Towards automatic segmentation of MS lesions in PD/T2 MR images

M. Stella Atkins, Mark S. Drew, and Zinovi Tauber
School of Computing Science, Simon Fraser University,
Vancouver, B.C. Canada V5A 1S6 *

ABSTRACT

Recognizing that conspicuous multiple sclerosis (MS) lesions have high intensities in both dual-echo T2 and PD-weighted MR brain images, we show that it is possible to automatically determine a thresholding mechanism to locate conspicuous lesion pixels and also to identify pixels that suffer from reduced intensity due to partial volume effects. To do so, we first transform a T2-PD feature space via a log(T2)-log(T2+PD) remapping. In the feature space, we note that each MR slice, and in fact the whole brain, is approximately transformed into a line structure. Pixels high in both T2 and PD, corresponding to candidate conspicuous lesion pixels, also fall near this line. Therefore we first preprocess images to achieve RF-correction, isolation of the brain, and rescaling of image pixels into the range 0-255. Then, following remapping to log space, we find the main linear structure in feature space using a robust estimator that discounts outliers. We first extract the larger conspicuous lesions which do not show partial volume effects by performing a second robust regression for 1D distances *along* the line. The robust estimator concomitantly produces a threshold for outliers, which we identify with conspicuous lesion pixels in the high region. Finally, we perform a third regression on the conspicuous lesion pixels alone, producing a 2D conspicuous lesion line and confidence interval band. This band can be projected back into the adjacent, non-conspicuous, region to identify tissue pixels which have been subjected to the partial volume effect.

Keywords: MR tissue analysis, Partial volume effects, Multiple Sclerosis, PD/T2, Robust statistics.

1. INTRODUCTION

Conspicuous multiple sclerosis (MS) lesions have high intensities in both dual-echo PD and T2-weighted MR brain images (see, e.g., Ref.[1]). Simple intensity histogram analyses of brain tissue can be used to identify the so-called typical "conspicuous" lesions, which are large enough not to suffer from reduced intensity due to partial volume effects. In this paper we set out a method to automatically locate the threshold in a combined and re-mapped PD-T2 feature space above which the pixels are likely to identify multiple sclerosis (MS) lesions, and to automatically identify pixels with likely partial volume effects. The partial volume effect at lesion borders means that they are darker in both spectra, and may be confused with other tissues.[2] Our method assumes that the tissue of interest has a low variability in PD and T2 signals, as is the case for conspicuous MS lesions.

The method consists of two stages: preprocessing and feature extraction. Preprocessing consists of RF-correction using Sled's method,[3] followed by isolation of the brain using Atkins' method[4] and scaling the brain pixels about their mean, to the range 0-255. Feature extraction is achieved by extracting first the larger conspicuous lesions which do not show partial volume effects, then using this to extract the darker subtle lesions and the conspicuous lesion borders. This method is an extension of our earlier work,[1] where we observe that lesion intensities change consistently in both PD and T2 images, although possibly at different rates.

Two-dimensional scatterplots of PD versus T2 may be re-mapped such that the scatterplot approximately forms a line. We have found that the regression lines determined in feature space, using a robust regression, are very similar for whole-slice pixels and for pixels corresponding to conspicuous lesions only. Therefore we regress on all pixels available by accumulating pixels from all slices, and linearize the feature space scatterplot by plotting log(T2) versus log(PD+T2). Using a robust Least Median of Squares (LMS) regression, we find a regression line that discounts outliers. We then perform a second LMS regression for 1D distances *along* the line. The LMS estimator concomitantly produces a threshold for outliers, which we identify with conspicuous lesion pixels in the high region.

For conspicuous lesion pixels only, we again determine a robust regression line, using a third regression. The 2D version of LMS automatically identifies an inlier region with threshold interval given in terms of Euclidean distance

* Email: {stella,mark,zinovi}@cs.sfu.ca Web: http://www.cs.sfu.ca

In *Medical Imaging 2000: Image Processing*, Kenneth M. Hanson, Editor,
Proceedings of SPIE Vol. 3979 (2000) • 1605-7422/00/$15.00

from the line; we call this interval a "beam". Partial volume effects can produce substantial errors, especially for larger slice thicknesses.[5] We use the beam to identify tissue pixels which have been subjected to the partial volume effect by projecting the beam back into the non-conspicuous lesion area of the scatterplot in feature space, as we introduced in Ref.[6].

This partial volume effect gradually increases with distance from the lesion centers, so a probability can be assigned to the pixels in the extended beam being identified as tissue-of-interest. Spatial information can be used to refine the tissue segmentation further, as in Ref.[1]. The method was tested on different sets of data with MS lesions which had been manually outlined. With this method we have shown a lesion volume correlation with manually outlined lesions of 0.97 on 8 data sets.

