České vysoké učení technické v Praze
Fakulta elektrotechnická

Czech Technical University in Prague
Faculty of Electrical Engineering

Dr. Ing. Jan Kybic

Kritéria podobnosti založená na vzájemné informaci vysoké dimenze pro registraci obrazů

High-dimensional mutual information image similarity criteria for image registration

Summary

We present new image similarity criteria based on high-dimensional mutual information using a nearest neighbor entropy estimation. Mutual information is the image similarity criterion of choice for inter-modality image registration. We propose to use mutual information with higher dimensional features. We show that these features are more powerful than standard scalar mutual information. Earlier attempts to use high dimensional features were limited to very low dimensions or used simplifying normality assumptions. The reason is a poor performance of a standard histogram estimator in higher dimensions. We use the Kozachenko-Leonenko estimator based on nearest neighbor distances, which is usable in higher dimensions.

The main drawback of this approach is that calculating the nearest neighbor (NN) distances is computationally expensive. That is why an essential part of our approach is an approximate nearest neighbor search algorithm tailored for this application. It is based on constructing a k-d tree augmented with tight bounding boxes, best bin first strategy, pruning, and incremental tree updating. Both the NN search algorithm and the estimator can also handle multiple points in the data. An experimental comparison of our nearest neighbor search algorithm with the state of the art ANN library shows that our method is superior for exact search and high number of points as well as for approximate search in small to moderate dimensions or when a fast approximation is needed.

Souhrn

V této práci představíme nová kritéra podobnosti obrazů založená na vzájemné informaci vysoké dimenze, používající odhadu entropie ze vzdáleností nejbližších sousedů v datovém souboru. Vzájemná informace je nejčastěji používané kritérium podobnosti pro registraci obrazů různých modalit. Novou myšlenkou je použití vzájemné informaci vysoké dimenze. Ukážeme, že na ní založená nová kritéria fungují lépe než standardní vzájemná informace používající pouze skalární intenzitu obrazu. Dřívější práce se omezovaly na vzájemnou informaci nízké dimenze nebo používaly zjednodušující předpoklad normálnosti hustoty pravděpodobnosti. Důvodem bylo špatné chování standardního odhadu pomocí histogramu ve vyšších dimenzích. V této práci proto používáme Kozačenko-Leoněnkův odhad založený na vzdálenosti nejbližších sousedů. který se ve vyšších dimenzích chová lépe.

Výpočetně nejnáročnější část postupu je nalezení nejbližších sousedů. Proto je důležitou částí naší metody algoritmus pro přibližné hledání nejbližších sousedů, vyvinutý zvlášť pro tuto aplikaci. Je založený na použití rozšířeného k-d stromu, prohledávací heuristiky nejlepšího uzlu, prořezávání a inkrementálních změn stromu. Jak algoritmus vyhledávání nejbližšího souseda, tak odhad entropie umí pracovat s násobnými body. Experimentální porovnání našeho algoritmu hledání nejbližších sousedů se špičkovou knihovnou ANN ukazuje, že náš algoritmus je v mnoha případech lepší, zejména v případě přesného vyhledávání pro velký počet bodů a pro přibližné a rychlé vyhledávání ve malých a středních dimenzích.

Klíčová slova: registrace obrazů, vzájemná informace, odhad entropie, nejbližší soused

Keywords: image registration, mutual information, entropy estimation, nearest neighbor

Contents

Chapter 1

High dimensional mutual information image similarity criterion

1.1 Image registration

We define image registration as a multidimensional optimization problem: given two images $\mathcal{F}$ and $\mathcal{G}'$ and a family of geometrical transformations parameterized by a finite number of parameters (such as translations, affine transformations, or general non-linear warpings), we search for a geometrical transformation T such that a warped test image $\mathcal{G} = \mathcal{G}' \circ \mathsf{T}$ is as similar as possible to a reference image $\mathcal{F}$ in the sense of maximizing an image similarity criterion J.

1.1.1 Image similarity criteria

One of the simplest criteria is the sum of square differences The SSD criterion is simple, fast, and optimal for additive i.i.d. Gaussian noise corruption. It can be extended to vector pixel values, e.g. for color images. Other alternatives are the sum of absolute differences (SAD) or normalized correlation (NCC) [1]. Normalized correlation allows for a linear dependence between the intensities in both images.

