The Mountain Habitats Segmentation and Change Detection Dataset

Frédéric Jean, Alexandra Branzan Albu, David Capson, Eric Higgs, Jason T. Fisher, Brian M. Starzomski
University of Victoria, Victoria, BC, Canada
{fjean,aalbu,capson,ehiggs,starzom}@uvic.ca, jason.fisher@albertainnovates.ca

Abstract

In this paper, we present a challenging dataset for the purpose of segmentation and change detection in photographic images of mountain habitats. We also propose a baseline algorithm for habitats segmentation to allow for performance comparison. The dataset consists of high resolution image pairs of historic and repeat photographs of mountain habitats acquired in the Canadian Rocky Mountains for ecological surveys. With a time lapse of 70 to 100 years between the acquisition of historic and repeat images, these photographs contain critical information about ecological change in the Rockies. The challenging aspects of analyzing these image pairs come mostly from the perspective (oblique) view of the photographs and the lack of color information in the historic photographs. The baseline algorithm that we propose here is based on texture analysis and machine learning techniques. Classifier training and results validation are made possible by the availability of expert manual ground-truth segmentation for each image. The results obtained with the baseline algorithm are promising and serve as a reference for new and improved segmentation and change detection algorithms.

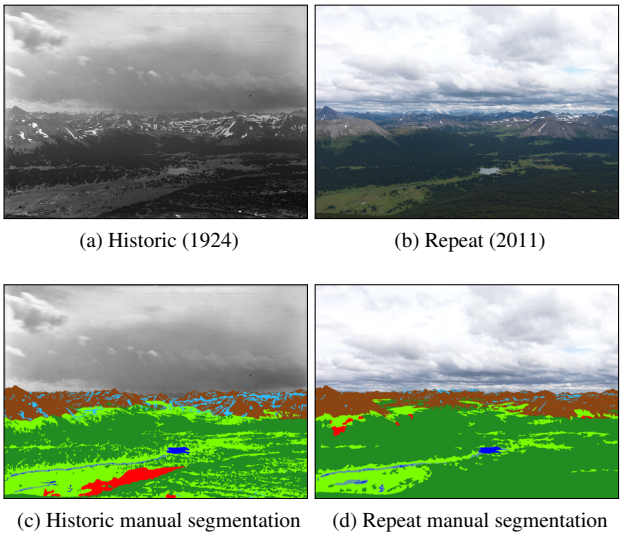

(a) Historic (1924) (b) Repeat (2011)

(c) Historic manual segmentation (d) Repeat manual segmentation

Figure 1: Example of historic and repeat photograph from the Kakwa-Willmore Interprovincial Park, along with their corresponding manual segmentation.

1. Introduction

Change detection in our biogeophysical environment using remotely sensed data provides valuable information that can be used to improve resource and environmental management and ultimately to better understand the human impact on ecological phenomena [11]. The process of change detection consists in identifying changes in phenomena or objects of interest over some temporal range. For instance, remote sensing data have been used to determine changes in land cover or land use, deforestation, regeneration and logging, forest fire and fire affected area, etc [19, 8]. For these types of applications, aerial and satellite imagery provide a useful source of remotely sensed data. The acquired images typically cover large areas, and the top view simplifies image analysis by minimizing the effects of perspective distortion. The main drawback to satellite imagery is the lack of temporal range; as this technique is relatively new, its temporal range is limited to the last 40 years. Also, old aerial photographs are not available for some areas of interest. In order to perform change detection over an extended time range, oblique ground-based photographs can be used, which are more commonly available and span a longer historic time interval. Oblique photography is a type of ground-based photography in which the camera axis is deliberately kept tilted from the vertical by a specified angle (90 degrees in most cases).

The Mountain Legacy Project (MLP) [15] hosts the world's largest collection of historic mountain photographs (about 140,000), which were taken between 1888 and 1958. These photographs originate from federal and provincial photo-topographic surveys, and thus their initial purpose was to create topographic maps of the mountainous regions of Canada. The photographs were used alongside survey measurements to create accurate topographic maps more

efficiently. The technology used at that time consisted of fixed lens box cameras that exposed images on glass plate negatives (approximately $102\,\text{mm} \times 152\,\text{mm}$), which offer exceptional resolution and longevity. Those glass plates are scanned on professional quality flat-bed scanners at 2400 dots per inch (dpi) producing 16 bits TIFF images, which are then cropped and inverted for the purposes of analysis. A subset of the historic photographs (more than 5,000) have been repeated between 2003 and 2012. Repeat photographs are high-resolution digital images (from 16 to 39 megapixels) acquired using the same viewpoint and field of view as their historical counterparts. They have been manually registered to their historical counterparts for comparison and change detection purposes. The resulting image pairs have been studied for climate change detection, as well as for a variety of research and management-focused projects on fire history, vegetation change, human activity and ecological restoration [20, 12].

