AviaTor

Save Time, Save Lives

Annual Report **2023**

AVIATOR PROJECT. COM

COLOPHON

AviaTor

AviaTor Annual Report
INHOPE (Co-office)
Bos- en Lommerplein 280

1055 RW Amsterdam
The Netherlands
aviatorproject.com

DEAR READERS,

2023 was a year of major developments in the fight against online sexual abuse of children. The EU proposed CSAM Regulation, the ideas about the EU Centre and the debates around privacy rights versus the protection of children were very much at the forefront of wide discussions within and outside the EU Member States.

The EU proposal (as it's known at the time of writing this report) is promising, although perhaps less extensive than the European Commission originally envisioned. It does highlight that the protection of children against online sexual abuse is a fight that constantly calls for attention and needs to be in the hearts and minds of all seeking a safer internet for children (and adults).

It also calls for involvement from all EU Member States. The EU-wide resilience against those abusing our children, online and offline, must be improved. Having the right tooling, knowledge, funding and people to do so is essential. AviaTor is one of these tools, and the large number of affiliates that have signed up for the project is proof that there is a need for these tools. In the future, the yet-to-be-established EU Centre could play a part in the further development, introduction and testing of such tools, as well as making them available for Member States law enforcement and other authorities. This would significantly improve the previously mentioned resilience and prevent criminals from hiding in countries that are still developing their capacity.

The EU proposed CSAM Regulation and the Digital Services Act can also help to make industry even more aware of their essential role in this fight. The distribution of known abusive material can and must be combated more effectively by companies, in which they should be expected to do this and be accountable.

Many comments on the new EU regulation discuss the role of preventing both child abuse and the distribution of child abuse material. It is important to continue to invest in this area and support initiatives. Knowledge and information can help children make decisions when interacting with others online. As we all know, there are risks associated with the use of the Internet and not everything is as innocent as it may initially seem.

For AviaTor, we look forward to the coming period of further development. Together with Europol and all our partners, we are trying to clarify how to optimise the procedure of receiving and processing CSAM reports: what the de-confliction of reports will mean and how we can increase that information flow to the Member States. AviaTor can play an important role in this, and it is therefore important that prolonging the EU funding for the next period of development is considered. This may strengthen the capacity of affiliate Member States significantly and prevent them from being overwhelmed with referrals and reports.

Phase 2 of the AviaTor project will end on September 1st 2024, and the project so far has shown that we can and have built a tool that helps our partners make the right decisions, prioritise based on scientific knowledge and pay attention to the health and wellness of LEAs. The Dutch Police has the project lead on the AviaTor project. Together with the Belgian Federal Police and a consortium of partners, we have the privilege of incorporating our wishes and requirements into the project as much as possible, resulting in an almost tailor-made product. It has proven to be successful, and we would like to thank the EU Commission for their trust, involvement and funding of the project. I think there has been excellent cooperation that can grow even further into an important instrument for the Member States. That growth - and therefore the support - remains important to continue to cope with the rapid developments in this field, such as AI applications. I have no doubt that we can and must tackle this together with all our partners, because only together will we be able to make a difference!

Ben van Mierlo

National coordinator for the fight against
Child Abuse Images and Transnational Child
Sex Offences Netherlands Police

Contents

Before We Start

Project AviaTor

Listening to AviaTor Users

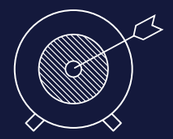

Obstacles in Becoming an AviaTor User LEA

Legal Checklist LAW

Federated Learning TECHNOLOGY

The Way Forward

Meet the Partners

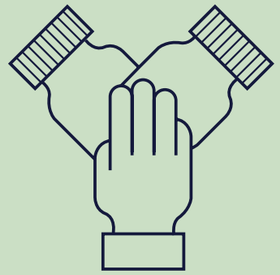

CHAPTER 01

Before We Start

Building innovative technological solutions to fight child sexual abuse and exploitation material

BEFORE WE START

Introduction

With the help of the European Union's Internal Security Fund - Police (ISFP) and a group of dedicated industry specialists and law enforcement officers, phase 1 of project AviaTor was successfully launched in 2019. Phase 2 started in 2021 and will end in 2024. The AviaTor project provides law enforcement agencies (LEAs) with innovative technological solutions in their fight against child sexual abuse and exploitation material (CSAM/CSEM).

The AviaTor tool processes and prioritises reports received from NCMEC, which are industry-submitted reports on suspected CSAM, also known as Cybertips. Project AviaTor was developed to support LEAs by improving efficiency in the handling of these reports. How AviaTor technology achieves this is through the application of visual intelligence, targeted online research, cross-matching and hash matching. Using artificial intelligence (AI), AviaTor provides LEAs with automation tools that help them prioritise, assess, and process reports, streamlining the pre-investigative process. With the aid of open-source intelligence (OSINT), AviaTor incorporates both AI-assisted categorisation and automated online research to enrich reports with risk profiles, such as identifying the occupation of reported offenders.

In 2023, the global number of incoming NCMEC reports has increased by 12% compared to the previous year, surpassing 36.2 million reports. The number of files included in the reports also increased by 19% from the previous year to more than 100 million[1] . These figures demonstrate how the production and distribution of CSAM remains a critical issue that requires global attention. It is indisputable that technology and its exponential growth have perpetuated the ease of access to, spreading and production of CSAM online, introducing more threats to children. To effectively mitigate CSAM and its impact, the need for improvements in report processing efficiency is more pertinent than ever. By leveraging advanced technological tools, there are several doors for opportunity. For example, automation allows the processing of reports at an increased pace, while also minimising the need for human analysts to review the reported illegal material.

The project, and its success, are largely attributed to AviaTor's community of committed experts and law enforcement officers who frequently exchange "know-how" and best practices. To better enhance AviaTor and its functionalities, a continuous stream of knowledge-sharing and communication must flow regarding progress and areas for opportunity. AviaTor fosters this exchange by hosting biannual Peer-to-Peer meetings and networking events, creating a platform to promote collaboration.

This is the third and final instalment of the annual reports published between 2022-2024. The annual reports serve to inform key stakeholders and the public about the development of the AviaTor project, relative to the objectives and outcomes of each phase. While this report allows us to look back at the developments and results in 2023, as we step into the final year of Phase 2, we aim to look forward.

AviaTor's future and sustainability have become our main focus, with our efforts directed towards maintaining the tool for current and future users. Highlighted are the efforts and future objectives to ensure the platform's sustainability in the future. This report will also cover key insights and trends using the data and statistics retrieved from AviaTor users, as well as focus on newly implemented features. Through the chapters, different key actors of the project will share their insights into AviaTor's progress and ways that we can better approach the processing of industry CSAM reports.

As we present this annual report, we celebrate the final steps of a project which has opened up the possibility for LEAs all over the world to significantly optimise their fight against the rapidly increasing CSAM production and distribution. We are grateful for the European Commission's support whose funding has been instrumental in the development and success of this tool. We are also thankful to all stakeholders, including affiliated law enforcement officers, whose expertise and knowledge have made AviaTor highly effective and accessible for LEAs within and outside the EU.

[1] https://www.missingkids.org/gethelpnow/Cybertipline/Cybertiplinedata

GOOD TO KNOW

Acronyms & Abbreviations

- **AI** - Artificial Intelligence
- **API** - Application Programming Interface
- **AVIATOR** - Augmented Visual Intelligence and Targeted Online Research
- **CSAM** – Child Sexual Abuse Material
- **CSEM** - Child Sexual Exploitation Material
- **CSE** - Child Sexual Exploitation
- **DFKI** - German Research Centre for Artificial Intelligence
- **DPIA** - Data Protection Impact Assessment
- **ESP** - Electronic Service Provider
- **EU** - European Union
- **EC** - European Commission
- **ESCO** - European Skills, Competencies, Qualifications and Occupations
- **GDPR** - General Data Protection Regulation
- **ICSE** - International Child Sexual Exploitation database
- **ISF** - Internal Security Fund (European Union)
- **ISFP** - Internal Security Fund - Police (European Union)
- **LEA** - Law Enforcement Agency
- **LFE** - Large File Exchange
- **NCMEC** - The National Centre for Missing and Exploited Children
- **NCMEC Reports** - reports on suspected CSAM coming from (mostly) US-based electronic service providers. Also called: Cybertips, NCMEC reports, Cybertipline reports, industry reports, and industry referrals
- **NLP** - Natural Language Processing
- **NPN** - National Police of the Netherlands
- **OSINT** - Open-Source Intelligence

CHAPTER 02

Project AviaTor

A look at the results of
building innovative new technology

OVERVIEW

Reflections on the AviaTor project

A lot has happened since the initial 2017 brainstorming session in INHOPE's Amsterdam boardroom. Ideas were formed between the Dutch and Belgian police on one side, and Web-IQ and ZiuZ Visual Intelligence on the other side, to build a system to process and prioritise NCMEC reports. We branded this system AviaTor, which stands for Augmented Visual Intelligence and Targeted Online Research. This expresses that by applying visual intelligence to the images and videos in a report and carrying out targeted online research on the reported person, the report is augmented with additional intelligence for the investigator to make an informed decision about the priority and processing of a report.

We executed the AviaTor phase 1 project that was funded by the European Commission under the ISFP program, which started in early 2019 and finished in 2021.

This project resulted in a tool that helps LEAs to process NCMEC reports more efficiently. Eleven LEAs joined the project and tested AviaTor, and many LEAs were interested in doing the same. Before AviaTor became available for use, most agencies were processing NCMEC reports manually.

2021 was not the end of the AviaTor project. Strengthened by the success of phase 1, we wrote a proposal to the European Commission for Phase 2 of the AviaTor project. This proposal had several objectives:

- Ensure that AviaTor is functionally complete and sustainable
- Develop advanced AI for text analysis and video analysis
- Drive collaboration among LEAs, Europol, INTERPOL and industry
- Onboard at least 25 LEAs using AviaTor
- Publish a yearly report
- Conduct a legal review of EU Legislation and policies influencing the project

This AviaTor Phase 2 project also received funding from the European Commission under the ISFP program.

With Phase 2 coming to a close in the last quarter of 2024, this third and final Annual Report is a good opportunity to reflect on the results of the project, the lessons learned and the sustainability of AviaTor.

Results

15 AviaTor version releases

When we started the project in Phase 1, we decided to release a new version with incremental functionality every three months. After every new version release, we organise online demonstrations and user group meetings to introduce the new functionalities to the users, providing opportunities for feedback and requests for additional functionalities. We did this to keep all users in a close loop. We kept this rhythm up until the last year of the project, with a planned 15 AviaTor releases over the two phases of the project. In the final year, we shifted the emphasis from building new functionalities to improving the reliability and stability of the system. This was necessary because of the frequent changes to the reporting layout and the way they are distributed by NCMEC and Europol.

Since AviaTor is installed as a stand-alone and the system is not connected between agencies or the developers, each update has to be installed separately for each user. This translates to 240 installations carried out by the ZiuZ Visual Intelligence Support department over time, to ensure each of the 16 LEA users had access to the newest version release of AviaTor.

20.000 Images annotated

By the end of the project, an estimated 20.000 images have been annotated by the Dutch and Belgian Hotlines according to key properties of the Universal Classification Schema .

These annotated images will be used to train a granular CSAM classifier for AviaTor to detect CSAM in new and unseen images and videos. This classifier can also be used outside of AviaTor by other LEAs, hotlines and potentially provider systems.

Text classifier

The goal of the AviaTor Text Classifier is to uncover keyword-based indicators of cyber-grooming and real-life interactions between offender and victim. Two features were developed in year 3:

The Text Flagger detects chats in which potential grooming language is used. The initial version was released in mid-2023 and supports chat logs that are included in one of the report's XML fields in both Dutch and English.

The Job Title Risk classifier flags jobs that potentially involve real-life access to children. In 2023 the second, multi-lingual version was developed, which will be integrated into a later AviaTor release. It uses a language embedding model to embed the job title found in the online profile of the reported offender. A classifier was trained on these embeddings, using hand-labelled job titles from the ESCO jobs and skills taxonomy dataset as training data.

