

Visualisation of the pattern of contrast enhancement in dynamic breast MRI

Andrew Mehnert
Centre for Sensor Signal and Information
Processing
The University of Queensland,
Brisbane, Queensland, Australia
mehnert@itee.uq.edu.au

Ewert Bengtsson
Centre for Image Analysis
Uppsala University
Uppsala, Sweden
ewert@cb.uu.se

Kerry McMahon and Dominic Kennedy
Queensland X-Ray
Greenslopes Private Hospital
Greenslopes, Queensland, Australia

Stephen Wilson and Stuart Crozier
School of Information Technology and
Electrical Engineering
The University of Queensland
Brisbane, Queensland, Australia

Abstract

A new pixel-mapping method for visualising contrast uptake in dynamic MR images of the breast is presented. The method reduces the sequence of images of a single spatial slice over time to a single colour-coded image. This is achieved by fitting a linear-slope model pixel-wise to the slice time series and using the fitted parameters to define HSV colour space coordinates. The model parameters are related to the shape of the signal intensity-time curve at each pixel. The effect is that pixels with rapid and significant initial postcontrast enhancement appear brighter and more saturated, whilst the nature and degree of intermediate and late postcontrast enhancement is reflected in the colour hue. Preliminary results are reported for six subjects with suspicious MRI findings subsequently confirmed by pathology. The results suggest that the method shows promise as a replacement for, or adjunct to, the review of the raw time series data and/or associated difference images in the clinical setting.

1. Introduction

Magnetic resonance (MR) imaging of the breast, before and after the administration of an extracellular gadolinium-containing contrast agent, can be used to detect and characterise breast diseases [1]. In particular the pattern of enhancement, i.e. the change in signal intensity over time, is an important criterion for the differentiation of malignant

from benign lesions. The patterns for most cancers show an early steep rise within five minutes of contrast-agent injection, followed by a plateau, and then washout, whilst those for benign lesions either do not enhance, or exhibit slowed continued enhancement with delayed washout [2]. A variety of methods for analysing the change in signal intensity over time have been reported in the literature including subjective (qualitative) classification of the shape of the signal intensity-time curve, measurement of simple quantitative parameters associated with the time-curve (e.g. percentage increase in signal intensity 90 s after administration of the contrast agent and the percentage increase achieved at the maximum signal intensity), pharmacokinetic modelling (parameters derived from compartmental models of dynamic contrast enhancement), and neural networks [1]. In the routine clinical setting, however, the most commonly adopted method is the qualitative approach [3]. Typically the clinician: (i) reviews images of the raw time series for each spatial slice, or of subtraction images (postcontrast minus precontrast), and identifies areas of suspicious enhancement; (ii) uses software produced by the MRI equipment manufacturer to select regions of interest (ROIs) and to plot their signal intensity-time curves; (iii) makes a visual assessment of the morphology and architecture of the suspicious lesions (as they appear in higher resolution anatomical images rather than the dynamic images); and (iv) combines this information together with patient history to classify the suspicious lesions.

There are essentially two approaches to the analysis and presentation of dynamic breast MRI data: ROI analysis

(region-based) and pixel-mapping (pixel-based) [1]. ROI analysis methods permit the user to select regions of interest and to plot the associated enhancement curves. Pixel-mapping methods, on the other hand, display quantitative enhancement information as a colourmap co-registered with an anatomical image. The enhancement curves generated by ROI methods have good signal-to-noise ratio but lack spatial resolution, are prone to partial volume errors, and are sensitive to ROI selection and placement (e.g. the method does not inherently take account of the heterogeneity of tumour enhancement) [1]. Pixel-mapping methods have the advantage of not requiring the user to select an ROI thus reducing the possibility that a diagnostically significant lesion is overlooked, and of introducing partial volume errors because of ROI misplacement. However, the disadvantage is that pixel-mapping methods are more sensitive to noise, and in particular to patient movement during the dynamic examination.

There are two basic approaches to pixel-mapping: (i) colour coding simple quantitative parameters associated with the enhancement curve for each pixel (e.g. FUNC-TOOL by GE Medical Systems, Milwaukee, USA); and (ii) fitting a model to each pixel time series and colour coding the fitted parameters ([4], [5], [6]). A variation on the latter approach is the three-time-point (3TP) method of [7]. The 3TP method generates a colourmap from the intensity values measured at three judiciously chosen time points: the precontrast time plus two postcontrast times. The intensity difference between the first two time points is coded by colour intensity and the change between the second and third is coded by colour hue (red, green, and blue). The selection of the three time points is determined using an algorithm based on the fitting of a pharmacokinetic model (Tofts model) to the data with two free parameters K and ν_1 ; the remaining parameter values are prescribed by the MR imaging parameters and the contrast agent dose [8]. The algorithm generates several two-axis (K on one axis and ν_1 on the other) colour calibration maps; one for each pair of post-contrast time points. The map that best divides the $K - \nu_1$ plane determines the optimal pair of postcontrast times.