2. AUTOMATIC THRESHOLD DETECTION

2.1. Feature Space Characteristics

In Ref.[1], candidate pixels in PD and T2 images were identified in terms of the ratio image PD/T2. Potential conspicuous, or so-called "hot" lesions were isolated in terms of ratios that fell between two fixed threshold values. These "conspicuous lesion pixels" were then further refined by noting that hot lesions are identified with high intensity profiles in both the PD and the T2 images.

In Ref.[6] scatter plots of PD versus T2 were used to form a threshold area in the PD versus T2 2-dimensional space such that high-high regions could be identified with conspicuous lesions. In that work, it was pointed out that conspicuous lesion pixels tend to form a fairly narrow line, or "beam", in the PD-T2 feature space. Therefore it was postulated that extending the beam back into a lower intensity PD-T2 region could possibly identify candidate pixels involving a partial volume effect. There, it was suggested that the beam, i.e., the line and its width, could be recovered from conspicuous lesion pixels via a Hough transform. Then by following the beam into the adjacent, non-conspicuous lesion, part of the scatter plot, tissue pixels which are subjected to the partial volume effect could be identified. Figs. 1(a,b) show the the PD and T2 images for one particular slice that involves "hot" lesions, and Fig. 1(c) shows the feature space, PD versus T2, with a simple threshold indicated by vertical and horizontal lines. (In all our figures, we have used reverse video to display images — therefore hot spots are dark.) As can be seen in Fig. 1(d), in fact a simple thresholding procedure with a chosen fixed threshold can successfully identify candidate conspicuous lesion pixels in a straightforward manner, for this MR slice.

Since the partial volume effect gradually increases with distance from lesion centers, in Ref.[6] a probability was also assigned to pixels in the extended beam area such that subtle lesions could be more closely identified. As well, spatial information, as in Ref.[1] , was used to refine the tissue segmentation further.

Here, we adopt the scheme of extending the beam into the non-conspicuous lesion region of the feature space, but concentrate on developing thresholds in PD and T2 automatically. To do so, first we transform the feature space via a nonlinear mapping (cf. Ref.[7]) such that the entire feature space plot resembles as much as possible a single linear structure. To ascertain just what remapping best accomplishes this task, we make use of a particular remapping and then employ a line finder from robust statistics to determine how good a fit the scatter plot is to a single line in the transformed space. The idea is to reduce a single MR brain slice, or for that matter the entire brain, to a line.

2.2. Data Re-mapping

In order to apply a "beam" method, whereby we extend the "arm" of conspicuous lesion pixels, seen in Fig. 1(c), into the partial volume effects region of the feature space, we first need to identify conspicuous lesion pixels. Since every slice may appear differently, and in fact we can expect to see no conspicuous lesion pixels at all in many slices, a useful approach is to use a robust regression (see, e.g., Ref.[8]) on *all* pixels available and thus produce a threshold for outlier points in the feature space that lie too far from the center, and are also too high both in PD and in T2.

The advantage of a robust regression is that such an estimation method is much less influenced by noise and also by legitimate pixels that truly are not governed by the hypothesized shape of the data in the feature space. The simplest regression we can carry out is fitting the data to a *line*. However, to that purpose it would be useful to first re-map the PD-T2 data so that all pixels more or less fall on a line.

By trial and error, we attempted to fit re-mapped data using several different data transforms. I.e., we performed a robust regression after remapping the data, and then examined the coefficient of determination R^2 to see which re-mapping gave the highest value. We found that the best re-mapping was given by plotting log(T2) versus log(AV),

where we denote the average of PD and T2 by AV≡(PD+T2)/2. Fig. 2(a,b,c,d) show two slices, in PD and T2, with the first slice displaying a "hot" lesion and the second slice displaying a subtle one.

Since we will identify outliers using a robust statistical estimator, it makes sense to collapse all pixels for all brain slices into a single 2D scatterplot. The rationale for this is that if there are indeed conspicuous lesion pixels, then the robust regression will identify them from contributions from one of the slices at least. On the other hand, if there are none, then the regression should be able to report that too. Fig. 2(e) combines all pixels from all slices, and shows that in fact it is not always a simple matter to select a threshold for identifying conspicuous lesion pixels. However, we notice that the re-mapping shown in Fig. 2(f) produces a much more linear structure in feature space.

In the re-mapped space, we propose carrying out a series of three regressions: **Step A:** the first regression is to find the line in Fig. 2(f); **Step B:** the second regression finds values *along the line* which are at the high end — these we identify with candidate conspicuous lesion pixels; and **Step C:** we then find the best line for the conspicuous lesion pixels alone, and project it through the remaining pixels to identify possible pixels displaying a partial volume effect.

2.3. Robust and LS Regression in 2D

In Fig. 2(f) we show the line found using a robust regression (**Step A**), and also show, dashed, the similar line found using ordinary least squares (LS) regression. A robust regression treats points that do not obey the hypothesized linear shape as outliers on top of a basic linear structure. Several robust models exist for discounting such outliers. Here we use a Least Median of Squares method.[9,10] A least squares method cannot always reliably find the underlying linear structure because even a single outlier far from the correct result can unduly affect the fit.

