Using controlled vocabularies to organise digital images for improved search results

Received (in revised form): 28th January, 2016

Ann Pool
graduated from the University of Washington in 2013 with a master's degree in library and information science. She has worked at Corbis as a search metadata technical specialist, focusing on using crowd-sourcing campaigns to improve the quality of search keywords associated with digital images.

Laura Horan
has been a search vocabulary specialist at Corbis, where she worked to improve search through taxonomy, metadata improvements and communication with contributors. She graduated with a master's degree in library and information science from the University of Washington in 2014.

Abstract Controlled vocabularies are essential tools for enabling search within a collection of assets. When constructing a controlled vocabulary for a collection of digital images, however, it is important to consider the visual nature of images. The most important principle when designing and using a controlled vocabulary to organise an image collection is that a term should be associated with an image if *and only if* a user searching the collection would expect and want the image to appear in the result set when searching for that term. This paper explores taxonomy development and maintenance for digital image collections, how hierarchical standards might vary between image and text-based collections, and the impact of keywords on precision and recall.

KEYWORDS: taxonomy, controlled vocabulary, search, image retrieval, image indexing

INTRODUCTION

The number of photographs being taken increases dramatically every year. It has previously been predicted that 1 trillion images would be taken in 2015 alone — an increase of nearly 300 per cent over the preceding five years.[1] This number is expected to rise to 1.3 trillion in 2017.[2] Although most of these images are taken by individuals for personal use, visual digital assets are also growing within organisations. From cultural institutions to stock photography companies, as more visual content becomes available, functional image search becomes even more important. In organisational settings, controlled vocabularies are excellent tools for accommodating the search and retrieval of images and other content. When developing a controlled vocabulary for organising and

accessing images, however, it is critical to build a vocabulary specific to the visual nature of images.

This paper explores best practices of developing and maintaining a controlled vocabulary within the context of digital image collections. The search and metadata needs of image collections can vary from text-based collections. In image search, it is important that the images retrieved visually reflect the search criteria of the user. This is especially apparent when thumbnails of the images are retrieved, such as within digital image collections. As a result, the visual nature of images should be considered when developing hierarchical relationships and determining the granularity of equivalence relationships.

This paper discusses the background of image retrieval and how the capabilities of modern search engines affect taxonomy design; explores how poor keywording and hierarchical relationships can affect search results; describes the potential pitfalls of working with batches of digital images; and discusses strategies to mitigate these problems.

LITERATURE REVIEW

Digital images are retrieved using one of two principal methods: content-based image retrieval (CBIR) and concept-based or text-based image indexing. In CBIR, artificial intelligence software uses attributes of the image itself to recognise its content. Machine learning for image recognition continues to improve. Researchers have developed software that can identify specific objects and activities such as games in images,[3] and algorithms have beaten human judges in an image recognition competition.[4] CBIR, however, remains largely in the experimental stage, and most real-world CBIR systems are limited to simpler functions such as facial recognition. CBIR systems also may challenge searchers, especially those who are less artistically inclined, if they rely on non-text-based queries such as sketches or example images.[5] Retrieving images based on complex, subjective or abstract characteristics continues to depend on text attributes associated with the image.[6]

Image attributes used in text-based indexing have been categorised and organised into hierarchies. Non-visual attributes include image titles and photographers, photographic and artistic techniques, and contextual attributes. Syntactic visual attributes include low-level visual properties of an image. Examples include colour, orientation and physical relationships between image elements. Semantic attributes include concrete or generic elements such as objects and people, and abstract concepts including emotions and human relationships. Generic attributes range from general to specific, from broad categories to proper names.[7]

Several studies of image retrieval have focused on targeted searches, either for images similar to a given example, or to satisfy subjects' own information needs.[8–10] In general, researchers have concluded that physical attributes, proper names and concrete terms are used in search queries more frequently than subjective terms, and concrete terms much more frequently than abstract terms.[11–13] Query length and specificity vary with information needs and content types[14] and as searches are refined to improve results, with searchers often beginning with more general terms and increasing specificity on later query iterations.[15] Searchers may begin with short, often single-term queries. These queries 'will be generally successful, although the amount of resulting images will often be overwhelming, discouraging, even frustrating sometimes'.[16] At this point, users often attempt to use the Boolean AND and NOT operators to narrow the result set on later iterations. Boolean queries, however, 'continue to demonstrate that Boolean is poorly understand [sic] by end users and that a little Boolean can be a dangerous thing'.[17]

