

The AI4Media project: Use of Next-generation Artificial Intelligence Technologies for Media Sector Applications

Filareti Tsalakanidou¹[0000-0002-5310-8045], Symeon Papadopoulos¹[0000-0002-5441-7341], Vasilios Mezaris¹[0000-0002-0121-4364], Ioannis Kompatsiaris¹ [0000-0001-6447-9020], Birgit Gray², Danae Tsabouraki³, Maritini Kalogerini³, Fulvio Negro⁴[0000-0002-1889-5485], Maurizio Montagnuolo⁴[0000-0002-6535-0322], Jesse de Vos⁵, Philo van Kemenade⁵, Daniele Gravina⁶, Rémi Mignot⁷, Alexey Ozerov⁸, Francois Schnitzler⁸, Artur Garcia-Saez⁹, Georgios N. Yannakakis¹⁰, Antonios Liapis¹⁰, and Georgi Kostadinov¹¹ [0000-0002-3465-8797]

¹ Information Technologies Institute, Centre for Research and Technology Hellas, 6th km Charilaou-Thermi Road, 57001 Thermi, Thessaloniki, Greece

² Deutsche Welle, Kurt-Schumacher-Straße 3, 53113 Bonn, Germany

³ Athens Technology Center (ATC), Rizariou 10, 15233 Chalandri, Athens, Greece

⁴ RAI - Radiotelevisione Italiana, Via G.C. Cavalli 6, 10138 Turin, Italy

⁵ Netherlands Institute for Sound and Vision, Mediaparkboulevard 1, Hilversum, Netherlands

⁶ modl.ai, Nørrebrogade 184, 1, 2200 Copenhagen, Denmark

⁷ IRCAM, 1 pl. Stravinsky, 75004 Paris, France

⁸ InterDigital, 975 av. de Champs Blancs, 35510 Cesson-Sévigné, France

⁹ Barcelona Supercomputing Center (BSC), c/Jordi Girona 31, 08034 Barcelona, Spain

¹⁰ University of Malta, Msida MSD 2080, Malta

¹¹ Imagga Technologies Ltd, 47A Cherni Vrah blvd, 1407 Sofia, Bulgaria

ikom@iti.gr

Abstract. Artificial Intelligence brings exciting innovations in all aspects of life and creates new opportunities across industry sectors. At the same time, it raises significant questions in terms of trust, ethics, and accountability. This paper offers an introduction to the AI4Media project, which aims to build on recent advances of AI in order to offer innovative tools to the media sector. AI4Media unifies the fragmented landscape of media-related AI technologies by investigating new learning paradigms and distributed AI, exploring issues of AI explainability, robustness and privacy, examining AI techniques for content analysis, and exploiting AI to address major societal challenges. In this paper, we focus on our vision of how such AI technologies can reshape the media sector, by discussing seven industrial use cases that range from combating disinformation in social media and supporting journalists for news story creation, to high quality video production, game design, and artistic co-creation. For each of these use cases, we highlight the present challenges and needs, and explain how they can be efficiently addressed by using innovative AI-driven solutions.

Keywords: Artificial Intelligence, Media, Democracy, Society, Use Cases, Social Media, Content Automation, Vision, Journalism, Co-creation, Gaming.

1 Artificial Intelligence in the Service of Media, Society and Democracy: Current Challenges & Opportunities

Following a series of breakthroughs in the field of Artificial Intelligence (AI), new technologies are emerging which are ushering a wave of innovations in all aspects of business and society, across industry sectors such as transportation, health or finance, for global goals such as fighting climate change, also impacting the media industry, journalism and politics [1]. In all these facets of economic and social life, AI is disrupting existing practices and creates opportunities. The exploitation of these opportunities brings significant socio-economic changes that necessitate a focus on issues of trust, ethics and accountability, besides the pursuit of technological excellence and financial profit.

