January 02, 2019

# Using machine learning to improve image storage and recall

Mary Witkowski

Kristina Bohl

Raja Ratna Murthy Ayyagar

# Using machine learning to improve image storage and recall


## ABSTRACT

Photographs taken using mobile devices are uploaded to image storage and sharing services at a resolution, server location, and upload priority that is blind to the importance of the photo to the user. This is sub-optimal, as there are images of great importance to a user (e.g., wedding photos) that are better uploaded at high priority and resolution. Similarly, there are images (e.g., screenshots, photos of receipts, etc.) that are not very important and can be uploaded at lower priority and resolution.

With user permission, the techniques of this disclosure use machine learning to predict the importance of an image captured by a user. The predicted importance is used to select a resolution, server location, and upload priority for the image in an online image-sharing service.



## KEYWORDS

- image storage
- photo backup
- photo library
- photo saliency
- cloud backup
- image retrieval
- semantic analysis
- image metadata
- machine learning

BACKGROUND

Photographs taken using mobile devices are uploaded to image storage and sharing services at a resolution, server location, and upload priority that is blind to the importance of the photo to the user. This is sub-optimal, as there are images of great importance to a user (e.g., wedding photos) that are better uploaded at high priority and resolution. Similarly, there are images (e.g., screenshots, photos of receipts, etc.) that are not very important and can be uploaded at lower priority and resolution.

DESCRIPTION

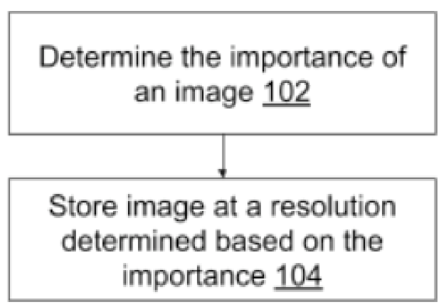


**Fig. 1: Storing an image at a resolution appropriate to its importance to the user**

Fig. 1 illustrates storing an image a resolution appropriate to its importance to the user, per techniques of the disclosure. The importance of an image is determined using machine learning (102), as further explained below. The image is stored (104) at a resolution appropriate to its importance. Storage resolution can manifest in the following ways.

High-value images

- **Framing:** High-value images include images of user experiences and memories that a user is likely to share, print, or view again later, e.g., wedding photos, a baby’s first steps, photos of the user with her friends at an event, etc.

- **Treatment:** An image that is determined to be of importance is stored at high resolution, possibly higher than the resolution indicated by a user setting, e.g., a default setting. This enables the user to print the photograph or project it on a screen.
- **Cost:** The additional cost for storing photos at high resolution can be offset by promoting the printing a photo album, e.g., in large-format prints. Additionally, the cost of high-resolution storage can be offset by low-resolution storage of photos of lesser importance.

Low-value images

- **Framing:** Images determined to be of relatively low importance or low complexity don't warrant full-resolution storage. Examples of such images include, e.g., functional photos with short shelf-life, e.g., images of receipts, event invitations that expire with the passage of the event, etc.
- **Treatment:** With user permission, images determined to be of relatively low importance are stored at lower resolution. Simple images, e.g., images that are not busy, have few features, or have large uniform areas are losslessly compressed.
- **Cost:** The reduced cost of storage of images of lesser importance is a benefit to both image-sharing service and the end-user. The user consumes less of a (typically) limited storage quota for these images of lesser importance.

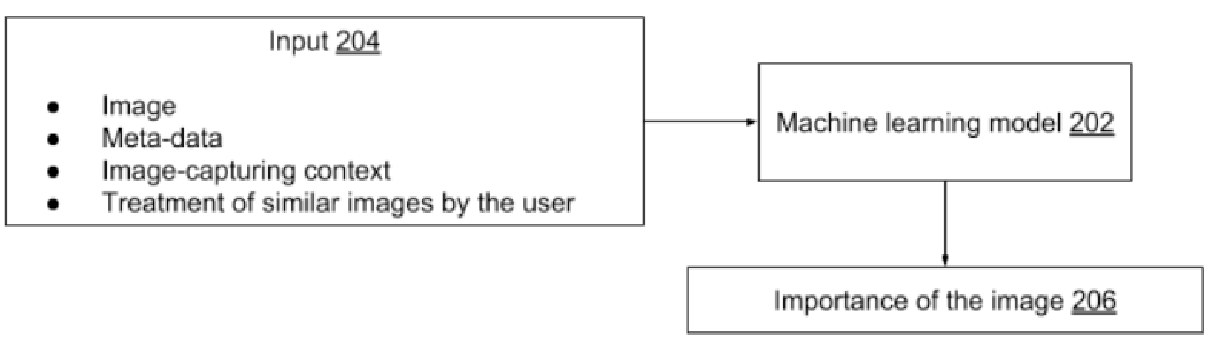


**Fig. 2: Determining the importance of an image**

Fig. 2 illustrates determining the importance of an image using a machine learning model. With user permission, a trained machine learning model (202) accepts as input (204) the following features: the image, metadata of the image, the image-capturing context, past treatment of similar images by the user, etc. The machine learning model performs various operations such as object detection, scene classification, semantic analysis, etc.

For example, an image of a wedding may be classified as such by the presence of the bride, the groom, the cake, etc. An image of a receipt or a screenshot may be classified as such using scene classification techniques. Image-capturing context includes, e.g., the location and direction of the camera, timing between image captures, etc. For example, a series of images captured in rapid succession may be indicative of an event, a vacation, etc. Past treatment of similar images by the user is used to inform the importance of the image, e.g., if the user viewed, accessed, or shared similar images in the past, then such images are deemed as important. To summarize, features such as the image, its metadata, the image-capturing context, past treatment of similar images, etc. are fed as input to the machine learning model that provides as output the predicted importance of the image (206) to the user.

