Matched filtering and multiple hypothesis tracking applied to C-fiber action potentials recorded in human nerves

Björn Hansson[a c], Clemens Forster[b], and Erik Torebjörk[c]

[a]Signals and Systems, Uppsala University, Sweden

[b]Institute of Physiology and Experimental Pathophysiology, University of Erlangen, Germany

[c]Department of Clinical Neurophysiology, Uppsala University Hospital, Sweden

ABSTRACT

We describe an application of multiple target tracking (MTT) to microneurography, with the purpose of estimating conduction velocity changes and recovery constants of human nerve C-fibers. In this paper, the focus is on the detection and the tracking of the nerve action potentials (APs). The subsequent parameter estimation is described only briefly. Results from an application of the tracking system on real data recorded in human subjects are presented.

Action potentials from C-fibers were recorded with a thin needle electrode inserted into the peroneal nerve of awake human subjects. The APs were detected by a matched filter constituting a maximum likelihood constant false alarm rate (ML-CFAR) detector. By utilizing the multiple hypothesis tracking (MHT) method, the detected APs (targets) in each trace were associated to individual nerve fibers (tracks) by their typical conduction latencies in response to electrical stimulation in the skin. The measurements were one-dimensional (range only), and the APs were spaced in time with intersecting, piecewise continuous, trajectories. The amplitude of the APs was varying slowly over time for each C-fiber and was in general different for different fibers. It was therefore incorporated into the tracking algorithm to improve its performance.

Keywords: Detection, MHT, target tracking, microneurography, C-fiber potentials

1. INTRODUCTION

In signal processing, problems arising from applications may generate challenges that the tools traditionally used within that application area are inadequate. We will here describe such a problem, arising from the need to study the stimulus-response characteristics of peripheral unmyelinated (C-) fibers in human skin nerves. Key tools for the solution were to be found in a seemingly unrelated area, namely radar tracking of multiple targets.

Neuronal activity is evoked by electrical stimuli delivered through electrodes positioned in the skin area innervated by the fibers of interest. A thin needle electrode is inserted transcutaneously into the nerve to record the action potentials (APs) of the C-axons.[1,2] The APs can be detected as extracellular spikes in this recording. However, the signal-to-noise-ratio (SNR) is rather poor.[3] Sometimes the amplitude of the APs under observation is in the same range as the peaks of the noise and special methods are required for the detection.

The sensory stimuli that are applied during the experiment usually excite several nerve endings. If their axons are close to each other in the nerve, it is likely that their spikes are recorded together by the intraneural microelectrode. Due to slight differences in conduction velocity of individual C-fibers, the action potentials will be spaced in time in the recorded signal, thereby allowing individual identification of separate C-fibers with characteristic response latencies.

In 1974, Hallin and Torebjörk introduced a method that reveals the excitation of a C-fiber using the slight decrease of its conduction velocity that occurs after the conduction of an AP. They called it the *marking phenomenon.*[4,5] The principle is that electrical stimuli are applied repetitively, at low frequency (0.25 Hz), into the

Further author information -
B.H. (correspondence): Email: Bjorn.Hansson@signal.uu.se; WWW: http://www.signal.uu.se/Staff/bh/bh.html
C.F.: Email: Forster@physiologie1.uni-erlangen.de
E.T.: Email: Erik.Torebjork@neurofys.uu.se

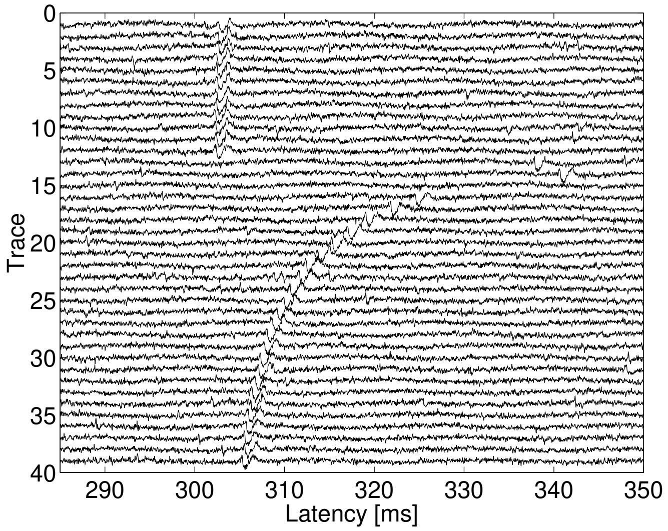

Figure 1. The APs of one C-fiber are shown. The responses are excited by electrical stimuli delivered at 0.25 Hz into the skin innervation territory of the fiber. Successive responses are displayed in traces from top to bottom. At trace 13, the unit is activated by a mechanical stimulus leading to a decrease in the conduction velocity, i.e. an increase in latency. Following this, the conduction velocity recovers gradually as indicated by the APs returning to the latency prior to the activation.