The mutual information image similarity criterion is used when the dependence between the image intensities is unknown, such as in the case of registering images from two different modalities [2, 3, 4, 5]. It can capture very general and not necessarily functional dependences. Mutual information can be defined as follows:

$$J_{\mathsf{MI}}(\mathcal{F}, \mathcal{G}) = I(\mathsf{F}, \mathsf{G}) = H(\mathsf{F}) + H(\mathsf{G}) - H(\mathsf{F}, \mathsf{G}) \tag{1.1}$$

where $H()$ stands for entropy and F and G are random variables corresponding to the two images. In most cases the scalar grayscale pixel intensities $f_i = f(\mathbf{x}_i)$, $g_i = g(\mathbf{x}_i)$ are used as descriptors (features); they are assumed to be independent realizations of the random variables F and G ($f_i \sim \mathsf{F}$, $g_i \sim \mathsf{G}$); the index i denotes a particular pixel. The joint random variable (F, G) corresponds to the concatenation of the pixel descriptors of the two images at the same location, with $(f_i, g_i) \sim (\mathsf{F}, \mathsf{G})$. To estimate the entropies, we shall use the sample sets $\mathbf{F} = \{f_i\}_{i\in\Omega}$, $\mathbf{G} = \{g_i\}_{i\in\Omega}$, $\mathbf{FG} = \{(f_i, g_i)\}_{i\in\Omega}$, where Ω is the set of all pixels.

1.1.2 Higher dimensionality MI criteria

We propose to use more general and more powerful d-dimensional feature vectors instead of using just a simple scalar pixel values ($d = 1$) as described above. Each sample vector $\mathbf{f}_i$ will correspond to one spatial location $\mathbf{x}_i$ of the image $\mathcal{F}$, and similarly for the image $\mathcal{G}$. Let us present two examples of criteria with higher dimensional features:

1) The 3D color component vector (for color images)

$$\mathbf{f}_i^{\mathrm{Co}} = (f^R(\mathbf{x}_i), f^G(\mathbf{x}_i), f^B(\mathbf{x}_i)) \tag{1.2}$$

   leads to a color MI *(CoMI)* criterion $J_{\mathsf{CoMI}} = I\big(\mathsf{F}^{\mathrm{Co}}, \mathsf{G}^{\mathrm{Co}}\big)$. It adapts automatically to any changes of the image colors, e.g. due to different (even radically different) lighting conditions. It is straightforward to extend J_{CoMI} to more dimensions, e.g. for multispectral imaging.

2) The neighborhood criterion $J_{\mathsf{NbMI}} = I\big(\mathsf{F}^{\mathrm{Nb}}, \mathsf{G}^{\mathrm{Nb}}\big)$ forms feature vectors of dimension $d = (2h+1)^2$ from pixel values in the neighborhood of a current location:

$$\mathbf{f}_i^{\mathrm{Nb}} = \big(f(x - \Delta_x, y - \Delta_y)\big)_{|\Delta_x|\le h,\ |\Delta_y|\le h} \tag{1.3}$$

   The variant presented here is for the 2D case and grey-level images but it can be obviously extended to 3D and color images. This criterion learns correspondences between image details such as peaks, ridges and transitions.

Some modest attempts to use high dimensional MI criteria for image registration have appeared in the literature: the second-order MI [6], regional MI [7], or combining intensity and gradient information [8]. However, these criteria either use very small dimensional features ($d = 2$), or their probability distribution is assumed to be Gaussian, which is a gross simplification. The main obstacle is the predominantly used standard histogram-based plug-in estimator [4] .

