Maximum Likelihood Joint Tracking and Association in a Strong Clutter

Leonid I. Perlovsky, *Senior Member, IEEE*, and Ross W. Deming, *Member, IEEE*

***Abstract*—We have developed an efficient algorithm for the maximum likelihood joint tracking and association problem in a strong clutter for GMTI data. The new tracker overcomes combinatorial complexity of tracking in highly-cluttered scenarios and results in about 20 dB (two orders of magnitude) improvement in signal-to-clutter ratio.**

***Index Terms*—Combinatorial Complexity, Ground Moving Target Indicator Radar, Tracking, Association, Clutter.**

I. INTRODUCTION

Performance of the state-of-the-art algorithms for tracking and association in strong clutter [1], as a function of Signal-to-Clutter Ratio (SCR), is significantly below the information-theoretic limit as indicated by the Cramer-Rao Bound (CRB) for tracking in clutter [2]. The reason for this underperformance is combinatorial complexity of algorithms. When clutter is strong, so that signals are below clutter, multiple associations between data and tracks have to be considered. The number of associations grows combinatorially with the number of data points. Therefore performance is limited by complexity of computations rather than by information in the data. Here we describe a non-combinatorial solution of the maximum likelihood joint tracking and association problem resulting in a significantly improved performance; it follows a discussion at [3].

Standard algorithms (such as Multiple Hypotheses Testing, MHT [4]) used in the current GMTI detection and tracking subsystems operate in a two-step process. First, Doppler peaks are detected that exceed a predetermined threshold. Second, these potential target peaks are used to initiate tracks. This two-step procedure is a state-of-the-art approach which is currently used by most tracking systems. The limitation of this procedure is determined by the detection threshold. If the threshold is reduced, the number of detected peaks grows quickly. Increased computer power does not help because the processing requirements are combinatorial in terms of the number of peaks, so that a tenfold increase in the number of peaks results in a billion fold increase in the required computer power.

I. LIKELIHOOD FOR JOINT TRACKING AND ASSOCIATION

Consider k GMTI radar scans, resulting in n = 1… N measurements $\mathbf{X}(n) = (x_n, y_n, a_n, D_n)$, where (x_n, y_n) are range and cross-range positions, a_n, is amplitude and D_n is Doppler. A likelihood of error measurement, $\mathbf{e}(n)$, is defined as follows. Error measurements are considered independent, therefore,

$$L(\{\mathbf{e}(n)\}) = \prod_{n \in N} \mathrm{pdf}(\mathbf{e}(n)). \qquad (1)$$

A pdf($\mathbf{e}$(n)) is defined according to multiple hypotheses (note a difference in terminology, in MHT algorithm a hypothesis is every *association* between data and track; we call a hypothesis a track or clutter with unknown parameters). The measurement $\mathbf{X}$(n) can originate from clutter, or from one of several moving objects. We number hypotheses h = 1…H; h=1 corresponds to clutter, h = 2…H correspond to H-1 tracks,

$$\mathrm{pdf}(\mathbf{e}(n)) = \sum_{h \in H} r(h)\, \mathrm{pdf}(\mathbf{e}(n)|h), \qquad (2)$$

where r(h) is an a priori probability that measurement n originates according to hypothesis h, and pdf($\mathbf{e}$(n)|h) is a conditional pdf for this hypothesis. Substituting eq.(2) into eq.(1), we obtain

$$L(\{\mathbf{e}(n)\}) = \prod_{n \in N} \sum_{h \in H} r(h)\, \mathrm{pdf}(\mathbf{e}(n)|h). \qquad (3)$$

