

Spare Time Activity Sheets from Photo Albums

Gabriela Csurka , Guillaume Bouchard and Marco Bressan

Xerox Research Centre Europe
6, Ch de maupertuis, 38240 Meylan, France
+33 (0)4 76 61 50 50
email: firstname.lastname@xerox.com
<http://www.xrce.xerox.com>

Abstract

Given arbitrary pictures, we explore the possibility of using new techniques from computer vision and artificial intelligence to create customized visual games on-the-fly. This includes coloring books, link-the-dot and spot-the-difference popular games. The feasibility of these systems is discussed and we describe prototype implementation that work well in practice in an automatic or semi-automatic way.

Keywords: segmentation, visual games, object categorization, inpainting

1 Introduction

The pencil and paper Spot-the-Difference game, coloring pages and related exercises, finding the impostor or the hidden objects, matching the right shadow are classical and popular recreations for children from many countries. Some of these spare time activities are preferred activities not only for children but is suitable for any age. They are often proposed by popular mag-

azines and journals.

Classically, most of these activities were available on cartoon-like images and drawings. Many of them recently became more popular with real images. With the growing number of digital cameras, a natural trend would be to apply these games to personal pictures. Nevertheless, currently the creation of them is done manually and requires several complex processes and steps.

In this paper, a set of automatic and semi-automatic creation tools is proposed to allow non-expert users to generate quickly such spare-time activity sheets from arbitrary images, generally photographs. We do not describe any novel computer vision techniques, but show how state-of-the-art image analysis and processing techniques can be adapted in this context.

In section 2, we describe a simple method to create games based on image contours which works well for simple images containing well contrasted and uncluttered objects with relatively uniform background and extend this method in order to cope with more complex scenes, allowing the creation of diverse activity sheets. Section 3 discusses the automatic creation of spot-the-difference games. Section 4 concludes the paper.

Digital Peer Publishing Licence

Any party may pass on this Work by electronic means and make it available for download under the terms and conditions of the current version of the Digital Peer Publishing Licence (DPPL). The text of the licence may be accessed and retrieved via Internet at <http://www.dipp.nrw.de/>.

First presented at the 4th International Conference on Computer Graphics Theory and Applications GRAPP 2009, extended and revised for JVRB.

2 Games based on image contours

Coloring images are generally simple black and white silhouette or border images with well separated regions, each corresponding to a different color. These images can also present several differences in style (see Figure 1) leading to different spare time activities such as unsupervised silhouette coloring, numbered/labeled region coloring, dot linking, etc. They are also often used in kindergarten and elementary schools, where the labeling has to be deduced as part

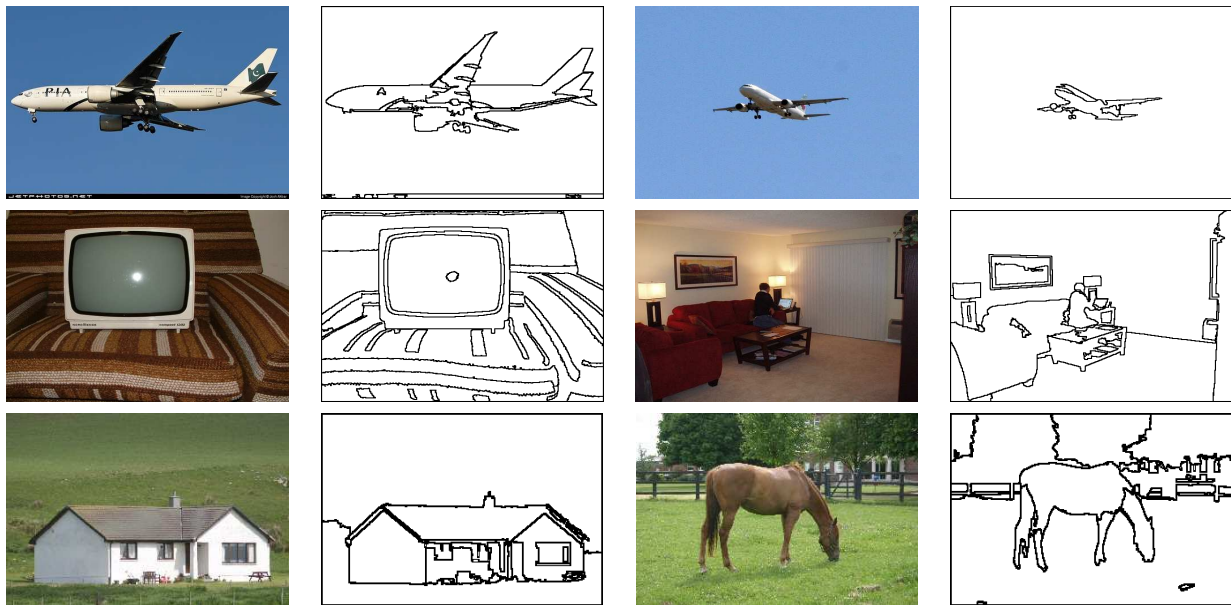


Figure 2: Example results for the basic approach.