innervation territory (receptive field) of the C-fiber under study. For each electrical impulse, one single spike is evoked and appears in the recording after a certain latency (Fig. 1, at 303 ms). For documenting the response characteristics of the C-fiber, physiological test stimuli are applied into the receptive field of the fiber. If such stimuli generate additional action potentials, the conduction velocity will decrease. In that case, the AP excited by the electrical stimuli will show a noticable increase in latency (Fig. 1, traces 13 to 40). This change in latency is used as a "marker" that the C-unit has responded to the applied physiological provocation. In addition, the quantity of the latency increase provides an estimate of the number of APs that were generated by the provocation.[6] To enhance the efficiency of these experiments, a computer supported recording system is used.[7]

In the human skin nerves, different types of C-fibers exist.[8] Recently, it has become evident that the their latency increase, due to a particular number of impulses, and the time course of their recovery, are different in different classes of C-fibers (to be submitted). This finding is very interesting as it may contribute to new insights into differential properties of membranes in different C-fiber classes in humans.

Today, the analysis of the recorded time courses is carried out manually; a work that is very time consuming. With the new findings in mind, it is believed that this type of analysis will become more important and therefore an automatic system that can be used to estimate the latency shift and the recovery constant quantatively has been developed.

2. ALGORITHM OVERVIEW

To estimate the characteristics of the latency time course, two major problems have to be solved: the detection of action potentials in noisy recodings and the discrimination of APs originating from different sensory units.

The discrimination problem could also be viewed as a target tracking problem, in which a particular action potential is tracked in the responses to the repetitive electrical stimuli. By exploiting the marking phenomenon, a reliable tracking algorithm is possible to derive.

Knowing the time course of the latency corresponding to a particular C-fiber unit, it is then straight-forward to fit a parametric model to the data.

By adding this together, we have developed an application that analyzes these experiments in three steps, see Fig. 2:

Figure 2. Sample results from the three steps of the algorithm. To the left, the detection is shown. The middle figure displays the five resulting tracks. To the right, parametric models have been fitted to these tracks.

Detection – Prior to any further processing, the action potentials must be detected. In our implementation, this is done by utilizing a matched filtering (MF) technique.

Tracking – Once the action potentials are detected, the discrimination is carried out. Although the discrimination may be easy for an experienced analyst, it is quite a difficult problem to solve automatically using a computer. By viewing the assignment as a target tracking problem, we solve it using the multiple hypothesis tracking (MHT) method.

Parameter estimation – Following the tracking/discrimination step, the latency shift and the recovery constant are estimated by fitting the latency model to the data. Currently, we use the simplex algorithm in combination with the least squares method.

The paper is organized as a description of the three main steps of the algorithm. Section 3 describes how the detection is done and how the optimal detector is found. Section 4 deals with the tracking of the APs using Kalman filtering and the MHT method, and describes how the amplitude of the action potentials is incorporated into the tracking logic. Section 5 outlines very briefly how the latency recovery is modeled and how its parameters are calculated. Section 6 illustrates the performance on actual recordings from awake human subjects in our research lab.

3. TARGET DETECTION

Detecting signals hidden in high levels of noise is a delicate task. If the signal is a member of a set of signals with known shapes, and if the color of the noise is known, then *matched filtering* constitues a standard signal processing technique for optimally enhancing and detecting the signal in noise.[9] In this section, a short derivation of the matched filter will be presented. When the signal is known and the noise is wide-sense stationary, the MF is the optimal detector, but it has turned out that its performance in this application is good even though these assumptions are not fulfilled.

3.1. Matched filtering

The recorded discrete-time data $\mathbf{w}(t)$ is hypothesized to take on one of the forms

$$H_0 : \mathbf{w}(t) = \mathbf{n}(t) \tag{1}$$
$$H_1 : \mathbf{w}(t) = \mathbf{s}(t) + \mathbf{n}(t) \tag{2}$$

where $\mathbf{w}(t)$, $\mathbf{s}(t)$, and $\mathbf{n}(t)$ are all p-dimensional vectors. They are time-reversed snapshots of the data, the signal, and the noise respectively, and are defined in the form

$$\mathbf{w}(t) = \begin{pmatrix} w(t) & w(t-1) & \dots & w(t-p+1) \end{pmatrix}^T . \tag{3}$$

It is now assumed that the signal $\mathbf{s}(t)$ contains well separated action potentials. As the APs are similar in shape but differ in amplitude, a p-dimensional *template* $\mathbf{s}$ and an *amplitude* γ are introduced to describe an AP present in the signal at a particular time t_d

$$\mathbf{s}(t_d) = \gamma\mathbf{s}. \tag{4}$$

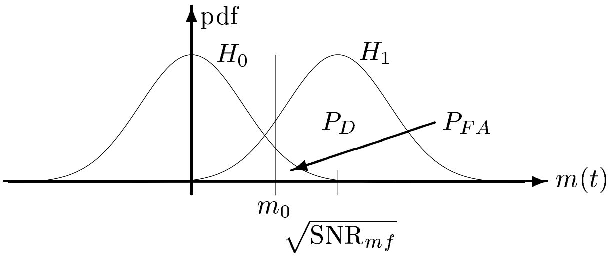

Figure 3. The probability density functions (pdf:s) of the two hypotheses. Under hypothesis H_0, the data contains zero mean noise only and its pdf is centered around zero. Under hypothesis H_1, the expectation of the MF peak output is equal to the square root of its signal-to-noise-ratio SNR_{mf}. The false alarm probability P_{FA} and the detection probability P_D are equal to the area under each pdf above the threshold m_0 as shown in the figure.