This product of sums contains H^N items, corresponding to all combinations of data and tracks or clutter (every data point could have originated according to any hypotheses). This huge number is the reason for combinatorial complexity of algorithms in the past. The above notation for conditional pdfs is a shorthand for pdf($\mathbf{e}$(n)| $\mathbf{M}_h$(n)), where $\mathbf{M}_h$(n) is a model predicting measurement $\mathbf{X}$(n); if this measurement originates according to hypothesis h, then $\mathbf{M}_h$(n) is an expected value of this measurement (when true parameter values are used),

$$\mathbf{M}_h(n) = E\{\mathbf{X}(n)|h\}, \qquad (4)$$

and

$$\mathbf{e}(n,h) = \mathbf{X}(n) - \mathbf{M}_h(n). \qquad (5)$$

We consider tracking short track segments, tracklets, along which velocities can be considered constant $\mathbf{V}_h$ = (V_{hx} , V_{hy}). Correspondingly, the complete model is

$$\mathbf{M}_h(\mathbf{S}_h,n) = (X0_h+V_{hx}t_n,\ Y0_h+V_{hy}t_n,\ a_h,\ D_h). \qquad (6)$$

Here parameters of the model, $\mathbf{S}_h$ = ($X0_h$, $Y0_h$, V_{hx}, V_{hy}, a_h, D_h); ($X0_h$, $Y0_h$) model an original position, (V_{hx}, V_{hy}) model velocity, (a_h, D_h) model amplitude and Doppler; t_n, is the known time counted from the first scan. Also,

$$V_{hx} = D_h. \qquad (7)$$

The unknown parameters also include r(h), parameters of conditional pdf, such as standard deviations or covariances, and the total number of track-models. Conditional pdf for clutter we define as uniform,

$$\mathrm{pdf}(\mathbf{X}(n)|1) = 1/\mathrm{par\_volume}, \qquad (8)$$

where par_volume is a volume of the parameter space, a product of ($S_{max} - S_{min}$) for all parameters. Conditional pdfs of tracks are defined as Gaussian; in view of hard boundaries S_{max}, S_{min}, this is an approximation; also parameters a_h are not likely to follow Gaussian distributions. In our practical cases this approximate treatment has been sufficient,

$$\mathrm{pdf}(\mathbf{e}(n,h)|h) = (2\pi)^{-2}(\det \mathbf{C}_h)^{-0.5} \cdot \exp\{-0.5\ \mathbf{e}(n,h)^T \mathbf{C}_h^{-1}\ \mathbf{e}(n,h)\}. \qquad (9)$$

We use diagonal covariance matrixes $\mathbf{C}_h = \mathrm{diag}(\sigma_{x_h}^2, \sigma_{y_h}^2, \sigma_{a_h}^2, \sigma_{D_h}^2)$; and, $\sigma_{x_h}^2 = \sigma_{D_h}^2$.

II. Dynamic logic

Now we describe a procedure to maximize the likelihood (1) while at the same time solving the association-assignment problem without combinatorial complexity. We call this procedure dynamic logic (DL) for reasons described in [5], [6]. It is similar to a procedure described in [7]; that publication also has given a detailed review of relationships of this technique to previous related publications, which we briefly summarize here. A fundamental idea of probabilistic association has originated from Bar-Shalom [8]. The general framework for the approach described in this paper was developed by Perlovsky in [9], [10], [11], [12], [13], [14], [15] and other publications listed in [15]. Perlovsky's DL tracker has much in common with a similar algorithm known as the probabilistic multihypothesis tracker (PMHT) [16], [17], [18] which was developed independently by Streit and Luginbuhl, and shares strong similarities with Avitzour's approach [19]. When applied to a benchmark multi-target tracking problem, it was found that the computational cost of PMHT has roughly the same order of magnitude as the cost of MHT and JPDAF [20] (combinatorial). The computational cost of the DL tracker scales only linearly with increasing numbers of data. The main idea of the DL procedure, which resulted in reduced computational complexity from combinatorial to linear, is coordination of the certainty of model parameters and certainty of assignment-associations as discussed later.

