Mixture Models for Estimating Maximum Blood Flow Velocity

Caren Marzban, PhD, Wenxiao Gu, MS, Pierre D. Mourad, PhD

Objectives—A gaussian mixture model (GMM) was recently developed for estimating the probability density function of blood flow velocity measured with transcranial Doppler ultrasound data. In turn, the quantiles of the probability density function allow one to construct estimators of the "maximum" blood flow velocity. However, GMMs assume gaussianity, a feature that is not omnipresent in observed data. The objective of this work was to develop mixture models that do not invoke the gaussian assumption.

Methods—Here, GMMs were extended to a skewed GMM and a nongaussian kernel mixture model. All models were developed on data from 59 patients with closed head injuries from multiple hospitals in the United States, with ages ranging from 13 to 81 years and Glasgow Coma Scale scores ranging from 3 to 11. The models were assessed in terms of the log likelihood (a goodness-of-fit measure) and via visual comparison with the underlying spectrograms.

Results—Among the models examined, the skewed GMM showed a significantly ($P < .05$) higher log likelihood for 56 of the 59 patients and produced maximum flow velocity estimates consistent with the observed spectrograms for all patients. Kernel mixture models are generally less "robust" in that their quality is inconsistent across patients.

Conclusions—Among the models examined, it was found that the skewed GMM provided a better model of the data both in terms of the quality of the fit and in terms of visual comparison of the underlying spectrogram and the estimated maximum blood flow velocity. Nongaussian mixture models have potential for even higher-quality assessment of blood flow, but further development is called for.

Key Words—blood flow; brain; head injury; noninvasive; transcranial Doppler ultrasound

Received May 28, 2014, from the Applied Physics Laboratory (C.M., P.D.M.) and Departments of Statistics (C.M., W.G.), Neurological Surgery (P.D.M.), and Bioengineering (P.D.M.), University of Washington, Seattle, Washington USA. Revision requested July 30, 2014. Revised manuscript accepted for publication April 30, 2015.

We thank Paul R. Illian and David Morison for contributions made during an early phase of this project. This work received support from the National Institutes of Health (grant R43NS46824-01A1), National Space Biomedical Research Institute (grant SMS00701-2009-513), and PhysioSonics Inc. Dr Mourad has a financial interest in PhysioSonics.

Address correspondence to Caren Marzban, PhD, Department of Statistics, University of Washington, Box 354322, Seattle, WA 98195-4322 USA.

E-mail: marzban@stat.washington.edu

Abbreviations

GMM, gaussian mixture model; KMM, kernel mixture model; LL, log likelihood; MGM, modified geometric method

doi:10.7863/ultra.14.05069

The importance of estimating blood flow in major cerebral arteries has been well documented.[1–5] The reduction of blood flow through a major cerebral artery, such as caused by cerebral vasospasm, can lead to a wide range of disorders; therefore, monitoring blood flow has important clinical consequences.[6] Transcranial Doppler ultrasound imaging is one of the methods for assaying blood flow.[7,8] Vasospasm leads to reduced blood flow but an increased blood flow velocity; therefore, estimating the maximum blood flow velocity is of particular importance.[9]

The time series of a flow velocity histogram is called a spectrogram, and the time series of the maximum flow velocity is called an envelope. Despite their clinical importance, neither the spectrogram nor the maximum flow velocity can be observed directly. From a statistical perspective, they must be considered population parameters

that are to be estimated from data. Moreover, the very notion of maximum flow velocity is ambiguous because theoretically flow velocity has no upper bound. Many envelope estimation algorithms have been proposed[10–15]; one that is frequently[16] used is the modified geometric method (MGM).[10] That method essentially searches for a "kink" in the cumulative histogram of flow velocity and defines the maximum flow velocity as the flow velocity at which the kink occurs. It provides reliable envelope estimates but lacks the important feature of allowing one to quantify the uncertainty in maximum flow velocity. By contrast, the mixture model described by Marzban et al[15] allows one to quantify flow velocity in terms of percentiles. For example, one may examine the envelope corresponding to the 90th, 95th, and 99th percentiles of flow velocity. Each of these provides a unique and well-defined measure of "maximum" flow velocity, and the ability to do this is the main advantage of mixture models. Furthermore, the maximum flow velocity as defined by the MGM is in close agreement with that produced by the 90th percentile envelope in a mixture model.[15] As such, mixture models enjoy two important features: (1) they produce envelopes corresponding to multiple percentiles of flow velocity; and (2) one of the percentiles (ie, the 90th) produces an envelope that is approximately equal to the envelope estimated by the MGM.