1.2 Nearest neighbor entropy estimation

Kozachenko and Leonenko (KL) [9, 10, 11] proposed to estimate the entropy $H(\mathsf{F})$ from the pairwise NN distances. Given a set of n samples S of a random variable F in $\mathbb{R}^d$ with a probability density distribution f, then the Shannon entropy of F

$$H(\mathsf{F}) = -\int_{\mathbf{x}\in\mathbb{R}^d} f(\mathbf{x}) \log f(\mathbf{x}) \mathrm{d}\mathbf{x} \tag{1.4}$$

can be estimated as

$$H_{\mathsf{KL}}(\mathbf{F}) = \frac{1}{n}\Big(\sum_{\mathbf{q}\in S} d \log \varrho_{\mathbf{q}}\Big) + \gamma + \log \frac{(n-1)\pi^{d/2}}{\Gamma(1+d/2)} \tag{1.5}$$

for the ℓ_2 NN distance $\varrho_{\mathbf{q}}$. We have also derived a variant for the ℓ_∞ NN distance $\varrho_{\mathbf{q}}$:

$$H_{\mathsf{KL}}(\mathbf{F}) = \frac{1}{n}\Big(\sum_{\mathbf{q}\in S} d \log \varrho_{\mathbf{q}}\Big) + \gamma + \log 2^n(n-1) \tag{1.6}$$

In both cases n is the number of samples $\mathbf{F}$, d is the dimensionality of the space, and $\gamma \approx 0.577$ is the Euler constant. The estimator (1.5) is asymptotically unbiased and consistent .

In our target application of image registration, the image intensity values are often quantized and conflicts (i.e. $\varrho_{\mathbf{q}} = 0$) happen frequently. We replace $d \log \varrho_{\mathbf{q}}$ by a thresholded version $d \log' \varrho_{\mathbf{q}}$ [12]

$$d \log' \varrho_{\mathbf{q}} = \begin{cases} d \log \varrho_{\mathbf{q}} & \text{for } \varrho_{\mathbf{q}} \geq \varepsilon \\ \log\big(\varepsilon^d / \chi_S(\mathbf{q})\big) & \text{for } \varrho_{\mathbf{q}} < \varepsilon \end{cases} \tag{1.7}$$

where the threshold ε corresponds to the measurement accuracy and $\chi_S(\mathbf{q})$ is the multiplicity of the point $\mathbf{q}$. We will call the modified estimator H_{KLD}. A more principled approach is found in [13].

A simple way to avoid the computational complexity of the all NN search for a high number of samples is the batch approach (KLBD): we randomly divide the n samples into groups of m samples and calculate the mean of the entropy estimates for each group. See Chapter 2 for a better but more complex method.

1.3 Experiments

The first experiment (not shown here) measures the bias and variance of the KLD, KLBD and histogram MI estimators for two 1D Gaus-

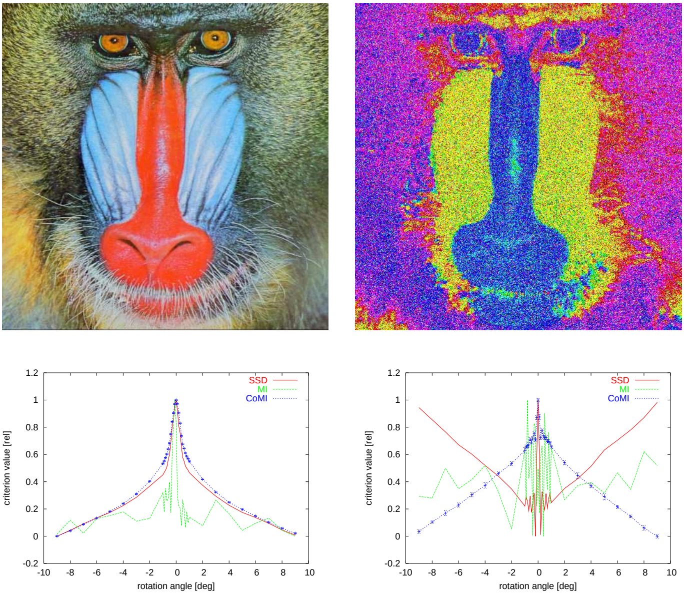

Figure 1.1: Negative *SSD*, *MI* and *CoMI (with* $M = 20$*)* criteria all work well when registering a *Mandrill* color image *(top left)* with itself. We show the criteria (rescaled to $[0, 1]$) as a function of the rotation angle *(top graph)*. One standard deviation is shown for the CoMI criterion, since the estimator is stochastic. When registering the original image with its color-modified and noisy version *(top right)*, the *SSD* and *MI* criteria break down, while the proposed *CoMI* criterion still gives correct results *(bottom graph)*.

sian random variables with varying dependency. We observe that all KL-based estimators have about the same variance, larger than the histogram estimator but still acceptable. The bias of the KLBD estimator decreases with M, for $M \geq 50$ it is already better than that of the histogram and it essentially vanishes for the KLD estimate.