The objective, then, is to determine whether there are action potentials $\mathbf{s}$ present or not in the measurement $\mathbf{w}(t)$. For a given matched filter impulse response

$$\mathbf{h} = (h(0) \quad h(1) \quad \ldots \quad h(p-1))^T \tag{5}$$

the decision that the prescribed signal template $\mathbf{s}$ is present is made whenever the filter output

$$m(t) = \sum_{i=0}^{p-1} h(i)w(t-i) = \mathbf{h}^T\mathbf{w}(t) \tag{6}$$

contains a peak that exceeds a given threshold level m_0. Assuming the noise to be white, the optimal impulse response $\mathbf{h}_o$ of the matched filter and the output signal-to-noise-ratio SNR_{mf} can be calculated in terms of the known signal vector $\mathbf{s}$ and the noise variance σ^2 according to[9,10]

$$\mathbf{h}_o = \frac{\mathbf{s}}{\sqrt{\sigma^2 \mathbf{s}^T\mathbf{s}}} \tag{7}$$

$$\text{SNR}_{mf} = \frac{\gamma^2}{\sigma^2}\mathbf{s}^T\mathbf{s} \tag{8}$$

where ^{T} denotes the vector transpose.

3.2. Detection performance

Assuming the noise $\mathbf{n}(t)$ is Gaussian, the MF output $m(t)$ is a stochastic variable with a Gaussian distribution and unit variance. The false alarm and detection probabilities, P_{FA} and P_D, can thus be calculated by means of the decision threshold m_0 using the normal distribution density function $\Phi(x)$, see Fig. 3

$$P_{FA} = 1 - \Phi(m_0) \tag{9}$$

$$P_D = 1 - \Phi\left(m_0 - \sqrt{\text{SNR}_{mf}}\right). \tag{10}$$

The false alarm probability P_{FA} is depending on the detection threshold m_0 only, thus resulting in a constant false alarm rate (CFAR) if the threshold is held constant. This is a desirable property in most applications.

3.3. Noise variance estimation

As the noise variance is unknown, and may change during the registration, it is estimated from the recorded data at each trace. The recordings may, however, contain hum from the power lines. In order to simplify the variance estimation and to avoid a biased estimate, the hum is removed using a notch filter. Assuming the noise variance to be constant during the trace, the maximum likelihood (ML) estimate of the noise variance $\hat{\sigma}^2$ is then calculated using the notched data samples $w_n(t)$ in one trace of length N

$$\hat{\sigma}^2 = \frac{1}{N}\sum_{t=0}^{N-1} w_n^2(t). \tag{11}$$

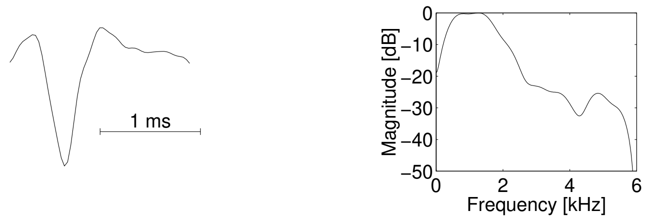

Figure 4. The template of the C-fiber action potentials to be detected, i.e. the time-reversed impulse response of the matched filter (left), and the corresponding Bode magnitude plot (right).

This is still a biased estimate, but the bias decreases as the width of the notch decreases. As most practical cases involve a narrow notch, the influence of the bias can be neglected and will not be considered further. Of more importance, on the other hand, is the possible influence of the action potentials present in the recorded data used for estimating the variance. Experience has, however, shown that they provide only a small contribution to the variance estimate.

Using the ML estimate of the noise variance, the impulse response and the SNR of the matched filter, Eqs. 7 and 8, are given by*

$$\mathbf{h} = \frac{\mathbf{s}}{\sqrt{\hat{\sigma}^2 \mathbf{s}^T \mathbf{s}}} \tag{12}$$

$$\widehat{\mathrm{SNR}}_{mf} = \frac{\gamma^2}{\hat{\sigma}^2} \mathbf{s}^T \mathbf{s}. \tag{13}$$

3.4. Tuning

The basic assumption in deriving the matched filter, is that the signal to be detected is known in advance which is not the case. It has turned out, though, that all C-fiber action potentials are similar in shape and differ mainly by a scalar factor. It is thus possible to use any, already recorded, action potential to tune the matched filter.

By aligning and averaging several traces of good quality, the influence of the noise and the background activity is reduced. To reduce it further in the region where no AP energy is expected, the resulting data is low pass filtered. In Fig. 4, the resulting signal template and, thus, the time-reversed impulse response of the matched filter, is shown together with the corresponding Bode plot.