DL is an iterative procedure, which starts with unknown values of model parameters and correspondingly large uncertainty of associations; this later requirement is achieved by setting standard deviation of parameters equal to one half of (max – min) value for this parameters (which are usually approximately known in an operation scenario). Taking these initial values of parameters, conditional probabilities eq.(9) are computed. Then association variables are computed; they are defined similarly to posteriori Bayes probabilities (for shortness, we use indexes n, h instead of the corresponding data $\mathbf{X}(n)$ and models $\mathbf{M}_h(n)$.

$$f(h|n) = r(h)\ \mathrm{pdf}(n|h) \Big/ \sum_{h' \in H} r(h')\ \mathrm{pdf}\ (n|h'). \qquad (10)$$

Although eq.(10) looks like posteriori Bayes probabilities, f(h|n) are not probabilities, since parameter values are incorrect; they can be called association variables or estimated probabilities of measurements n originating from tracks (objects or hypotheses) h.

The next step is to estimate parameters, using these estimated association probabilities. The following equations are used for r(h),

$$r(h) = \sum_{n \in N} f(h|n)/N. \qquad (11)$$

This gives an estimated average ratio of data points assigned to track h (or clutter h=1) to the total number of data points N. This and other parameter estimation equations look simpler with the following notation:

$$\langle \ldots \rangle_h = \sum_{n \in N} f(h|n)\ (\ldots)_n. \qquad (12)$$

Then eq.(11) can be rewritten as

$$r(h) = \langle 1 \rangle_h / N. \qquad (13)$$

Other parameter estimation equations, at each iteration, are computed as

$$a_h = \langle a_n \rangle_h. \qquad (14)$$

$$Y0_h \langle 1 \rangle_h + V_{yh} \langle t_n \rangle_h = \langle Y_n \rangle_h,$$
$$Y0_h \langle t_n \rangle_h + V_{yh} \langle t_n^2 \rangle_h = \langle Y_n t_n \rangle_h. \qquad (15)$$

$$X0_h \langle 1 \rangle_h + V_{xh} \langle t_n \rangle_h = \langle X_n \rangle_h,$$
$$X0_h \langle t_n \rangle_h + V_{xh} (\langle t_n^2 \rangle_h + c \langle 1 \rangle) =$$
$$= \langle X_n t_n \rangle_h + c \langle D_n \rangle_h. \qquad (16)$$

Here, $c = \sigma_{xh}^2 / \sigma_{Dh}^2$. For the unknown parameters, $Y0_h$ and V_{yh}, eq.(15) is a two-dimensional linear system of equations; similarly eq.(16) is a two-dimensional linear system of equations for $X0_h$ and V_{xh}. Standard deviations for each parameter s are estimated, as follows:

$$\sigma_{hs}^2 = \langle (X_s(n) - M_{hs}(n))^2 \rangle_h \qquad (17)$$

DL consists in iterative computations of eqs.(10) through (17). While estimated parameters are far from true values, models do not match data, standard deviations are large, and associations f(h|n) are “flat”: small numbers for many combinations of n and h (including incorrect ones), any data point has a nonzero assignment to any track (or clutter). Nevertheless, even with these poor initial associations, parameter values improve on every iteration according to a theorem proven in [15]: Likelihood (1) grows on every iteration and the DL procedure converges (local vs. global convergence is discussed below). As parameter values converge close to their true values, standard deviations converge to small values close to the sensor errors. Association variables converge close to true probabilities, close to 1 for n and h pairs corresponding to data n originating from object h, and to 0 otherwise (this last statement is true to the extent that the information contained in the data is sufficient for track separability and data association). This DL process from vague to crisp associations is characteristic of DL [15].

The computational complexity of the DL procedure described above is proportional to the number of data points and the number of tracks, const*N*H. The const here accounts for the number of iterations, and for complexity of procedures described by (10) through (17). Typical numbers are discussed in the next section. The principal theoretical moment is that this number is linear in *N* and in *H*, rather than combinatorial, $\sim H^N$ like in MHT.