This human-centric and trustworthy ethical brand of AI is particularly relevant to the media sector. Digital media permeates most aspects of human and social activity and is intertwined with information exchange and knowledge transfer. Machine vision and visual content understanding were some of the first fields to exhibit significant breakthroughs in the evolution of AI, including advances in audio/music analysis and generation, text and language analysis, and modeling of social trends. The media market is already benefiting from AI-based support across the value chain: for media newsgathering, production, distribution, and delivery as well as audience analysis. This includes a range of tools and services for processes such as information analysis, content creation, media editing, content optimisation, audience preference analysis, and recommender systems [2].

Furthermore, it has become apparent that society and politics are increasingly affected by AI developments. There is strong concern that the combination of the power held by the major, globally operating social media platforms and the large-scale automation capabilities associated with these platforms offered by AI technology could prove detrimental to individuals, society and democracy [3]. For instance, AI-powered bot collectives were found to have played an important role in shaping the results of US Elections in 2016 and the Brexit referendum [4] while the recent advent of synthetic media technology is posing new risks to citizens' trust in online content.

AI technologies are expected to disrupt the media industry through advances in content synthesis, analysis, and distribution, but also by offering new deeper insights into the complex and rapidly evolving social processes that unfold online and offline by sensing citizen activities, interests and opinions [2]. AI technology could help shape the democratic role of the media by enabling new ways of being informed, of deliberation, political participation, and decision making. In addition, AI technologies could support the relationship between media providers and their audiences, helping to align with the needs of media users and citizens. The use of AI can also cut down operating costs and ultimately free up resources for more in-depth and quality journalism. In addition, the use of AI can also create opportunities for the better realization of public values, such as media diversity, freedom of expression, and inclusiveness.

Motivated by the challenges, risks and opportunities that a wider use of AI brings to the media sector and society, the EU-funded AI4Media project [5] aspires to i) deliver the next generation of core AI technologies, ii) ensure the embedding of ethi-

cal and trustworthy AI into future AI deployments, and iii) reimagine AI as a human-centered, trusted and beneficial enabling technology for media and society.

AI4Media aims to advance the AI state-of-the-art in four key technology areas, with a view to unify a currently fragmented landscape of media related AI and offer solutions for the ever-growing needs and challenges of media organizations:

- *New learning paradigms and distributed AI*: AI models that adapt and learn on-the-fly, models that teach each other, algorithms that learn from limited data; Distributed AI and AI systems operating on the Edge.
- *Explainability, robustness and privacy in AI*: Federated learning and privacy-enhancing technologies ensuring data privacy; Technologies that protect information systems from malicious attacks and increase robustness; Approaches for evaluation, systematization, and explainability of AI.
- *Content-centered AI*: Multimedia metadata extraction, summarization, and clustering; Automatic audiovisual content generation and enhancement; Linguistic analysis; Learning with sparse data and transfer learning.
- *Human- and Society-centered AI*: Technologies for multimedia manipulation detection; Political opinion mining, disinformation detection, and local news detection; User perception measurement algorithms for personalization and bias detection; Privacy-preserving content recommendation.

The developed technologies will be tested in the context of seven industrial use cases that focus on the media sector, with relevance for both society and the economy. They cover different forms of media, such as digital content, broadcast video, film, audio, music, games, and virtual environments. Content areas include news stories, social media, media archives, and user-generated content. The use cases also address a variety of societal and democracy related topics such as disinformation, public discourse, media moderation, consumer privacy, creativity, and human-machine interaction. The seven use cases will be realized through close collaboration between AI researchers and European media organizations or content related companies.

The rest of the paper is organized as follows: Section 2 discusses the AI technologies explored in AI4Media that focus on the media sector. Section 3 presents the seven AI4Media use cases. Finally, section 4 draws the conclusions.

2 AI technologies for the Media Sector

AI4Media will conduct research in cutting-edge AI areas, focusing on tools and applications for the Media Industry. In the following, we briefly discuss the spectrum of AI technologies that will be examined in the context of the project and the use cases.