This paper presents a new pixel-mapping method for visualising significant contrast uptake in dynamic MR images of the breast. The method is based on the direct visualisation of the parameters associated with a pixel-wise fit of a *linear-slope* model to each slice image series. The model parameters can be easily related to the shape of the enhancement curve (specifically the nature and degree of early postcontrast enhancement, and of intermediate to late post-contrast enhancement). The method requires no calibration or selection of threshold parameters. Additionally the method utilises order statistic filtering to improve robustness to small in-slice movement. The new pixel-mapping method effectively reduces the sequence of images of a sin-

gle slice over time to a single colour coded image. The colour coding of each pixel is performed with respect to the HSV [9] colour model and encodes the shape of the enhancement curve. Preliminary results are reported for six subjects with suspicious MRI findings that were subsequently verified by pathology: three with benign lesions and three with malignant lesions. The results indicate that the proposed pixel-mapping method is a valuable visualisation tool that can assist the clinician with the identification of suspicious lesions. The method shows promise as a replacement for, or adjunct to, the review of the raw time series data and/or associated difference images.

2. Materials and methods

2.1. Image database

Image data from six subjects was used for this study. The data originates from routine breast MRI examinations performed by Queensland X-Ray, Greenslopes Private Hospital, Greenslopes, Queensland, Australia in the last four years. MRI examinations, of a single breast, were performed on a 1.5 T Signa EchoSpeed (GE Medical Systems, Milwaukee, USA) using an open breast coil which permitted the subject to lie prone. A 3D dynamic scan using an SPGR sequence of TE = 1.5 ms, TR = 5.4 ms, 10 degree flip angle, and acquisition matrix size 256×256 interpolated to 512×512 (ZIP512) was typically used. Gadopentate dimeglumine, 0.2 mmol/kg, was administered manually at a rate of about 3 ml/s. The number of sagittal slices per volume acquired for each subject depended on the size of the breast and ranges from 22 to 56. The number of volumes per scan for each subject, including one precontrast volume, ranges from 7 to 11. Slice thicknesses, with 50% overlap (ZIP2), range from 4.5 to 5 mm. The resulting slice images are of size 512×512 pixels with an 8-bit per pixel intensity scale.

The six subjects were deliberately chosen: three examples of enhancing lesions subsequently confirmed to be malignant, and three of enhancing lesions subsequently confirmed to be benign. The MRI finding of the respective radiologist as well as the subsequent pathology for each of the subjects are shown in Table 1. The pathology together with screen captures of the ROIs selected by the radiologist (including the corresponding enhancement curves produced using FUNCTOOL) provided the *ground truth* for the data. A sample screen capture and associated enhancement curve are shown in Figure 1.

2.2. Slice data normalisation

For the purposes of this study only the dynamic series for each slice containing an ROI was used; i.e. one series of images of a particular slice over time for each subject. The

Table 1. MRI findings and pathology for the subjects in this study.

Subject	MRI finding	Pathology
1	8 mm lesion	malignant: invasive ductal carcinoma grade 2
2	5 mm lesion	malignant: ductal carcinoma
3	8 mm lesion	malignant: invasive ductal carcinoma grade 3
4	16 mm × 11.8 mm × 11.5 mm lesion	benign: fibrocystic change
5	small enhancing lesion	benign
6	focal area of suspicion < 3 mm	benign: atypical ductal hyperplasia

time interval between the acquisition of successive post-contrast slices is a fixed value for each subject. However, in practice the clinician acquires several precontrast volumes but retains only one of these (typically the one yielding the least amount of motion artefact in the difference images) for the purpose of constructing an enhancement curve. A consequence of this is that for any given slice in space, the difference between the acquisition time for the precontrast image and the first postcontrast image depends on which precontrast volume is chosen. This is illustrated (red overlay) in Figure 1; the width of the interval A is different to that of B to I. In this study, this anomaly was corrected by setting the time stamp of the precontrast slice to be that of the first postcontrast slice minus the fixed postcontrast slice interval. In addition, all of the times were offset so that the precontrast acquisition time is zero.

