Annex 1.1

AI4Media – Open Call #2 Challenges

Project Title	AI4Media - A European Excellence Centre for Media, Society and Democracy
Contract No.	951911
Instrument	Research and Innovation Action
Thematic Priority	H2020-EU.2.1.1. - INDUSTRIAL LEADERSHIP - Leadership in enabling and industrial technologies - Information and Communication Technologies (ICT) / ICT-48-2020 - Towards a vibrant European network of AI excellence centres
Start of Project	1 September 2020
Duration	48 months

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 951911

Index of Contents

1 Open Call Challenges

The AI4Media – Open Call #2 is structured around challenges that have been designed to provide value to the AI4Media ecosystem and contribute to the richness of the AI4EU platform. These challenges complement or expand on research being addressed in the project or are aligned with a selection of the project's use cases (see Figure 1).

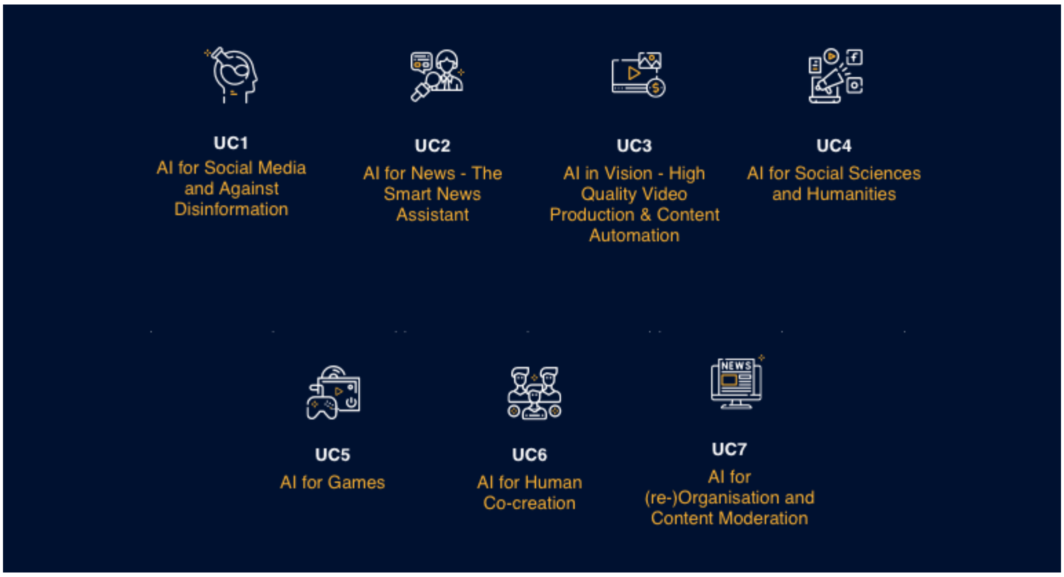

Figure 1. AI4Media uses cases.

When preparing a proposal, the applicant must select the track and challenge to which the proposal will be submitted. A total of 20 challenges have been defined (Table 1) - 12 for the Research track and 8 for the Application track - including an open challenge for both tracks where proposals for other research ideas and application solutions can be submitted.

With the objective of ensuring a distribution of the number of funded sub-granted projects per predefined challenges (C1-R to C11-R and C1-A to C6-A), there is a maximum number of proposals (2) to be selected per challenge. If there are multiple challenges without proposals that guarantees this distribution, primacy will be given to the proposals with the highest score, regardless of the challenge and including those submitted to the open challenges.

Table 1. Challenges of the AI4Media - Open Call #2

Code	Title	Proposals to be selected
RESEARCH TRACK CHALLENGES		
C1-R	Reinforcement learning and the challenge of generalisation	Maximum 2 for each challenge
C2-R	Evolutionary learning and the challenge of evaluating quality	
C3-R	Scarce data and the cross-media transfer of knowledge	
C4-R	New decentralised collaborative learning paradigms for AI	
C5-R	Quantum Reinforcement Learning for classical data processing	
C6-R	Frugal cross-modal representation for media research	
C7-R	Broadening the spectrum of Interpretable AI	
C8-R	Standardising AI datasets and model creation	
C9-R	Novel AI-powered enablers for social media research	
C10-R	Social media alert system to avoid the attention gathering loop	
C11-R	Representative and Inclusive Depictions of AI	
C12-R-OPEN	Open Research Challenge	
APPLICATION TRACK CHALLENGES		
C1-A	Realising a hybrid AI application in AI4EU Experiments	Maximum 2 for each challenge
C2-A	AI global support to informative content production	
C3-A	New learning methods for music overcoming scarce data	
C4-A	AI for suggesting visually appealing images based on text	
C5-A	Recommender Systems to Support Exploratory Research with Media	
C6-A	AI for automated testing and sound synthesis	
C7-A-OPEN	Open Application Challenge	