However, the specific type of mixture model developed by Marzban et al[15] assumes that the distribution of flow velocity, at any point in time, as measured by a transcranial ultrasound device, is gaussian. The agreement between an envelope produced by such a gaussian mixture model (GMM) and that produced by the MGM suggests that either the distribution of flow velocity truly is gaussian, or that the envelope is not sensitive to violations of that assumption. As shown below, the distribution of flow velocity varies with time; at certain times, it is gaussian (or at least, near gaussian), while at other times it is not. The deviations from normality manifest themselves mostly as a skew. For this reason, one of the models described here assumes that the underlying distributions of flow velocity are skewed gaussian. Data are used to estimate the mean, the variance, and the shape (or skew) parameter of the distributions. Here, this model is referred to as a skewed GMM.[17–19] Another model discussed (briefly) does not assume a parametric form for the distribution of flow velocity at all; instead, the distribution is inferred by using kernel estimates.[20–24] Such kernel mixture models (KMMs) are flexible in that they accommodate a wide range of distributions (including gaussian). However, kernel methods involve a smoothing parameter whose value is a priori unknown. Generally, extreme values of the parameter correspond to either a smooth distribution (eg, gaussian) or a highly irregular distribution. One often uses some criterion (such as the maximization of likelihood) to optimize the smoothing parameter.

In this study, the skewed GMM and KMM were developed for estimating the envelope for each of 59 patients in our data set. Although quantitative measures for goodness of fit were used to compare the various models, qualitative comparisons were also used. This practice is necessary because, as the true envelope is not directly observable, the estimated envelope is assessed by visually comparing it with the underlying spectrogram.

Materials and Methods

Data

The data for this work were collected from a preclinical study involving 59 patients with closed head injuries from multiple hospitals in the United States, with ages ranging from 13 to 81 years and Glasgow Coma Scale scores ranging from 3 to 11. A PowerMode 100 transcranial Doppler device (Spencer Technologies, Northborough, MA) was used at the Harborview Medical Center (Seattle, WA), Columbia University Medical College (New York, NY), and the University of Texas Medical School (Houston, TX). The device has an ultrasound carrier frequency that ranges between 1.75 and 2 MHz and pulse repetition frequencies of 8000 to 10,000 Hz. The transcranial Doppler probe was placed on the temple of the patient by the sonographer after application of ultrasound gel. In all cases, the sonographer sought to produce an optimal spectrographic signal, defined here as one with a maximum systolic flow velocity in the main branch of the middle cerebral artery. To do so, the sonographer manipulated the angle of insonation of the middle cerebral artery by the probe, the power of the ultrasound, as well as the depth within the brain from which the Doppler spectral data were obtained. Further details on this data set can be found in previous studies by Marzban et al.[15,25] The former used only the ultrasound-derived transcranial Doppler spectra for the purpose of developing the aforementioned GMMs for envelope estimation. In the latter, the data were used to develop a model of arterial blood pressure, which can in turn be used to predict intracranial pressure. In accordance with the Institutional Review Board for each hospital, informed consent was obtained from all patients or their families.

With clinically approved transcranial Doppler units, blood flow velocity in the middle cerebral artery was measured with Doppler ultrasound. Data were collected for periods from 5 to 30 minutes. All retrospective data pro-

cessing and analysis were conducted at the Applied Physics Laboratory of the University of Washington. Although the Doppler spectral time series were initially sampled at 125 Hz, they were down-sampled at 40 Hz. This resolution was sufficient to resolve each patient's systolic rise, diachrotic notch, and diastolic minimum. For each patient, envelopes were produced over a fixed duration (118 sampled points, or about 3 minutes). This duration is sufficiently long to include several cardiac cycles while allowing details of the envelopes to be visually evident.

Data Processing and Statistical Analysis

In previous work, it was shown that the GMM is a natural extension of the MGM in that the former provides for multiple estimators of maximum flow velocity based on percentiles of the probability density function of the flow velocity. It was found that the envelope corresponding to the 90th percentile of the probability density function is comparable with the envelope estimated by the MGM.[15] As such, the GMM is superior (to the MGM) in that it allows one to estimate the highest blood velocities, and effectively their uncertainty, through multiple percentiles of the probability density function of flow velocity. The MGM is included in the analysis here only for the purpose of providing a comparison with the GMM. In fact, since the GMM outperforms the MGM, the skewed GMM and KMM were both compared with the GMM. The KMM leads to envelopes that are moderately superior to those based on the skewed GMM but not for all patients; for this reason, the KMM is discussed only briefly and only for the purpose of discussing future work.

As explained in Appendix 1, in a GMM, the distribution (more accurately, the probability density function) of flow velocity for each patient and at any time is assumed to be a linear combination of two gaussians. The coefficients, called mixing coefficients, as well as the parameters of the gaussians are estimated from data. In the study by Marzban et al,[15] it was shown that the two gaussians represent the "signal" and the "background" portions of the distribution, respectively. Those, then, allow one to use the upper percentiles of the former to quantify the maximum flow velocity. The left panel of Figure 1 shows an instance of the signal gaussian (blue), the background gaussian (red), the 90th, 95th, and 99th percentiles of the former (blue vertical lines), and the maximum flow velocity according to the MGM (black vertical line). The data (black circles) are from a single patient and at a given time.

