videos of the different masking techniques, see the .mp4 files at <https://github.com/MaskAnyone/MaskAnyone>

Video Inpainting	Attempts to remove the person entirely, using surrounding data to fill in the background.	High, avoids residual artifacts but is computationally intensive.	
Table 2. MaskAnyone Strategy Elements (Hiding Approach)			

Note that video inpainting methods are resource-intensive and may require powerful hardware for timely execution. The integrated STTN model is limited to fixed resolutions, but alternative models supporting different resolutions are available.

Video Masking Strategy: Masking (S2)

MaskAnyone provides various masking options to preserve key information while maintaining privacy. These are summarized in Table 3 below and include:

Strategy	Description	Effectiveness	Sample Illustration
Skeleton	Captures body pose using 33 key points, rendering a skeleton overlay to preserve motion context.	High for motion context preservation.	
Face-Mesh	Extracts and renders 478 facial landmarks, preserving facial features and expressions when faces are adequately sized in the frame.	High for facial detail preservation.	
Holistic	Combines pose, face, and hand landmark detection in a heuristic-based implementation.	Comprehensive for full-body interaction contexts.	A standalone program yet to be integrated into MaskAnyone.
Face Swap	Swaps the subject's face with a selected target face, maintaining expressions while hiding identity.	High for identity concealment while preserving expressions.	
Rendered Avatar	Utilizes Blender and the BlendARMocap plugin to transform motion coordinates into a 3D avatar.	High for anonymity, maintains motion fidelity.	
Blendshapes Facial Avatar	Uses MediaPipe Face-Mesh to apply blendshapes on a 3D face, effectively preserving facial features.	High for facial expressions and details.	
Table 3. MaskAnyone Strategy Elements (Masking Approach)			

Voice Masking (S3)

Voice masking is a crucial extension of MaskAnyone's video masking capabilities. We offer three primary approaches to audio handling, as enumerated in Table 4 below:

Strategy	Description	Effectiveness	Sample Demonstration
Preserve	Retains the original audio and is suitable only for video masking.	Limited privacy, maintains original audio.	N/A
Remove	Eliminates all audio for maximum privacy.	High for privacy, low for utility.	N/A
Switch	Transforms the original voice to a target voice.	Balances privacy with utility and engagement.	N/A

Table 4. MaskAnyone Strategy Elements (Audio)

These features were implemented using MoviePy for voice extraction, RVC for voice conversion, and FFmpeg for audio-video merging. Various pre-trained target voices are available, and the system is extensible. We are also exploring other voice masking techniques, such as those from the Voice Privacy Challenge. End-users are advised to manually evaluate the privacy level/requirements before sharing.

Preliminary Evaluation - Towards a Robust Evaluation Framework

Evaluating video masking tools is crucial for determining the ethical benchmarks reached by behavioral science research. However, masking aims to balance privacy preservation and utility retention. Therefore, robust masking must be evaluated on both accounts, and we propose an automatic evaluation methodology. Automated evaluations offer the advantages of standardization and scalability, which is particularly important for defending against computational re-identification methods. We also consider human evaluations an important complementary form of assessment that caters to behavioral researchers' specific needs and accounts for the subjective nature of privacy and utility. Our techniques may still be deemed effective in scenarios where the trade-off between privacy and utility is acceptable to human evaluators, even if not to machines.

Automatic Video Evaluation

Utility: Our evaluation strategy is divided into privacy preservation and utility retention. For utility, we focus on a specific use-case: emotion classification. We apply our face-swapping technique on selected videos with precise headshots and then use an existing emotion classification model to categorize emotions in these videos. We introduce an agreement score that measures the concordance between the emotional states classified in the original and anonymized videos. This score serves as a proxy for utility retention, with a high agreement indicating practical preservation of emotional cues. We sourced five images and three videos for this preliminary evaluation, with subjects explicitly facing the camera. Face-swapping was also applied to generate 15 test samples. The insights from our survey and user studies underscore the importance of the usability of the toolkit. Moreover, the issues raised by the participants relate to the utility of the data after it has been masked, depending on the research context and strategies selected to preserve utility. At the same time, we learned that fully anonymizing videos is not 100% guaranteed; however, the tradeoff can be balanced sensibly as more researchers can reproduce existing research when they can access the data. This suggests avenues for further refinement of our proposed methodology.