1.1 Research track challenges

1.1.1 C1-R: Reinforcement learning and the challenge of generalisation

Context:

Reinforcement learning has proven its ability to reach superhuman performance in multiple different tasks. The problem is that most of the advancements take a long time to train on a single task and in the end, they don't perform well on similar tasks without retraining. For example, running a trained model on a different environment than the one it was trained on usually results in a poor performance and the problem becomes increases when we change the task completely. This issue has induced the need to discover better methods and new ways to train these models such that the models learn a general strategy instead of memorising all the possible situations in the current task. Learning a generic strategy will enable these models to work between different tasks with a decent performance without the need for retraining and enable the models to master these tasks with a small amount of fine tuning.

Challenge:

To explore the use of reinforcement learning in other domains than video games (where it is most common), with the aim to overcome current generalisation issues so that trained agents can be used to address different problems and in situations beyond those they were initially trained in. The final goal is, therefore, the development of transferable and generalisable agents that are easy to integrate, are more accessible, and cheaper to train.

1.1.2 C2-R: Evolutionary learning and the challenge of evaluating quality

Context:

Quality-diversity algorithms have been recently introduced to the evolutionary computation (EC) literature as a way of handling deceptive search spaces. The goal of these algorithms is "*to find a maximally diverse collection of individuals (with respect to a space of possible behaviours) in which each member is as high performing as possible*". The inspiration for such approaches is natural evolution which is primarily open-ended—unlike the objective-based optimization tasks to which EC is often applied. While the rationale of open-ended evolution has been previously used as an argument for genetic search for pure behavioural novelty, quality diversity algorithms re-introduce a notion of (localised) quality among individuals with the same behavioural characteristics.

Challenge:

In freeform and subjective domains such as media, AI art, and game content, the definition of quality is minimal (e.g., a level that can be completed, a canvas that has at least some colours) while the evaluation of quality is ill-defined. Diversity is similarly ill-defined, as simple heuristics like colour similarities between two artworks superficially address aesthetic novelty. The exploration of a quality-diversity algorithm is thus at worst unguided (with poor quality and diversity heuristics) or, at best, iteratively adjusting as new knowledge is produced and aesthetic preferences are learned by the AI. Future research trends in this vein for the media sector will have to address four main issues: (a) the type of representations that evolution can explore, (b) the way in which quality and diversity are calculated, (c) ways of modelling designers in order for the AI to produce more personalised artefacts, and (d) interaction paradigms for the users to be able to view, control, and make use of the generated artefacts.

1.1.3 C3-R: Scarce data and the cross-media transfer of knowledge

Context:

Deep learning-based algorithms for multimedia content analysis need a large amount of annotated data for effective training. Having a dataset of insufficient size for training usually leads to a model which is prone to overfitting and performs poorly in practice. But in many real-world applications of multimedia content analysis, it is not possible or not viable to gather and annotate such a large training data sets. This may be due to the prohibitive cost of human annotation, ownership/ copyright issues of the data, or simply not having enough media content of a certain kind available.

Challenge:

There are applications where there is a lack of high-quality annotated training sets and to create high-performance AI algorithms for these scenarios is a challenging issue. The challenge is to explore the issue of scarce data and develop approaches to facilitate the cross-media transfer of knowledge (e.g., knowledge acquired in image-analysis to text). Cross-modal solutions are being used to allow searching for images using text, or vice versa. The following questions emerge: Can existing methods be extended to cross-modal transfer learning? For instance, is it possible to learn how to extract knowledge from images, once we know how to extract knowledge from text? Furthermore, how can foundation models help with the challenge of scarce data learning?