The generalization from a GMM to a skewed GMM is relatively straightforward in that each gaussian in the mixture model is replaced by a skewed gaussian (Appendix 1). This process introduces only 1 more parameter: namely, the shape parameter, which affects the skew of the gaussian. The difference in the models may be visualized by the example in Figure 1. The middle panel of Figure 1 shows the various components. The data are for the same patient and same time as in the left panel. Note the evident skew in the signal and the background components. The skewed GMM algorithm is implemented by an R package called mixsmsn,[18,19] available online from the Comprehensive R Archive Network.[26]

In the KMM the distribution of flow velocity is assumed to be a linear combination of two distributions, but unlike the GMM and skewed GMM, each distribution is estimated by using kernel methods (Appendix 1). The kernel method followed here uses the expectation maximization algorithm for fitting multivariate nonparametric mixtures with com-

Figure 1. Distribution (more accurately, probability density function) for the signal (blue) and background (red), as determined by the GMM, skewed GMM, and KMM. The blue vertical lines denote the 90th, 95th, and 99th percentiles of the signal flow velocity (FV), and the black vertical line marks the maximum flow velocity according to the MGM.

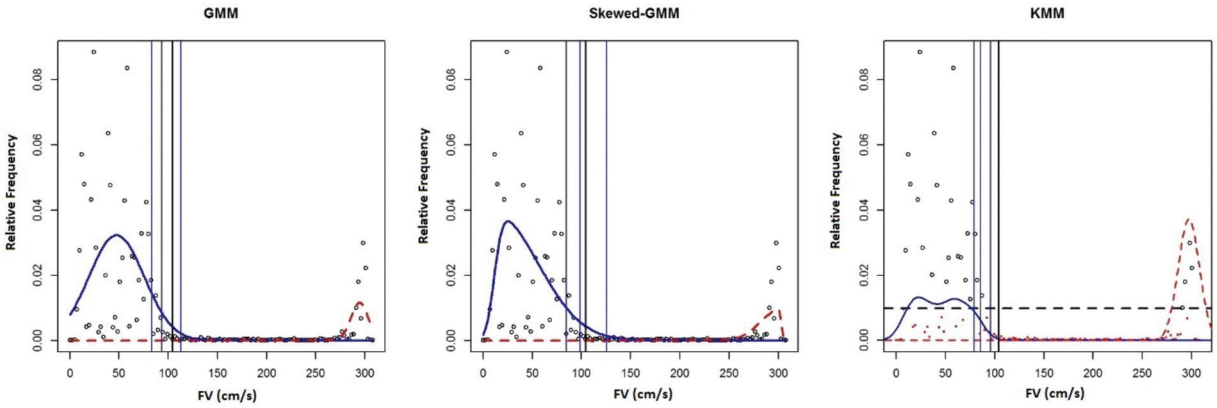

pletely unspecified component densities,[19–21] except for a conditional independence assumption described by Benaglia et al.[23] The kernel is the standard normal density function. The algorithm is implemented in an R package called mixtools,[23–25] available from the Comprehensive R Archive Network.

The right panel of Figure 1 shows the resulting distributions. Note that the distributions are now more irregular, capturing the nonsmooth structure evident in the data (black circles). Appendix 1 explains the manner in which the smoothing parameter is set.

Results

As mentioned previously, the comparison of the various models examined here has both a quantitative component as well as a qualitative (visual) component. The "visual envelope" suggested by a visual inspection of the spectrogram is contingent on the color coding underlying the spectrogram.[15] Given that the true envelope is unknown, the visual envelope can be considered yet another estimate, and despite its qualitative nature, the visual envelope provides a background against which all of the other estimates may be viewed. Then a model is declared as useful if the estimated envelope is visually consistent with the underlying spectrogram. Another qualitative criterion that is applied to assess the usefulness of the models is the consistency (in time) with which the estimated envelope agrees with the spectrogram. For example, if the estimated envelope has stand-alone peaks or troughs interspersed throughout the time window of the spectrogram, then the model is deemed not useful, or at least worthy of further research.

The spectrograms and envelopes for the GMM, skewed GMM, and KMM for a single patient are shown in Figure 2. (The corresponding patient is different than that for Figure 1; the reason why different patients were used for different figures was to illustrate different characteristics of the various methods.) For this particular patient, it can be seen that the 3 percentiles according to the GMM (left panel) are close to one another and generally lower than the envelope based on a visual inspection of the spectrogram. By contrast, the 3 envelopes based on the skewed GMM are more widely spread and more consistent with the spectrogram (middle panel). The KMM's envelopes (right panel) have similarities with and differences from the GMM and the skewed GMM. For example, they underestimate the flow velocities (as for the GMM) but are relatively spread out (as for the skewed GMM). These patterns are generally true across many patients. Also, note the complete disagreement between the spectrogram and the 99th percentile envelope at the time just before 20 (0.025 seconds). This type of failure is characteristic of the KMM, and it is discussed further in the "Discussion" section.

Another noteworthy feature in Figure 2 is that GMM envelopes are generally lower than the aforementioned visual envelope but only for relatively high flow velocity values. This characteristic is evident in the manner in which the envelope is "below" the highest flow velocities in the spectrogram but only when flow velocities are highest. The skewed GMM solves that problem. It is a consequence of the fact that low flow velocity values are consistent with gaussian distributions, but for high flow velocity values, the distributions are more skewed.

The time dependence of the shape of the distributions can be confirmed by examining the values of the shape parameter as a function of time. Figure 3 shows the time series for the estimated shape parameter for the signal (left panel) and background (right panel). It can be seen that for most of the time series, the values of the shape parameters are approximately 0; ie, the distribution of the flow

Figure 2. Spectrograms (color background) and 90th, 95th, and 99th percentile envelopes (black) according to the GMM, skewed GMM, and KMM for a single patient.

