



HAL
open science

Bayesian inference of visibility in fog and smoke artificial conditions from 3D-LiDAR point clouds

Karl Montalban, Dinesh Atchuthan, Christophe Reymann, Alexandre Ravet, Paul-Edouard Dupouy, Nicolas Rivière, Simon Lacroix

► To cite this version:

Karl Montalban, Dinesh Atchuthan, Christophe Reymann, Alexandre Ravet, Paul-Edouard Dupouy, et al.. Bayesian inference of visibility in fog and smoke artificial conditions from 3D-LiDAR point clouds. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 2023, 308, pp.108666. 10.1016/j.jqsrt.2023.108666 . hal-04171913

HAL Id: hal-04171913

<https://laas.hal.science/hal-04171913>

Submitted on 26 Jul 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial 4.0 International License

Bayesian inference of visibility in fog and smoke artificial conditions from 3D-LiDAR point clouds

Karl Montalban^{a,b,c}, Dinesh Atchuthan^a, Christophe Reymann^a, Alexandre Ravet^a,
Paul-Edouard Dupouy^b, Nicolas Rivière^b, Simon Lacroix^c

^a*EasyMile, 21 Boulevard de la Marquette, Toulouse, 31000, France*

^b*ONERA/DOta, Université de Toulouse, Toulouse, 31055, France*

^c*LAAS-CNRS, Université de Toulouse, CNRS, 7, Avenue du Colonel Roche, Toulouse, 31031, France*

Abstract

3D-LiDARs are heavily impacted by a degraded visual environment (DVE) like rain, fog and smoke which limits their use for perception algorithms. The capacity to retrieve information about the environmental conditions from an embedded sensor can be an asset to improve autonomous driving performances. False positive artifacts in the point clouds caused by aerosols and hydrometeors particles tend to cause perception issues and thus need filtered out. However, those artifacts can also be used as valuable information to infer weather properties and maybe improve filters. This article proposes a Bayesian inference model which can classify discrete values of visibility using 3D-LiDAR point clouds. Gamma and Log-normal distributions are used to model the distance distributions of the noise points and the model is extended using the Random Finite Set (RFS) formalism with the Poisson and Binomial RFS models. Experiments in artificial fog and smoke conditions are presented and the classification model is trained and tested independently for each experiment. The used point clouds are extracted from specific parts of the field-of-view that can be used to generalize the proposed method to any outdoor scenario. The model shows good classification results with increased performances when the RFS extension is used.

Keywords:

3D-LiDAR, Degraded Visual Environment, Visibility estimation, Bayesian inference

1. Introduction

Light Detection And Ranging (LiDAR) is a key technology to provide 3D measurements of the environment, in the form of point clouds, and has become a critical sensor for autonomous driving [1, 2]. But LiDARs are impacted by Degraded Visual Environment (DVE) conditions [3, 4, 5, 6, 7]: the interactions between the laser beams emitted by a LiDAR and particles in the atmosphere can produce false positive points and occlusions (false negatives) [8]. This can significantly alter the quality of the point clouds, and thus the overall autonomous vehicles behavior, possibly impeding safety and availability. Denoising methods can identify some of false positives in the point clouds caused by DVE particles [9], but they could be improved by the quantitative knowledge of the DVE conditions. This knowledge is also essential to efficiently adapt a vehicle navigation, e.g. speed, braking distances, or parameters of data processing algorithms [10]. Weather monitoring platforms can provide large-scale DVE information about areas the vehicle is operating in, but it is the *local* conditions, in the direct vicinity of the vehicle, that matters. One can consider the use of weather sensors mounted on the vehicle itself (precipitation or visibility sensors) but there is high interest in estimating these

conditions from a single 3D-LiDAR sensor already embedded on the vehicle [11]. Indeed, the weather information retrieved can be directly correlated to degradation levels of the point cloud data and used to identify sensor and system limitations, which are valuable for perception algorithms. An end-goal could be to filter weather artifacts in the point clouds using denoising algorithms with parameters trained on the inferred weather properties while giving confidence levels as part of the results. In addition, as each LiDAR scan can be considered individually, it ensures redundancy and avoids cross-dependencies which is mandatory for the development of safety critical systems.