Ménard and Smithglass[18] categorised search interfaces and controlled vocabularies as a function of the types of organisations employing them. Libraries are most likely to use traditional controlled vocabularies such as the Library of Congress Subject Headings (LCSH) and the Getty Research Institute's Art & Architecture Thesaurus. Choi and Hsieh-Yee,[19], however, found that mismatches often occur between LCSH terms and those used in image searches. The more technical terms may be useful for specialists, but others need everyday terminology. Museums also use traditional vocabularies but are moving towards user-supplied keywords or selectable categories that relate directly to the museum's collection. Browsing also is emphasised in museum image collections. Image search engines, including those employed by stock photography websites, provide simple, web search-like query interfaces. These systems depend on the ability to bridge user queries with the system's vocabulary, and are not particularly conducive to browsing. Stock photography sites in particular 'are structured outside of traditional approaches to descriptive methodologies and controlled vocabularies. Images are organised by categories and with standardised vocabularies in a taxonomic structure unique to each site, but with common functionalities across the subtype'.[20] While basic principles of information organisation still apply to commercial organisations' digital image collections, the implementation of those principles is adapted extensively to meet the needs of those companies and their customers.

Studies of text-based image indexing and retrieval have focused primarily on small groups of users performing targeted image searches. Patterns of search behaviour such as query length, categories of terms used and iterations have been analysed to infer how people search for images and how well the results meet their needs. This paper provides recommendations for improving search results based on applying one primary principle: an image should be associated with a search term if *and only if* a user searching the collection would expect and want the image to appear in the result set when searching for that term.

CREATING A TAXONOMY FOR DIGITAL IMAGES

A well-designed and maintained taxonomy provides a solid foundation for organising images to facilitate retrieval of relevant results. Although the same basic principles used in designing controlled vocabularies for organising documents apply to images as well, those principles need to be applied in ways that take the nature of visual information into account. Capabilities of online taxonomy management tools and search engines also must be taken into account when designing controlled vocabularies for online retrieval.

DESIGNING HIERARCHICAL RELATIONSHIPS

When a user searches for a term with narrower terms, modern search engines typically return results associated with those narrower terms along with those associated with the search term itself. Therefore, determining hierarchical relationships in a taxonomy to be used for classifying images for online retrieval requires constant consideration of the question, 'Will a user searching for the broader term also want to see results for the narrower term?' A hierarchical relationship should be created if and only if the answer to that question is yes.

In its guidelines for controlled vocabularies, the National Information Standards Organisation (NISO) recognises three types of hierarchical relationships in taxonomies. Generic or 'isA' relationships link a class with its members. For example, generic relationships exist between *mammal* and *dog*, and between *dog* and *dachshund*.

Instance relationships, another type of 'isA' relationship, link general categories with specific instances of those categories. The relationship between US *states* and *Alabama* is an instance relationship, as is that between *rivers* and *Mississippi River*. Whole/part relationships apply to terms that can be broken down into sub-divisions. Examples can be found in domains including anatomy, such as *hand* and *finger*, and geography, such as *Texas* and *Austin*.[21]

Whole/part relationships often do not translate well into hierarchical relationships in an image taxonomy. A person searching for *hand*, for example, may not be interested in images of index fingers or thumbnails. Another problematic area in whole/part relationships is geographic locations. An image that effectively represents a specific location may not be a strong depiction of a larger geographic region containing that location. For instance, an indoor scene in a seafood restaurant in Seattle is probably not the first thing that comes to mind for a customer looking for images of Seattle. Geographic locations can be particularly problematic for digital images because cameras frequently apply locations to images as they are taken. This may result in the unintentional application of location-related keywords to images that do not represent the location well. On the other hand, sometimes whole/part relationships do work well for image searches. Rooms within special-purpose buildings are one example: hospital rooms may be desirable when searching for *hospital*, and similarly classroom scenes may work well as results for *school*. Awareness of likely user needs and expectations will help determine whether specific whole/part relationships are appropriate for a particular image collection.

Generic and instance relationships are more likely than whole/part relationships to work well in image taxonomies but still require thought. For instance, a search for *animals* can be expected to retrieve not only images that have been associated with *animals* but also with *two animals* and *ring-tailed lemur*, among many others. The threshold for a sufficiently relevant image can, however, vary at different levels of specificity. Someone searching for *animals*, for instance, probably wants to see images where animals are the primary focus. However, a person looking for a *West Highland terrier* or a *ring-tailed lemur* may be happy with any image containing that particular breed or species, even in the background. Similarly, terms linked by instance relationships may not be given the same threshold of relevancy in the eyes of a customer. As with generic relationships, a broader term often needs to be a more prominent part of an image than does a narrower term. A person searching for *island* probably wants to see all or most of an island surrounded by water, while someone looking for *Easter Island* may well be happy with anything taken at that location.

