Case study: Integrating artificial intelligence metadata within Paramount's digital asset management system

Received (in revised form): 12th November, 2020

Dony West
Executive Director, Digital Archive, Paramount Pictures, USA

Dony West is the Executive Director of Paramount Pictures Archive's asset management group, where he oversees all stills, artwork and document management, workflows and preservation efforts for Paramount Theatrical, Paramount Television and Paramount Animation divisions. He has a BFA in still-based media studies from Columbus College of Art & Design and is pursuing an MFA at the Academy of Art San Francisco.

Paramount Pictures, 5555 Melrose Ave, Balaban bld Office #120, Hollywood, CA 90038, USA
Tel: +1 323 956 8380; E-mail: dony_west@paramount.com

Caitlin Denny
Senior Media Archivist, Paramount Pictures, USA

Caitlin Denny is the Senior Media Archivist at Paramount Pictures, where she heads the preservation of photochemical photography collections. She has a master's degree in library and information science from UCLA, and a BFA degree in media arts from California College of the Arts. She has exhibited her film and video work internationally and has curated exhibitions in California, New York and France. She is co-creator of the former online-based art gallery JstChillin, a landmark net-art project that brought critical attention to internet-based artwork.

Paramount Pictures, 5555 Melrose Ave, Balaban bld Office #120, Hollywood, CA 90038, USA
Tel: (323) 956-5160; E-mail: caitlin_denny@paramount.com

Rebecca Ruud
Media Archivist, Paramount Pictures, USA

Rebecca Ruud is a media archivist at Paramount Pictures, where she is responsible for the preservation of born-digital content. Prior to this, she was the head of the archives at Whittier College. She is co-author of 'Images of America: Whittier' (Arcadia Publishing). She obtained her first master's degree in Shakespeare and Renaissance history from the University of Birmingham, UK. She obtained a second master's degree in library and information science from UCLA, where she also received her BA in history with a concentration in museum studies.

Paramount Pictures, 5555 Melrose Ave, Balaban bld Office #120, Hollywood, CA 90038, USA
E-mail: rebecca_ruud@paramount.com

Abstract The Stills Archive team at Paramount Pictures Asset Management Group is responsible for the preservation and digitisation of all photography and artwork created in conjunction with each title produced by the studio. The studio's digital asset management (DAM) system is a customised solution built on a trusted and widely used asset management platform. By partnering with a third-party artificial intelligence (AI) service, the DAM system creates AI metadata to coexist with human-generated metadata. Over the past year, Paramount's Stills Archive team has experimented with this new metadata

technology with the goal of streamlining workflow procedures in preservation and reducing the time spent on research requests. This paper outlines the vision, goals, implementation process and challenges associated with this endeavour.

KEYWORDS: artificial intelligence, computer learning, metadata, photography, celebrity detection, stills archive

INTRODUCTION

Paramount Pictures Corporation (PPC) has been making films for over a century, with such recognisable titles as 'Wings' (1929), 'The Godfather' (1972), 'Raiders of the Lost Ark' (1981), 'Forrest Gump' (1994), 'Braveheart' (1995), franchises like 'Star Trek' and 'Transformers' and many more. What is often overlooked, however, is that for every great movie, there is a talented unit photographer capturing amazing behind-the-scenes images that document the making of the film. These photographers often end up capturing moments that make their way into the marketing campaigns for a film, emphasising their contribution to cinema history.

PPC's photography collection consists of over 11 million individual prints, slides, negatives and born-digital files. Over the past eight years, the studio has invested considerable effort into digitising, processing and ingesting stills into its digital asset management (DAM) system. At the time of writing, the studio has over 3 million fully processed files in the system. The newest version of this DAM system works in collaboration with an artificial intelligence (AI) service platform to create amazing machine-learning initiatives to explore and potentially integrate within current processing and preservation workflows. Paramount is now incorporating some of these offerings, such as celebrity, object and face detection, into its DAM system.

With an ever-growing collection, the need to explore more automated solutions for adding descriptive metadata for enhanced searchability is ongoing. Since 2008, the studio has manually applied metadata and keywords to every production photograph using the IPTC Core metadata schematic. Once ingested into the DAM system, the data is extracted from the attached. xmp file and mapped to correlating fields. Over the past three years, the archive has refined this taxonomy and standardised its controlled vocabulary, but it is still a very manual effort as the team of technicians and librarians process assets at the item level. Recent progress in facial recognition and object detection software is showing massive potential for time-saving results, however, and the team has been actively testing this solution.