Figure 1: Different types of activity sheets.

of an exercise e.g. a mathematical formula.

Typically, most of these activities were available on cartoon-like images and drawings. Transforming a printed photograph into a drawing suitable for any of these activities requires a complex manual process with multiple steps. With the popularity of digital photography, it is natural to use a digital technique to simplify this process.

The main challenge we address here is to obtain coloring pages from the arbitrary types of images children might be interested in coloring or filling i.e. photographs and cartoons. The coloring page creation can be seen as a particular case and application of photographic stylization and abstraction. Indeed, they

share many common components, such as building edge maps and image segmentation, even face detection for non-realistic rendering [Bro07], and therefore these components used or developed in the former field can be re-used to inspire and to improve the coloring page creation. The former techniques are not specifically adapted to the coloring page creation and even less to derive the divers activity sheets. Their aim is to obtain painting rendering of the images [DS02, OK06], stained glass effect [Mou03] or to enhance the compression rate for visual communication efficiency [WOG06]. They generally combine the luminance edge maps with the “abstracted colors” of the image which mutually compensate the visual imperfection of both of them, leading to a paintings rendering stylized effect.

We propose a method to generate coloring book pages automatically from user photos [CB09]. The only other approach we are aware of with this objective is the a recent commercial product Kidware Photo Color Tool (<http://www.kidware.net>). The results of the latter system are weighted edge maps similar to our texture edges, its main drawback being that many of the regions are not closed and not suitable for the derived activity sheets. As coloring images are mostly designed for young children, simplicity is desirable. Preferred images are black and white silhouette or border images with well separated regions, each corresponding to a different color. Another advantage of our system compared to it is that we address a wider



Figure 3: Example results for the system with addition of texture edges. In second column the results of the basic system B, in the third column the edge map obtained by the DoG algorithm normalized to $[0, 1]$ and in the last column the weighted combination of them.

range of derived activities.

We first propose a relatively simple system to obtain automatically coloring pages from a natural or cartoon-like image which nevertheless goes beyond a simple weighted edge map. As many segmentation techniques, the main challenge is to obtain connex regions that do not miss important object boundaries and at the same time do not over-segment the image into small areas difficult to color. The main steps of our system are (see for details [CB09]): Color Conversion (allows metric-based processing such as clustering), Edge-Preserving Low-pass Filtering (reduces image noise [GBT00, PM91]), Image Segmentation (region clustering using Normalized Cuts [JJ00] or Mean Shift [CM02] to take into account the spatial closeness of the pixels and therefore lead to more compact segments), Region Merging (since we intentionally over-segment the image in the previous step) and Edge Di-

latation (to get thicker region borders).

In our experiments, we used Mean Shift with a flat kernel and low color and spatial bandwidths ($\sigma_s, \sigma_r \in [5, 10]$). The bandwidth parameter allows handling the coarseness similarly in different images without specifying the exact number of clusters to be found in the image. In order to ensure we do not miss any perceptually important boundary, we intentionally use a low bandwidth to over-segment the image. Concerning the region merging step, we used a set of simple rules that take into account both spatial and perceptual information:

1. If the area of the region is below a given threshold ($T_1 = 0.0005\%$ of the image area), it will be absorbed by the most similar neighbor, independently of the color difference between them.
2. If the area of the region is above T_1 but below a

second threshold $T_2 > T_1$ ($T_2 = 0.05\%$ of the image area) the region is merged with its most similar neighbor only if their color similarity is below a threshold that depends on the color variance of the image.

3. If the area of the region is above T_2 the region is kept unchanged.

In both cases the color similarity is computed as a combination of distances in chrominance and luminance space, giving a higher weight (importance) to the distance in the chrominance space. This merging algorithm is applied iteratively until no modification is made or the maximum number of iterations is achieved. Alternatively, more complex region merging criteria could be applied, such as minimal cost edge removal in the corresponding region adjacency graph [HEMK98].

Figure 2 shows a few examples obtained on real images with this basic system. We see that this simple approach is suitable for images where regions are easily distinguishable and each region shows fairly uniform colors. It gives satisfactory results in many cases. However, it can also lead to less satisfactory results as scenes become more complex (as in images of the second column in Figure 3).