1.1.4 C4-R: New decentralised collaborative learning paradigms for AI

Context:

Decentralised edge computing provides a promising alternative to centralised cloud processing infrastructures for addressing increased data scale and privacy concerns. Existing approaches for AI at the edge focus on (a) federated learning schemes, which are mostly centrally coordinated, (b) the reduction of pretrained AI resources to fit edge devices (e.g., mobile phones), or (c) moving closer to data by performing collaborative learning between edge servers and gateways, with which devices connect at the local level. However, there is still a lack of approaches for harnessing decentralised networks of devices through truly decentralised collaborative learning schemes.

Challenge:

Collaborative learning paradigms could be boosted in terms of efficiency by integrating edge device computations. This could also provide increased privacy, for example by not letting sensitive data leave devices while still using them to learn. To achieve this, several barriers need to be overcome. First, in the real-world, devices tend to exhibit irregular connectivity and poor communication efficiency, both of which are often critical to existing distributed and decentralised learning protocols (e.g., gossip learning). Furthermore, there usually exists strong heterogeneity between individual device data that needs to be properly addressed during learning. In a similar vein, devices are heterogeneous in terms of hardware, which make it hard (if feasible at all) to adapt and deploy the latest AI advances. Therefore, a great challenge is to devise new truly decentralised collaborative learning paradigms that can robustly harness massive numbers of edge devices and compete in terms of scale and accuracy with large, centralised learning schemes.

1.1.5 C5-R: Quantum Reinforcement Learning for classical data processing

Context:

Quantum computation emerges as a revolutionary technology, with the expectation to have deep impact on fundamental and applied aspects of computation, ranging from exponential speedups of complex computations to a reduced cost and power consumption compared to major supercomputers. Detecting practical applications of this novel methodology is a complex task that currently can only be addressed using small Quantum devices already working in laboratories, combined with high performance simulations.

Challenge:

The challenge is to define in a robust way the computing capabilities of quantum computers against current technologies in practical problems, defined by datasets of classical data or containing quantum states. These problems must be selected to be of practical interest yet needing a large computational power. Among these applications, one identifies machine learning related operations as those interesting to boost with advanced technologies such as quantum computing. A particular problem is to match Quantum Machine Learning tasks to data processing problems and assess the advantage over alternative processing methods.

1.1.6 C6-R: Frugal cross-modal representation for media research

Context:

Multimedia content covers all modalities (e.g., visual, audio, text, etc.) and each modality carries its own specific piece of information. To gather all these pieces together and explore the full scope of information all at once, each modality requires to be represented in a common description space where they can all be combined and compared. Recent advances open new opportunities for developing systems to analyse and retrieve multimedia content, with direct applications in the media industry. The current focus of researchers in the field is to design gigantic new models called transformers as unifying models that receive all available modalities of multimedia content and provide solutions for many different tasks all at once.

Challenge:

The challenge is to address the potential of integrating deep learning in designing a new generation of search engines. Two factors must be considered regarding transformer architectures: the resources required both in terms of data to train the models and in terms of computation and energy consumption. Projects should first design and train sophisticated models that require less annotations while combining several modalities, and second integrate external knowledge to reduce the cost of building these models.

1.1.7 C7-R: Broadening the spectrum of Interpretable AI

Context:

The inner workings of machine learning and deep learning are rather opaque, complicated, and indecipherable with accurateness. As the complexity of the proposed algorithms grow, it is complicated to analyse and predict the reliability of the algorithm, and to understand the reasons for success and failure of the model's predictions. This can become increasingly worrying in the context of biased predictions, unfair treatment of subjects and high-risk failures in high stakes.

Challenge:

To widen the horizon of interpretable AI approaches, extending post-hoc approaches to yet under-explored tasks such as structured prediction, regression, and time-series analysis. To provide "glass box" models that can be read, explained, understood, and reliably employed by the humans in the loop of the media industry.

1.1.8 C8-R: Standardising AI datasets and model creation

Context:

In the past few years, deep learning has become the de-facto approach in the multimedia field, spanning all types of applications and covering each major type of data, i.e., visual, textual, and audio. Research in the AI field faces several challenges when it comes to working with datasets. Many deep learning theories offer almost ideal results, but they are impossible to put into practice. Similarly, many deep learning models work well on specific tasks, but their creators cannot offer a sound explanation for choosing a specific setup. Finally, there is a lack of pre-trained models for certain tasks, which excludes low-resourced AI researchers from working on the problem and leads to unnecessary resource consumption in order to retrain models from scratch.