BACKGROUND

Paramount began using DAM products in 2008 to store digitised film stills and born-digital photography collections. Over the years, the studio has been on the frontlines of the evolution and implementation of how (when used correctly) metadata and talent 'tagging' can add rich value to media collections. In 2012, a company initiative expanded the database to support Paramount's catalogue titles going back to the company's first feature title, 'Queen Elizabeth' (1912), meaning the studio now had a single system that stored digital files for titles spanning over 100 years.

In late 2018, following a software upgrade to its DAM system, PPC gained the ability to integrate innovative AI technologies, and the in-house IT team began to build AI services into the DAM system. The main objective of this integration was to enhance the already detailed metadata with a focus on object and celebrity detection.

PPC technicians upload images to the DAM system with 17 IPTC Core fields (see Appendix A), with film talent, such as main actors and directors, tagged in the 'Keywords' field. Other data, such as landmarks or objects, however, are left untagged — despite the fact that many image requests that come into the Stills Archive department are for objects or landmarks. To address this, AI can help to fill in the gaps for object detection and certain other useful identifiable information. The AI platform can also capture such data as age and gender, face count, emotion, colour and text detection and safe content judgment.

In early 2020, with the assistance of the Paramount IT team, the archival photography database was consolidated into the marketing photography database to provide a collection of some 4.5 million photographs, artwork, documents and trailers.

The addition of machine-learned metadata does not alter the current photography preservation workflows. After manually entering IPTC Core based metadata, the team uploads images to the PPC instance of its DAM system, and the metadata is automatically mapped to searchable and filterable values. The images are then passed through the AI platform, analysed through their proprietary dataset, and metadata results fed back into the DAM system, populating new fields separate from the IPTC Core metadata. The process runs concurrently with import processing. The PPC IT teams are currently running a backlog of images on the DAM system through the software and new AI-based metadata is showing up in its system daily. At the time of writing, the PPC IT team has completed the AI powered metadata on 355 of 645 titles on its DAM system.

Along with potentially aiding in internal image request queries, the AI-powered metadata provide a great way for internal and external users of the PPC DAM system (of which there are over 1,500) to search for a broader range of terms outside of the system's controlled vocabulary. While the results are not always helpful or conclusive, they have greatly increased search result success when it comes to detecting celebrities (eg Brad Pitt) and objects (eg aeroplane).

Any decision to apply primary descriptive metadata manually usually requires the subject being tagged to have previously negotiated approval/kill rights to a predefined percentage of images containing their likeness. This will include only a small percentage of the cast and possibly any notable producers or film makers. Using AI, it is possible to expand scope to include a wider selection of individuals that might be requested in the future.

TESTING AI CONTENT

The studio's internal and external image requests provided a huge impetus for the team to move forward with integrating AI-powered metadata. Over the years, image requests have varied from 'Get me art for my wall. Anything green' to 'We need images of people laughing' and 'Paul Newman's Cadillac in Hud'. To fulfil these requests, the team has historically relied mostly on personal knowledge of the collections, which at times has led to countless hours of tracking down boxes or looking through tens of thousands of photos. AI-powered metadata can now help point the team in the right direction for these less straightforward image requests. It is by no means a catch-all solution, but it has the potential to get archive staff closer to the answers in less time than before.

The following sections describe some of the challenges and benefits relating to AI-powered metadata with respect to the new fields created by the platform.

Celebrity detection and caption text generator

One of the most anticipated features of the AI tool is celebrity detection. If this tool

could identify at least every main talent and director in photographs, this would cut the digital processing time by 50 per cent. Celebrity detection is used to create metadata identifying actors, crew members, politicians, athletes, historical figures, fictional characters and even internet celebrities. It works by detecting a face within any ingested asset and comparing it with the AI dataset of faces considered celebrities. The term 'celebrity' is of course subjective, meaning some of the results have been mixed. In this respect, the fact that the dataset pushing the AI metadata is human-driven cannot be overstated.

When the team began working with the AI service, its celebrity dataset included 100,000 faces. At the time of writing, the dataset contains well over 1 million faces. The manual IPTC Core keywording of talent is mapped into the DAM system field 'Keywords', while metadata created by AI through the celebrity detection tool manifest in a separate 'Celebrity' field.