An example image of a historic photograph and its corresponding repeat photograph are shown in Figures 1a and 1b, respectively. Since the two photographs were not taken in the same lighting conditions, and were not initially acquired using the same imaging process (glass plate emulsion vs digital camera), they cannot be easily compared for changes on a pixel per pixel basis. Thus, change detection must to be performed on a higher level by first segmenting both images according to predefined categories and then by comparing the resulting segmented images. Figures 1c and 1d show a manual segmentation of the historic and repeat images. Here, the categories are defined as habitat types and are depicted using different colors.

The manual segmentation of image pairs is an extremely time consuming process (more than 2 hours per image for an expert annotator). Given the number of actual and future image pairs to analyze, it has become obvious that this process should be automated as much as possible. However, the automated segmentation and analysis of this kind of images are very challenging tasks. The challenging aspects come mainly from the perspective (oblique) view of the photographs in combination with the mountainous topography of the scene (i.e. non-planar scene under perspective projection). Other challenging aspects include the lack of color information for the historic photographs and the occasional but visible damages to the glass plates (mostly scratches and localized emulsion defects). Several of those challenging aspects are not found in usual remote sensing datasets. Indeed, the vast majority of remote sensing datasets are based on satellite or aerial imagery, for which perspective projection is negligible given the top-down view and the large difference between the scene depth and the distance of the camera from the scene (i.e. weak perspective projection).

We present in this paper a novel and challenging remote sensing dataset along with a baseline algorithm for habitats segmentation and change detection. The dataset consists of historic and repeat image pairs along with their corresponding manual segmentation into 8 habitat categories. The proposed baseline algorithm is based on texture analysis and machine learning techniques. The algorithm uses a subset of the image pairs for training, and the remaining image pairs for testing. Results of the baseline algorithm are promising and will be useful as reference and basis for comparison for future segmentation and change detection algorithms.

The remainder of this paper is structured as follows. Related work is presented in Section 2. The dataset and the proposed baseline algorithm are described in Sections 3 and 4, respectively. Experimental results are discussed in Section 5, and Section 6 draws conclusions and outlines future work directions.

2. Related Work

Change detection techniques in remote sensing are classified mainly in two groups [6]: pixel-based and object-based. Pixel-based techniques aim at directly measuring the changes between bitemporal or multitemporal images on a pixel per pixel basis, and in most cases do not consider the spatial context around the pixels. On the other hand, object-based techniques aim first at defining objects in the images, then use these objects to perform change detection. Most approaches in the literature are pixel-based techniques that perform image transformations to obtain a map of binary changes (change vs no change) or a matrix describing the category transitions (from-to). For instance, vegetation indexes [22, 17], change vector analysis [16], principal component analysis [2], and texture-based analysis [21] are used to perform direct change detection between temporal images. The approaches that generate the matrices of changes usually perform post-classification comparison [23, 3], that is, the pixels in each image are first classified in one of the predefined categories and the resulting segmented images are compared. Approaches using object-based techniques usually focus on detecting specific changes, for instance detecting the presence or absence of vehicles [14] or changes in grassy strips [9].

An implicit assumption made by change detection approaches is that the temporal images have been acquired by the same type of sensors within a period of time that is relatively short compared to the lapse of time separating the historic and repeat images that we use in this paper. Several approaches are thus designed to be used with specific images, for instance satellite imagery such as Landsat TM and ETM+ [22, 17, 16, 2, 23, 3], SPOT-5 [2], or satellite and aerial color/gray-scale images [21, 14, 9]. The historic and repeat images presented here show quite different pixels intensities and texture for a given unchanged

Meta-category	Category	ID	Color
Forest(F)	broadleaf/mixedwood (B-MW)	1	orange
	coniferous forest (CF)	2	dark green
	regenerating areas (RA)	8	red
Non-forest (NF)	ice/snow (I-S)	4	light blue
	sand/gravel/rock (S-G-R)	16	brown
	upland herbaceous (UH)	32	light green
	water (WT)	64	blue
	wetland (WL)	128	turquoise

Table 1: Habitat categories and meta-categories.

area. This is mainly due to the different imaging sensors used for acquisition, the difference in weather conditions, and the perspective projection along with the non-planar nature of the scenes. Therefore, direct pixel-based change detection techniques are inappropriate for this type of imagery. Approaches that would be appropriate for the proposed dataset are most likely post-classification change detection approaches where the historic and repeat images are segmented independently according to predefined categories, and compared for change detection using the resulting segmentations.

