

Success stories from AI4Media Open Call #2

Empowering research and
applications in Artificial
Intelligence for the Media industry





The main objective of the AI4Media project is to establish a network of Artificial Intelligence (AI) excellence that will bring together the currently fragmented European AI landscape in the field of media and will foster collaborations between academia and industry as well as with the European AI community.


AI4Media has delivered a competitive funding programme to attract and involve researchers and SMEs to the AI4Media Network, to boost cross-border and cross-sector innovation within media sectors. The funding programme had a total of €1 million to fund 20 innovative projects focusing on AI-driven research or applications for the media sector.

Open Call #2 overview

The objective of the AI4Media – Open Call #2 was to engage research and industry actors that develop and integrate applied research in the field of AI, to develop new research and applications for AI, and contribute to the enrichment of the pool of research and technological tools to be made available on the AI-on-Demand (AIoD) platform.

The AI4Media Open Call #2 targeted two types of projects, divided into two tracks:

- ➔ **Research track:** targeted academia and researchers working in the AI field to increase the value of the AI4Media ecosystem by developing and integrating new research in AI media domains. Research track projects ran for 12 months and each received up to €50.000.
 - ➔ **Application track:** targeted entrepreneurs, companies (micro-SMEs, SMEs, mid-caps) to submit innovative AI applications for the media sector by building on the research coming out of AI4Media or by adopting AI solutions from other sectors to the media sector. Application track projects ran for nine months and each received up to €50.000.
- 




The AI4Media – Open Call #2 launched on 29 September 2022 and closed on 30 November 2022. A total of 95 proposals were submitted to the open call, 63 addressing the Application track and 20 addressing the Research track. Of the 95 submissions, 54 proposals were submitted by industry entities (e.g. start-ups, SMEs) and 12 from research organisations. The remaining 29 were submitted by individuals, secondary and higher education institutions, and other entities. Submissions were received from 24 different countries. Spain was the most represented country with 13 submissions, followed by Germany and the United Kingdom with nine submissions each. There were several countries with a minimum of five submissions (e.g., Greece, Netherlands, Portugal, and Romania).

Open Call #2 programme

The AI4Media – Open Call #2 funding programme launched on 1 March 2023 with the five research and five application projects.

In addition to the €50,000 to fund their activities, each project received mentorship support provided by the AI4Media partners. The 10 third parties funded also participated in a Bootcamp that provided them with training on specific topics aligned with their interests and needs, including marketing and communication, funding opportunities, business acceleration, and technology transfer and IPR.

The five application and five research tracks successfully ended their projects in November 2023 and February 2024, respectively, having provided relevant advances to the AI and media domains, and valuable contributions to the AI and media-driven ecosystem.



CAMOUFLAGE

Full title of the project

Controllable anonymization through diffusion-based image collection generation

Project tagline/ slogan

Beyond Faces: Comprehensive attribute-preserving image anonymization with latent diffusion models

Project Track

Research

Funding period

1 March 2023 - 28 February 2024

Beneficiary lead name

Politecnico di Torino

Type of organisation

University

City & Country

Turin, Italy

Website

<https://www.polito.it/>

Main contact

Lia Morra

Position in the organisation

Assistant Professor, Tenure Track

Other people involved in the project

Luca Piano

Pietro Basci

LinkedIn profile

<https://www.linkedin.com/in/liamorra/>

Brief description of the project

Social media generates a tremendous amount of visual material. However, privacy regulations impose significant restrictions to data collection and sharing for researchers in social sciences, digital humanities and artificial intelligence. CAMOUFLAGE leverages recent advances in controllable latent diffusion models to generate completely synthetic versions of a target image. Specifically, CAMOUFLAGE extracts non-sensitive data from the original image in order to constrain a diffusion model to preserve the composition of the image, as well as key properties of the subjects depicted. At the same time, CAMOUFLAGE prevents image and facial recognition systems from retrieving the original image, thus protecting peoples' privacy.

What was the challenge?

An ambitious project like CAMOUFLAGE presented several challenges: first, how to design the artificial intelligence components needed to generate an image with similar, but not too similar, properties; second, how to properly validate the output. Indeed, classic image anonymisation such as blurring or pixelization clearly maintain a link to the original image, but make them difficult to interpret. The more we replace part of the original images with synthetic content, the more difficult it becomes to determine which aspects should be reproduced, which should be changed, and how the resulting image would be interpreted by a human or a machine.

What results did you deliver (the solution)?

We delivered two versions of the CAMOUFLAGE system. CAMOUFLAGE-Base exploits a combination of pre-trained ControlNets and introduces an anonymisation guidance based on the original image, while CAMOUFLAGE-Light trains a lightweight IP-Adapter to encode key elements of the scene and facial attributes of each person. The former achieves stronger anonymization, while the latter generally preserves image content better and simultaneously reduces the inference time by 75%. Compared to the state-of-the-art, CAMOUFLAGE can anonymize complex scenes by introducing variations in the faces, bodies, and background elements, and can prevent re-identification from both faces and backgrounds.

What is the impact of your project for the AI and media ecosystem?

CAMOUFLAGE will provide open source code and models, making both versions freely available to the research community to use and study. Besides advancing the state-of-the-art on image anonymization, the resulting images could be used to enhance training dataset for machine learning models, making them more fair and inclusive.

Organisation logo / project logo



Main contribution/ value from the AI4Media project/ programme

Participating in the AI4Media programme as a funded project and associate member provides access to a network of excellent research centres. It provides a variety of training opportunities for both PhD candidates and SME. In addition, the AI4Media programme offers several opportunities and channels for dissemination of project activities.

What you did and achieved that would have been difficult/ impossible without this funding?

Open-call fundings are excellent opportunities for early career researchers, those who are starting a new group, or a novel research line, for which securing larger funding at the European level would be difficult. This funding allowed our group to start a new line of research on controllable image synthesis.

What's next for the project?

We are currently working towards disseminating CAMOUFLAGE results to the research community, not only in the artificial intelligence and multimedia field, but also towards the social science and digital humanities. We are working with researchers in visual semiotics from the FACETS team (<https://www.facets-erc.eu/>) at the University of Turin to further validate the system, while preparing to release the code. CAMOUFLAGE will be presented at the Biennale Tecnologia (biennaletecnologia.it), at the IEEE Games, Entertainment and Multimedia conference (<https://www.ieee-gem2024.org/>), and many more venues to come. At the research level, we are looking forward to open new research lines in controllable image synthesis, beyond anonymization.



| Real image (top) versus anonymised image by CAMOUFLAGE (bottom)

Relevant lessons learnt

CAMOUFLAGE has been an intense experience. Thanks to the project, we acquired new team members, new competences and skills, and strengthened our collaboration with the University of Turin. The short time frame forced us to be very selective and focused on the objectives we wanted to achieve, but we definitely needed to set aside some time after the end of the project to properly disseminate the results.

Total FSTP funding

€49,750

Brief testimonial about participation in the programme, including interaction with AI4Media and coaches, if you'd recommend such programmes to others.

Our experience with the AI4Media programme was definitely positive. Even though some initial themes were defined and our project was in my opinion a good fit, the submission itself was evaluated based on its scientific outlook and overall quality, rather than strictly adhering to a predefined theme.

Everything was handled with great care and professionalism, from the selection process, to the timely and useful feedback provided by the coaches. I would definitely recommend the programme and participate again if the opportunity arises.