In practice, adjustments made to an image taxonomy to improve the precision of search results often tend to flatten the hierarchy when compared with a taxonomy used for organising documents. In other cases, however, relationships may be found between terms that do not fit any of the accepted guidelines for hierarchical relationships, but do improve results for image searches. Often these would be classified as associative relationships between related terms. For example, *Christmas* and *Christmas tree* normally would be considered related terms, but it may be that images with Christmas trees would satisfy a typical user's desire for images of Christmas. Creating a hierarchical relationship in an image taxonomy might be beneficial in a case like this. It is also well worth considering cases in which a relationship is not hierarchical in a technical sense, but is thought of as hierarchical in the way that things are used or perceived on an everyday basis. For example, peanuts are legumes rather than nuts, but they are much more likely to be found in a can of mixed nuts than in a bean salad. Therefore, *peanut* would more likely be appropriate as

a narrower term to *nut* than to *legume* in a taxonomy used to organise a collection of images used for non-scientific purposes.

In summary, when creating or maintaining an image taxonomy, it is critical to understand the users and what they are looking for, and to create relationships between terms that facilitate the retrieval of relevant images, even if those relationships do not fit the typical standards for controlled vocabularies. Appropriate terms and their relationships will differ greatly, for example, in a collection of food-related images used in a university food science department compared with the same collection used in a cooking website.

EQUIVALENCE RELATIONSHIPS AND GRANULARITY

Written and spoken language is rich with levels of meaning and intensity. Absolute synonyms, or words with exactly the same meaning in all contexts, are exceedingly rare at best. Synonymy is, therefore, considered as a scale ranging from cognitive synonyms, or words with identical descriptive meanings, to near-synonyms whose meaning and sense overlap. Thesauri and dictionaries of synonyms typically include near-synonyms in their listings.[22] In taxonomies, equivalence relationships are created between synonyms, with one term treated as the preferred term and other equivalent terms as non-preferred terms, which refer to the same concept. NISO recommends that:

> 'terms should be treated as near-synonyms only in subject areas that are peripheral to the domain of the controlled vocabulary. When concepts can be distinguished in the controlled vocabulary domain with sufficient precision to justify their representation as separate terms, they should be individually defined and retained.'[23]

This recommendation may be considered when designing controlled vocabularies for digital images, as long as the idea of a domain within the context of visual information is well defined and understood. Near-synonyms that can readily be differentiated in written work may be indistinguishable in images. This is particularly likely with abstract terms. It may be possible to write scope notes that prescribe the use of *happiness*, *cheerfulness* and *joy* to represent different shades of meaning, but that does not guarantee that people can tell the difference between happy, cheerful and joyful people in an image collection. If those subtleties are deemed irrelevant to the collection and its users, equivalence relationships between these terms would be appropriate. On the other hand, near-synonyms that may be used interchangeably in writing can have significantly different connotations when used to describe images. For instance, *serene* and *easygoing* are near-synonyms, but each likely conveys a different mental image to a person searching for images to convey a mood. A serene wilderness landscape and a group of easygoing teenagers watching a movie can be expected to satisfy sharply different information needs.

These considerations apply to more concrete terms as well. For example, *sea* and *ocean* represent different geographic concepts; seas are smaller than oceans and surrounded at least partially by land. Photographs of oceans and seas may, however, appear the same. Whether they are interchangeable depends primarily on user needs, and in particular on intended uses of the images. To an advertiser looking for images of children playing in the surf, the precise geographic definition would not matter. For a controlled vocabulary organising a collection of stock images used in marketing, treating these terms as synonyms and creating an equivalence relationship between them makes sense. Recall would be improved and excessive complication would be avoided. To researchers in the fields of oceanography and geography, however, the difference between these concepts would indeed be significant.

To achieve an acceptable level of precision when searching image collections in these domains, controlled vocabularies do need to keep these terms separate. As with all other design decisions, anticipated uses of the collection and needs of its users must always be taken into account when determining the granularity of terms in a controlled vocabulary.