Targeted Online Research

Targeted Online Research enriches reports with pieces of information from online sources such as the occupation listed on the reported person's profile, as mentioned above. These online checks have been updated several times in 2023 due to changes on the platforms. A lot of time was invested in assessing the non-technical feasibility of automated online checks, including legal and other challenges.

As a result of this assessment, users of the AviaTor system in countries which restrict automated online checks now have access to on-demand Targeted Online Research. The on-demand version allows for the same checks to be performed manually via a separate user interface. This development started in 2023 and is ongoing.

A benefit of the stand-alone version is that new OSINT checks are made available to users sooner, who can then request for certain checks to (also) be integrated into AviaTor, following the legal guidance from AviaTor partner Timelex.

Standardised video workflow

Approximately 50% of the media files included in NCMEC Referrals are videos. Unfortunately, the processing of a video takes more human and computer resources than an image does. Therefore, an effective and efficient video workflow became a primary focus in 2023.

In this workflow, a video file is first matched on a file level, which is a relatively 'cheap' operation. If a file is known, it can be classified automatically based on the prior classification. If a file is not known, the video is dissected into shots and PDQ hashes are calculated for the shots. These shots are matched with the shots already known in the system. Known shots can be classified automatically based on the prior classification. For unknown shots, a summary of the shot will be created for a quick review by the investigator. Selected frames of the shot will be classified by the CSAM classifier to determine the likelihood that the shot contains CSAM.

Face detection and grouping

The original planning of Phase 2 of the AviaTor project included the detection of faces in images and videos, and the grouping of similar faces. Although this functionality was realized in the project, the functionality has not been integrated into the system due to ongoing discussions about the lawful use within law enforcement agencies.

Test deep fake detection

The original planning of Phase 2 also included a test with deep fake detection. However, since the start of the project, the quality of AI-generated imagery has taken a huge stride forward. We concluded that doing a limited test in the AviaTor project will not contribute to the solution of a problem that requires a thorough scientific approach and continuous improvement in detection technology.

Connection with GRACE project

The GRACE project, funded under the Horizon 2020 Research program, ran in parallel with the AviaTor project and had similar goals. Our intention was to implement a connection with the GRACE system running at Europol to enable peer-to-peer communication between AviaTor users and to feedback results from AviaTor to Europol. Since GRACE as a system did not materialise during the project period, we could not implement this connection. GRACE instead became a specialised box with almost 40 tools, that can be integrated in existing LEA systems.

Collaboration with Europol

The methods by which EU national LEAs receive industry CSAM reports varies. Some LEAs receive reports directly from NCMEC, while an increasing number of EU Member States receive NCMEC reports through Europol. To optimise the latter process, Europol has taken several steps in the recent years to improve their report distribution methods. During Phase 2 of the AviaTor project, Europol started implementing additional functionalities in their processing. This functionality is implemented before they are transferred through the Large File Exchange (LFE) to LEAs that get their Reports through Europol.

The intended functionality is threefold:

- Cross reference NCMEC referrals with other databases. Any hit information is added to the report and forwarded by Europol to a LEA.
- Calculate a report priority based on the reported content that is sent together with the report to a LEA.
- Create an API for LEA systems to automate downloading of reports from Europol by LEA.

Europol and the AviaTor project team have set up regular meetings to make sure that LEAs using AviaTor can benefit from the new Europol functionality. This resulted in the following AviaTor functionalities:

- Data enrichments by Europol are visualised the AviaTor user interface and can be easily reviewed by the investigator handling the report.
- Any indicators that may influence priority can be included in the score calculation in AviaTor. This gives the investigator the option to either combine the scores, set up their own scoring rules in AviaTor score or copy the Europol value.
- The API is still under development on the Europol side, but we are committed to implementing the communication through the API as soon as it becomes available. This will take away the burden of manually downloading the reports through the large file exchange, thus saving the investigator considerable time.

Capacity building and dissemination

With INHOPE in the lead, we set ambitious goals to:

- Increase the reach of AviaTor to a minimum of 25 LEAs
- Promote collaboration amongst LEAs and with industry

We reached these goals by organising biannual peer-to-peer learning sessions designed for LEAs using AviaTor to exchange best practices and learn from each other. There were lively discussions on several topics including the use of AI and OSINT, prioritising NCMEC reports and hacked Facebook accounts, as well as the drafting of DPIAs to get permission to use AviaTor.

We held a networking event (AviaTor Seminar) halfway through the project, which focused on exchanging information with other stakeholders on how to prioritise reports, and there will be a closing event (AviaTor Forum) where law enforcement, industry, hotlines and other key stakeholders will come together to talk about the future of the workflow of law enforcement in processing CSAM reports, with a specific focus on innovation, legislation and cooperation.

Our efforts have led to all European LEAs having access to AviaTor and they will be supported in the implementation of the tool in their agency. Our goal is for AviaTor to become the de facto standard for processing NCMEC reports, and around AviaTor, a community of police investigators will develop that promotes best practices and peer-to-peer learning. This community provides support to LEAs implementing a workflow for processing industry reports and joining the network of AviaTor users. This network will have the capacity to handle the ever-growing number of reports and potentially additional reports resulting from future EU legislation involving mandatory CSAM reporting by service providers.

Support the EU Strategy to fight child abuse

We support the EU Strategy to fight child abuse in several ways. We strengthen law enforcement by providing, at no initial cost, a system to efficiently and effectively process NCMEC referrals. With AviaTor, LEAs will create a workflow that can handle the influx of reports adequately, even if the proposed EU CSAM Regulation comes into force and the EU Centre is established. In addition, industry efforts will be galvanized by making specific AviaTor components, such as the AI classifier and matching technology, available to prevent the creation and circulation of CSAM on their platforms. Finally, we actively promoted best practices by organising half-yearly partnering meetings for all LEAs using AviaTor to exchange knowledge and experience in the processing of NCMEC reports.

In conclusion

We have achieved the majority of the goals we defined at the start of the project. While the connection with the Grace project did not turn out to be feasible, we replaced this functionality by focusing on the best possible integration with the new Europol functionality through their Large File Exchange (LFE). LEAs that receive CSAM reports through Europol will directly benefit from this. Some of the AviaTor users who receive their reports directly from NCMEC have indicated interest in this functionality too. The direct connection to the Europol LFE, through an API, that will be realised in AviaTor later this year will bring additional time savings to the users. This functionality will likely overcome the historical latency in report delivery from NCMEC to Europol and expedite the investigation of high-priority reports.

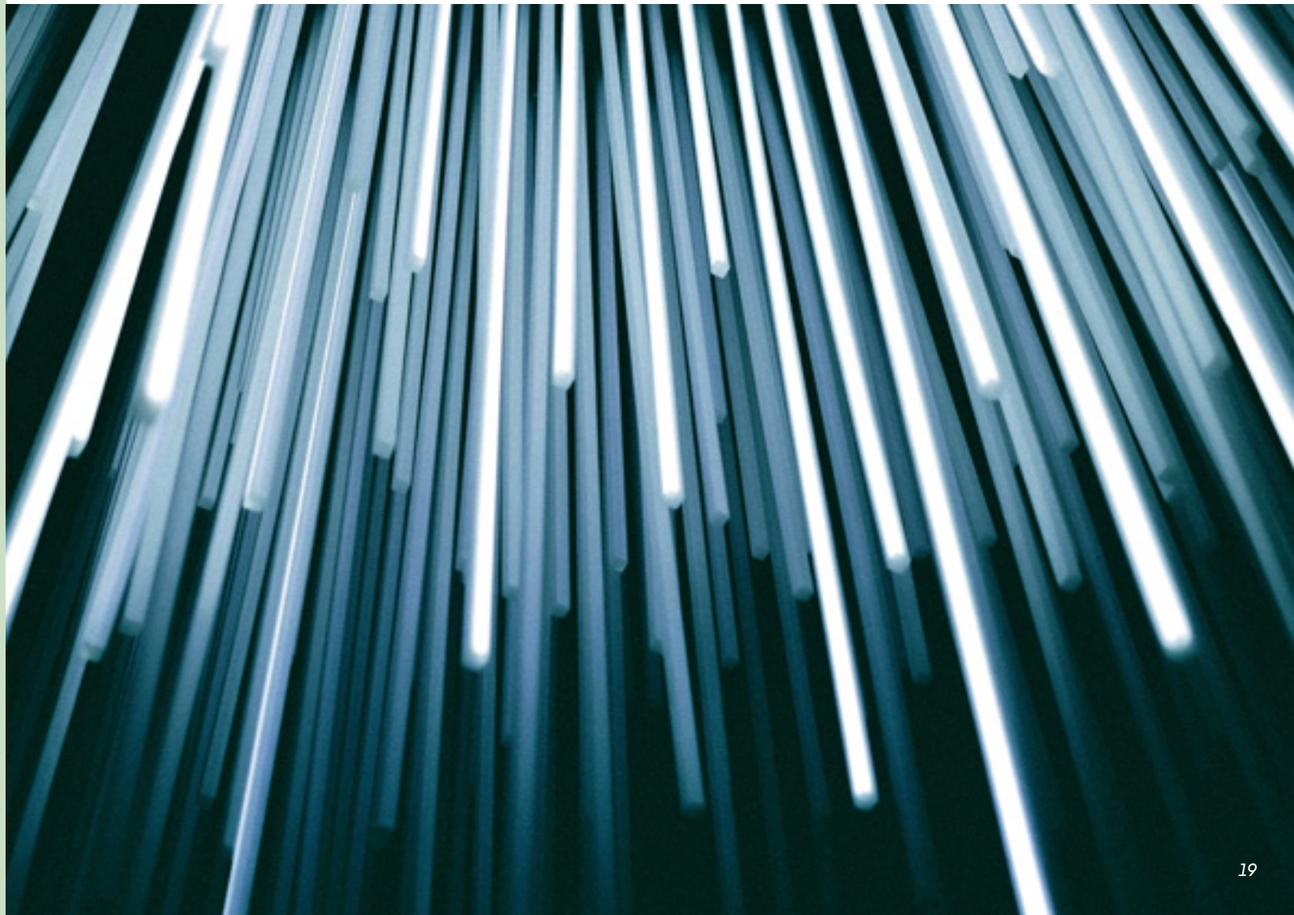

DEVELOPMENT

Technological progress and key developments

Technological progress

Even though the AviaTor tool is designed to process and sort NCMEC reports containing CSAM, it is crucial to know that these reports are not accessible to the AviaTor developers. Only law enforcement agencies (and most hotlines) are permitted access to these reports, which contain potentially illegal content. The feedback from LEAs who test and use the AviaTor tool is therefore particularly important to the development team. They are the only ones who can test the software following its intended use.

The fact that AviaTor is installed as a standalone application is another crucial detail. Every LEA user has a specially configured version, as opposed to a single, universal version. They have no direct communication with the development team and are not connected.

The development team requests that the AviaTor user group submit regular feedback and requests for software improvements. The development team's main priorities include completing all modification requests that law enforcement submits to them and providing the major advancements detailed below.

Key developments

During the last year of AviaTor's phase 2 (Sep 2023 – Sep 2024), the development team focused on the following key developments.

1 Rollout of the AI classifier for text analysis

2 More advanced targeted online research

3 Creating a more granular CSAM classifier and applying this classifier on video

4 Creating face detection and grouping

5 Making AviaTor functionality complete and implementing new requirements

6 Integrating with Europol EU Cares system

Rollout of the AI classifier for text analysis

The new Text Flagger plugin facilitates the detection of coercion, grooming or threats in reported text messages within the NCMEC report. It does this by flagging any keywords that indicate these risks. The keywords for Dutch and English were provided by the Dutch National Police. The plugin assigns a score to a report based on how many unique risk keywords it contained.

Further work on the text classifier

↗ Support for other languages by having participating LEAs submit keyword lists in other languages and integrating these into the application.