People involved in the project



CLIP Lens

Full title of the project

CLIP models, Looking for Enhanced New Systems

Project tagline/ slogan

Filling data gaps with synthetic images for multiple data-scarce scenarios

Project Track

Research

Funding period

1 March 2023 - 28 February 2024

Beneficiary lead name

CENTIC (Asociación Centro Tecnológico de las T.I.C. de la Región de Murcia)

Type of organisation

Non-profit business association

City & Country

Murcia, Spain

Website

<https://centic.es/>

Main contact

Raquel Espinosa

Position in the organisation

Data Scientist

Other people involved in the project

Javier Abellán

José Miguel Bolarín

LinkedIn profile

<https://www.linkedin.com/company/centic/>

Twitter profile

<https://twitter.com/CenticTwitt/>

Organisation logo / project logo



Brief description of the project

CLIP Lens addresses data scarcity in various domains using cutting-edge technologies like few-shot classifiers (CLIP), and robust generative data augmentation (DreamBooth and Stable Diffusion) to generate synthetic images. Additionally, LoRA has been used to reduce the computational costs of the customised model. This methodology is applied across different use cases such as clothing items, professions, and objects. Image quality was assured by using quantitative metrics, such as CLIP score and CLIP zero-shot, and qualitative metrics, through a survey to determine the quality of the images generated.

What was the challenge?

The main challenge was the training of customised models from 4 or 5 example images. These models had to generate realistic images with reduced computational time and cost. Initially, the use of Stable Diffusion alone was able to generate images, but it could sometimes lead to overfitting and the training time was very high. However, the introduction of DreamBooth and LoRA solved these problems. In addition, finding example images was also costly, as there were not a large number of datasets that met the requirements of the project.

What results did you deliver (the solution)?

The main result of the project has been the proposal, evaluation and validation of a methodology capable of generating realistic and quality images from a very small sample of images. For the three proposed use cases, three datasets with relevant images have been created and published. These results have been validated quantitatively by using metrics such as CLIP score and CLIP zero-shot and qualitatively with a survey (conducted among AI experts and non-experts) about the quality of the generated images.

What is the impact of your project for the AI and media ecosystem?

On the AI side, it has been proven that the use of LoRA is able to generate customised images without having to re-train the original Stable Diffusion model. On the media ecosystem side, this methodology is capable of both generating images from other sample images and from a prompt. Thus, media professionals will be able to generate realistic images, for example, to accompany news stories, without having to worry about copyrights or without having to do a time-consuming search.

Main contribution/ value from the AI4Media project/ programme

One of the main contributions that AI4Media had in this project was the figure of a coach. Having an expert in the area to investigate has made our work easier and his advice has helped us to improve the project, especially on the technical side. In addition, the bootcamp organised by AI4Media allowed us to learn more about other aspects of the business such as communication, product development, marketing tools or European funding opportunities, among others.

What you did and achieved that would have been difficult/impossible without this funding?

Without funding, forming a team skilled in synthetic image generation would have been quite challenging. The financial support facilitated the training of some team members in the technologies used in the project, such as DreamBooth, Stable Diffusion or LoRA. This has allowed us to have a deeper knowledge of the subject and to have been able to successfully develop the methodology proposed in the project. It has been possible to establish a state of the art with all the technologies relevant to date in the context of image generation.

What's next for the project?

Over the next year, the project will follow several strategic activities to ensure effective use of its results. In the first three months, detailed documentation of the methodology and employed technologies will be compiled, followed by an analysis of resulting datasets. Months four to six will prioritise enhancements in image quality to broaden applicability. In the subsequent three months, efforts will focus on identifying potential partners for collaboration. Finally, the project will develop a roadmap for long-term sustainability, encompassing maintenance, updates, and adaptation to market trends and technological advancements.



| Series of images generated with the CLIP Lens project technology

Relevant lessons learnt

Integrating diverse technologies such as DreamBooth, Stable Diffusion, and LoRA allowed us to learn the importance of leveraging multiple tools to enhance the robustness and flexibility of synthetic image generation systems. Also, ensuring the quality of synthetic images was paramount. We learned the importance of applying rigorous validation and quality control measures to maintain realism and relevance, avoiding the propagation of inaccurate or biased data. Lastly, we gained insights into techniques for reducing computational time and costs without compromising on the quality of generated images.


Total FSTP funding

€48,125

Brief testimonial about participation in the programme, including interaction with AI4Media and coaches, if you'd recommend such programmes to others.

The interaction with both AI4Media and our coach has been incredibly enriching and useful. Not only did we gain valuable insights into cutting-edge AI technologies, but we also received personalised guidance and support that improved our project. The collaborative environment fostered by the programme encouraged creativity and innovation, and we felt genuinely supported every step of the way. We would recommend this programme to anyone looking to delve into the world of AI and media, especially startups or SMEs. In addition, its philosophy of contributing to a repository of public datasets and software is very inspiring.

People involved in the project



ELMER

Full title of the project

Efficient Long-term Multimodal Video Retrieval

Project tagline/ slogan

To develop an efficient system for content retrieval that can handle multimodal audio, image, text, and video data for long video.

Project Track Research

Funding period

1 March 2023 - 28 February 2024

Beneficiary lead name

University of Surrey

Type of organisation

University

City & Country

Guildford, UK

Website

<https://www.surrey.ac.uk/>

Main contact

Andrew Gilbert

Position in the organisation

Associate Professor of Machine Learning

Other people involved in the project

Edward Fish

Jon Weinbren

LinkedIn profile

www.linkedin.com/in/andrew-john-gilbert

Brief description of the project

The objective of this project was to design an efficient content retrieval system that can process multimodal data, including audio, image, text, and video, specifically footage that extends beyond a duration of 10 seconds. To achieve this, a zero-shot learning technique that employs pre-trained networks was investigated, taking into account various modes of data.

What was the challenge?

Integrating multiple visual, auditory, and temporal features for diverse video classification and retrieval tasks poses a considerable challenge that has yet to be resolved. Furthermore, the issue's complexity is compounded by the need to address missing or occluded data across multiple modalities. To address this challenge more efficiently, reducing the resource requirements regarding the dataset and computation power is essential. This is achieved through leveraging pre-trained models.

What results did you deliver (the solution)?

The project created two submissions of knowledge contributions. The first submission leveraged audio and video to facilitate the querying of locations of activities in home and sports videos, enabling the identification and retrieval of examples of cutting carrots or hitting a tennis ball, for instance. The second contribution involved retrieving activities in videos by employing significantly less training data to generalise the approach. This approach proved more effective in capturing the general characteristics of activities, thus mitigating the risk of overfitting to a few examples of activities that otherwise happen when only a few training videos are used.

What is the impact of your project for the AI and media ecosystem?

During the funding period, these solutions were submitted and presented at two major machine learning conferences, NeuIPS'23 and ECCV'24. The code has also been released to further enable other researchers to build and develop our two ideas. The team has joined the AI4Media ecosystem as associate members.

What you did and achieved that would have been difficult/impossible without this funding?

AI4Media provided the staffing costs to enable a researcher to work part-time on this project over the year, allowing the researcher to develop cutting-edge solutions to the problem.

What's next for the project?

We plan to investigate new avenues in our research, particularly regarding audio integration in multimodal video-temporal action localisation. During the project, we discovered that audio in video footage, especially highly edited ones, can be counterproductive and distracting instead of helping us understand the footage. This was an unexpected finding, so we plan to explore ways to incorporate audio using the modal gating or optical transport research done later in the project, but currently only with video. However, this is dependent on securing further funding.