The number of tracks is estimated as follows. The algorithm starts with 1 *active* track model, which parameters are updated from iteration to iteration. In addition, the algorithm keeps one (or several) *dormant* track, which parameters are not updated, except for r(h). On the 1st iteration all standard deviations are large and the track and clutter models have low but nonzero associations to all data points. After few iterations the active track standard deviations become smaller, it is stronger associated with some returns and weaker with others. According to eq.(11), sum total associations of every data point n, $\Sigma_h f(h|n) = 1$; therefore some associations for dormant tracks grow. After r(h) for a dormant track exceeds a predetermined threshold, this track is activated and its parameters are updated. Similarly, if r(h) falls below the threshold, the track is eliminated. In this way as many tracks are activated as justified by the data. This procedure may lead to too many active tracks. When standard deviations approach sensor error values and association variables approach 0s and 1s, extra tracks tend to converge either on top of each other, or to one or two data points, these are pruned. Upon convergence (likelihood eq.(1) increase between iterations become less than a threshold), a detection measures is computed for each track; it is defined as a local log-likelihood ratio computed using returns within two standard deviations {*n'*} of each track:

$$LLR(h) = \sum_{n' \in N'} [\ \ln pdf(n'|h) - \ln pdf(n'|1)\ ]. \quad (18)$$

Tracks with LLR(h) exceeding a predetermined threshold are declared detections.

Convergence of this iterative procedure to a local maximum of similarity measure (1), as mentioned, was proven in [15]. Such local convergence usually occurs within relatively few iterations; a typical example in the next section took 20 iterations. Since similarity is a highly non-linear function, regular convergence to the global maximum can not be expected. The local rather than global convergence sometimes presents an irresolvable difficulty in many applications. In the presented method, this problem is resolved in several ways. First, the large initial standard deviation of the similarity measure smoothes local maxima. Second, tracks are pruned and activated as needed. Therefore if a particular real track is not "captured" after few iterations, it will be captured at a later iteration, after a track-model activation. Third, if a spurious track is declared detected, or a real track is missed, these errors will be self-corrected at a later stage of a system operation, when detected track segments or tracklets are connected into longer tracks (system operation procedures are beyond the current communication).

III. Tracking example and ROC

An application example of the above described DL tracker is illustrated in Fig. 1, where detection and tracking are performed for targets below the clutter level. Fig. 1(a) shows true track positions in a 0.5km * 0.5km data set, while Fig. 1(b) shows the actual data available for detection and tracking. In this data, the target returns are buried in the clutter, with signal-to-clutter ratio of about –2dB for amplitude and –3dB for Doppler. Here, the data are displayed such that all six revisit scans are shown superimposed in the 0.5km * 0.5km area, 500 pre-detected signals per scan, and the brightness of each data sample is proportional to its measured Doppler value. Figs. 1(c)-1(h) illustrate the dynamics of the algorithm as it adapts during increasing iterations; the brightness is proportional to association variables, which for this display purpose are computed not just for **X**(n) but for all pixels (resulting in a smooth image shape). Only association variables for active track models are shown. Fig. 1(c) shows the initial vague track-model, and Fig. 1(h) shows track-models upon convergence at 20 iterations. Between (c) and (h) the DL tracker automatically decides how many track-models are needed to fit the data, and simultaneously updates the track parameters and association variables. There are two *types* of models: one uniform model describing clutter (it is not shown), and linear track-models, which uncertainty changes from large (c) to small (h). In (c) and (d), the DL tracker fits the data with one model, and uncertainty is somewhat reduced. Between (d) and (e) the DL tracker uses more than one track-model and decides that it needs two models to 'understand' the content of the data. Fitting with 2 tracks continues until (f); between (f) and (g) a third track is added. Iterations stop at (h), when similarity stops increasing. Detected tracks closely correspond to the truth (a).

FIG.1 GOES HERE

Fig. 1. Detection and tracking three targets in clutter using DL: (a) true track positions in 0.5km * 0.5km data set; (b) actual data available for detection and tracking. DL iterations are illustrated in (c) – (h), where (c) shows the initial, uncertain model and (h) shows the models upon convergence after 20 iterations. Note the close agreement between the converged models (h) and the truth (a).