To attenuate noise and to compensate for small in-slice movements (on the order of one or two pixels) each slice image within each volume was filtered using a 3×3 order statistic filter (also called a *rank* filter or operator) [10] defined to replace the value of the central pixel in a 3×3 sliding window with the third largest value. This filter was chosen in preference to a mean or median filter because these filters are more likely to miss or diminish the response of small enhancing areas. It was chosen in preference to a maximum filter because the maximum filter is prone to select impulse-type noise artefacts.

Finally, each filtered postcontrast slice was subtracted (pixel-wise) from its corresponding filtered precontrast slice, and the precontrast slice pixels set to zero. The resulting intensity values thus represent relative MR units (i.e. relative to the precontrast values). This ensures that en-

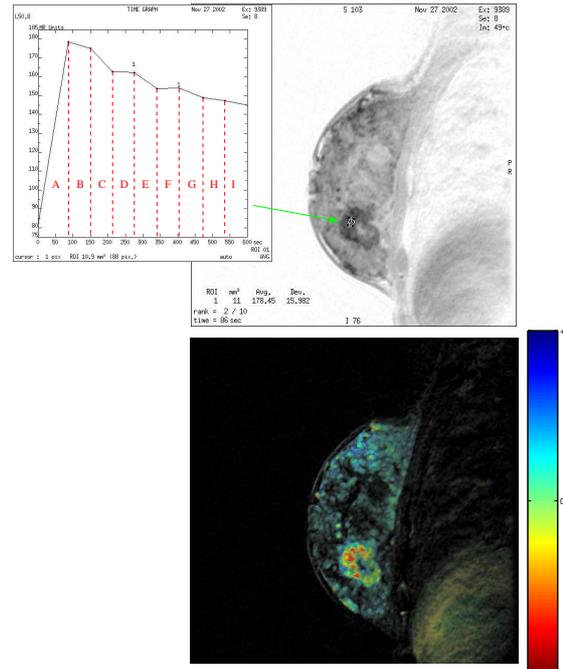


Figure 1. Top: Clinician-traced ROI for subject 1 and the corresponding enhancement curve produced using FUNCTOOL (Note: The red overlay is not produced by FUNCTOOL. Refer to the text for an explanation). Bottom: Proposed HSV visualisation.

hancement curves for individual pixels begin at $(0, 0)$.

2.3. Pixel-wise model fitting

In the work of Kuhl et al. [11] three basic types (shapes) of enhancement curve were identified as shown in Figure 2. Type I curves are characterised by rapid early postcontrast rise followed by a continued straight line or curved rise, type II by a rapid initial rise followed by a plateau, and type III curves by a rapid initial rise followed by washout. This characterisation suggests a very simple model of enhancement based on two piece-wise line segments: the first segment describes the early postcontrast rise and the second describes the continued uptake (positive slope), plateau (zero slope), or washout (negative slope). This model is known as the *linear-slope* model in the plant and soil sciences [12]. Given a random sample of $i = 1, \dots, n$ observations on the intensity response Y_i of a given pixel at a corresponding time t_i , and assuming that the intercept of the first line segment is zero (as must be the case for the normalised data), the model has the form:

$$E[Y_i] = \begin{cases} \beta_1 t_i & \text{if } t_i \leq \alpha, \text{ and} \\ \beta_1 \alpha + \beta_2 (t_i - \alpha) & \text{if } t_i > \alpha, \end{cases}$$

where $E[Y_i]$ is the mean or expectation of the random variable Y_i , β_1 is the slope of the first line segment, α is the point (time) at which the two line segments meet, and β_2 is the slope of the second line segment. This model is not linear in its parameters (because of the product of α and β_2) and hence cannot be fitted using linear least squares (LLS). Rather it is necessary to use a non-linear least squares (NLS) algorithm such as the *Levenberg-Marquardt* or *Trust-Region* algorithms [13]. In contrast to LLS, NLS algorithms are iterative requiring the specification of initial parameter estimates [14]. For this study the Trust-Region algorithm, as implemented in MATLAB (The MathWorks, Inc., Natick, MA, USA), was used to fit the linear-slope model to the enhancement curve of each pixel using the following initial parameter estimates: $\hat{\alpha} = t_2$ (the first post-contrast time), $\hat{\beta}_1 = y_2/t_2$ (the slope of the line from the origin and joining the observed value at the first postcontrast time), and $\hat{\beta}_2 = 0$ (assumes the second line segment has no slope). Another issue with NLS algorithms is that there is no guarantee of convergence. Hence in this study the convergence status of each pixel-wise model fit was recorded. For the data used in this study the Trust-Region algorithm never failed to converge. Two examples of the fitted model are shown in Figure 2.