3. Dataset Overview

The proposed dataset consists of 60 pairs of historic and repeat images that have been manually segmented by field experts. The images and their manual segmentations are released under a *Creative Commons License* and can be freely downloaded from an open access repository [7]. The image resolutions range from 16 to 24 megapixels. Each pair of historic and repeat images is manually registered (aligned) in order to allow for change detection. Each image in a pair comes with a color and a gray-scale segmentation mask obtained from the manual segmentation process. The color segmentation mask is mainly useful for display purposes, whereas the gray-scale segmentation mask is used for selecting pixels of a given category in the corresponding image.

Table 1 presents the eight categories that were defined according to the broad categories of habitats present in the Rocky Mountains. Each category has been assigned a unique color as well as a power-of-two category ID, which corresponds to the pixel values in the gray-scale segmentation mask. Using power-of-two category IDs allows for easy grouping of categories to simplify the segmentation problem or to perform hierarchical segmentation. The category ID zero (black) is a special category that represents uncategorized pixels (e.g. anything not related to habitats). This category is represented by transparent black pixels in the color segmentation mask. The eight habitat categories are also grouped into two meta-categories that represent *forest* and *non-forest* habitats. This habitat grouping into two meta-categories is relevant since there is much interest in the rate at which the tree line may advance or forest patterns may change in response to climate change [4].

4. Proposed Baseline Algorithm

A visual analysis of the image pairs reveals that the pixel intensities alone are not enough to distinguish between the different categories. Since the visual differences between the habitat categories seem to come from complex patterns of pixel intensities, the proposed algorithm is based on texture analysis. We propose here an algorithm that will serve as a baseline for habitats segmentation of the images in the proposed dataset. The algorithm is used with two different texture descriptors: Histograms of Local Binary Patterns (HLBP) and Histograms of Oriented Gradients (HOG). We have chosen these texture descriptors since they are both robust to changes in intensities.

4.1. General Procedure

Given the high resolution of the images in the proposed dataset, each image is divided into blocks of $K \times K$ pixels. For instance, an image of size $M \times N$ is divided into $\lfloor M/B \rfloor \times \lfloor N/B \rfloor$ blocks. The baseline algorithm segments the images by assigning a habitat category to each block of pixels in an image. The ground-truth habitat category for a given block of pixels is defined as the category that has the largest number of pixels within the corresponding block in the manual segmentation image.

The first step of the algorithm consists in partitioning the dataset into a *training set* and a *test set*. Texture features are then extracted from the image pairs in the *training set* to generate feature vectors according to the specific texture descriptor in use. For repeat images, a gray scale version of the image is used to extract the features so that the same algorithm can be applied to both historic and repeat images. Each feature vector is associated with the ground-truth habitat category that is defined at the corresponding block location. The extracted feature vectors and their corresponding habitat category are then used to train two independent classifiers, that is, one for the historic images, and one for the repeat images. This is justified by the difference in the visual appearance of the habitats between the historic and repeat images, which can be explained by the difference between the imaging processes (i.e. glass plate emulsion vs digital camera). Once the classifiers have been trained, the images in the *test set* are segmented and the generated segmentations are compared to the manual segmentation to assess the algorithm performance.

4.2. Histograms of Local Binary Patterns

The Local Binary Pattern (LBP) [18] is a simple yet powerful texture descriptor that is invariant to monotonic intensity changes. For a given position in an image, a binary

pattern is computed by considering P evenly spaced positions x_p for a circular neighborhood of radius R. The LBP code at a given position is computed as

$$\text{LBP}_{P,R} = \sum_{p=1}^{P} s(x_p - v)\, 2^{p-1}, \qquad (1)$$

where x_p is the gray scale intensity at position p in the circular neighborhood, and $s(x)$ is defined as 1 if $x \geq 0$, and as 0 otherwise. The original LBP operator [18] uses $v = x_0$ (the value of the center pixel), whereas the version of the LBP descriptor presented in [10] uses $v = \mu = \frac{1}{P}\sum_{p=0}^{P} x_p$, that is, the mean gray scale value (including the central position). Here we use $v = \mu$ since better performance and more consistent results have been reported [10].