In this example, target signals are below clutter. A single scan does not contain enough information for detection. Detection should be performed concurrently with tracking, using several radar scans, and six scans are used. In this case, a standard multiple hypothesis tracking, evaluating all tracking association hypothesis, would require about 10^{5000} operations, a number too large for computation. Therefore, existing tracking systems require strong signals, with about a 15 db signal-to-clutter ratio [1]. DL successfully detected and tracked all three targets and required only 10^6 operations, achieving about 18 dB improvement in signal-to-clutter sensitivity.

A detailed characterization of performance requires operating curves (ROC), plots of probability of detection vs. probability of false alarm, computed for various signal-to-clutter ratios, densities of targets, target velocities, and other scenario parameters. Such detailed characterization is beyond the scope of this communication. Instead, Fig.2 illustrates three ROCs for selected parameter values.

.

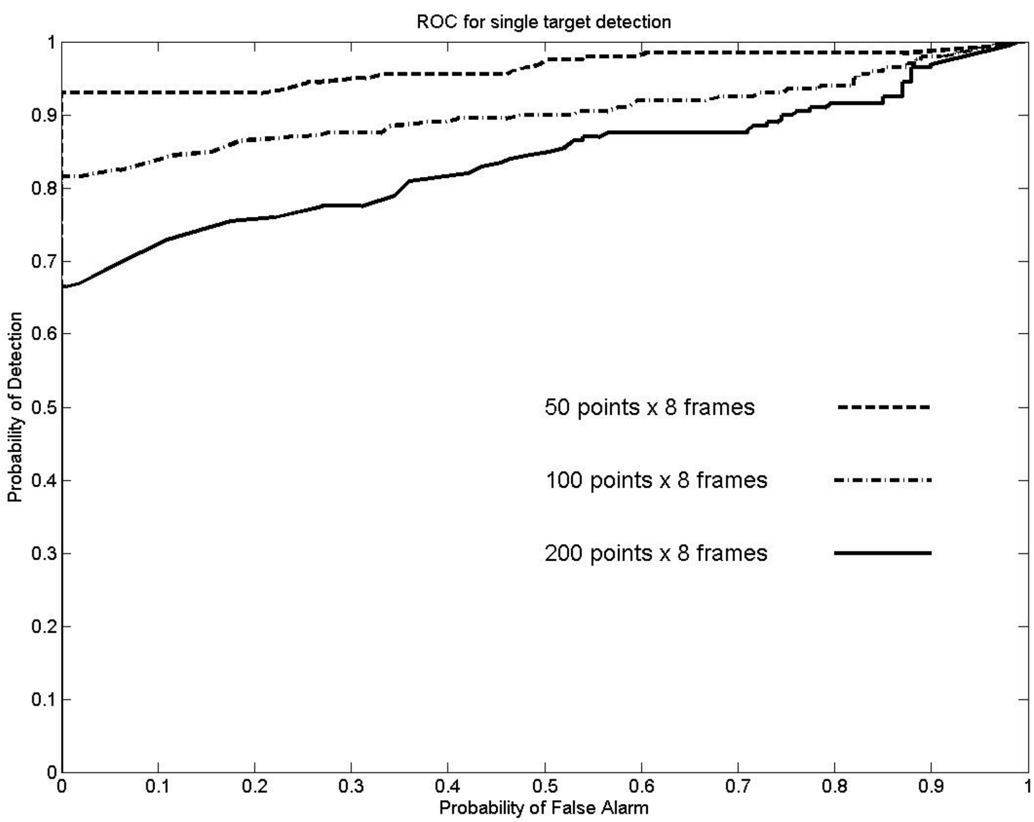

Fig. 2. Three ROC curves for different clutter levels: 50, 100, and 200 pre-detected signals per frame. 8 frames are used (total of 400, 800, and 1600 clutter signals per 8 target signals). Signal to clutter ratio, S/C, is defined as a signal strength divided by standard deviations taken as a sum of clutter and target standard deviations: S/C = $[(\sigma_C + \sigma_T)]$. S/C is 1.7 for amplitude and 2.0 for Doppler.

IV. CONCLUSION

The paper presented a maximum likelihood solution for tracking in clutter, while avoiding combinatorial complexity.