It should be noted that a more complex model of enhancement, the biexponential model (a two-compartment pharmacokinetic model [15]), was initially considered in this study. The model is defined: $E[Y_i] = \alpha_1 e^{-\beta_1 t_i} + \alpha_2 e^{-\beta_2 t_i}$. However, although the model can be convincingly fitted to time curves of pixels in enhancing regions, in many areas of non-enhancing tissue and in air it either fails to converge outright or does not do so within a fixed number of iterations. In the latter case the resulting parameter estimates typically have extreme values making interpretation difficult. Another issue with the biexponential model is that it is a four-parameter model and it is more difficult to visualise four-dimensional data than three. For these reasons the biexponential model was not used in this study.

2.4. Interpretation and visualisation

The three-parameter linear-slope model above suggests that a way to visualise the model fit at each pixel is as a colour specified with respect to a three-dimensional colour coordinate system such as that defined by the RGB or HSV colour models [9]. A naive visualisation can be achieved using the parameters α , β_1 , and β_2 as RGB or HSV colour-space coordinates. The problem with this approach is that the dynamic range for these parameters varies from subject to subject (this is in part a consequence of the variability in tissue-MR interaction between patients [1]). As a consequence the meaning of the various colours is difficult to interpret and even more difficult to compare between sub-

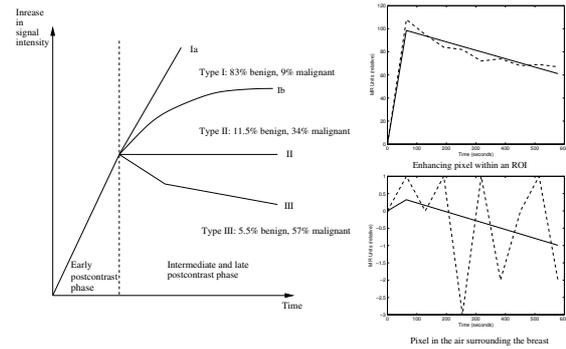


Figure 2. Left: Three types of signal intensity-time curves and the respective proportion of benign and of malignant lesions that exhibit each shape-type [11]. Right: Two examples of the linear-slope model fit (solid line) to the normalised slice data (dashed line) for subject 1.

jects.

A better approach is to colour code the shape of the enhancement curve at each pixel. The product $\alpha\beta_1$ (the height at the join point) is a measure of the degree of early postcontrast enhancement. The slope β_2 is a measure of the nature (i.e. continued rise, plateau, or washout) and degree of the intermediate and late postcontrast enhancement. These quantities can be visualised simultaneously in HSV colour space as follows. The saturation (S) and lightness (V) coordinates can be used to encode the product $\alpha\beta_1$ (early postcontrast enhancement) and the hue (H) component can be used to encode β_2 . The resulting plot will then show brighter and more saturated pixels in areas of rapid early postcontrast enhancement, and the colour hue will indicate the rate of intermediate and late postcontrast enhancement. There are, however, three problems with this approach. Firstly, if β_2 is simply scaled to $[0, 1]$ then the hue associated with the value zero may be different for different slices (either from the same subject or for another subject). Secondly, in the HSV colour model as the value H varies from 0 through to 1, the hue progresses from red through orange, yellow, green, blue, magenta, and back to red. This means that when visualising β_2 the colour red can occur at both extremely positive and extremely negative values (see Figure 3). Thirdly, the dynamic range of the product $\alpha\beta_1$ varies from individual to individual and can be greatly influenced by extreme values (e.g. due to background noise and motion artefact).

A solution to the first problem is to clamp zero slope to the middle of the H range. A solution to the second problem is to remap the hue scale to obtain hues that range from red at one extreme (washout) through green (plateau) to blue

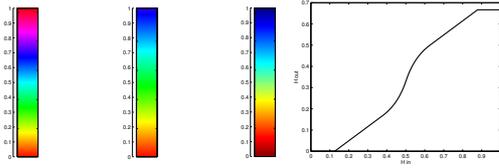


Figure 3. Hue colour scales. From left to right: the full range of H values, H values in the interval $[0, 0.7]$ linearly scaled to the interval $[0, 1]$, and H values described by the function shown and then linearly scaled to the interval $[0, 1]$.

at the other extreme (continued rise). Two possibilities are shown in Figure 3. The non-linear remapping is the better solution because it gives a better gradation of hues between the red and blue extremes (note the wide band of green hues in the middle of the truncated HSV scale). The third problem can be overcome, or at least diminished, by constraining the visualisation to only those pixels for which $\beta_1 > 0$ and $\alpha \geq 0$. An example of the proposed HSV visualisation method is shown in Figure 1.

