© IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

This material is presented to ensure timely dissemination of scholarly and technical work. Copyright and all rights therein are retained by authors or by other copyright holders. All persons copying this information are expected to adhere to the terms and constraints invoked by each author's copyright. In most cases, these works may not be reposted without the explicit permission of the copyright holder.

Evaluation of Super-Resolution Methods in the Context of Colonic Polyp Classification

M. Häfner
Department for Internal Medicine,
St. Elisabeth Hospital,
Vienna, Austria

M. Liedlgruber and A. Uhl and G. Wimmer
Department of Computer Sciences,
University of Salzburg,
Salzburg, Austria
Email: {mliedl,uhl,gwimmer}@cosy.sbg.ac.at

***Abstract*—In this work we investigate whether it is possible to improve the results of an automated classification of colonic polyps by using super-resolution algorithms on endoscopic video sequences.**

For this purpose we apply different super-resolution methods to endoscopic sequences and use a set of feature extraction methods for the classification of the SR reconstruction results. We then compare the results obtained from these experiments against the classification results based on original low-resolution frames and against classification rates based on upscaled versions of low-resolution frames.

We show that, at least for the set of super-resolution methods and feature extraction methods evaluated, applying super-resolution methods to the low-resolution frames has no significant impact on the resulting overall classification results.

***Keywords*—Endoscopy, Super-resolution, Polyps, HD, Classification**

I. Introduction

Throughout the past years a variety of different methods for an automated classification of colonic polyps based on endoscopic images has been developed. The majority of these works is based on traditional endoscopes. But there also exists work which is based on imagery acquired using an endoscope with high magnification capabilities (e.g. [1], [2], [3]). One advantage of such endoscopes is that they allow to inspect the colonic mucosa in a magnified manner, thus revealing the fine surface structure of the mucosa as well as small lesions. However, throughout the last few years high-definition (HD) endoscopes got more and more popular. While this type of endoscopes provides a roughly four times higher image resolution as compared to many zoom-endoscopes, they are often not providing magnification capabilities.

One possible way to unveil more details from such HD images would be to use super-resolution (SR) algorithms. In literature there already exists a work which evaluates the application of an SR method to wireless capsule endoscopy video frames [4]. In this work the authors test their algorithm on low-resolution (LR) images generated from a single video frame by shifting it into different directions and downscaling the shifted frames. This, however, does not reflect a realistic application scenario.

In more recent work, SR algorithms are evaluated on endoscopic images in a more realistic setup [5], [6]. Based on patches, extracted from HD images, high-resolution (HR) images are created, using a set of different SR algorithms. The quality of the SR reconstruction is then assessed using different image quality metrics.

(a) Zoom (b) HD (c) SR method

Fig. 1. Illustration of the difference between two different imaging modalities in (a) and (b). And a region of the same size extracted from the outcome of an SR method applied to (b) in (c).

In this work we adapt the basic experimental setup from [5] (i.e. this work is also based on sequences of patches extracted from successive HD video frames). But instead of using a quality metric to assess the quality of the SR results, the HR images are subject to image classification in this work. This way we are able to determine whether applying SR algorithms to our images is beneficial when it comes to a classification of endoscopic images.

Figure 1 shows two tubulovillous adenoma, one captured with a zoom-endoscope and one captured with a HD endoscope without zoom. We immediately notice the dramatic difference between these images in terms of the details visible. It is also obvious that we can not expect the HR images to be comparable to the ones obtained with zoom-endoscopes. However, it can be expected that we at least obtain more details.

The remaining part of this work is organized as follows: In Section II we briefly highlight the challenges we face when applying SR algorithms to endoscopy videos, followed by a description of the SR algorithms and feature extractions methods evaluated in Section III. In Section IV we describe the experimental setup used and present the results obtained. We conclude the paper in Section V.

II. Super-Resolution in Endoscopy

Throughout literature many SR algorithms are evaluated on artificially generated LR images only. That is, although real-world video test sequences are available, the respective

978-1-4799-3990-9/14/$31.00 © 2014 IEEE

sequences are subject to blur and downsampling to generate LR frames (e.g. [4]). These frames are then used to reconstruct an HR image. While this is a practical way to assess if an algorithm works (i.e. an accurate quality assessment is possible since the HR ground truth is available), this hardly matches real-world scenarios.