Privacy: We categorize our techniques into hiding and masking methods for privacy preservation. Hiding techniques primarily focus on object detection models, such as YOLOv8, to identify silhouettes, bounding boxes, and facial features. Their efficacy is usually assessed using the mean average precision (MAP) metric. In our case, the face-specific model achieved a MAP of 38%, and person detection models yielded MAP

scores ranging from 37.3% to 53.9%, depending on the version used. Masking techniques, on the other hand, are evaluated through re-identification tasks. Due to the complexity of these tasks, a specialized dataset and model are often required. We fine-tuned a Vision Transformer and triplet-loss-based Re-ID model using a subset of the Celeb-A dataset. Preliminary results indicate a low precision score of 0.0017%, suggesting practical privacy preservation. However, these results also emphasize the need for further research to validate whether the low precision score indicates strong privacy or a model limitation.

Automatic Audio Evaluation

We have not performed a voice masking evaluation, but we outline a planned methodology for future research focusing on utility and privacy. Utility assessment in voice masking can be complex, as it varies by use case. However, general metrics like Word-Error Rate (WER) and Pitch Correlation could be considered. WER can be computed using an automatic speech recognition system (ASR) like ESPnet or OpenAI's Whisper. Pitch Correlation, significant for fields like psycholinguistics, can be measured using the Pearson Correlation Coefficient between the input and output signals. The Librispeech dataset, widely used in ASR evaluations, can be harnessed for this purpose. For privacy, we propose a re-identification or automatic speaker verification (ASV) approach for privacy evaluation, using Equal Error Rate (EER) as the metric. EER effectively gauges the ability to re-identify individuals in audio data and is implemented in popular toolkits such as Kaldi. While the Librispeech dataset could potentially be used for privacy evaluation, the suitability of other datasets should also be explored.

Human Evaluation

Human evaluation is pivotal in privacy and utility preservation and is the most challenging part of such design science projects. We conducted a user study involving Ph.D. students ($n=15$). The students were divided into two groups to serve as both control and treatment groups. Initially, each group was shown unmasked celebrity videos to establish a recognition baseline. They were then shown masked videos and asked to identify the individuals. One unmasked video was included to assess the impact of initial biasing. The results are also visualized here on the GitHub page². Notably, face-swapping techniques significantly reduced re-identification scores. For instance, none could correctly identify Jackie Chan, possibly due to the choice of the replacement face or cultural factors. Priming also had a notable impact, as demonstrated by the difference in identification rates for the same video of Steph Curry when presented as either masked or unmasked. These findings underscore the necessity for further research to validate and improve upon our masking techniques. We also conducted two more user feedback studies with researchers ($n=16$) at another European university. We can report that data stewards and early career researchers registered interest in subsequently participating in a more robust evaluation of the toolkit. We have also taken some of their feedback to streamline the masking strategies and added a login/database management option to manage access rights between researchers in the same faculty/lab to masked data.

Discussion

Maintaining privacy and utility in video masking is challenging and context-dependent. Different masking and hiding techniques may excel in privacy but limit utility in specific research settings. For this reason, we have developed MaskAnyone, a platform that combines masking and hiding techniques to tailor to the specific needs of researchers and other users. This masking tool also invites a discussion and a need to validate what combination of features constitutes a certain level of privacy and utility. In this paper, we have also proposed novel methods for evaluating those dimensions. We hope to contribute to such evidence-based measures and become better situated within ethical regulatory frameworks like GDPR. A vital issue is that it is unclear what counts as a proper reduction of identifiable information or anonymization. Human evaluation experiments might be insightful in this regard, but they come with limitations, too. Using celebrities to assess identifiability leverages background knowledge but likely inflates re-identification scores relative to real-world use cases. This is because researchers who share data often do not know the persons who are part of the original data recorded; the risk of identification is low. Legislation in the European Union is also evolving on this matter, such that although identification might be possible in

² <https://github.com/MaskAnyone/MaskAnyone>

principle if, in practice, data users are not likely to ascertain someone's identity, then a reduced risk of privacy violation is implied (EUR-Lex 2023). Further discussion points for our proof-of-concept evaluation methods include that providing a larger pool of identification choices may further reduce the likelihood of correct identification, as evidenced by our last test video, where no options were given. Important considerations include the cost and the typically small sample sizes of the resulting Human evaluations. Finally, specific to our experiment, our small sample size limits the generalizability of our findings.