The meaning of the LMS fit can be understood as follows. Consider a 1–dimensional collection of data x_j. The LS estimate of x_j is its average. To find the LMS estimate, order the data from low to high and successively bracket half the data, from 1 to $n/2$, then from 2 to $(n/2+1)$, etc. The smallest gap $x_j - x_{n/2+j-1}$ is the one sought; the LMS fit is the midpoint of that interval.

For straight line regression with the model $y = mx + b$, find that strip covering half of all the data points that has the narrowest width in the y–direction. The middle of that strip is the LMS fit. A theorem[8] limits the number of random pairs of data needed to perform the optimization.

The standard approach to outlier detection[9] is to consider the robust dispersion estimate

$$s_0 = K \left(1 + \frac{5}{N-p} \sqrt{\underset{i}{\text{med}}\ r_i^2} \right) \tag{1}$$

for N cases and p independent variables. Here, r_i is the residual r for the i^{th} case, $r = y - mx - b$, and, in terms of the Gaussian probability function Φ, $K = 1/\Phi^{-1}(0.75) \simeq 1.4826$.

Then an pixel is *accepted* as corresponding to the model if $|r_i/s_0| \leq 2.5$; else the point is an outlier and is rejected. Thus the utility of the LMS method is that it concomitantly produces a confidence interval beyond which one may identify values as outliers. For the LMS method, as much as 50% of the data can be arbitrarily corrupted without affecting the fit, provided that there are indeed data points that do fit the chosen model.

The linear structure of log(T2)-log(AV) plots is borne out by the value of the robust version of the coefficient of determination, R^2 =90.9%. In comparison, R^2 for other re-mappings were: T2–PD: 58.4%; T2–AV: 90.6%; log(T2)–log(PD): 61.2%, T2–log(AV): 87.7%.

In Fig. 2(f), the robust and standard LS estimator regression lines are very close to each other; nevertheless, use of a robust estimator in Step A has the important side-benefit of automatically producing a dispersion estimate and 2D band within which we can be confident that the data does indeed follow the line.

Moreover, a robust regression is in fact necessary for the next step, Step B, which consists of finding a threshold. In Step B, we make use of the confidence interval from Step A to guide the second regression. And for the final step, extending the "beam" back through non-conspicuous lesion pixels (Step C), we again need a 2D robust regression. The final regression provides an identification of possible partial volume effect pixels.

Let us denote by $\mathcal{T}_A$ the 2D threshold region established by Step A.

2.4. Robust Regression in 1D — Threshold Detection

For Step B, let us now perform a second transform on the data, moving the intercept to zero and rotating so the regression line becomes the new abscissa. Then with (x, y)=(log T2, log AV), if the regression yields intercept and slope $\{i, s\}$, we translate $y \rightarrow y - i$ and rotate by $\theta = -\arctan(s)$. The resulting rotated data is displayed in Fig. 3(a).

To find a threshold, we look for pixels that are too far from the *mode* of x values in the rotated data. Regardless of the *appearance* of the data in feature space (as in Fig. 3(a), the data may be widely dispersed), a *robust* estimator can find the correct value of the mode of the data — i.e, where the data is most concentrated.

The identification of outliers of the model line in rotated log(T2)–log(AV) feature space is equivalent to the special case of 1D location, with $p \equiv 1$ in eq. (1). Fig. 3(b) shows outliers detected for the data considered.

Let us denote by $\mathcal{T}_B$ the threshold established by Step B. Call $\mathcal{T}_B^U$ and $\mathcal{T}_B^L$ the upper and lower ends of the accepted region in Fig. 3(b), respectively. For the image corresponding to Fig. 3(b), $\mathcal{T}_B$ equals (PD,T2) = (187.0, 216.6), when transformed back to non re-mapped coordinates. (But of course we cannot simply use these limits as a simple threshold since we need to consider distance along a line in log-log space.)

2.5. Final Regression in 2D — Conspicuous Pixel Detection

We can more accurately determine the "beam" corresponding to conspicuous lesion pixels, that we wish to project back through the rest of the slice, if we carry out a regression on the candidate conspicuous pixels by themselves. Doing so will automatically establish a confidence interval beyond which pixels are flagged as outliers.

However, we have more than statistical knowledge about conspicuous pixels — we identify them with high-high regions of the feature space. Therefore here we modify the definition of outliers given following eq. (1) by *accepting* pixels as conspicuous provided the residual r is negative: i.e., we accept as conspicuous pixels lying above the main regression line, as shown in Fig. 3(c).