2.1 Visual Enhancement

In this section, we propose extensions to the basic system to improve the quality of the coloring pages.

Adding Texture Edges One of the main difficulties of obtaining an acceptable coloring page for a complex scene images is that generally there are several objects/elements of the scene for which the level of “interesting” details can vary a lot. In a coloring page application, not every detail will require the same level of attention, e.g. a human face or the leaves or branches of a tree. An automatic system that has no knowledge about the image content will handle these regions in a similar fashion. The global parameters can be tuned to increase or decrease details, but the same criterion will be applied to all regions.

In order to handle this, we propose a solution based on texture/luminance edges. To extract the luminance edges we use the Difference of Gaussians (DoG) algorithm as it approximates well the Laplacian of Gaussian, known to lead to a weighted edge map. Furthermore, the DoG is believed to mimic well how neural processing in the retina of the eye extracts details

from images destined for transmission to the brain. After elimination of small edges and isolated dots in the DoG map, these luminance edges are combined with the region boundaries obtained in section 2. We used a weighted combination giving a higher weight¹ to the original coloring page borders in order to let them guide the coloring.

The main role of this combination is visual enhancement, however in many cases it can also compensate for missing borders between regions that were wrongly merged either by the low level segmentation or by the region merging step. These cues further help children to better understand the content if they do not have the original model.

Figure 3 shows a few examples of coloring pages with and without adding these texture edges. The final results are still not perfect compared to what a human would do manually, however it seems that children can cope well with these small imperfections (see Figure 14).

Extracting ridges and valleys [TL04] can provide an alternative or a complement to this approach as artists frequently use it in their work.

Semantic Content Analyses It is clear that if the system has some further knowledge about the semantic content of the regions then it can automatically handle those regions accordingly. For example, it can increase the weights or the thickness of the object’s border, merge regions within the same semantic region, add luminance edges or ridges only to the regions of interest, etc. We can also alternatively fill some of the regions with the original content as in Figure 4 based on human skin detection.




Figure 4: Example results for the system with addition of skin detection and filling with original content.

In the last few years there were many publications and an increasing interest on semantic segmentation of images, i.e. assigning each pixel in an image to

¹In our experiments we used $\max(E, DoG^\lambda)$ with $\lambda = 0.4$, where the original edge map E is binary and DoG was normalized to have values between $[0,1]$.

one of a set of predefined semantic classes. This is a supervised learning problem in contrast to the “classical” unsupervised low-level segmentation. In the coloring page application, if the image contains main object (often centered) and simple background classical foreground-background separation techniques are well adapted such as GrabCut [RKB04] with manual (e.g. drawing a box or a contour) or automatic (a centered box in the image) initialization. For more complex images, semantic based image segmentation techniques are more adapted. The idea is that the image is partitioned generally in more than two semantically labeled (meaningful) regions. [SWRC06, YMF07, CP08]. We used the last approach [CP08] in our experiments referred as CBIS in which follows.

The integration of these techniques with the coloring page system can hence further enhance the visual quality of the coloring page (see Figures 4 and 5), but they are of particular interest for the derived activity sheets as we will see.




Figure 5: The sheep class mask (2nd image) obtained by CBIS was used to eliminate the background from the coloring page obtained by the basic system (3rd image). Furthermore, the low level segmentation was replaced by the high level masks boundary and texture edges added inside the relevant region (4th image).

2.2 Diverse Activity Sheets

Region Labeling In contrast to the unsupervised case, the idea here is that the child has to follow some rules to color each region. It can be simply the recognition of some letters as in the 2nd example of Figure 1), or it can be more complex such as mathematical or logical formulas. The latter are often used in kindergarten and elementary school with pedagogical purposes.

With our system, this can be done automatically, be-

cause (1) we have closed regions and (2) we have a representative color of each region (mean color, cluster center or mean shift mode). We can therefore select a set of standard colors (e.g. using the well-known, standard NBS-ISCC color name dictionary ², to find for each region the selected standard color which is closest to its representative color and plot the corresponding letters, formulas, shape, etc (depending on the children’s age) onto it. Eventually, on the border or next to the image, the legend is printed with the labels and the corresponding color (see Figure 6).




Figure 6: An example of a labeled coloring page.

Link the Dots A second popular activity sheet example is the link-the-dots sheet (see 5th example in Figure 1). These sheets are also often used by kindergarten as they help children to learn number ordering and the alphabet.