Future opportunities and ethical challenges

However, when considering these preliminaries, human evaluations and automated techniques of the sort we have proposed may be crucial to ensure a multi-faceted assessment of masking techniques. MaskAnyone offers a range of options to researchers, allowing customization based on specific needs for privacy and utility. However, several untapped scenarios offer avenues for future development. The toolkit can expand on voice customization features for multiple individuals in a single video and allow for masking self-defined objects or persons. Moreover, since the toolkit can effectively anonymize video call recordings offline (see, e.g., link), it could also be implemented as part of the recording process. Such an online audio-visual masking would ensure that identifiable data is never stored outside the masking process). This potential expansion of the utility of MaskAnyone could be significant in CCTV contexts where a lot of audio-visual data is collected that is not necessarily for purposes of identification and widespread masking while preserving other information (number of people, types of activities, group forming) could change surveillance practices for the better. These uncovered use cases and possible extensions indicate room for the tool's evolution to meet diverse privacy and utility needs. And the best part? Since MaskAnyone is open source, it invites you, the audience, to shape its potential further. This first version of MaskAnyone has several limitations. Factors such as lighting, camera angles, and the presence of multiple subjects can influence anonymization accuracy and sometimes lead to identity leaks. Additionally, fine details like micro-expressions may also be compromised. Another aspect to consider is the toolkit's resource intensiveness, which could limit its accessibility. High computational demands and a 50GB storage requirement could pose barriers, although a lightweight installation is available to mitigate these issues. Scalability could be an issue, especially when handling large datasets or multiple users, leading to performance bottlenecks. Ethical and legal considerations add further complexities. Potential misuse of technology and compliance with data protection laws are issues that we are looking to address in close engagement with data stewards and ethics personnel. Several technical enhancements are planned to refine masking techniques and explore more robust machine-learning solutions. Such enhancements will focus on making the toolkit more robust, interactive, and user-friendly, and we are exploring incorporating identity and access management tools like Keycloak. We have ongoing collaborations with researchers and data stewards so that the toolkit's development follows real-world needs and to explore how masking toolkits can become integrated with ethical data archiving and sharing procedures. We also plan to enhance the data utility aspect of the toolkit by incorporating modules for multimodal behavior analyses and classification (multimodal analytics), which may include gesture detection (Ripperda et al. 2020) and analyses (Dang and Buschek 2021; Zeng et al. 2023), and a summary of prosodic markers in the audio before voice masking (ref). One other line of application of our toolkit that is timely relates to analyzing deep-fake videos and AI-generated content, seeing the developments of SORA (Ho et al. 2022; openai 2024; Yan et al. 2021)

Performance and Scalability

We assessed the performance of the MaskAnyone toolkit across various local setup scenarios to establish a baseline understanding of its processing capabilities. Using MediaPipe's Pose Landmarker, we observed that processing time scales linearly with video length while memory usage remains constant, regardless of the video's duration. Additionally, with an increasing number of subjects in a video, the memory usage goes up and processing time increases proportionally, all other factors being equal. Resolution had a minor impact on memory consumption, though processing time increased slightly with higher resolutions, likely due to the additional overhead in decoding and scaling larger frames. These tests were conducted on an Intel i7 9750H laptop CPU, demonstrating that MaskAnyone can perform effectively even on slightly older hardware. On more modern CPUs, like the Ryzen 5 7600X, we observed a significant increase in processing speed, achieving over twice the frame rate. These findings underscore the scalability of MaskAnyone beyond

local setups (i.e servers on university research clusters), with further potential for optimization, especially when utilizing GPU acceleration. While the results from these tests are promising, we plan to conduct a more comprehensive benchmark study using our GPU-based server to further optimize and validate the toolkit's performance. This future work will allow us to explore more advanced masking techniques and ensure that MaskAnyone can handle a wide range of datasets and research needs effectively, even as AI technologies and processing demands continue to evolve. The MaskAnyone toolkit is designed to meet the diverse needs of researchers working with various audiovisual datasets. We offer two versions: A local version, suitable for small to medium datasets and single-user environments, and a scalable server version, designed for enterprise-level, multi-user research requiring robust processing capabilities. MaskAnyone has been benchmarked on a high-spec research laptop, processing 30 minutes of HD video in about 2 hours while never pushing the CPU load beyond 20%. However, its scalability is limited for larger datasets or multiple videos. In contrast, the scalable version, optimized for servers with GPUs, can process up to 10 hours of HD video within 24 hours, handling large datasets with minimal performance degradation. Future benchmarking across different server configurations will further optimize its performance for diverse research environments.