To guide the regression, we perform Step C, i.e., 2D regression on conspicuous lesion pixels only, by utilizing only conspicuous pixel data that passed thresholds $\mathcal{T}_A$ and $\mathcal{T}_B$. Let us denote by $\mathcal{T}_C$ the 2D strip produced by this final regression: this beam is shown in Fig.3(c).

Pixels lying in this beam which are above $\mathcal{T}_B^U$ can be identified as conspicuous. Mapping conspicuous pixels back to the 2D MR images yields an identification of candidate conspicuous lesion pixels. These pixels may need to be further culled by other methods (cf. Ref.[1]). Here, an image mask pixel is set to 1 if the corresponding candidate conspicuous lesion pixel is 9-connected.

Figs. 3(d,e) show mask images for these areas in the same slices shown in Figs. 2(a,b) and (d,e), respectively. As can be seen, the method does a reasonable job of finding potential conspicuous lesion pixels.

2.6. Identification of Partial Volume Effects

For each MR slice, we project the beam shown in Fig.3(c) back through the remainder of the slice, disqualifying pixels above $\mathcal{T}_B^U$ and below $\mathcal{T}_B^L$.

We should also like to further process candidate partial volume pixels spatially by associating a probability to a location based on its proximity to a conspicuous pixel. Here we adopt the probability developed in Ref.[6]: we form a product image of the conspicuous pixel mask times the original PD image, and then convolve with a Gaussian. Then we may identify as pixels having a partial-volume effect those non-conspicuous pixels that are inside the beam, and have high enough probability in the convolved image. Figs. 4(a-d) show the convolved conspicuous image and associated partial volume image mask, for the same MR slices as in Figs. 2(a,b) and (c,d).

3. ALGORITHM

In summary, the algorithm can be stated as:

1. **Step A**:
   (a) Combine PD and T2 data from all slices into two "images".
   (b) Find robust 2D best fit line through log(T2) vs. log(AV) data. This gives $\mathcal{T}_A$.

2. **Step B**:

   (a) Translate and rotate to make best-fit line the new abscissa.

   (b) Find robust 1D location estimate: this gives $\mathcal{T}_B$. Call $\mathcal{T}_B^U$ and $\mathcal{T}_B^L$ the upper and lower ends of the accepted region in Fig. 3(b), respectively. Pixels above $\mathcal{T}_B^U$ are candidate conspicuous lesion pixels.

3. **Step C**: Find the beam: for pixels that are non-outliers with respect to $\mathcal{T}_A$ (i.e., pixels that are near the line), for *outlier* pixels with respect to $\mathcal{T}_B^U$ (i.e., high-high pixels), find robust best-fit 2D line through conspicuous pixels; this gives 2D threshold region $\mathcal{T}_C$ for the beam, as in Fig. 3(c).

4. For each slice separately,

   (a) Identify conspicuous lesion pixels: rotated (log T2, log AV) pairs that lie above $\mathcal{T}_B^U$ and in the beam $\mathcal{T}_C$. Cull conspicuous pixels that are not 9-connected, with results as in Fig.3(d,e).

   (b) Identify pixels having partial volume effects: Make a probability density out of the PD image by multiplying it by the conspicuous-lesion mask image and convolving with a Gaussian. Identify as pixels having a partial-volume effect those non-conspicuous pixels that are inside the beam, and have high enough probability of being close to a conspicuous pixel. Here the Gaussian is 9×9 with standard deviation 2, and high probability is taken as being above the 80th percentile.

4. NO IDENTIFICATION OF CONSPICUOUS LESIONS IF NONE EXIST

An important property of robust methods is that outliers are flagged only if such actually exist. Fig. 5 uses data for which few if any conspicuous pixels exist. Step A produces a line (and the robust method will produce a low R^2 if there is not a good line); Step B does find some high- and low-end pixels; but Step C finds essentially no candidate conspicuous lesion pixels, since we use data that is both inlier data from Step A and outlier data from Step B as input to the final regression. I.e., if slices are combined to find the beam, and if at least one slice does contain enough conspicuous lesion pixels, the beam will correctly identify any candidate conspicuous lesion pixels in any slice. But using data with essentially no conspicuous lesion data means that we attempt to form the beam using almost no data. Here, the number of pixels identified as high-high within the confidence interval for the final regression is small — and in fact no pixels survive our 9-connected screening step. Thus the robust-based algorithm finds conspicuous lesion pixels if they exist, and indicates none if they do not exist.

5. DISCUSSION

The beam method set out here, with automatic identification of thresholds, was found to produce good results for cases studied to date. In a related study Ref.[6], the beam method was shown to produce a lesion volume correlation with manually outlined lesions of 0.97 on 8 data sets. Robust methods in general mitigate the effects of noise and other errors, e.g. the effect of repositioning on volume measurements.[5] The results produced for the present method are promising. Further work will be aimed at identifying how the contents of the beam in feature space, along the length of the beam as it proceeds from the conspicuous high-high end into the remaining data area, is mapped into lesion, grey matter, and CSF areas in slices. This should give insight into how the robust method might be best utilized to identify lesion pixels and how schemes for inclusion or exclusion of data may further guide the statistical method toward the best accuracy.