The main idea is to take a single object boundary using one of the semantic segmentation techniques mentioned above, sample dots on the edges, label them with letters, numbers or formulas following the contour and eventually delete the original contour.

The dots on the boundary can be sampled uniformly or obtained by more complex algorithms that seek for corners and inflexions points such as chain code detection [LS90], local contours [RLGR⁺02], direct estimation of the curve and its high curvature points [CS99, HK05]. We used a local implementation of the popular Curvature Scale Space (CSS) based corner detector [AMK99, HY04] to obtain a set of dots on the object boundary (see Figure 7).

Object Discovery through Coloring Finally, a third type of activity sheet is to discover hidden objects through coloring. The objective can be a simple foreground or background coloring, labeled by dots as in 4th example in Figure 1; or more complex, where a set of colors has to be used and the individual labels are mathematical or logical formulas (as in the 6th example of Figure 1).

²<http://www.anthus.com/Colors/NBS.html>

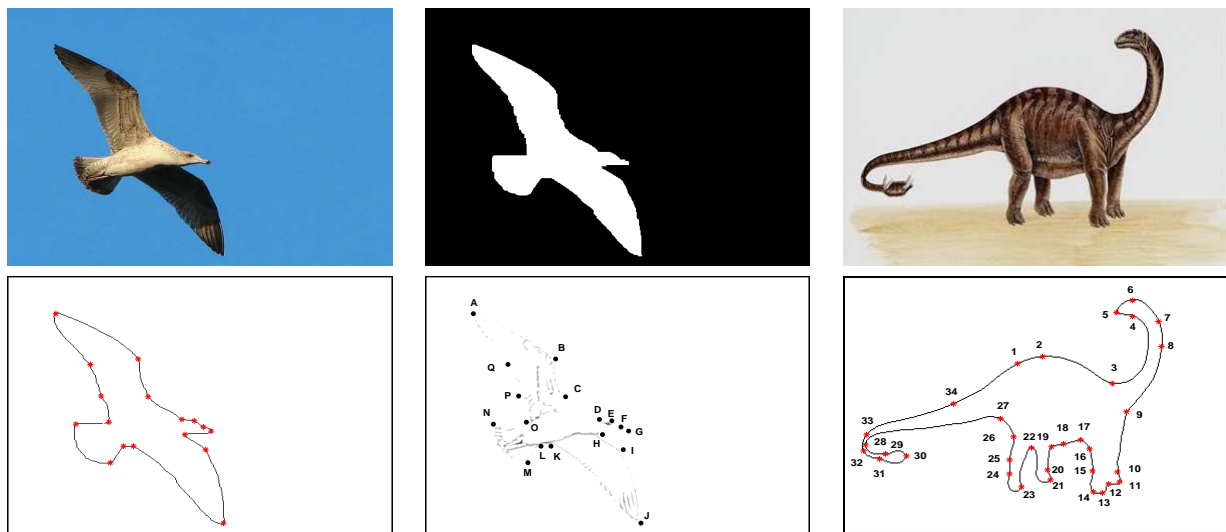


Figure 7: Example of automatically obtained follow-the-dot examples. In the first case we used the bird class mask obtained by CBIS and the dots were obtained by CSS corner detector on the boundary of the mask. We also added the DoG edges inside the object region to enhance the final result. In the second case, we initialized the GrabCut [RKB04] with a box centered in the middle of the image and made foreground/background separation. We show the original object boundary on the results for pure visualization purposes.

These sheets can also be derived from our solution when we have the knowledge of foreground/background or alternatively the semantic regions. In these cases, the idea is to either use the over-segmentation we already have in step 3 (section 2) or we can combine the high level segmentations with some random partitioning of the image. Finally, the dots/formulas can be added automatically to sub-regions according to their semantic meanings. In our experiments (see Figure 8) we obtained the semantic meanings of the regions with the CBIS, considering all Pascal VOC 2007 [EGW⁺07] classes as relevant regions.

2.3 Interactive coloring book generator

We can envisage the integration of the system with any photo editing tool or interactive coloring systems by adding different interactivity levels:

1. At a first level of interaction the system allows the user (child) to select or to upload a photo. The photo is automatically processed and a set of coloring pages (with or without texture edges) and activity sheets are proposed.
2. A second level of interaction would be designed for older children or parents allowing to modify/adjust some of the parameters of the system. However, the interaction with the parameters has

to be user friendly, such as choosing between less or more details, thinner or thicker edges, adding or not texture edges, adding letters or formulas, etc. Then the corresponding interior parameters are adjusted accordingly.