3.5. Action potential detection

The spectral content of the spikes is concentrated to the frequency range between 500 Hz and 1500 Hz. The matched filter will constitute an optimized bandpass filter with passband in this frequency range. The filter output is smooth and the time instants where there are spikes present can be found using a simple peak finding algorithm that reports all peaks above the given threshold m_0.

4. TARGET TRACKING

Once the action potentials are detected, the APs corresponding to a particular sensory C-unit should be grouped together. This grouping could be done in various ways. If the shape of the APs were different for different units, a bank of matched filters could be used for the grouping, where each matched filter is individually optimized to detect APs originating from one of the sensory units only.

As mentioned in Sect. 3, the action potentials are, however, similar in shape so another method must be used. Due to the repetitive stimulation and due to different conduction velocity in different units, the latency contains information about which APs originate from the same unit.

*Introducing the ML estimate of the variance implies, formally, that a Kelly test[10] should be used instead of the generalized likelihood ratio test (GLRT) used above. Moreover, the MF output has then a t-distribution instead of a normal distribution. However, the number of data points used in the estimations is so large, that these effects have a negligible effect on the result.

The experienced eye easily exploits this information in subsequent traces and forms tracks of APs that belong to each other. We try to mimic this effect by solving the association problem using the well known multiple hypothesis tracking method.[11]

Although the tracks can be found using the latency information only, the tracking performance can be improved if more information can be exploited. Even though the difference in shape is too small to be useful in the discrimination of the APs, their amplitude is in general different for APs originating from different C-fiber units.* For practical reasons, the peak value of the matched filter output is used instead of the amplitude estimate.

The MHT method, outlined below, is a bayesian probabilistic approach to the tracking problem. For each trace in the recorded data, the detected action potentials are collected by the tracking system. At a given time and a given set of detected action potentials, there will be a number of plausible ways to combine the APs into tracks. Instead of choosing only the most probable partitioning after each trace, the MHT method will generate a number of candidate hypotheses to be evaluated later as more data are received. Thus, the probability to choose the correct partitioning of the data into tracks and false alarms is increased.

In order to evaluate the probability of each hypothesis, a model of how the AP latencies change from trace to trace must be known. In Figure 1, it can be seen that the latency is either constant or exponentially decreasing. Both cases can be modeled by an exponential model. A straight-forward way to incorporate this model in the tracking method is by means of a Kalman filter, which is outlined next.

4.1. Kalman Filter

Fundamental in any tracking system is the prediction and filtering logic. The two alternative methods used today are the Kalman filter,[12] and the filter based on interacting multiple models (IMM).[13] The best performance is often provided by the IMM method, but, in the application under consideration, the Kalman filter is used as a starting point due to its simplicity. It is easy to extend this to an IMM based filtering method, if required.

4.1.1. System model

As mentioned above, we assume that the latency follows an exponential trajectory described by

$$y_{lat}(t) = y_0 + Ae^{-\alpha_0(t-t_0)}, \quad t \geq t_0 \tag{14}$$

where y_0 is the steady state latency, A is the latency shift due to stimulation, α_0 is the recovery coefficient, and t_0 is the time of excitation. Using a state space approach, two states can appropriately model the time course of the latency shift.

Assuming further that the SNR at the matched filter output is uncorrelated with the latency time course and slowly varying, a random-walk would be appropiate to describe the MF output SNR.

On this basis, a continous time stochastic model, describing the time course of the latency and the matched filter SNR, can be found. By sampling this model, the discrete time state space model is, after series expansion and truncation, given by

$$\begin{aligned}
\mathbf{x}(k+1) &= \mathbf{F}\mathbf{x}(k) + \mathbf{v}_1(k), \quad k \geq k_0 = \frac{t_0}{T} && (15)\\
\mathbf{F} &= \begin{pmatrix} 1 & T & 0 \\ 0 & a & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad a = e^{-\alpha T} && (16)\\
\mathbf{Q}_1(k) &= E\mathbf{v}_1(k)\mathbf{v}_1^T(k) && (17)\\
&= \begin{pmatrix} \frac{1}{3}T^3 & \frac{1}{2}T^2 & 0 \\ \frac{1}{2}T^2 & T & 0 \\ 0 & 0 & \rho \end{pmatrix} \begin{pmatrix} \sigma_v^2(k) & 0 & 0 \\ 0 & \sigma_v^2(k) & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad 0 < \rho \ll 1 && (18)
\end{aligned}$$

where $\mathbf{x}(k)$ is the 3-dimensional state vector consisting of the latency, its derivative, and the square root of the SNR at the matched filter output. The constant T is the period between traces, i.e. the period of the repetitive electrical

*The reason that APs from different C-fiber units have different amplitude could be explained by the different distances between the recording electrode and the nerve fibers.

stimulation. The transition matrix $\mathbf{F}$ describes the dynamics of the discrete time system, in which α is an a priori value of the recovery time constant of the latency. The three-dimensional vector $\mathbf{v}_1(k)$ represents process noise. It is introduced to capture errors in the approximate dynamic model. The process noise is modeled as zero-mean, white noise with covariance matrix $\mathbf{Q}_1(k)$. The function $\sigma_v^2(k)$ describes the modelling error of the latency due to deviations of the constant α from the true value α_0. In order to describe small changes in the SNR at the matched filter output, a drift rate ρ is introduced.