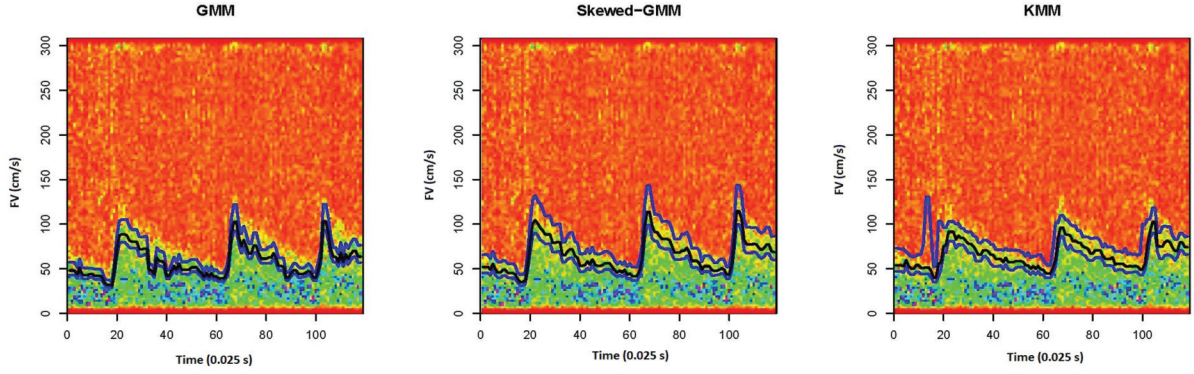

velocity is gaussian. However, there are also times at which the shape parameters are unambiguously not 0. The times correspond to regions of the spectrogram where flow velocities take their highest values. Note that when the shape parameters are not 0, they are positive for the signal and negative for the background components of the mixture model. (A positive shape parameter value corresponds to a right-skewed gaussian.)

As seen for the patient corresponding to Figure 3, the shape parameter is 0 for most times. However, that is not true for all patients; Figure 4 shows a histogram of the estimated shape parameters for a different patient. Note the bimodal structure; the left mode corresponds to the background, whereas the right mode is for the signal. It follows that the background component is consistently left skewed (with shape parameters in the –30 range), and the signal component is either gaussian (corresponding to the peak at 0) or right skewed (with shape parameters between 0 and +30).

Figures such as those in Figure 2 are useful for assessing the quality of an envelope. However, they were impractical for addressing the quality of the envelopes for all 59 patients. Consequently, although we performed this visual inspection for all patients, the results are not shown here. Instead, one can rely on the goodness-of-fit measure itself, not the envelope, to compare the various models developed here. There are numerous quantities that can be used for measuring how well a model fits the data. One commonly used measure is the log likelihood (LL). Given that the fitting of the model is performed at each time, there exists a time series for the LL as well (not shown). For the purpose of comparing two models, however, it is sufficient to examine the difference between LL values of the models. Therefore, for each patient, the histogram, or box plot, of this quantity provides a visual assessment of the relative performance of the models. If two models are comparable in terms of goodness of fit, then the box plot will be centered on the number 0. Otherwise, the models can be deemed different. It is important to point out that the LL assesses the goodness of fit of the mixture model used for estimating an envelope, not the quality of the estimated envelope itself; an assessment of the quality of the estimated envelope would require the true envelope, which is unobservable. The reason for evaluating the models here in terms of LL and the quality of envelopes is to allow for the possibility that one model may outperform another model in terms of LL but not necessarily in terms of the quality of the envelope; or vice versa.

Figure 5 shows a box plot of LL (skewed GMM) minus LL (GMM) for all 59 patients. (Recall that the variability of these box plots is due to time.) The fact that most box plots are "above" the vertical line at 0 implies that for most of the patients, the LL of the skewed GMM is generally higher than that of the GMM across time. In other words, on average (across time), the skewed GMM is a better model of the data than the GMM. A few exceptions are patients 19, 41, and 53, for whom the box plots are mostly centered around 0, and as such, there is no significant difference between the GMM and skewed GMM. To assess whether the skewed GMM has higher LL values than the GMM, a 1-sided *t* test was performed. All of the *P* values were found to be less than .05, except those for patients 19,

Figure 3. Time series of the estimated shape parameter for the signal and background distribution for a single patient.

