

Rigid Registration of Echoplanar and Conventional Magnetic Resonance Images by Minimizing the Kullback-Leibler Distance

Salil Soman^{1,2,3}, Albert C.S. Chung⁴, W. Eric L. Grimson¹,
and William M. Wells III^{1,2,5}

¹ MIT Artificial Intelligence Laboratory, Cambridge, MA USA

² Surgical Planning Laboratory, Brigham & Women's Hospital, Boston, MA USA

³ UMDNJ – Robert Wood Johnson Medical School, NJ USA

⁴ Dept. of Computer Science, Hong Kong University of Science & Technology, HK

⁵ Harvard Medical School, Boston, MA USA

sal@ai.mit.edu

Abstract. Functional Magnetic Resonance Imaging (fMRI) studies are derived from a time series of Echo-Planar images (EPIs). Compared to conventional Magnetic Resonance Images (MRIs), EPIs are of relatively poor quality for discerning anatomic features and are often registered with corresponding MRIs to map brain activity to neuroanatomy. In this paper we demonstrate the utility of a technique to register an EPI-MRI pair by minimizing the discrepancy between its joint intensity probability mass function (PMF) and a previously learned one for a properly registered EPI-MRI pair, using the Kullback-Leibler Distance (KLD). In probing experiments Joint Entropy (JE) and Mutual Information showed significant bias relative to KLD along the axial direction and JE along a rotation axis. A comparison of searches using random starting poses showed KLD to have lower final pose errors than JE. Results of variation on parameters of the KLD based EPI-MRI registration technique are also presented.

1 Introduction

Researchers often use functional Magnetic Resonance Imaging (fMRI) to study brain activity. These studies produce spatial activation maps through analysis of a time series of Echo-Planar Images (EPIs) taken of the patient's head as the subject responds to specific tasks or stimuli. While EPIs provide brain activity information on a time scale adequate for making inferences about regional brain function, they are of relatively poor quality for discerning anatomical features. Because there are a variety of tasks that are facilitated by the ability to correlate features like motor function to specific neuroanatomy (e.g. planning neurosurgical tasks or performing neuroscience experiments), researchers often desire to register EPIs to their corresponding Magnetic Resonance Images (MRI). This registration allows brain activity noted in regions of EPIs over successive volumes of the time series to be attributed to anatomic regions of the brain. Registration techniques which employ Mutual Information (MI) provide a partial solution to this problem, but they show limitations in capture range for rigid transformations. In this paper we demonstrate the utility of a method that leverages *a priori* domain knowledge to register an EPI –MRI pair by utilizing infor-

mation about a properly registered EPI-MRI pair. The technique searches the space of transformations of the EPI of a test image pair for a pose that results in a joint intensity probability mass function (PMF) that most resembles a learned one for a properly registered training pair. The difference between the PMFs is determined using the Kullback-Leibler Distance (KLD), an information theoretic similarity measure commonly used in machine learning and information theory [1]. The KLD registration approach has previously been used to solve challenging angiographic registration problems [2, 3].

In fMRI studies a time series of EPIs of a patient's head are taken and regions of the EPIs are analyzed for changes over time [4]. Regions of successive EPIs where intensity changes correlate with the experimental protocol are indicative of brain activity. Brain activation indirectly causes intensity variations in the EPI images by way of the Blood Oxygenation Level Dependent (BOLD) effect [5]. Time series of EPIs must capture successive images at intervals of approximately once every two seconds to be able to reflect these rapid occurring changes in blood flow. Conventional MRIs are acquired over a period of some minutes per image capture, and provide more detailed images of brain anatomy [6]. Registering a set of EPIs to a set of corresponding MR images allows researchers to correlate areas of apparent activity in the EPI images with specific brain anatomy.

It is widely recognized that EPI images can contain significant distortions in addition to the intensity voids caused by magnetic susceptibility effects [6]. A definitive solution to the EPI/MRI registration problem will likely address this issue directly, by simulating the effect, or by field mapping at the time of acquisition. Nevertheless, we feel that in the interim, a robust method of rigid registration would be well received by the research community that uses fMRI.

One pragmatic approach to the MRI / EPI registration problem is to acquire a conventional MRI scan just before the EPI images for use as a registration reference (RR) that may be registered to other conventional MRI using, for example, MI. This method depends on the EPI and RR being in correspondence by the design of the scanning protocols. While the method is viable, it does require an additional scan, and it is not usable retrospectively for data that was acquired without the RR. In addition, there is frequently residual misregistration between the EPIs and the RR that needs to be corrected manually.