REFERENCES

1. K. Krishnan and M.S. Atkins. Segmentation of multiple sclerosis lesions in MRI - an image analysis approach. In *Medical Imaging 1998*, pages 1106–1116. SPIE Vol. 3338, 1998.
2. B. Johnston, M.S. Atkins, B. Mackiewich, and M. Anderson. Segmentation of multiple sclerosis lesions in intensity corrected multispectral MRI. *IEEE Transactions on Medical Imaging*, 15(2):154–169, April 1996.
3. J.G. Sled, A.P. Zijdenbos, and A.C. Evans. A nonparametric method for automatic correction of intensity nonuniformity in MRI data. *IEEE Trans. Med. Imag.*, 17:87–97, 1998.
4. M.S. Atkins and B. Mackiewich. Fully automatic segmentation of the brain in MRI. *IEEE Trans. Med. Imag.*, 17:98–107, 1998.

5. M.J. Firbank, A. Coulthard, R.M. Harrison, and E.D. Williams. Partial volume effects in MRI studies of multiple sclerosis. *Mag. Res. Imaging*, 17:593–601, 1999.
6. M.S. Atkins and Z. Tauber. Overcoming partial volume effects in MR tissue analyses. In *Proc. of the International Society for Magnetic Resonance in Medicine*, 1999.
7. H.S. Zadeh, J.P. Windham, D.J. Peck, and A.E. Yagle. A comparative analysis of several transformations for enhancement and segmentation of magnetic resonance image scene sequences. *IEEE Trans. Med. Imag.*, 11:302–318, 1992.
8. P. J. Rousseeuw and A. M. Leroy. *Robust Regression and Outlier Detection*. Wiley, 1987.
9. P. J. Rousseeuw. Least median of squares regression. *J. Amer. Stat. Assoc.*, 798:871–880, 1984.
10. P. Meer, D. Mintz, D. Y. Kim, and A. Rosenfeld. Robust regression methods for computer vision: a survey. *Int. J. Comput. Vision*, 6:59–70, 1991.

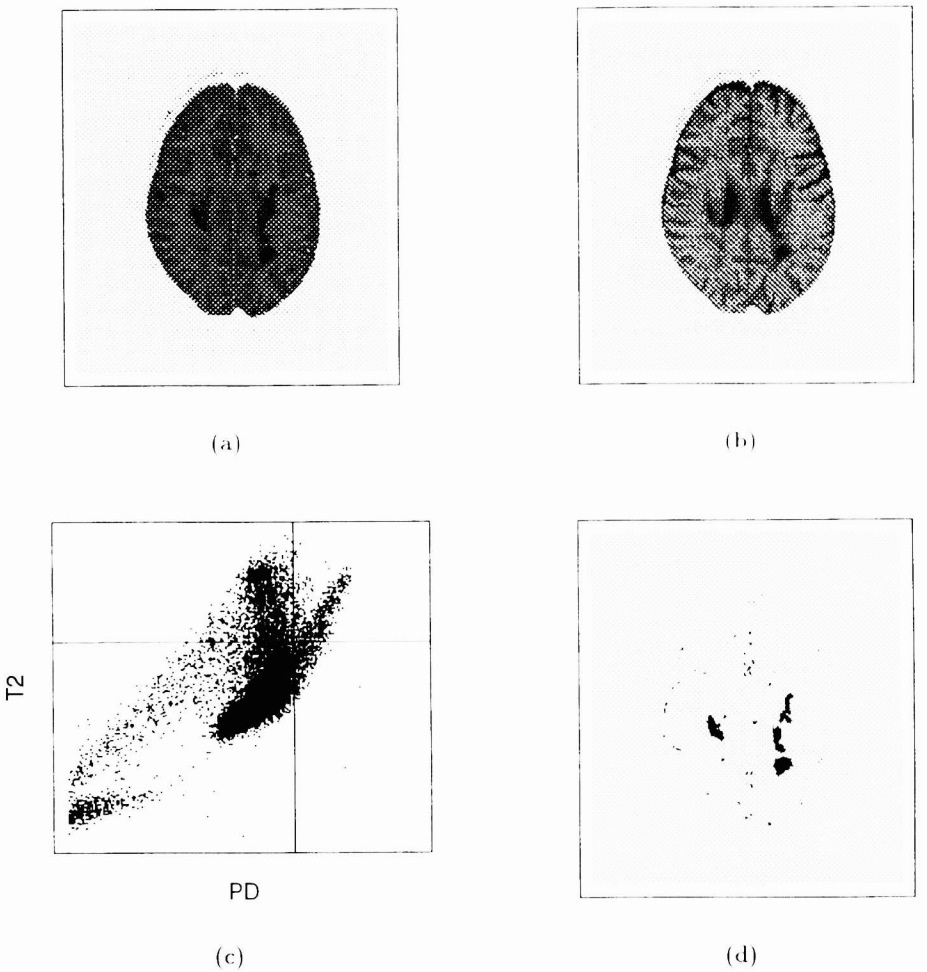

Figure 1. (a): PD image. (b): T2 image. (c): PD – T2 feature space: simple threshold indicated by lines. (d): Simple thresholding: (PD $> \mathcal{T}$) && (T2 $> \mathcal{T}$), where here $\mathcal{T} \equiv 165$.