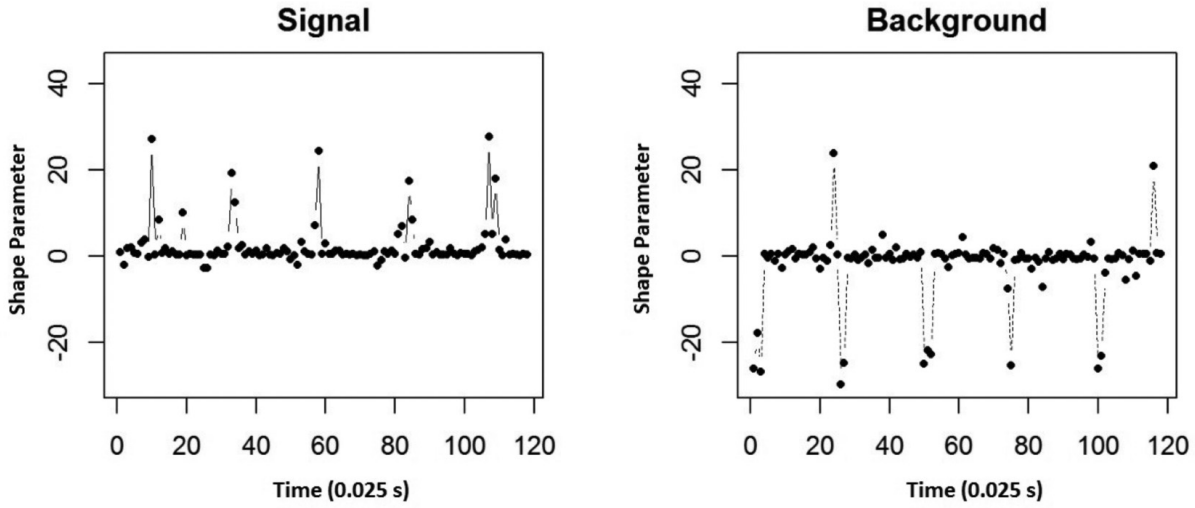

41, and 53 (shaded box plots in Figure 5). (The assumption of normality for the t test is confirmed by a visual inspection of the normal quantile-quantile plot.[26])

The comparison between the KMM and GMM can be done in the same fashion. A box plot for the difference between their LL is shown in Figure 6. Given the systematic presence of the boxes to the right of the vertical line at LL = 0, it is evident that the KMM has a generally higher LL than the GMM. (All of the *P* values were <.05.) It is important to emphasize that this conclusion is not unexpected. The KMM allows for more flexible distributions and, consequently, leads to higher LL values than the GMM. The more important question is whether the KMM also leads to better envelopes than the GMM, and the answer to that is no. The KMM envelopes for all 59 patients were compared with the corresponding envelopes from the GMM (both skewed and not) in terms of a visual comparison of the envelopes, a scatterplot of two envelopes (eg, Figure 5 of Marzban et al[15]), and a comparison of scalar summary measures (eg, Figure 6 of Marzban et al[15]). Although the details of the comparison are not shown here, the results suggest that the KMM envelopes are no better than those of the GMM and are often comparable with those produced by the skewed GMM. The KMM is also not as consistent (across time) as the GMM or skewed GMM, as seen by the "peak" in the right panel of Figure 2. Therefore, although there is sufficient evidence to advocate the use of the skewed GMM, that is not the case for the KMM. Further discussion of the KMM is provided in the next section.

Figure 4. Histogram of the estimated shape parameter for a single patient. The left (right) mode in this histogram corresponds to the background (signal) component of the mixture model.

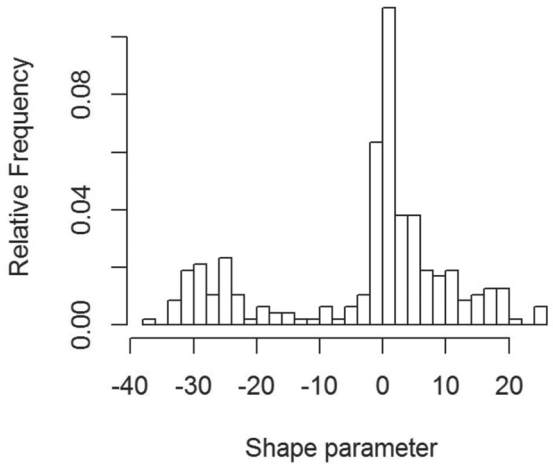

Discussion

In an earlier study, it was shown that the GMM has numerous advantages over the MGM for estimating the envelope. Here, it has been shown that a skewed GMM has the same qualitative advantages as the GMM (eg, allowing for a percentile-based notion of maximum flow velocity) but also outperforms it in terms of both the quality of the envelope (assessed through a visual comparison of the spectrogram and the envelopes) and a quantitative comparison (in terms of the goodness of fit). One may wonder whether the higher goodness-of-fit values for the skewed GMM (compared with the GMM) may be a consequence of overfitting. This concern is unwarranted because the skewed GMM has only 2 additional parameters: the shape parameters for the signal and background components. As such, the total number of independent parameters to be estimated is only 7 (mean, standard deviation, and shape parameter for each of the signal and background components, plus 1 for the mixing coefficient). There exists a large body of theoretical work on the relationship between the number of parameters and the number of observations; eg, see the Vapnik-Chernovenkis dimension in the work by Hastie et al[27] (page 210). On the practical side, one guideline that is often followed is from the classic book by Ryan[28] (page 20), who recommends at least 10 times as many observations as the number of parameters. Here, the data used for estimating the 7 parameters number in the hundreds; therefore, overfitting is not a concern.

An issue that arises in the mixture model approach to envelope estimation is the identifiability of the two components with signal and background. More specifically, a mixture model with two components will identify two components in the distribution of flow velocity values; however, an additional criterion must be introduced to decide which of the components should be identified as the signal. This factor is important because it is the percentiles of the signal component that lead to the envelope. Visually, the signal component can be identified as the one on the left, but quantifying what is meant by "on the left" is nontrivial. The approach adopted here for identifying the signal and background components of the mixture model is outlined in Appendix 2. That proposal is somewhat ad hoc, so it would be desirable to develop a more statistically sound method for identifying the signal component of mixture models.