The number of distinct patterns (LBP codes) that can be obtained is 2^P. One way to reduce the number of codes is to consider only "Uniform-2" patterns [18], which are the patterns having no more than two bit transitions ($0 \rightarrow 1$ and $1 \rightarrow 0$). Patterns having more than two transitions are grouped under a special LBP code, which reduces the total number of codes to $P(P-1)+3$. Another improvement is the "Rotation Invariant Uniform-2" patterns [18], where the total number of codes is reduced to $P+2$ by grouping all rotations of a "Uniform-2" pattern under the same LBP code.

Once an LBP code has been obtained for each pixel in an image, a histogram of LBP (HLBP) codes is be computed within each image block. The normalized version of the histogram (sum of bins equal to 1) is then used as a feature vector for the corresponding image block.

4.3. Histogram of Oriented Gradients

A Histogram of Oriented Gradients [1] is computed within an image block by accumulating the gradient magnitude into accumulator bins according to the gradient orientations. For a given pixel in a block, the basic idea is to add the gradient magnitude to the histogram bin that corresponds to the gradient orientation for that pixel. The B bins are evenly spaced in the range 0°–180°(unsigned gradient orientation). As in [1], we use the operator $[-1, 0, 1]$ to compute the gradient magnitude and orientation at each pixel. In order to reduce aliasing, the gradient magnitude at a given pixel is distributed to the three surrounding blocks according to a trilinear interpolation scheme.

The normalization process described in [1] accounts for local contrast and intensity variation in the image. This normalization process is performed within cells of $C \times C$ blocks, and each cell is normalized by concatenating the histogram of each block into a vector, which is then normalized by its Euclidean norm. The histogram of each block is extracted from this normalized vector. Since multiple histograms might be obtained for a given block (when the block belongs to multiple cells), we use the average of the histograms for that block. This normalization process has been tested in this paper along with a standard histogram normalization for each block (sum of bins equal to 1).

5. Experimental Results

This paper reports results from a preliminary study that aims at identifying the *forest* class with respect to everything else (*non-forest*). Focusing on *forest* class allows us to identify the most significant ecological changes in the habitats. These changes are driven by various processes such as plant succession, climate change, and land conversion. Manual segmentation of all habitat categories is made available in order to enable other researchers to address either binary (*forest* vs *non-forest*) or multi-category classification problems.

The dataset consists of 60 image pairs numbered from 1 to 60. Our proposed baseline approach uses a partition of the dataset into a training set of 13 image pairs and a test set of 47 image pairs. This relatively high number of training image pairs allows us to extract tens of thousands of feature vectors, which provide a wide variety of exemplars to train the classifiers. This variety is needed because of the presence of large texture variations within each image caused by the perspective projection and the different lighting conditions (e.g. shadows cast by clouds). The following image pairs have been selected to be part of the training set since they include a balanced proportion of *forest* and *non-forest* pixels: 24, 29, 31, 35, 42, 43, 45, 47, 49, 50, 51, 57, and 58.

The training and classification processes are performed independently on the historic and repeat images using Support Vector Machine (SVM) classifiers. In order to limit the training time, a maximum of 15,000 feature vectors are retained for each category. The SVMs are trained using a 3-fold cross validation on the retained feature vectors and a Gaussian radial basis function as the kernel. The kernel scaling factor γ and the trade-off parameter C are optimized by performing a grid search over $\gamma \in \{0.1, 0.01, 0.001, 0.0001\}$ and $C \in \{1, 10, 100, 1000, 10000\}$.

A total of 42 experiments is performed to test several combinations of the HLBP and HOG texture descriptor parameters. First, both texture descriptors were tested with three different block sizes: $K \in \{24, 32, 40\}$. For HLBP, the number of positions around the center pixel has been fixed to $P = 8$ to limit the dimensionality of the histograms. We performed experiments by varying the radius ($R \in \{1, 2\}$) and different patterns: Rotation Invariant Uniform-2 patterns (RIU2, 10 codes), Uniform-2 patterns (U2, 59 codes), and all patterns (ALL, 256 codes). In the case of HOG, the experiments are performed by varying the number of bins ($B \in \{9, 18\}$), by performing interpolation (I) or no interpolation (N-I), and by normalizing with