Organisation logo / project logo



Relevant lessons learnt

The addition of multiple modalities, such as audio, alongside video, as a means of improving performance is not always practical. This is primarily due to the presence of distractions within audio. Notably, many sporting event videos lack localised audio information or include music and narration that are tangential to the video content. Such audio components may have a detrimental impact on performance. Therefore, it is imperative to carefully consider the appropriateness of adding audio to video content, with due consideration of the potential for audio to introduce distractions and hinder performance.

Total FSTP funding

€50,000

Brief testimonial about participation in the programme, including interaction with AI4Media and coaches, if you'd recommend such programmes to others.

The guidance provided by AI4Media during our engagement proved to be an invaluable asset to our team's efforts. Our assigned coach offered a wealth of insightful queries and constructive suggestions, which played a critical role in further fortifying and consolidating the two major conference submissions we had under development. We are grateful for the assistance provided by the AI4Media team, and we look forward to the opportunity to leverage their expertise again in the future.




Illustration of our approach using multiple prompts to describe and locate the activity Golf Swing in the video

People involved in the project



HoloNeXT

Full title of the project

Holographic transmission and Neural radiance fields for a novel Xr production Tool

Project tagline/ slogan

Automatically create and deliver new integrated 3D media content using novel neural and volumetric technologies

Project Track

Research

Funding period

1 March 2023 - 28 February 2024

Beneficiary lead name

i2CAT Foundation

Type of organisation

Research and innovation centre

City & Country

Barcelona, Spain

Website

<https://i2cat.net>

Main contact

Ivan Huerta

Position in the organisation

Research Line Lead

Other people involved in the project

Gianluca Cernigliaro

Sergio Montoya

Leonel Toledo

LinkedIn profile

<https://www.linkedin.com/company/i2cat/>

Twitter profile

<https://twitter.com/i2cat>

Organisation logo / project logo



Brief description of the project

Modern media are transitioning from a single-dimensional nature becoming more dynamic, interactive, and including novel technologies such as AI and XR. However, those are still limited to traditional channels and media domains. Thus, there is a need to address the challenges for producing content in this mutable environment allowing users to experience media in different platforms and offering access and transitions between different media contents and formats. Our proposed solution explored novel Neural and Volumetric technologies to provide a solution capable of automatically delivering new integrated 3D media content and formats for XR TV production. This solution provided a novel XR Media Production Tool integrating the two volumetric/XR technologies represented by 1) Neural Radiance Field scene modelling and 2) holographic real-time video volumetric transmission. The proposed solution was demonstrated in a news media XR production use case where a news presenter was teleported to a 3D reconstruction.

What was the challenge?

The main challenge was to provide a novel XR production tool integrating two volumetric/XR technologies represented by 1) Neural Radiance Field scene modelling and 2) holographic real-time video volumetric transmission. Additionally, a secondary challenge was to Advance the State of the ART (SoA) of automatic 3D neural radiance field volumetric scene modelling based on a collection of 2D images, video stream, or RGBD sensors.

What results did you deliver (the solution)?

Firstly, research and develop a new NeRF called RefMesh for the automatic training and creation of photometric scenes using NeRF, providing a measurable improvement in terms of capturing and scene modelling time, and render quality when compared to the state of the art, capturing the scene in less than 15 minutes and modelling it in less than 2h. Secondly, integrate this novel NeRF technology with a Holoportation pipeline within the XR Media Production tool, providing a measurable improvement in terms of resolution and frame rate when compared to the state of the art with a frame rate of 30 fps and resolution of 300k voxels.

What is the impact of your project for the AI and media ecosystem?

The project produced an impact on:

1. **Content automation and high quality video production for broadcasting** by improving the state of the art in neural based content creation and by providing volumetric live streaming tools for more immersive media tools.
2. **The Gaming industry** by i) including a novel an immersive representation for the players that will now be holoported within the game and ii) using NeRF for novel, AI based 3D content creation with photo-realistic quality.

3. **Human Co-Creation** by providing a final tool where holoportation and NeRF based content creation are integrated in a tuneable VR environment.

Main contribution/ value from the AI4Media project/ programme

Firstly, provided the funds to conduct research and advance the state of the art in scene modelling using NeRFs and volumetric live communications. Secondly, appointed an administrative contact (monitor) during the project implementation phase, who assisted with organising deliverables and conducted review meetings at the end of each sprint. Thirdly, assigned a coach to support and motivate the project team, providing expert feedback on progress and evaluating deliverables. Fourthly, organised a boot camp with valuable information on various aspects essential for achieving project success. Finally, in the future, we expect AI4Media to provide the network and opportunities to further develop the research of this project and facilitate connections with companies interested in the technology.

What you did and achieved that would have been difficult/ impossible without this funding?

The funds allowed us to research and advance the state-of-the-art technologies in order to pave the way to novel AI based XR Media Production Tools. In these tools, i) NeRF will make available the creation of new content in an automatic or semi-automatic manner, focusing on reducing the time required for capturing and scene modelling, and maximising the render quality of the modelled scene. And ii) Holoportation will enable volumetric live communications in a more immersive manner.

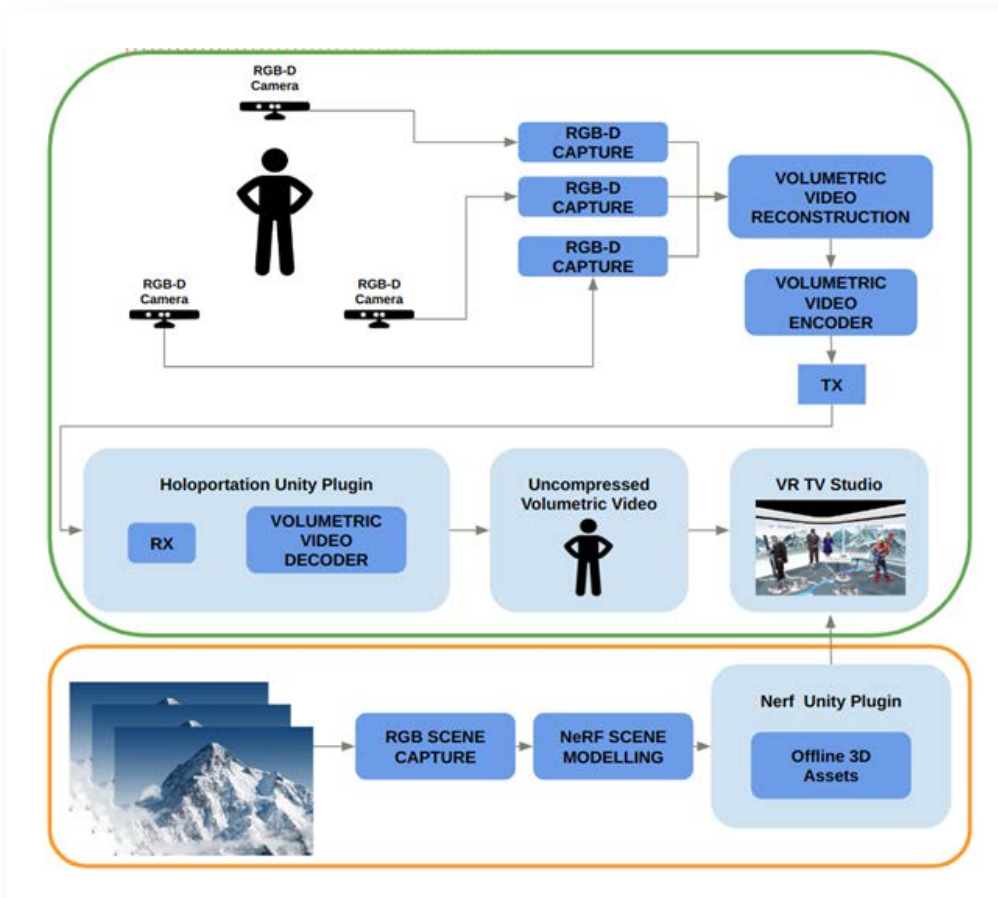


Illustration of the HoloNext architecture

What's next for the project?