3. Finally a highest level of interaction could allow the user to edit the obtained coloring page by some interactive tools such as an erasing tool (the two regions separated by the selected edges will be merged automatically), a Dot Adding Tool (to complement follow-the-dot pages), or a Region Filling Tool (by the texture edges, the mean color value or the original content)³, as shown on Figure 9.

3 Spot-the-difference

The basic idea of a spot-the-difference game is simple: given two similar images, the player is asked to find the differences that have been voluntarily introduced by the artist/designer. The creation of spot-the-difference images is a complex process that requires

³The main goal of such interaction can be visual quality enhancement but also can provide extra fun to the children. Indeed, in the case of interactive coloring of the page this can also be seen as a “magical pencil” that allows the child to fill image regions with the original content of the image instead of coloring it.

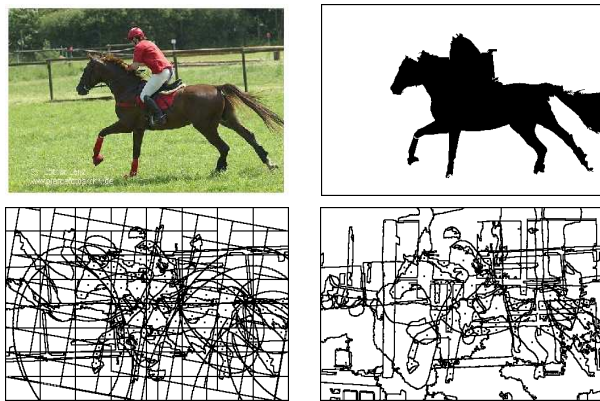


Figure 8: Example results for hidden object sheet. We used two different image partitioning strategy (many others can be used): partitioning the image by random parallel lines and random ellipses (third image) and respectively combining a few existent coloring pages as we have closed regions for them (fourth image). In both cases the dots were added to the regions that had a minimum of 70% overlap with the estimated “horse + person” mask (second image).

several manual steps. These steps generally requires some advanced expertise in digital image retouching. However, the current result from the computer vision research communities can be used to simplify this step. In this section, we propose an automatic or semi-automatic system that would help a non-expert user to generate quickly a pair of spot-the-difference images from arbitrary images, generally photographs.

The main idea is as follows. Starting from a given image I , our goal is to find a transformation T so that $(I, T(I))$ is a pair of spot-the-difference images, meaning that I and $T(I)$ are dissimilar only of a set of K local regions of the image⁴. We assume that the transformation $T = T_1 \circ T_2 \circ \dots \circ T_K$ is a composition of K local transformations T_1, T_2, \dots, T_K . We propose to sequentially introduce the local transformations in the following algorithm:

- 1 initialize $J_0 = I$,
- 2 for $k = 1, \dots, K$, choose a random local transformation T_k such that:
 - (a) its support does not overlap with the previous local transformations T_1, T_2, \dots, T_{k-1} .
To ensure this we create a modification

⁴In France, K is classically equal to 7 so that the game is known under the name “the seven differences”.



Figure 9: Example results for the system interactive region filling with original content.

mask, in which the morphologically dilated version of all the previous supports $(1, \dots, k - 1)$ are added (see Figure 13).

- (b) optionally, verify some constraints on the quality of the modification.
- (c) apply the transformation $J_k = T_k(J_{k-1})$.

The key part of the algorithm is the generation of the local transformations. They are randomly generated by (1) choosing a type of transformation (object insertion, color change, etc.), (2) selecting the boundary of the modified region (such the mask of the object in the case of object insertion) and (3) the parameters of the transformation (the scale of the inserted object, the modified color, etc). Inspired by existing spot-the-difference games, we identified the following transformation types:

- *Geometric transformation.* It can be a rigid local transformation, such as a rotation, translation, scaling, stretching, or a non-rigid transformation (e.g. a change of human posture).
- *Object insertion.* Segmented objects from a given database (there are huge available databases of isolated objects) can be inserted in an image. An important technical point is the use of a class-based segmentation algorithm to insert an object at the right position. E.g. Figure 12 the insertion of a car is only possible on a region segment classified as road as shown on Figure 12. The insertion can be photo realistic (e.g the hand in [BRB10]) or non-photorealistic one (e.g. the inserted car in Figure 12). Further analysis of the scenes, with some 3D and prior knowledge or an