	Experiments (descending MCC order)	MCC	F_1-score Forest	F_1-score Non-forest
Repeat	HLBP $K = 40$, $R = 2$, ALL	0.7028	0.8801	0.8185
	HLBP $K = 40$, $R = 1$, ALL	0.6919	0.8783	0.8113
	HLBP $K = 40$, $R = 1$, U2	0.6901	0.8735	0.8113
	HLBP $K = 40$, $R = 2$, U2	0.6891	0.8719	0.8108
	HLBP $K = 32$, $R = 2$, ALL	0.6850	0.8734	0.8080
	HLBP $K = 32$, $R = 1$, ALL	0.6784	0.8724	0.8035
	HLBP $K = 32$, $R = 1$, U2	0.6736	0.8673	0.8016
	HOG $K = 40$, $B = 9$, C-N, I	0.6670	0.8694	0.7958
	HOG $K = 40$, $B = 18$, C-N, I	0.6645	0.8665	0.7952
	HOG $K = 32$, $B = 18$, C-N, I	0.6623	0.8638	0.7945
	HLBP $K = 32$, $R = 2$, U2	0.6602	0.8591	0.7941
	HOG $K = 40$, $B = 9$, H-N, I	0.6598	0.8665	0.7915
	HOG $K = 40$, $B = 18$, H-N, I	0.6576	0.8623	0.7915
	HOG $K = 32$, $B = 9$, C-N, I	0.6569	0.8635	0.7907
	HOG $K = 40$, $B = 18$, C-N, N-I	0.6553	0.8601	0.7905
	HOG $K = 40$, $B = 9$, C-N, N-I	0.6531	0.8613	0.7885
	HOG $K = 32$, $B = 9$, H-N, I	0.6512	0.8613	0.7872
	HOG $K = 24$, $B = 18$, C-N, I	0.6502	0.8580	0.7878
	HOG $K = 32$, $B = 18$, H-N, I	0.6492	0.8588	0.7866
	HOG $K = 32$, $B = 18$, C-N, N-I	0.6484	0.8567	0.7867
Historic	HOG $K = 40$, $B = 9$, C-N, I	0.6364	0.8398	0.7953
	HOG $K = 32$, $B = 18$, C-N, I	0.6344	0.8381	0.7945
	HOG $K = 40$, $B = 9$, H-N, I	0.6339	0.8403	0.7929
	HOG $K = 32$, $B = 18$, H-N, I	0.6270	0.8376	0.7886
	HOG $K = 32$, $B = 9$, C-N, I	0.6262	0.8368	0.7886
	HOG $K = 40$, $B = 18$, C-N, I	0.6247	0.8320	0.7901
	HOG $K = 40$, $B = 9$, C-N, N-I	0.6245	0.8351	0.7883
	HOG $K = 40$, $B = 18$, H-N, I	0.6184	0.8310	0.7858
	HOG $K = 40$, $B = 18$, C-N, N-I	0.6180	0.8260	0.7877
	HOG $K = 32$, $B = 9$, H-N, I	0.6168	0.8328	0.7832
	HOG $K = 24$, $B = 18$, C-N, I	0.6154	0.8303	0.7837
	HOG $K = 24$, $B = 9$, C-N, I	0.6076	0.8293	0.7777
	HOG $K = 24$, $B = 18$, H-N, I	0.6059	0.8288	0.7766
	HOG $K = 40$, $B = 9$, H-N, N-I	0.6045	0.8259	0.7773
	HOG $K = 24$, $B = 9$, H-N, I	0.6024	0.8292	0.7731
	HOG $K = 32$, $B = 9$, C-N, N-I	0.6017	0.8269	0.7742
	HOG $K = 32$, $B = 18$, C-N, N-I	0.5979	0.8189	0.7757
	HOG $K = 40$, $B = 18$, H-N, N-I	0.5979	0.8178	0.7764
	HOG $K = 32$, $B = 9$, H-N, N-I	0.5922	0.8246	0.7673
	HOG $K = 32$, $B = 18$, H-N, N-I	0.5848	0.8166	0.7666

Table 2: The best 20 experiments for historic and repeat images. See text for parameters notation.

the original cell normalization method (C-N, cells of 2×2 blocks), or standard histogram normalization (H-N). There are thus 18 and 24 experiments for HLPB and HOG, respectively, and each of these were performed independently on the historic and repeat images. Given that all the images were systematically taken with a camera that had a level tilt and roll, the upper half of the images is ignored in the experiments as it contains mostly sky pixels (the horizon line passes approximately by the image center point).

Table 2 presents the 20 best experiments for the images in the test set. The results for all of the 42 experiments are provided in the supplemental material and in [7]. For each experiment, the computed segmentation for a given image is compared to the corresponding manual ground truth segmentation. The Matthews Correlation Coefficient (MCC) and the F_1-score of the *forest* and *non-forest* meta-categories are computed for each experiment by considering comparisons performed over all the test images. The MCC [13] is a useful statistic for 2-category problems as it accounts for category imbalance. An MCC value of 1.0 represents a perfect classification, whereas a value of 0.0 represents a classification that is no better than random, and a value of -1.0 represents a systematic misclassification. There is a slight imbalance between the *forest* and the *non-forest* categories in the proposed dataset. Pixels manually categorized as the *forest* category account for 59.2% of the pixels over all repeat images and 53.1% of the pixels over all historic images. The category imbalance can be much greater within each image, where pixels from either category can account for as much as 99% of the categorized pixels.