Future research will include feature-added tracking when - in addition to amplitude, position, and velocity - other characteristics of received signals are also used for improved associations between signals and track models. DL can naturally incorporate this additional information. Since association neural weights in DL are functions of object models (1) any object feature can be included into the models and will be used for signal-model associations.

Other sources of information can be included. For example, coordinates of roads can be easily incorporated into the DL procedure. For this purpose road positions should be characterized by a probability density, depending on the known coordinates and expected errors. Then similarities (2) can be modified by multiplying them by the probability densities of roads.

ACKNOWLEDGEMENT

The authors are thankful to AFOSR for supporting part of this research under the Lab. Task 05SN02COR, PM Dr. Jon Sjogren. LP is thankful to participants of the ECE Colloquium, University of Connecticut at Storrs for detailed discussion.

Leonid Perlovsky (M'86-SM'81) Ph.D.'74 in theoretical and mathematical physics, Institute for Mathematics, Novosibirsk and Joint Institute for Nuclear Research, Dubna, Moscow, Russia.

He is Visiting Scholar at Harvard University, Cambridge, MA and Principal Research Physicist and Technical Advisor at the Air Force Research Laboratory, Hanscom AFB, MA, Program Manager for DOD Semantic Web program and other research projects. From 1985 to 1999, he served as Chief Scientist at Nichols Research. Previously Professor at Novosibirsk University and New York University, principal in commercial startups developing tools for text understanding, biotechnology, and financial predictions. His company predicted the market crash following 9/11 a week before the event, detecting activities of Al Qaeda traders, and later helped SEC tracing the leads. He delivered invited keynote plenary talks and tutorial lectures worldwide, published more then 270 papers, 7 book chapters, and authored a monograph "Neural Networks and Intellect," Oxford University Press, 2001 (currently in the 3rd printing). Two more books are scheduled to appear in 2007.

Dr. Perlovsky is a member of INNS, Chairs IEEE Boston Computational Intelligence Chapter, organizes conferences on Computational Intelligence, serves as Associate Editor for IEEE Transactions on Neural Networks, Editor-at-Large for "Natural Computations" and Editor-in-Chief for "Physics of Life Reviews." He received the IEEE Distinguished Member of Boston Section Award 2005; the US AF Research Lab Charles Ryan Memorial Award for Basic Research 2007; and International Neural Network Society Gabor Award for achievement in engineering/applications 2007.

Ross Deming (M'05) received the B.S. degree from Cornell University in 1985, the M.S. degree from the University of Vermont in 1993, and the Ph.D. degree from Northeastern University in 1996, all in electrical engineering.

He was previously employed at General Electric, Inc. and Teradyne, Inc., as an electrical engineer, and at Nichols Research Corp. and Witten Technologies, Inc., as a senior scientist. Currently, he is a consultant for General Dynamics, Inc., where he is working on radar and optical signal processing problems for the Air Force Research Laboratory.

While at Witten Technologies (WTI), he was part of a small team that won, among other awards, the Wall Street Journal's 2004 Technology Innovation Award (1st place software category, honorable mention overall) for its technology that creates detailed images of objects and conditions underground using radar. He is also coauthor on a patent for this radar system.

FIGURE CAPTIONS

Fig. 1. Detection and tracking three targets in clutter using DL: (a) true track positions in 0.5km * 0.5km data set; (b) actual data available for detection and tracking. DL iterations are illustrated in (c) – (h), where (c) shows the initial, uncertain model and (h) shows the models upon convergence after 20 iterations. Note the close agreement between the converged models (h) and the truth (a).

Fig. 2. Three ROC curves for different clutter levels: 50, 100, and 200 pre-detected signals per frame. 8 frames are used (total of 400, 800, and 1600 clutter signals per 8 target signals). Signal to clutter ratio, S/C, is defined as a signal strength divided by standard deviations taken as a sum of clutter and target standard deviations: S/C = [$(\sigma_C + \sigma_T)$]. S/C is 1.7 for amplitude and 2.0 for Doppler.