Economic Analysis

We recognize that the economic viability of MaskAnyone is crucial for its adoption across institutions with varying resources. Ongoing work as part of a partnership with a Dutch university will ensure to include a detailed analysis comparing the costs of different MaskAnyone versions. The local version, suited for individual researchers, requires minimal hardware, making it highly accessible. In contrast, the scalable server version, designed for large-scale, multi-user deployments, may require significant initial investment in advanced servers or cloud infrastructure. While the standard version's low setup costs enhance accessibility, the enterprise version's higher operational costs, including maintenance and potential cloud fees, are offset by its ability to streamline data processing and reduce manual anonymization efforts. We anticipate that the long-term benefits—such as enhanced privacy, regulatory compliance, and improved data sharing—will outweigh the initial investment, particularly for larger institutions. Future studies will continue to refine this analysis to ensure MaskAnyone remains a cost-effective solution across diverse research environments.

Regulatory Compliance and Ethical Considerations

The MaskAnyone toolkit has been developed with a strong emphasis on regulatory compliance and ethical considerations, particularly in the context of research involving personal data. It aligns with key regulations like the General Data Protection Regulation (GDPR) in the European Union by ensuring data minimization, supporting the "right to erasure," and facilitating data portability while maintaining transparency. Additionally, the toolkit is designed to uphold ethical research practices, especially in studies involving vulnerable populations, by integrating privacy-preserving techniques that adhere to guidelines set by institutional review boards (IRBs) and ethics committees.

Planned Study on Masking Effectiveness/Advanced Detection Strategies

We plan to evaluate the MaskAnyone toolkit's effectiveness against advanced image detection technologies by processing a curated selection of videos and images using face-swapping and blurring techniques. These masked outputs will be analyzed with the Google Cloud Vision API to assess the extent of de-identification, using confidence scores as the primary metric. We expect face-swapping to significantly reduce detection scores, though some structural similarities may still be detectable. Gaussian blurring should further obscure features, particularly facial landmarks. The refinement of MaskAnyone, includes the development of advanced masking techniques and multi-stage processes. Additionally, principles from deception detection, such as pattern recognition, multi-layered defence, and adversarial testing, will be explored to enhance the toolkit's robustness and adaptability in real-time (Twyman et al. 2014)

Conclusion

Audio-visual data is crucial in behavioral sciences and linguistics, offering valuable insights into human behavior and multimodal language use. However, this data also presents ethical and privacy challenges,

particularly under frameworks like GDPR. Design science tools like MaskAnyone can help integrate ethical procedures into audiovisual data archiving, reducing the need to store fully identifiable data and enabling safer data sharing. Masking toolkits can also be used in academic settings to address privacy concerns during conferences and research dissemination. By standardizing masking strategies, tools like MaskAnyone can support compliance with privacy regulations, foster responsible data sharing, and enhance research integrity. Our toolkit addresses privacy needs while mitigating potential misuse, such as the creation of deepfakes, and aims to align with ethical guidelines in socio-behavioral research, ultimately improving the management and sharing of sensitive audiovisual data.

Acknowledgements

We would like to express our sincere gratitude to the editors and reviewers for their insightful feedback and guidance throughout the revision process. We thank the participants of the summer school who contributed their time and perspectives, helping to refine the MaskAnyone toolkit during its early stages.