Second, we evaluate the new proposed MI-based image registration criteria *CoMI* (1.2) and *NbMI* (1.3) on real images by a standard rotation experiment [14]: Starting from two perfectly aligned images of size 512^2 pixels, we rotate one of them by $\pm 10°$ (using linear interpolation), crop the images to a fixed size to avoid the influence of the background and evaluate the criterion, showing also the standard deviation for the stochastic KLBD estimator.

First of all we compare the CoMI (1.2), the vector SSD, and stan-

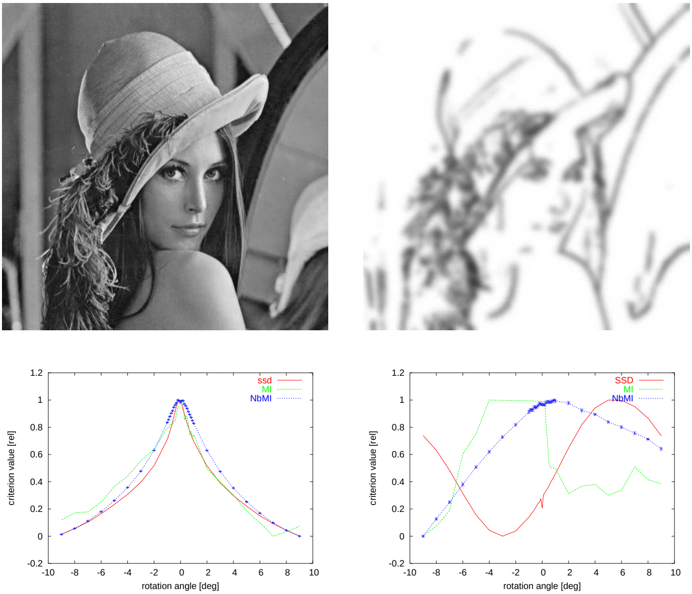

Figure 1.2: The correct rotation angle can be determined from any of the negative *SSD*, *MI* and *NbMI* criteria when registering a B&W *Lena* image *(top left)* with itself. We show the criteria (rescaled to $[0, 1]$) as a function of a rotation angle *(top graph)*. One standard deviation is shown for the NbMI criterion, since the estimator is stochastic. When registering a blurred version of the original image with a blurred version of its edges *(top right)*, the *SSD* and scalar *MI* are clearly inadequate, while the proposed *NbMI* criterion allows for the correct angle to be reliably detected *(bottom)*.

dard scalar gray-scale histogram-based MI criteria when registering a color Mandrill image with itself (Fig. 1.1), to verify that all criteria work well in this simple case. We then modify the colors in one of the images (by increasing the saturation and brightness and rotating the colormap) and add some i.i.d. Gaussian noise to individual color components. This confuses the SSD criterion beyond usability and the standard MI is only slightly better, while the CoMI still provides correct and almost undisturbed results. We also observe that the uncertainty due to the stochastic character of the estimator (one standard deviation shown) is below the level of changes we need to detect for a registration accuracy that can be realistically expected, (i.e. around 1 pixel $\sim 0.2°$).

We perform the same kind of experiment comparing the grayscale

SSD and MI criteria and the NbMI criterion (1.3) with neighborhood size $h = 2$, leading to $d = 25$ dimensional features and estimation of the joint entropy $H(\mathbf{FG})$ in dimension $2d = 50$. First we register a gray-scale Lena image with itself, to find that all criteria work well. Then we register a low-pass version of the image with a smoothed Sobel-detected edges from the same image. The SSD and scalar MI criteria are useless in this case, showing only irrelevant oscillations due to the global orientation of the edges, while the NbMI criterion identifies the correct alignment flawlessly.