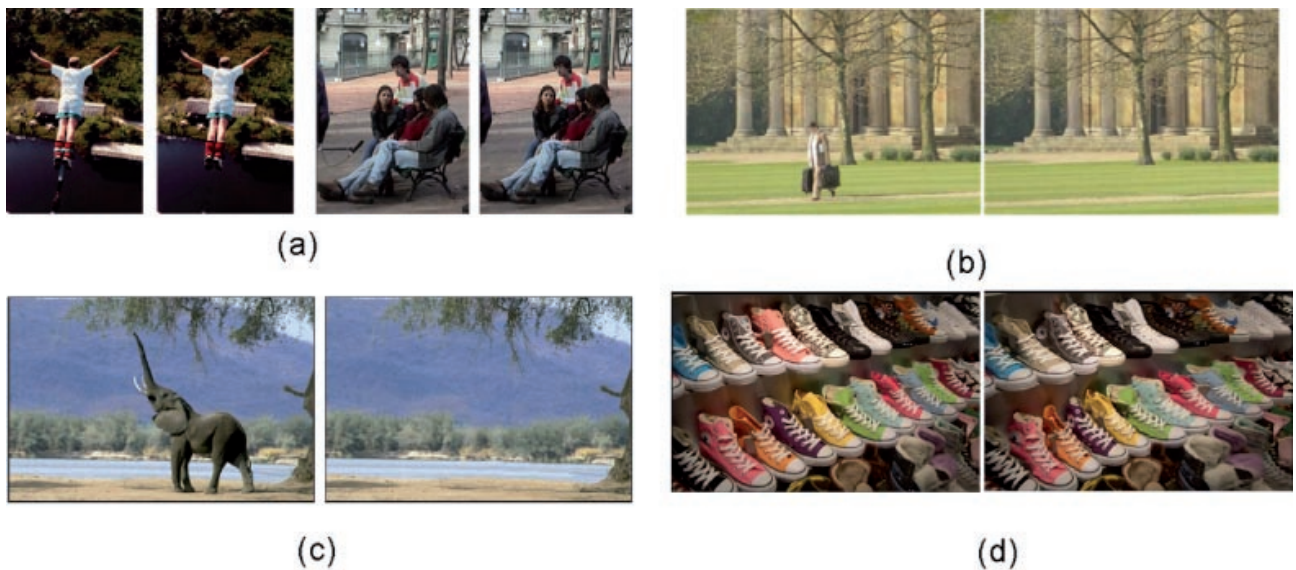


Figure 10: In [BSCB00] the inpainting technique is based on smoothly propagation of the surrounding areas information in the isophotes direction (a). In [CPT04] the authors propagate not only texture, but also structure information using a block-based sampling process (b). [DCoY03] and [AS07] proposes iterative inpainting approaches. While in [DCoY03] the removed region are completed by adaptive image fragments, based on their visual similarity (c), in [AS07] the object removal is done by repeatedly carving out seams based on image energy function (d).

interactive system can help to make the insertion more photo realistic if requested.

- *Deletion of an object or replacement of a region.* There were recent methods proposed that address the problem of object removal or replacement in different application scenario [BSCB00, DCoY03, CPT04, AS07, HE07]. We show a few remarkable examples from the literature to illustrates the strength and the appropriateness of this approach for our system (e.g. Figure 10 or Figure 11).
- *Color aspect change.* This may include the replacement or modification of the colors (e.g. add more red, diminish the green, etc) to a region, transfer the colors from another region (e.g. in the neighborhood), the modification of the brightness, saturation, etc of a given region. Again the knowledge of the “semantic content” of these region can be helpful to make the changes more photo realistic.

Interactivity By default, the transformation choice is fully random. To allow more advanced customization of the generated images, we can allow the user



Figure 11: The facade of the right building was automatically replaced by finding similar image regions in a huge database [HE07].

to act on the current image: before validating (or not) the modification, the user can play with the proposed modifications. In the case of image insertion, the user can modify the scale, orientation and blending factor. In the case of color segment modification, the user can modify the contrast, luminosity or gamma value of the selected fragment. He can also modify (extend, move or reduce) the shape of the selected region. He can further select an object for deletion or a region for replacement. This type of interaction may be very interesting in a real application, as it enables



Figure 12: Example of output from our system for $K = 3$ transformations (here object insertion). Last image corresponds to the semantic class based level segmentation obtained with the CBIS proposed in [CP08].

the user to actively participate in the creation of the spot-the-difference pair, while maintaining a significant speedup in the process, compared to the manual introduction of the changes. This enable a rapid choice of modification, and compared to the manual system, the user does not have to worry about choosing a right modification and does not need specific photofinishing skills. This typically leads to a win of several minutes in the creation process.