New learning paradigms. The AI explosion was initially spearheaded by Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs) applied to well-known classification tasks. Building on this, we focus on developing faster and more accurate learning models by examining: techniques for lifelong learning to allow updating system capabilities without forgetting past knowledge [6]; new spaces and operators for unsupervised manifold learning to exploit non-annotated data; trans-

fer learning to allow reusing models trained for different tasks [7]; transfer neural architecture search to find suitable DNN architectures fast by leveraging past experience in different datasets [8]; and deep Quality Diversity methods for DNNs to improve quality and diversity of generated outcomes in gaming [9]. We leverage Quantum architectures and test the efficiency of Quantum-inspired ML algorithms dealing with large datasets. We also move AI from centralized cloud architectures to networks of heterogeneous devices, to allow lightweight devices sitting on the Cloud Edge to run their own DNN models [10]. In this direction, we explore distributed learning and decentralized graph mining.

Trustworthy AI. To ensure trustworthy AI for media applications, AI robustness, explainability, privacy and fairness are explored. We focus on enhancing system resilience towards adversarial attacks like poisoning or evasion [11]; improving AI interpretability [12] with regard to the use of DNNs for tasks like deepfake detection; improving the privacy aspect of recommendation systems by combining de-centralized methods for federated learning [13], differential privacy and encryption; and improving fairness through detection and mitigation of bias in recommendation systems [14]. These contributions will be supplemented by a new benchmarking framework aiming to evaluate effectiveness and reproducibility of AI systems.

Multimedia content analysis and creation. The aim is to provide improvements on the state of the art, such as development of better automatic video summarization algorithms for news agencies based on Generative Adversarial Networks (GANs) [15] and on combination of knowledge representation and deep learning; and to design better audio generation algorithms for games, exploiting DNNs and domain-specific knowledge for music annotation and audio analysis [16]. In addition, we explore vision-based DNNs for automatic online cinematography based on UAV footage [17], while also developing image enhancement techniques for 360 videos, thermal imagery, depth, etc. We also improve visual detection tasks by using GANs and 3D virtual world simulators to tackle data scarcity [18]. Finally, we improve Word Embeddings for NLP algorithms [19] to provide efficient multi-language models.

Human-centered AI. Online platforms have become a central source of information, shaping the democratic debate and society [3]. To support journalists in providing trustworthy information and citizens in being adequately informed outside of filter bubbles, AI4Media: tackles multimedia content manipulation (deepfakes) by combining text verification based on graph mining [20], audio analysis, and video analysis based on extended GANs [21] and novel video pre-processing approaches [22]; promotes healthy political debate by including diversity and novelty in recommender systems to avoid over-personalisation [23], and by modeling disinformation and political polarization in national and local scale; and finally, analyzes user perception of social media through a combination of user studies and content analysis [24].

3 The AI4Media Use Cases

AI4Media has designed seven use cases that demonstrate the use of AI in the media industry, including aspects of human-centric, ethical and trustworthy AI (see Fig. 1).

The AI4Media use cases (UC) are inspired by market needs, emerging opportunities, and a range of industry challenges, and they highlight how AI applies throughout the media and content value chain and how different types of media players aim to address user and business needs with novel AI solutions.

The use cases deal with a broad range of media processes that will be optimized and enhanced by the application of state-of-the-art AI technologies. The industry partners gather a set of user requirements for new AI functionalities with a view to upgrade available tools already used by the media sector with AI features developed by the project (see section 2). In the following, we present the seven use cases.




Fig. 1. The AI4Media use cases.

3.1 UC1: AI for Social Media and Against Disinformation

While disinformation activities and the volume of related verification work increase, the workflow for verifying digital content items and detecting disinformation remains complex, time-consuming and often requires specialists. The current generation of support tools lack functions for identifying the most recent forms of manipulated content and easily understanding disinformation patterns in social media. In addition, technologies used for content manipulation and disinformation have become more advanced, also involving synthetic media. There is a need for upgraded, easy-to-use tools with trusted AI functions, as more journalists and other researchers will be involved in content verification and disinformation detection tasks. In order to ensure acceptance of AI-based tools and the feasibility of their implementation, users need to judge aspects of trustworthiness with regard to machine-generated results and predictions, such as *Explainability*, *Bias Mitigation*, *Robustness* or *Ethical Compliance*.