Chapter 2

Approximate best bin first k-d tree all nearest neighbor search

2.1 Nearest neighbor search

In this chapter we describe a practical all-NN search algorithm that can be used together with the KL entropy estimator for image similarity evaluation in image registration. This application has a number of specific requirements which we believe cannot be found simultaneously in any existing method — this justifies developing a specialized algorithm:

(i) The datasets are large and the all-NN search is run once for each iteration of the optimizer. Therefore speed is very important. The number of data points n (corresponding to the number of pixels) is typically between 10^5 and 10^7. The typical number of dimensions d is between 4 and 100.

(ii) Due to quantization, the dataset can contain the same point several or many times. This can be for example caused by a homogeneous background in the images. The all-NN search algorithm must handle these cases efficiently and report the data point multiplicities.

(iii) The algorithm can take advantage of the fact that the query set is identical to the set of data points and both are known in advance.

(iv) As exact all-NN search is likely to be too slow for moderate to large n and d, we also need the ability to calculate an approximate solution for a given time budget.

(v) The data points depend continuously on a geometrical transformation being controlled by an optimizer. We can therefore assume that especially at later stages of the optimization, the changes of data point values between iterations will be small. Our all-NN search algorithm can take advantage of this fact.

2.2 Extended k-d tree

We use an extended version of the k-d tree data structure [15, 16, 17]. A k-d tree stores data points in its leaves. Each node represents a hyperinterval (an axis aligned hyperrectagle), which we shall call a *loose bounding box* (LBB) of the node.

As a novelty, we also maintain and store a tight bounding box (TBB) for each node. The TBB of a node Q is the smallest hyper-interval (in the sense of inclusion) containing all points in the subtree of Q.

For the purpose of the dynamic tree update procedure we introduce a parameter $\delta \in \langle 0, 0.5 \rangle$, limiting each child subtree to contain at most $(\frac{1}{2} + \delta)n_p$ points, where n_p is the number of points of its parent subtree.

2.2.1 Building the tree (`BuildTree`)

The tree building algorithm is standard, based on recursive splitting. We start with the entire input multiset S as the root node and the whole space R^D as its loose bounding box LBB. Then, recursively from the root, for each node Q we choose the splitting dimension m as the longest edge of the TBB. The splitting value ξ is the median of $\{x_m; \mathbf{x} \in Q\}$ [18]. In this way, the tree is balanced with respect to the number of points in each subtree at the same level.

2.2.2 Nearest neighbor search (`SearchTree`)

We loop through all leaves and through all points in each leaf. Each point $\mathbf{q}$ acts as a query points and we find its nearest neighbor $\hat{\mathbf{q}}$. The basis for the search is the *BBF (best bin first)* tree traversal [19], with pruning, using a lower bound η_X of the distance from the query $\mathbf{q}$ to yet unexplored points reachable from a node X. Additionally, instead of starting from the global root, the search starts from the leaf Q containing the query point $\mathbf{q}$.

The approximative search is controlled by a parameter V that bounds the number of visited points. If this number is exceeded, the search is stopped and the best result so far is reported.