↗ Support for scanning report attachments.

↗ More advanced targeted online research

The OSINT plugin finds information online indicating possible risks to children posed by a report. One way it does this is by classifying any job title found for the subject into high or minimal risk for real-life access to children.

The first version of this job title classifier used a resource list of high-risk Dutch and English job titles and comparing them to the job titles found online. This system had several disadvantages:

↗ Only an exact string match was considered a positive classification.

↗ The job title list needed to be curated by hand.

↗ There were no support for other languages.

To improve this, we developed a second version that uses a language embedding model to embed the job title found online. We then trained a classifier on these embeddings, using hand-labelled job titles from the ESCO jobs and skills taxonomy [1] dataset as training data.

Because this dataset is multi-lingual, we were able to train a multi-lingual model with support for the following languages: Arabic, German, Greek, English, French, Hungarian, Dutch, and Portuguese.

We compared this model to the old, string-match model, using an ESCO-based dataset[5] of 23.000 job title samples (21.000 low-risk and 1400 high-risk), using 80% for training and 20% for testing. The result showed an increase in F1 score from 0.49 (strict keyword match) to 0.89 (DistilBERT embedding model[6]). Even when using a more lenient keyword match (the keyword substring appearing anywhere in the job title string), the keyword match F1 score was still only 0.65. Therefore, the new job title classifier is a clear improvement.

The F1 scores varied across the languages from 0.96 (AR) to 0.85 (HU) but were in all cases significantly higher than for the best keyword matcher (strict or lenient).

Creating a more granular CSAM classifier and applying this classifier on video

In the past year, an open-source annotation tool was deployed at Offlimits in Amsterdam and at the Belgian Federal Police in Brussels. Until now, about 8000 images were annotated using an annotation schema based on a draft version of the Universal Classification Schema. The currently available annotations were statistically analysed, resulting in the insight that the source data is statistically unbalanced concerning (combinations of) attributes, such as ethnicity. Initial experiments on classifier training were conducted, which will be further pursued in the coming months. The application of the AI classifier to video frames has been scheduled to take place after concluding the main phase of the classifier training.

Creating face detection and grouping

A prototype tool to search for and filter on similar faces in images was developed and validated on a real case of about 2 million images by a European LEA partner. Further development of this functionality is not foreseen in the near future due to multiple reason including the AI image classifier and its application to videos receiving higher priority.

Making AviaTor functionality complete and implementing new requirements

The focus for this year was to stabilise the AviaTor ecosystem and create a reliable platform for triaging NCMEC reports.

Since many AviaTor countries receive referrals via Europol, integration with that process was also a main focus point in 2023. We had mutually beneficial discussions of expertise with Europol IT specialists, through virtual and in-person meetings, which resulted in Europol installing a version of AviaTor for testing purposes. A number of features enabling process integration have already been implemented in AviaTor and work on further compatibility is ongoing.

Additional AviaTor updates implemented in 2023 include:

- Local policy settings can trigger data anonymisation after a configurable retention time
- Expanded media details overview with explicit scoring explanation
- Visualisation of the results from visual analysis plugins
- Visualisation of the results from report analysis plugins
- Expanded cross-matching
- Better logging and explaining of failed imports
- Processing of multiple changes in cybertip layout
- More user management functionality
- Changes to the configuration are now validated before they are applied
- Events can now be triggered for a subset of reports
- Made several settings available via configuration UI
- Several health checks, system stability updates and container start optimizations
- Offline installer created to allow installation without an internet connection
- Version information now displayed in the "About" page

Integrating with Europol EU Cares system

Since many reports are processed directly by Europol, it was essential for us to establish a strong relationship and a mutually beneficial discussion of expertise with their IT specialists. As a result, we had virtual and in-person meetings, which resulted in Europol installing a version of AviaTor and actively testing it.

Other progress we made so far in the cooperation:

- Support for importing the new format from the Europol EU Cares system has been implemented in AviaTor which helps to prevent errors on missing files during import.
- Currently the integration of the extra information that is provided by Europol using the hit list is being implemented and should find its way in the upcoming Release 14. Scoring rules can be applied based on the information added by Europol.
- Further development will include a direct connection to the EU Cares LFE to directly insert reports without the need to copy them to the AviaTor server manually.

[5] https://esco.ec.europa.eu/en/classification/occupation_main

[6] https://huggingface.co/docs/transformers/model_doc/distilbert

MARKETING

Aviator Events & Communications

Our outreach efforts have been crucial in amplifying the impact of Aviator tool and bringing together a network of practitioners specialised in fighting CSAM. By organising various in-person events, we aimed to encourage relevant stakeholders to share their best practices and insights between each other.

These gatherings have been pivotal for strengthening our collective effort against CSAM, fostering a collaborative environment where knowledge and strategies can be exchanged to advance our common goals.

Annual Campaign

AviaTor's 2024 campaign launched in April 2024. The campaign focuses on AviaTor's impact on law enforcement workloads, its user growth and expansion, as well as its potential shortcomings for future users. The target audience is law enforcement agencies not yet working with AviaTor. Even though the main focus lies within the European Union, it is important for the sustainability of the AviaTor database that LEAs from outside the EU are also onboarded. This will also strengthen the law enforcement network of AviaTor users. In addition, the campaign aims to grow awareness of the issue at hand (high number of NCMEC reports and limited LEAs capacity to process the reports) and bring together different stakeholders during the AviaTor Forum to discuss this topic.

Aviator Forum

Our most anticipated event is the Aviator Forum, which will be held on June 26th, 2024 in Brussels, Belgium. This event is mainly organised towards law enforcement, current and prospective AviaTor users, hotline analysts, legislators, hotline analysts, academics, NGOs and other third parties.

The theme for the Forum is "Reshaping the reporting workflow of law enforcement in tackling CSAM" which focuses on three pillars, which are Innovation, Legislation and Cooperation. The innovation pillar focuses on how AviaTor and other innovative tools can improve the efficiency of LEAs' workload. The Legislation pillar discusses the implications of the upcoming changes in the EU legislation towards the battle against CSAM and how this impacts law enforcement. The Cooperation pillar shows the high level of interactive segments of this event, with several demos, workshops and a panel.

Attendees will be able to join forces in trying to solve fictional use cases in a demo of the AviaTor database and there will be a closed interactive demo of the ARICA project for LEAs only.

Alongside these themes, the speakers will also cover our trends and progress of AviaTor within the last year using data and information retrieved from LEAs users – and the official presentation of this report.

Peer-to-Peer Learning event

The AviaTor Peer-to-Peer events are held bi-annually, in which AviaTor users can learn from fellow law enforcement officers. This event gives a chance for knowledge exchange and "know-how" on tools and techniques to optimise user experience. Internally, this event also serves as a feedback tool and points us towards what features are working well and what areas need readjusting.

The fifth and final AviaTor Peer-to-Peer event took place in May 2024 in Amsterdam. This event focused on best practices of prioritisation, operational information exchange, and legal risk assessment management. Attending the event were speakers from INHOPE, Web-IQ, INTERPOL, Timelex, The National Police of the Netherlands (NPN), and ZiuZ Visual Intelligence. The event included interactive sessions on Operational Information Exchange where various national law enforcement agencies presented their case-studies and know-how.

Trending topics from Peer-to-Peer learning sessions

During the AviaTor Peer-to-Peer Learning sessions real case examples are shared in a confidential setting between AviaTor project partners and affiliate LEA's. The most prompting issues and topics are discussed among the participants such as OSINT, platform-specific tips & tricks, and retracted referrals. At the recent session, the trend of hacked accounts was more specifically addressed.

A number of examples of hacked accounts started making the news some 3-4 years ago, of people who had lost access to their Facebook or Google account and were accused of uploading CSAM. They had been automatically locked out, likely after a hash match triggered the platform's safety mechanism. This also leads to a NCMEC Cybertip to their country's law enforcement. Multiple people have sought public attention for claiming to be falsely accused of distributing CSAM.

In reality, their accounts were compromised by a third party. By uploading known CSAM, and knowing that this leads to an automatic ban, hackers manage to lock out the genuine owner of the account while they gain control of, for instance, a linked Facebook business account with monetary value.

As there can be quite some time between such an event and a law enforcement investigation, and it is not immediately apparent from the information contained in the report, several referrals have led to law enforcement investigations and even house searches, before it became clear that the alleged uploaders were the victim of hacking.

Retracted referrals have a similar impact. AviaTor users reported an increase in referrals that were later retracted by the ESP. In multiple cases, these had already been actioned by law enforcement.