Many medical image registration problems have been solved using the Mutual Information approach [7, 8, 9]. The MI approach to registration seeks a transformation which maximizes the statistical dependence among the two images, without regard to particulars of the relationship. There are strong similarities between the MI approach and that of minimizing the joint entropy (JE). The objective functions share a joint entropy term, which is responsible for a pronounced extremum at the correct pose. The MI approach uses additional individual entropy terms that can enhance long range capture in some applications. In the experience of the authors, and as we will demonstrate below, the widely-used MI registration approach can, however, perform poorly on the EPI/MRI registration.

2 Methods

In this section we describe the objective function search and probing experiments reported in the results section.

2.1 Overview of Registration Experiments

In a training operation an EPI and its corresponding MRI are registered by an expert, blurred to 4 levels, and joint intensity PMFs are computed for all 5 of the EPI-MRI pairs. These PMFs will be referred to as Aligned Joint Intensity PMFs (AJPs).

In the registration phase, a test EPI-MRI pair is blurred to the same four levels that were used on the training pair. Then, starting with the most blurred image (blur level 4 in the figures), the algorithm searches for a rigid transformation of the test pair's EPI that results in a test pair joint intensity PMF that has a minimal KLD from the AJP for that blur level. Once this pose is found, the algorithm repeats this search on the test pair's EPI blurred to the highest blur level that has not yet been searched, using the final pose from the previous search as the start pose. The search continues until a pose is found to minimize the KLD between the test pair's PMF and the AJP of the non-blurred training pair. This final pose is returned as result.

2.1.1 Image Preparation and Selection

All image data in our experiments were acquired on a GE 1.5 Tesla Signa system equipped with the HORIZON hardware/software package. The MRI data in our experiments is a standard gradient echo (SPGR) scan that consists of 124 1.5 mm thick slices, FOV 24cm, image resolution 256 X 256 pixels. The EPI data sets contain 21 contiguous 7mm-slice images. The functional images are acquired in an auditory experiment using the EPIBOLD pulse sequence with the following set of parameters: TE=50msec, TR=3sec, FOV 24cm, image resolution=64 x 64 pixels.

EPI and MRI volumes were rigidly aligned for good visual agreement throughout, the EPI volume was re-sampled, using tri-linear interpolation, into the lattice of the MRI data, and corresponding 256 X 256 slice pairs were generated to be used as 2D test and training pairs.




Fig. 1. MR and Echo-planar Image Pairs from A Single Acquisition Session

2.1.2 Blurring the Images

Four rotationally symmetric Gaussian low-pass filters of the following sizes and standard deviations (σ) were created to blur the original images: **Level 1**- 20x20 pixels, σ : 3, **Level 2** – 40x40 pixels, σ : 5, **Level 3** – 80x80 pixels, σ : 8, **Level 4** – 100x100 pixels, σ : 10.




Fig. 2. (Left to Right) – Increasing strength blur applied to an Echo-planar image

2.1.3 Image Transformations

All transformations of the EPI images were performed using a rigid transformation model, with rotations occurring around the center of the image [10].

2.1.4 Histogramming to Estimate Joint Intensities

For the KLD objective function, a histogram was computed on an image pair by rastering over the EPI and MRI simultaneously and counting the co-occurrence of intensities in regions where the images overlapped. The co-occurring pixel intensity i from the EPI and j from the MRI led to incrementing bucket (i,j) in a 256×256 matrix.

The matrix was then normalized to convert the histogram into a PMF. Because it was possible for some buckets to have counts of zero after normalization, an ξ value was added to all buckets in the matrix before normalization. Varying ξ between 1×10^{-6} and 1×10^{-23} had no significant difference in the results of X probing experiments (see Objective Function Probing Experiments section). The experiments presented in this paper used an ξ of 1×10^{-23} .

The histogramming techniques use for KLD and the JE or MI probes differed slightly. The JE and MI objective functions were computed in the space of the fixed image, with the moving image zero padded as necessary. Histograms for the MI and JE method were calculated with and without interpolation, respectively.