By using separate fields, the team can compare the AI-powered metadata against the metadata provided manually. A great example of the difference in keyword results can be seen by doing an 'all content' search and comparing numbers of images tagged with Abe Vigoda, the beloved actor from 'The Godfather' franchise.

The manually tagged metadata resulted in 119 images of the actor; in contrast, the AI-powered metadata fetched only 22 images from the same set of photographs. To see the difference between manual and AI-powered metadata on Paramount's DAM system, see Appendices B and C.

To illustrate this contrast better, the team used a small subset of 19 AI-powered values within the 'Celebrity' field, and retrieved the following statistics from this subset:

- 16 per cent of the AI 'Celebrity' tags were duplicates of the manually added 'Keywords' tags;
- 32 per cent of the AI 'Celebrity' tags were new and correct; and
- 52 per cent of the AI 'Celebrity' tags were new and incorrect.

Although the success rate is not outstanding, it is promising. Most of the new and correct 'Celebrity' tags created by the AI platform were of crew members on set and political figures at events — people that the librarians would not normally tag in the 'Keywords' field. For example, the AI tagged New York City Mayor Abraham Beame on the set of 'Three Days of the Condor' (1975) — someone who had not previously been manually tagged. The AI celebrity detection was also able to tag the same celebrity across their lifetime, recognising Gloria Swanson from her 1924 role in 'Manhandled' as well as in the 1950 film 'Sunset Boulevard'.

Still, there are lesser known talent or non-starring actors that rely on manual tagging based on the team's personal entertainment and film knowledge. At the end of the day, the Stills Archive is part of a business and has deadlines to meet. Because of this, the team does not have the time to tag most crew members. As information professionals, though, they understand the power of metadata and try their hardest to tag women and people of colour on film sets. A Stills Archive team favourite 'Keywords' tag is of Catherine Coulson, most recognisable as the Log Lady in David Lynch's 'Twin Peaks' television series, showing up as First Assistant Cameraperson on 'Star Trek II: The Wrath of Khan'. This is someone who would not normally be listed in the set of talent to tag, but the team's personal knowledge was used to identify and tag her. She was not tagged by the AI platform.

After observing the capabilities of the celebrity detection AI tool, the team found a perfect project to make the best use of it. Over the past four years, the Stills Archive team has collected over 200 orphaned photographs — prints, negatives and transparencies of unknown production titles and unidentified talent. The PPC librarians decided to upload these digitised orphaned

images to the DAM system — something that would not normally have been done prior to the integration of the AI platform. The intention in doing this was to retrieve AI-powered 'Celebrity' values in order to resolve the mystery photographs. In the absence of an AI-powered title detection tool to help identify the orphaned images, the celebrity detection tool must suffice.

Prior to ingesting the orphaned photographs into the studio's DAM system, the team decided to pool the depths of the hardworking Paramount archivists' knowledge. On a processing table and light box, the team laid out all the orphaned photographs and had a communal identification session. Almost 50 per cent of the images were correctly identified by the PPC archival staff. The last 115 of the orphaned images were then scanned, digitally processed and uploaded into the PPC DAM system. Of these 115 ingested images, six images were successfully identified, two images were incorrectly identified and 107 were unidentified images (possibly due to factors beyond the dataset, such as dim lighting and obscured faces). This result represents a 5.2 per cent successful identification rate for the celebrity detection tool.

There was one more tool discovered in this process as useful for identifying talent in images — the AI-powered 'Caption Text' field. This field generates a single-sentence description of what the AI platform defines as the content of the photograph. Why this separate tool can identify names while the celebrity detection tool misses them is unknown at this time. Using this option, four more orphaned images were successfully identified.

AI-powered celebrity detection cannot stop at the retrieval of a name. This project proved what had already been assumed — that archivists give context to machine-learning processes through research and knowledge. To realise its full potential in the archive, AI must be supported by information professionals of all kinds.

Object detection

Perhaps the most used category behind celebrity detection is object detection. Most of the image requests, if not for a specific film or star, are for objects in films. These requests have ranged from dogs, to helicopters in a specific title, cars and even food items.

Using the built-in search bar modelled after search engines, researchers can keyword search for such objects across the entire PPC DAM system. For example, in a request for a specific model of car, the team was able to type 'car', 'automobile' and other synonyms into the search bar; most of the results that were returned were indeed cars, but the query also retrieved images of horses, carriages and various images of people sitting on something.

Likewise, when a publisher requested images of actors with cocktails for a drink recipe book, the team searched for the term 'cocktail', but most of the results were of cocktail dresses. More successful results were obtained using the terms 'glass' and 'bar'.