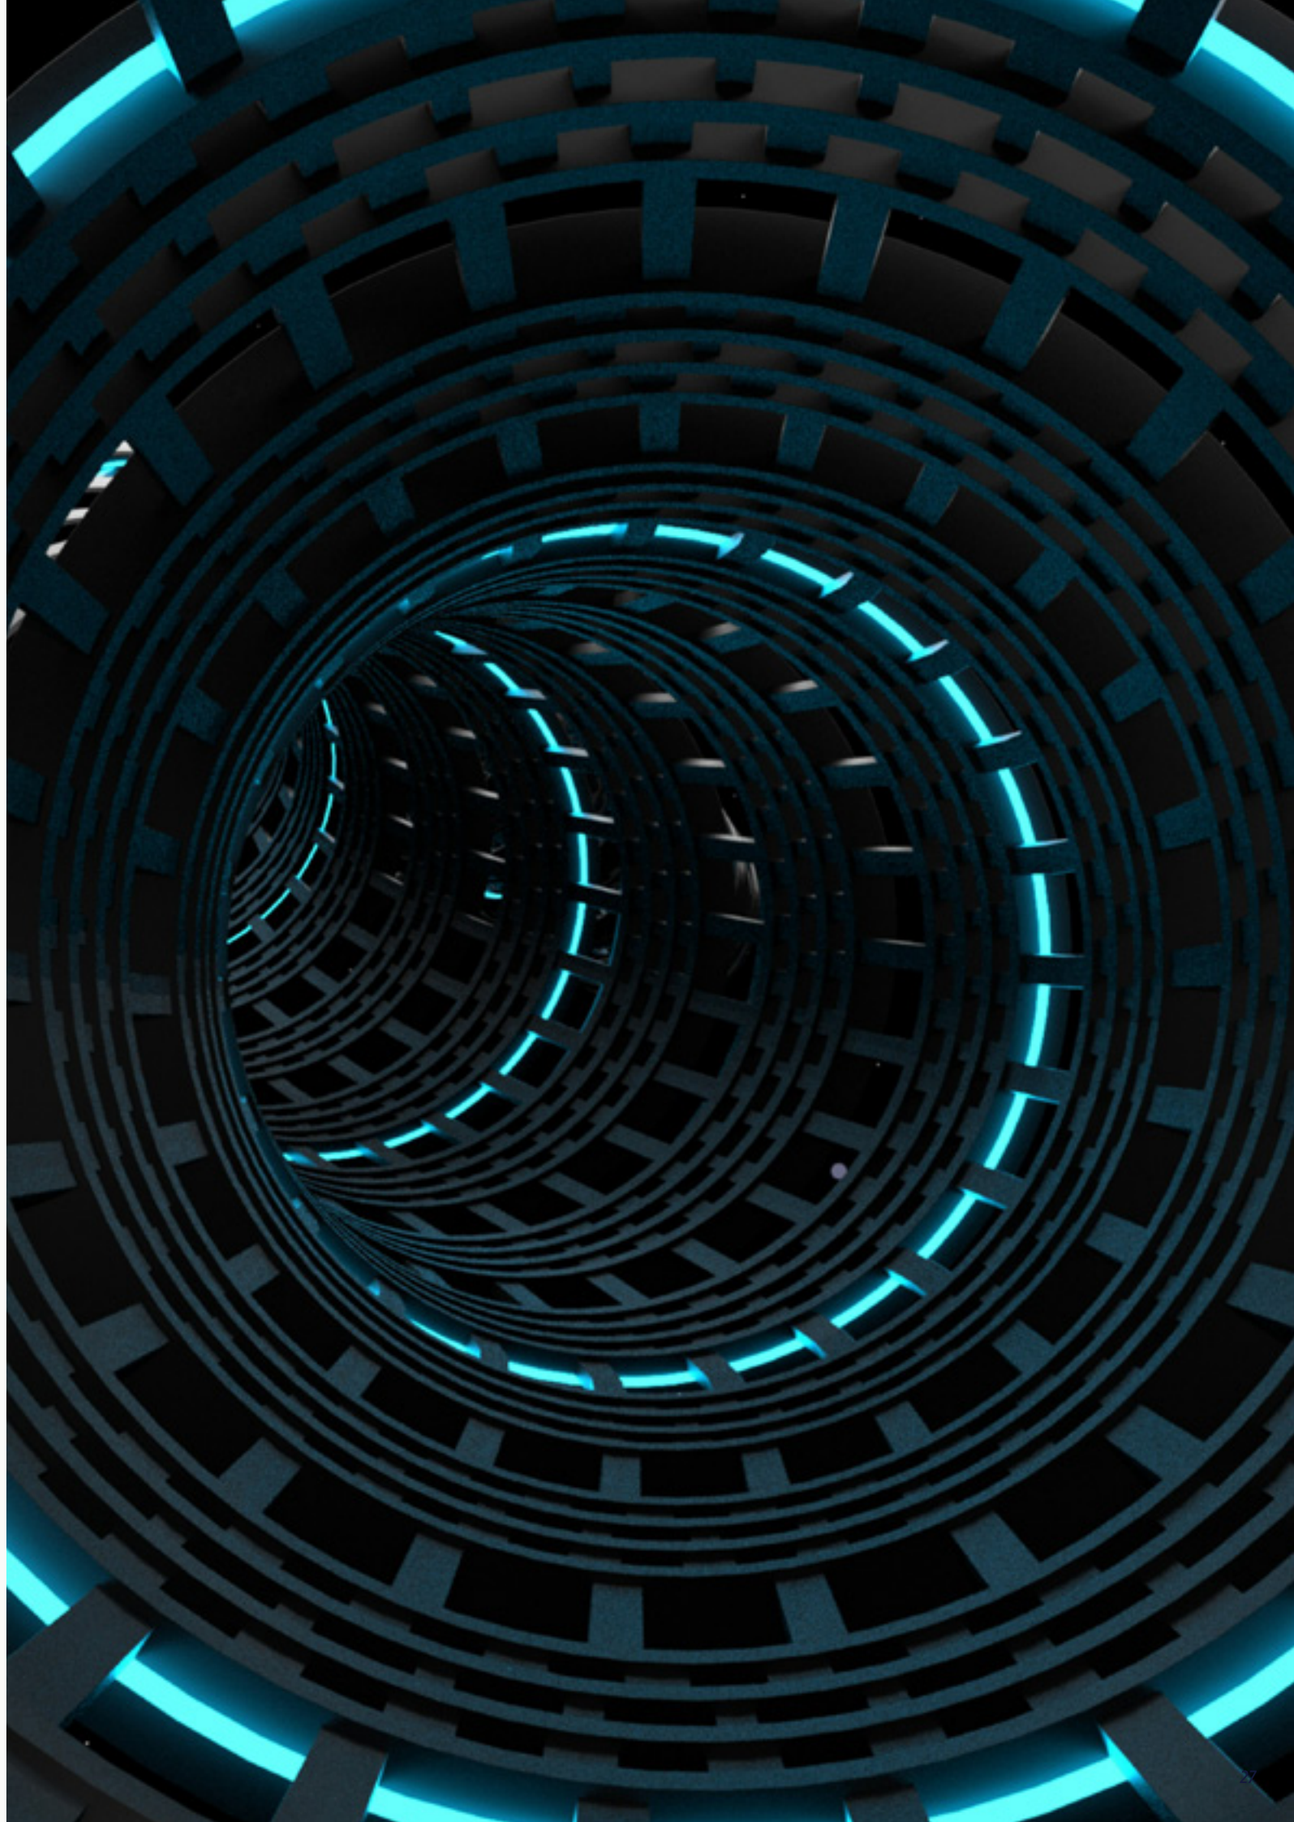

CHAPTER 03

Listening to AviaTor Users

Insights into the experience of our users

USER SUPPORT

AviaTor support system

One of the main challenges LEAs face in report processing is that, with thousands of reports waiting to be analysed, most time is spent on identifying illegal imagery within the content. Based on this insight, AviaTor created its tagline of "Save time, save lives". The AviaTor tool aims to reduce the amount of time spent on report analysis by automating the process using AI and OSINT. This allows for LEAs to devote their time and energy to prioritising reports and cases, identifying and prosecuting offenders, as well as rescuing victims of child sexual abuse and exploitation material.

This section covers the results of 17 interviews that were conducted with LEA users of AviaTor to gain insight into the efficiencies and complexities of the program.

Obstacles to overcome before implementing AviaTor – Results of the AviaTor affiliates survey

The 2023 Annual Report provided a first quantification of the impact AviaTor has on the time spent by law enforcement investigators to assess NCMEC referrals and to find pertinent referrals faster. In that 'Deep Dive' article we zoomed in on three countries using AviaTor and looked at the time-saving potential of deduplication, cross-matching and other AviaTor functionality.

To assess to what extent AviaTor is actually saving time in practice, the project team held interviews with all current AviaTor affiliates: law enforcement agencies who have signed up since 2019 to get a local installation of AviaTor.

Getting an AviaTor installation has however proven to be an arduous process for a number of affiliates due to several circumstances which will be discussed in this article.

We say that "AviaTor can be up and running tomorrow". The technical installation for a new affiliate takes about 1-2 hours and is done by local personnel, with remote assistance from AviaTor's support team. However, it is hardly ever that easy, as internal processes or technical, local issues only arise when installing AviaTor locally or at the start of testing. Some AviaTor affiliates have not yet been able to start using AviaTor even six months after joining the project.

To find out which problems they encounter, meetings were set up with 17 affiliates to survey them about bottlenecks in the process of getting AviaTor up and running in their department before, and after, installation.

Remote support

The technical support team (3 FTE) have three main tasks: to assist with installations, to investigate and fix errors, and to roll out new releases. While the project is working on installation packages whereby no real-time support is needed, only in an ideal world are no bug fixes necessary nor do errors occur.

The favourable news is that 100% of the affiliates are (very) happy with the support provided by the technical team. This is no small feat, since the team has no direct access to any of the installations, introducing several complexities.

The support team have to rely on the AviaTor users to initiate contact when something is not working as expected. In phase 2 of the project, we therefore initiated proactive regular communication with the affiliates.

Analysis of errors is also much more complicated for remotely installed systems, as it is not always immediately clear whether a problem is related to the AviaTor system, or to hardware and operating systems that are locally supplied, installed, and maintained.

Time-saving

When asked about time-saving, all affiliate LEAs who use AviaTor in production say, on average, it saves at least 20-30% of their time spent on triage and assessment pre-AviaTor. They expect even more time-saving and value with added functionality.

Bottlenecks in the onboarding process

The onboarding process for new affiliates usually starts with internal procedures that must be completed before they can install AviaTor. An impact assessment is often required before a new database can be used. Agencies' IT departments may impose additional restrictions before a new supplier can be added to the portfolio. Some features of AviaTor require a (temporary) outside connection, which in turn may require additional signoffs. When the red tape is gone, the technical installation of AviaTor takes about 1-2 hours, and subsequent releases are usually faster.

We asked the affiliates about other steps they have (had) to take to start using AviaTor:

- The project initially speaks to potential users of the system. Usually, the investigators who want to start using the software are not at the decision-making level, so having to convince management of the need for the system is not always easy. They are asked to make a presentation, with external consultation potentially involved, followed by another round of questions that, altogether, take up time.
- A comprehensive legal assessment.
- Large questionnaires that need to be completed by the AviaTor project team.
- Requesting permission to create a new database that contains AI and OSINT.
- Hardware may need to be acquired.
- Technical issues experienced during installation.
- Country-specific requirements. Among these are some major features that will take up a lot of development time, whereas AviaTor development is in principle driven by the majority.

- Integration with other software/systems if an agency has systems in place for part of the process.
- A testing phase with evaluation before taking any system into production.
- Not joining as an affiliate due to lack of time for feedback and peer-to-peer meetings.
- A general lack of manpower and/or time causing delays in the onboarding process.

A certain number of respondents giving feedback also pointed out that the installation lacked more specific details that would have improved the clarity of the installation instructions. Modifications are being made after careful analysis of the feedback.

Various measures are under consideration to help investigators in this process.

When bottlenecks are overcome and documents signed off, AviaTor is installed and configured to the local situation. As legislation differs among countries, settings must be configured locally. These settings pertain to different considerations, such as data retention rules, classification schemas, phone number format or scoring rules.

The complex process of configuring the scoring rules

Once up and running, most affiliates find AviaTor easy to use - with a big exception for the scoring mechanism. Based on the survey results, all affiliates deemed this to be the most complex part. Scoring involves determining potential indicators for high-priority situations and translating these factors into scoring rules for AviaTor. Factors for prioritisation are country-specific; therefore, the project leaves the configuration up to the agencies and does not prescribe any rules.

To our knowledge, no validated model is available for AviaTor that can be used for the prioritisation of NCMEC reports. Many of the LEAs still utilise their own programmed spreadsheets, according to the information we got from speaking with them. Existing risk models for prioritisation of online CSE cases can only be used to assess identified individuals. However, upon receipt of an NCMEC report, the identity and location of the reported offender are usually not yet known.

It is difficult to recognise potential indicators and to translate (enriched) NCMEC data to potential child abuse situations, especially when the picture is often incomplete. '[Two] reports can look nearly identical to a law enforcement officer. Investigating both, however, could yield very different results: one may reveal no further illicit activity, while the other could uncover evidence of hands-on abuse'. To make steps, peer-to-peer learning events are organised, to facilitate the exchange of expertise, as well as consultations with outside experts who may be able to provide guidance.

To conclude this section, the data presented has supported AviaTor's positive impact on LEAs' workload and efficiency with report processing. Moving forward, a focus for the AviaTor team is to better help streamline the installation process, for instance by creating a pre-installation checklist that includes mandatory legal and technical assessments, making it more straightforward for users, while also considering solutions and opportunities for users to learn (i.e., knowledge-sharing seminars) and be guided on complexities, enhancing their user experience.

[7]Grossman, S., Pfefferkorn, R., Thiel, D., Shah, S., DiResta, R., Perrino, J., Cryst, E., and Stamos, A. (2024). The Strengths and Weaknesses of the Online Child Safety Ecosystem. Stanford Digital Repository. Available at https://purl.stanford.edu/pr592kc5483. https://doi.org/10.25740/pr592kc5483

LAW ENFORCEMENT

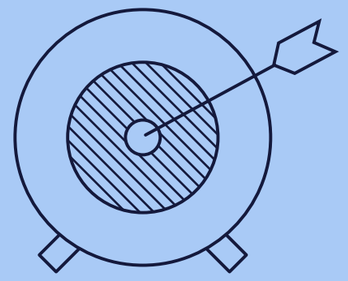

CHAPTER 04

Obstacles in becoming an AviaTor User

An interview with Law Enforcement

INTERVIEW

LEA

Tackling obstacles & challenges in becoming an AviaTor User

AviaTor is a great tool for Law Enforcement to optimise their work in processing industry reports on CSAM. However, the journey for LEAs to adopt AviaTor can be both time-consuming and challenging. It involves progressing through several key steps, starting with expressing initial interest in the AviaTor database, followed by the installation and implementation phases. The goal is to achieve full operational integration of AviaTor within the law enforcement unit. Each step is a substantial undertaking, requiring careful consideration and execution.

To gain more insight into this process, we spoke to **Yves Goethals, Judicial Commissioner of the Belgian Federal Police in Brussels** and **Detective Inspector Darren Young from the Metropolitan Police in London**. Two vastly different national law enforcement agencies that joined the AviaTor project at different stages.