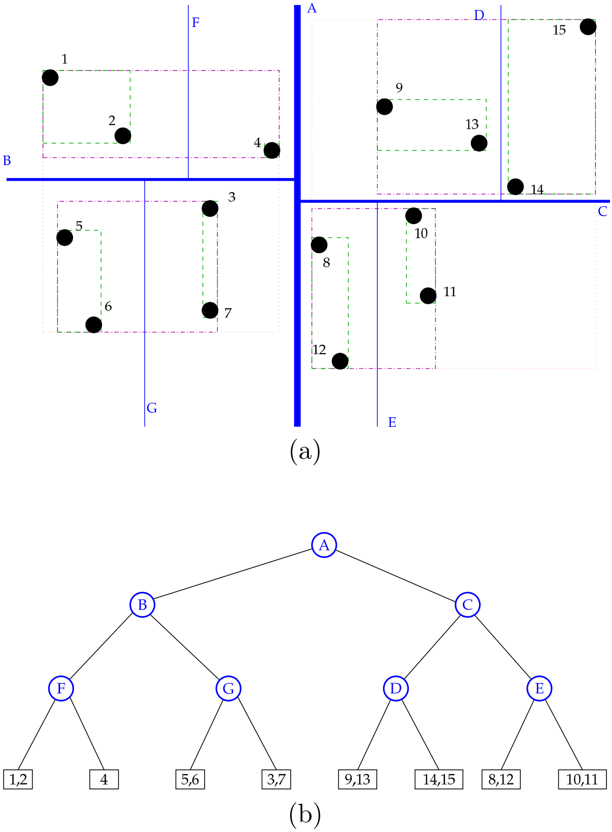

Figure 2.1: A k-d tree example in a 2D space. (a) Solid blue lines marked by uppercase letters represent the splitting hyperplanes hierarchically subdividing the space into loose bounding boxes (LBB) corresponding to tree nodes. Dashed lines show the tight bounding boxes (TBB). (b) The k-d tree itself, with round non-leaf nodes marked by the corresponding dividing hyperplanes and rectangular leaf nodes each containing a set of data points denoted by numbers.

Figure 2.2: Relative mean square error of entropy estimation versus elapsed time for BBF and ANN methods and the KL NN entropy estimator. There were $n = 10^5$ normally distributed points in dimensions $d = 3 \sim 20$. The curves of the same color correspond to the same d and are meant to be compared, solid lines correspond to the BBF method and dashed lines to the ANN method. Each point in the graph is a mean of 100 runs.

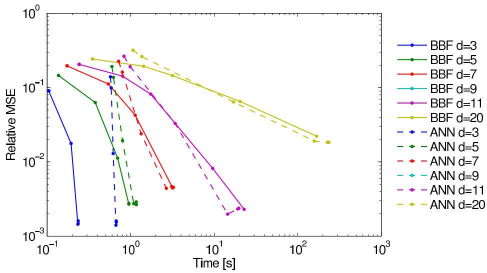

2.2.3 Updating the tree (`UpdateTree`)

The update routine detects points from the tree which have moved out of the LBBs of their original leaves and attributes them to the appropriate new leaves. The TBBs of affected nodes are updated. Finally, parts of the tree violating the balance condition are rebuilt. The method consists of two depth-first recursive traversals of the tree.

2.3 Experiments

We have implemented the algorithm in C++ and performed several experiments on an Intel 1.8GHz PCs with 2GB of RAM running Linux to test its practical properties and its performance against some alternative approaches. We are comparing our results (the all-NN search algorithm, denoted BBF) against the brute force algorithm (refered to as 'brute') with time complexity $O(N^2)$ and a state-of-the-art approximate NN search implementation in the ANN library by Arya et al.[20] (referred to as ANN), which uses a balanced box decomposition (BBD) tree. For comparison, we have also implemented the n times repeated NN search using our BBF approach, denoted BBF NNN.

The first experiment compares the elapsed time for the three subquadratic methods (ANN, BBF NNN, BBF) on the number of points n

Table 2.1: The total time in seconds to find all NNs as a function of the number of points n and dimension d for the three subquadratic methods. The shortest time for each n and d is typeset in bold.

n	d	1	2	3	4	5	10	15	20
	BBF	**0.01**	**0.01**	**0.04**	**0.07**	0.12	0.84	2.52	7.74
10^4	NNN BBF	0.02	0.03	0.06	0.10	0.16	0.97	2.67	8.01
	ANN	0.04	0.04	0.06	0.08	**0.12**	**0.80**	**2.51**	**7.43**
	BBF	**0.07**	**0.13**	**0.50**	**0.91**	**1.58**	**20.32**	**92.46**	**370.54**
10^5	NNN BBF	0.34	0.48	0.91	1.70	2.75	26.32	97.56	372.00
	ANN	0.50	0.60	0.93	1.48	2.26	25.04	107.24	434.44
	BBF	**1.10**	**1.91**	**5.33**	**10.22**	**17.56**	**335.54**	**2562.78**	—
10^6	NNN BBF	6.10	7.77	11.75	22.44	34.81	417.34	2654.09	—
	ANN	8.16	9.05	13.04	20.66	30.64	364.80	2989.96	—

for different dimensions d. We can see that for $n \geq 10^5$, our BBF method outperforms the ANN for all d, the difference is very pronounced in low dimensions (a factor of 7 for $d = 1$) and decreases with increasing d. This seems to be a general pattern: in higher dimensions all methods struggle and their differences get smaller.