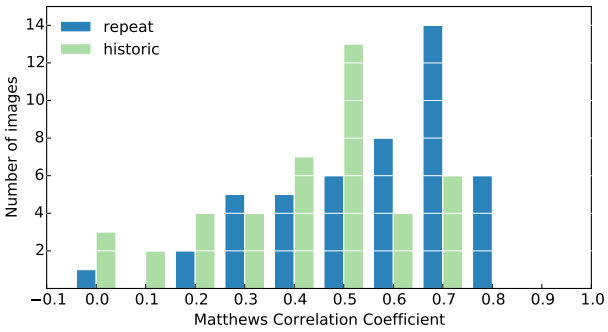

Figure 2: MCC distribution for the best experiments.

It is possible to see that the HLBP texture descriptors perform slightly better for the repeat images, whereas all experiments with the HOG descriptors are systematically better than the HLBP descriptors on the historic images. This confirms that the experiments should be performed independently on the historic and repeat images. A trend that can be observed in these results is that the performances are better when using block size $K = 40$ and $K = 32$ for both HLBP and HOG texture descriptors. Also, the results are consistently better for the repeat images than for the historic images, which can be explained in part by the differences in quality between the repeat and historic images. It is worth mentioning here that all the HLBP experiments using Rotation Invariant Uniform-2 (RIU2) patterns have a poorer performance, with MCC values varying from 0.14 to 0.27 and from 0.02 to 0.09 for repeat and historic images, respectively (see supplemental material). This shows that the textures contained in the dataset are not well represented by a rotation-invariant texture descriptor.

Figure 2 presents the distribution of MCC values in the case of the best repeat and historic experiments, that is, experiment "HLBP $K = 40$, $R = 2$, ALL" and "HOG