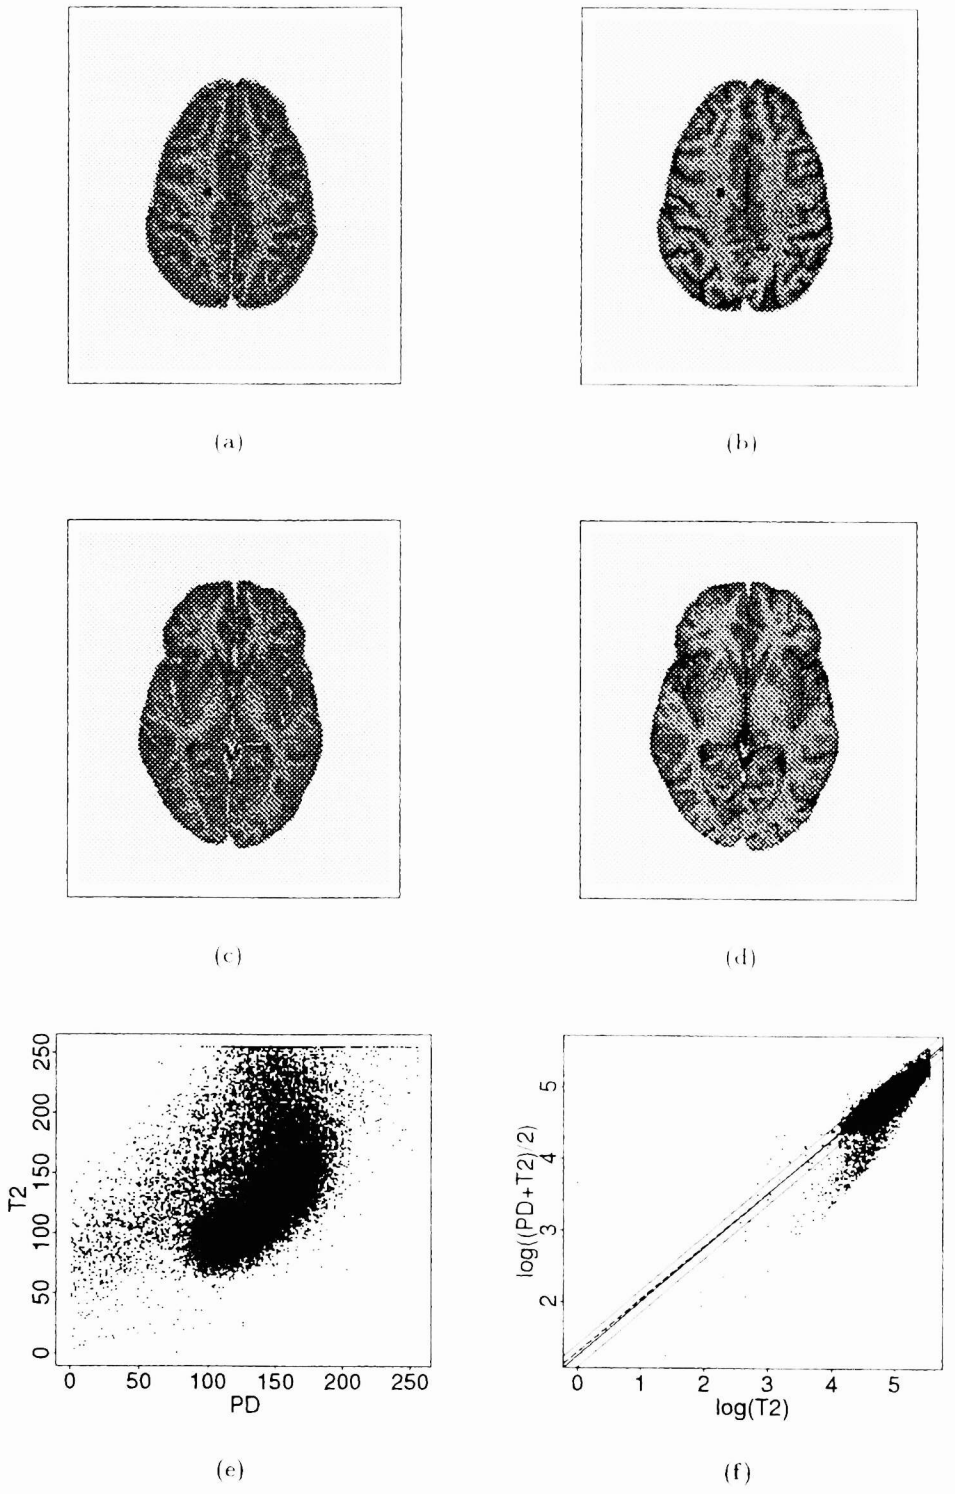

Figure 2. (a): PD image, with conspicuous lesion: "Patient 303", slice 16. (b): T2 image, with conspicuous lesion. (c): PD image, with subtle lesion: "Patient 303", slice 10. (d): T2 image, with subtle lesion. (e): PD – T2 feature space. (f): log T2 – log AV feature space, with AV=(PD+T2)/2. Solid line: robust regression; dashed line: least squares regression; dotted lines: robust confidence interval $\mathcal{T}_{\mathcal{A}}$.