4.1.2. Measurement model

The measurement model describes how the measurements are collected and is defined as

$$\mathbf{y}(k) = \mathbf{C}\mathbf{x}(k) + \mathbf{v}_2(k) \tag{19}$$

$$\mathbf{C} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix} \tag{20}$$

where $\mathbf{y}(k)$ is a two-dimensional measurement vector containing the latency time and the matched filter peak output value. The measurement matrix $\mathbf{C}$ is thus a mapping of the state vector to the measurement vector. The two-dimensional vector $\mathbf{v}_2(k)$ is the measurement noise modeled as zero-mean, white noise with a constant covariance matrix $\mathbf{Q}_2$, that is tuned by the operator.

4.1.3. Kalman filter algorithm

We assume that the state space model defined in Eqs. (15)-(20) describes the second order statistics of $\mathbf{y}(k)$ correctly. Under this assumption, the Kalman filter recursion[12] is the linear filter which provides the optimal mean-squared error one-step prediction $\hat{\mathbf{x}}(k+1|k)$. The filter is initiated using the first two measurements $\mathbf{y}(k_0)$ and $\mathbf{y}(k_0+1)$.

4.2. Multiple Hypothesis Tracking

Multiple hypothesis tracking (MHT) is recognized as the theoretically best approach to multitarget tracking problems with respect to tracking performance. In applications with heavy clutter† and high traffic densities,‡ its performance is outstanding in comparison with other methods, e.g. nearest neighbor (NN) correlation or joint probabilistic data association (JPDA).[14] We therefore select this method for the current implementation, as good results could be hard to obtain in cases where the tracking situation is complex.

Generally, the MHT method is composed of five steps; four that are carried out for each trace and one that is performed as all traces have been processed:

1. collect the *observations* (APs and false alarms).
2. generate all plausible candidate *hypotheses* by associating each observation to existing tracks and by creating new tracks (containing that observation only).
3. evaluate the probability of all hypotheses, i.e. calculate their *scores*.
4. calculate target predictions for each track.
5. when all traces have been processed, the most probable partioning of the data, i.e. the hypothesis with the highest score, is chosen to be the correct result.

The drawback of MHT is that the number of computations and the memory requirements are extensive. This is of less concern in our application, as the tracking is done off-line. Furthermore, the computation complexity can be held at a reasonable level by limiting the number of hypotheses handled. There are several methods available, and three of them (gating, pruning, and combining) are used in the present application and have worked very well.

†The term clutter refers to targets that are of no interest to the operator and that often degrade the performance of the tracking system. In radar air surveillance, for example, clutter is primarily referring to ground echoes. In our context, clutter would be spurious APs that are uncorrelated with other APs.

‡In our case, the term means that different sensory units have similar latencies.

4.2.1. MHT implementation

The current implementation is essentially based on the MHT method described in Ref. 11. As described in the original work, the computational load and memory requirements of the MHT method are limited by gating, pruning, and combining. *Gating* is used to limit the number of hypotheses generated by not associating a particular observation to a track if it is "far" from the predicted target position of that track. *Pruning* and *combining* are both limiting the number of hypotheses maintained. The former discards unlikely hypotheses and the latter merges similar hypotheses.

With respect to the basic algorithm, following extensions should be pointed out:

- as the maximum likelihood gate given in the original work tended to be too large, i.e. generating many false observation-to-track associations, a small fix gate is used instead.
- only new hypotheses are considered for pruning in the first step of the two-step method. This is to prevent premature deletion of hypotheses not yet updated by the observations in the current trace.
- in the combining process, only tracks are considered where similarity is evaluated using the N-scan criterion. Should more than one hypothesis contain the same set of tracks, all but the most probable one are deleted.

5. PARAMETER ESTIMATION

Throughout the registrations, it has turned out that the spike latencies follow an exponential function during the recovery phase. Practical experience has also shown that fitting an exponential curve to the data works well, even for a rather small number of points.

In accordance with the assumed time course of the latency in Equation (14), the model to be fitted to data is

$$y(k) = y_0 + Ae^{-\alpha(k-k_0)T}, \quad k \geq k_0 \tag{21}$$

where k_0 and k is the initial and the consecutive trace numbers, respectively. The parameter A is the latency shift, α is the recovery constant, and T is the time between each trace.

The model in Equation (21) includes one parameter (α) which enters non-linearly and two parameters (y_0 and A) which enter linearly. The parameters can be estimated by an iterative method in which the non-linear term is estimated using the simplex method. At each step in the simplex algorithm, the two linear terms are estimated using the least squares method, and the error norm is returned to the simplex algorithm.