Since our aim is to reveal new details in endoscopy videos, it is not meaningful to evaluate SR algorithms on artificially generated LR frames. Hence, we directly apply different algorithms to HD videos frames. Doing so, we face different problems:

- **Compression artifacts:** Although there exists work which specifically aims at SR for videos (e.g. [4]), these algorithms are quite often evaluated only on uncompressed sequences. Since HD videos would require a fairly high amount of storage if stored uncompressed, they are usually compressed. This comes at the price of sometimes clearly noticeable compression artifacts (e.g. 8×8 DCT blocks). If the compression level is too high, reconstructing additional information might be challenging, if not impossible.

- **Lack of aliasing artifacts:** When generating LR images artificially, aliasing artifacts may get introduced synthetically (if the LR simulation process does not take care of this). While this facilitates the SR reconstruction, it is hard to tell whether an algorithm eventually works just because of artificially introduced aliasing artifacts. The images we are using in this work show a lack of aliasing artifacts. One reason for this is that the videos are compressed.
  In addition, aliasing artifacts only appear if a signal is sampled using a sampling frequency which is too low. Since, however, in endoscopic images there are usually no sharp edges (i.e. high frequency content) which may result in aliasing artifacts due to undersampling, it is clear that our images do not expose clearly noticeable aliasing. Another cause may be the presence of noise and blur, caused by the sensor and small camera movements.

- **Complex motion:** An accurate motion estimation is very crucial for a high visual quality of the SR reconstruction result. Since in endoscopic videos we are facing highly complex motion (e.g. position-variant transformations and parallax effects) simple motion models are not sufficient to describe the motion between successive HD endoscopy video frames.
  In this work we use the optical flow estimation by Black and Anandan [7], which is part of the implementation available for the work in [8]. While being more complex these methods are also more versatile when it comes to the estimation of arbitrary complex motion between images. This is mainly due to the fact that optical flow methods allow to estimate local motion, while simpler methods usually work well only with global motion. In our application scenario this task is hindered to some extent by compression artifacts and a lack of aliasing artifacts.

III. Methods Evaluated

While other types of SR-algorithms exist too, we restricted our experiments to reconstruction-based algorithms. In the following we briefly describe the SR algorithms and the feature extraction methods which have been employed for the experiments in this work.

A. SR Algorithms

In the following y_k denotes the k-th LR image from the input sequence and $\hat{X}_n$ denotes the HR estimate after the n-th iteration of the respective iteration.

- **Iterative Back Projection (IBP):** The Iterative Back Projection method has been proposed in [9] and was chosen for our experiments due to its simplicity and intuitive nature. Simply stated, this method computes the pixel-wise difference between y_i and $\hat{X}_n$ after applying the respective warp, smoothing, and downsampling. The difference image is then upsampled, followed by computing the gradient image, and warping back the gradient image to the image space of $\hat{X}_n$. The final update for $\hat{X}_n$ is obtained by summing up the gradient images for all y_i in a pixel-wise fashion and adding the resulting image, multiplied by a constant factor, to $\hat{X}_n$.

- **Robust Super-Resolution (ROBZ):** This method, proposed in [10], is basically a modification to the IBP method. Instead of summing up the single gradient images, the authors propose to compute a pixel-wise median to obtain the update weight for each pixel. By changing the IBP algorithm this way, outlier pixels are removed. Such outliers might arise, for example, due to an inaccurate motion estimation.

- **Projection Onto Convex Sets (POCS):** The idea of POCS was introduced to image processing by the work in [11]. The key idea of POCS-based SR algorithms is to express every piece of prior knowledge about the solution as a constraint in image space. More specifically, the solution is constrained by convex sets which, according to the prior knowledge available, impose restrictions on a HR estimate in order to be a valid one. The experiments in this work are based on the POCS-approach proposed in [5] as this method has been developed in the context of endoscopic imaging.

- **Regularized Super-Resolution (RSR):** The RSR method used in this work was proposed in [12]. Since the SR reconstruction problem is an ill-posed one [13], regularized approaches aim at finding the desired HR image in the space of possible solutions by imposing one or more constraints on the SR reconstruction. The algorithm proposed in [12] is in some way similar to the IBP method described, as it also aims at minimizing the error between an observed LR image y_k and a simulated LR image. But besides a different cost function, the approach in [12] uses an additional regularization constraint to compensate for the ill-posedness nature of SR reconstruction problems. The constraint used is termed as bilateral total variation (BTV), which penalizes the total variation within an image with a spatial decaying effect.