Related work. Various authors focused on the recognition of weather conditions using automotive LiDAR point clouds. [11, 12, 13, 14] suggest a statistical analysis of the water particles detections and estimate the nature of the weather conditions (clear, rain, fog, snow) using machine learning techniques, *i.e.* KNN, SVM with different feature vectors, or convolutional neural networks. The latter require large amount of training data, and lack a robust qualification of their results, which is a limitation to operational safety of autonomous driving. A probabilistic hierarchical Bayesian model is proposed in [15] to predict rainfall intensity. The model is trained using point cloud

data acquired in artificial rain and precipitation data acquired with a disdrometer to assess the rainfall ground truth. Artifacts caused by rain are used to train and test the model. The proposed method yields a rainfall measurement error of 2.89 mm/h, similar to the measurement error of the disdrometer itself. The results presented are promising, but the model might be too tightly fitted to the specific experimental setup and therefore not suitable to external conditions. Indeed, unlike realistic environments, the considered scene only includes the walls of the chamber and no parts of the field-of-view (FOV) is free of target. This impacts the resulting point clouds, as explained later in section 2.1. And, the production of rain with nozzles at high intensity is known to produce not representative rainfalls, closer to heavy rain showers than to spatially scattered rain droplets [8].

In fog conditions, [16, 17] present methods to evaluate visibility using the physical equations of light extinction described by the Mie theory and visual observations. This physics-based method is interesting but the full-wave LiDAR signals are most of the time not available on commercial automotive LiDAR systems. In [18, 19], a machine learning approach focuses on LiDAR point clouds taken in artificial fog conditions. A statistical analysis of the impacts of fog on range measurements is used to build a Gaussian process regression model. Given a certain visibility, the model predicts the range at which an object can be detected by the LiDAR. The end of the works of [12] proposes an estimation of the optical visibility during artificial fog conditions using a convolutional neural network architecture and LiDAR data. The algorithm considers the full point cloud with associated feature vectors to classify visibility from 5 m to 100 m. Their model shows relatively good performances except for visibilities above 60 m because their number of available data samples is reduced. Though, using the full point cloud makes the results highly dependent on the scene which, in the case of an artificial chamber, is quite different from outdoor scenarios. Also, no error metric is given concerning the visibility classification to compare the results.

Optical visibility or meteorological optical range, in *meters*, is defined by the International Civil Aviation Organization as "the greatest distance at which a black object of suitable dimensions, situated near the ground, can be seen and recognized when observed against a bright background" – see details in [20].

Contribution. This article describes a probabilistic framework based on a Bayesian inference model which aims to directly infer visibility from point clouds obtained by an automotive 3D-LiDAR sensor. The objective is to increase the capacity of autonomous vehicles to estimate the DVE conditions so as to yield a more robust autonomous navigation capacity. The approach is a naive Bayes supervised classification scheme. A series of visibility classes are defined by contiguous visibility distance intervals and, a labelled dataset acquired in controlled conditions consti-

tute the learning base. The dataset is composed of sets of LiDAR echoes' distances generated by DVE particles and labelled with visibility values. It is exploited to estimate the probability of each visibility class using the density likelihood distribution of the distances of these LiDAR echoes. The likelihoods are modeled as Gamma or Log-normal probabilistic distributions. Then, the Random Finite Set (RFS) formalism (with the Poisson and Binomial RFS models) accounts for the cardinality of the data and allows to capture the evolution of the number of echoes in the sets. As the number of noise point evolves regarding to changing visibility, the RFS formalism improves the inference results. At runtime, the acquired point cloud is analysed in order to detect echoes caused by the DVE conditions, and the application of the Bayes rule allows to estimate partial probabilities of each visibility class. The learning step of the optimal parameters for the different likelihood distributions is done using a Markov Chain Monte Carlo technique (MCMC) and the Metropolis-Hastings algorithm. Finally, classification results are analysed in artificial fog and smoke conditions and accuracy of the various likelihood models are compared.