Implemented system Our current prototype includes only color changes and object insertions. The color changes correspond to random changes of the saturation, contrast and color replacement of random connex regions of the image. The difficulty of the game can be tuned by setting different levels of color change. Insertions were randomly selected in a small database of categorized objects with transparency masks. The insertion of objects are conditioned on the category of background, so that boats are inserted on water, cars on roads, trees on grass, etc. The last image in Figure 12 shows that our system properly identified the categories road (yellow), sky (light blue) and buildings (brown). We also see that the person inserted in the top left window was due to an error of the class-based segmentation algorithm (the window is categorized as road). As this examples show, there is further need not only for better region categorization but also a better understanding of the image and its elements, such as image depth, orientation and size of a vehicle compared to other objects, etc.

An example of game generated using in addition interactivity with our proposed method is shown on Figure 13.

4 Conclusion

We propose two systems that partially or completely automate the creation of popular visual games and demonstrate their technical feasibility. The main originalities and advantages of the proposed systems are that the initial image is arbitrary, and that the activity sheets can be created automatically. We also propose different levels of interactivity that depend on the skill requirements we wish to impose.

The approaches are simple and hence could be easily integrated in any photo editing or online web sites. As there is no real ground truth neither benchmark data it is difficult to establish the best parameter set of the system. In our experiments we processed hundreds of images of the Pascal VOC 2007 Challenge as we had for them the CBIS estimates and visually compared them to tune the exemplary parameters reported and used to obtain the images shown in the paper. They are probably not the best choices, and, as a future work, it would be necessary to do intensive user preference studies to better establish the correct values of these parameters.

As a proof of concept, we gave a set of automatically created color pages to a few children (see some of them in Figure 14) who accepted⁵ and enjoyed coloring them for us.

References

- [AMK99] Sadeh Abbasi, Farzin Mokhtarian, and Josef Kittler, *Curvature scale space image in shape similarity retrieval*, *Multi-media Systems* 7 (1999), no. 6, 467–476, ISSN 1432-1882.

⁵We would like to acknowledge Anton (11 years), Gabriel (9 year), Johanna (8 years), Elisabeth (7 years) and Mikhaël (5 years) for their contribution.




Figure 13: Example of spot-the-difference image pair generated using the semi-automatic system. The bottom image shows the dilated mask of the inserted objects.

- [AS07] Shai Avidan and Ariel Shamir, *Seam Carving for Content-Aware Image Resizing*, ACM Transactions on Graphics **26** (2007), no. 3, ISSN 0730-0301.
- [BRB10] BRB, *Online games of the Birmingham Royal Ballet*, 2010, <http://www.brb.org.uk/3571.html>, Last visited July 30th, 2010.
- [Bro07] Stephen Brooks, *Mixed Media Painting and Portraiture*, IEEE Transactions on Visualization and Computer Graphics **5** (2007), no. 13, 1041–1054, ISSN 1077-2626.
- [BSCB00] Marcelo Bertalmio, Guillermo Sapiro, Vicent Caselles, and Coloma Ballester, *Image Inpainting*, Proceedings of the 27th annual Conference on Computer Graphics and Interactive Techniques, 2000, pp. 417–424, ISBN 1-58113-208-5.
- [CB09] Gabriela Csurka and Marco Bressan, *Spare Time Activity Sheets from Photo Albums*, International Conference on Graphics Theory and Applications, 2009, pp. 156–163, ISBN 978-989-8111-67-8.
- [CM02] Dorin Comaniciu and Peter Meer, *Mean shift: A robust approach toward feature space analysis*, IEEE Transactions on Pattern Analysis and Machine Intelligence PAMI 2002 **24** (2002), no. 5, 603 – 619, ISSN 0162-8828.
- [CP08] Gabriela Csurka and Florent Perronnin, *Object Class Localization and Semantic Class Based Image Segmentation*, BMVC, 2008.
- [CPT04] Antonio Criminisi, Patrick Pérez, and Kentaro Toyama, *Region Filling and Object Removal by Exemplar-Based Image Inpainting*, IEEE Transactions on Image Processing **13** (2004), no. 9, 1200 – 1212, ISSN 1057-7149.
- [CS99] Dmitry Chetverikov and Zsolt Szabo, *A simple and efficient algorithm for detection of high curvature points in planar curves*, Workshop of Austrian Pattern Recognition Group, 1999.
- [DCoY03] Iddo Drori, Daniel Cohen-or, and Hezy Yeshurun, *Fragment-based image completion*, Proceedings of the 29th annual Conference on Computer Graphics and Interactive Techniques, 2003, pp. 303 – 312, ISBN 1-58113-709-5.
- [DS02] Doug DeCarlo and Anthony Santella, *Stylization and Abstraction of Photographs*, Proceedings of the 29th annual Conference on Computer Graphics and Interactive Techniques, 2002, pp. 769 – 776, ISBN 1-58113-521-1.
- [EGW⁺07] M. Everingham, L. Van Gool, C. Williams, J. Winn, and A. Zisserman, *The PASCAL Visual Object Classes Challenge*, 2007, <http://www.pascal-network.org/challenges/VOC>, Last visited July 30th, 2010.
- [GBT00] Carsten Garnica, Frank Boochs, and Marek Twardochlib, *A new Approach to Edge-preserving Smoothing for Edge Extraction and Image Segmentation*, International Archives of Photogrammetry and Remote Sensing (Amsterdam), vol. XXXIII, Part B3, 2000, pp. 320–325.