CONTROLLED VOCABULARIES: CHANGING OVER TIME

Controlled vocabularies are ever-evolving documents. As collections grow and new topics are represented, there will be a continued need for new terms and relationships within a controlled vocabulary.[24] The iterative nature of controlled vocabularies may also include evaluating non-preferred terms on existing concepts. Non-preferred terms that were appropriate when a concept was created may become a liability as the controlled vocabulary evolves. For example, in 2013, 'royal baby' was used frequently leading up to and after the birth of Prince George. Although somewhat ambiguous, this term was synonymous with Prince George and was an appropriate non-preferred term to refer to him. In 2015, however, 'royal baby' also started being used in relation to the birth of Princess Charlotte. It became important to remove 'royal baby' from the Prince George concept, so that pictures of his sister would not map to the concept for Prince George. Controlled vocabularies are living documents, and it is important to keep them up to date, both by creating new concepts and making sure older concepts in the vocabulary stay current.

Depending on the context of image search, an organisation might want to prioritise recent images in search results. Commercial and current event content is often time-sensitive, whereas recency might be less important for other content types, such as historical or documentary content. While tuning search to prioritise recent content is primarily handled by search engine settings, it is important to consider how the controlled vocabulary may affect search results for concepts that evolve visually over time. For example, cell phones have changed significantly over the past 20 years. A cell phone from 1995 looks quite different from a 2005 cell phone, and both are visually distinct from a 2015 cell phone. Although images of any of these phones are appropriate results for a search for *cell phone*, a stock photography company may want to prioritise cell phones from the past couple of years because they are more likely to sell to commercial photography customers. If search results for a time-sensitive concept look stale, it may be worth evaluating the hierarchy of this concept. In older hierarchies, there may be some non-preferred terms on concepts that are no longer current. For example, *flip phone* may be a non-preferred term on *cell phone*, because ten years ago when the concept for *cell phone* was created they were considered synonyms. Creating a separate concept for *flip phone* that is not narrower to *cell phone* will prevent these dated images from returning in a search for *cell phone*. Because it is stored in a separate concept, customers searching for *flip phone* will have more precision in searching for these images. As discussed earlier, flatter hierarchies can be an advantage in image search. This may be a strategy worth exploring for managing topics where the visual style has changed drastically over time.

CONTROLLED VOCABULARIES AND SEARCH

Controlled vocabularies are utilised to better enable findability of images within a collection; often this is achieved through search. Search functionality is often very specific to an organisation, depending on which search engine is used, how it is tuned and whether it includes full text search. Search priorities may also vary due to several

factors, including the size and audience of the collection. Smaller collections may prioritise recall over precision, preferring to return broader search results (example: a search for *happiness* may also return images of *laughter*). Owing to their size, larger collections may prefer to prioritise precision. Internal users may be more acquainted with the system and have a better understanding of how to search to accomplish their goals. In addition, an internal collection may be curated to better reflect the needs of the specific organisation. Search targeted at external users can be challenging to improve owing to a limited understanding of the needs and search behaviour of these users.

PRECISION AND RECALL

Keywording and search are inherently connected, and the perspective and search behaviour of the user should inform keywording. If a system caters primarily to external users, however, it may be difficult to truly understand the goals of these users. Search logs are helpful in determining what users of a website are searching for, but it can be difficult to identify whether or not the results are meeting the user's expectations. Even if the results are topically relevant to a search, 'users of image retrieval systems often find it frustrating that the image they are looking for is not ranked near the top of the results they are presented'.[25] In many image collections, searches for certain concepts return a large number of images — most of which will be buried several pages deep in the search results.

As image collection size grows, delivering relevant search results becomes increasingly important. Although this quote is from 1998, Anna Bjarnestam provides an excellent example of how many users still approach image search today:

> 'In contemporary stock photography users talk about pictures using words rather than shapes, forms, intensity, texture and colours. A client may ask for a picture of a happy person of a certain age and gender, who is sitting by a table holding a cup of coffee in one hand and a newspaper in the other, and the whole image should signal relaxation.'[26]

This combination of abstract and concrete terms suggests that users often approach a search with a specific idea of what they are looking for. A search for 'relaxed 35-year-old woman coffee newspaper', however, may not return what the customer expects. Website analytics data can be helpful in determining how users interact with a website but do not provide the full context of why a user clicked on one image over another, or why they abandoned their search entirely. Returning relevant search results is an important goal for image collections, but it is difficult to know what will be relevant to each individual user.