The work completed will be leveraged for future national and European calls, where we will continue with the research on NeRF for fast and accurate scene modelling and with the development of more accurate and faster real-time volumetric live communications.

Relevant lessons learnt

There is still a significant gap between the photometric results obtained using neural rendering technology (NeRFs) and the 3D reconstructions derived from them. Obtaining an accurate mesh quickly from a series of in-the-wild images, with sufficient quality for media production, remains an ongoing research area.

Total FSTP funding

€50,000

Brief testimonial about participation in the programme, including interaction with AI4Media and coaches, if you'd recommend such programmes to others.

Participating in this open call was truly a rewarding experience and an excellent programme. It provided us with the opportunity to conduct research freely, evaluating results without the obligation to provide a final solution. This allowed us to delve deeper into media technologies that would otherwise be challenging to explore due to the absence of a clear client. The programme was well-structured, and the administrative contact was always approachable and understanding, particularly during delays in deliverables. The coach offered valuable advice, despite it not being his primary expertise. I would highly recommend this programme to others, and I intend to apply for future calls if given the opportunity.

People involved in the project



JECT-CLONE

Full title of the project

JECT – Creative Landscapes of News

Project Track

Application

Funding period

1 March 2023 – 30 November 2023

Beneficiary lead name

JECT.AI Limited

Type of organisation

Start-up

City & Country

London, United Kingdom

Website

Ject.ai

Main contact

Neil Maiden

Position in the organisation

Chief Product Officer

Other people involved in the project

Konstantinos Zachos

Kostas Petrianakis

Fotis Paraskevopoulos

Twitter profile

@ject_ai

Brief description of the project

The aim of the JECT-CLONE project was to deliver new computational creativity capabilities as a software-as-a-service (SaaS) that will autonomously generate novel themes, angles and voices for stories and send them regularly using existing channels to journalists and editors who are subscribed to the service. It sought to provide an innovative new form of smart assistance to newsrooms.

What was the challenge?

Journalism is one of the creative industries, yet journalists and editors often fail to embrace its full creative remit. Consequently, different digital creativity support tools for journalists have emerged. One, called JECT.AI, provides creative intelligence to journalists. However, the current version of JECT.AI only generates creative recommendations in response to stories already being written. Feedback from newsrooms has revealed the need for a more proactive service that autonomously generates recommendations for new stories and investigations yet to be written.

What results did you deliver (the solution)?

The JECT-CLONE project delivered a first version of the planned software-as-a-service as part of the JECT.AI offering. This service allows journalists and editors to configure their service preferences via a web-based tool, then sends emails to them each morning, depending on their configured requirements. The service launched with two algorithms to explore themes and angles for stories across languages and based on converging news stories. The service was launched to enable evaluations with pilot users, and has elicited constructive formative feedback and plaudits.

What is the impact of your project for the AI and media ecosystem?

The JECT-CLONE project has demonstrated the need for such a service – one that avoids yet another tool on the desktops of journalists. It simply integrates into their email inbox. The first evaluations with pilot users have revealed the potential of the algorithms to deliver novel insights to journalists and editors. The project also revealed the markets for greatest potential are outside of Europe, e.g., MENA and India.

Main contribution/ value from the AI4Media project/ programme

The main contribution of the AI4Media programme was the no-stake funding that was made available to recruit the specialist team needed to design and implement the JECT-CLONE service. The funding and project duration were sufficient to enable the development of the first version of the service. The AI4Media programme also provided some new connections to media networks in Europe of value to the JECT.AI start-up.

Organisation logo / project logo



What you did and achieved that would have been difficult/impossible without this funding?

In simple terms, the design and implementation of the JECT-CLONE service as planned would not have been possible without the AI4Media funding. Few alternative funding routes that met JECT.AI's expectations were available.

What's next for the project?

Development-wise, the new software service will continue to be extended and revised. Two new algorithms will be added to the service, to provide more capabilities and choices to users. The web application is being refined to respond better to emerging journalist needs and expectations. The format and styling of the emails is being refined. At the same time, we will hold a more substantial product launch in Q1 of 2024 to one of the target markets, to encourage users to sign up and use the service.

Relevant lessons learnt

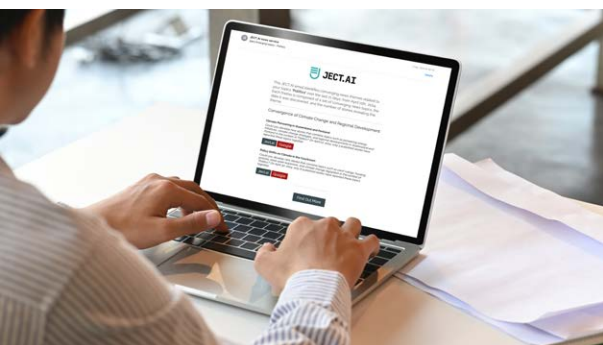
Many of the lessons have been technical – e.g., the new technical developments that will enable the JECT.AI team to stand out from its competitors. Other perhaps more important lessons relate to markets and business opportunities. The market analyses undertaken during the project revealed greater opportunities for JECT-CLONE services beyond the EU, in emerging markets with less digital legacy in newsrooms. The other major lesson was advice from business colleagues and one investor to seek income streams outside of news and media – clients more able to pay regularly for such services.

Total FSTP funding

€49,299

Brief testimonial about participation in the programme, including interaction with AI4Media and coaches, if you'd recommend such programmes to others.

Overall, participation in the AI4Media programme was very positive, in that it allowed JECT.AI to pursue one of its medium-term aims for product development without the need to dilute the effort or the product for shorter-term aims. This pursuit of the longer-term JECT.AI vision was important to the company. Therefore, if you are a media business with a strong vision to pursue using public funds, then the AI4Media programme is one that we can recommend. The AI4Media programme representatives and coaches were responsive before, during and after scheduled meetings, and were proactive in proffering advice to the JECT.AI team. Indeed, the structure of the programme and its sprints provided a sound external framework for managing the project and its target outcomes. Again, if you are seeking such structure and guidance, then we can recommend the AI4Media programme.



| Illustration of the JECT.AI software interface

magnet

Full title of the project

Automatic Recommendation of In-Context Media Content to Support Exploratory Research in Journalism

Project tagline/ slogan

let the knowledge come to you
no more scanning through your documents
or scouring the net to find information -
magknet pulls the relevant knowledge
towards you and accelerates access to
meaningful information

Project Track

Application

Funding period

1 March 2023 - 30 November 2023

Beneficiary lead name

Inknow solutions, lda

Type of organisation

SME

City & Country

Lisboa, Portugal

Website

www.inknow.pt

Main contact

Rui Neves-Silva

Position in the organisation

Partner

Other people involved in the project

Paulo Pina

Alexandre Silva

LinkedIn profile

<https://www.linkedin.com/company/9405252>

Organisation logo / project logo

The logo for 'inknow' features the word 'inknow' in a lowercase, sans-serif font. The 'i' is stylized with a red dot and a red underline that extends to the right. The 'n' has a red underline that extends to the right. The 'o' has a red underline that extends to the right. The 'w' has a red underline that extends to the right.

