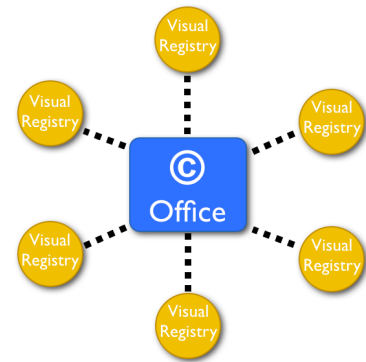


Competitive Visual Registries for Copyrighted Works

Creating a visual registry system for copyrighted works will be essential to reducing or eliminating the problem of orphan works in the future. The updates currently being made to the U.S. Copyright Office's publicly accessible registry were not designed to address this particular problem, nor does the Copyright Office have expertise in visual recognition technology. Budgetary requirements for the Office to create a visual registry on its own also may be politically infeasible. As an alternative, private entities could create their own visual registries, subject to the accreditation of the Copyright Office.

The Model: A Copyright Office Visual Registry Accreditation System

The Domain Name System, or DNS, provides a useful model for a visual registration system; the Copyright Office should consider modeling a system of independent visual registration databases on a similar accreditation-type model. The market for registering a domain name is a competitive one. Registrars are independent companies which are accredited through Internet Corporation for Assigned Names and Numbers, or ICANN, to register individuals' and companies' domain names. Each registry company shares the information with ICANN and other registry companies.



This proposal would bestow the responsibility upon the U.S. Copyright Office to accredit Visual Registries to act as a databases and search engines for visual works. Our proposal for the accreditation agreement would include the following:

- Each Visual Registry would be permitted to advertise that it is accredited by the Copyright Office.
- Copyright owners need only register a work once with the Copyright Office. A Visual Registry must provide the option during the visual registration of a work to associate the visual registration with a pre-existing official copyright registration, or make an official registration.

- The Copyright Office will specify what data the Visual Registry is required to collect with each visual work.
- A Visual Registry will allow free public and machine readable searches of its entire database.
- A Visual Registry must have a robust search engine to enable the public to determine the copyright owner for a given visual work if that work is located in the Visual Registry's database. Such a service must include a machine readable digital image recognition search engine where a user can provide an image and the search engine can search for matches within the database to determine the potential copyright owner.
- Public access to registry images for purposes of conducting a visual search shall only include the ability to view low-resolution images with a visual watermark bearing the copyright owner's name and basic contact information, and, if available, copyright registration number.
- The visual works owners would maintain all copyright in their own registered works.
- A Visual Registry must maintain the registry data as long as it remains accredited, and when accreditation expires or is terminated, the full registry data must be transferred to another accredited visual registry or the Copyright Office.
- A Visual Registry must require the registrants to agree to provide accurate information under penalty of law, that the registrant must to the best of his knowledge be the copyright owner of the work and that the work does not infringe on anyone else's copyright.
- A Visual Registry may provide additional value-added services to registrants, to differentiate itself from other Visual Registries. While a Visual Registry may offer these services to off-set the cost of the official copyright registration fee to the copyright owner, the fee that the Copyright Office sets for registration must be paid in full to the Copyright Office.

To address the issue of older works that have already been registered, but have not been entered into a visual registry to be indexed, legislation would be needed to allow competitive registries equal access to digitize the Copyright Office's copyright deposits.

Once the domain name registration system was opened to competition, dozens of registries have entered the market. The benefits to the public have been clear - a registration that once cost \$40 per name can be had for as little as \$6.95. There is every reason to think that a competitive visual registration system will flourish as well.

The State of Visual Recognition Technology

The current state of online search for registered copyrights of visual works is text-based: searchers can try to find works by author, title, copyright claimant, registration number, etc. An improved system set to debut in the near future allows searching of keywords associated with a copyright record. There are at least two problems with the Copyright Office's current and near future system: 1. keyword searches may often fail to retrieve the sought-after image; 2. the results of a search of the Copyright Office's registry only provides basic text description of a work, not an actual visual display of found records, which makes it all the more difficult for a searcher to discover the desired work.