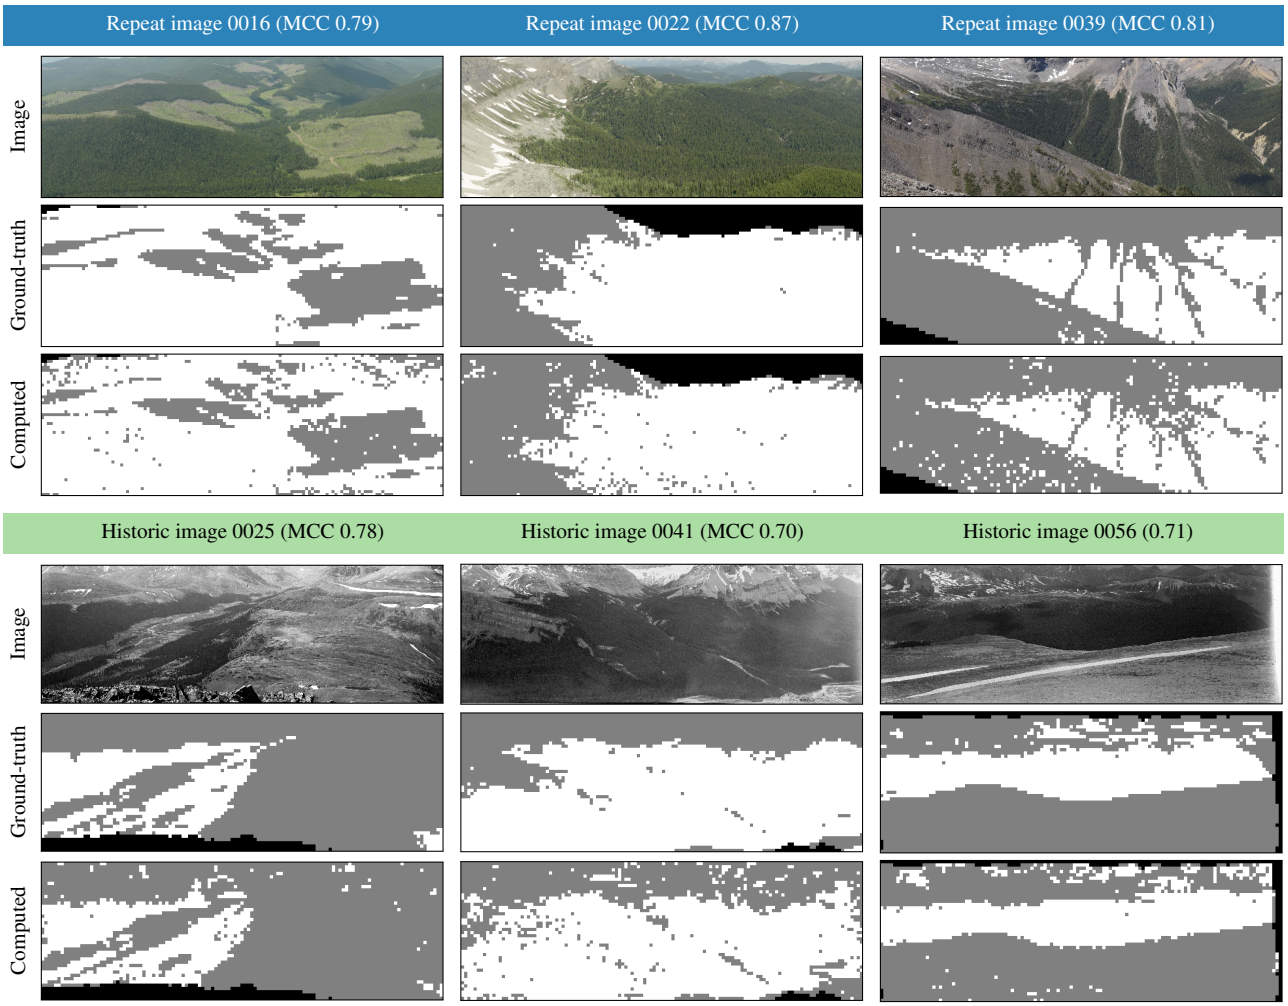

Figure 3: Example of segmentations obtained for the best experiments. White represents *forest*, gray represents *non-forest*, and black represents uncategorized pixels. The MCC is given for each image.

$K = 40$, $B = 9$, C-N, I", respectively (see Table 2). The average MCC is 0.609 for repeat images and 0.430 for historic image, whereas the median MCC is 0.676 for repeat images and 0.469 for historic image. The distribution shows that for a few images, the baseline algorithm generates segmentations that are no better than random (i.e MCC close to 0.0). This is case for the test image pair 54, for which both the repeat and historic images possess a small number of *forest* pixels mostly localized in the same image region. Those pixels were misclassified by the baseline algorithm, whereas most *non-forest* pixels where classified correctly.

Figure 3 presents examples of repeat and historic image segmentations for the best experiment mentioned above. Each example presents the original bottom half of an image, its corresponding manual ground-truth segmentation, and the computed segmentation. The *forest* and *non-forest* meta-categories are represented by white and gray pixels, respectively. Black pixels correspond to image areas for which manual ground-truth segmentation was not performed (i.e. uncategorized pixels). Those areas where thus not considered in the computed segmentation. One should note that the "blocky" appearance of the segmentation images comes from the division of the image into blocks of pixels on which the histograms of LBP or oriented gradients are computed. The overall results are promising given the difficulty of performing texture-based segmentation on these images. Indeed, the variability of the textures within each category makes the segmentation more challenging in the proposed dataset than in other existing databases for texture analysis [5] (e.g. Brodatz, CUReT, VisTex, etc.).

6. Conclusion and Future Work

This paper proposes a an novel and challenging dataset for mountain habitats segmentation and change detection.

We also propose and test a baseline algorithm that provides promising results on the dataset. We hope that the dataset and the results obtained from the baseline algorithm will be useful to other researchers for developing and evaluating the performance of new segmentation and change detection algorithms.

Future research will focus on testing and developing algorithms for segmenting images according to the 8 categories available in the proposed dataset. More image pairs along with their manual segmentation will also be added to the dataset as soon as the manual segmentation becomes available.

Acknowledgments

We are grateful for funding from the Pacific Institute for Climate Solutions, Alberta Tourism, Parks, and Recreation, Alberta Innovates – Technology Futures, Alberta Environment and Sustainable Resource Development, and from the Natural Sciences and Engineering Research Council of Canada.

Historical photographs are courtesy of Library and Archives Canada; Repeat photographs are courtesy of the Mountain Legacy Project.

We would like to thank Morgan Thompson and Jessica McCarron for manually labeling the historic and repeat photographs.

References

[1] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, volume 1, pages 886–893, June 2005.

[2] J. S. Deng, K. Wang, Y. H. Deng, and G. J. Qi. PCA-based land-use change detection and analysis using multitemporal and multisensor satellite data. *Int. J. of Remote Sens.*, 29(16):4823–4838, 2008.