Outline. The next section depicts the proposed methodology to directly infer the visibility of the environment from 3D-LiDAR data. It first presents how the data necessary to the inference is extracted from point clouds, then depicts the Bayesian formalism adopted, and finally details the various analytical probability distributions used to model the data. Section 3 presents the two different experimental setups used to train and test the approach with artificial fog and smoke conditions. Section 4 presents and analyses experimental results and, a discussion concludes the paper.

2. Bayesian inference model

2.1. Methodology

A 3D-LiDAR detects points from the laser beams it sends out if the full-wave signals, received after backscattering in the environment, contain peaks with an amplitude above a threshold [21]. Then, the sensor determines if one or multiple points from these peaks are produced (one point for each peak and depending on the multi-echo capacity). If multiple peaks are found, the order of selection is their magnitude (in most systems). In this study, only the strongest echo is considered for each laser shot. Solid targets in the path of a laser beam are likely to return higher energy signals because of their reflectivity (in very low visibility conditions, the obscurant can return more signal than a solid target at given distance and reflectivity [22]). On the other hand, DVE particles yield scattering interactions and tend to return lower energy signals. This implies that solid targets in a LiDAR frustum can prevent false positive echoes from signals with lower amplitude (e.g. caused by raindrops, fog or smoke particles).

Conversely, a frustum free of any target is more likely to produce such echoes and, here, we call this a *free-sky* frustum (see fig.1).




Figure 1: *Free-sky* frustum designates any FOV that is free of obstacle within the range of view of the LiDAR sensor. All the 3D-LiDAR echoes acquired in the *free-sky* frustum are generated by atmospheric particles.

Our approach relies on these characteristics of LiDARs to isolate the echoes generated by the DVE conditions from the ones generated by actual targets. It is indeed the density and distance repartition of the echoes generated by the DVE conditions that is representative of the visibility conditions and, on the basis from which we can estimate these conditions. To select atmospheric echoes, for both learning and run-time cases, we consider the data that belong to target-free frustums, denoted as “free-sky” frustums, that are all due to atmospheric particles.

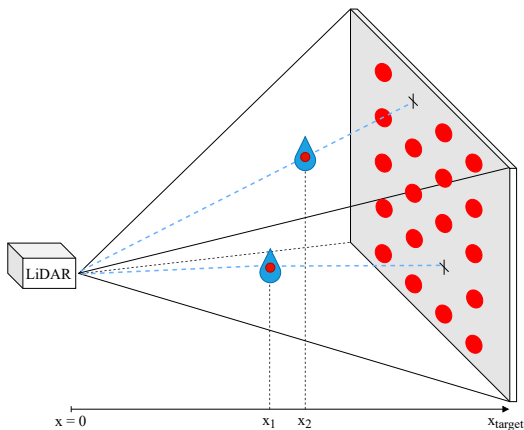


Figure 2: Schematic representation of the extraction of points on a solid target and inside the LiDAR to target frustum. x represents the distance of detected echoes (red circles) to the LiDAR sensor.

The presence and knowledge of the parts of the LiDAR FOV that do not contain any target can be ensured in autonomous driving operational conditions, in which the vehicles evolve in precisely mapped environments. Alternatively, in the absence of such frustums, the knowledge

of some targets nature and positions can be exploited to distinguish atmospheric echoes from target echoes – this is the way we proceed in one of our experiments (section 3.2 and illustrated on fig.2).

2.2. Bayesian inference

The inference model is a naive Bayes classifier [23, 24, 25]. Its goal is to infer discrete classes of visibility V from independent features – in our case, series of echoes e (fig.3). It is a supervised model because a labelled dataset \mathbf{Y} is used¹, and it is generative as it allows to generate samples of distances e from classes of visibility V or model parameters θ , as shown on fig.4.




Figure 3: PGMs of the naive Bayes classifier model for visibility classification using independent LiDAR echoes, related to eq.1. (a): developed model, (b): same model using the plate notation.