This use case from Deutsche Welle and ATC leverages AI technologies to improve support tools used by journalists and fact-checking experts for digital content verification and disinformation detection. New AI-based features will be made available within two existing journalism tools: *Truly Media* [25] (web-based platform for collaborative verification) and *TruthNest* [26] (Twitter analytics and bot detection tool).

Two topics are covered: i) verification of content from social media with a focus on *synthetic media detection*, and ii) detection of communication narratives and patterns that are related to *disinformation*. Another aspect is the exploration of *Trustworthy AI* in relation to these topics and the specific needs of media organisations.

The detection and verification of synthetic media refers to both synthetically generated media items and synthetic elements within media items. This provides contextual information for the verification process, helping journalists and researchers to understand specific manipulation activities. The more advanced synthetic media be-

comes, the more difficult and time-consuming it will be to detect it. While there are useful and entertaining applications for synthetic media, it is also used for targeted, malicious disinformation. Synthetic media items misused for disinformation purposes are colloquially known as *Deepfakes* [27]. All types of content can be synthetically generated (or manipulated): text, images, audio, video and elements within AR/VR experiences. *Synthetic Characters* can be either fictional or with high levels of likeness to a person, its voice and personality.

Another goal is to provide AI-based tools for better understanding communication patterns in social media that are related to disinformation. This includes the detection of topical disinformation campaigns or stories, which may entail false or distorted information. The latter can emerge dynamically, or be steered by certain social media actors and networked communities in order to achieve a specific objective.

We address these challenges by collecting specific user needs in the area of content verification and disinformation detection that might be solved by advanced AI functionalities, including needs for associated Trustworthy AI features. Once respective AI tools have been developed in AI4Media, they will be tested against end user and business needs within Truly Media and TruthNest (see Fig. 2).




Fig. 2. AI technologies for content verification and disinformation detection developed in AI4Media will be integrated in existing journalism tools like Truly Media.

3.2 UC2: AI for News - The Smart News Assistant

Journalists face a challenging environment where the amount of incoming content is ever increasing while the need to publish news as fast as possible is extremely pressing. On top of that, journalists need to ensure that published content is trustworthy and relevant for its audience. The use case focuses on the concept of a Smart News Assistant, i.e. a tool that will support journalists in the creation of news stories by providing a variety of AI-enabled functionalities for story production and development, story curation and publication, and audience engagement.

Story development will be facilitated by providing automatic story suggestions based on analytics of past audience engagement and trend detection in social media, as well as by providing real-time suggestions during article preparation, e.g. suggest-

ing relevant articles or video content from news archives and warning about fake news. Functionalities for automated news articles summaries, automated article indexing and annotation for easier search, automated content curation and adaptation to different audience needs, and personalised news experiences based on past behaviour of news consumers will be supported to enhance story curation. Finally, semi-automated moderation of online discussions about news articles, including polarisation detection and classification of disagreements, and automatic management of audience feedback will be supported to facilitate audience interaction and engagement.

3.3 UC3: AI for High Quality Video Production & Content Automation

This use case aims at supporting broadcasters' newsrooms (and in general information and entertainment production) in reporting unexpected events like natural disasters (e.g., floods, earthquakes). Generally, whenever similar events occur, newsrooms need to readily publish and update fresh news in a very short timeframe. It is evident that the "time" variable plays a crucial role: providing and getting news on time can make all the difference and the entire production process is stressed. Together, content quality and reliability must always be properly and readily verified before flowing into the production chain. AI4Media aims to broaden the horizons along three perspectives involved in such just-in-time, high-quality production process (Fig. 3).

First, *timely coverage* of unexpected events must be ensured. Several tools for professional news search and filtering are available, but they lack in providing new and engaging information as soon as an event occurs. However, auxiliary sources (e.g., social media and video sharing platforms) provide a plethora of information such as the nature of the disaster, affected people's emotions and relief efforts. Unfortunately, that content cannot be used indiscriminately due to e.g., quality, privacy and reliability issues. AI tools for content quality enhancement (e.g., denoising, dehazing, inpainting), privacy protection (e.g., identity obfuscation) and manipulation detection are the key to overcome such obstacles. Furthermore, AI computer vision technologies can speed up the creation of metadata, e.g., descriptive information about actions, places and things identifiable within the content, making it easier to be searched.