2.1.5 Kullback-Leibler Distance (KLD)

Given the learned \hat{P} and the pose dependent observed P_o^T joint intensity PMFs, the Kullback-Leibler distance between the two PMFs is given by [11, 12]

$$D(P_o^T \parallel \hat{P}) = \sum_{i_1, i_2} P_o^T(i_1, i_2) \log \frac{P_o^T(i_1, i_2)}{\hat{P}(i_1, i_2)} \quad (1)$$

2.1.6 Interpolation vs. Rounding

Histograms were generated by applying a standard rounding function to image intensities. For comparison, we also implemented an interpolation function that spread counts of intensity co-occurrences over the buckets corresponding to the combinations of floors and ceilings of each intensity value.

2.1.7 Optimization of the Transformation T - Downhill Simplex Searches

The goal of the KLD based registration is to find the optimal transformation \hat{T} by minimizing $\hat{T} = \arg \min_T D(P_o^T \parallel \hat{P})$ - the difference between the observed P_o and the expected \hat{P} joint intensity PMFs.

For the registration experiments presented, the value of KLD or JE was minimized using the Nelder-Mead Downhill Simplex (direct search) method with a multi-resolution strategy [13, 14]. This approach does not use derivatives.

One hundred start poses were randomly chosen, where the X, Y, and θ parameters were varied within boundaries implied by Figure 3 to cause the majority, but not all downhill searches using KLD and JE to succeed. The parameters ranged from: -45 to 55 for X, -25 to 35 for Y, and -50° to 130° for θ . The KLD searches were then carried out as described in the overview part of the Methods section. The JE searches were performed in the same manner, except that the training phase was not carried out, and for each evaluation of the objective function, only the training image pair (transformed and / or blurred appropriately) was evaluated.

2.2 Objective Function Probing

Two EPI-MRI slice pairs were selected from a single acquisition set to be used as training and test pairs to perform rigid transformation probes (see Figure 1). Both image pairs were blurred to the four levels described earlier, and the AJPs for the blurred and non-blurred training pairs were computed. The probes then consisted of performing the following for the image pairs at all levels of blurring: transforming the test pair's EPI in fixed steps along a single axis (X, Y, or θ), computing the joint intensity PMF for the transformed EPI and its corresponding MRI blurred to the same level, and then computing the KLD between this PMF and the AJP for that blur level.

The X probe shifted the test pair EPI along the X axis from -65 to 65 pixels in 1 pixel increments, the Y probe performed the same shift along the Y axis, and the θ probe rotated the test EPI from -180 to 180 degrees in two degree increments.

2.2.1 Intensity Scaling Image Pairs from Different Acquisitions

To account for inter-acquisition intensity variations for training and test pairs, experiments were performed in which an image dependent scaling factor was applied to all intensities during histogramming. The factor consisted of the number of buckets used divided by the intensity of the brightest pixel occurring in that image.

2.2.2 Pseudo 3D Experiment – Histogramming over Multiple Slices

For these experiments, X probes were performed as previously described, with the exception that every step requiring a single image pair used a set of images. Registered EPI-MRI pairs from a single time point were divided into two consecutive sets – anterior and posterior halves. Two probes were then performed, each in which one set was used for training and the other as a test pair. Generating the training AJPs entailed blurring all of the images in the training set, and then histogramming over all intensity pairs occurring in areas of image pair overlap. Computing the test pair PMF required transforming all EPIs of the test set and then histogramming cumulatively as described for generating the AJPs (See Figure 6).

3 Results

3.1 X Probe Comparisons: KLD, JE, MI

Joint Entropy [9], Mutual Information [15, 8] and KLD probes of X, Y or θ axes were performed on a test pair of images (Pair 12). MI & JE's definitions are given by

$$MI = \sum_{i_1, i_2} P_o^T(i_1, i_2) \log \frac{P_o^T(i_1, i_2)}{P_o^T(i_1)P_o^T(i_2)}, \text{ and } JE = -\sum_{i_1, i_2} P_o^T(i_1, i_2) \log P_o^T(i_1, i_2), \quad (2)$$

where $P_o^T(i_1)$ and $P_o^T(i_2)$ are the marginal distributions.

In Figure 3 the plots show the variations in the objective functions for each transformation of the EPI of the test pair for each of the blur levels. Negated MI is presented to facilitate comparison with JE and KLD, which both show minima near proper registration. The transformation resulting in the pose chosen by the expert is indicated by the vertical bar plus symbols. The square on each line indicates the pose resulting in the lowest value for the objective function. Joint Entropy showed significant bias on the Y and θ probes, while MI showed significant bias on the Y probe.