Fig.3a shows the Probabilistic Graphical Model (PGM) representation of a naive Bayes model applied to our problem where i.i.d. features $\{e_1, e_2, \dots, e_D\}$ are used to classify a visibility class V . A PGM representation gives visual insight about the causal dependencies between the random variables involved in a probabilistic model, the theory behind PGMs is given in [23]. Naive Bayes models are based on the assumption of independence between the features so that the joint model can be expressed as $P(V, e_1, e_2, \dots, e_D) = P(V) \prod_{d=1}^D P(e_d|V)$. LiDAR echoes are considered independent because the laser beams are fired independently, and each laser emitter is paired optically with a unique detector. For each visibility class, D features (echoes in our case) can be used for inference and D can vary depending on the instance of the class V .

Let $\mathbf{Y} = \{(\mathbf{E}_n, \mathbf{V}_n)\}$ for $n \in [1, N]$, be the labelled dataset which consists of distances $\mathbf{E}_n = \{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_D\}_n$ and discrete optical visibilities \mathbf{V}_n . At run-time, when operating in DVE conditions, the set of echoes distances $E = \{e_1, e_2, \dots, e_D\}$ from a free-sky sensor frustum is acquired (these points are noise artifacts created by the DVE conditions). Using Bayes’ theorem, the probability of having a visibility value V knowing the received set E and the dataset \mathbf{Y} is $P(V|E, \mathbf{Y})$:

¹**boldface** means that the characters belong to the labelled dataset, following the convention of [24].

$$P(V|E, \mathbf{Y}) = \frac{f(E|\mathbf{Y}, V)P(V)}{\sum_{i=1}^N f(E|\mathbf{Y}, V_n)P(V_n)} \quad (1)$$

where $f(E|\mathbf{Y}, V)$ is the likelihood of the echoes distances E for a specific visibility V , and $P(V)$ is the prior probability on the discrete visibility class.

Without any knowledge on the status of visibility, $P(V)$ is considered uniform. Information available from a local weather station could be included here and improve the precision of the inference. Additionally, a recursive Bayesian filter would improve inference results over time by using the posterior probability at a certain time step k as the prior probability of the next time step $k + 1$ [26].

We now introduce model parameters, as the vector θ , used in the model to parametrize the different likelihood distributions. In the following, θ remains a generic hidden variable but more information is given in section 2.3. Through marginalization over θ , $f(E|\mathbf{Y}, V)$ becomes :

$$f(E|\mathbf{Y}, V) = \int_{\theta} f(E|\theta) f(\theta|\mathbf{Y}, V) d\theta \quad (2)$$

Fig.4 adds θ to the PGM representation of the model. It shows that distance samples e can be generated from the random variable θ and $f(E|\theta)$ can be computed as in eq.5.




Figure 4: PGM of the naive Bayes classifier model for visibility classification using independent LiDAR echoes, with the addition of the model parameters θ .

The Probability Density Function (PDF) $f(E|\theta)$ from eq.2 is the likelihood of the echoes given the parameter θ . Given that each echo feature is independent, the likelihood $f(E|\theta)$ is the product of probabilities of receiving each echo following the likelihood distribution L ,

$$f(E|\theta) = \prod_{e \in E} L(e|\theta) \quad (3)$$

The second part of eq.2 is the PDF $f(\theta|\mathbf{Y}, V)$ which represents the learned model parameters of the likelihood distributions generated by the echoes observed at a given visibility V . For each visibility class \mathbf{V}_n in the training set, each set of echoes \mathbf{E}_n is paired with \mathbf{V}_n in \mathbf{Y} , so $f(\theta|\mathbf{Y}, V)$ becomes

$$f(\theta_n|\mathbf{Y}, V = \mathbf{V}_n) = f(\theta_n|\mathbf{E}_n, V = \mathbf{V}_n) \quad (4)$$

This means that θ_n only depends on the part of the training dataset \mathbf{Y} acquired at the specific visibility V_n . We use Bayes' theorem again to expose the likelihood of \mathbf{E}_n knowing θ_n ,

$$f(\theta_n|\mathbf{E}_n, \mathbf{V}_n) = \frac{f(\mathbf{E}_n|\theta_n, \mathbf{V}_n)P(\theta_n|\mathbf{V}_n)}{\int_{\theta_n} f(\mathbf{E}_n|\theta_n, \mathbf{V}_n)P(\theta_n|\mathbf{V}_n)d\theta_n} \quad (5)$$