Brief description of the project

The magnet project developed a tool to support journalists to access the best information, particularly the ones kept in their archives, related with the current situation they might be handling.

The objective of this work is to walk towards the vision of an intelligent system that is aware of the current activity of journalists and, automatically and in a timely manner, delivers exactly the set of items they would select if they had all the time needed to explore and pick the items of interest. The selected acronym - magnet - describes this vision, that all relevant items are attracted by the specific needs of a particular journalist on a mission.

With this purpose in mind, the magnet project developed, tested, and validated an application prototype (TRL 7) to demonstrate the benefits of the approach and its implementation.

The demonstration serves to attract the attention of investors who can further support bringing the application to market as a complete and qualified product, upon completion of the project.

What was the challenge?

The main challenge of the project was to balance the aspects of usability, computational load, information about explainability and transparency of the recommended items. The project had a very clear concept from the start that was described in an initial specification document. From that concept, the project rolled as planned.

What results did you deliver (the solution)?

The magnet project developed an application that is available online (SaaS) at: www.magknet.net

A demonstration of the application can be seen at:
https://www.youtube.com/watch?v=TK8E_cSz8tY

What is the impact of your project for the AI and media ecosystem?

The magnet prototype is currently available on-line.

The magnet project makes available the developed dataset that define the fundamental topics and the context ontology, both in English.

The results of this work are included in the proceedings of the KES 2024 Conference on Human Centred Intelligent Systems, describing the approach and algorithms.

On a more conceptual level, this project contributed to alternative approaches to AI focusing on transparency and user-centric control of knowledge provision.

Main contribution/ value from the AI4Media project/ programme

The main contribution from the AI4Media project was the financial contribution allowing a small company as inknow to allocate efforts to develop the magnet prototype.

Furthermore, the periodic support with coaching sessions and reviews helped the project to be kept on track, achieving all promised results and benefiting from excellent feedback from experienced researchers.

What you did and achieved that would have been difficult/impossible without this funding?

We wouldn't say it's impossible, but for sure our company would not be able to develop this prototype, tested and validated, in 9 months without the AI4Media support. And for sure not with this level of quality.

What's next for the project?

We are currently exploring the use of the application for a broader market than journalism. Broader in diversity of possible uses but targeting niche markets. There is a clear alternative exploitation route of exploiting the results of the project within the creative media industries outside the news sector, e.g., marketing.

Relevant lessons learnt

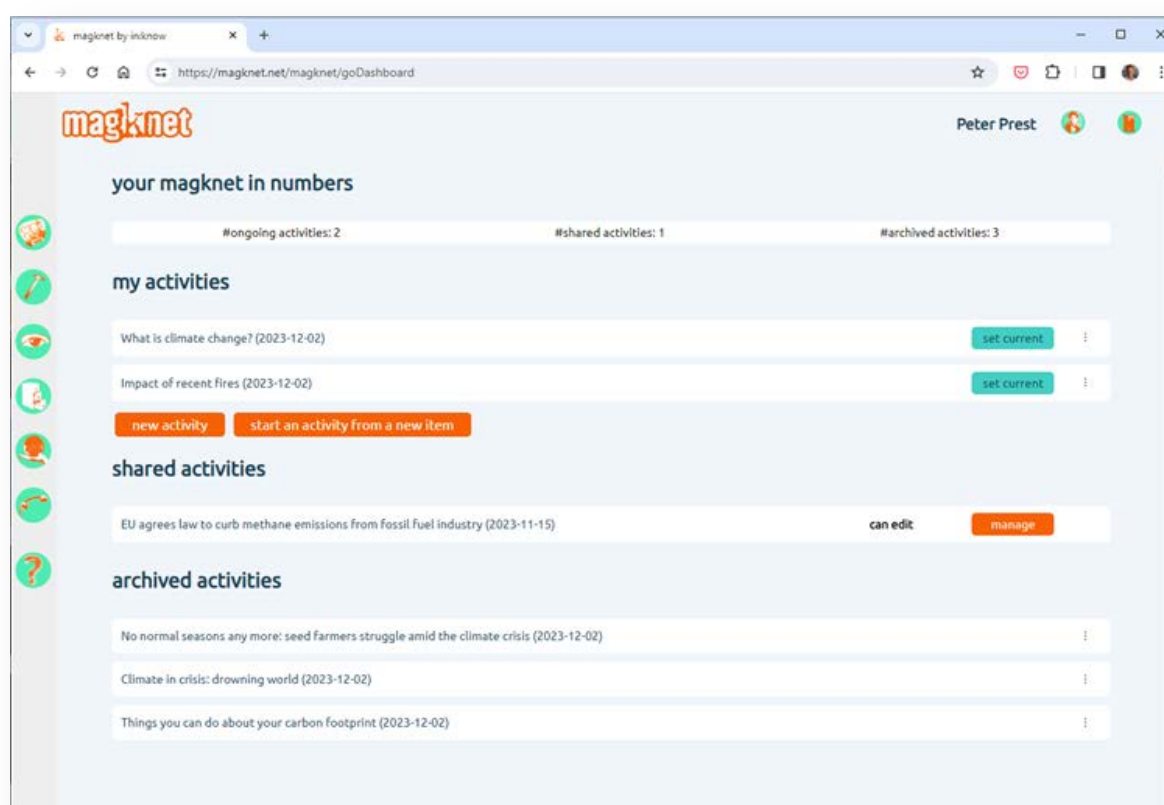
The most important lesson is to start validating the concept with the intended target group as soon as possible, even just using mock-up demos of the application.

Total FSTP funding

€49,000

Brief testimonial about participation in the programme, including interaction with AI4Media and coaches, if you'd recommend such programmes to others.

We highly recommend participation in these programmes. They are very focused in obtaining a specific result e.g. application. In the particular case of magnet, we have been supported by a great team, very professional and friendly, providing excellent feedback. Of course, for a small company like we are, the financial support provided by the AI4media project was key in putting the magnet prototype up in such a short period. We are grateful for all of that.



| Magknet application initial dashboard

MBD

Full title of the project

MindBugs Discovery

Project tagline/ slogan

The MindBugs project aims to address misinformation and disinformation issues by creating a knowledge graph that depicts the structure of misleading information.

Project Track

Application

Funding period

1 March 2023 - 30 November 2023

Beneficiary lead name

Tech Wave Development SRL

Type of organisation

SME

City & Country

Cluj-Napoca, Romania

Website

<https://discovery.mindbugs.ro/>

Main contact

Ioana Cheres

Position in the organisation

CEO

Other people involved in the project

Mihai Topor

Oana Cervinski

Ioana-Alexandra Baragan

LinkedIn profile

<https://www.linkedin.com/company/mindbugs-official/>

Organisation logo / project logo



MINDBUGS

Brief description of the project

The MindBugs project aimed to address misinformation and disinformation issues by creating a knowledge graph that depicts the structure of misleading information, providing specialists and the general public with insights into the world of deception. The pipeline for the knowledge graph involves extracting and tagging fake news data with the help of deep learning and then connecting the extracted information with existing ontologies. The pipeline integrates machine learning and symbolic artificial intelligence to provide high-quality insights into the hidden structure of misinformation.