For our experiments the initial HR estimate $\hat{X}_0$ is set to an upscaled version of y_1 in case of IBP, ROBZ, and POCS. For RSR, $\hat{X}_0$ is set to the pixel-wise mean of all LR images after registration and upscaling. In addition, in case of RSR, $\hat{X}_0$ is subject to a regularized deconvolution to cope with noise and blur.

Since all the SR methods evaluated work in an iterative manner, we employ the adaptive termination criterion proposed in [5] to decide upon termination of the iterative process.

Figure 2 shows an example outcome of the different SR methods. In this figure the top image shows an example LR image (i.e. y_1) and a region (denoted by the red square in the top image) after upsampling, using bicubic interpolation. The remaining images show the same region from the HR images, resulting from the SR algorithms.

B. Feature Extraction

In this work we evaluated the following set of feature extraction methods for a subsequent classification. These feature extraction methods represent a subset of features we already used successfully in the past for the classification of endoscopic imagery.

- **DT-CWT:** The Dual-Tree Complex Wavelet Transform is used with six scales and six orientations. Based on the absolute values of the detail subband coefficients, the statistical features mean and standard deviation are computed [14]. This process is repeated for each color channel and the resulting values are concatenated to obtain the final vector.
- **Gabor-Classic:** The Gabor Wavelet Transform is used with three scales and six orientations, the mean and standard deviation of the coefficient magnitudes within a subband are used as features [14]. Similar to DT-CWT, the final feature vector is obtained by concatenating the feature vectors generated for each color channel.
- **SSF:** This method has been specifically designed for the classification of colonic polyps. SSF analyzes the shape of connected components (blobs) from images (after a conversion to grayscale) segmented by a variation of the fast level lines transform. But in contrast to the algorithm proposed in [15], in this work we use a slightly modified algorithm, which prevents merging of blobs. The final feature vector of an image consists of the histograms computed from three shape features (convex hull feature, skeletonization feature, and perimeter feature) and a contrast feature extracted from the blobs [15].
- **EF:** After a grayscale conversion of the input image, this method aims at finding regions which correspond to pits, as typically observed on a colonic mucosa. Based on these regions, various different shape features are extracted (e.g. mean pit area and mean perimeter across all pits detected in an image) [16]. While the original work in [16] is based on a feature selection, the experiments in this work have been carried out without such an optimization in order to avoid any overfitting.
- **LBP:** Based on a grayscale image, this operator generates a binary sequence for each pixel by thresholding the neighbors of that pixel by the center pixel value. The binary sequences are then interpreted as numbers (i.e. the LBP numbers). In our case the $LBP_{8,1}$ operator has been used (i.e. eight neighbors with a radius of one). Once all LBP numbers for an image are computed, a histogram based on these numbers is generated and used as feature vector.

While the DT-CWT and the Gabor-Classic method operate in a multi-resolution fashion, the LBP operator works on small pixel neighborhoods only. The methods EF and SSF, in contrast, are specifically designed to analyze shapes. Since the main goal of applying the SR methods is to obtain more detailed images, it is interesting to see, which of these methods is able to exploit additional details.

IV. Experimental Setup and Results

A. Experimental Setup

Each LR sequence used in this work is based on four successive frames taken from 62 videos acquired during colonoscopy sessions between the years 2011 and 2013 at the Department for Internal Medicine (St. Elisabeth Hospital, Vienna) using a HD colonoscope (Pentax HiLINE HD+ 90i Colonoscope) with a resolution of 1280×1024 pixels. In order to acquire the videos, 37 patients underwent endoscopy.

Lesions found during colonoscopy have been examined after application of dye-spraying with indigocarmine, as routinely performed in colonoscopy. Biopsies or mucosal resection have been performed in order to get a histopathological diagnosis. In addition to the topical staining, the Pentax i-SCAN image enhancement has been enabled (i.e. i-SCAN mode 3, which enhances the visibility of pit pattern and vascular features).

To reduce the computational demand for the SR methods we chose positions from which we manually extracted 256×256-pixel patches which serve as LR images (the position remained the same in case of a single sequence). For the SR reconstruction we use an upscaling factor of two. Since the LR images used are color images, we apply the SR algorithms only to the intensity component in the CIELAB color space. The color components of the HR images are obtained by a simple bicubic upscaling of the first frame from the respective LR sequence.