3. Results: Comparison with clinically marked ROIs

Each slice corresponding to an ROI in Table 1 was colour-coded using the proposed HSV visualisation method. In all six cases the ROI marked by the radiologist coincides with the most prominent cluster of pixels in the corresponding HSV map. Moreover the hues associated with these clusters are indicative of the nature of enhancement in the intermediate and late postcontrast phase: red hues for pixels with a high degree of washout (indicative of malignancy), blue hues for pixels with significant continued enhancement (typical of benign lesions), and green hues for pixels with plateau. The result for subject 1 is shown in Figure 1 and shows a strong correlation between the ROI marked by the clinician and the orange/red spots prominent in the HSV visualisation. Interestingly, at least two other smaller clusters of *hot* pixels, adjacent to the ROI, appear suspicious. Results for another three subjects are shown in Figure 4. Again, in the case of subject 2 several other smaller clusters of hot pixels appear suspicious. In the case of subject 4 several adjacent clusters of light-green pixels appear to be focal areas of benign enhancement. In the case of subject 5 the clusters of hot pixels in the lower right are located within the liver and not the breast tissue and so are not relevant. The smaller focal areas in yellow at the top of the breast, however, appear suspicious.

4. Summary and conclusion

We have presented a novel pixel-mapping method for visualising the pattern of contrast uptake in dynamic breast MRI. Each slice pixel is colour-coded to reflect the shape of its signal intensity-time curve. This is done by fitting a linear-slope model to each slice pixel and expressing the associated parameters that describe the nature and degree of early, and of intermediate to late postcontrast uptake as coordinates in HSV colour space. The effect is that pixels with rapid and significant initial uptake appear brighter and more saturated, whilst the nature and degree of the intermediate to late postcontrast enhancement is reflected in the particular colour hue. We applied the method to data from six subjects—three with benign lesions and three with malignant lesions—and confirmed that the most prominent clusters of pixels apparent in the HSV visualisation coincide with the ROIs of suspicious lesions selected by the radiologist. The results suggest that the method shows promise as replacement for, or adjunct to, the review of the raw time series data and/or associated difference images.

The efficacy of the proposed method needs to be evaluated on a larger database. This will be the subject of further work.

References

- [1] R. Warren and A. Coulthard, eds., *Breast MRI in practice*. London: Martin Dunitz, 2002.
- [2] S. Sinha and U. Sinha, “Functional magnetic resonance of human breast tumors: Diffusion and perfusion imaging,” in *Techniques in bioinformatics and medical informatics* (F. Valafar, ed.), vol. 980 of *Annals of the New York Academy of Sciences*, pp. 95–115, New York, N.Y.: The New York Academy of Sciences, 2002.
- [3] M. D. Schnall, “Breast MR imaging,” *Radiologic Clinics of North America*, vol. 41, pp. 43–50, 2003.
- [4] U. Hoffmann, G. Brix, M. V. Knopp, T. Hess, and W. J. Lorenz, “Pharmacokinetic mapping of the breast: A new method for dynamic MR mammography,” *Magnetic Resonance in Medicine*, vol. 33, no. 4, pp. 506–514, 1995.
- [5] G. J. Parker, J. Suckling, S. F. Tanner, A. R. Padhani, P. B. Revell, J. E. Husband, and M. O. Leach, “Probing tumor microvasculature by measurement, analysis and display of contrast agent uptake kinetics,” *Journal of Magnetic Resonance Imaging*, vol. 7, no. 3, pp. 564–574, 1997.

