Optimal data fusion for thunderstorm risk assessment by the saddlepoint method

Christian Musso, Aurélie Bouchard, Magalie Buguet

ONERA – The French Aerospace Lab
Palaiseau, France
{christian.musso, aurelie.bouchard, magalie.buguet}@onera.fr

***Abstract*—Thunderstorm risk assessment is a key point in the framework of the air traffic management and aircraft safety. We consider a relevant list of 10 meteorological parameters characterizing the state of the atmosphere from the standpoint of lightning strike risk. Each parameter makes a binary decision by performing a likelihood test based on its local threshold. The optimal decision fusion is a weighted sum of discret random variables not identically distributed. We show that the law of this sum can be accurately estimated by the saddlepoint approximation. We propose a method to maximize the detection probability with respect to the local thresholds, subject to the constraint of a global false alarm probability. An illustration of the application of this method to thunderstorm risk assessment has been done on the Gulf of Mexico.**

***Index Terms*—thunderstorm risk, data fusion, thresholds optimization, saddlepoint approximation, detection probability optimization**

I. Introduction

Lightning risk is one of the risks for aviation, when aircraft fly near or inside a thunderstorm (the main lightning generator on Earth). In recent years, different methods to assess the risk of thunderstorms have been developed based on different mathematical concepts such as belief function [1], machine learning [2], ensemble prediction model [3] or based on the definition of new index using thresholds [4]. In this study, a new method is proposed for thunderstorm risk assessment based on optimal data fusion. This method has the specificity of being able to set a global false alarm probability P_{fa}.

The paper considers the problem of the fusion of binary decisions H_1 or H_0. Each local decision is made by performing a likelihood ratio test (LR) based on a local measurement and on a local threshold. The problem of optimal decision fusion has been formulated in [5], [6] in terms of weighted sum of incoming decisions. This statistic is a weighted sum of discrete variables not identically distributed. Under each hypothesis, we show that the law of this statistic can be accurately estimated by the saddlepoint method. We present an efficient method to maximize the global detection probability P_d with respect to the local thresholds subject to a global P_{fa} constraint. Local decisions are then made based on these optimized thresholds. By performing a LR test on the weighted sum of the local decisions, the final decision is made. Thanks to the saddlepoint approximation, the global P_d and P_{fa} are estimated accurately. Similar approaches have been proposed in the context of distributed data fusion [7] and in the context of serial distributed decision fusion [8]. Saddlepoint approximation has also been applied in a centralized detection network where the raw data are transmitted to the fusion center [9], in the context of parallel fusion network [10] and in the field of channel coding [11].

In section II, we recall the optimal decisions fusion rule based on the Neyman-Pearson test. In section III, we describe the saddlepoint approximation and we adapt it to estimate the law of a sum of discrete random variables not identically distributed. In section IV, a method of optimization is described to maximize the global detection probability with respect to the local thresholds. Section V is devoted to the application of the proposed method to thunderstorm risk assessment using real data from meteorological parameters.

II. Risk assessment methodology

A. Likelihood ratio test for binary decision

Under the hypothesis H_1 and under H_0, denoting the distribution of the measurement z_i from the i-th meteorological parameter as $\mathbb{P}(z_i|H_1)$ and $\mathbb{P}(z_i|H_0)$ respectively, the likelihood ratio (LR) test, referred as the Neyman-Pearson test [12], is expressed as follows:

$$\Lambda(z_i) = \frac{\mathbb{P}(z_i|H_1)}{\mathbb{P}(z_i|H_0)} \underset{H_0}{\overset{H_1}{\gtrless}} \eta_i. \tag{1}$$

We denote the local decisions as $X_1, \ldots, X_n$:

$$\begin{cases} X_i = 1 \iff \Lambda(z_i) \geq \eta_i \\ X_i = 0 \iff \Lambda(z_i) < \eta_i. \end{cases} \tag{2}$$

The conditional laws of X_i are assumed to be known for any value the threshold η_i:

$$\begin{cases} \mathbb{P}_0(X_i = 0) \overset{\text{def}}{=} \mathbb{P}(X_i = 0|H_0) = 1 - P_{fa}(i) \overset{\text{def}}{=} 1 - \alpha_i \\ \mathbb{P}_1(X_i = 1) \overset{\text{def}}{=} \mathbb{P}(X_i = 1|H_1) = P_d(i) \overset{\text{def}}{=} \mu_i. \end{cases} \tag{3}$$