As described above, the smoothing parameter of the KMM determines the shape of the kernel estimate of the probability density function. Here, an algorithmic criterion is used to set the values of that parameter, but the criterion optimizes the goodness of fit (ie, LL). As such,

the criterion does not automatically lead to improved envelopes. An instance of this is shown in Figure 2, where the KMM envelope displays an abrupt peak at approximately time 15 (0.025 sec), which is clearly inconsistent with the observed spectrogram. This is one of the main reasons why the KMM is not wholeheartedly advocated here.

Another reason is more computational: the procedure for estimating kernels is rather computer intensive and, therefore, slow. In fact, for this study, the KMM was developed on data for which the lowest 5% of flow velocity values

Figure 5. Box plot of LL for the skewed GMM minus LL for the GMM for all 59 patients. The shaded boxes indicate $P > .05$.

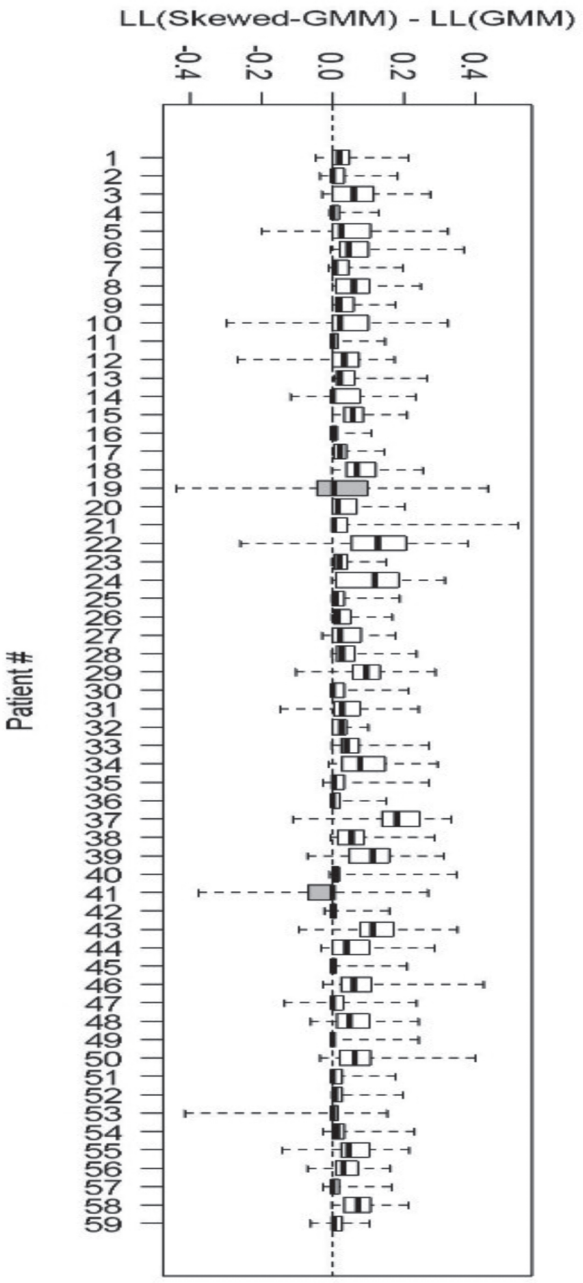

Figure 6. Same as Figure 5 but for LL for the KMM minus LL for the GMM.