[3] A. Ghosh, N. S. Mishra, and S. Ghosh. Fuzzy clustering algorithms for unsupervised change detection in remote sensing images. *Inform. Sciences*, 181(4):699–715, 2011.

[4] J. Grace, F. Berninger, and L. Nagy. Impacts of climate change on the tree line. *Annals of Botany*, 90(4):537–544, Oct. 2002.

[5] S. Hossain and S. Serikawa. Texture databases – a comprehensive survey. *Pattern Recogn. Lett.*, 34(15):2007–2022, Nov. 2013.

[6] M. Hussain, D. Chen, A. Cheng, H. Wei, and D. Stanley. Change detection from remotely sensed images: From pixel-based to object-based approaches. *ISPRS J. of Photogramm. and Remote Sens.*, 80:91–106, June 2013.

[7] F. Jean, A. Branzan Albu, D. Capson, E. Higgs, J. T. Fisher, and B. M. Starzomski. The Mountain Habitats Segmentation and Change Detection Dataset. Dataset on ZENODO. `http://dx.doi.org/10.5281/zenodo.12590`, November 2014.

[8] X. Juan, L. Zhen, L. Liping, T. Bangsen, and S. Zili. Land cover classification of polarimetric SAR images for the yellow river delta based on support vector machine. In *Proc. Int. Conf. Comput. Vis. in Remote Sens.*, pages 256–261, 2012.

[9] A. Lefebvre, T. Corpetti, and L. Hubert-Moy. Object-oriented approach and texture analysis for change detection in very high resolution images. In *IEEE Int. Geosci. and Remote Sens. Symp.*, volume 4, pages IV – 663–IV – 666, July 2008.

[10] L. Liu, L. Zhao, Y. Long, G. Kuang, and P. Fieguth. Extended local binary patterns for texture classification. *Image Vis. Computing*, 30(2):86–99, Feb. 2012.

[11] D. Lu, P. Mausel, E. Brondzio, and E. Moran. Change detection techniques. *Int. J. Remote Sens.*, 25(12):2365–2401, 2004.

[12] I. S. MacLaren, E. Higgs, and G. Zezulka-Mailloux. *Mapper of Mountains: M.P. Bridgland in the Canadian Rockies, 1902-1930*. Univ. of Alberta Press, 1st edition, Dec. 2005.

[13] B. W. Matthews. Comparison of the predicted and observed secondary structure of t4 phage lysozyme. *BBA Protein. Struct.*, 405(2):442–451, Oct. 1975.

[14] O. Miller, A. Pikaz, and A. Averbuch. Objects based change detection in a pair of gray-level images. *Pattern Recognition*, 38(11):1976–1992, Nov. 2005.

[15] Mountain Legacy Organization. The Mountain Legacy Project. `http://www.mountainlegacy.ca`, 2014.

[16] K. Nackaerts, K. Vaesen, B. Muys, and P. Coppin. Comparative performance of a modified change vector analysis in forest change detection. *Int. J. of Remote Sens.*, 26(5):839–852, 2005.

[17] M.-L. Nordberg and J. Evertson. Vegetation index differencing and linear regression for change detection in a swedish mountain range using landsat TM and ETM imagery. *Land Degradation & Development*, 16(2):139 – 149, 2005.

[18] T. Ojala, M. Pietikainen, and T. Maenpaa. Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *IEEE Trans. Pattern Anal. Mach. Intell.*, 24(7):971–987, July 2002.

[19] C. Petit and E. F. Lambin. Integration of multi-source remote sensing data for land cover change detection. *Int. J. of Geographical Inform. Science*, 15(8):785–803, 2001.

[20] J. Rhemtulla, R. Halls, E. Higgs, and S. Macdonald. Eighty years of change: vegetation in the montane ecoregion of Jasper National Park. *Canadian J. Forest Research*, 32:2010–2021, 2002.

[21] D. Tomowski, M. Ehlers, and S. Klonus. Colour and texture based change detection for urban disaster analysis. In *Urban Remote Sens. Event*, pages 329–332, Apr. 2011.

[22] E. H. Wilson and S. A. Sader. Detection of forest harvest type using multiple dates of landsat TM imagery. *Remote Sens. of Environ.*, 80(3):385–396, June 2002.

[23] F. Yuan, K. E. Sawaya, B. C. Loeffelholz, and M. E. Bauer. Land cover classification and change analysis of the twin cities (minnesota) metropolitan area by multitemporal landsat remote sensing. *Remote Sens. of Environ.*, 98(23):317–328, Oct. 2005.