3.2 Objective Function Downhill Simplex Searches

Figure 4 summarizes the results of 100 Downhill simplex searches using KLD and JE, with the final pose error shown as a function of the start pose error. The start and final pose errors were defined as $\sqrt{Error_x^2 + Error_y^2 + Error_\theta^2}$. These searches began with the test image pair blurred at the highest blur level and the EPI of the test pair transformed to the start pose. For a majority of the starting points, the KLD objective function returned a pose with a lower Final Pose Error than the JE based method. A final pose error threshold was selected to classify searches into successful and unsuccessful. This threshold selected the best-scoring significant groups of results. The average final pose for successful searches was then calculated. On average, successful KLD trials converged to pose parameters closer to those chosen by the expert ($x=0, y=0, \theta=0$) than the JE trials, and showed smaller variation. JE was significantly biased in the Y axis, and showed its largest variation in the Y axis pose parameter. Figure 5 shows the overlay of the test EPI transformed to the pose chosen by the expert with its corresponding MRI (left image), the overlay for the test EPI transformed to the average pose parameters suggested by KLD (middle image) and by JE (right image). Visibly, the JE image is most discrepant from the expert's pose, with its registration aligning some brain activation to the scalp. The KLD based method showed a pose similar to the one chosen by the expert.

3.3 Effect of Different Training Pairs on KLD

We examined the effect of choice of training pair on KLD X probes. One EPI-MRI pair was designated as a test pair, and then two KLD X probes were performed, each

using a different EPI-MRI training pair (sampled from the same EPI-MRI volume pair). Both probes showed a minimum near the pose selected by the expert.

3.4 KLD Using Training and Test Pairs from Different Acquisitions

KLD X probes using a test pair from a similar but different EPI-MRI volume pair than the training pair, were performed directly and with the use of an image dependent scaling factor. Both probes yielded minima near the expert’s pose for the non-blurred images, and showed minimal variations of minima for KLD at the four blur levels. The non-scaled probe showed more minima at the expert’s pose, but the probe that scaled the image intensities showed less variation in the position of its minima. The final pose of either method would likely agree with the expert’s.

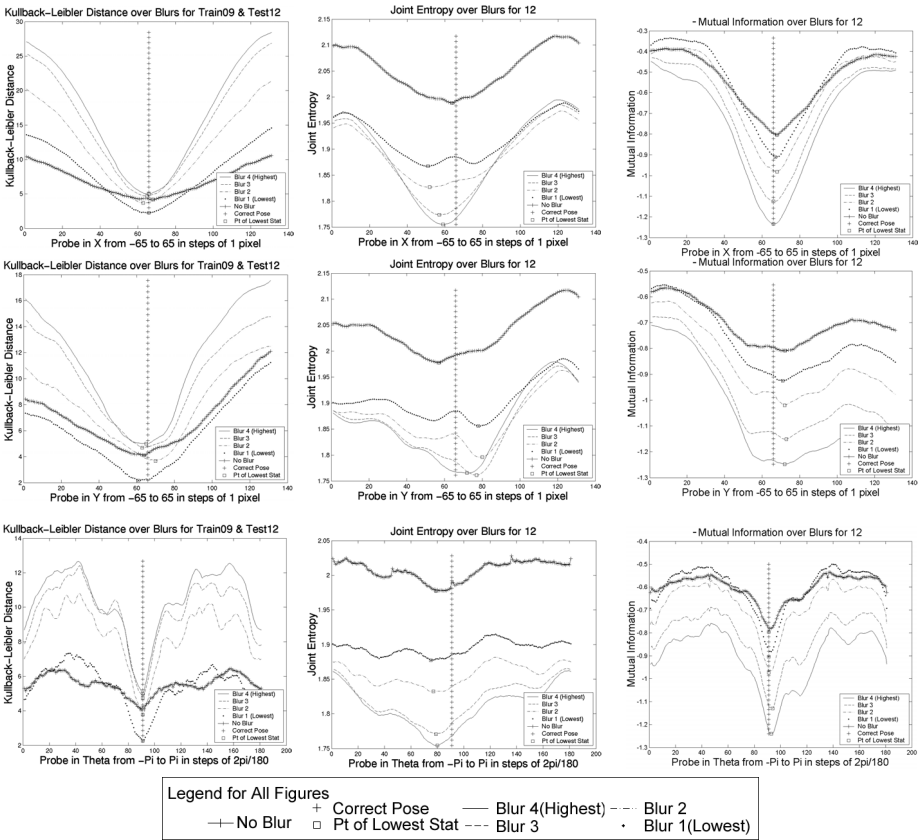


Fig. 3. Comparisons of X, Y, and θ Probes for registering an image pair using (left column) the Kullback-Leibler Distance (KLD), (middle column) Joint Entropy (JE), and (right column) Negated Mutual Information (MI). JE displayed significant bias in Y and θ , while MI shows significant bias along the Y axis