Second, *broadcaster's archives* are a valuable source of information about how the landscape was in the past. However, archive material might be impaired and inadequately documented. Thus, appropriate AI-based technologies are needed for video restoration (e.g., artifacts reduction, colourisation) and metadata extraction (e.g., image and video captioning). Furthermore, drone technology can provide unprecedented viewing experiences, e.g., by gathering data from inaccessible or arduous areas. Archives and drones might work in synergy and even amplify each other's results. Archive material might help in pre-setting operational drone's parameters (e.g., flight paths, altitude, speed) while drone camera footage might be used to assist the video restoration process, e.g., by providing chroma information for video colourisation.

At last, *open data, professional newsgathering services and broadcast archives* are the main data sources that could be used to enrich the storytelling with e.g., graphics, statistics and reports about past events related to the present one. AI may help the

editorial staff analyse data and identify patterns and insights from these sources, thus reducing the effort for data scraping, aggregation, mining and transformation.

Results of AI4Media are expected to bring benefit along the whole value chain, including journalists, editorial staff, TV production, rights and privacy management departments, as well as media researchers. Finally, to the audience, who will enjoy always new, modern, original and high-quality content.



Fig. 3. AI4Media aims to empower the whole news production process: from news analysis to high-quality and impactful content creation.

3.4 UC4: AI for Social Sciences and Humanities

Scholars from the humanities and social sciences have been adopting digital methods into their respective fields. Large, publicly accessible datasets allow researchers to distill trends beyond individual occurrences of media phenomena. This use case provides researchers (and by extension investigative journalists) with practical methods to sift, connect and analyze various data and media collections in search of factual responses to broad societal research questions. Our aim is to provide the AI-based tools that facilitate identification of patterns or new research questions in aggregated, multi-modal collections. Based on the AI tools, we envision an infrastructure that will facilitate gathering and preparation of data to support *macro-level analysis*, e.g. the role of political and gender bias in media programs, and *micro-level analysis*, e.g. close reading of specific programs, bias effects of speaker selection, and topic presentation. This includes the use of Computer Vision analysis and Natural Language Processing, but goes beyond detection of individual features, to equip researchers with reliable and explainable analyses and visualisations at aggregated levels.

This use case focuses on issues of framing, representation and bias: topics of great interest to social sciences and humanities. However, the definition and operationalization of such complex societal issues is itself subject to debate. Scholars will develop their own understanding and identify more specific instances of these broad phenomena, e.g. the degree to which talk show hosts interrupt female versus male speakers, or the representation of violence in news channels covering a conflict. To facilitate such research with AI-based tools requires a great degree of flexibility and configurability. At the core, the tools should enable researchers to perform basic operations: detect entities in various modalities within large multi-modal datasets. But for example, what exactly the tools detect, which confidence scores for results are to be used and in which combinations entities should occur should be configurable by the researchers.

Finally, academic research has the highest standards when it comes to transparency of methodologies, fairness of representation in training sets and research reproducibility. Therefore, this use case will also evaluate the pre-conditional requirements for trustworthy AI-based tooling, which allows scholars to confidently deploy the tooling.

3.5 UC5: AI for Games

Over the past decade, the game industry has undergone drastic changes. There has been an increase of online games, digital distribution platforms, mobile and social games. Digital games are now services that rely on continuous developing and adding new content over the lifetime of the product in a winner-takes-all market. Before launch, optimizing a game for the market is crucial for success.

AI can provide a radically new approach for testing and optimizing games by letting developers test thousands of different configurations and cases instead of a few hundred or less. It can replace or augment existing practices by providing product evaluations faster than current methods. Additionally, AI may even provide methods for generating content automatically. The game industry is seeing an increased need for automation and improved efficiency. Finally, AI has the capacity to provide new gaming interactions. For the past decade, new control devices have been appearing with the aim to replace the traditional gamepad and provide more natural interaction.