The optimal detection rule based on the local decisions $(X_1, \ldots, X_n)$ is given by the LR test:

$$\frac{\mathbb{P}(X_1 = x_1, \ldots, X_n = x_n|H_1)}{\mathbb{P}(X_1 = x_1, \ldots, X_n = x_n|H_0)} \geq \eta \implies H_1. \tag{4}$$

We assume that the measurements z_i, conditioned on each hypothesis, are independent. The global likelihood ratio test is then expressed as follows:

$$Z = \sum_{i=1}^{n} \log \left[\frac{\mathbb{P}_1(X_i = x_i)}{\mathbb{P}_0(X_i = x_i)} \right] \geq \log \eta \implies H_1 \tag{5}$$

Since the variables X_i take the values $x_i = 0, 1$, we have:

$$\begin{cases} \mathbb{P}_1(X_i = x_i) = \mu_i^{x_i}(1-\mu_i)^{1-x_i} \\ \mathbb{P}_0(X_i = x_i) = \alpha_i^{x_i}(1-\alpha_i)^{1-x_i}. \end{cases} \tag{6}$$

By replacing these terms in (5), we obtain:

$$\begin{aligned} Z &= \sum_{i=1}^{n} \log \left[\frac{1-\mu_i}{1-\alpha_i} \right] + \sum_{i=1}^{n} \log \left[\frac{\mu_i(1-\alpha_i)}{\alpha_i(1-\mu_i)} \right] X_i \\ &= \text{const} + \sum_{i=1}^{n} a_i X_i, \end{aligned} \tag{7}$$

where

$$a_i = \log \left[\frac{\mu_i(1-\alpha_i)}{\alpha_i(1-\mu_i)} \right]. \tag{8}$$

As a result, by denoting $Y_i = a_i X_i$, the LR test (4) test relies on the sufficient statistic $\bar{Y} = \frac{1}{n}\sum_{i=1}^{n} Y_i$. Hence, the optimal final decision is given by the following test:

$$\text{LR test} : \begin{cases} \bar{Y} = \dfrac{1}{n}\displaystyle\sum_{i=1}^{n} a_i X_i \geq \gamma \implies H_1 \\ \bar{Y} = \dfrac{1}{n}\displaystyle\sum_{i=1}^{n} a_i X_i < \gamma \implies H_0. \end{cases} \tag{9}$$

The threshold γ is determined by a fixed global false alarm probability P_{fa}, namely:

$$P_{fa} = \mathbb{P}\left(\frac{1}{n}\sum_{i=1}^{n} Y_i \geq \gamma | H_0 \right) = \mathbb{P}_0(\bar{Y} \geq \gamma). \tag{10}$$

We observe that $\bar{Y}$ is a weighted sum of independent discrete random variables not identically distributed. Indeed, the law of Y_i depends on (α_i, μ_i) (3). The issue is to approximate the law of $\bar{Y}$ especially in the tail areas since the global false alarm probability P_{fa} (10) will take small values.

III. Saddlepoint approximation

Since $\bar{Y}$ is a weighted sum of n discret random variables (9), the central limit theorem (CLT) is not suitable to estimate the law of $\bar{Y}$. The CLT gives poor estimations in the tail areas especially in terms of relative errors. On the contrary, as we will see below, the saddlepoint (SP) approximation can be adapted to accurately estimate the cumulative density function (CDF) of $\bar{Y}$, in particular in the tail area.

A. Introduction to saddlepoint approximation

Originally, the saddlepoint method [13] has been developed to approximate the distribution of the mean of i.i.d. continuous random variables Z_i: $\bar{Z} = \frac{1}{n}\sum_{i=1}^{n} Z_i$. The thesis [14] provides a complete overview of the saddlepoint method. Theoretical aspects are given in [15]. The probability density function (pdf) of $\bar{Z}$ is approximated as follows:

$$f_{\bar{Z}}(\bar{z}) \approx \left(\frac{n}{2\pi K''(T_0)} \right)^{\frac{1}{2}} \exp\left[n(K(T_0) - T_0\,\bar{z} \right], \tag{11}$$

where $K(t) = \log \mathbb{E}\left[e^{tZ} \right]$ is the cumulant generating function (CGF) of Z_i and where T_0 is the saddlepoint, solution of the equation: $K'(t) - \bar{z} = 0$. A very accurate approximation of the CDF, even in the tail area and for small values of n, is the following [16]:

$$\mathbb{P}[\bar{Z} \geq \bar{z}] \approx 1 - \Phi(u) + \phi(u)\left(\frac{1}{t} - \frac{1}{u} \right), \tag{12}$$