6. EXPERIMENTAL RESULTS

As mentioned in the introduction, the objective of the application is to detect weak action potentials, combine them into tracks, and estimate the latency shift parameters and the recovery constants.

In this section, we show the performance of the matched filter detector and the MHT tracking method when applied on real recordings from human subjects.

6.1. The matched filter

The recorded data is noisy and the question is how well the matched filter can detect low amplitude action potentials. Previously, a non-linear filter was used.[15] It was called a "noise-cut" filter, as it amplified only the part of a signal that was above a certain magnitude and thus removed the middle, noisy, portion of the signal.

The noise-cut filter works well in detecting the spikes, but has the drawback of not beeing useful for detecting APs with the same magnitude as that of the noise itself. In order to improve upon this performance, the matched filter has to detect APs that are of the same magnitude as the noise.

In Fig. 5, a sample recording is shown. It is a part of the fifth trace corresponding to the tracking result shown in Fig. 7. As can be seen in the figure, the SNR in the matched filter output is improved as compared to the original data. With the detections from several traces (see Fig. 7), it can be concluded that there are at least four action potentials present in the trace (approxiamately at 442 ms, 450 ms, 456 ms, and 466 ms) and that even the low amplitude APs could be detected without decreasing the detection threshold down to the noise level.

From the figure, it is clear that the notch filter must be used in order to not get a biased noise variance estimate with the estimator in Equation (11).

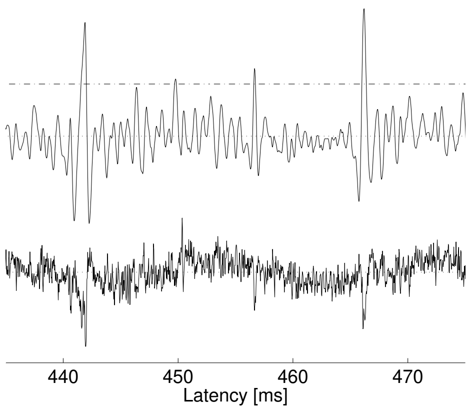

Figure 5. The output of the matched filter (top) when applied on a part of trace 5 in the recording in Fig. 7 (bottom). With the same threshold setting, four detections are reported. Note also the hum in the recorded data.

6.2. Multiple hypotheses tracking

In this section, the performance of the MHT method is evaluated. In Sect. 1, it was mentioned that recent findings imply that the quantative analysis will become more important and that there is a need for an automatic system facilitating this analysis. As the tracks are easily found by an experienced user, it is important that the tracking result is consistently close to the desired. Otherwise, the user could just as well manually mark the APs that should make up a track.

To facilitate the interpretation of the results, a simple example will be shown first, followed by a more realistic scenario.

6.2.1. A single unit recording

In order to show the principles, an example involving only one, high amplitude, C-fiber unit is shown in Fig. 6. First, the unit is inactive and its latency is constant at about 303 ms. At trace 13, the unit is activated by a mechanical stimulus and the latency increases dramatically. Following this, the latency slowly recovers to the level prior to the activation.

As shown in the figure, the algorithm, not surprisingly, tracks most of the APs originating from the unit. Note, however, that the detection threshold is rather low ($m_0 = 5$). The tracking could have been simplified by choosing a higher threshold. For example, if the threshold would have been set to 11, still below the lowest MF peak value of the APs originating from the unit, only the APs marked by crosses in the figure would have been detected and processed.

The good result despite the low threshold is a strength that is important in more complex situations.

6.2.2. A multi unit recording

A more realistic example will now be shown, where several C-units are recorded and have crossing trajectories. Two of the units are inactive during the recording, and one of the units (drawn with a thick line) is strongly activated prior to trace 12. Moreover, there are two spontaneously active sympathetic C-units that can be recognized by their more irregular behaviour, see Fig. 7.

As can be seen in the figure, the tracking result is good for the activated unit as well as the inactive ones. Moreover, the potential importance of the amplitude information should be obvoius, as the latency trajectory of the active unit crosses the two sympathetic units, but their amplitudes differ.

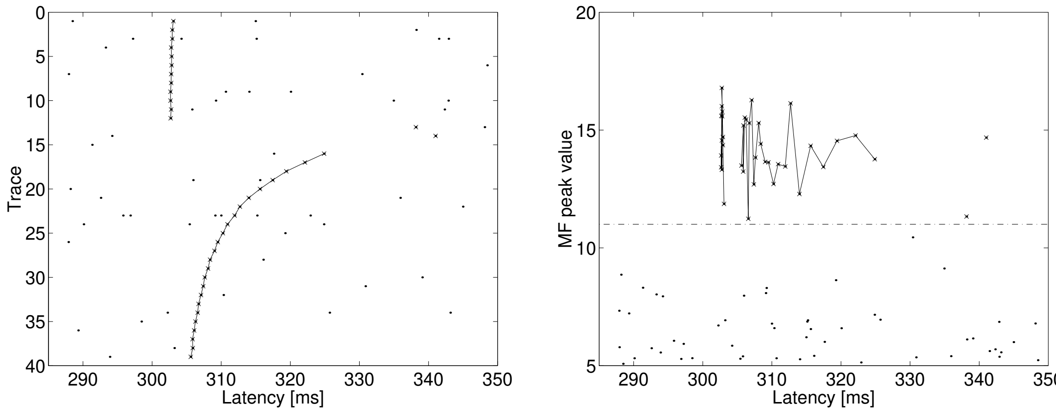

Figure 6. The figure shows, in two perspectives, the tracks that have been confirmed by the tracker when applied to a sample recording with only one active C-fiber unit. In the left plot, the APs with an MF peak output above 11 are marked (x). In the right plot, samples above the detection threshold from all the traces are displayed with their amplitudes as a functon of the latency. The level 11 is here marked with a dash dotted line. Compare with Fig. 1.