References

- Abay, N. C., Zhou, Y., Kantarcioglu, M., Thuraisingham, B., and Sweeney, L. 2019. “Privacy Preserving Synthetic Data Release Using Deep Learning,” in *Machine Learning and Knowledge Discovery in Databases, Lecture Notes in Computer Science*, M. Berlingerio, F. Bonchi, T. Gärtner, N. Hurley, and G. Ifrim (eds.), Cham: Springer International Publishing, pp. 510–526. (https://doi.org/10.1007/978-3-030-10925-7_31).
- Abney, D. H., Dale, R., Louwerse, M. M., and Kello, C. T. 2018. “The Bursts and Lulls of Multimodal Interaction: Temporal Distributions of Behavior Reveal Differences Between Verbal and Non-Verbal Communication,” *Cognitive Science* (42:4), Wiley-Blackwell Publishing, pp. 1297–1316. (<https://doi.org/10.1111/cogs.12612>).
- Balci, K. 2005. “Xface: Open Source Toolkit for Creating 3D Faces of an Embodied Conversational Agent,” *Lecture Notes in Computer Science* (3638), Springer Verlag, pp. 263–266. (https://doi.org/10.1007/11536482_25/COVER).
- Bazarevsky, V., Grishchenko, I., Raveendran, K., Zhu, T., Zhang, F., and Grundmann, M. 2020. BlazePose: On-Device Real-Time Body Pose Tracking. (<https://arxiv.org/abs/2006.10204v1>).
- Benson, C. H., Friz, A., Mullen, S., Block, L., and Gilmore-Bykovskiy, A. 2020. “Ethical and Methodological Considerations for Evaluating Participant Views on Alzheimer’s and Dementia Research,” *Journal of Empirical Research on Human Research Ethics*. (<https://doi.org/10.1177/1556264620974898>).
- Bishop, L. 2009. “Ethical Sharing and Reuse of Qualitative Data,” *Australian Journal of Social Issues* (44:3), pp. 255–272. (<https://doi.org/10.1002/j.1839-4655.2009.tb00145.x>).
- Bradski, G. 2000. “The OpenCV Library,” *Dr. Dobb’s Journal of Software Tools*.
- Cgtinker. 2021. Cgtinker/BlendArMocap: Realtime Motion Tracking in Blender Using Mediapipe and Rigify. (<https://github.com/cgtinker/BlendArMocap>).
- Cienki, A. 2016. “Cognitive Linguistics, Gesture Studies, and Multimodal Communication,” *Cognitive Linguistics* (27:4), Walter de Gruyter GmbH, pp. 603–618. (<https://doi.org/10.1515/cog-2016-0063>).
- Dang, H., and Buschek, D. 2021. “GestureMap: Supporting Visual Analytics and Quantitative Analysis of Motion Elicitation Data by Learning 2D Embeddings,” in *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, , May 6, pp. 1–12. (<https://doi.org/10.1145/3411764.3445765>).
- D’Errico, F., Poggi, I., Vinciarelli, A., and Vincze, L. 2015. “Conflict and Multimodal Communication: Social Research and Machine Intelligence,” *Confl. and Multimodal Communication: Soc. Research and Machine Intelligence*, Springer International Publishing, p. 479. (<https://doi.org/10.1007/978-3-319-14081-0>).
- EUR-Lex. 2023. “EUR-Lex - 62020TJ0557 - EN - EUR-Lex,” Law. (<https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A62020TJ0557>).
- Fang, F., Wang, X., Yamagishi, J., Echizen, I., Todisco, M., Evans, N., and Bonastre, J.-F. 2019. Speaker Anonymization Using X-Vector and Neural Waveform Models, *International Speech Communication Association*, pp. 155–160. (<https://doi.org/10.21437/SSW.2019-28>).
- Gafni, O., Wolf, L., and Taigman, Y. 2019. “Live Face De-Identification in Video,” *Proceedings of the IEEE International Conference on Computer Vision (2019-October)*, Institute of Electrical and Electronics Engineers Inc., pp. 9377–9386. (<https://doi.org/10.1109/ICCV.2019.00947>).