Φ and ϕ stand for the standard normal CDF and for the pdf respectively, where $u = \text{sign}(T_0)\sqrt{2n}\,(T_0\,\bar{z} - K(T_0))^{1/2}$ and $t = T_0\,(n\,K''(T_0))^{1/2}$. When $\bar{z}$ is close to the mean of $\bar{Z}$, T_0 is close to 0 and equation (12) should be replaced by the following one [15]:

$$\mathbb{P}[\bar{Z} \geq \bar{z}] \approx \frac{1}{2} - \frac{1}{6\sqrt{2\pi n}} \frac{K'''(T_0)}{K''(T_0)^{\frac{3}{2}}}. \tag{13}$$

B. Application to the estimation of the law of $\bar{Y}$

Saddlepoint approximations of i.i.d discrete random variables are similar to that of continuous random variables [17]. The discrete random variables, $Y_i = a_i X_i$, conditioned on each hypothesis, are assumed independent (5) but they are not identically distributed. To derive an approximation of $\mathbb{P}(\bar{Y} \geq \bar{y})$ by the saddlepoint method, one can prove that it is right to apply the formulae (12) with $K(t) = \frac{1}{n}\sum_{i=1}^{n} K_i(t)$ where $K_i(t) = \log \mathbb{E}\left[e^{tY_i} \right]$ is the CGF of $Y_i = a_i X_i$. The expression of $K(t)$ is the following:

$$\begin{cases} K(t) = \dfrac{1}{n}\displaystyle\sum_{i=1}^{n} \log\left(1 - \alpha_i + \alpha_i\, e^{a_i t} \right) \quad \text{under } H_0 \\ K(t) = \dfrac{1}{n}\displaystyle\sum_{i=1}^{n} \log\left(1 - \mu_i + \mu_i\, e^{a_i t} \right) \quad \text{under } H_1. \end{cases} \tag{14}$$

The derivatives $K'(t)$, $K''(t)$ and $K'''(t)$ involved in the SP approximation (12), (13) can be straightforward calculated. For each hypothesis H_i, the saddlepoint T_0 is the solution of $K'(t) - \bar{y} = 0$, which can be found easily since $K(t)$ is convex [18].

The following example illustrates the accuracy of the saddlepoint approximation. For this, we will require the Receiver Operating Characteristic (ROC) curves of the meteorological parameters connecting the local P_d to the local P_{fa}: $\mu_i = \text{ROC}(\alpha_i)$. They are described in section V. We propose to estimate $\rho_0 = \mathbb{P}_0(\frac{1}{n}\sum_{i=1}^{n} a_i X_i \geq \gamma)$ for several γ by the SP approximation (12) with $n = 10$. The coefficients a_i (8) and the law of X_i (3) depend on (α_i, μ_i) marked with an asterisk

(*) in the ROC curves (Fig. 3). To assess the accuracy of the SP method, we estimate empirically ρ_0 by sampling 10^7 Bernoulli variables: $X_i \sim \text{Bernoulli}(\alpha_i)$. The latter estimation is considered as exact. As we can see on figure (1), the SP approximation is accurate and smooth for the whole range of the threshold γ.

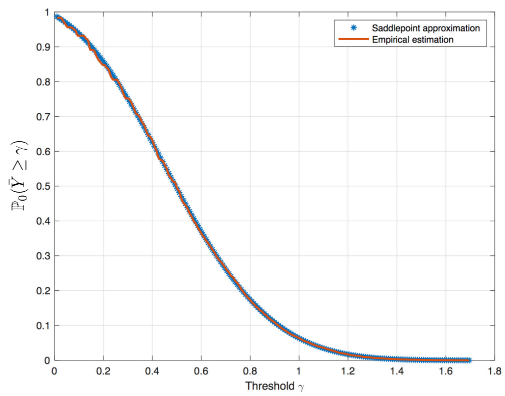

Fig. 1: SP approximation and empirical estimation of ρ_0

IV. Detection probability optimization

A. Local threshold optimization

We aim to optimize the global detection probability P_d with respect to the local thresholds η_i (1) related to local decisions X_i (2). This is equivalent to maximize P_d with respect to the local false alarm probabilities α_i (3). This optimization is subject to the constraint on the global false alarm probability P_{fa} (10). Hence, the optimization problem leads to:

$$\begin{cases} \max\limits_{(\alpha_1 \ldots, \alpha_n)} \mathbb{P}\left(\frac{1}{n}\sum_{i=1}^{n} a_i X_i \geq \gamma | H_1\right) \\ \text{subject to} \\ \mathbb{P}\left(\frac{1}{n}\sum_{i=1}^{n} a_i X_i \geq \gamma | H_0\right) = P_{fa}, \end{cases} \quad (15)$$

where the binary random variables X_i are defined in (3). Under each hypothesis, the distribution of $\frac{1}{n}\sum_{i=1}^{n} a_i X_i$ depends only on $(\alpha_1 \ldots, \alpha_n)$ since the local detection probability μ_i (3) is a function of α_i thanks to the ROC curves (Fig. 3). We propose a formulation in which the constraint (15) is included in the objective function to be maximized. Assessment of this function is done in three stages.

Objective function formulation

- Compute the coefficients a_i (8)

$$(\alpha_1, \ldots, \alpha_n) \rightarrow (a_1, \ldots, a_n) = g_0(\text{ROC}) \quad (16)$$

- Compute the threshold γ satisfying the global P_{fa} constraint

$$\begin{aligned} (a_1, \ldots, a_n) &\rightarrow \mathbb{P}_0[\bar{Y} = \frac{1}{n}\sum_{i=1}^{n} a_i X_i \geq \gamma] = g_1(\gamma) \\ g_1(\gamma) = P_{fa} &\rightarrow \gamma = g_2(P_{fa}, \alpha_1, \ldots, \alpha_n) \end{aligned} \quad (17)$$

- Compute the objective function G

$$\gamma \rightarrow P_d = \mathbb{P}_1[\frac{1}{n}\sum_{i=1}^{n} a_i X_i \geq \gamma] \stackrel{\text{def}}{=} G[P_{fa}, \alpha_1, \ldots, \alpha_n]. \quad (18)$$

All the probabilities involved in the evaluation of the objective function formulation are estimated by the SP approximation (12) with $K(t)$ described in (14). The coefficients a_i (8) are computed thanks to the ROC curves (Fig. 3). The function g_2 is computed by solving $g_1(\gamma) = P_{fa}$, which is an easy task, since $g_1(\gamma)$ is continuous (due to the SP approximation) and strictly decreasing. The constrained optimization (15) is equivalent to the following unconstrained maximization:

$$(\hat{\alpha}_1, \ldots, \hat{\alpha}_n) = \underset{(\alpha_1, \ldots, \alpha_n)}{\arg\max} \; G[P_{fa}, \alpha_1, \ldots, \alpha_n], \quad (19)$$

where $G(.)$ is defined in (18). By construction, the optimum $(\hat{\alpha}_1, \ldots, \hat{\alpha}_n)$ satisfies the global P_{fa} constraint (15). The optimization algorithm used is the interior-point algorithm implemented in MATLAB. For n=10, its computation cost is about 4 min on a 2.5 GHz Intel Core i5 processor. We note that this method is carried out offline.

B. Optimization implementation

The ROC curves of the meteorological parameters (θ_i) are provided to the local thresholds optimizer (Fig. 2). From these curves and from the given global P_{fa}, the optimizer gives the optimized coefficient $\hat{\alpha}_i$ (19) to each local detector related to θ_i and provides the global threshold $\hat{\gamma} = g_2(P_{fa}, \hat{\alpha}_1, \ldots, \hat{\alpha}_n)$ (17) to the fusion center. Then, the thresholds $\hat{\eta}_i$ (1) are computed according to $\hat{\alpha}_i$. The decisions X_i are the results of the local LR tests (2) with the measurements z_i and the thresholds $\hat{\eta}_i$. Now, the fusion center gives the final decision (9). The meteorological parameters may vary from mesh to mesh (section V-C). The optimization must then be updated, giving new local thresholds $\hat{\eta}_i$. The optimal $\hat{\alpha}_i$ for $P_{fa} = 0.1$ and $n = 10$ are shown in figure (3).