As part of the project team and one of the key partners of the AviaTor project, the Belgian Federal Police was one of the first users of AviaTor and has had a significant impact on the project and the development of AviaTor itself. Being the first to test and use AviaTor, the challenges for the Belgian police mainly lay within the extensive Belgian laws around data privacy that come with additional mandatory paperwork.

The London Metropolitan Police, on the other hand, is one of the latest users of AviaTor. By joining the project later on, they received a version of the AviaTor database tool that already went through several rounds of updates. More database functionalities however also lead to more technical challenges before the tool can be fully operational. They faced a relatively long trajectory leading up to the implementation of the AviaTor tool.

Hurdles on the road to becoming an AviaTor user

Can you describe any initial challenges or hesitations you faced when considering the adoption of AviaTor?

Yves: My only hesitation was the fact that it was a project. Not every project is able to deliver what they promised. But seeing the partners involved I gained a lot of trust in this project. The Belgian Federal Police needed a tool that could assist with the prioritisation of NCMEC reports, so this made the choice to participate rather easy.

Darren: We faced mostly internal challenges in terms of data protection, impact assessments and installing the software.

Technical challenges

Have you encountered any technical difficulties while setting up or using AviaTor?

Yves: I needed some support when it came to finding the right hardware to use. The setup itself, however, was quite straightforward as we were one of the first users testing the first version of AviaTor. The slow build-up made it easy for everyone to get used to the system.

Darren: The hardware we use for AviaTor we already owned, so setting up and installing was not difficult. When we started the initial testing of the tool we ran into some minor bugs that needed tweaking for us to be able to make AviaTor operational. We are working with the developers to address these issues and find solutions.

Are there any technical features you wish AviaTor had to make your experience smoother?

Yves: I would like to have the ability to make changes to the different criteria ourselves without having to ask the developers. Many different countries use AviaTor and I can imagine they all have different needs, so making this functionality customisable by the user could be beneficial.

Darren: There are some minor tweaks I would like to see implemented. For example; I would like to move from one referral to the next by using a simple “Previous – Next” functionality.

Privacy assessment

What do you think about the clarity and comprehensiveness of AviaTor's privacy assessment?

Yves: We are working with information that belongs to a police investigation and the clarity of the privacy assessment is comprehensive e nough and follows the regulations and rules around privacy and GDPR.

Darren: I don't think AviaTor has that in itself. Our internal governance unit assesses all tools we use and decides if the tool is safe enough to hold information. It took us 9 months to get this signed off by the Metropolitan Police internally. It took this long because it was difficult to clarify the concept and usage of AviaTor. To get approved, the safety of the tool is measured by looking at GDPR compliance, the AI and machine learning that is involved, where we store the data, what type of room it is stored in, who has access, etc. Many aspects come into the mix.

Were there specific privacy-related concerns that you faced during your onboarding or use of AviaTor?

Yves: AviaTor arrived at a time in Belgium when we just had controversy over the possible use of Clearview AI software by law enforcement. Clearview is a facial

recognition company whose algorithm matches faces to a database of over 20 billion images collected from the Internet, including social media applications. The software was deemed unlawful in 2022 by the Belgian Supervisory Body for Police Information, which led to high awareness among law enforcement about the use of open-source intelligence in the context of the applicable general legal framework and data protection law in particular.

Thankfully, we didn't face many issues with the first version of AviaTor, since this was a completely closed fully stand-alone system which was not connected to the data system of LEAs. The hashes being used in AviaTor weren't an issue either, as we are only working on ceased material and no open-source material. As long as the material that is being used is ceased or reported and found as illegal material - we comply with the law. But we are not allowed to use publicly available material online.

Legal Assessment

Were there any legal implications or challenges you came across when implementing AviaTor?

Yves: As a project partner, we had to make an assessment at the start of the project. Once the project was accepted, so was the tool. The only challenge we encountered is having to provide a very detailed explanation of which parameters we would use for prioritisation. As AviaTor is a prioritisation tool, naturally that means that certain cases will be deprioritised. We had to explain how this process would work and how we would handle that remaining data. But it was immediately accepted once we explained our method.

Bureaucracy

Did any bureaucratic processes affect your implementation of AviaTor?

Yves: Not more than any other project. Bureaucracy rules everywhere, so it's the same process for all projects, and we don't remember any challenges specific to Aviator.

Darren: There weren't really any bureaucratic delays.

Lack of Capacity

Were there instances when you felt a lack of capacity?

Yves: The system itself operates in an easy and straightforward way that doesn't demonstrate any capacity issues. You are only really confronted with your lack of capacity once you begin handling reports that have been chosen by the system.

The biggest issue is having to go through reported material that is not illegal. Last year we received almost 20.000 reports, with only 40% illegal material. And it takes a long time to look through these reports. Unfortunately, this issue cannot be solved with tools like AviaTor. When you ask the tool to sort for unknown material all legal content will rank high because they are unknown images and, therefore not in the database. And this increases the workload significantly.

Darren: We are not using the tool operationally yet. So, once we get a stable version installed, we will see if we encounter any capacity issues. So far, it's just me using the tool, once I'm happy with everything I will hand it over to a team who will use it on a daily basis. In the UK 43 different police forces will be able to benefit from it. Once we are using it in London, we will demo it to the other forces.

Data Privacy Impact Analysis (DPIA)

How would you describe your experience with making a DPIA?

Yves: I made a DPIA for every tool we have, including AviaTor, which will have to be adapted for phase two. The process is quite straightforward, and the approval process is fast and uncomplicated.

I love small-scale projects like AviaTor and CPORT because if an issue arises at any point, you can just pick up the phone and call your developers. It's easy to get in touch with people if you need help. That's the advantage of a small-scale project.

Darren: The DPIA took 9 months to get done. I had to fill out the assessment form and submit it to another department for review. In the UK the assessment form is an extensive document and the process of getting it approved by multiple departments is time-consuming.

Developer Support

How do you feel about the support you have received from AviaTor Developers?

Yves: I believe the outreach has to come from us because we are the ones who use the tool. So, if we notice anything that needs adaptation, we need to communicate it to the developers. When we do, they always try to help us.

The tool is being used in many countries and has to work for different legislations. So, it's great that it is customisable in the way that you can decide to not use certain criteria. However, I wish there was also an option to add certain criteria ourselves applicable to our national circumstances. This would also help to ease the workload for the developers.

Darren: We are all spread around the globe, and getting support remotely can be difficult. There is currently a need for several devices to communicate with the developer's team due to the use of different servers, which is a bit problematic. It would be useful to have some type of remote support in the future.

Management Support

How supportive was your management team during the adoption and continued use of AviaTor?

Yves: We had no problems; our director was very supportive. The only hesitation he expressed was that sometimes we have to be careful with projects, as it can happen that they promise more than they are able to deliver. But he trusted my lead to make this work.

Darren: They are desperate for us to start using it.

Financial support

Did you face any financial constraints or challenges when implementing or using AviaTor?

Yves: No, because we are partner to the AviaTor project and therefore the implementation and testing costs are covered by the AviaTor Funds.

Darren: We did not run into any challenges because AviaTor does not cost money at the moment. The hardware necessary for AviaTor we already had in our property. If this changes in the future and AviaTor costs more money this would not be an issue, as long as AviaTor can deliver what we hope it can in terms of prioritisation and reducing workload

Have you experienced any other challenges you want to note?

Yves: Nothing to note for now. The only challenge will arise when the project finishes as it is currently unclear what will be the possibilities in terms of technical support after the end of the project.

Darren: The next challenge for us will be when AviaTor becomes operational, we will be redesigning our workflow to implement AviaTor. But hopefully that will be quite straightforward.

Darren Young

Detective Inspector at the Metropolitan Police in London

Yves Goethals

Judicial Commissioner of the Belgian Federal Police in Brussels

LAW

CHAPTER 05

Legal Checklist

A Legal checklist for Law Enforcement Agencies when using Open-Source Intelligence (OSINT)

OSINT CHECKLIST

LAW

Legal Checklist

One of the important ways in which LEAs maintain law and order and obtain information about (potential) crimes, is by being present in and among the public. Historically, this public sphere has of course mostly consisted of our towns, streets, squares and parks, but in the last few decades, it expanded into the virtual realm and entails publicly accessible websites, social media and other online areas.

Inherently to their nature, these online areas are particularly useful for obtaining information about individuals and how they communicate and interact with one another. They serve as an important source of intelligence about potentially criminal behaviour, irrespective of whether that behaviour itself took place in the physical or the virtual world. The intelligence derived from information obtained from publicly available sources is commonly referred to as "open-source intelligence" or "OSINT".

For many LEAs, the use of OSINT has become an important part of their overall information-gathering toolbox. The virtual world indeed presents opportunities for obtaining information about individuals with ease and on a scale which are incomparable to what is possible in the real world. This, in turn, has motivated legislators to introduce a range of different legal instruments meant to protect individuals' fundamental rights, including a right to privacy and the protection of personal data, online. Generally speaking, these instruments do not prevent law enforcement from using OSINT, but they do create additional limitations and constraints to prevent overzealous police work.

While the main instruments protecting the fundamental rights of individuals online have been introduced and are managed at the European level, the rules governing law enforcement use of OSINT are largely left to the discretion of national legislators. Therefore, there is no one single European answer to the question under which conditions and constraints law enforcement is allowed to use OSINT. That said, there are certain commonalities discernible across Member States regarding such conditions and constraints. This contribution provides an overview of the different practical aspects law enforcement should consider or check before and during the acquisition of open-source intelligence. Additionally, this contribution will look into the legal requirements that need to be taken into account when storing the data collected through OSINT in a police database.

Acquisition of AviaTor data: legal requirements when using the OSINT component

This section will investigate the opportunities of LEAs in the European Union to gather information which is publicly accessible on the internet (social media, darknet, forums, etc.). This aspect is important to consider, as in most cases the information that is included in an NCMEC report will not be sufficient to appropriately prioritize reports. LEAs will therefore need to rely on other information sources. Upon receipt of an NCMEC report, investigators will be further interested in, for example, search terms that are used by the suspect in the report on Google or other websites (such as porn websites), chat messages that were exchanged on non-E2E communication platforms, posts on social media, active use of any darknet websites, etc.

The information that is collected through OSINT is used in AviaTor to prioritise the NCMEC reports more accurately. However, this type of investigation is subject to strict restrictions due to the level of intrusiveness on the right to privacy of the person mentioned in the NCMEC report. In this section we will give a general overview of the legal requirements that need to be taken into account by LEAs when they are using OSINT investigative techniques:

01

Assess the applicable legal basis for using OSINT

In many jurisdictions, the criminal procedural law grants investigatory powers to LEAs for real-life observations and/or information-gathering but does not specify the possibility of using this investigatory power in an online setting. However, doctrine and case law have recognised that these investigatory powers can also be used in an online setting, provided that the same legal restrictions as applicable in an offline setting are respected.

In general, criminal procedural law in Member States will grant a general investigatory power to LEAs to "detect crimes and gather evidence1" and to perform their task of "upholding law and order2". In essence, these general legal bases will allow police forces to perform their tasks without the need to obtain specific authorisation from a judicial or other competent body. For OSINT investigations, this more general legal basis can be used in cases where the interference with the fundamental right of privacy is limited (see also point 2 below).

02

Check whether a formal authorisation from a judicial or other competent body is required

In many jurisdictions, the competence of law enforcement to use OSINT follows from the general task of upholding law and order, as explained above. Such a general task requires law enforcement to have competencies which it can exercise without prior authorisation from a judge, public prosecutor or other role within the judiciary. Nonetheless, it should be verified whether a particular use of OSINT is to be construed as exercising special investigatory powers which go beyond the normal observation of public places and which would require prior authorisation.

In cases where the impact on the investigated person is more grave, i.e. in the case of systematic observations of behaviour on the internet, the police forces will need to rely on a more specific legal basis (i.e. to fulfil the requirement of the foreseeability in law).