The second experiment compares the performance of the ANN and BBF methods for approximative search by examining their time-accuracy trade-off. The accuracy of the NN search is evaluated using the reported approximate NNs to estimate the entropy of a given set of points using the KL NN entropy estimator for $n = 10^5$ normally and isotropically distributed points in dimensions $d = 3 \sim 20$, which allows the entropy to be calculated analytically. We can see (Figure 2.2) that our BBF method outperforms ANN for small d. Starting from about $d = 7$, BBF is better for operating points privileging shorter times and larger errors, while ANN performs better at longer times and smaller errors.

We have also performed tests with additional distributions, finding that for $n = 10^6$ and $d = 2 \sim 20$, BBF in almost all cases (26 out of 28) outperformed ANN.

Finally, we have evaluated the effectiveness of the tree update operation by taking uniformly distributed points from $[-1; 1]^d$ with $d = 5$, perturbing them by additive uniformly distributed noise from $[-\sigma, \sigma]^d$ and compared the speed of finding all exact NNs by updating a tree built for the original dataset and by building the tree from scratch. We have found that it is indeed always advantageous to update the tree instead of rebuilding it anew, as long as we allow some unbalance δ. The time savings in terms of the total time (tree building/update and search) are relatively modest but the savings in terms of the update versus build times are already more important, around 50%. This is relevant as in practice mostly approximative search will be used which reduces the search times by one or several orders of magnitude, making the build or update times dominate.

2.4 Conclusions

In the first chapter we address an import part of image registration algorithms, the similarity criteria used to evaluate the quality of the alignment. We propose to extend standard scalar mutual information similarity criteria to use higher dimensional features. In particular, we propose a color based criterion and a criterion taking into account a small neighborhood around each pixel. In contrast to classical methods, the proposed criteria can adapt to and quantitatively evaluate a much wider range of possible dependencies between the images being registered, resulting in a more robust registration. The key factor for successful evaluation of high-dimensional mutual information is a Kozachenko-Leonenko entropy estimator, based on evaluating nearest neighbor distances.

The second chapter attempts to solve the problem of efficiently finding a nearest neighbor for each point from a given set, as this is the computational core of the Kozachenko-Leonenko estimator employed in Chapter 1. We have developed an approximate all nearest neighbor search algorithm, based on an extended k-d tree and best bin first search. It is a general algorithm, capable of dealing with a large number of points in high dimensions. It takes advantage of the fact that all query points are known in advance and that the set of points often changes only very slowly between subsequent evaluations of the criterion. It can also successfully deal with multiple points — no other algorithms combining these features is known to us. We have compared our implementation of the algorithm with a state-of-the-art library ANN for approximate nearest neighbor search with favorable results.

For more details, please see my habilitation thesis or my publications.

References

[1] Barbara Zitová and Jan Flusser, "Image registration methods: a survey," *Image and Vision Computing*, , no. 21, pp. 977–1000, 2003.

[2] Paul Viola and William M. III Wells, "Alignment by maximization of mutual information," *International Journal of Computer Vision*, , no. 2, pp. 137–154, 1997.

[3] J.P.W. Pluim, J. B. A. Maintz, and M. A. Viergever, "Mutual-information-based registration of medical images: A survey," *IEEE Transactions on Medical Imaging*, vol. 22, no. 8, pp. 986–1004, Aug. 2003.

[4] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens, "Multimodality image registration by maximization of mutual information," *IEEE Transactions on Medical Imaging*, vol. 16, no. 2, pp. 187–198, Apr. 1997.

[5] Philippe Thévenaz and Michael Unser, "Optimization of mutual information for multiresolution image registration," *IEEE Transactions on Image Processing*, vol. 9, no. 12, pp. 2083–2099, Dec. 2000.

[6] D. Rueckert, M. J. Clarkson, D. L. G. Hill, and D. J. Hawkes, "Non-rigid registration using higher-order mutual information," in *Proceedings of SPIE Medical Imaging 2000: Image Processing*, 2000, pp. 438–447.

[7] D. B. Russakoff, C. Tomasi, T. Rohlfing, and C. R. Jr. Maurer, "Image similarity using mutual information of regions," in *Proceedings of the 8th European Conference on Computer Vision (ECCV)*, T. Pajdla and J. Matas, Eds. May 2004, number 3023 in LNCS, pp. 596–607, Springer.