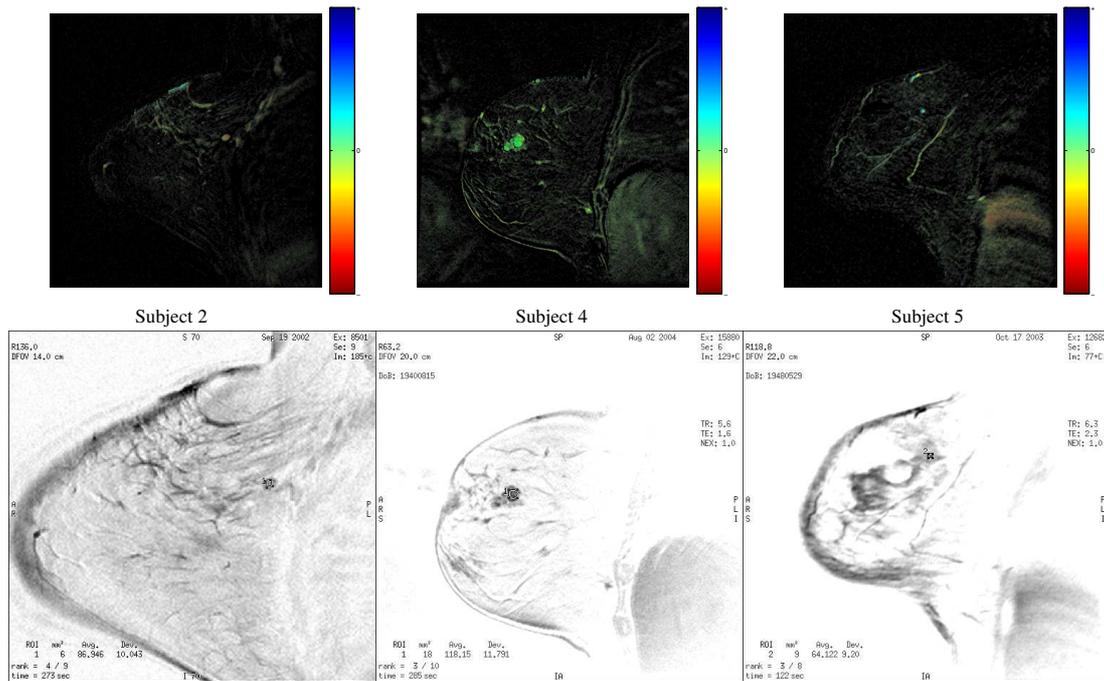


Figure 4. Comparison of the proposed HSV visualisation with the ROI screen captures made by the radiologist for three of the six subjects.

- [6] A. T. Agoston, B. L. Daniel, R. J. Herfkens, D. M. Ikeda, R. L. Birdwell, S. G. Heiss, and A. M. Sawyer-Glover, "Intensity-modulated parametric mapping for simultaneous display of rapid dynamic and high-spatial-resolution breast MR imaging data," *RadioGraphics*, vol. 21, no. 1, pp. 217–226, 2001.
- [7] H. Degani, V. Gusis, D. Weinstein, S. Fields, and S. Strano, "Mapping pathophysiological features of breast tumors by MRI at high spatial resolution," *Nature Medicine*, vol. 3, no. 7, pp. 780–782, 1997.
- [8] E. Furman-Haran, D. Grobgeld, F. Kelcz, and H. Degani, "Critical role of spatial resolution in dynamic contrast-enhanced breast MRI," *Journal of Magnetic Resonance Imaging*, vol. 13, no. 6, pp. 862–867, 2001.
- [9] J. D. Foley, A. van Dam, S. K. Feiner, and J. F. Hughes, *Computer graphics: Principles and practice*. The Systems Programming Series, Boston: Addison-Wesley, second ed., 1996.
- [10] H. J. A. M. Heijmans, *Morphological Image Operators*. Advances in Electronics and Electron Physics. Supplement 25, Boston: Academic Press, 1994.
- [11] C. K. Kuhl, P. Mielcareck, S. Klaschik, C. Leutner, E. Wardelmann, J. Gieseke, and H. H. Schild, "Dynamic breast MR imaging: Are signal intensity time course data useful for differential diagnosis of enhancing lesions?," *Radiology*, vol. 211, no. 1, pp. 101–110, 1999.
- [12] O. Schabenberger and F. J. Pierce, *Contemporary statistical models for the plant and soil sciences*. Boca Raton, FL: CRC Press, 2002.
- [13] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical recipes in C: The art of scientific computing*. Cambridge: Cambridge University Press, second ed., 1992.
- [14] W. N. Venables and B. D. Ripley, *Modern applied statistics with S-PLUS*. Statistics and Computing, New York: Springer, third ed., 1999.
- [15] S. B. Hladky, *Pharmacokinetics*. Manchester: Manchester University Press, 1990.