03

Verify whether the online resources are truly public and limit the use of OSINT to such publicly accessible resources

A key characteristic of OSINT is that the intelligence is derived from information acquired from public sources. Whether a resource can be considered public will depend on national law and/or interpretations by national courts. Examples of elements which may be relevant to consider whether a resource is public, are:

- Registration requirements (e.g. do you need to create a profile to access the website?)
- Identity or device verification steps (e.g. will the resource check whether you are who you say you are, to check whether you are an authorised user?)
- Accessibility on invitation only (e.g. can the resource only be accessed by individuals who have been personally invited?)
- Obfuscation of access links (e.g. can the resource only be accessed by those who already know the resource's URL, which is not known to anyone else or by search engines?)
- Available data access mechanisms (e.g. is the resource accessible to a public API which anyone can use?)
- Communicated access conditions and intended target audience (e.g. do the terms and conditions of the resource clearly stipulate that the resource is not meant to be used by everyone or for particular law enforcement purposes?)

All of the examples above may, from time to time, be construed as limiting who has access to the resource and as such its public nature. Hence, the LEAs need to assess at the national level to what extent such elements impact the qualification of the resource as public or non-public.

04

Security measures used by websites, social media, hosting companies and other public online resources should at all times be respected and observed

The "open" or "public" nature of a resource can never be forced. This means that law enforcement can never use measures to circumvent or force through the security measures which a resource has put in place, even if those measures are incredibly simplistic (e.g. login name and password are both "admin").

05

Refrain from indiscriminately (bulk) scraping whatever information can be found

Most law enforcement agencies are not allowed to go on so-called "fishing expeditions". This means conducting investigations with no clear indications of any wrongdoings by someone, with the sole purpose of trying to find incriminating evidence. The temptation to engage in such fishing expeditions is particularly present in OSINT. By its digital and accessible nature, it is abundantly easy to first collect a seemingly innocuous dataset which later is subjected to advanced analysis techniques. Hence, law enforcement should always remain targeted in the way OSINT capabilities are used (see also below).

06

The purpose of using OSINT must be strictly limited in aiding LEAs to prioritise reports, based on pre-determined criteria

As stated above the use of OSINT can be based on the general legal basis of 'upholding law and order', and will therefore not need a prior authorisation by a judicial or other competent body provided that the interference with fundamental rights (i.e. right to privacy) is limited. Therefore, the OSINT use must be limited to targeted information-gathering for prioritisation reasons. Using the intelligence for other purposes could be construed as an illegal fishing expedition.

07

Do not engage in in-depth observation of one or more specific individuals (privacy infringement should be kept to a minimum)

This requirement again ties back to the need to use the OSINT capabilities in a targeted way. Observing a specific person's activity over a longer period and combining data found on a multitude of internet resources, will inherently be more privacy-invasive than targeting the OSINT research to certain specific darknet websites or other internet sources that are mainly used by offenders.

08

Do not use misleading or enticing undercover techniques when obtaining OSINT

The interference is not considered limited when a police officer uses aliases to first befriend the suspect, and then communicate with the suspect on online forums or social media. For these kinds of investigative techniques, LEAs will not be able to use the general legal basis of upholding law and order, as they will be able to have a somewhat complete insight into the suspect's private life. This goes further than solely observing the publicly accessible information of the suspect. Again, this does not mean that these types of investigations are illegal per se, but it does mean that LEAs will need to obtain authorisations from a judicial or other competent body. It goes without saying that in any case, LEAs cannot use any entrapment techniques to induce the suspect to commit further offences (i.e. by asking the suspect to share CSAM).

09

Verify whether an external service provider can be used to perform OSINT activities

Some jurisdictions explicitly prohibit that police data can be accessed by external service providers (so also the provider of an OSINT tool). LEAs need to assess whether such a prohibition exists in their local jurisdiction. In general, the following elements need to be considered by LEAs when using external service providers:

- Verify the contractual requirements that need to be put in place (license agreements, data processing agreements, etc.);
- Some jurisdictions specifically provide localisation requirements, i.e. that the data cannot leave the jurisdiction or that if it can leave the jurisdiction, it is subject to very strict conditions;
- Some jurisdictions explicitly prohibit reliance on service providers that are established outside the European Economic Area;
- Due to the very sensitive nature of the data that is going to be processed, it is of paramount importance that a vendor's due diligence process is established (i.e. by using pre-determined selection criteria, vendor security questionnaires, etc.).

Storage of AviaTor data: rules on setting up and maintaining a police database

After the above-mentioned 'information gathering phase', it is clear that the information will remain in AviaTor. Consequently, this means that there will be the creation of a police database, triggering a number of rules which are (mostly) national law specific. Overall, the following legal requirements can be discerned:

- Some jurisdictions require a prior notification of a new (special) police database to the Supervisory Authority (under the LED).
- Identify the technical requirements for the creation of such a police database, including:
  a) whether the database can be hosted locally or in the cloud;
  b) whether specific police resources need to be leveraged (e.g., an existing police database with similar functionalities and purposes);
  c) which security measures are to be put in place (e.g., integration in the existing information security architecture, data classification, logging, etc.).
- Verification of the role-based access requirements to the newly created database

For each police officer that needs access to the AviaTor database, the assigned access permissions need to be adapted to their role and function.

- Performance of a prior Data Protection Impact Assessment (DPIA)

The obligation to perform a DPIA is included in the data protection framework set out in the Law Enforcement Directive ("LED"). A DPIA is necessary when a processing activity is likely to result in a high risk to the rights and freedoms of natural persons. It must be clear that in AviaTor, a large amount of highly sensitive data will be processed. Sensitive data includes data concerning a person's sexual interest and data concerning vulnerable children, or when OSINT shows particular criminal activity on darknet websites. Moreover, the risk to the rights and freedoms of natural persons will be high, as false positives could lead to someone being falsely accused of possessing or disseminating CSAM. Based on this assessment, performing a DPIA will be necessary.

- Maintain log files regarding actions taken concerning the data in the AviaTor database

Any action performed with the AviaTor tool must be sufficiently logged, i.e. access to reports, changes to reports, deletion of reports, etc.

- Define the data retention periods of the data in the AviaTor database following local laws or sectoral laws

Data retention rules are often specific to Member State legislation. For example, IP addresses can only be stored for a limited period and need to be deleted or anonymised following local data retention laws.

- LEAs need to take into account specific rules on linking and/or correlating police databases

It must be noted that when LEAs decide to correlate and/or link the data in AviaTor to data which are stored in other police databases, specific rules may exist in national law that need to be considered. In essence, these rules often ensure that, through correlating or linking police data in different LEA databases, the data remains only accessible to the police officers on a "need-to-know basis". The access rights and data retention rules should be respected and in line with the purposes that were initially established for each database.

Concluding remarks

This contribution has provided LEAs with an initial legal checklist when using the OSINT component of AviaTor and subsequently, when storing this data in police databases. However, it must be highlighted that this checklist is a starting point for LEAs to – together with their legal department and/or data protection officer – assess the local legal requirements that will be applicable when using the OSINT components and when creating a new police database.

TECHNOLOGY

CHAPTER 06

Federated Learning

Annotating images to train AI models

TRAINING AI

TECHNOLOGY

Testing Federated Learning

One of the goals of the AviaTor project is to assist police agencies in the ranking and triage of NCMEC reports. Such NCMEC reports are suspected to contain CSAM, usually consisting of a series of images or videos. These reports have high priority for police forces due to the seriousness of the crime, all the while being extremely difficult, time-sensitive, and labour-intensive to process. Therefore, AviaTor employs AI methods to detect relevant content to help police investigators in this work. AviaTor is currently being tested in several different EU countries. While each country uses their own AviaTor system, it would be a natural fit to join forces by allowing the AI models to communicate with each other.

↗ **Annotation efforts AI models require high-quality training data to be trained for a specific task (e.g., image classification).**

Due to the sensitive nature of CSAM, this kind of data is not freely available for training AI models. Therefore, images from NCMEC reports were provided by the Dutch police for manual annotation, which was performed by the Dutch hotline Offlimits. Similarly, the Belgian Federal Police provided images for annotation to the Belgian hotline Child Focus.

Annotators were provided with a detailed annotation guideline to identify important concepts. The annotation process was supported by the annotation software Label Studio. Offlimits and Child Focus have rules and guidelines in place that stipulate the working conditions under which the data annotation has to take place. The annotation concepts were developed after meetings with partners from law enforcement, after which we mostly adopted a pre-release draft version of the Universal Classification Schema 1 to define several concepts, such as:

- Estimated age of all depicted persons
- Ethnicity of the victim
- Gender of the victim
- Type of sexual interaction (if any depicted)
- Amateur or professional/studio setting
- Modification (e.g., anonymisation or meme) of the image

As of January 25th, 2024, this led to the annotation of approximately 8,000 images by Offlimits. Annotation at Child Focus started slightly later, therefore no numbers will be provided in this report. The median annotation time was 22.4 seconds per image. All concepts, except the studio setting, allowed multiple annotations. For example, an image potentially contains more than one person and therefore allows one to annotate several maturity levels or skin colours.

Individual distributions are depicted in Figure 1. For most categories, we observed strong biases. For example, the majority of images depicting potential victims were white females. A large proportion of analysed images (46.8 %) depicted potentially adult persons, followed by the second highest category of pre-pubescent children (32.0%). The first category indicates the difficulty posed to the everyday work of police officers in assessing the legal traceability of NCMEC reports.

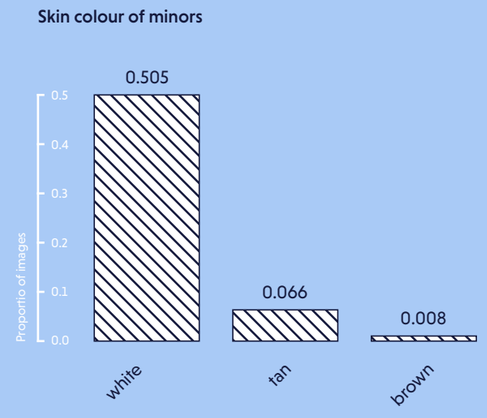

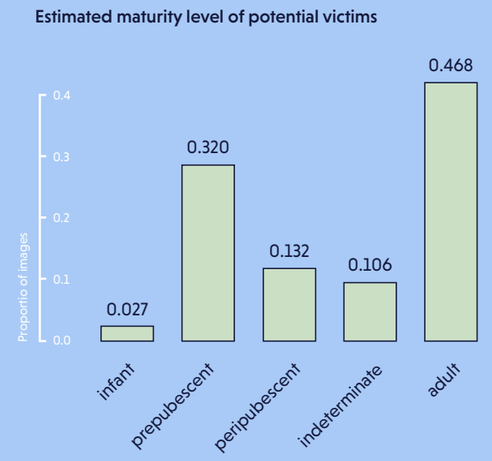

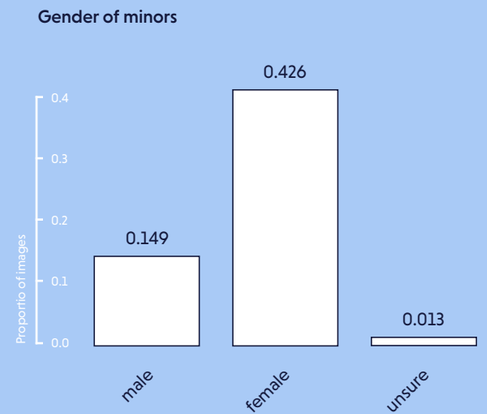

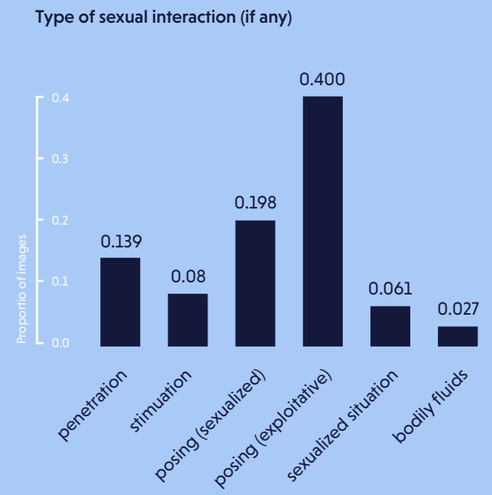

Classical approach for image classification using machine learning

This sub-section briefly explains the most frequently used strategy to train a state-of-the-art classifier on comparably little training data. In most cases, image classifiers are trained by adapting an existing state-of-the-art model. The original model is then usually trained on a relatively large image dataset, such as ImageNet. ImageNet, with more than 14 million labelled images, is probably the most frequently used large dataset for pre-training a model for image-related tasks, such as image classification. The general assumption is that the initially trained model implicitly learns to automatically discover and extract meaningful features or 3 representations from the training data. This process is frequently referred to as representation learning.