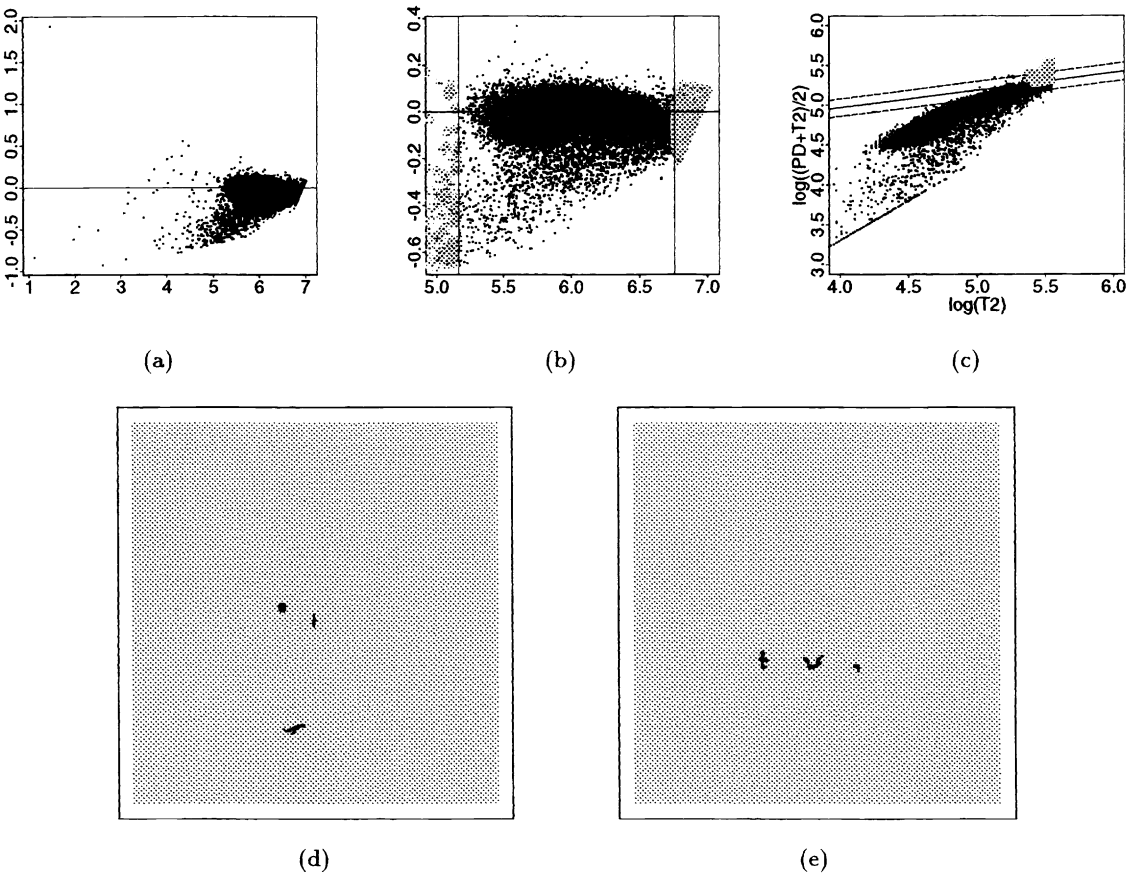

Figure 3. (a): Rotated feature space points, utilizing data from all brain slices. (b): Outlier detection (very dark image pixels are not shown, for display). Outliers are shown circled. The outlier region is $\mathcal{T}_B$. (c): Inlier beam ($\mathcal{T}_C$) arises from regression using conspicuous pixels only. (d): Candidate conspicuous pixels for slice showing "hot" lesion. (e): Candidate conspicuous pixels for slice showing subtle lesion.