The AI platform also tags landmarks. The AI landmark recognition tool can tag everyday landmarks like 'school', 'theater' or 'airport'. It also recognises miniature models, like the airport model used in the film 'Airplane!'. Specific landmarks like the Eiffel Tower or Empire State Building are detected as 'skyscraper', 'building' and 'tower' in the 'Tag Name' field but are identified by their formal name in the 'Caption Text' field. The AI-powered metadata generated using the object detection tool populates two fields: 'Object Detection' and 'Caption Text'. See the 'Tags' field in Appendix C for an example.

To get some statistics behind the findings, PPC librarians ran a keyword search for 'cat' across assets belonging to the title 'Breakfast

at Tiffany's', in order to find images of Orangey the cat — the feline star of both television and feature films. The results came back with nine images. When all photos from the title were examined, however, a total of 57 images of Orangey were identified, indicating a success rate of 6 per cent for the AI object detection tool. The eight successful AI retrieved images (one result was of a sweater) were all unobscured photographs of the cat with few actors in the shot. The other 49 images found without the object detection tool were of the cat from the side, with other animals or in motion.

A similar search for 'Eiffel Tower' yielded images of the actual tower, images of a long dinner table from above, a forest fire, a long window shot from below and even the Television Tower in Berlin. The PPC title with the most photographs containing the Eiffel Tower is 'Mission: Impossible — Fallout'. Searching across all assets within this title, the AI platform returned 171 images of the landmark, while human-powered identification found 754. This yields a 4 per cent success rate for the AI object detection tool. Again, the AI found great images, but it was far from comprehensive.

When retrieving results from the AI platform, it is thus necessary to examine the preceding and successive images surrounding the correct AI results.

The object detection tool also includes the field 'Safe Content Type', which designates an image as suggestive, adult or undefined. As displayed in Appendix C, the 'Safe Content Type' field provides an excellent example of the subjective nature of AI. Examples of what AI thinks might be suggestive or adult include violence and blood, people kissing and nudity (perceived or actual). Generally, the AI's confidence level (a percentage value assigned by the AI platform to rate its confidence in its data) on the 'Safe Content Type' field is below 80 per cent, meaning that it cannot be used in any effective manner.

Finally, the object detection tool is used to find colour in images, manifesting in the 'Color' field, as seen in Appendix C. Although these are generally accurate, they are rarely used. When an internal requestor was searching for photographs of special effects during production, the PPC team keyword searched for 'green' to find photographs containing green screens, a tool often used in filming special effects. This feature is also useful for finding specific costumes in films. If one keyword searches for 'yellow', the results might include Cher's iconic plaid skirt suit in 'Clueless' or even the character Bumblebee of 'Transformers' fame.

The AI-powered object detection tool has proven very effective for locating 'things' in studio's photograph collection. It helps to narrow down which titles to explore more closely and points the research in the right direction, all while dramatically cutting the time spent searching.

Face detection

The AI platform also analyses human faces in photographs to detect age, gender and face count. As with the object detection tool, the results from this analysis are helpful but ultimately not comprehensive. Age and gender analysis have been shown to be subjective, but sometimes useful in a broader sense. Age analysis, manifested in the PPC DAM system's 'Age' field, works well when looking for either children or adults (see Appendix C). In age-based searches, the filtering option helps to narrow down results. To use this feature, the team starts by conducting a search of every asset in the photograph collection, and from there filtering within the 'Age' field to restrict the search to a specific age range, such as 1–11 or 55+. The face detection tool does not accurately differentiate between a person in their 30s or 50s, but it works for well for extreme age differences.

Gender analysis has similar issues with AI platform's lack of nuance in identification. There is clearly a gender bias present in the data, but the data is generally accurate (see Appendix C). Males with longer hair and male actors of the 1920s, who often wore heavy makeup on set, often get misidentified as females. The same is true for people costumed against gender stereotypes. In this next example, the team used photographs from two films where male actors were dressed in feminine attire. The AI-powered metadata tagged the actors as female in one of the films and as male in the other. In an opposite scenario, women dressed in masculine clothing did not seem to confuse the AI platform as they are generally tagged as women. In a time of gender fluidity, trans visibility and better representation of the queer community in entertainment, this tool will become either passé or completely inaccurate. The validity of this tool will be interesting to track moving forward as films increasingly become more diverse.