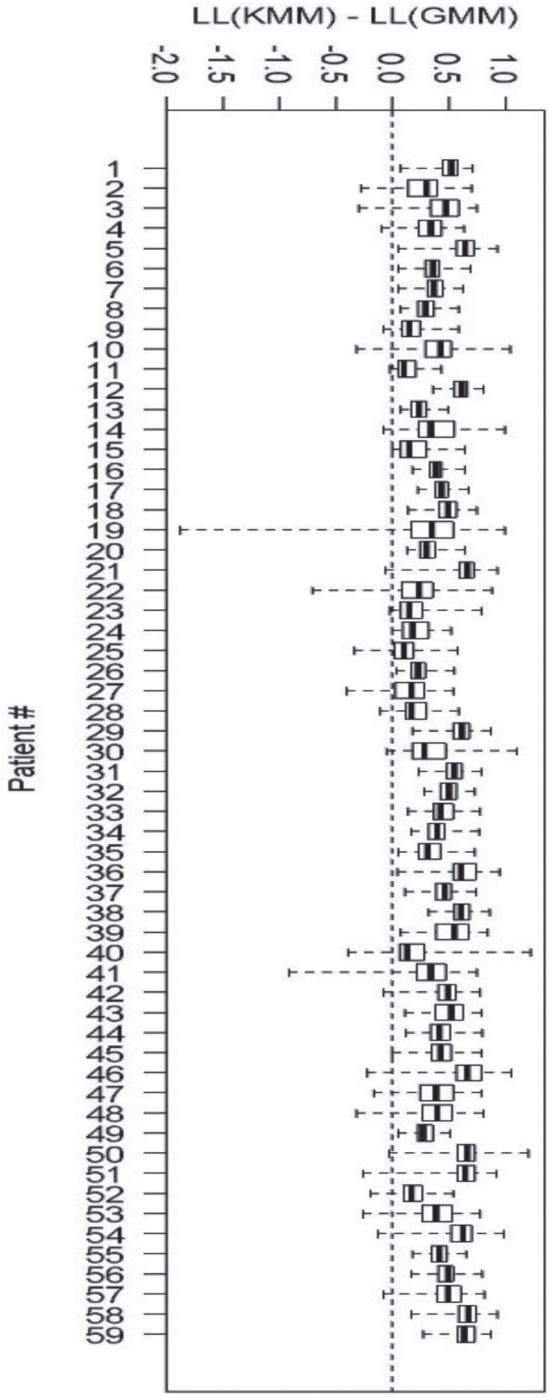

were trimmed (deleted). This process substantially increases computational speed but introduces an ad hoc parameter (ie, the 5%). In short, more work is required before one can benefit from the flexibility the of KMM as a useful envelope estimation method.

Appendix 1: Mixture Models

In a mixture model, one generally assumes that the data are generated from an underlying probability density function $[f(x)]$ that can be written as a linear combination of more well-known probability density functions. For example, for the case of two components (as considered in this study), the probability density function is written as

$$f(x) = \sum \lambda_k f_k(x), \quad (1)$$

where λ_k are called mixing coefficients and must sum to 1. For GMMs, $f_k(x)$ is given by

$$\phi\left(\frac{x-\mu}{\sigma}\right),$$

where μ and σ^2 are mean and variance parameters, and $\phi(.)$ is the standard gaussian probability density function. Consequently such a mixture model has 5 independent parameters: the mean and variance parameters for each of the two gaussian probability density functions and one for the mixing weights, which must be estimated from data.

In a skewed GMM, the component probability density functions are assumed to be given by

$$f_k(x) = 2\phi\left(\frac{x-\mu}{\sigma}\right)\Phi(\alpha x), \quad (2)$$

where $\Phi(x)$ is the gaussian cumulative probability distribution, and α is called the shape (or skew) parameter.[19] Large positive values of the shape parameter lead to a probability density function with a more pronounced tail for large values of x (ie, a right-skewed distribution), whereas large negative values of α have the same effect but for large negative values of x (ie, a left-skewed distribution). In this application, the skewed GMM, therefore, has 2 more parameters than the 5 parameters in the GMM. Estimated values of the shape parameter are shown in Figure 3, and the skewed probability density functions are shown in the middle panel of Figure 1. The presence of the shape parameter allows one to fit a more general class of data than can be fit with GMMs.

Even more general data can be fit with kernel methods, in which the component probability density functions in the mixture model are themselves assumed to be a sum of the form

$$f_k(x) = \frac{1}{nh}\sum K\left(\frac{x-x_i}{h}\right), \quad (3)$$

where x_i are the observed data; n is the sample size; and the function $K(.)$, which is called the kernel, can be any function that integrates to 1; commonly, and in the present application, $K(.)$ is taken to be the standard gaussian probability density function.[27] The additional parameter h, called the bandwidth, effectively determines the smoothness of the probability density function; a small value of h will lead to a highly "spiky," nonsmooth probability density function, whereas larger values of h will lead to smoother, more gaussian-looking, probability density functions. For this study, its value was determined as follows: The default choice of the bandwidth h minimizes a quantity called asymptotic mean integrated squared error[29] and is given by

$$h = 0.9n^{-1/5}\min\left\{\mathrm{SD}, \frac{\mathrm{IQR}}{1.34}\right\}, \quad (4)$$

where SD and IQR are the standard deviation and interquartile range of all n observations, respectively. It has been shown that this estimate is somewhat smaller than the true optimal value,[24] so the value chosen here is taken to be larger than the default value, which in turn leads to smoother fits. As a result, at each time, and for each patient, $h = 10$ is used, unless equation 4 gives a larger value, in which case the larger value is selected. The fits are visually inspected to ensure that the chosen value of the bandwidth does not lead to overfitting.

Appendix 2: Identifying the Signal Distribution

As described above, the mixture model identifies two components in the distribution of flow velocity values. However, the task of identifying which of the components is the signal (and which is the background), is nontrivial. Here, the following 3 criteria are used to derive an algorithm for the identification task: The signal component is more likely the one with (1) a lower mean, (2) a higher peak, and (3) a larger standard deviation. The means, peaks, and standard deviations of the two components are denoted as (μ_1, μ_2), (p_1, p_2), and (s_1, s_2), respectively. The ratios are defined as $a = \mu_1/\mu_2$; $b = p_1/p_2$; and $c = s_1/s_2$. As such, the signal component is more likely to be the component labeled 1 when a is small and b and c are large. Therefore, if $x = (1/a) \times b \times c$ is larger than 1, then the component labeled 1 should be identified as the signal. Similarly, $x < 1$ indicates that the component labeled 1 should be identified as the background.