Subsequently, the pre-trained model is then applied to the relevant task (e.g., CSAM classification), where usually substantially less training data is available. However, the model still relies on the image representations learned during the first task. This strategy, known as fine-tuning or transfer-learning, has been shown to substantially lower the required amount of annotated training data to build a good AI model and leads to better results than training on the small dataset alone. However, it is still the case that more training data usually results in better and more robust models. To this end, data augmentation can be used to artificially increase the number of training instances by performing transformations on the data (e.g., brightness and contrast adjustments, rotations, noise injection, ...).

Motivation for federated learning Manual annotation is not only time- and labour-intensive, but due to the sensitivity of the content, nearly impossible to outsource.

Therefore, the most likely option for acquiring more training data is to integrate the annotation process into the day-to-day work of police officers working (and therefore assessing) with NCMEC reports. In order to increase the amount of available training data for the model, it would be great to incorporate such byproducts from different police agencies. However, this poses serious legal problems, as CSAM data cannot be easily shared (even for LEAs) across national borders, or in some cases federal borders.

One promising way to train models with more data, without having access to the data, is federated learning. Federated learning allows a model to be collaboratively trained across multiple decentralized devices holding local data samples without exchanging data. This ensures data sovereignty and protects data privacy. The federated learning paradigm is in stark contrast to the regular machine learning paradigm, which aggregates all the data in a central place for training.

Federated learning implements usually the following steps:

1. Initialisation: A global model is created and sent to all the devices participating in the federated learning process.

2. Local Training: Each device fine-tunes the model locally using its own data. The training is done on the device itself, and only the model updates (not the raw data including CSAM) are sent back to the central server.

3. Aggregation: Individual updates (Step 2) are collected by a central server and are then used to update the global model (Step 1). This global model now incorporates insights from all the local datasets without compromising individual data privacy.

4. Iteration: Steps 2 and 3 are iteratively repeated. The global model is sent back to the devices, and the process continues with each device updating the model based on its local data.

Experimental setup

In the AviaTor project, we will execute a preliminary study to evaluate the potential of federated learning in the context of CSAM classification. In this study, we will compare supervised learning with federated learning.

Models for image classification

Over the last years, several neural models for image classification have been proposed. Recent models are either using a Vision Transformer model or some sort of convolution. There is still ongoing debate about which of the two architectures is better suited for image classification.

In real-world applications, we do not only consider performance but the complexity of a model. Usually, the more parameters a model has, the higher the burden for computation power. Therefore, in several applications, well-established general-purpose models like VGG, ResNet, Inception, or DenseNet are used. In our experiments, we will select an appropriate model architecture to perform our experiments. Following common practices, we will split the annotated data into training, development and test-data. This split will be performed on premises (i.e., at Offlimits and Belgium police).

Baseline

To evaluate the impact of federated learning, we need to define a baseline first. In this setup, the model has only access to the data from one physical entity. For each entity, we will fine-tune a supervised image classification model on the training data and evaluate it on the test data. This means that we specify the image classifications on the testing data, and then feed it into the model. During the training phase, relevant parameters will be monitored on the development set. Hence, trained models have only access to their "personal" training data and cannot exchange information.

Federated Learning

In the second training regime, we will implement a federated learning scenario. In this setup, both models will again be fine-tuned on their training data individually. However, to use the advantage of additional training data, the individual gradient updates will be communicated between the two models via a central mediator. For the implementation of a federated learning strategy consider using the FLOWER framework.

One difficulty poses the implementation of a communication mode due to the high-security needs of the individual premises. If such practical considerations make a realistic federated learning experiment unfeasible, we may choose to experiment with a single location by splitting up the data set between two client nodes that are connected to the central server directly via Ethernet cable. In any case, the practical steps will be discussed together with the relevant project partners.

Conclusion and future of federated learning

In this chapter, we exemplified the advantages of federated learning in our scenario and explained the ongoing experiments. We believe techniques, such as federated learning, can help overcome the data sparsity problem to train a strong CSAM classifier in the future.

An interesting direction for any future experiments would be to investigate methods that deal with mixed label spaces to address differences in labels and the definitions thereof. This would allow each local model to use additional labels. For example, some agencies might be interested in extending our pre-defined set of labels with individual categories.

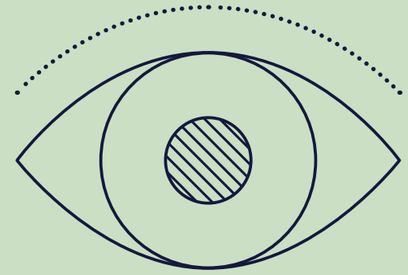

CHAPTER 08

The Way Forward

The future of AviaTor after the project ends

SUSTAINIBILITY

The Future of AviaTor

Whenever we give an AviaTor demo to an interested LEA, we also inform them that the project is coming to an end. The obvious questions we get are: But what happens after the project? Can I keep using AviaTor? Will it be supported? These are, of course, very legitimate and relevant questions. And in this closing article, we will address them.

Sustainability

Since the European Commission funded AviaTor for two consecutive periods, for which we are very grateful, it put a lot of emphasis on creating a reliable sustainability plan. To this end, we drafted a concept sustainability plan halfway through the project. A final sustainability plan is due at the end of the project.

During the project, we worked to create the best starting point for sustainability from a technical perspective. We created a clear separation of concerns between the development partners in the project: Web-IQ and ZiuZ. As a result, the maintenance and support of the AviaTor system can be facilitated by ZiuZ, avoiding the need for LEAs to deal with multiple partners to resolve issues. ZiuZ can also maintain the AI CSAM classifier, which is part of AviaTor. Web-IQ can take responsibility for the text classifier and targeted online research, which can be supported and maintained independently of the AviaTor system.

We also optimised the installation and upgrade of the AviaTor system to take place in a very limited time frame and we intend to improve this even further by using tools like Ansible to configure and deploy AviaTor.

And finally, in the last phase of the project, we focus on the reliability and performance of AviaTor.

Current users

At the end of the project, we expect to have between 25 and 30 LEAs testing and/or using AviaTor. For these LEAs to be able to continue the use of AviaTor, we must provide several of the services as explained in the following table:

Service	Provider	Funding
Implement new generic user requirements	ZiuZ	Annual service fee
Maintain the current functionality of AviaTor	ZiuZ	
Implement new generic user requirements	ZiuZ	
Bug resolution	ZiuZ	
Help desk	ZiuZ	
AI classifier for imagery	ZiuZ	
AI classifier for text	Web-IQ	
Targeted online research (including support, maintenance, new features and sources, and training)	Web-IQ	
Implementation of specific user requirements	ZiuZ	Contract
Purchase and maintenance of hardware	LEA	P.M.
Maintenance of system software (Ubuntu, Docker, etc.)	LEA	

Following the end of the project (Sept. 2024) until the end of 2026, the annual service fee will be based on an open calculation at cost price that will be shared with the project partners and affiliates. For this calculation, we will review the maintenance and support cost over the project and estimate the effort necessary for the implementation of new generic user requirements.

The cost for the "Targeted online research" will be handled between each LEA and Web-IQ individually, and will initially be free until the end of 2023. Additional costing arrangements may be put in place in case of very high volumes.

Request for the implementation of specific user requirements will be discussed between individual LEAs and ZiuZ or Web-IQ and agreed upon in a separate contract.

New users

For new users, we will follow the same calculation principle but at an economic rate.

Additional funding

We consider the period after the project until the end of 2026 as a transition from an EU-funded project to a commercial product that fits a real need from law enforcement customers. The EU funding made it possible to develop the product and will keep the product affordable for all users in the long run. In the calculation described above, the essentials are included to keep AviaTor running reliably for current and new users.

There are also activities that we want or need to carry out, that are not included in the annual fee calculations. For these activities, we will try to find additional funding.

Examples of these activities are:

Onboarding

It can take a LEA up to 1.5 years to get AviaTor installed after the initial interest is shown. This is mostly due to the decision hierarchy and the decision-making process in place. The fact that AviaTor is free software does not change this. Often, a DPIA has to be carried out, and the fact that AviaTor includes AI and collects OSINT can contribute to a further slowdown in the approval process.

The effort that ZiuZ and Web-IQ spend to support new LEAs during the onboarding process is not included in the annual fee calculation.

Training

AviaTor is easy to use! That's the feedback we have received from a survey of our users. However, we also know users struggle with setting up the scoring rules that determine the priority of reports in the system. There are extensive rules that can be configured per individual LEA for reports, reported persons and imagery contained in the reports. Since ZiuZ and Web-IQ do not (even) have access to actual reports, it is difficult to assist LEAs in determining the right configuration for their specific situation.

With additional funding, we could set up a specific curriculum to address this 'training' gap in collaboration with external partners like the European Union Agency for Training Law Enforcement (CEPOL) or the European Cybercrime Training and Education Group (ECTEG).

Peer-to-peer learning

The bi-annual peer-to-peer learning events organised by INHOPE were highly appreciated by LEAs using AviaTor. During these events, participants discussed best practices for processing NCMEC Industry reports, the usage of AviaTor and emerging trends and challenges in reports.

We would like to continue these events with the guidance of one or more experienced LEAs to help LEAs that are relatively new to processing NCMEC reports to set up their workflow.

Annotation of images and text and federated learning

In the context of the newly proposed EU legislation to fight child sexual abuse online, there is a need for reliable and unbiased AI classifiers to detect CSAM in previously unseen imagery. To train these classifiers, large and balanced sets of annotated images are needed. This annotation needs to take place in line with what is considered illegal in EU legislation.

We would like to build on the experience gained in the AviaTor project with the annotation of images by the Dutch and Belgian hotline, according to the principles of the Universal Classification Schema, and expand this to all users of AviaTor. The annotation itself could be built into AviaTor's workflow to limit the extra work for LEAs. Federated learning techniques can be used to aggregate the resulting country-specific classifiers to create one universal EU CSAM classifier.

By guaranteeing the sustainability of AviaTor and maintaining a support system around AviaTor, we stay true to our motivation and two main goals set at the start of the project:

To develop an automation and intelligence system that can help prioritise NCMEC reports and reduce the manual labour necessary to process them.

To avoid that the volume of reports leads to situations where reports can only be reviewed superficially or not at all, which might result in urgent and impactful cases being missed and leaving a child in harm's way.

Jos Flury

Project Executive, ZiuZ Visual Intelligence

CHAPTER 09

Meet the Partners

An overview of all the partners involved in this project

THE TEAM

Meet the Partners

For AviaTor to successfully reach its objectives, the project has partnered with leading experts in their respective fields. Despite its small team, there has been quick and remarkable growth in the four years that AviaTor has been active.

AviaTor partnered with ZiuZ Forensic and Web-IQ due to their extensive expertise in visual intelligence and OSINT. These companies have worked together to provide the advanced, cutting-edge AviaTor tool.

Several European LEAs partnered with AviaTor, with the National Police of the Netherlands having led the project since 2019. Additionally, the Belgian Police are a crucial partner within the AviaTor project. Through these partnerships, the LEAs involved have demonstrated their commitment to the cause, contributing greatly through expertise/knowledge sharing and the dedication of resources towards the project.