- Gregori, A., Amici, F., Brilmayer, I., Ćwiek, A., Fritzsche, L., Fuchs, S., Henlein, A., Herbolt, O., Kügler, F., Lemanski, J., Liebal, K., Lücking, A., Mehler, A., Nguyen, K. T., Pouw, W., Prieto, P., Rohrer, P. L., Sánchez-Ramón, P. G., Schulte-Rüther, M., Schumacher, P. B., Schweinberger, S. R., Struckmeier, V., Trettenbrein, P. C., and von Eiff, C. I. 2023. “A Roadmap for Technological Innovation in Multimodal Communication Research,” in *Lect. Notes Comput. Sci.* (Vol. 14029 LNCS), Duffy V.G. (ed.), Springer Science and Business Media Deutschland GmbH, pp. 402–438. (https://doi.org/10.1007/978-3-031-35748-0_30).
- Grishchenko, I., and Bazarevsky, V. 2020. MediaPipe Holistic – Simultaneous Face, Hand and Pose Prediction, on Device – Google Research Blog. (<https://blog.research.google/2020/12/mediapipe-holistic-simultaneous-face.html?m=1>).
- Hasan, R., Hassan, E., Li, Y., Caine, K., Crandall, D. J., Hoyle, R., and Kapadia, A. 2018. “Viewer Experience of Obscuring Scene Elements in Photos to Enhance Privacy,” in *Conference on Human Factors in Computing Systems - Proceedings* (Vol. 2018-April). (<https://doi.org/10.1145/3173574.3173621>).
- Hevner, A. R., March, S. T., Park, J., and Ram, S. 2004. “Design Science in Information Systems Research,” *MIS Quarterly* (28:1), Management Information Systems Research Center, University of Minnesota, pp. 75–105. (<https://doi.org/10.2307/25148625>).
- Ho, J., Chan, W., Saharia, C., Whang, J., Gao, R., Gritsenko, A., Kingma, D. P., Poole, B., Norouzi, M., Fleet, D. J., and Salimans, T. 2022. Imagen Video: High Definition Video Generation with Diffusion Models, arXiv. (<https://doi.org/10.48550/arXiv.2210.02303>).
- Hsu, W. N., Bolte, B., Tsai, Y. H. H., Lakhotia, K., Salakhutdinov, R., and Mohamed, A. 2021. “HuBERT: Self-Supervised Speech Representation Learning by Masked Prediction of Hidden Units,” *IEEE/ACM Transactions on Audio Speech and Language Processing* (29), Institute of Electrical and Electronics Engineers Inc., pp. 3451–3460. (<https://doi.org/10.1109/TASLP.2021.3122291>).
- InsightFace. 2023. InsightFace: An Open Source 2D&3D Deep Face Analysis Library. (<https://insightface.ai/>).
- Jarolimkova, A., and Drobikova, B. 2019. “Data Sharing in Social Sciences: Case Study on Charles University,” in *Information Literacy in Everyday Life, Communications in Computer and Information Science*, S. Kurbanoglu, S. Špiranec, Y. Ünal, J. Boustany, M. L. Huotari, E. Grassian, D. Mizrachi, and L. Roy (eds.), Cham: Springer International Publishing, pp. 556–565. (https://doi.org/10.1007/978-3-030-13472-3_52).
- Jeng, W., He, D., and Oh, J. S. 2016. “Toward a Conceptual Framework for Data Sharing Practices in Social Sciences: A Profile Approach,” *Proceedings of the Association for Information Science and Technology*. (<https://doi.org/10.1002/pra2.2016.14505301037>).
- Johannesson, P., and Perjons, E. 2021. “Ethics and Design Science,” in *An Introduction to Design Science*, P. Johannesson and E. Perjons (eds.), Cham: Springer International Publishing, pp. 185–197. (https://doi.org/10.1007/978-3-030-78132-3_13).
- Kasturi, R., and Ekambaram, R. 2014. “Person Reidentification and Recognition in Video,” *Lecture Notes in Computer Science* (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics) (8827), Springer Verlag, pp. 280–293. (https://doi.org/10.1007/978-3-319-12568-8_35/COVER).
- Khasbage, Y., Carrión, D. A., Hinnell, J., Robertson, F., Singla, K., Uhrig, P., and Turner, M. 2022. “The Red Hen Anonymizer and the Red Hen Protocol for De-Identifying Audiovisual Recordings,” *Linguistics Vanguard*, Walter de Gruyter GmbH. (<https://doi.org/10.1515/LINGVAN-2022-0017/MACHINEREADABLECITATION/RIS>).
- Kim, J., Kong, J., and Son, J. 2021. “Conditional Variational Autoencoder with Adversarial Learning for End-to-End Text-to-Speech,” *Proceedings of Machine Learning Research* (139), ML Research Press, pp. 5530–5540.
- Kim, J. W., Salamon, J., Li, P., and Bello, J. P. 2018. “Crepe: A Convolutional Representation for Pitch Estimation,” *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings* (2018-April), Institute of Electrical and Electronics Engineers Inc., pp. 161–165. (<https://doi.org/10.1109/ICASSP.2018.8461329>).
- Kong, J., Kim, J., and Bae, J. 2020. “HiFi-GAN: Generative Adversarial Networks for Efficient and High Fidelity Speech Synthesis,” *Advances in Neural Information Processing Systems* (2020-December), Neural information processing systems foundation. (<https://arxiv.org/abs/2010.05646v2>).
- Korshunov, P., and Marcel, S. 2018. DeepFakes: A New Threat to Face Recognition? Assessment and Detection. (<https://arxiv.org/abs/1812.08685v1>).