Where face recognition can be helpful is with the face count tool. Although not always correct, face count has assisted in finding groups of people for image requests (see Appendix C). If someone is looking for a crowd scene or even a scene with just two people, filtering the collection by face count helps narrow down the search parameters. In a real-life example, a franchise production was looking for an image of the entire cast from one of its previous films. At the time, PPC did not have AI so the Stills Archive team had to track down boxes and look for an image of the entire cast. Had the archive team had AI back then, they could simply have searched for a face count of seven. This would have cut the workload down significantly. Another use for this tool is when there is a need to find images without humans. In these cases, the researcher can filter the collection by entering 0 in the 'Face Count' field.

Facial recognition, while not perfect, does help to narrow down results for choices between two options: male or female presenting, children or adults and multiple faces or no faces.

Text detection

The final key tool within the AI platform is the text detection tool. The preferred choice in AI integration uses optical character recognition (OCR) to analyse text. This is very good at recognising simple typed fonts but is not accurate with hand-created fonts or handwriting. Indeed, images with handwriting or sign painting are almost always incorrect or do not even register as text. The hope is that this will eventually improve as the archive has many handwritten notes on documents that would be valuable to have searchable.

The text detection tool can even make a guess at the language used in an identified text. As an example, a photograph of director chairs with names printed on the seat cloth was successfully identified and transcribed by the text detection tool. This tool, manifested in the 'Text on Image' and 'Image Text Language' fields, is hardly ever used by the Stills Archive team (see Appendix C).

Overall, the AI face detection, object detection and text detection tools are useful for getting an archivist closer to finding the answers they are looking for. AI tools can help find images of the set by searching for zero faces, objects for an image request, or even a green picture to hang over someone's couch; however, these tools are not sufficiently accurate to forgo human interaction. Archivists are crucial to the quality control of the AI platform's findings and help to fill in the gaps when the AI is inaccurate.

Summary

Overall, the AI face detection, object detection and text detection tools are useful for getting an archivist closer to finding the answers they are looking for. AI tools can help

find images of the set by searching for zero faces, objects for an image request, or even a green picture to hang over someone's couch; however, these tools are not sufficiently accurate to forgo human interaction. Archivists are crucial to the quality control of the AI platform's findings and help to fill in the gaps when the AI is inaccurate.

LESSONS LEARNED AND A LOOK TO THE FUTURE

As with any AI project, the base dataset module is of extreme interest. This forms the core knowledge for the AI system, determining the usefulness and accuracy of the platform. Because employees have created a servicing AI platform, however, there are inherent biases present in the dataset. For example, a woman wearing a white dress in one photograph has been tagged by the AI as a 'bride'. Given that brides do not wear white in many non-Western cultures, this highlights a potential Eurocentric bias in the dataset. Indeed, the AI has returned countless results that pivot around a Western cultural viewpoint, often reinforcing a patriarchal structure.

Unfortunately, neither the Stills Archive nor its internal IT teams have access to the AI dataset or glossaries. However, were it possible to develop a way to let users interact with the dataset, a crowd-sourced AI correction project could help diversify the metadata retrieved. This is the kind of solution utilised by Microsoft, which regularly collects information based on client image collections in order to refine its own dataset.

The film 'Nacho Libre' has uncovered numerous issues with the software. For example, when searching all assets for the film for the keyword 'luchador', the search failed to return any results. When searching for the English word 'wrestler', however, the system returned many correct results. Interestingly, one of these images was also tagged as 'Sanshou', a form of Chinese kickboxing. This shows that the dataset has a wide variety of sports and sports language — just not 'luchador'.

When examining 'Nacho Libre' for text detection results, the team found an image of Spanish language printed on a wall that was not detected. This could be because it was in a language other than English or because it was painted on the wall — for now, the reason remains a mystery.

As for the face detection tool, while the facial recognition software is getting more accurate as more images are fed through the machine-learning process, a problem arises when actors don masks or special effects makeup. For example, when Jack Black was in his luchador mask, the AI completely failed to identify him. This presents a real problem for science fiction or horror makeup.

A final lesson learned is that the AI is not accurate when the subject of the photograph is in profile. The AI can usually not detect the celebrity, give an age or gender or generate an accurate face count. Because so many of photographs in the archive depict people in some degree of profile, this issue would appear to account for the underwhelming percentage of successful celebrity and object detection using the AI platform. Without an interactive user correction method, AI metadata search results may not get much better than they are right now.