Furthermore, AviaTor partnered with Timelex and The German Research Centre for Artificial Intelligence (DFKI). Timelex is a highly specialised, legal partner that has contributed both legal and ethical knowledge throughout the duration of the project. The DFKI is the largest, independent AI research centre in the world and has provided technological innovation and expert research to AviaTor. Because of this, AviaTor has expanded its project capabilities and legal framework. Lastly, INHOPE is a key player in the fight against CSAM online by promoting collaboration in providing solutions to online safety issues.

The project's impact and success is largely attributed to our dedicated partners. These companies have leveraged their skill and expertise to provide a strong solution for LEAs all around the globe.

The National Police of the Netherlands

The National Police of the Netherlands (NPN) has been at the forefront of the AviaTor project through their leadership. Since being one of the first LEAs to use the original version of the AviaTor tool, the NPN has been managing the functionality process of the tool. They are responsible for defining, prioritising, and approving functionality, as well as providing the datasets needed to train and enhance the AI of the AviaTor tool. Being one of the two main practitioners of AviaTor project, their role is crucial in providing feedback to the team on shortcomings that LEAs may face, helping strengthen user experience and the tool.

The Belgian Federal Police

The Belgian Federal Police have also been a part of the project since 2019 when they were one of the first LEAs to start using the AviaTor tool. Similar to the NPN, they have been active participants in the project by providing feedback and insights to the team regarding AviaTor tool or user experience development. They are the other main practitioners of the project, alongside the NPN.

ZiuZ Visual Intelligence

ZiuZ Visual Intelligence provides innovative, high-grade visual intelligence solutions to assist in forensic investigations, also aiming to find new technologies to further develop their product and services. Through their product, ZiuZ Visual Intelligence enhances LEAs capabilities to analyse and assess vast visual datasets in child abuse investigations. The company also works closely with universities, NGOs, companies and research institutes.

Web-IQ

Web-IQ provides expertise and knowledge in OSINT technology, offering advanced solutions for important societal issues such as child abuse, human trafficking and fraud. To effectively tackle CSAM, Web-IQ believes in LEAs having access to the best intelligence tools and data.

They are a private sector partner of the Virtual Global Taskforce, which is the international collaboration of LEAs, NGOs, and industry partners that protect the safety of children from sexual exploitation, both online and offline. In AviaTor, they are responsible for enriching reports, building the initial user interface and developing technology for text analysis.

DFKI – The German Research Centre for Artificial Intelligence

DFK specialises in ground-breaking "human-centric AI" research and real-life applications of AI. They strongly focus on the research and application of AI within a societal context, targeting important issues such as climate change or social injustices.

Through their impact, they have initiated, realised and supported efforts towards developing reliable and trustworthy AI for society's benefit. Within AviaTor, DFKI is responsible for offering insights into dataset collection and creation, alongside approaches to machine learning.

Timelex

Founded in 2007, Timelex is a niche law firm leading in the legal aspects of information technology (IT), privacy and data protection (GDPR), intellectual property and media and electronic communications.

Operating from the EU capital, Timelex works with a large network of leading, global law firms in the aim to match law and innovation. They are a top-tier, world-renowned law firm ranked highly by feedback from international law firm rankings, clients and other national and foreign lawyers. Within AviaTor, Timelex offers legal guidance and instruments to the project.

INHOPE

INHOPE is the leading, global network of hotlines in the fight against CSAM, consisting of 53 hotlines in 48 countries (as of April 2024). INHOPE provides the public with an anonymous approach to reporting child abuse material online. These reports are reviewed and classified on illegality by INHOPE-trained content analysts. Reports containing illegal content are then sent to LEAs, and a Notice and Takedown order will be sent to the hosting provider of the content, removing the imagery from the Internet as soon as possible.

Within AviaTor, INHOPE is responsible for the organisation of the project: marketing and communication, website development, campaign creation, organising capacity-building events, and the annual report. Furthermore, they provide relevant feedback from the new working process and tooling from an international perspective.

This project was funded by the European Union's Internal Security Fund – Police

The content of this annual report represents the views of the author only and is their sole responsibility. The European Commission does not accept any responsibility for use that may be made of the information it contains.

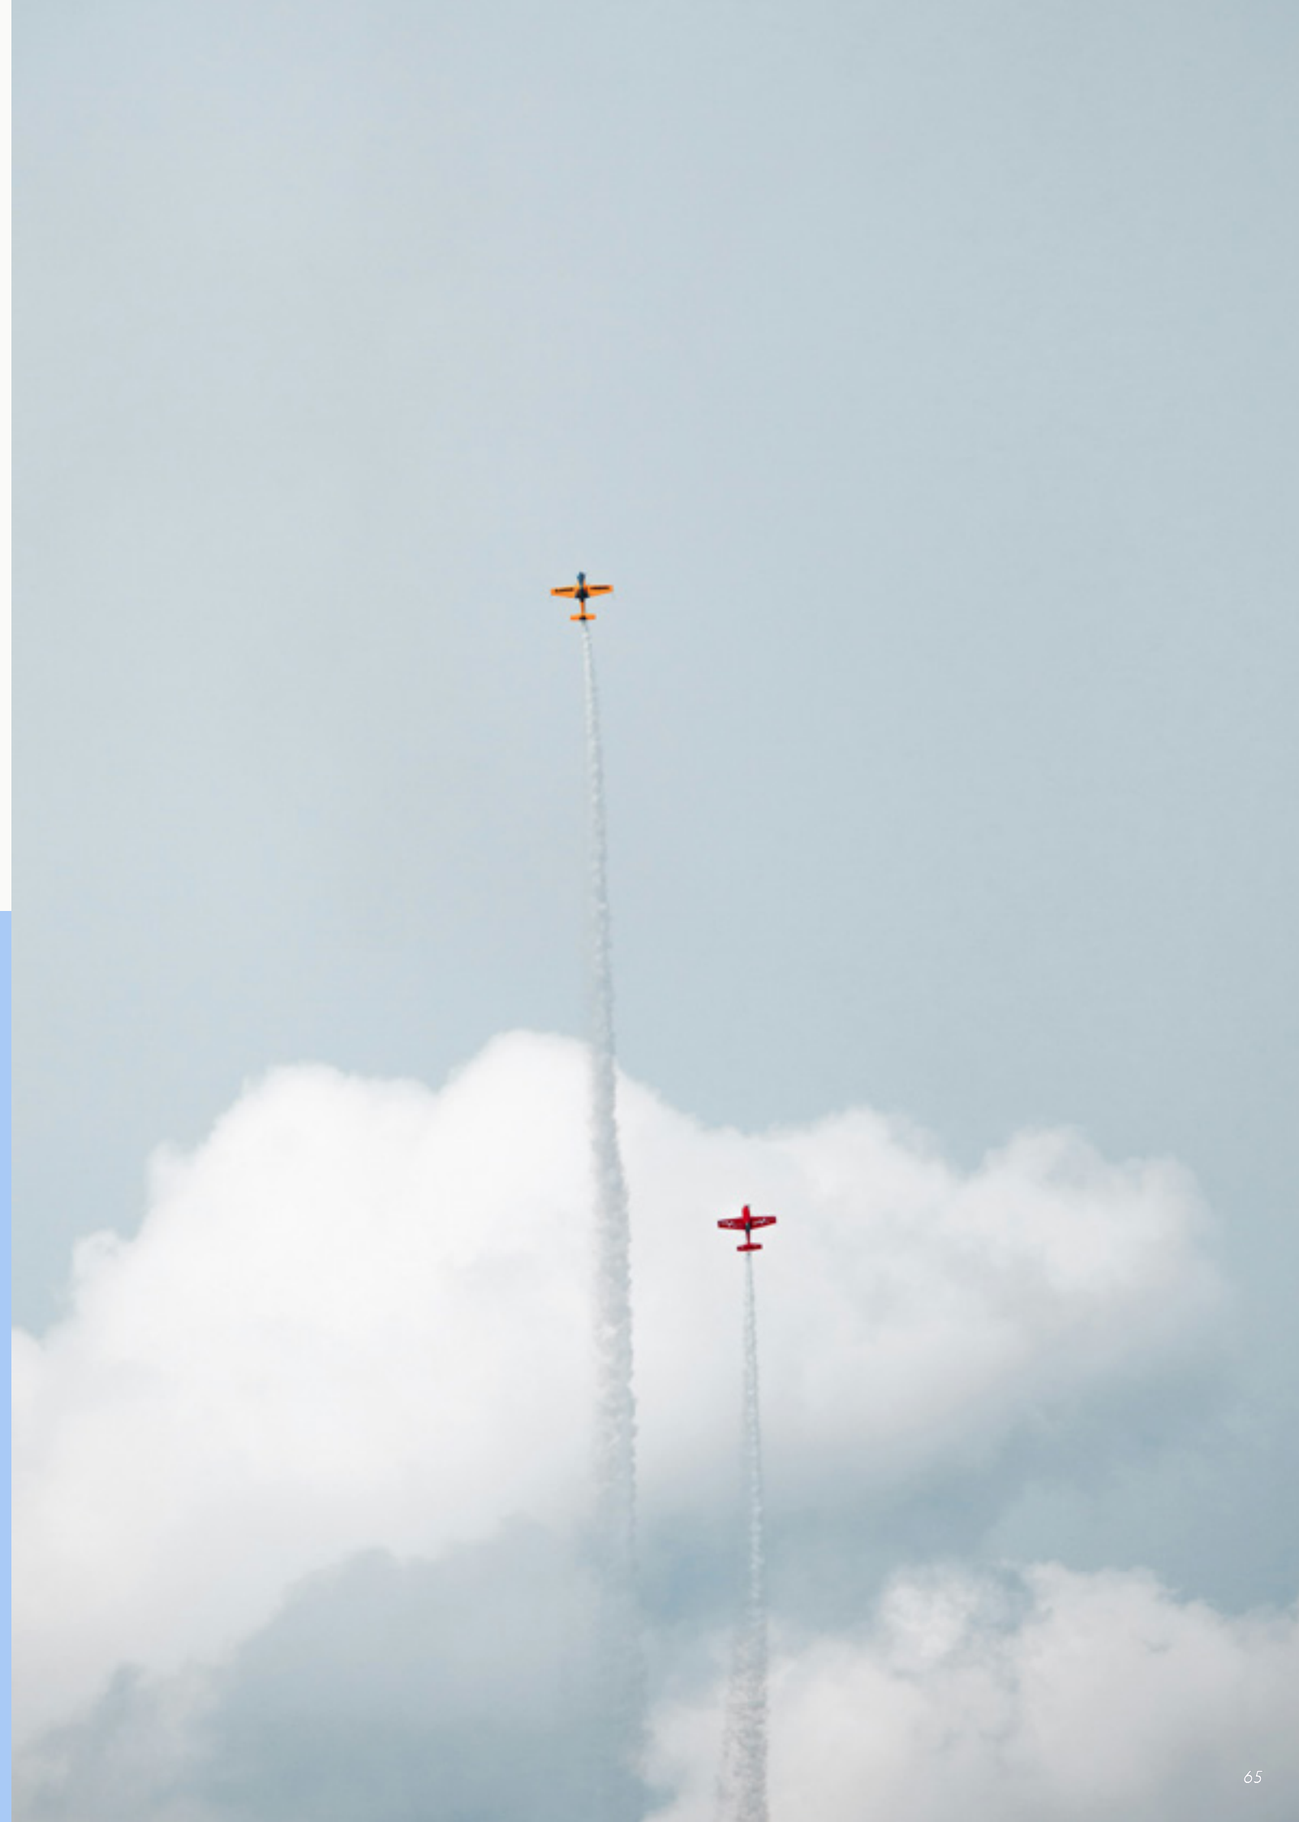

AviaTor

Annual Report of 2023
by INHOPE Association

Learn more and support us at
aviatorproject.com