The ground truth for the LR sequences and, as a consequence, for the SR reconstruction results is given in Table I. As we notice from this table, we carry out a classification between non-neoplastic and neoplastic polyps. While a more fine-grained classification would theoretically be possible, this would lead to rather unstable results due to the quite limited number of LR sequences available. Since different types of lesions may develop inside the colon of a single patient such a patient may appear in more than one class. Hence, the total number of patients in Table I is slightly higher (48) as compared to the number of patients who underwent colonoscopy (37).

For the classification we employ the k-NN classifier with different choices for k (i.e. $k = 1, \ldots, 10$). We chose this

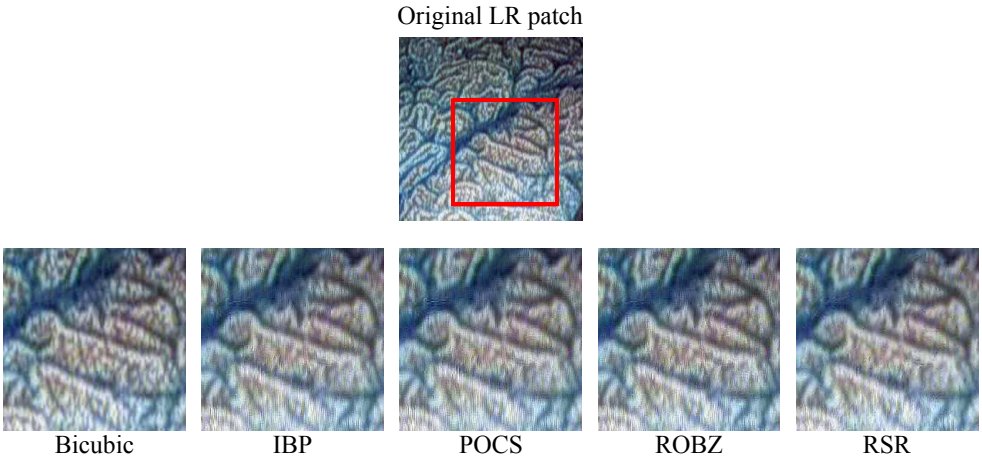

Fig. 2. An example outcome of the different SR algorithms.

TABLE I. GROUND TRUTH INFORMATION FOR THE LR SEQUENCES USED IN OUR EXPERIMENTS.

	Non-neoplastic	Neoplastic	Total
LR sequences	19	43	62
Patients	13	35	48

rather simple classifier since our goal is not to obtain the highest possible rates but to compare rates between different classification scenarios. In addition, in the past we already demonstrated that the k-NN classifier performs quite well in case of endoscopic image classification.

To estimate the overall classification accuracy for the different combinations of SR methods and feature extraction methods we use leave-one-patient-out cross-validation (LOPO-CV). In this setup one image out of the images to be classified is considered as an unknown image. The remaining images are used to train the classifier (omitting those images which originate from the same patient as the image left out). The class of the unknown image is then predicted by the system. These steps (training and prediction) are repeated for each image, yielding an estimate of the overall classification accuracy. While the specificity and sensitivity values would be of interest too, the quite limited number of images used would render the respective numbers meaningless. We therefore decided to omit them in this work.

Since we want to know whether the SR algorithms facilitate the classification of our images we investigated two reference classification scenarios: in the first scenario the first frame of each LR sequence is used for a classification. In the second scenario the first LR frame of each sequence is subject to upscaling using an upscaling factor of two and bicubic interpolation. The resulting images are then used for a classification. Choosing the first LR frame in both cases can be justified by the fact that the first frame is also used as the reference frame when applying the SR algorithms.

It must be pointed out that, by applying the different SR methods and by considering the two reference scenarios, we end up with a total of six different image databases. Each of these databases is then subject to a separate classification (i.e. the images used for training and classification always belong to the same database).