To address these issues, this use case aims to advance the game testing, design and interaction through AI algorithmic innovations. It focuses on three topics:

“Automated testing for Games”: AI agents can be used to automatically detect bugs and glitches in video games by autonomously playing them and compiling test reports as human testers would do. Video game testing is a time-consuming process that traditionally has been entrusted to human testers. However, the future of the game market sees an increasing adoption of the game-as-a-service model, where new content has to be provided to players on a monthly basis. Therefore, quality assurance becomes a vital part of the game development cycle, as the new content produced has to be continuously tested.

“Improved music analysis and synthesis for Games”: This aims to help developers to choose or conceive background music or sounds. In video games, the audio track has the role to strengthen the atmosphere of the present scene, and it has an impact on the emotion felt by the user. Developers can either choose music from a catalog or delegate this task to music composers or sound designers; but automatic AI tools may help, especially in the case of dynamic changes and procedural content generation.

“Natural game interaction through edge analysis of camera stream”: With cloud gaming, end users can play games without dedicated hardware and directly interact with games running on the cloud and streamed to their home. A Smart TV equipped with a standard RGB camera coupled with AI on the edge will replace the gamepad for natural game interactions. The camera stream will be analyzed with AI models that dynamically adapt to the available resources and offloaded to the edge. The output of the analysis (e.g. human body detection, gestures) will be interpreted into game commands within stringent latency requirements and ensure the privacy of end users.

3.6 UC6: AI for Human Co-creation

Content generation is an exhaustive project often developed by a single or a reduced group of people. Many of the tasks involved require a high level of craftsmanship, and as such are available to a reduced number of experts after extensive training. In an era where content is consumed in a greedy manner by a gigantic audience, content creators face the necessity of novel tools to assist them in their creative process.

Music creation uses novel technologies in many ways, and a wide range of techniques are employed by composers routinely. These technologies have been evolving dramatically in recent years, with the introduction of AI into music generation. While this is a great opportunity for creators, it requires an increasing level of technical proficiency, out of reach of non-experts. A novel industry may emerge from this set of possibilities and requirements, giving each user a new personal experience of media addressed particularly to each taste. Authors need novel tools to develop their work and expand their creativity beyond the constraint of potentially automatized tasks.

Our goal is to reduce the gap between novel content creation techniques, in particular raw audio music, and artists/creators. The content may be also created and manipulated in real time according to external inputs. Tasks related to content creation are of different complexity, some of which are already performed in an automated way. In this context, novel tools may contribute to an efficient creation task, where efforts of the artist or creator are focused on deeply creative tasks, relying less on critical parts to be performed by an AI assistant.

The elaboration of creative assistants has a wide application to many different fields of creation. While visual arts are suited to the application of AI tools based on CNNs, recent developments on RNNs or similar techniques allow the efficient processing of music, speech, or video. In this context, the scope of interest of these techniques will soon cover the complete spectrum of content production.

This approach has potential impact in a broad range of disciplines, where the combination of data sources could add value. Beyond visual arts, we find many areas of human creation where language, image, sound, and other information are combined in complex ways, such as games. Tools to generate, modify and synchronise all these data sources are in need to automatically ensure quality and consistency of novel digital content. In this use case, we collaborate with *music composers* with previous experience with the use of AI to help us define new work methodologies and new ways of using AI in their creation. In particular, we focus on coupling raw audio CNNs with GANs trained with mood and emotional labels, lyrics, and other data extracted from music knowledge. Also, general purpose RNNs used in speech synthesis will be explored to process RAW audio from music tracks. The aim is to produce a novel way of content co-creation based on AI tools that reach a wide range of content.

3.7 UC7: AI for (re-)organisation and content moderation

Media companies have accumulated vast digital archives and collections of images and videos over the years. Since these collections have been gradually and iteratively built over time, often by different departments and units of media companies, they

usually have little or no metadata such as tags, categories, and other types of annotations. This lack of coherent media asset organisation tailored to the media company business and services precludes the successful monetisation of these media assets and the creation and offering of new services. In addition, both big traditional media companies and more so digital media platforms combine in their collections both media content, created by these companies, but increasingly also User-Generated Content (UGC). Such hybrid media archives need advanced Content Moderation (CM) solutions, often working in real time to safeguard viewers and meet laws and regulations.