Even if the Copyright Office were to open its registry database to indexing services like the image searches provided on the web by Google and Yahoo!, a searcher would have similar results because each is only searchable using text. This limitation is the result of the lack of information in the Copyright Office's database and how the search engines are engineered. Generally, when a searcher types in 'elephant' into Google's Image Search, it will return a page with images of elephants, and other seemingly random occurrences. The undesired results occur because the search engine examines the meta-tags (short descriptions written by the person who uploaded the image that tell the search engine what the picture holds) associated with the image, the words surrounding the image on the webpage, and any other information the search engine can gather from the context of the web page. The search engine does not directly analyze the visual characteristics of the image.

A text-based visual search is perfectly adequate when searching for general images of a specific topic, however, if a searcher require a higher level of accuracy when searching for a specific image with specific characteristics, a text-based search may not be his best option. For instance, if a searcher is a museum curator looking for a work of art similar in style to something in your gallery, a text-based search is not going to be of much use – most descriptions or meta-tags contain a few words of description, not the precise details she would need to locate a similar work. If a photographer is looking for unlicensed use of his work, there will be few helpful text descriptions – an infringer is not going to help the copyright owner find his or her illegal activity.

Shape, texture, color, brightness, and proportions define images, and while text can limit the scope of the search, the image's inherent traits are necessary to find similar items. A visual registry uses these specific traits to locate similar items. Instead of typing in text, a searcher would upload a comparative image into the visual registry. That image would be analyzed and used to find similar works.

Visual registries and content-based image retrieval systems are needed.

The growth of the World Wide Web coupled with broadband connections has led to the increase in the use of images online by users for both professional and private uses. The world is going online, one picture at a time: digital cameras abounding, online photo-albums becoming the defacto way to store pictures, and museums and artists realizing the benefits of putting their collections online. However, this process of digitization comes with certain growing pains – namely, how to organize those images into a system that allows them to be searched and accessed accurately. Content-based image retrieval systems are the solution to this problem.

How the technology works.

A content-based image retrieval (CBIR) system generally uses image processing, image analysis, and computer vision technology to compare features of the image, itself as the basis for the search. The user uploads an example, either a sketch or an image to the system. A basic CBIR system will then use an algorithm to convert color, texture or shape into a mathematical measurement. The measurements are used to compare the features of the image entered by the user to the images in the system's database. More complex CBIR systems go beyond these features to include scene recognition, object recognition, and systems that 'learn' – basic neural networks.

When looking at color, most CBIR systems follow the same basic idea. When an image is entered into the database, the system will map it onto a grid. Each square on the grid, or pixel, will be of a specific color. The system computes the proportion of pixels of each color within the image and stores that information in the database. When the user enters an image or chooses a color from a palette, the system analyzes it in the manner described above, and pulls out the images that have a similar proportion of colors. This technique is often combined with the use of

other features to ensure more than a color match. In addition, the image can be split into regions, with the analysis above performed on each region.

Texture comparison is accomplished through a comparison of the brightness of a set of pixels. By analyzing how the light bounces off two disparate points, it is possible to determine features such as granularity, contrast, pattern variations, and coarseness. Texture is integrated into a CBIR system in a similar manner to color – the user can either choose a specific texture from a palette, or the user can upload a picture that the system analyzes.

Shape retrieval is a slightly easier problem for CBIR systems to resolve – shape is a defined concept, as opposed to texture, which has infinite variability. A CBIR system will determine a set of characteristics for each object within every image in the database. Size and orientation are not taken into account in order to allow scaled images to be recognized. Shape retrieval systems use a variety of approaches in deciding which characteristics will be taken into account. One specific algorithm compiles a list of each edge of an object, and puts a direction on each edge with relation to the other edges of the object. When given a new image, the system will apply the algorithm, get a new list of edges with their respective directions, and compare that list with the ones already stored in the database.

Another approach involves comparing the spatial location of objects within the images. In other words, the system will look at the position of one object within one image, and then look for images within its database that have that same object in a similar position. Alone, spatial location is not the ideal method, but when combined with other approaches, it can be highly useful.

Currently deployed technologies in the market.

Researchers generally identify four stages when looking at innovation cycles for a new technology: (1) research, (2) early commercialization, (3) niche-specific products, and (4) widespread proliferation. During the research and early commercialization stages, experts add knowledge to the field, and expertise is required to use any products that are developed. During the niche-specific product and widespread proliferation stages, the technology becomes more transparent – the user simply uses the technology, without knowing the details of how it works.