- Lander, K., Bruce, V., and Hill, H. 2001. "Evaluating the Effectiveness of Pixelation and Blurring on Masking the Identity of Familiar Faces," *Applied Cognitive Psychology* (15:1). ([https://doi.org/10.1002/1099-0720\(200101/02\)15:1<101::AID-ACP697>3.0.CO;2-7](https://doi.org/10.1002/1099-0720(200101/02)15:1<101::AID-ACP697>3.0.CO;2-7)).
- Liu, K., Perov, I., Gao, D., Chervoniy, N., Zhou, W., Zhang, W., Liu, K., Perov, I., Gao, D., Chervoniy, N., Zhou, W., and Zhang, W. 2023. "Deepfacelab: Integrated, Flexible and Extensible Face-Swapping Framework," *PatRe* (141), Elsevier Ltd, p. 109628. (<https://doi.org/10.1016/J.PATCOG.2023.109628>).
- Meshcapade. 2023. Meshcapade | The Digital Human Company. (<https://meshcapade.com/>).
- Mroz, S., Baddour, N., McGuirk, C., Juneau, P., Tu, A., Cheung, K., and Lemaire, E. 2021. "Comparing the Quality of Human Pose Estimation with BlazePose or OpenPose," *BioSMART 2021 - Proceedings: 4th International Conference on Bio-Engineering for Smart Technologies*, Institute of Electrical and Electronics Engineers Inc. (<https://doi.org/10.1109/BIOSMART54244.2021.9677850>).
- Nautsch, A., Jasserand, C., Kindt, E., Todisco, M., Trancoso, I., and Evans, N. 2019. "The GDPR & Speech Data: Reflections of Legal and Technology Communities, First Steps towards a Common Understanding," *Proceedings of the Annual Conference of the International Speech Communication Association*, INTERSPEECH (2019-September), International Speech Communication Association, pp. 3695–3699. (<https://doi.org/10.21437/INTERSPEECH.2019-2647>).
- Nirkin, Y., Keller, Y., and Hassner, T. 2019. "FSGAN: Subject Agnostic Face Swapping and Reenactment," *Proceedings of the IEEE International Conference on Computer Vision (2019-October)*, Institute of Electrical and Electronics Engineers Inc., pp. 7183–7192. (<https://doi.org/10.1109/ICCV.2019.00728>).
- Openai. 2024. "Video Generation Models as World Simulators," <https://Openai.Com/Research/Video-Generation-Models-as-World-Simulators>. (<https://openai.com/research/video-generation-models-as-world-simulators>, accessed April 27, 2024).
- OpenCV. 2023. Releases - OpenCV. (<https://opencv.org/releases/>).
- Owoyele, B., Trujillo, J., Melo, G. de, and Pouw, W. 2022. "Masked-Piper: Masking Personal Identities in Visual Recordings While Preserving Multimodal Information," *SoftwareX* (20), Elsevier B.V., p. 101236. (<https://doi.org/10.1016/j.softx.2022.101236>).
- Patino, J., Tomashenko, N., Todisco, M., Nautsch, A., and Evans, N. 2020. "Speaker Anonymisation Using the McAdams Coefficient," *Proceedings of the Annual Conference of the International Speech Communication Association*, INTERSPEECH (3), International Speech Communication Association, pp. 1958–1962. (<https://doi.org/10.21437/Interspeech.2021-1070>).
- Pavlakos, G., Choutas, V., Ghorbani, N., Bolkart, T., Osman, A. A., Tzionas, Di., and Black, M. J. 2019. "Expressive Body Capture: 3D Hands, Face, and Body from a Single Image," *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (2019-June)*, IEEE Computer Society, pp. 10967–10977. (<https://doi.org/10.1109/CVPR.2019.01123>).
- Redmon, J., Divvala, S., Girshick, R., and Farhadi, A. 2016. "You Only Look Once: Unified, Real-Time Object Detection," in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (Vol. 2016-December). (<https://doi.org/10.1109/CVPR.2016.91>).
- Resnik, D., Antes, A., and Mozersky, J. 2024. "Should Researchers Destroy Audio or Video Recordings?," *Ethics & Human Research* (46), pp. 30–35. (<https://doi.org/10.1002/eahr.500205>).
- Reynolds, D. 2015. "Gaussian Mixture Models," *Encyclopedia of Biometrics*, Springer, Boston, MA, pp. 827–832. (https://doi.org/10.1007/978-1-4899-7488-4_196).
- Ripperda, J., Drijvers, L., and Holler, J. 2020. "Speeding up the Detection of Non-Iconic and Iconic Gestures (SPUDNIG): A Toolkit for the Automatic Detection of Hand Movements and Gestures in Video Data," *Behavior Research Methods* (52:4), Springer, pp. 1783–1794. (<https://doi.org/10.3758/s13428-020-01350-2>).
- "Roadmaps from the Three Thematic DCCs – Digital Competence Centres | NWO." (n.d.). (<https://www.nwo.nl/en/researchprogrammes/implementation-plan-investments-digital-research-infrastructure/roadmaps-three>, accessed April 30, 2024).
- Rokoko. 2023. Intuitive and Affordable Motion Capture Tools for Character Animation. (<https://www.rokoko.com/>).
- Su, Z., Liu, W., Yu, Z., Hu, D., Liao, Q., Tian, Q., Pietikäinen, M., and Liu, L. 2021. "Pixel Difference Networks for Efficient Edge Detection," in *Proceedings of the IEEE International Conference on Computer Vision*. (<https://doi.org/10.1109/ICCV48922.2021.00507>).
- "Thematic Digital Competence Centers | NWO." 2024. , November 15. (<https://www.nwo.nl/en/calls/thematic-digital-competence-centers>, accessed April 30, 2024).