Fig. 4. Imagga's CM platform will be enriched with AI functionalities delivered by AI4Media.

To address these challenges, this use case utilizes several AI-enabled tools such as visual tagging, categorization, and content moderation to facilitate:

a) *Automated (re)organisation of large media collections of photos and video:* Techniques for mapping existing taxonomies and ontologies used by media companies, to restructure them to more optimal ones, and training of specialized models for such companies using state-of-the-art CNNs for media asset categorisation will be used. The know-how of consortium partners with personal and enterprise photo collection organisation will be extended to mixed collections of photos and video.

b) *Automated and human-in-the-loop moderation of user-generated media content:* We will also use existing CM solutions that combine automated detection of diverse and customizable types of inappropriate content, such as weapons, drugs, nudity, pornography, etc. with a CM platform and option to add internal or external data teams to verify inappropriate content flagged by AI algorithms. As part of this use case, Imagga's CM platform [28] (Fig. 4) will be tested and further optimised with mixed and hybrid media collections and platforms where UGC can come both as standalone new visual content (photos, video) or as comments to existing media assets, such as films, news, etc. Additional focus will be on implementing a live-stream optimization module for optimizing costs and precision of automated real-time moderation. This will meet the growing demand of media companies to have scalable real-time content moderation to safeguard users and meet various regulations.

Additionally, the two demonstrators presented in this use case will be tested by target groups of real users from media companies in Living Lab trials, in order to validate their performance and help tailor them further to meet user requirements.

4 Conclusions

This paper offers an introduction to the AI4Media H2020 project. AI4Media aims to deliver the next generation of core AI technologies, including new AI learning paradigms and distributed AI; ensure the embedding of ethical and trustworthy AI into future AI deployments by exploring technologies for AI explainability, robustness and privacy; and reimagine AI as a human-centered, trusted and beneficial enabling technology for media and society. To this end, AI4Media introduces use cases focusing on the media sector that cover different forms of media and digital content while examining various societal and democracy related topics, such as disinformation, public discourse, media moderation, consumer privacy, creativity, and human-machine interaction. AI4Media is expected to build on the rich and diverse European media culture and its role in advancing public values, such as democracy, inclusiveness, political participation, diversity, and self-expression and thus contribute into strengthening Europe's position in the global AI research landscape.

Acknowledgment. This work was supported by European Union's Horizon 2020 research and innovation programme under grant agreement No 951911 - AI4Media.

References

1. Russel, S., Norvig, P.: *Artificial Intelligence: A Modern Approach*. Pearson (2020).
2. AI and the Media: Too Hot, Too Cold, Just Right? A Mapping of Artificial Intelligence Applications, <https://tinyurl.com/7erhbyky>, last accessed 2021/03/09.
3. Sunstein, C.R.: *# Republic: Divided democracy in the age of social media*. Princeton University Press (2018).
4. Smialek, J.: Twitter Bots Helped Trump and Brexit Win, Economic Study Says, Bloomberg article, <https://tinyurl.com/3p4x38uu>, last accessed 2021/03/09.
5. AI4Media Website, <https://ai4media.eu/>, last accessed 2021/03/09.
6. Parisi, G.I., Kemker, R., Part, J.L., Kanan, C., Wermter, S.: Continual lifelong learning with neural networks: A review. *Neural Networks* 113, 54-71 (2019).
7. Bengio, Y., Louradour, J., Collobert, R., Weston, J.: Curriculum learning. In: 26th International Conference on Machine Learning Proceedings, pp. 41-48. ACM (2009).
8. Wistuba, M.: XferNAS: Transfer Neural Architecture Search. In: Hutter, F., Kersting, K., Lijffij, J., Valera, I. (eds.) *European Conference on Machine Learning and Knowledge Discovery in Databases, LNCS*, vol. 12459, pp. 247-262. Springer, Heidelberg (2020).
9. Gravina, D., Liapis, A., Yannakakis, G.: Quality Diversity through surprise. *IEEE Transactions on Evolutionary Computation*, 23(4), 603-616 (2019).
10. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D.: Going deeper with convolutions. In: *Conf. on Computer Vision and Pattern Recognition Proceedings*. IEEE (2015).