In order to be able to assess whether the classification results based on the HR images are statistically significant different as compared to the reference scenarios outlined above, we employ McNemar's test [17]. For two methods M_1 and M_2 this test statistic keeps track of the number of images which are misclassified by method M_1 but classified correctly by method M_2 (denoted by n_{01}) and vice versa (denoted by n_{10}). The test statistic, which is approximately Chi Square distributed (with one degree of freedom), is then computed as

$$T = \frac{(|n_{01} - n_{10}| - 0.5)^2}{n_{01} + n_{10}}. \quad (1)$$

From T the p-value can be computed as

$$p = 1 - F_{\chi_1^2}(T) \quad (2)$$

where $F_{\chi_1^2}$ denotes the cumulative distribution function of the Chi Square distribution with one degree of freedom. The null-hypothesis H_0 for McNemar's test is that the outcomes of M_1 and M_2 lead to equal error rates. Given a fixed significance level α, there is evidence that the methods M_1 and M_2 produce significantly different results if $p < \alpha$. As a consequence we can reject the null-hypothesis H_0. Throughout this work we chose a significance level of $\alpha = 0.05$. This implies that, if M_1 and M_2 are significantly different, there is a confidence level of 95% that the differences between the outcomes of the methods are not caused by random variation.

B. Results

Figure 3 gives an overview of the results from our experiments (grouped by feature extraction methods in Fig. 3(a) and grouped by SR algorithms in Fig. 3(b)). In these figures the solid black line shows the mean over the overall classification rates obtained with the different choices for k for one specific combination of SR algorithm and feature extraction method. The shaded area shows the minimum and maximum overall rates over all values for k for one combination. "Normal" and "Bicubic" denote the cases where the original and upscaled images are used for classification, respectively (i.e. the reference scenarios as described in Section IV-A).

From Fig. 3 we notice that the mean classification rates are in general lower when classifying upscaled images instead of

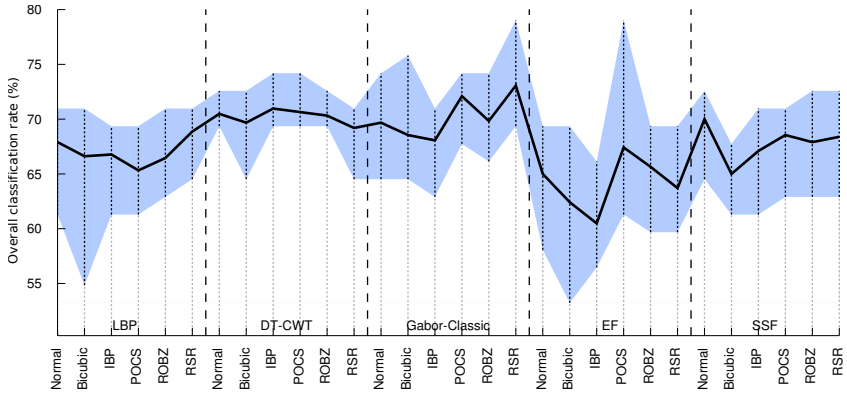

(a) Grouped by feature extraction methods.

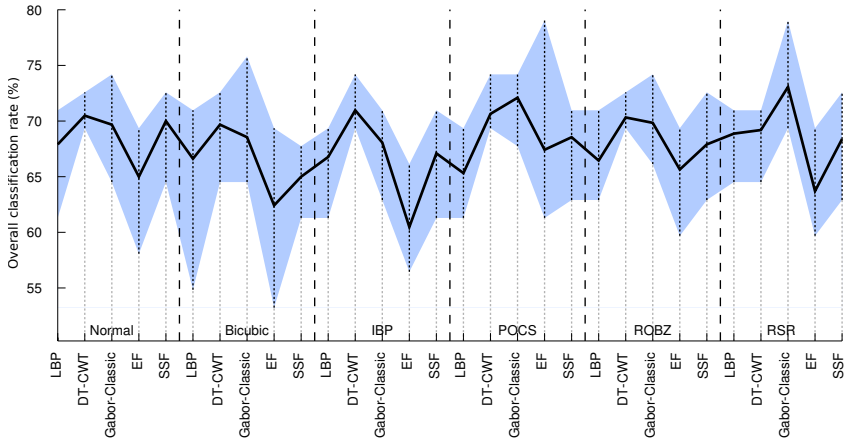

(b) Grouped by SR algorithms

Fig. 3. Overview of the results from our experiments.

the original images. This, however, is no surprise since details get blurred during the bicubic interpolation. When applying SR algorithms, the mean rates are mostly lower as compared to the rates for "Normal" or "Bicubic" in case of the LBP feature. Only the RSR method is able to yield a slightly higher mean overall classification rate. In case of the DT-CWT feature the mean rates are roughly equal. The two remaining feature show a slightly different behavior: in case of the Gabor-Classic feature at least applying POCS and RSR results in a higher classification rates as compared to "Normal" and "Bicubic". When using the EF feature, "Normal" yields the lowest mean overall rate as compared to the other feautures. In most other cases the mean rates are lower or roughly equal. Only the POCS SR method is able to yield a higher mean overall classification rate. The SSF feature, in contrast, behaves different. That is, "Normal" yields the highest mean overall rate, in all other cases the mean overall rates are slightly lower or roughly equal.