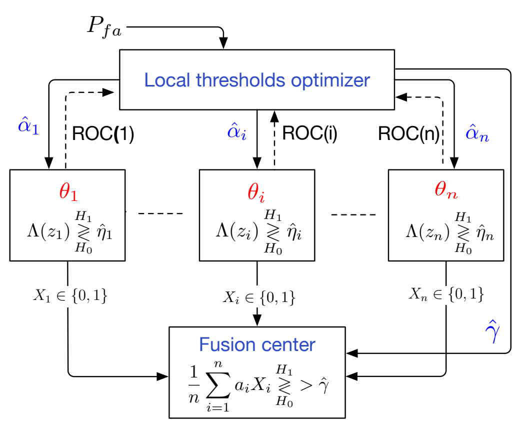

Fig. 2: Final decision implementation.

V. Experimental results

In order to illustrate the potential of the proposed method, it has been applied on a real meteorological dataset.

A. Meteorological parameters

The case of May 9, 2016 on the Gulf of Mexico was chosen. The dataset used in this application comes from the output of the numerical weather prediction model, Action de Recherche Petite Echelle et Grande Echelle (ARPEGE), developed by Météo-France in collaboration with the European Centre for Medium-Range Weather Forecasts (ECMWF) [19]. ARPEGE is a stretched global and spectral general circulation model with a horizontal resolution ranging from 7.5 km over Europe to 36 km at the opposite side of the globe in its 2015 version [20]. Numerical weather prediction models solve the equation describing the evolution of the atmosphere, with an assimilation of observations. In addition, physical parameterizations are included in this model to take into account meteorological phenomenon not resolved by the model. Thus, the model provides an estimate of the future state of the atmosphere. For this application, the estimation of the state of the atmosphere at 09:00 UTC on May 9, 2016, based on informations produced at 00:00 UTC on May 9, 2016, has been chosen.

A relevant list of 10 meteorological parameters has been compiled to characterize the state of the atmosphere from the standpoint of lightning strike risk. The following 10 parameters provide informations on the dynamical, thermodynamical or electrical state of the atmosphere. More details on theses parameters can be found in [1].

- The convective nebulosity (NEBCON) indicates the percentage convective cloud in a mesh
- The convective precipitation (PRECIP), expressed in kg/m^2, is the estimate of 3-h accumulated precipitation induced by a convective process
- The total water content (TWC) and the total ice water content (TIWC), expressed in (kg/m^2), are respectively the sum of the water (resp. ice) content of all layers of the atmosphere in a given mesh of the model
- The vertical velocity at the isotherm $0°C$ (WISO), expressed in (m/s), provides information on the dynamic component inside the cloud
- The maximum of relative humidity in the atmospheric column at a given location (RH) expressed in $\%$
- The lifted index (LI), expressed in degrees Kelvin, developed by Galway [21], characterizes the instability of the atmosphere
- The convective available potential energy (CAPE), expressed in (J/kg), is the potential energy available to an air parcel to lift up from the free convection level [22]
- The temperature at the top of the convective cloud (CTT) is expressed in degrees Kelvin
- The difference between the altitude of the bottom of the convective cloud and the altitude of the $0°C$ isotherm (ALTB) expressed in meters

In the present application, we aim to estimate thunderstorm risk. Consequently, two states are defined: H_0 is associated with a meteorological situation without thunderstorm and H_1 is a meteorological situation with thunderstorm. The first step in applying the method was to obtain the distribution of each of the meteorological parameters with and without a thunderstorm. To calculate these distributions, data have been cumulated over the Atlantic North (only meshes above the sea are considered), with ARPEGE model, during 5 years from 2012 to 2016. The distinction between the two cases (H_0 and H_1) is made using the lightning location information detected by the World Wide Lightning Location Network (noted WWLLN) [23]. The ability to discriminate between the 2 hypotheses of these distributions is illustrated in the following ROC curves.

B. Receiver Operating Characteristic curves

The ROC curves, deduced from meteorological parameters distributions, are computed as follows. Given $P_{fa}(i)$, the threshold η_i is calculated such that $\mathbb{P}_0(\Lambda(z_i) \geq \eta_i) = P_{fa}(i)$ (1). The detection probability is then obtained by computing $P_d(i) = \mathbb{P}_1(\Lambda(z_i) \geq \eta_i)$. The ROC curves of the 10 meteorological parameters, as well as the optimal $\hat{\alpha}_i$ (19), are presented in the following figure. As we can see, the parameters ALTB, RH and CTT are the least efficient to discriminate the 2 hypothesis.