Similarly to the PDF $f(E|\theta)$, $f(\mathbf{E}_n|\theta_n, \mathbf{V}_n)$ can be modelled by the likelihood distribution with independent events so that :

$$f(\theta|\mathbf{E}_n, \mathbf{V}_n) = \frac{\prod_{e \in \mathbf{E}_n} L(e|\theta_n)P(\theta_n, \mathbf{V}_n)}{\int_{\theta_n} \prod_{e \in \mathbf{E}_n} L(e|\theta_n)P(\theta_n, \mathbf{V}_n)d\theta_n} \quad (6)$$

Eq.(6) can then be used directly to learn the distribution of θ for each visibility class. Because the integral of the denominator does not have an analytical closed form, we use a sampling technique to produce a numerical approximation on the distribution of the parameters and best approximate $f(\theta|\mathbf{E}_n, \mathbf{V}_n)$. Markov Chain Monte Carlo (MCMC) methods are statistics algorithms for sampling from probability distributions where random walk methods create Markov Chains and appropriate distributions proportional to given functions (in our case, $\prod_{e \in \mathbf{E}_n} L(e|\theta)$ from eq.6). The Metropolis Hastings algorithm is used with a Gaussian proposal density and a rejection rule [23]. The results consist in a collection of learned parameter samples for each visibility instance. To use these samples and compute the probability of a visibility class (eq.1 and 2), we use Monte Carlo integration [27]. This technique allows to compute an integral using random discrete samples at which the integral is evaluated. In our case, random samples of θ are taken from the sampled representation of $f(\theta|\mathbf{Y}, V)$ (learned with the MCMC method) and injected into $f(E|\theta)$ so that eq.2 can be calculated.

The schematic diagram of fig.5 shows the overall framework with both learning and testing phases applied to visibility classification. The learning block represents the approximation of the parameters distribution as described previously. During the test phase, frustum echoes E are received. Using eq.2, the likelihood probability $f(E|\mathbf{Y}, V)$ of each class V is calculated with Monte Carlo integration using the learned parameters θ for this class and the product of likelihood (eq.3). Finally, eq.1 is used to compute the probability of each visibility class using Bayes' theorem.

2.3. Likelihood distributions

This section describes the different likelihood distributions used in the model to compute the probability of eq.3.

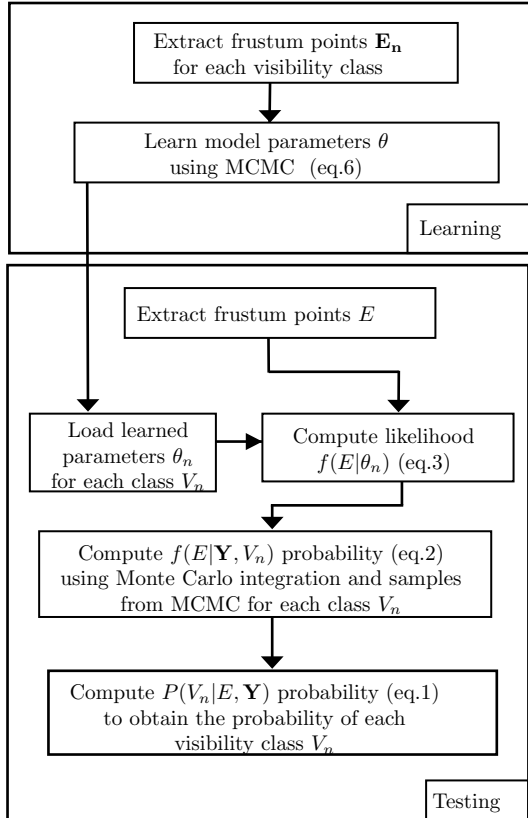


Figure 5: Schematic diagram of the Bayesian inference framework.

Gamma and Log-normal distributions are used to represent the statistical repartition of our echoes' distances. The Random Finite Set (RFS) formalism with the Poisson RFS and Binomial RFS is then introduced to extend the initial likelihood with the capture of the cardinality of the sets containing the echoes. For each combination of likelihood, the vector θ parametrizes $f(E|\theta)$.