Alternatively, or in addition to machine learning, heuristics can be used to predict the importance of an image. In this manner, the techniques of this disclosure utilize user-permitted factors about image importance to infer optimal storage footprint. The techniques can also be used to identify one or more images suitable for printing.

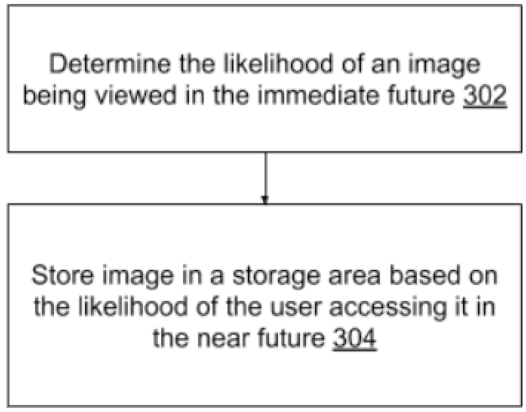


**Fig. 3: Storing an image in an area appropriate to the likelihood of the user accessing it in the near future**

Fig. 3 illustrates storing an image in a storage area based on the likelihood of being viewed or shared by the user in the near future, per techniques of this disclosure. The likelihood of an image being viewed in the immediate future is determined (302), e.g., by machine learning techniques, as further explained below. The image is then stored (304) in an area appropriate to its likelihood of being viewed, accessed, or shared by the user in the near future.

For example, if the image is of a wedding or other event, it is likely to be viewed and/or shared in the immediate future (e.g., upon capture, or upon upload), e.g., within the next forty-eight hours. Such images are either stored locally for quicker access, or in non-local (cloud-based) storage that enables fast serving and recall. As another example, if the image is that of an event flyer whose date (as indicated on the flyer) has passed, then the image is archived, e.g., at a storage location that is cheaper or farther away from the user, since there is low likelihood of such an image being viewed in the near future.

The likelihood of an image being viewed or accessed by the user in the near future is determined by machine learning techniques in a manner similar to that shown in Fig. 2. In particular, the trained machine learning model detects dates imprinted on the image that can be indicative of expiry. For example, an image of an event flyer may have date that has passed; an image of a receipt may have a date that is more than, e.g., sixty days, the time within which a receipt is typically submitted as part of an expense report; etc.

The machine learning model analyzes available information regarding the image to determine whether the user may access the image soon, e.g., to share the image. Information that can be used includes, e.g., an image that includes a date that has not yet passed; an image that includes a date that has passed but is still of relevance, e.g., an image of a receipt that is within a sixty-day window for submitting as part of an expense report; an image that is suggested as part of a "remember this day" event; etc. Some factors may indicate that the user may not access the image soon, e.g., an image that matches archival criteria (e.g., it is a document, receipt, screenshot, etc.); an image that includes content that is not typically shared by users; etc.

In this manner, the techniques of this disclosure enable storage of images in an optimal storage area such that images of importance or immediate relevance are quickly retrievable. Additionally, by archiving images that are of lesser immediate relevance, storage costs are reduced.

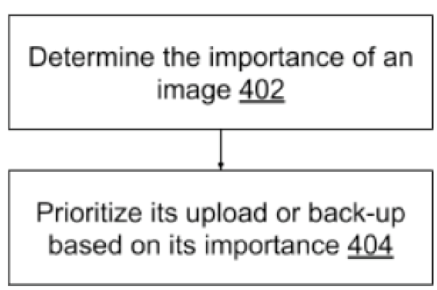


**Fig. 4: Determining position of an image in the backup queue**

Fig. 4 illustrates prioritizing the backup (e.g., by uploading to an image-sharing service) an image based on its importance to the user, per techniques of this disclosure. The backup or upload queue order is determined based on the importance of images. The importance of an image is determined (402) using machine learning techniques similar to those described in Fig. 2. Images that are identified as important are backed-up (404) at a relatively higher priority, e.g., placed at the top of a queue of images to be backed up. Images of lesser importance occupy other, later positions in the queue.

For example, consider a user that takes a photo of a receipt earlier in a day and then a photo of her daughter at her graduation ceremony. The graduation ceremony image, being of greater importance, is assigned higher priority in the backup queue, even though it was taken later than the image of the receipt. In this manner, important images are made available sooner for consumption in search, creation, sharing, or viewing on secondary devices, thereby improving user experience.

Further to the descriptions above, a user may be provided with controls allowing the user to make an election as to both if and when systems, programs or features described herein may enable collection of user information (e.g., information about a user’s social network, social actions or activities, profession, a user’s preferences, or a user’s current location), and if the user is sent content or communications from a server. In addition, certain data may be treated in one or more ways before it is stored or used, so that personally identifiable information is removed. For example, a user’s identity may be treated so that no personally identifiable information can be determined for the user, or a user’s geographic location may be generalized where location information is obtained (such as to a city, ZIP code, or state level), so that a particular location of a user cannot be determined. Thus, the user may have control over what information is collected about the user, how that information is used, and what information is provided to the user.

## CONCLUSION

With user permission, the techniques of this disclosure use machine learning to predict the importance of an image captured by a user. The predicted importance is used to select a resolution, server location, and upload priority for the image in an online image-sharing service.