The initiative united artists, journalists, and programmers against disinformation and equips the public with information that is currently only accessible to specialists

What was the challenge?

The challenge taken by the MindBugs project was to deploy a hybrid AI pipeline on an AIoD platform where all components are reusable. This hybrid system would utilize both machine learning and symbolic AI to extract and tag data related to fake news and then connect this information to existing ontologies. The goal is to reveal the complex structure of misinformation and make this insight accessible not only to specialists but to the general public as well.

What results did you deliver (the solution)?

The MindBugs project has delivered a suite of tools and insights to combat misinformation. Firstly, comprehensive articles detailing the structure of misleading information, as deciphered through the knowledge graph, has been published for the specialist community. Additionally, a GitHub repository has been established, providing full documentation and code, ensuring transparency and facilitating further development by the community. Secondly, the AI pipeline, which is at the core of this initiative, has been fully integrated into the AIoD platform, enhancing the European AI ecosystem by offering a tool that is both accessible and practical for a wide range of users.

What is the impact of your project for the AI and media ecosystem?

The MindBugs project has made significant contributions to the AI and media ecosystem by deploying a robust pipeline on the AI-on-Demand Platform, making its code and data publicly available for community use and research. The project's efforts have fostered a collaborative environment through an open-source GitHub repository, a custom local development orchestrator, and the publication of scientific blogs that elucidate the system's impact and data analysis processes. These efforts not only enhance the AI4Media community's capabilities but also strategically empower stakeholders to adapt and utilize these resources to combat disinformation effectively.

Main contribution/ value from the AI4Media project/ programme

The AI4Media project's main contribution for our project lies in its facilitation of collaborative work on AI technologies. By providing a space and connections for collaboration and innovation, it enabled the development and deployment of advanced, adaptable AI tools. These tools have been instrumental in empowering researchers, journalists,

and AI practitioners with robust resources to address the challenges of misinformation. The project's emphasis on open-source sharing and community engagement further enhanced our project's value, making cutting-edge AI research and tools accessible to a wider audience.

What you did and achieved that would have been difficult/impossible without this funding?

The funding was instrumental in transforming our tool into a specialist-grade resource. Direct access to AIoD developers and experts was crucial; without their support, deploying the tool would have been very unlikely. Additionally, collaboration with our monitor enhanced the tool's journalistic relevance, ensuring its utility in real-world media applications. Discussions with our coach on technical aspects allowed for sophisticated development, fine-tuning the tool to meet the complex needs of specialists. Without this funding, such comprehensive development, integration, and expert collaboration would have been difficult, if not impossible, solidifying the project's success.

What's next for the project?

The next steps for the project involve several key initiatives: validating the technology with fact-checking journalists to ensure it meets day-to-day operational needs, improving the user interface and user experience (UI/UX) to enhance usability, and commercializing the product. Furthermore, the new MindBug feature will be integrated into an AR game, broadening the project's application scope. The immediate goals include finalizing a simple and clear user interface with aggregated content from six fact-checking websites and acquiring five paying customers by the end of May 2024. This customer feedback will inform the development of a full marketing strategy and the onboarding process for new users.

Relevant lessons learnt

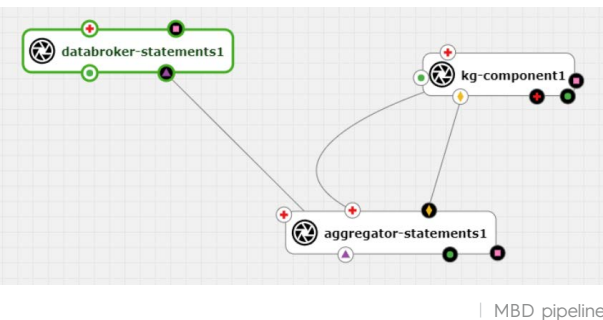
The key lesson learned was the value of cross-domain collaboration and iterative validation of technology with relevant stakeholders. Engaging with different specialists, such as journalists, AI experts, and developers, at each stage of the project was crucial. This collaborative approach ensured that the developed tools were not only technically sound but also practically useful in real-world scenarios. Working closely with stakeholders across various domains allowed for the technology to be continuously tested and refined, ensuring it met the diverse needs of its users.

Total FSTP funding

€50,000

Brief testimonial about participation in the programme, including interaction with AI4Media and coaches, if you'd recommend such programmes to others.

Participating in the AI4Media program was a rewarding experience for our team. We found the resources and support provided throughout the program extremely helpful. The wealth of knowledge and expertise within the community was a standout aspect, offering us insights, guidance and direct access to the people and resources we needed. Interaction with AI4Media and our coaches played a significant role in our project's success. They were knowledgeable and always ready to assist, which greatly enhanced our learning and development process. Admittedly, the initial phase of the program did require managing a substantial amount of reporting, which was a bit overwhelming at times. Overall, we would recommend programs like AI4Media to others. The combination of expert guidance, community support, and comprehensive resources makes it an excellent opportunity for growth and innovation, especially for those looking to make a meaningful impact in their field.



| MBD pipeline

NLMIE

Full title of the project

Natural Language Media Indexing Engine

Project tagline/ slogan

Search and find - anything - intuitively

Project Track

Application

Funding period

1 March 2023 - 30 November 2023

Beneficiary lead name

Kaspar ApS

Type of organisation

Limited company

City & Country

Denmark

Website

www.kasparai.com

Main contact

Piotr Winiwicz

Position in the organisation

Product Owner

Other people involved in the project

Mads Damsbo, CEO
Esbern Kaspersen, CTO
Sofie

LinkedIn profile

<https://www.linkedin.com/company/kasparai/>

Organisation logo / project logo

0 1 0 0
1 0 1 1
0 1 0 0
1 0 0 0

KASPAR.

Brief description of the project

A Natural Language Media Indexing Engine (NLMIE) that can analyse images and their relation to natural language text, and an Application Programming Interface (API) which will allow easy integration into the pipelines of content management systems (CMS), media production studios, and film and TV archives. We will also integrate the NLMIE into our existing Kaspar online tool.

What was the challenge?

AI global support to informative content production; Film and media archives and traditional pre-editing processes waste time with manual moment retrieval.

What results did you deliver (the solution)?

We delivered a Mac OS app called K1 that works through Adobe Premiere Pro to analyse and index both audio and visual features of local footage, and allow for fast intuitive natural language search and retrieval.

What is the impact of your project for the AI and media ecosystem?

The app is freely available for the media industry and is allowing users to deploy and take advantage of novel ML models, directly in their media production pipeline.

Main contribution/ value from the AI4Media project/ programme

The main contribution was the cash grant and guidance from industry professionals.

What you did and achieved that would have been difficult/ impossible without this funding?

It would have been difficult to have time for the necessary research and development needed for the innovation.

What's next for the project?

We will update K1 to K2 which includes more specialised ML models, and a conversational interface. Also we will expand to windows and Avid Media Composer.