Challenge:

The challenge is to address AI datasets and models from one or several of the following perspectives:

- **Low dataset diversity**: Most AI algorithms have a limited applicability because there are no large and diverse enough datasets that correspond to their needs. This leads to training done on other datasets than what is specifically needed for the task at hand. As a result, the obtained models are biased, and encode a representation of the target task which is only partially adequate.
- **Complex data collection and annotation processes**: Creating a relevant dataset involves a great deal of effort from the involved team. Gathering the dataset samples is just the start of a tedious process and, even at this point, critical problems may arise. It is important to propose dataset creation pipelines which optimize the collection and annotation phases, while considering dataset representativeness, size, and diversity.
- **Closed datasets and models**: The high costs of creating a good dataset and pre-trained model may deter researchers from freely sharing their work. This is, most often, the case for companies involved in research that want to ensure an advantage over their competitors. The public availability of datasets and pre-trained models has the potential to lower the access barrier for low-resource research groups.
- **Inadequate consideration of ethical and legal issues**: Important aspects concerning the ethical creation and implications of a dataset (e.g., harmful content, bias, privacy leaks) or model (e.g., discriminatory behaviour) as well as legal issues (e.g. copyright) are often overlooked by dataset creators, which limits their reusability and decreases trust in the AI community.

1.1.9 C9-R: Novel AI-powered enablers for social media research

Context:

Social media platforms such as Facebook, Twitter, and LinkedIn have been around for many years and have marked indelibly our online communication patterns. Due to their popularity, these platforms are generating large quantities of low-cost data, which can be mined with AI and ML technologies. For example, social media companies monetize social data by creating user behaviour profiles and selling them to advertisers. Many other applications have been proposed, which have clearer societal value and don't compromise user privacy, but their development faces barriers. These include:

- **The existence of walled-garden data**. Since social media companies profit and gain competitive advantage from their platforms' data, they tend to keep them private, in so-called private data silos. Third party access is possible but at a high price or through severely restricted APIs for commercial and non-commercial purposes respectively. This poses significant obstacles to small research groups and innovators working on social media analysis.
- **The inconsistent quality of social data**. Social media are replete with content of questionable nature and provenance such as rumours, fake news, and misinformation.

If not handled properly, this content can lead ML algorithms to erroneous conclusions and insights.

- **The multilingual and multimodal nature of social media**. Due to the global scale and visual character of social media, social data come into multiple languages and modalities (text, image, video, etc). In contrast, these are frequently overlooked by the developed ML models, which have a strong bias for English text. As AI applications are gaining traction, more effort is needed towards a linguistically inclusive and multimodal AI.

Challenge:

Based on the identified barriers, we propose the following research areas to be addressed:

- **The generation of synthetic social data to address the scarcity of relevant data because of social media walled gardens**. Synthetic data have been successfully used for training ML models in computer vision and can alleviate the scalability and privacy issues of deep learning and real user data collection. In addition, the data generators can be customised to diversify data, protect under-represented samples, and mitigate bias. This research area thus targets generative models that can substitute real user data to assist the work of small research groups and innovators on AI-based social network analysis. The developed models will be evaluated in terms of realism (how accurately they represent real data), bias (how fairly they represent data categories, especially sensitive ones), and privacy (how they hide personally identifiable information appearing in their training data from their output). The synthetic data is not restricted to a specific modality.
- **New models and approaches for the automatic credibility assessment of social media data**. As the quality and accuracy of social media information is questionable, assessing its credibility via automated means becomes critical. This research area thus targets novel techniques for automatic information verification such as false information detection, fact-checking, and authoritativeness scoring. The proposed techniques can exploit both intrinsic (derived from the data) and extrinsic (derived from external online sources) information, and the most comprehensive solutions will be evaluated highest. The solutions are not restricted to a specific data modality, in fact, multimodal approaches are highly encouraged.
- **New AI models and use cases for social media analysis in non-English social media**. As most ML models cater primarily for English, non-English social media and user groups are left behind in the AI revolution, especially in the case of minority languages. Additionally, without multilingual support, AI applications miss the opportunity to cross-pollinate knowledge from multilingual content, which is more valuable than the knowledge derived from content of a single language. This research area thus targets multilingual models and use cases that consider multiple idioms and cater to non-English social media. The proposed solutions will be evaluated in terms of number of supported languages, language inclusiveness, and the capacity to extract added value from

multilingual content. The accurate treatment of informal language and slang will be highly evaluated.