Without the ability to change the dataset and add terminology like 'luchador', some archival collections may find the software less useful. To make this a better functioning system, the software needs to learn from the manual metadata as well as human users learn from the AI metadata. An additional issue with the AI integration into the DAM system is that the team cannot edit the metadata within the system. If a woman is tagged as a man or an actor is tagged with an incorrect name, librarians currently cannot change the metadata. This leads to inaccuracies when users search the PPC collections in its DAM system. Because of these inaccuracies, the

metadata generated by the AI platform is currently viewable only by internal Paramount employees with DAM-related administration roles. An additional stop-loss measure was to separate the AI-powered fields from the functional IPTC Core metadata.

In addition to the PPC stills collection, the archive team is in the process of expanding AI integration to business units and colleagues within Paramount's Asset Management group. Audio, trailers, documents, artwork and even low-resolution proxy files of video masters have proven interesting in early-stage results. It is truly powerful and exciting to see a video run through these same AI-powered features described above. Although still developmental, there are some reliable time-saving tools that can assist with searching across video content. AI functionality integrated with the PPC moving image collection can aid in searching for key talent or specific dialogue. Third-party AI tools can find the talent or the line of dialogue and highlight the timecode on the video with surprisingly accurate results. If results continue to be positive, archivists may one day be able to search confidently through entire films without having to watch each scene. In the future, the team hope to do more analysis on the speech-to-text functionality of the product.

CONCLUSION

As the team continues to explore the possibilities of AI integration, juxtaposed with human-created metadata, the mood remains hopeful but sceptical. While the software capabilities will evolve as the technology improves, archivist oversight will remain essential for the foreseeable future. To harness the strengths of both archivist and AI services, software companies (that provide AI solutions) should begin to incorporate user feedback regarding their glossaries and datasets. Numerous software and technology companies offer solutions for full reliance on AI — and all have notable strengths and weaknesses in the technical areas described. The best option might be a hybrid solution that combines multiple AI services in conjunction with human oversight. As with most successful and long-term solutions, the call for human–computer partnership is pivotal for success. Within the PPC stills team, the consensus is that one cannot rely fully on machine learning — yet. AI is not the panacea. Archivists must take the information AI provides and provide context through research.

Appendix A: IPTC core fields

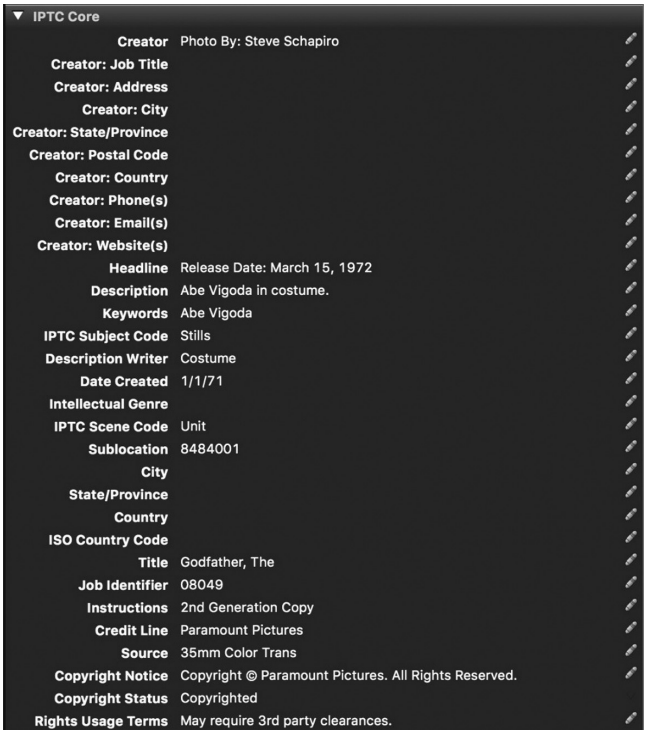

IPTC Core	
Creator	Photo By: Steve Schapiro
Creator: Job Title	
Creator: Address	
Creator: City	
Creator: State/Province	
Creator: Postal Code	
Creator: Country	
Creator: Phone(s)	
Creator: Email(s)	
Creator: Website(s)	
Headline	Release Date: March 15, 1972
Description	Abe Vigoda in costume.
Keywords	Abe Vigoda
IPTC Subject Code	Stills
Description Writer	Costume
Date Created	1/1/71
Intellectual Genre	
IPTC Scene Code	Unit
Sublocation	8484001
City	
State/Province	
Country	
ISO Country Code	
Title	Godfather, The
Job Identifier	08049
Instructions	2nd Generation Copy
Credit Line	Paramount Pictures
Source	35mm Color Trans
Copyright Notice	Copyright © Paramount Pictures. All Rights Reserved.
Copyright Status	Copyrighted
Rights Usage Terms	May require 3rd party clearances.