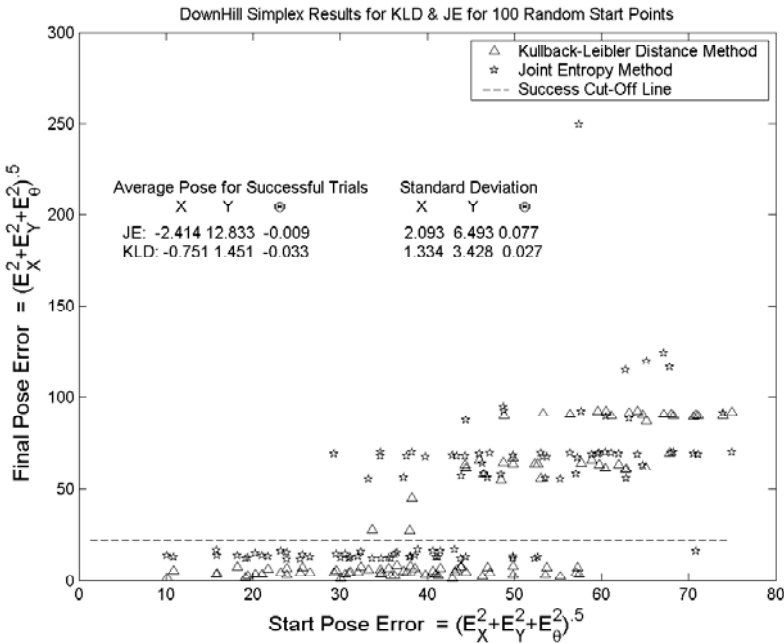


Fig. 4. Results of KLD and Joint Entropy Driven Downhill Simplex Searches for a test image given 100 random start poses for the test EPI. Trials were classified as being successful if the final pose error fell under a threshold. The KLD based successful searches showed average pose parameters closer to the expert’s chosen pose than the Joint Entropy based method




Fig. 5. Significant EPI intensities overlaid on corresponding MRI. (Left) Hand registration of the EPI-MRI pair by an expert, (middle) EPI transformed to the average final pose parameters of successful KLD method trials for downhill simplex, and (right) EPI transformed according to the average final pose parameters of successful Joint Entropy based method

3.5 Effects of Bucket Size and Interpolation on KLD

Non-integer intensity values were encountered during histogramming, and were rounded to allow bucketing. To see if rounding was leading to loss of important information, KLD X probes using a bi-linear interpolation of the intensity value were performed. The interpolation method did not show significant improvement of the

objective function. Similarly, increases in bucket size from 32 to 256 showed relatively small improvements in the smoothness or minima of the KLD X probes.

3.6 Combining Histograms across Multiple Slices – Pseudo 3D X Probes

Figure 6 shows the two KLD X pseudo 3D- probes performed. The left plot used the posterior half of the image set for training, and the anterior as the test set. The right plot used the opposite configuration. In general, both plots show much less variation in minima at all blur levels than the X probes performed using single training and test pairs. The right plot showed more uniform minima at the correct pose than the probe depicted on the left.