Relevant lessons learnt

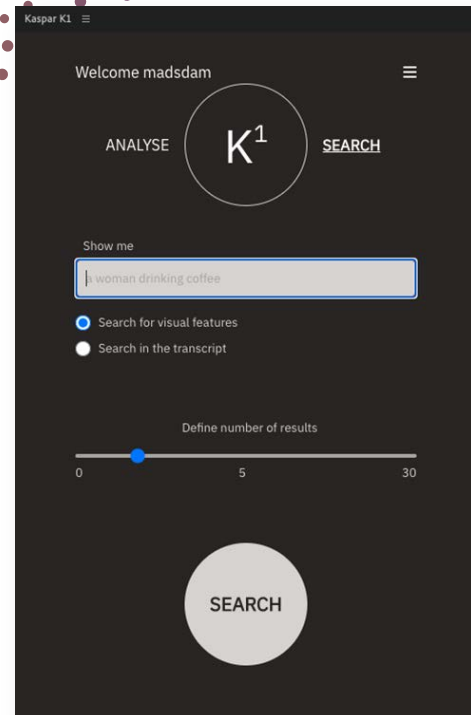
How to deploy an application locally for MacOS. How to create and utilise a vector based indexing engine.

Total FSTP funding

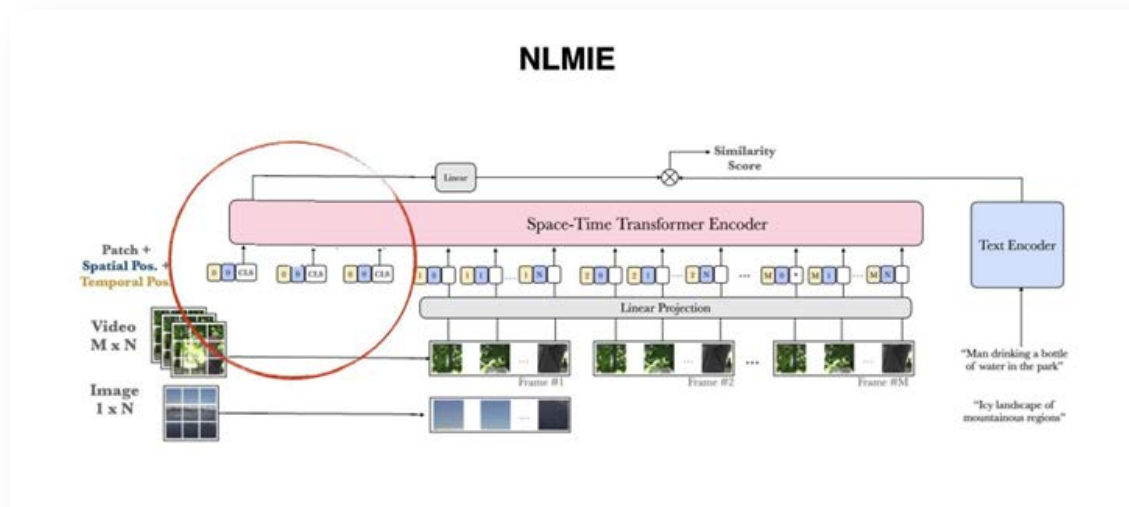
€50,000

Brief testimonial about participation in the programme, including interaction with AI4Media and coaches, if you'd recommend such programmes to others.

We were very pleased with the programme and the experts. We sometimes felt that reporting was consuming too much time, and would wish for more freedom to just work.



| Illustration of the solution interface



| NLMIE architecture

People involved in the project



VIREO

Full title of the project

Visually appealing Image
Recommendation based on Article
Content using Artificial Intelligence

Project tagline/ slogan

VIREO: Transforming News and Media
with AI Innovation - Faster, Engaging, and
Visually Appealing Stories for Journalists
and Readers

Project Track

Application

Funding period

1 March 2023 - 30 November 2023

Beneficiary lead name

Human Opsis

Type of organisation

SME

City & Country

Patras, Greece

Website

<https://humanopsis.com/>

Main contact

George E. Raptis

Position in the organisation

Director

Other people involved in the project

Christina Katsini

Filio Vogiantzi

Vasilis Theodorou

Brief description of the project

The main objective of the VIREO project was to develop a system that adopts artificial intelligence (AI) methods to analyse the textual content of an article and recommend a collection of images that would best match the article. The VIREO project aims to benefit both professionals (e.g., content writers, journalists) and readers, both stakeholders of the News and Media ecosystem: the professionals quickly select visually appealing images to create engaging and captivating stories, and the readers are provided with enhanced reading experience, increased engagement, and better recall.

What was the challenge?

The challenge that VIREO aimed to address was to develop image recommendations based on text to create visually appealing stories and eventually save time for professionals in the news and media industry (e.g., journalists). This challenge aligns with the "C4-A: AI for suggesting visually appealing images based on text" challenge identified in the AI4Media project for AI use cases for news and AI for (re-) organisation and content moderation.

What results did you deliver (the solution)?

VIREO delivered an integrated digital solution comprising three key components:

- Article Processing Tool for preprocessing and summarising articles, making natural language computer-readable and understandable. This tool involves advanced algorithms for vectorization, natural language understanding, semantic analysis, and syntactic and etymological similarity.
- Image Processing Tool for image preprocessing, translating images into vectors. The tool utilises image captioning algorithms to generate textual descriptions, image tags, and captions for semantic analysis.
- Recommendation Engine that matches and suggests a set of images based on a text, employing AI to save time, enhance storytelling, and improve the experience for journalists and readers.

What is the impact of your project for the AI and media ecosystem?

VIREO impacts the AI and media ecosystem by addressing the need for efficient news content creation in the evolving media landscape. By integrating AI techniques, VIREO enables journalists to select visually appealing images based on article content, saving time and enhancing engagement for readers. It also contributes to AI4Media's goals by advancing applied research in media AI, fostering industry-academia interactions, and enriching the AI4EU platform. VIREO's deployment benefits the media industry by saving time and boosting productivity. It aligns with European values of ethics and trustworthiness in AI deployments, positioning AI as a crucial enabler for societal and media services.

Organisation logo / project logo





Main contribution/ value from the AI4Media project/ programme

AI4Media made significant contributions to VIREO's success through:

- Financial support, which was vital for VIREO's progress and achieving its milestones;
- Expert (coach) guidance and support, providing strategic advice, industry insights, and feedback, enhancing the project's direction and ensuring alignment with industry needs;
- Interdisciplinary collaboration supporting interactions and expanding networks between diverse stakeholders within AI and News and Media ecosystems;
- Access to the AI-on-Demand Platform and related communities, enabling us to share and gain insights of resources related to AI-based projects and knowledge applied to the News and Media sector.

What you did and achieved that would have been difficult/ impossible without this funding?

The financial support from AI4Media was essential in achieving critical milestones for VIREO. It enabled developing and integrating sophisticated AI algorithms for article processing, image vectorization, and real-time recommendation. The funding enabled in-depth research and implementation of innovative mechanisms for text understanding, image caption embeddings, and accurate matching processes. With this funding, AI4Media made the VIREO approach for developing image recommendations based on text to create visually appealing stories and save time for professionals in the news and media industry easier to deploy. Moreover, the support allowed for interdisciplinary collaboration, which might have been difficult to achieve without dedicated funding.

What's next for the project?