We see CIBR in transition between stages 3 and 4. From well-known corporations like IBM to web startup Riya.com, CBIR is moving rapidly into public use. Based on a conservative assessment of the market, We believe a visual registry is a viable option.

A few of the more mature examples of CBIR systems are described below.

Nevenvision

Nevenvision was recently acquired by Google (the implications of one of the pre-eminent tech corporations of the 21st century seeing the value in CIBR should be obvious). Nevenvision combined cellular phones and CIBR in an innovative system called iScout that allows users to obtain information about a landmark, movie, bracelet, car – anything that can be photographed through a cellular phone camera and sent to Nevenvision’s CIBR system. For instance, if the user was interested in a movie, they could snap a picture of the billboard. The picture would be sent to the Nevenvision CIBR systems. After identifying the image and any objects of interest within the image, the system would transmit results – movie times, reviews, perhaps a trailer – back to the user. iScout was used in Germany for a Coca-Cola campaign, where users were asked to take pictures of Coke cans and send them in for prizes. Nevenvision’s servers would first rotate the image until the CIBR system could focus and pull out certain ‘interest points’ – characteristics that allowed the CBIR system to recognize the incoming image first as a can or a bottle, and then look for the Coca-Cola image on the bottle or can itself. The interest points are then used to extract a set of numbers that act as a descriptor of the local neighborhood of the interest point – for instance, the curl of the C in ‘Coca-Cola’. The CBIR system then compares the set of numbers derived from the incoming image against the set of numbers derived from the image that exists in the database. If the sets are similar enough, the image is accepted as a Coca-Cola bottle. The ad campaign was a massive success, with participation much higher than the usual marketing campaigns that involved cellular phones because of the integration of a CBIR system with the phone’s camera.

Nevenvision also licensed its facial recognition and image recognition software to Docomo, Japan's largest mobile operator. The new phone models will have facial recognition software that will allow secure payment transactions. Various Docomo phone have cash-card

functionality that allows the phones to be used in a manner similar to a credit card. The technology was deemed trustworthy.

Google's acquisition of Nevenvision is no surprise – the search engine giant recently took its Picasa picture organizing software online. A CIBR system embedded into Picasa would make it superior to its competitors, boosting Google's user base.

Interwoven

Interwoven's MediaBin Asset Server is a ready-to-use CBIR system. It is specifically designed to help companies manage and locate their digital assets after those assets had been distributed. Interwoven's CBIR system does not do a pixel to pixel comparison. Instead, it compares sections of pixels with one another. The benefit of this technique is that it gives the user the ability to search for specific features rather than an overall image. The CBIR system is searchable by shape, texture, color, or resemblance to another image. It is not keyed to a specific geometry (in contrast to facial recognition systems cannot be used for general image recognition purposes, as the algorithms are specified for face specific objects). In addition, Interwoven's CBIR system can emphasize certain search mechanisms according to user preference – for instance, if the user feels that returning triangular shapes is more important than returning black, yellow and green colors, the shape search can be weighted so that the CBIR system will focus primarily on that field.

Interwoven's system is used by several companies -- including Walt Disney International, Thomas Kinkade, and Ventura Foods -- to track digital assets including images, video, power point presentations, and graphics.

Like.com

CBIR systems are not just the domain of large corporations; Like.com provides an example of a CBIR system that ordinary consumers can use.

Like.com is an image search engine, but unlike any other on the Web – it is a true CBIR system. For instance, the user begins with a picture of a purse. Like.com's CBIR system operates on individual aspects of the purse – its handle, its curves, or a specific jeweled exterior. As in other CBIR systems, these features are represented by a set of numbers. What is special

about Like.com is that the user can focus on these aspects – highlight the handle, and the engine will return purses with similar handles. This is accomplished by comparing the particular set of numbers that represent the handle with similar sets of numbers in the database. The CBIR system also allows users to choose from among colors and shapes, or a combination of the categories. At the moment, Like.com does not have an option for users to upload their own images, but that functionality is expected to be added within a few months' time. In the meantime, the shopping experience is remarkably smooth and enjoyable.

Visual registries are a viable option.

Like.com's entry into the marketplace, alongside the presence of Google, IBM, Interwoven, and numerous other competitors, is an indication that CBIR systems are in the 4th phase of the innovation cycle – widespread proliferation. The technology exists and is already on the market. A visual registry is more than an option – it's a reality.