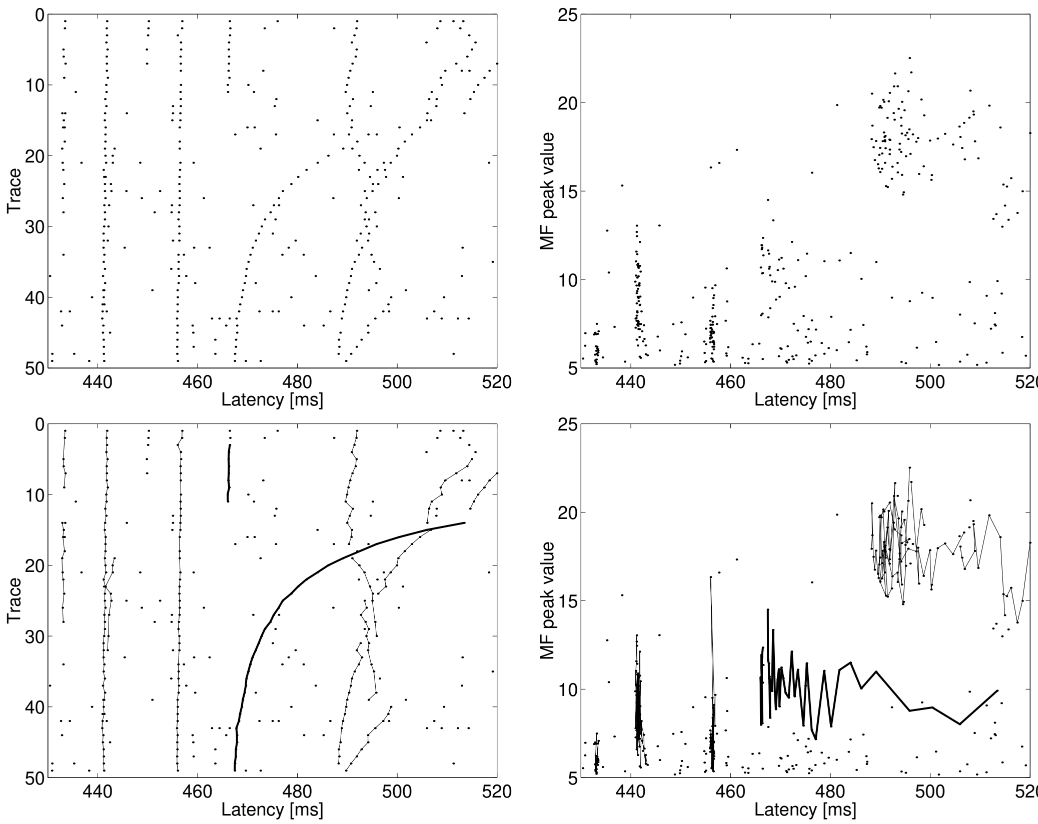

Figure 7. An example of a result of the current tracking algorithm, with five recorded units where the trajectories cross. The detected APs only are plotted as a function of latency both with respect to trace number and with respect to MF peak value (top). The tracking result is plotted similarly, where the activated unit is marked with a thick line. It can be seen that the method handles crossing tracks well if the amplitude of the APs differ significantly (bottom).

7. DISCUSSION

An application of matched filtering and multiple hypothesis tracking applied to human nerve C-fiber action potentials, aiming at quantatively estimating conduction latency shifts and post-stimulus recovery constants of human nerve C-fibers, has been presented. The focus in this paper was on the detection and the discrimination of the action potentials originating from different C-fiber units.

The APs were successfully detected, by means of a matched filter, constituting an optimal ML-CFAR detector. Even APs with amplitude in the order of the noise, were detected.

When deriving the matched filter, the noise was (erroneously) assumed to be white. Further investigations should be carried out where the importance of a more realistic noise model is sought out. One could, for example, estimate the noise covariance matrix in AP free sections of the recording, and use this estimate when calculating the matched filter impulse response. In the practical cases considered so far, the performance of the matched filter has been good despite this deficiency of the algorithm.

The discrimination between APs originating from different sensory units was carried out using the MHT method as described in Ref. 11. Only some minor changes were introduced to adapt the method to the application described. The results are promising and correspond well with what an experienced analyst would consider to be a "correct" result.