When looking at Fig. 3 we also immediately notice that in most cases there is a high variation in the classification rates over the different choices for k. This implies that the selection of k has a great impact on the resulting classification rates. This, however, may be attributed to the quite limited number of LR sequences available for our experiments. The combinations of feature extraction method and SR method, which yield the highest overall classification rates, are RSR/Gabor-Classic and POCS/EF.

Table II shows the classification results in a more detailed manner. This table shows the mean overall classification rates over all choices for k along with the respective standard deviations. From this table we notice that the results for "Bicubic" are consistently lower as compared to the results for "Normal", which is something, as already indicated above, that one might expect. One more important thing we notice is that the result differences between "Normal"/"Bicubic" and the SR methods are in most cases rather small (i.e. below the respective standard deviation). This indicates that the differences between the classification rates are negligible when comparing "Normal"/"Bicubic" and the SR methods.

The results shown in Table III underpin this observation. This table shows whether the differences between "Normal" and "Bicubic" are statistically significant. To perform the analysis behind Table III we fixed the value for k to 3, since for this value the overall classification rates are in most cases higher for the SR methods as compared to "Normal" and "Bicubic". In this table a check mark indicates a significant difference and the sign given in brackets shows whether the respective result is significantly higher (+) or significantly lower (-) as compared to either "Normal" or "Bicubic". From this table we notice that there is no significant difference between the overall classification rates for "Normal"/"Bicubic"

TABLE II. MEAN CLASSIFICATION RESULTS (OVER ALL CHOICES FOR k) ALONG WITH THE RESPECTIVE STANDARD DEVIATIONS (GIVEN IN PERCENT).

	Normal	Bicubic	IBP	POCS	ROBZ	RSR
LBP	67.9 ± 2.8	66.6 ± 5.0	66.8 ± 2.5	65.3 ± 2.9	66.5 ± 2.4	68.9 ± 2.0
DT-CWT	70.5 ± 1.1	69.7 ± 2.1	71.0 ± 1.7	70.6 ± 2.1	70.3 ± 1.1	69.2 ± 2.1
Gabor-Classic	69.7 ± 3.0	68.5 ± 3.8	68.1 ± 2.1	72.1 ± 1.7	69.8 ± 3.0	73.1 ± 2.8
EF	65.0 ± 3.9	62.4 ± 4.4	60.5 ± 3.3	67.4 ± 5.5	65.6 ± 2.8	63.7 ± 2.8
SSF	70.0 ± 2.9	65.0 ± 2.4	67.1 ± 3.1	68.5 ± 2.6	67.9 ± 3.5	68.4 ± 3.2

TABLE III. SIGNIFICANCE TEST RESULTS FOR OUR EXPERIMENTS WITH $k = 3$.

	IBP	POCS	ROBZ	RSR	IBP	POCS	ROBZ	RSR
	Normal				Bicubic			
LBP	✗	✗	✗	✗	✗	✗	✗	✗
DT-CWT	✗	✗	✗	✗	✗	✗	✗	✗
Gabor-Classic	✗	✗	✗	✗	✗	✗	✗	✗
EF	✗	✗	✗	✗	✗	✗	✗	✗
SSF	✗	✗	✗	✗	✗	✗	✗	✗

and the results obtained after applying an SR algorithm.

V. CONCLUSION

We investigated whether applying SR methods to LR sequences from endoscopic videos has an impact on the outcome of the classification of the images.

The classification results obtained indicate that - at least for the set of SR methods and feature extraction methods evaluated - applying SR methods has no real impact on the resulting overall classification results.

This has also been supported by a statistical test, which showed that there are no statistically significant differences between the classification of original frames and frames after applying SR methods.

In future work we therefore will focus on using a larger set of methods (SR methods as well as feature extraction methods). It will also be important to base the experiments in future work on a larger set of LR sequences to be able to make more solid statements about the usefulness of SR methods in the context of an automated classification of colonic polyps. Another interesting question is, whether the SR methods would be useful in case of the classification of uncompressed data.