THE SEARCH IMPACT OF OVER-KEYWORDING

Keywording images is by nature subjective. Keywords are used to describe both the *ofness* and *aboutness* of an image.[27] However, what level of detail is appropriate? Should keywords identify every object in an image, even if some are not significant to what the image is about? For example, should an image of a kitchen have *kettle* as a keyword, even if the kettle is one of several objects in the kitchen and is not a significant aspect of the image? A rule of thumb for this has been to consider if an image would be a great search result for the keyword in question: Would this image be a great search result for *kettle*? The answer to this question is inherently subjective. The aboutness of an image is also subjective, as it can be difficult to determine whether an image reflects abstract conceptual keywords like emotions and themes. As a result, images are often over-keyworded. This is a common problem in image collections, as 'many image collections have a tendency to apply

large numbers of keywords to their images without much thought as to whether the keywords will help or hinder the image searcher'.[28] Over-keywording increases recall but contributes to search results being less precise.

Compounding the problem, many image collections contain externally generated keywords. This is particularly true in stock photography, as many companies rely on contributor-supplied keywords. Stock photography images are often submitted to multiple companies, all with varying keywording standards and expectations. It may be difficult to ensure that contributor images are adequately keyworded, owing to the lack of understanding of a company's keywording system and the fact that image keywords may not be specifically targeted to this particular company's standards. Some contributors over-keyword in the hope that their image will show up in additional searches, resulting in additional sales. The greater effect of over-keywording, however, is diluted search results, leading to a poor search experience.

Digital images are particularly susceptible to over-keywording because of the way that they are created, processed and ingested into asset management systems. An editor working with a batch of related images may be tempted to apply keywords *en masse* to the entire batch to save time and effort when ingesting the batch into a digital asset management system. However, it is extremely unlikely that each of these keywords will be appropriate for every image in the batch. A photo shoot with a family in a park, for example, might include a few images that strongly depict togetherness, others that convey playing and still others showing eating. Applying all of these keywords to the entire batch would result in over-keywording and later frustrate end users with irrelevant search results. The effort required to select appropriate keywords for individual images will pay off in an improved search experience.

Similar to over-keywording, another common problem is keywording for end use. Keywording for end use occurs when a keyword describing a potential use of an image is applied to an image that does not visually reflect that concept. For example, an image of a bouquet might have keywords like *Mother's Day* or *Valentine's Day* applied, even though the image does not contain any contextual details suggesting either of these holidays. Keywording for end use dilutes search results by causing irrelevant images to be returned for these concepts. While keywording for end use is well intended, trying to anticipate how an image is going to be used can create confusing search results for customers.

TRAINING AND DOCUMENTATION

Many of the search problems outlined in this paper are a result of misunderstandings around keywording standards. Communicating and enforcing keywording standards, including number of keywords and types of keywords (descriptive versus end-use), should encourage higher-quality keyword submissions and create a better search experience. Even in settings where keywording is done in-house, training and documentation are still relevant. Encouraging consistent metadata makes it easier to improve search by reducing irrelevant keyword usage and creating taxonomy mappings to help mitigate ambiguous terms.

SUMMARY AND CONCLUSION

The needs of image taxonomies are highly specific to the format, users and scope of the collection. Within most image collections, however, it is important to consider the visual nature of images in developing a strategy for taxonomy development and maintenance.

Image taxonomies may require different hierarchies versus their text-based

equivalents. While 'isA' relationships and generic relationships generally work well with image taxonomies, whole/part relationships may cause problematic search results. Digital image collections may benefit from a flatter hierarchy, which can result in greater search precision. When determining the depth of a taxonomy, it is important to consider the needs and perspectives of the users. Organisations should strive to create relationships that support the retrieval of images relevant to their users, even if these relationships do not reflect typical controlled vocabulary standards.

The granularity of equivalence relationships in taxonomies varies based on the collection and the needs of its users. It is worth considering whether near-synonyms that may be used interchangeably in writing are equivalent visually. This can vary based on the needs of a collection's specific audience.

Taxonomies are constantly evolving and it is important to keep them up to date. This involves adding new concepts and relationships, as well as maintaining previously developed concepts in the vocabulary. Certain topics may evolve visually over time; for example, images related to technology may begin to look dated. If search results look stale, it is important to evaluate hierarchies and synonyms to make sure that they are current and not contributing to dated search results.

Keywording has a close relationship with precision and recall in search. Occurring frequently in digital image collections, over-keywording and keywording for end use can contribute to diluted search results. When possible, it is important to encourage image keywords that reflect significant aspects and themes that are visible in an image. When image keywording does not occur in-house, developing standards that limit the number of keywords applied to each image and communicating how to identify relevant keywords can be helpful in improving image metadata and search results.