For completeness, other methods should be investigated as well, which has not been done. The characteristics of the existing methods known from literature, suggest, however, that no other tracking method would be able to perform equally well. The low track maintenance with tracking based on nearest neighbor (NN) is not acceptable. Moreover, the typical track switching behaviour of the joint probabilistic data association (JPDA) method is not acceptable either.

The MHT method has, in principle, no *inherent* limitation in performance as long as its computational and storage requirements are fulfilled. These requirements could be exhaustive, though, but in the cases considered so far they have not been a limiting factor.

One major drawback of the current implementation, however, is that the filtering performance is sensitive to errors in the model parameter α, which represents an a priori value of the recovery time constant of the latency. It is believed, though, that by introducing robust filtering methods[16,17] or by replacing the Kalman predictor with an IMM based predictor, using multiple values of α, the filtering could be made less sensitive to differencies in the recovery constant.

ACKNOWLEDGMENTS

The authors would like to thank Prof. Mikael Sternad for valuable comments on this work and the manuscript. Comments and suggestions from colleagues at Signals and Systems, Inst. of Physiology and Experimental Pathophysiology, and Dep. of Clinical Neurophysiology have been very useful to improve the quality of this work.

Financial support was obtained from the Swedish Medical Research Council, Proj. no. 5206, the Deutsche Forschungsgemeinschaft (SFB 353), and a Max Planck Research Award to Erik Torebjörk.

REFERENCES

1. R. G. Hallin and H. E. Torebjörk, "Afferent and efferent C units recorded from human skin nerves in situ," *Acta Soc. Med. Ups.* **75**, pp. 277–281, 1970.
2. H. E. Torebjörk and R. G. Hallin, "C-fibre units recorded from human sensory nerve fascicles in situ," *Acta Soc. Med. Ups.* **75**, pp. 81–84, 1970.
3. C. Forster and M. Schmelz, "New developments in microneurography of human C fibers," *News in Physiological Sciences* **11**, pp. 170–175, 1996.
4. H. E. Torebjörk and R. G. Hallin, "Responses in human A and C fibres to repeated electrical intradermal stimulation," *J. Neurol. Neurosurg. Psychiatry* **37**, pp. 653–664, 1974.
5. M. Schmelz, R. Schmidt, M. Ringkamp, C. Forster, H. O. Handwerker, and H. E. Torebjörk, "Limitation of sensitization to injured parts of receptive fields in human skin C-nociceptors," *Exp. Brain Res.* **109**, pp. 141–147, 1996.

6. M. Schmelz, C. Forster, R. Schmidt, M. Ringkamp, H. O. Handwerker, and H. E. Torebjörk, "Delayed responses to electrical stimuli reflect C-fiber responsiveness in human microneurography," *Exp. Brain Res.* **104**, pp. 331–336, 1995.
7. C. Forster and H. O. Handwerker, "Automatic classification and analysis of microneurographic spike data using a PC/AT," *J. Neurosci. Methods* **31**, pp. 109–118, 1990.
8. R. Schmidt, M. Schmelz, C. Forster, M. Ringkamp, H. E. Torebjörk, and H. O. Handwerker, "Novel classes of responsive and unresponsive C nociceptors in human skin," *J. Neurosci.* **15**, pp. 333–341, 1995.
9. J. A. Cadzow, "Matched filters," in *Foundations of Digital Signal Processing and Data Analysis*, pp. 442–461, Macmillan, New York, NY, 1987.
10. A. Steinhardt, "Adaptive multisensor detection and estimation," in *Adaptive Radar Detection and Estimation*, pp. 91–160, Wiley, New York, NY, 1992.
11. S. S. Blackman, *Multiple-Target Tracking with Radar Applications*, Artech House, Dedham, MA, 1986.
12. S. Haykin, *Adaptive Filter Theory*, Prentice Hall, Englewood Cliffs, NJ, second ed., 1991.
13. D. Lerro and Y. Bar-Shalom, "Interacting multiple model tracking with target amplitude feature," *IEEE Trans. on Aerospace and Electronic Systems* **29**(2), pp. 494–508, 1993.
14. M. de Feo, A. Graziano, R. Miglioli, and A. Farina, "IMMJPDA versus MHT and Kalman filter with NN correlation: performance comparison," *IEE Proc.-Radar, Sonar Navig.* **144**(2), pp. 49–56, 1997.
15. H. E. Torebjörk, "Afferent C units responding to mechanical, thermal and chemical stimuli in human non-glabrous skin," *Acta physiol. scand.* **92**, pp. 374–390, 1974.
16. M. Sternad, K. Öhrn, and A. Ahlén, "Robust $\mathcal{H}_2$ filtering for structured uncertainty: the performance of probabilistic and minimax schemes," in *Proc. of 3rd European Control Conf.*, vol. 1, pp. 87–92, (Rome, Italy), 1995.
17. K. Öhrn, *Design of Multivariable Cautious Discrete-Time Wiener Filters: A Probabilistic Approach*. PhD thesis, Uppsala University, Uppsala, Sweden, 1996.