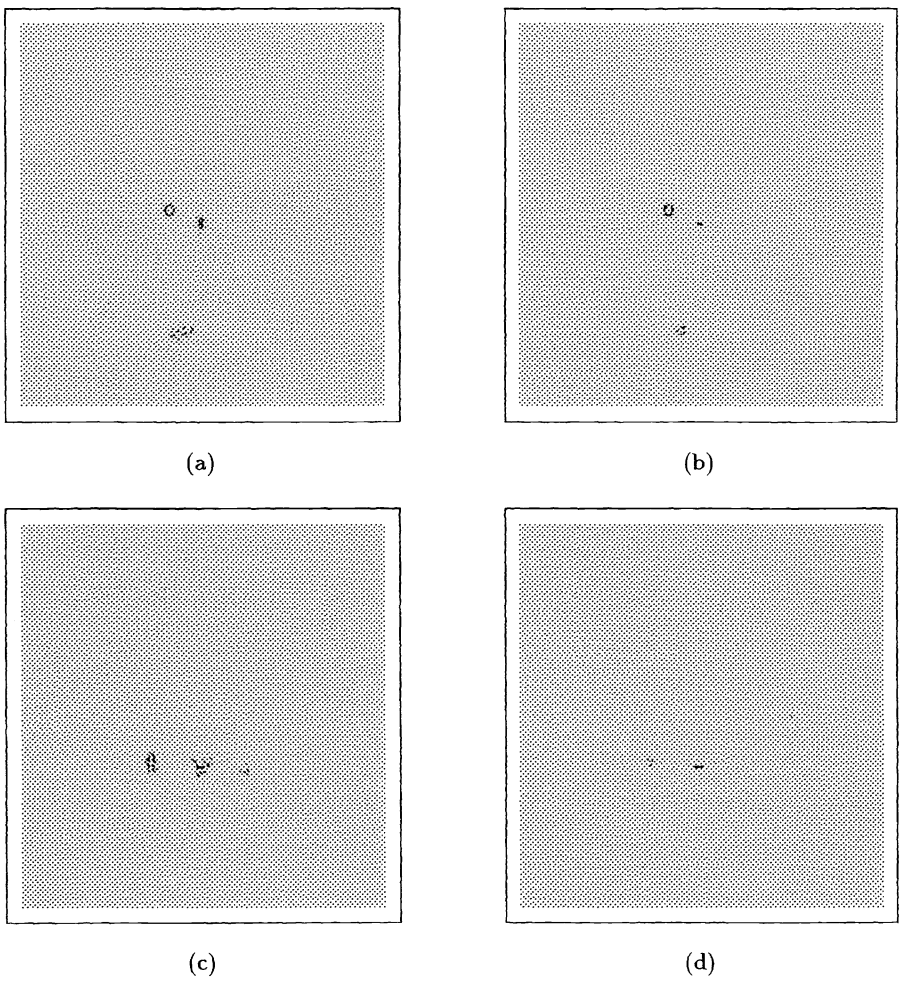

Figure 4. (a): Convolution of product of conspicuous pixel mask with PD image, for slice showing hot lesion. (b): Partial-volume pixels are identified as those in conspicuous pixel beam, with high probability in Fig. (a). (c): Convolution image for slice showing hot lesion. (d): Partial-volume pixels.

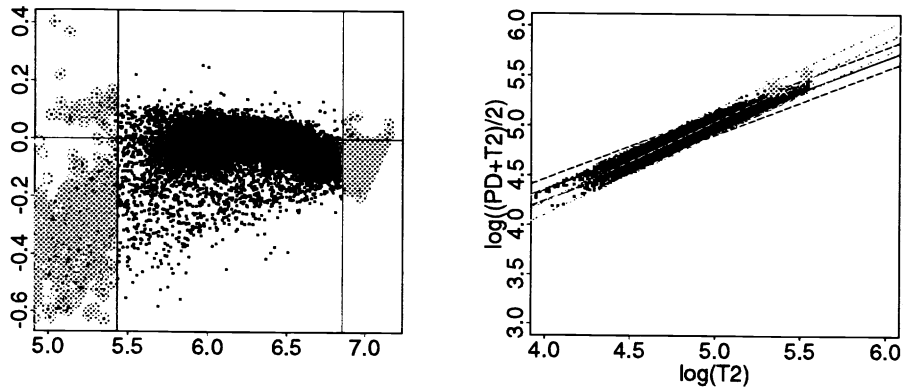

(a) (b)

Figure 5. (a): Data with negligible conspicuous lesion data: Step B produces $\mathcal{T}_{\mathcal{B}}$. (b): The Step C regression for candidate conspicuous lesion pixels uses inlier data with respect to the confidence interval $\mathcal{T}_{\mathcal{A}}$ for Step A, and outlier data with respect to $\mathcal{T}_{\mathcal{B}}$. Regressions for both Step A and Step C are shown in the figure. Only a few pixels are identified as conspicuous, and none survive further spatial processing.