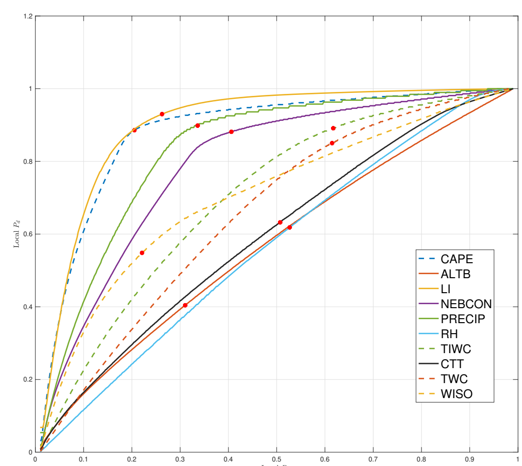

Fig. 3: ROC curves of the 10 meteorological parameters. Optimal $(\hat{\alpha}_i, \hat{\mu}_i)$ (*) for global $P_{fa} = 0.1$.

C. Thunderstorm detection performance

The geographical domain study extends from $[15°N; 40°N]$ in latitude and $[100°W; 50°W]$ in longitude, with a mesh of $0.25° \times 0.25°$. For each mesh of the grid, firstly non-physical data are removed from the dataset. The methodology is then applied. Depending on the availability of the meteorological parameters, the number of parameters used in a mesh varies between 5 and 10. As a result, local thresholds optimization (Section IV-A) must be performed for each mesh. At the output of the the global LR test (9), a binary decision is obtained for each mesh. By setting the theoretical P_{fa} to 10% (15), we obtain the results shown in figure 4. On the map, a mesh is colored in red if a risk is estimated and in grey if not. In addition, the green dots indicates the lightning flashes detected by the WWLLN. This map highlights different areas of thunderstorm risk.

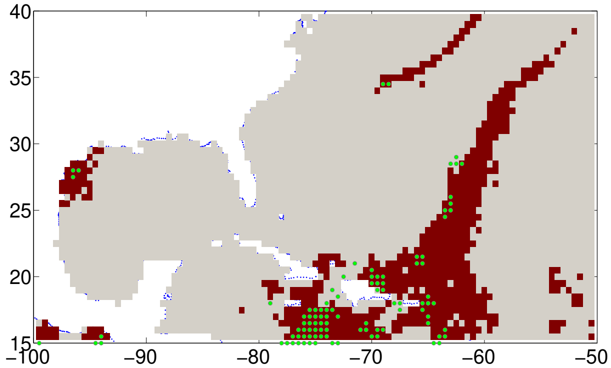

Fig. 4: Map of decision for the May 9, 2016 at 09:00 UTC. In gray, the decision is equal to 0, which means no risk of thunderstorm. In red, the decision is equal to 1, which means a risk of thunderstorm. The green dots indicate lightning strokes detected by the WWLLN.

For each mesh, two informations are compared: the lightning detection by the WWLLN and the decision obtained at the output of the methodology. The geographical domain is composed of 3655 meshes. Among them, lightning flashes have been detected in 83 meshes. Statistical results are presented in table I. By setting a global false alarm probability to 10% (15) at the input of the method, we obtain a actual global false alarm rate of 14.7% and an actual detection rate of 84.3%.

	Correct detections	False alarms
Number of meshes	70	526
Percentage	84.3 %	14.7 %

TABLE I: performance of the method on real data.

VI. Conclusion

An optimal data fusion methodology, based on the saddlepoint approximation, has been developed. The advantage of this method is to set a global false alarm probability while maximizing the global detection probability by optimizing the local thresholds related to the meteorological parameters. This method has been applied to a real dataset composed of 10 meteorological parameters extracted and derived from outputs of a numerical weather prediction model. Its application shows good performance in thunderstorm risk assessment. Nevertheless, improvements of the optimization algorithm could be done in order to reduce the computational time. In addition, further works will take into account possible correlations between the parameters.

Acknowledgments

The authors wish to thank the World Wide Lightning Location Network (https://www.wwlln.net), a collaboration among 50 universities and institutions, for providing the lightning location data used in this paper. The authors gratefully acknowledge Météo-France for providing the ARPEGE data.

References

[1] A. Bouchard, M. Buguet, A. Chan-Hon-Tong, J. Dezert, and P. Lalande, “Comparison of different forecasting tools for short-range lightning strike risk assessment,” 2022.

[2] J. Leinonen, U. Hamann, U. Germann, and J. R. Mecikalski, “Nowcasting thunderstorm hazards using machine learning: the impact of data sources on performance.,” 2021.