- Tomashenko, N., Srivastava, B. M. L., Wang, X., Vincent, E., Nautsch, A., Yamagishi, J., Evans, N., Patino, J., Bonastre, J. F., Noé, P. G., and Todisco, M. 2020. "Introducing the VoicePrivacy Initiative," *Proceedings of the Annual Conference of the International Speech Communication Association, INTERSPEECH (2020- October)*, International Speech Communication Association, pp. 1693–1697. (<https://doi.org/10.21437/INTERSPEECH.2020-1333>).
- Tomashenko, N., Wang, X., Vincent, E., Patino, J., Srivastava, B. M. L., Noé, P. G., Nautsch, A., Evans, N., Yamagishi, J., O'Brien, B., Chanclu, A., Bonastre, J. F., Todisco, M., and Maouche, M. 2022. "The VoicePrivacy 2020 Challenge: Results and Findings," *Computer Speech & Language* (74), Academic Press, p. 101362. (<https://doi.org/10.1016/J.CSL.2022.101362>).
- Twyman, N. W., Lowry, P. B., Burgoon, J. K., and Nunamaker, J. F. Jr. 2014. "Autonomous Scientifically Controlled Screening Systems for Detecting Information Purposely Concealed by Individuals," *Journal of Management Information Systems* (31:3), pp. 106-137. (<http://www.jstor.org/stable/43590295>).
- Wang, R., Ding, Y., Li, L., and Fan, C. 2020. "One-Shot Voice Conversion Using Star-Gan," *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings (2020-May)*, Institute of Electrical and Electronics Engineers Inc., pp. 7729–7733. (<https://doi.org/10.1109/ICASSP40776.2020.9053842>).
- Xu, Y., Zhang, J., Zhang, Q., and Tao, D. 2022. "ViTPose: Simple Vision Transformer Baselines for Human Pose Estimation," *Advances in Neural Information Processing Systems* (35), Neural information processing systems foundation. (<https://arxiv.org/abs/2204.12484v3>).
- Yan, W., Zhang, Y., Abbeel, P., and Srinivas, A. 2021. VideoGPT: Video Generation Using VQ-VAE and Transformers, arXiv. (<https://doi.org/10.48550/arXiv.2104.10157>).
- Yu, T., Feng, Runseng, Feng, Ruoyu, Liu, J., Jin, X., Zeng, W., and Chen, Z. 2023. Inpaint Anything: Segment Anything Meets Image Inpainting. (<https://arxiv.org/abs/2304.06790v1>).
- Zeng, H., Wang, X., Wang, Y., Wu, A., Pong, T. C., and Qu, H. 2023. "GestureLens: Visual Analysis of Gestures in Presentation Videos," *IEEE Transactions on Visualization and Computer Graphics* (29:8), pp. 3685–3697. (<https://doi.org/10.1109/TVCG.2022.3169175>).
- Zeng, Y., Fu, J., and Chao, H. 2020. "Learning Joint Spatial-Temporal Transformations for Video Inpainting," *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* (12361 LNCS), Springer Science and Business Media Deutschland GmbH, pp. 528–543. (https://doi.org/10.1007/978-3-030-58517-4_31).
- Zhao, J., and Zhang, H. 2022. "Thin-Plate Spline Motion Model for Image Animation," *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (2022-June)*, IEEE Computer Society, pp. 3647–3656. (<https://doi.org/10.1109/CVPR52688.2022.00364>).