References

1. Aaslid R, Nornes H. Musical murmurs in human cerebral arteries after subarachnoid hemorrhage. *J Neurosurg* 1984; 60:32–36.
2. Aaslid R, Huber P, Nornes H. Evaluation of cerebrovascular spasm with transcranial Doppler ultrasound. *J Neurosurg* 1984; 60:37–41.
3. Aaslid R, Huber P, Nornes H. A transcranial Doppler method in the evaluation of cerebrovascular spasm. *Neuroradiology* 1986; 28:11–16.
4. Aaslid R. Transcranial Doppler assessment of cerebral vasospasm. *Eur J Ultrasound* 2002; 16:3–10.
5. Armonda RA, Bell RS, Vo AH, et al. Wartime traumatic cerebral vasospasm: recent review of combat casualties. *Neurosurgery* 2006; 59:1215–1225.
6. Babikian VL, Wechsler LR, Toole JF (eds). *Transcranial Doppler Ultrasonography*. Boston, MA: Butterworth-Heinemann; 1999.
7. Sloan MA, Wozniak MA, Macko RF. Monitoring of vasospasm after subarachnoid hemorrhage. In: Babikian VL, Wechsler LR, Toole JF (eds). Transcranial Doppler Ultrasonography. Boston, MA: Butterworth-Heinemann; 1999:109–128.
8. Tegeler CH, Rotanakorn D. Physics and principles. In: Babikian VL, Wechsler LR, Toole JF (eds). *Transcranial Doppler Ultrasonography*. Boston, MA: Butterworth-Heinemann; 1999:3–11.
9. Tsivgoulis G, Alexandrov AV, Sloan MA. Advances in transcranial Doppler ultrasonography. *Curr Neurol Neurosci Rep* 2009; 9:46–54.
10. Fernando KL, Mathews VJ, Clark EB. A mathematical basis for the application of the modified geometric method for maximum frequency estimation. *IEEE Trans Biomed Eng* 2004; 51:2085–2088.
11. Güler I, Hardalac F, Kaymaz M. Comparison of FFT and adaptive ARMA methods in transcranial Doppler signals recorded from the cerebral vessels. *Comput Biol Med* 2002; 32:445–453.
12. Moraes R, Adin N, Evans DH. The performance of three maximum frequency envelope detection algorithms for Doppler signals. *J Vasc Invest* 1995; 1:126–134.
13. Östlund N, Yu J, Karlsson JS. Improved maximum frequency estimation with application to instantaneous mean frequency estimation of surface electromyography. *IEEE Trans Biomed Eng* 2004; 51:1541–1546.
14. Rubin JM, Bude RO, Fowlkes JB, Spratt RS, Carson PL, Adler R. Normalizing fractional moving blood volume estimates with power Doppler US: defining a stable intravascular point with the cumulative power distribution function. *Radiology* 1997; 205:757–765.
15. Marzban C, Illian PR, Morison D, Mourad PD. A double-gaussian, percentile-based method for estimating maximum blood flow velocity. *J Ultrasound Med* 2013; 32:1913–1920.
16. Pereira M, Freire M. *Biomedical Diagnostics and Clinical Technologies: Applying High-Performance Cluster and Grid Computing*. Hershey, PA: IGI Global; 2010:263.
17. McLachlan GJ, Peel D. *Finite Mixture Models*. Hoboken, NJ: John Wiley & Sons; 2000.
18. Basso R, Lachos V, Cabral C, Ghosh P. Robust mixture modeling based on scale mixtures of skew-normal distributions. *Comput Stat Data Analysis* 2010; 54:2926–2941.
19. Prates M, Cabral CH, Lachos V. mixsmsn: fitting finite mixture of scale mixture of skew-normal distributions. *J Stat Software* 2013; 54:1–20.
20. Dinov ID. Expectation maximization and mixture modeling tutorial. California Digital Library, Statistics Online Computational Resource website. http://repositories.cdlib.org/socr/EM_MM. December 9, 2008.
21. Dempster AP, Laird NM, Rubin DB. Maximum likelihood from incomplete data via the EM algorithm. *J R Stat Soc Ser B* 1977; 39:1–38.
22. Benaglia T, Chauveau D, Hunter DR, Young D. mixtools: an R package for analyzing finite mixture models. *J Stat Software* 2009; 32:1–29.
23. Benaglia T, Chauveau D, Hunter DR. An EM-like algorithm for semi- and non-parametric estimation in multivariate mixtures. *J Comput Graph Stat* 2009; 18:505–526.
24. Benaglia T, Chauveau D, Hunter DR. Bandwidth selection in an EM-like algorithm for nonparametric multivariate mixtures. In: *Nonparametric Statistics and Mixture Models: A Festschrift in Honor of Thomas P. Hettmansperger*. Singapore: World Scientific Publishing Co; 2011:15–27.
25. Marzban C, Illian PR, Morison D, et al. A method for estimating zero-flow pressure and intracranial pressure. *J Neurosurg Anesthesiol* 2013; 25:25–32.
26. R Development Core Team. R: a language and environment for statistical computing. R Project website; 2011. http://www.R-project.org.
27. Hastie T, Tibshirani R, Friedman J. *The Elements of Statistical Learning*. New York, NY: Springer; 2001.
28. Ryan T. *Modern Regression Methods*. Hoboken, NJ: John Wiley & Sons; 2009.
29. Silverman BW. *Density Estimation*. London, England: Chapman & Hall; 1986.