[3] F. Bouttier and H. Marchal, “Probabilistic thunderstorm forecasting blending multiple ensemble.,” vol. 72, pp. 1–20, 2020.

[4] Y. Yair, B. Lynn, C. Price, V. Kotroni, K. Lagouvardos, E. Morin, A. Mugnia, and M. D. C. Llasat, “Predicting the potential for lightning activity in mediterranean storms based on the weather research and forecasting (WRF) model dynamic and microphysical fields.,” vol. 115, pp. 13 p., 2010.

[5] P. K. Varshney, “Multisensor data fusion,” *Electronics & Communication Engineering Journal*, vol. 9, no. 6, pp. 245–253, 1997.

[6] S. H. Javadi and A. Peiravi, “Fusion of weighted decisions in wireless sensor networks,” *IET Wireless Sensor Systems*, vol. 5, no. 2, pp. 97–105, 2015.

[7] Z. Chair and P. K. Varshney, “Distributed Bayesian hypothesis testing with distributed data fusion,” *IEEE transactions on systems, man, and cybernetics*, vol. 18, no. 5, pp. 695–699, 1988.

[8] R. Viswanathan, CA. Thomopoulos, and R. Tumuluri, “Optimal serial distributed decision fusion,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 24, no. 4, pp. 366–376, 1988.

[9] C. Musso, E. Jay, and J. P. Ovarlez, “Saddlepoint approximation applied to fusion in multiple sensor and to detection in clutter,” in *Proc. Sixth Int Conf. Information Fusion*, 2003.

[10] S. Aldosari and J. MF. Moura, “Detection in sensor networks: the saddlepoint approximation,” *IEEE Transactions on Signal Processing*, vol. 55, no. 1, pp. 327–340, 2006.

[11] G. Vazquez-Vilar, A. G. Fabregas, T. Koch, and A. Lancho, “Saddlepoint approximation of the error probability of binary hypothesis testing,” in *2018 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2018, pp. 2306–2310.

[12] J. Neyman and E. S. Pearson, “Ix. On the problem of the most efficient tests of statistical hypotheses,” *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, vol. 231, no. 694-706, pp. 289–337, 1933.

[13] H. E. Daniels, “Saddlepoint approximations in statistics,” *The Annals of Mathematical Statistics*, pp. 631–650, 1954.

[14] S. Buitendag, *The saddle-point method and its application to the hill estimator*, Ph.D. thesis, Stellenbosch: Stellenbosch University, 2016.

[15] H. E. Daniels, “Tail probability approximations,” *International Statistical Review/Revue Internationale de Statistique*, pp. 37–48, 1987.

[16] R. Lugannani and S. Rice, “Saddle point approximation for the distribution of the sum of independent random variables,” *Advances in applied probability*, vol. 12, no. 2, pp. 475–490, 1980.

[17] J. G. Booth, P. Hall, and A. TA. Wood, “On the validity of edgeworth and saddlepoint approximations,” *Journal of Multivariate Analysis*, vol. 51, no. 1, pp. 121–138, 1994.

[18] R. W. Butler, *Saddlepoint approximations with applications*, vol. 22, Cambridge University Press, 2007.

[19] P. Courtier, C. Freydier, J. F. Geleyn, R. Rabier, and M. Rochas, “The ARPEGE project at Météo-France.,” in *Proceedings of 1991 ECMWF Seminar on Numerical Methods in Atmospherics models*, 1991, pp. 193–231.

[20] J. Pailleux, F. Geleyn, R. El Khatib, C. Fisher, M. Hamrud, J. N. Thepaut, F. Rabier, E. Andersson, D. Salmond, D. Burridge, A. Simmons, and P. Courtier, “Les 25 ans du système de prévision numérique du temps IFS/ARPEGE.,” vol. 89, pp. 18–27, 2015.

[21] J. G. Galway, “The lifted index as a predictor of latent instability.,” 1956.

[22] R. A. Jr. Holton, *An introduction to dynamix meteorology, 4th Edition.*

[23] A. R. Jacobson, R. Holzworth, J. Harlin, R. Dowden, and E. Lay, “Performance assessment of the world wide lightning location network (WWLLN) using the Los Alamos sferic array (LASA) as ground truth,” *Journal of Atmospheric and Oceanic Technology*, vol. 23, no. 8, pp. 1082–1092, 2006.