11. Yuan, X., He, P., Zhu, Q., Li, X.: Adversarial examples: Attacks and defenses for deep learning. *IEEE Trans. on Neural Networks and Learning Systems* 30 (9), 2805 - 2824 (2019).
12. AI Explainability 360 Website, <http://aix360.mybluemix.net/>, last accessed 2021/03/09.
13. Jalalirad, A., Scavuzzo, M., Capota, C., Sprague, M.: A Simple and Efficient Federated Recommender System. In: 6th IEEE/ACM International Conference on Big Data Computing, Applications and Technologies Proceedings, pp. 53-58. ACM (2019).
14. Ntoutsis, E., Fafalios, P., Gadiraju, U., Iosifidis, V., Nejdil, W., Vidal, M. E., ..., Staab, S.: Bias in data-driven artificial intelligence systems - An introductory survey. *WIREs Data Mining and Knowledge Discovery*, 10(3), (2020).
15. Apostolidis, E., Metsai, A., Adamantidou, E., Mezaris, V., Patras, I.: A Stepwise, Label-based Approach for Improving the Adversarial Training in Unsupervised Video Summarization. In: 1st Int. Workshop on AI for Smart TV Content Production, Access and Delivery Proceedings, pp. 17-25. ACM (2019).
16. Pons, J., Nieto, O., Prockup, M., Schmidt, E. M., Ehmann, A. F., Serra, X.: End-to-end learning for music audio tagging at scale. In: 19th International Society for Music Information Retrieval Conference Proceedings, pp. 637-44 (2018).
17. Nägeli, T., Meier, L., Domahidi, A., Alonso-Mora, X., Hilliges, O.: Real-time planning for automated multi-view drone cinematography. *ACM Trans. on Graphics* 36(4), (2017).
18. Reed, S., Akata, Z., Mohan, S., Tenka, S., Schiele, B., Lee, H.: Learning what and where to draw. In: 30th International Conference on Neural Information Processing Systems Proceedings, pp. 217-225. ACM (2016).
19. Mikolov, T., Sutskever, I., Chen, K., Corrado, G., Dean, J.: Distributed representations of words and phrases and their compositionality. In: 26th International Conference on Neural Information Processing Systems Proceedings, pp. 3111-3119. ACM (2013).
20. Zellers, R., Holtzman, A., Rashkin, H., B. Yonatan Bisk, Farhadi, A., Roesner, F., Choi, Y.: Defending Against Neural Fake News. In: Annual Conference on Neural Information Processing Systems Proceedings, pp. 9051-9062 (2019).
21. Mercier, G., Markatopoulou, F. Cozien, R., Zampoglou, M., Apostolidis, E., Metsai, A. I., Papadopoulos, S., Mezaris, V., Patras, I., and Kompatsiaris, I.: Detecting Manipulations in Video. In: Video Verification in the Fake News Era, pp. 161-189. Springer (2019).
22. Charitidis, P., Kordopatis-Zilos, G., Papadopoulos, S., Kompatsiaris, I.: Investigating the Impact of Pre-processing and Prediction Aggregation on the DeepFake Detection Task. arXiv preprint: <https://arxiv.org/abs/2006.07084> (2020).
23. Kaminskis, M., Bridge, D.: Diversity, serendipity, novelty, and coverage: a survey and empirical analysis of beyond-accuracy objectives in recommender systems. *ACM Trans. on Interactive Intelligent Systems* 7(1), (2016).
24. Koelstra, S., Patras, I.: Fusion of facial expressions and EEG for implicit affective tagging. *Image and Vision Computing* 31(2), 164-174 (2013).
25. Truly Media homepage: <https://www.truly.media/>, last accessed 2021/03/09.
26. TruthNest homepage: <https://www.truthnest.com/>, last accessed 2021/03/09.
27. Mirsky, Y., Lee, W.: The creation and detection of deepfakes: A survey. *ACM Computing Surveys* 54(1), 1-41 (2021).
28. Imagga Content Moderation Platform: <https://imagga.com/content-moderation-platform>, last accessed 2021/03/09.