1.1.10 C10-R: Social media alert system to avoid the attention gathering loop

Context:

During recent decades, political and social communication has been transformed by digital technologies, and particularly, social media. Among the challenges of modern communication is the economy of attention. Citizens have a limited capacity for consuming information, creating a coveted (both by producers and consumers) market of audiences. The rise of the attention economy plays a big role in current democracies, and the role of participation by citizens in democracy. People are more informed and empowered than ever, having powerful tools at their disposal for defining and disseminating their opinion. However, citizens are also subject to an unprecedented amount of information, sometimes due to poor quality, designed mainly for the purpose of gathering attention.

Challenge:

To explore a social media recommendation system for alerting users when unconsciously falling into an attention gathering loop, applied either into text-based social media platforms or applied into multi-modal based (text & image) social media platforms.

1.1.11 C11-R: Representative and Inclusive Depictions of AI

Context:

Images that are currently used to depict AI in the mainstream media often lean on visual tropes such as glowing brains, white robots and men in suits interacting with holograms. These images are harmful on several levels. Firstly, they are not representative of the majority of current AI research and application areas, such as natural language processing, content classification techniques or recommender systems. Secondly, they reinforce damaging gender and ethnic stereotypes. Finally, they create misleading understanding of AI's capabilities among the public and can sow fear of the technology. Initiatives such as Better Images of AI have shown that there is much more research needed to understand how to effectively communicate about AI through images.

Challenge:

Since effective communication about AI is essential for its uptake and acceptance by media organisations and the public, this challenge invites interdisciplinary approaches for producing more representative and inclusive images depicting AI. In their projects, participants are invited to address all the following:

- Conduct research around the criteria for producing more representative and inclusive depictions of AI by involving stakeholders from research, industry, and creative industries.

- Propose methods and conduct activities for creating such images (for instance, design workshops, generative AI models).
- Produce a set of openly licensed images that could be used by AI researchers and the media industry to communicate about AI.

1.1.12 C11-R-OPEN: Open Research Challenge

Innovative research ideas that do not fall within the previously listed research challenges can be submitted to the Open Research challenge so long as they are aligned with one of the following topics, while also targeting the media sector:

- AI research & technologies of tomorrow.
- AI multimedia applications of tomorrow.
- Trustworthy AI: future trends for robust, interpretable, privacy-preserving, and fair AI.
- AI data, benchmarks, and open repositories: towards AI democratisation.
- AI applications and solutions for the media sector: imagining the future of next-gen media.

1.2 Application track challenges

1.2.1 C1-A: Realising a hybrid AI application in AI4EU experiments

Context:

The AI4EU Experiments platform allows the easy design and deployment of AI pipelines consisting of predefined modules. For example, such a pipeline may first transcribe spoken information into written text and then recognize the named entities such as persons, locations, and organisations in that text. Many of the modules for AI4EU Experiments are publicly available and are intended to be re-used for different pipelines and in different contexts. In this way, the marketplace of AI4EU Experiments modules serves as an open AI repository.

In many areas of application, machine learning ("AI") methods outperform classical algorithms by far. However, there are many computational tasks where machine learning methods must be combined with symbolic reasoning to leverage their full potential. This approach is called hybrid AI and, as it turns out, plays an important role for many applications in the media sector; see below for some examples.

Challenge:

The challenge consists in the realisation of a hybrid AI application on AI4EU Experiments. More concretely, the requirements are as follows:

- The main outcome must be a working pipeline on AI4EU Experiments.
- This pipeline must combine machine learning modules with modules for symbolic reasoning.
- The pipeline and the modules therein must be provided under an open-source licence.

- It must be possible to run the pipeline in the AI4EU Experiments playground.
- The pipeline must offer a web UI which allows basic handling and progress tracking, depending on the implemented functionality.
- The functionality provided by the pipeline must be suitable for the media sector.