[8] Josien P. W. Pluim, J. B. Antoine Maintz, and Max A. Viergever, "Image registration by maximization of combined mutual information and gradient information," *IEEE Transactions Med. Imag.*, vol. 19, no. 8, Aug. 2000.

[9] L. F. Kozachenko and N. N. Leonenko, "On statistical estimation of entropy of random vector," *Probl. Inf. Trans.*, vol. 23, no. 9, 1987, (in Russian).

[10] Jonathan D. Victor, "Binless strategies for estimation of information from neural data," *Physical Review E*, vol. 66, no. 5, pp. 051903(15), Nov. 2002.

[11] J. Beirlant, E. J. Dudewicz, Györfi L., and E. C. van der Meulen, "Nonparametric entropy estimation: an overview," *International J. Math. Stat. Sci.*, , no. 6, pp. 17–39, 1997.

[12] Jan Kybic, "Incremental updating of nearest neighbor-based high-dimensional entropy estimation," in *ICASSP2006*, Pierre Duhamel and Luc Vandendorpe, Eds., Toulouse, France, May 2006, pp. III–804, IEEE, DVD proceedings.

[13] Jan Kybic, "High-dimensional entropy estimation for finite accuracy data: R-NN entropy estimator," in *IPMI2007: Information Processing in Medical Imaging, 20th International Conference*, Nico Karssemeijer and Boudewijn Lelieveldt, Eds., Berlin, Heidelberg, Germany, July 2007, pp. 569–580, Springer.

[14] M. Unser, P. Thévenaz, and L.P. Yaroslavsky, "Convolution-based interpolation for fast, high-quality rotation of images," *IEEE Transactions on Image Processing*, vol. 4, no. 10, pp. 1371–1381, October 1995.

[15] Jon Louis Bentley, "Multidimensional binary search trees used for associative searching," *Commun. ACM*, vol. 18, no. 9, pp. 509–517, 1975.

[16] F. P. Preparata and M. I. Shamos, *Computational geometry: An introduction*, Texts and Monographs in Computer Science. Springer-Verlag, 1985.

[17] R. Sedgewick, *Algorithms*, Addison-Wesley, Reading, MA, 1988.

[18] Jerome H. Freidman, Jon Louis Bentley, and Raphael Ari Finkel, "An algorithm for finding best matches in logarithmic expected time," *ACM Trans. Math. Softw.*, vol. 3, no. 3, pp. 209–226, 1977.

[19] Jeffrey S. Beis and David G. Lowe, "Shape indexing using approximate nearest-neighbour search in high-dimensional spaces," in *Proceedings of Conference on Computer Vision and Pattern Recognition*, June 1997, pp. 1000–1006.

[20] Sunil Arya, David M. Mount, Nathan S. Netanyahu, Ruth Silverman, and Angela Y. Wu, "An optimal algorithm for approximate nearest neighbor searching in fixed dimensions," *Journal of the ACM*, vol. 45, no. 6, pp. 891–923, 1998.

Dr. Ing. Jan Kybic

http://cmp.felk.cvut.cz/~kybic
Born on 18^{th} January 1974 in Prague, Czech Republic.

2010–2011	Visiting researcher at CVLAB, EPFL, Switzerland
since 2003	Senior researcher at CMP, FEE, CTU Prague
2001–2003	Post-doc position at INRIA, France.
1998–2001	PhD in biomedical image processing at EPFL, Switzerland
1994–1998	Master (Engineer) in control engineering from CTU Prague, with honors

Associate Editor for *IEEE Transactions on Medical Imaging*, senior member of IEEE, reviewer for many medical imaging and image processing journals and conferences.

Teaching Digital Image Processing and Medical Imaging courses. Supervising three PhD students, two PhD students and 10 Master students have already successfully graduated. Collaborated on six projects funded by Czech grant agencies and participates in one European training network.

I have (co)authored 15 articles in peer-reviewed journals, one book, one book chapter, and about 60 conference publications. I have about 220 citations (excluding autocitations).

My hobbies are music (piano), photography, travelling, literature, cinema, hiking, volleyball, skiing, sailing. Besides Czech, I speak English, French and Spanish fluently, and I also know German, Italian and Russian.