ACKNOWLEDGMENTS

This work is partially funded by the Austrian Science Fund (FWF) under Project No. TRP-206.

REFERENCES

[1] I. N. Figueiredo, P. N. Figueiredo, G. Stadler, O. Ghattas, and A. Araújo, "Variational image segmentation for endoscopic human colonic aberrant crypt foci," *IEEE Transaction on Medical Imaging*, vol. 29, no. 4, pp. 998–1011, Apr. 2010.

[2] M. Häfner, M. Liedlgruber, A. Uhl, A. Vécsei, and F. Wrba, "Delaunay triangulation-based pit density estimation for the classification of polyps in high-magnification chromo-colonoscopy," *Computer Methods and Programs in Biomedicine*, vol. 107, no. 3, pp. 565–581, Sep. 2012.

[3] J. J. W. Tischendorf, S. Gross, R. Winograd, H. Hecker, R. Auer, A. Behrens, C. Trautwein, T. Aach, and T. Stehle, "Computer-aided classification of colorectal polyps based on vascular patterns: a pilot study," *Endoscopy*, vol. 42, no. 3, pp. 203–207, Mar. 2010.

[4] K. Duda, T. Zieliński, and M. Duplaga, "Computationally simple super-resolution algorithm for video from endoscopic capsule," in *Proceedings of the International Conference on Signals and Electronic Systems (ICSES'08)*, 2008, pp. 197–200.

[5] M. Häfner, M. Liedlgruber, and A. Uhl, "POCS-based super-resolution for HD endoscopy video frames," in *Proceedings of the 26th IEEE International Symposium on Computer-Based Medical Systems (CBMS'13)*, 2013, pp. 185–190.

[6] ——, "Super-resolution techniques evaluated in the context of HD endoscopic imaging," Department of Computer Sciences, University of Salzburg, Austria, http://www.cosy.sbg.ac.at/research/tr.html, Tech. Rep. 2013-04, 2013.

[7] M. Black and P. Anandan, "The robust estimation of multiple motions: Parametric and piecewise-smooth flow fields," *Computer Vision and Image Understanding*, vol. 63, pp. 75–104, 1996.

[8] S. Deqing, S. Roth, and M. Black, "Secrets of optical flow estimation and their principles," in *Proceedings of the 2010 IEEE Conference on Computer Vision and Pattern Recognition (CVPR'10)*, 2010, pp. 2432–2439.

[9] M. Irani and S. Peleg, "Improving resolution by image registration," *CVGIP: Graphical Models and Image Processing*, vol. 53, no. 3, pp. 231–239, Apr. 1991.

[10] A. Zomet, A. Rav-Acha, and S. Peleg, "Robust super-resolution," in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, (CVPR'01)*, no. 1, 2001, pp. 645–650.

[11] D. C. Youla, "Generalized image restoration by the method of alternating orthogonal projections," *IEEE Transactions on Circuits and Systems*, vol. 25, no. 9, pp. 694–702, 1978.

[12] S. Farsiu, M. D. Robinson, M. Elad, and P. Milanfar, "Fast and robust multiframe super resolution," *IEEE Transactions on Image Processing*, vol. 13, no. 10, pp. 1327–1344, Oct. 2004.

[13] S. Baker and T. Kanade, "Limits on super-resolution and how to break them," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 24, no. 9, pp. 1167–1183, Sep. 2002.

[14] M. Häfner, R. Kwitt, A. Uhl, A. Gangl, F. Wrba, and A. Vécsei, "Feature-extraction from multi-directional multi-resolution image transformations for the classification of zoom-endoscopy images," *Pattern Analysis and Applications*, vol. 12, no. 4, pp. 407–413, Dec. 2009.

[15] A. Häfner, A. Uhl, and G. Wimmer, "A novel shape feature descriptor for the classification of polyps in HD colonoscopy," in *Proceedings of the 3rd International MICCAI Workshop on Medical Computer Vision (MICCAI-MCV'13)*, ser. Springer LNCS, vol. 8331, 2013.

[16] M. Häfner, A. Gangl, M. Liedlgruber, A. Uhl, A. Vécsei, and F. Wrba, "Endoscopic image classification using edge-based features," in *Proceedings of the 20th International Conference on Pattern Recognition (ICPR'10)*, 2010, pp. 2724–2727.

[17] B. Everitt, *The Analysis of Contingency Tables*. Chapman and Hall, 1977.