REFERENCES

[1] J. Jones, J. Bradstreet, M. Kozak, T. Hughes, and M. Blount, "Ground moving target tracking and exploitation performance measures," Pentagon Report A269234, approved for public release, 2004.
[2] L. I. Perlovsky, "Cramer-Rao bound for tracking in clutter and tracking multiple objects," Pattern Recognition Letters, vol. 18, pp. 283-288, 1997.
[3] L. I. Perlovsky, "Maximum likelihood joint association, tracking, and fusion in strong clutter," Department of Electrical and Computer Engineering Seminar, University of Connecticut at Storrs, 6 March, 2009.
[4] R. A. Singer, R. G. Sea, and R. B. Housewright, "Derivation and evaluation of improved tracking filters for use in dense multitarget environments," IEEE Trans. Information Theory, vol. IT-20, pp. 423-432, 1974.
[5] L. I. Perlovsky, "Toward Physics of the Mind: Concepts, Emotions, Consciousness, and Symbols," Phys. Life Rev. 3(1), pp.22-55, 2006.
[6] B. Kovalerchuk and L. I. Perlovsky, "Dynamic Logic of Phenomena and Cognition," World Congress on Computational Intelligence (WCCI). Hong Kong, China, 2008.
[7] R. Deming, J. Schindler, L. I. Perlovsky. (2008). Multitarget/Multisensor Tracking using only Range and Doppler Measurements, IEEE Transactions on Aerospace and Electronic Systems, in print.
[8] Fortmann, T., Bar-Shalom, Y., and Scheffe M. Sonar tracking of multiple targets using joint probabilistic data association, IEEE J. Oceanic Eng., OE-8 (1983), pp. 173-183.
[9] Perlovsky, L.I. (1987). Multiple Sensor Fusion and Neural Networks. DARPA Neural Network Study, MIT/Lincoln Laboratory, Lexington, MA.
[10] Perlovsky, L.I. (1988). Neural Networks for Sensor Fusion and Adaptive Classification. First Annual International Neural Network Society Meeting, Boston, MA.
[11] Perlovsky, L.I. (1990). Development of Neural Networks for Adaptive Classification and Sensor Fusion. Ballistic Missile Discrimination Workshop, MIT/Lincoln Laboratory, Lexington, MA.
[12] Perlovsky, L.I. (1991). Tracking Multiple Objects in Visual Imagery. Conference on Neural Networks for Vision and Image Processing, Tyngsboro, MA.
[13] Perlovsky, L.I. & McManus, M.M. (1991). Maximum Likelihood Neural Networks for Sensor Fusion and Adaptive Classification. Neural Networks 4 (1), pp. 89-102.
[14] Perlovsky, L.I., Schoendorf, W.H., Tye, D.M., Chang, W. (1995). Concurrent Classification and Tracking Using Maximum Likelihood Adaptive Neural System. Journal of Underwater Acoustics, 45(2), pp.399-414.
[15] L. I. Perlovsky, Neural Networks and Intellect, New York, NY: Oxford Univ. Press, 2001.
[16] Streit, R. L., and Luginbuhl, T. E., Maximum likelihood method for probabilistic multi-hypothesis tracking. In Proceedings of SPIE International Symposium, Signal and Data Processing of Small Targets 1994, Vol. 2335-24, Orlando, FL, Apr. 5-7, 1994.
[17] Willett, P., Ruan, Y., and Streit, R. PMHT: Problems and some solutions. IEEE Transactions on Aerospace and Electronics Systems, 38 (2002), 738-754.
[18] Ruan, Y., and Willett, P. Multiple model PMHT and its application to the second benchmark radar tracking problem. IEEE Transactions on Aerospace and Electronics Systems, 40 (2004), 1337-1347.
[19] Avitzour, D. A maximum likelihood approach to data association. IEEE Transactions on Aerospace and Electronics Systems, 28 (1992), 560-565.
[20] Ruan, Y., and Willett, P. Multiple model PMHT and its application to the second benchmark radar tracking problem. IEEE Transactions on Aerospace and Electronics Systems, 40 (2004), 1337-1347.