References

1. Heyman, S. (2015) 'Photos, photos, everywhere', *New York Times*, 23rd July, available at: http://www.nytimes.com/2015/07/23/arts/international/photos-photos-everywhere.html (accessed 27th June, 2016).
2. *Ibid*.
3. Markoff, J. (2014) 'Researchers announce advance in image-recognition software', *New York Times*, 18th November, available at: http://www.nytimes.com/2014/11/18/science/researchers-announce-breakthrough-in-content-recognition-software.html?_r=0 (accessed 27th June, 2016).
4. Johnson, R.C. (2015) 'Microsoft, Google beat humans at image recognition', *EE Times*, 18th February, available at: http://www.eetimes.com/document.asp?doc_id=1325712 (accessed 27th June, 2016).
5. McDonald, S. and Tait J. (2003) 'Search strategies in content-based image retrieval', in 'Research and Development in Information Retrieval: Proceedings of the 26th Annual International ACM SIGIR Conference; 28th July to 1st August', ACM, Toronto and New York, pp. 80–87.
6. Ménard, E. and Smithglass, M. (2014) 'Digital image access: an exploration of the best practices of online resources', *Library Hi Tech*, Vol. 32, No. 1, pp. 98–119.
7. Goker, A., Ridley, D. and Davies, J. (2009) 'Information Retrieval: Searching in the 21st Century', Wiley, Hoboken, NJ.
8. Choi, Y. and Hsieh-Yee, I. (2010) 'Finding images in an online public access catalogue: analysis of user queries, subject headings, and description notes / Le repérage d'images dans un catalogue en ligne à accès libre: analyse des requêtes des utilisateurs, des vedettes-matière, et des notes descriptives', *Canadian Journal of Information and Library Science*, Vol. 34, No. 3, pp. 271–295.
9. Choi, Y. (2010) 'Analysis of image search queries on the web: query modification patterns and semantic attributes', *Journal of the American Society for Information Science and Technology*, Vol. 63, No. 7, pp. 1423–1441.
10. Chung, E. and Yoon, J. (2010) 'An exploratory analysis on unsuccessful image searches', *Proceedings of the American Society for Information Science and Technology*, Vol. 47, No. 1, pp. 1–2.
11. Choi, ref. 9 above.
12. Chung and Yoon, ref. 10 above.
13. Goker *et al.*, ref. 7 above.
14. Choi, Y. (2010) 'Investigating variation in querying behavior for image searches on the web', *Proceedings of the American Society for Information Science and Technology*', Vol. 47, No. 1, pp. 1–10..
15. Choi, Y. (2011) 'Query modification patterns and concept analysis of web image queries', *Proceedings of the American Society for Information Science and Technology*, Vol. 48, No. 1, pp. 1–3.
16. Ménard and Smithglass, ref. 6 above.
17. Jörgensen, C. and Jörgensen P. (2005) 'Image querying by image professionals', *Journal of the American Society for Information Science and Technology*, Vol. 56, No. 12, pp. 1346–1359.
18. Ménard and Smithglass, ref. 6 above.

19. Choi and Hsieh-Yee, ref. 8 above.
20. Ménard and Smithglass, ref. 6 above.
21. National Information Standards Organization (2005) 'Guidelines for the Construction, Format, and Management of Monolingual Controlled Vocabularies', NISO Press, Bethesda, MD.
22. Stanojević, M. (2009) 'Cognitive synonymy: a general overview', *Linguistics and Literature*, Vol. 7, No. 2, pp. 193–200.
23. National Information Standards Organization, ref. 21 above.
24. Slawsky, D. (2007) 'Building a keyword library for description of visual assets: thesaurus basics', *Journal of Digital Asset Management,* Vol. 3, No. 3, pp. 130–138.
25. Hare, J., Sinclair, P., Lewis, P. and Martinez, K. (2008) 'Giving order to image queries', in: Gevers, T., Jain, R.C. and Santi, S. (eds) 'Multimedia Content Access: Algorithms and Systems II'; available at: http://eprints.soton.ac.uk/265187/1/iRankr.pdf (accessed 27th June, 2016).
26. Bjarnestam, A. (1998) 'Text-based hierarchical image classification and retrieval of stock photography', paper presented at the Challenge of Image Retrieval Conference, 5th February, available at: http://www.bcs.org/upload/pdf/ewic_im98_paper4.pdf (accessed 27th June, 2016).
27. Slawsky, ref. 24 above.
28. Hare *et al.*, ref. 25 above.