- [HE07] James Hays and Alexei A. Efros, *Scene Completion Using Millions of Photographs*, ACM SIGGRAPH 2007 papers, no. 4, 2007, ISSN 0730-0301.
- [HEMK98] Kostas Haris, Serafim N. Efstratiadis, Nicos Maglaveras, and Aggelos K. Kat-saggelos, *Hybrid image segmentation using watersheds and fast region merging*, IEEE TIP **7** (1998), no. 12, 1684–1699, ISSN 1057-7149.
- [HK05] Simon Hermann and Reinhard Klette, *Global Curvature Estimation for Corner Detection*, Tech. Report 171, Communication and Information Technology Research, The University of Auckland, 2005.
- [HY04] Xiaochen Chen He and Nelson H.C. Yung, *Curvature Scale Space Corner Detector with Adaptive Threshold and Dynamic Region of Support*, ICPR, 2004, pp. 791–794, ISBN 0-7695-2128-2.
- [JJ00] Shi Jianbo and Malik Jitendra, *Normalized Cuts and Image Segmentation*, PAMI **22** (2000), no. 8, 731, ISBN 0-8186-7822-4.
- [LS90] Hong-Chih Liu and Mandyam D Srinath, *Corner detection from chain code*, Pattern Recognition **23** (1990), no. 1-2, 51–68, ISSN 0031-3203.
- [Mou03] David Mould, *A stained glass image filter*, 14th Eurographics Workshop on Rendering, 2003, ISSN 1727-3463, pp. 20–25.
- [OK06] Adriana Olmos and Frederick Kingdom, *Automatic non-photorealistic rendering through soft-shading removal: a colour-vision approach*, ICVVG, 2006.
- [PM91] Pietro Perona and J. Malik, *Scale-space and edge detection using anisotropic diffusion*, PAMI **12** (1991), no. 7, 629–639, ISSN 0162-8828.
- [RKB04] Carsten Rother, Vladimir Kolmogorov, and Andrew Blake, *GrabCut: Interactive Foreground Extraction using Iterated Graph Cuts*, SIGGRAPH, 2004, pp. 309–314.
- [RLGR⁺02] P. Reche-Lopèz, Cristina Urdiales Garcia, Antonio Bandera Rubio, Carmen de Trazegnies, and Francisco Sandoval Hernández, *Corner detection by means of contour local vectors*, Electronics Letters **38** (2002), no. 14, 699–701, ISSN 0013-5194.
- [SWRC06] Jamie Shotton, John Winn, Carsten Rother, and Antonio Criminisi, *TextonBoost: Joint Appearance, Shape and Context Modeling for Multi-Class Object Recognition and Segmentation*, ECCV, vol. Lecture Notes in Computer Science Vol. 3951, 2006, ISBN 978-3-540-33832-1.
- [TL04] Thanh Hai Tran Thi and Augustin Lux, *A Method for Ridge Extraction*, ACCV, 2004.
- [WOG06] Holger Winnemöller, Sven. C. Olsen, and Bruce Gooch, *Real-Time Video Abstraction*, SIGGRAPH, 2006, pp. 1221–1226, ISBN 1-59593-364-6.
- [YMF07] Lin Yang, Peter Meer, and David J. Foran, *Multiple class segmentation using a unified framework over mean-shift patches*, CVPR, no. 1-8, 2007, ISBN 1-4244-1179-3.

Citation
Gabriela Csurka, Guillaume Bouchard and Marco Bressan, <i>Spare Time Activity Sheets from Photo Albums</i> , Journal of Virtual Reality and Broadcasting, 7(2010), no. 10, October 2010, urn:nbn:de:0009-6-26703, ISSN 1860-2037.




Figure 14: Example of coloring pages colored by Anton (11 years), Gabriel (9 year), Johanna (8 years), Elisabeth (7 years) and Mikhaël (5 years).