Appendix B: Manual metadata

Embedded Camera Data

Pixel Dimension X	6415	Pixel Dimension Y	4184
Camera Orientation	Horizontal	Resolution X	600
		Resolution Y	600
Color Space	Uncalibrated		

Image Info

Creator	Photo By: Steve Schapiro	Image Date Created	03/12/2019
Scene	Unit	Image Location	8484001
Content Headline	Release Date: March 15, 1972	Category	Stills
Title	Godfather, The	Production Number	08049
Condition Issues	No Visible Defect	Credit Provider	Paramount Pictures
Original	35mm Color Trans	Rights Usage Terms	May require 3rd party clearances.
Keywords	Marlon Brando,Abe Vigoda,	Copyright Notice	Copyright © Paramount Pictures. All Rights Reserved.
Artist Photographer	Photo By: Steve Schapiro	FileSize	62 MB
		FileName	Godfather_Still_00432.tif
Color Mode	RGB	Release Year	1972

Appendix C: AI metadata

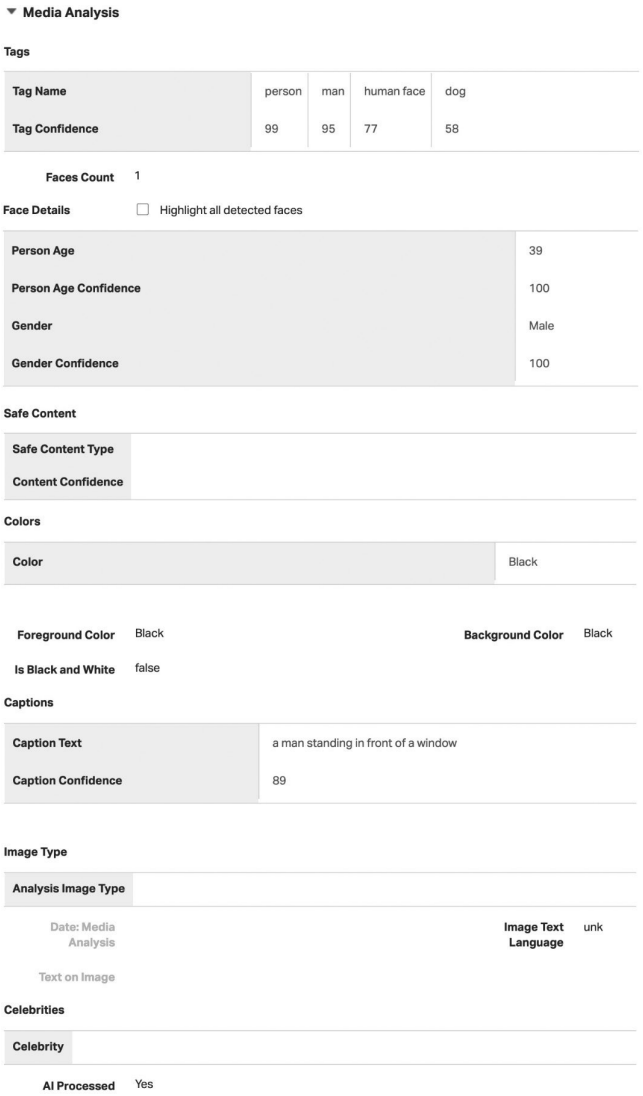

▼ Media Analysis

Tags

Tag Name	person	man	human face	dog
Tag Confidence	99	95	77	58

Faces Count 1

Face Details ☐ Highlight all detected faces

Person Age	39
Person Age Confidence	100
Gender	Male
Gender Confidence	100

Safe Content

Safe Content Type	
Content Confidence	

Colors

Color	Black

Foreground Color Black — Background Color Black

Is Black and White false

Captions

Caption Text	a man standing in front of a window
Caption Confidence	89

Image Type

Analysis Image Type	

Date: Media Analysis — Image Text Language unk

Text on Image

Celebrities

Celebrity	

AI Processed Yes