The following steps for the VIREO project involve advancing its technology components to achieve TRL7+, demonstrating operational viability in a real-world environment. Activities include: i) refining algorithms and models for article processing, image vectorization, and recommendation engine integration; ii) incorporating VIREO in the work pipeline of News and Media agencies, such as integrating it in popular content management systems, (e.g. WordPress); iii) performing real-world, large-scale user studies to validate and finetune the workflow of VIREO; iv) moving towards exploitation and commercialization.

Relevant lessons learnt

From a technical perspective, the importance of robust natural language understanding in article processing was emphasised, effective image recommendation depends on accurate image caption embeddings, and seamless integration of AI techniques is essential for a fast and accurate recommendation engine. From a more project-oriented perspective, AI4Media helped us improve VIREO's capabilities through a collaborative environment and ethical AI practices, supporting and enhancing the project's visibility and communication.

Total FSTP funding

€49,675

Brief testimonial about participation in the programme, including interaction with AI4Media and coaches, if you'd recommend such programmes to others.

It was a great pleasure working with the AI4Media team and experts during the implementation of the VIREO project. We want to thank our three coaches from VRT. Their guidance and support were invaluable; we were lucky to have them on this journey. Their feedback and insights helped us align the VIREO solution directly with the industry's needs, perform meaningful evaluations, and shape this project's future directions. We also want to thank the F6S team for their ongoing support and to-the-point comments during the project implementation. They were there when we needed them, helped us address our concerns/challenges, and did it quickly. AI4Media was our first experience with FSTP projects; it was highly positive, and we recommend such programs to others.



**People involved
in the project**



VolEvol

Full title of the project

Generation of Meaningful Representations of Volume Data Through Evolutionary Learning

Project tagline/ slogan

Development of techniques for the automatic generation of sets of quality and diversified images from 3D volume data sets.

Project Track

Research

Funding period

1 March 2023 - 28 February 2024

Beneficiary lead name

"Gheorghe Asachi" Technical University of Iasi, Romania

Type of organisation

University / Academic Institution

City & Country

Iasi, Romania

Website

<https://www.tuiasi.ro/?lang=en>

Main contact

Marius Gavrilescu
marius.gavrilescu@academic.tuiasi.ro

Position in the organisation

Lecturer, PhD

Other people involved in the project

Sabina-Adriana Floria
Florin Leon
Lavinia-Eugenia Ferariu
Cristian Nicolae Buțincu

Brief description of the project

We developed techniques for generating collections of high-quality, diverse images from volume data. We combined volume rendering with evolutionary optimization in order to explore the space of viewer properties and color/opacity mappings that result in visual representations of the volume. To this end, we identified multiple quality and diversity criteria based on combinations of data, surface and geometry-based features of the volume structures, and established optimization objectives and a diversity space based on these criteria. We used a MAP-Elites evolutionary method to generate rendering parameters that result in sets of illustrative, diverse images of the data.

What was the challenge?

We addressed the challenge C2-R: *Evolutionary learning and the challenge of evaluating quality*. Briefly, this involved using an evolutionary algorithm to search for multiple locally-optimal solutions in a feature space whose dimensions account for diversity. In our case, we searched for sets of parameters which, when supplied to a volume rendering method, would result in sets of quality representations of 3D data. The most challenging aspect has been formulating properties of the data that could be objectively-measurable, which would allow the automatic generation of images of perceptual and subjective quality.

What results did you deliver (the solution)?

We developed a technique that combines a volume renderer, an evolutionary optimizer and multiple feature extraction methods to automatically generate sets of quality and diverse images from volume data sets. We developed a software prototype that specifically implements this technique, which is currently freely-available under an open-source licence. We also extensively documented our work in great detail in the multiple reports and deliverables developed throughout the project. We mainly disseminated our results through conference papers, presentations, demo videos and a user study.

What is the impact of your project for the AI and media ecosystem?

Our original aim has been to facilitate the visual exploration of volume data in a manner that is accessible outside of specific technical/research fields. Volume visualisation is a highly-complex field of information visualisation and data analytics, with multiple applications in medicine, various industries, media, game development and art. VolEvol contributes to automating volume visualisation using a learning approach that allows for the fast, effective generation of relevant representations of volume data sets. We address a broad range of users by developing volume visualisation techniques that are not domain specific and don't require much pre-existing technical knowledge of the data.

Organisation logo / project logo





Main contribution/ value from the AI4Media project/ programme

Conducting our work within the framework of AI4Media has been beneficial from multiple perspectives: the funding constituted a motivating factor and provided support for disseminating our work through scientific publications and conference presentations; being part of AI4Media facilitated collaboration and knowledge exchange among the parties involved, enhancing the potential for innovation and meaningful contributions. Additionally, the association with AI4Media lends credibility and visibility to the research, facilitating dissemination of findings and potential adoption by potentially-interested media representatives/industrial partners. Furthermore, this collaboration enabled us to tackle a complex challenge and contribute meaningfully to the corresponding research field.

What you did and achieved that would have been difficult/ impossible without this funding?

It is unlikely that our team would have tackled the project topic entirely on its own, without a call for proposals and an actual project as a supporting framework. With VolEvol we achieved automated generation of high-quality, diverse images from volume data sets through the fusion of volume rendering and evolutionary optimization. This innovative approach streamlined the exploration process, enabling users across various domains to efficiently represent intricate details within volume structures. Without this research, manually navigating the complexities of volume data visualisation would have been laborious and time-consuming, hindering broader accessibility.

What's next for the project?

We intend to further improve and extend the techniques and software already developed within VolEvol, in directions such as: refinement of the optimization techniques used within our pipeline; integration of machine learning techniques as complementary optimizers; incorporating user interaction and feedback into the image generation pipeline; extending our methods to domains other than volume visualisation; submission of further publications and project proposals in the directions already pursued within VolEvol. In this context, we welcome future collaborations with our AI4Media contacts, as well as with the other potentially-interested sub-grantees.

Relevant lessons learnt

The lessons learned from the VolEvol project involve the importance of interdisciplinary collaboration and the value of automation in complex data visualisation tasks. The project highlighted the efficacy of combining evolutionary optimization with volume rendering to streamline the image generation process. Furthermore, we emphasised the importance of user-centric design, which proved critical in ensuring the accessibility and usability of our results across data from diverse domains. Additionally, the project highlighted the significance of scalability in handling large, complex volume data sets effectively. These aspects demonstrated the need for continued innovation and refinement in leveraging AI-driven approaches for advancing visualisation techniques.

Total FSTP funding


€50,000

Brief testimonial about participation in the programme, including interaction with AI4Media and coaches, if you'd recommend such programmes to others.

Participating in the AI4Media programme was a positive experience that significantly contributed to improving our expertise in the related research directions. The opportunity to collaborate within AI4Media provided valuable insights and improved our capability for innovative thinking. The AI4Media staff were highly-professional in all interactions. The guidance and support from our coach throughout the programme was essential for navigating the various challenges and maximising the impact of our project. Consequently, we fully recommend such programme to others seeking to develop innovative, multidisciplinary research. The collaborative environment, fast effective feedback and professionalism demonstrated by the AI4Media team offered a significant opportunity for growth and learning. Engaging with AI4Media not only enhanced our project's outcomes but also expanded our network and opened doors to future collaborations. Our participation in the programme was highly rewarding, and we look forward to applying the knowledge and skills gained to future developments.



Examples of images generated automatically by our software prototype using volume data sets from climatology and medical imaging



Our Consortium



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