Applications in the media sector where hybrid AI pipelines can be useful include, but are not limited to, the following areas:

- Searching items in large collections based on a combination of criteria, for example:
    - Searching for images of a certain person having a certain emotional expression
    - Searching for photos of a certain landmark taken during a certain season or time of day
    - Searching for images showing a certain number of objects from a certain category
- Mapping results from machine learning modules to predefined vocabularies or taxonomies, for example:
    - Classification of news items like texts, photos, graphics, audio, or video files according to the IPTC media topics
    - Mapping of results from topic modelling or semantic tagging modules to some controlled journalistic vocabulary, for example IPTC media topics
- Adjusting machine learning methods to limited areas by filtering the results symbolically, for example when applying video segmentation to a certain TV format
- Combining machine learning with knowledge graphs

As part of this challenge, it will be necessary to design suitable APIs for the modules in the AI4EU Experiments pipeline which is to be created. Since this is, to our best knowledge, the first time that such APIs are designed for hybrid AI pipelines in AI4EU Experiments, the resulting APIs themselves and some best practices regarding the design of APIs for hybrid AI pipelines should be properly documented.

Further resources:

- AI4EU Experiments platform: *https://aiexp.ai4europe.eu/#/home*
- Container Specification: *https://github.com/ai4eu/tutorials/tree/master/Container_Specification*
- Source Code Tutorials: *https://github.com/ai4eu/tutorials*
- YouTube Playlist: *https://www.youtube.com/playlist?list=PLL80pOdPsmF6s6P6i2vZNoJ2G0cccwTPa*
- IPTC media topics: *https://www.iptc.org/?_standard=media-topics*

1.2.2 C2-A: AI global support to informative content production

Context:

Nowadays, informative content (e.g., news, talk shows, magazines, documentaries) is made readily available and accessible on multiple sources: online, radio, television, social networks, etc. Information is continuously produced by large groups of professional journalists and communication experts. The overall process can benefit from AI tools, in the generation/ transport, indexing and retrieval and publishing phases.

Challenge:

The challenge is to develop new methods and tools that contribute to the advancement of one or more of the following informative content production areas:

- **Efficient video compression**: Video compression is a critical component of Internet and broadcast video delivery. Video content has reached 82% of internet traffic volume today (according to Cisco), hence the interest in efficient video encoding technologies that can handle tomorrow's bandwidth-hungry video services (4K video, high frame rate, immersive content, etc.). However, since resolution and frame rates are increasing at the same time, existing video coding standards used in Internet streaming or aerial broadcasting are no longer sufficient for some applications. Over the past years, the research community has investigated the recent developments in Artificial Intelligence (AI) and Machine Learning (ML). The challenge is to overcome 'classical algorithms' limitations and deliver industry-leading video quality considering bandwidth constraints and hardware efficiency.
- **Media indexing and retrieval**: It is a hard task for computer systems to analyse, understand and represent media content. Neural networks have been applied to this task and in several domains (among which e.g., image classification and ASR), but they need more investigation and trial in the professional media domain. Furthermore, the potential of existing media archives and datasets in this domain must be fully exploited. In this context, a specific and highly challenging task is that of media clipping, i.e., the capability of a system to identify meaningful clips or segments of a larger media item that have a coherent semantic structure.
- **AI-assisted multiplatform publishing**: modern digital publishing requires that content items are adapted and published rapidly and efficiently to reach the users of the various platforms. To achieve this, AI-based summarisation techniques are needed, as well as content generation tools capable of producing publishable images, text, and videos from a variety of inputs (natural language descriptions, tabular data, sample content).

1.2.3 C3-A: New learning methods for music overcoming scarce data

Context:

In the early stages of automatic music processing, such as in musical representation or sound signal modelling, ad hoc methods and expert algorithms were explicitly designed by the researchers or engineers to achieve specific tasks. Nevertheless, these approaches were not sufficient to explain more complex concepts of music, such as the musical genre or the emotion provided by a song. Therefore, since almost two decades, machine learning has been used to build models able to automatically learn a musical concept, without having to formalise it mathematically. For some years now, deep learning has made it possible to go even further. Many ad hoc methods, which worked fine, have been overtaken by deep neural network models.

Challenge:

One standard problem in machine learning, and especially in deep learning, is the lack of data. For most of the applications in music processing, the creation of such datasets is very difficult and can be expensive. The challenge is to develop or to adapt for the music new learning methods able to be trained with scarce music data.