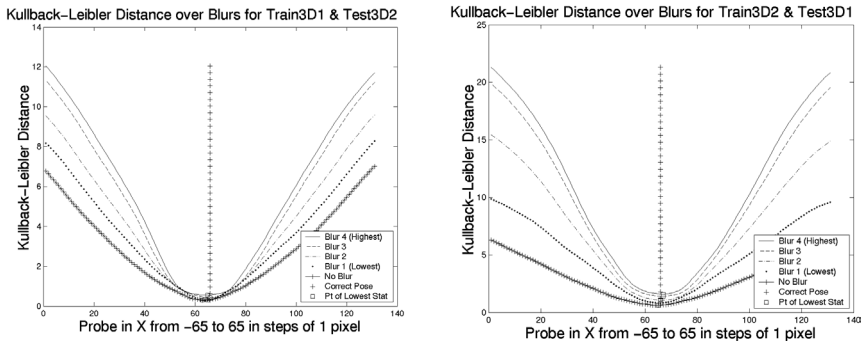


Fig. 6. Pseudo 3D Experiment Results – modified KLD X probes in which all the EPI-MRI pairs from one time point were divided into two sets (anterior and posterior halves), and information from one set was used for training, while the other was used as the test set. Intensity pair information from all the image pairs in a set was combined to make the histograms for these probes. Probes where the anterior half (right) and the posterior half (left) of the image set was used

4 Discussion

While MRI/EPI registration is a 3D-3D problem, in the present work, informative results have been obtained from a series of 2D experiments on representative slices. Clear shortcomings of the MI and JE objective functions were shown, especially in the axial direction, which are unlikely to be ameliorated by 3D registration. Variations on bucket size and the use of rounding or interpolation did not appear to significantly degrade performance of the KLD objective function. Similarly good results were achieved when using images from two different acquisition sets as training and test pairs to perform registration. The pseudo-3D experiments provide good indication that the KL approach will be useful for the full 3D-3D problem.

5 Future Work

This technique needs to be assessed within a three dimensional framework, where probing experiments and searches are performed over all six degrees of freedom.

Acknowledgements

Salil Soman is supported in part by a Medical Informatics training grant from the National Library of Medicine, and a grant from the Pfizer Corporation. We thank Cindy Wible for providing the image data. The following grants supported the work reported in this document - NLM, P41RR13218, P01CA67165, R21CA98449, and ERC9731748. AC would like to acknowledge support from RGC HKUST6209102E and SSRI01/02.EG22 grants.

References

1. Bishop, C. M. (1995). *Neural networks for pattern recognition*. Oxford, New York, Clarendon Press; Oxford University Press.
2. Albert C. S. Chung, William M. Wells III, Alexander Norbash, W. Eric L. Grimson: Multi-modal Image Registration by Minimizing Kullback-Leibler Distance. *MICCAI* (2) 2002: 525-532
3. Ho-Ming Chan, Albert C.S. Chung, Simon C.H. Yu, Alexander Norbash and William M. Wells III: Multi-modal image registration by Minimizing Kullback-Leibler distance between expected and observed joint class histograms. To appear in *CVPR 2003*
4. KK Kwong et al. Dynamic Magnetic Resonance Imaging of Human Brain Activity During Primary Sensory Stimulation. *Proc. Natl. Acad. Sci USA*, June 15, 1992.
5. S. Ogawa and T. M. Lee, "Magnetic resonance imaging of blood vessels at high fields: In vivo and in vitro measurements and image simulation," *Magn. Reson. Med.*, vol. 16, pp. 9-18, 1990.
6. P. Jezzard and S. Clare. Sources of distortions in functional MRI data. *Human Brain Mapping*, 8:80-85, 1999.
7. Paul Viola , William M. Wells, III, Alignment by Maximization of Mutual Information, *International Journal of Computer Vision*, v.24 n.2, p.137-154, Sept. 1997
8. W.M. Wells, P. Viola, Viola P, Atsumi H, Nakajima S, Kikinis R. Multi-Modal Volume Registration by Maximization of Mutual Information. *Medical Image Analysis*, 1(1):35-51,1996.
9. F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, P. Suetens, Multimodality image registration by maximization of mutual information , *IEEE transactions on Medical Imaging*, vol. 16, no. 2, pp. 187-198, April 1997
10. J. Maintz, and M. Viergever. A survey of medical image registration. *Medical Image Analysis*, 2(1):1-36, 1998.
11. T.M. Cover and J.A. Thomas. *Elements of Information Theory*. John Wiley & Sons, Inc., 1991.
12. S. Kullback. *Information Theory and Statistics*. Dover Publications, Inc., 1968
13. Nelder, J. A., and Mead, R., A Simplex Method for Function Minimization, *Computer Journal*, Vol. 7, pp. 308-313, 1965.
14. W.H. Press, S.A. Teukolsky, et al. *Numerical Recipes in C, 2nd Edition*. Cambridge University Press, 1992.
15. Umaki, S. http://www.engineering.uiowa.edu/~aip/248_s02_solutions/soumik_ukil/hw3, May 2003