1.2.4 C4-A: AI for suggesting visually appealing images based on text

Context:

Nowadays, journalists are working in an increasingly challenging environment. The rise of (social media) platforms has led to the need to publish as fast as possible and to be present on an increasing amount of those platforms to reach the target audience. More people, especially young people, are using social media as a news source. Thus, news professionals need to adapt their publication strategies and publish on these platforms to match media consumers' expectations.

Challenge:

The challenge is to develop image suggestions based on text to create visually appealing stories and save time. Specifically, the following topics are of interest:

- **Natural Language Understanding**: Currently, there is no algorithm available that quite captures the concept from a text except from being rather general (e.g., a text can be about 'war', but which aspect is highlighted?).
- **Image suggestions**: Subsequently, those general concepts yield visual content that is not very appealing, surprising, or inspiring. As this should be the eyecatcher to convince someone in a spit second to continue reading the article, it's not on par yet with a human curating the right image. Questions that come to mind when investigating this further: *Does AI understand what the correct image is? What defines a correct image?*

1.2.5 C5-A: Recommender Systems to Support Exploratory Research with Media

Context:

Recommender systems are by now widely used on content platforms as well as professional media production applications. Could they equally support scholars from the fields of media studies, journalism and digital humanities who conduct research with media content? The potential of this application area is yet to be explored. Given the increasing amount of multimodal data available for investigation, recommender systems could be employed to enhance research workflows by suggesting content items from a variety of sources. For such applications, it seems particularly promising to combine approaches for collaborative filtering (e.g., based on user feedback or user preferences) and content-based filtering (e.g., exploiting concepts derived via automatic metadata extraction). If relevant, it is also important to address possible privacy requirements, using appropriate measures and technologies.

Such recommender systems are expected to help researchers navigate large data collections, invite them to learn about a topic from different perspectives and provide a way to explore content from both mainstream media and user-generated content in one place. They would be particularly relevant for exploratory research stages where a researcher is trying to define their research question or find an angle they would focus on.

Challenge:

This challenge asks participants to design a prototype of a recommender tool designed to support research in media studies, journalism, or digital humanities. Participants should show that they have or are able to gain access to adequate datasets (regarding both data quantity and variability) and should try to include both professional media and user-generated content. A combination of content-based and collaborative filtering should be used. Participants should also be open to communicate and cooperate with AI4Media project partners and/or use AI4Media results to achieve the aforementioned goals.

1.2.6 C6-A: AI for automated testing and sound synthesis

Context:

The game industry has changed their business model and the types of games they create compared to early days of the industry. New games start adopting the idea of being a service instead of being a product. This means the game life cycle last as long as the company exists. Companies need to produce new content on regular intervals to keep the current audience happy, attract new audiences, and stay ahead of their competition. Companies must make sure the new content fits the current player skills, and that it will provide new experience for the user. They also need to make sure that the new content won't break the game. All these advances increase the production cost of games to an unsustainable level such that a bad game/update release from a company could cause companies to go bankrupt.

Challenge:

The challenge is to reduce the testing and production cost of new games. The challenge is divided into two sub-use cases (note: applicants should select one sub-use case):

- **Automated testing for games**: the amount of time wasted by QA testers repeating a small task can be optimised using AI agents. The QA tester's responsibility is to set up the test and its condition and the AI will cover the space to find different types of glitches and errors that can exist in the game.
- **Music Synthesis for games**: creating new music takes a long time. Most of that time is wasted in trying to find inspiration of what to create such that it is unique and fits the game mood. Having a tool that can generate short music samples based on a certain

mood can help music composers as an inspiration to create a unique music track for the current game.

1.2.7 C7-A-OPEN: Open Application Challenge

Innovative ideas that do not fall within the previously listed application challenges can be submitted to the Open Application challenge so long as they are aligned with one or more of the AI4Media use cases (Figure 1). Solutions must focus on being close to market and therefore target a minimum TRL 7 at the end of the sub-granted project. Applicants are invited to review the use cases *here*.

AI4media

ARTIFICIAL INTELLIGENCE FOR THE MEDIA AND SOCIETY

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 951911

info@ai4media.eu www.ai4media.eu