2.3.1. Gamma and Log-normal distributions

The Gamma distribution is introduced in the context of autonomous vehicles sensors in [28] to model the distance repartition of echoes from LiDAR point clouds in snow-fall conditions. No physical explanation is given for this choice, but it is showcased that it fits well the data. Also under snow conditions, [29] exposes similar distance distributions of echoes, but suggests to use the Log-normal distribution. They also state that the shapes of the distance distributions arise from the product between an "optical detection" function and a "building shielding effect" because the sensors are located at a building window. [22] studies the properties of light reflected from fog-particles and objects captured by single-photon avalanche photodiode detectors (SPAD) in a 3D-LiDAR design. They showcase that the time profiles of fog and objects laser returns respectively have Gamma and Gaussian distributions. This allows them to extract real targets inside fog and reconstruct the 3D scene.

Considering the application of laser pulses in fog or smoke conditions, the resulting signal on the LiDAR detector is a combination of the laser pulse shape, considered as Gaussian, an atmospheric extinction function, which can be modelled as a decreasing exponential function and, of the detector response. This yields a product similar to a Gaussian function restricted to the positive domain and, therefore with a longer decreasing tail. In the absence of a more refined model and according to the literature, both Gamma and Log-normal distributions seem valid candidates to model the distribution of fog or smoke echoes returned by a LiDAR. The PDF of a noise echo at range x which follows a Gamma distribution with the shape-scale parametrization $\theta = (\gamma, \beta)$ is expressed as :

$$L_G(x) = \frac{x^{\gamma-1} \exp(-x/\beta)}{\Gamma(\gamma)\beta^\gamma} \quad (7)$$

with $\Gamma(\gamma)$ the Gamma function evaluated at γ and γ is the shape parameter and β is the scale parameter. The PDF of a noise echo at range x which follows a Log-normal distribution is parametrized by $\theta = (\mu, \sigma)$ with mean μ and standard deviation σ and is defined by :

$$L_{LN}(x) = \frac{\exp(-(\ln(x) - \mu)^2/2\sigma^2)}{x\sigma\sqrt{2\pi}} \quad (8)$$

2.3.2. Random Finite Sets

The Random Finite Set (RFS) theory is a mathematical framework originated from the Point Process theory [30] which is used in information fusion and multi-target detection and tracking [31]. A RFS $X = \{x_1, x_2, \dots, x_n\}$ is a random variable which takes values as unordered finite sets. The purpose of an RFS is to capture the probability state of elements contained in a set as well as the cardinality of the set, noted $|X| = n$. The cardinality is random and modelled by a discrete distribution $\rho(n) = P\{|X| = n\}$. The elements in the set are modelled by $L_\theta(x_1, x_2, \dots, x_n)$, the joint distribution of the elements of X , parametrized by θ . Common probabilistic descriptors apply to it, such as the PDF. Following Mahler's approach [32] and finite set statistics, the PDF of a RFS is described as follows:

$$f(\{x_1, x_2, \dots, x_n\}) = n! \cdot \rho(n) \cdot L_\theta(x_1, x_2, \dots, x_n) \quad (9)$$

Our Bayesian inference model can then be extended, using the RFS formalism, by taking into consideration the number of echoes contained in our point clouds, in addition to the shape of the distribution (Gamma or Log-normal). We now consider our set of echoes E from section 2.2 as an RFS :

$$E = \{e_1, \dots, e_n\} \quad (10)$$

where a sufficient number of echoes $n > n_{lim}$ is received and e is an echo distance. The following describes respectively the Poisson, Bernoulli and Binomial RFS models which are later used for the inference model as likelihood

functions $L(e|\theta)$ of eq.3. Their discrete cardinality distribution is first given before the overall PDF is deduced from Eq.9 [33].

- Poisson RFS

The Poisson distribution is a discrete probability distribution which describes the probability of a given number of events occurring in a fixed interval of time or space, parametrized by the Poisson rate l . The Poisson RFS is the most common point process for spatial point patterns, it uses the Poisson distribution as a distribution over its cardinality. The number of LiDAR echoes received from weather particles can be modelled using a Poisson distribution. The cardinality of a Poisson point process is Poisson distributed with rate l and likelihood distribution of the model L_θ . In this case, the model parameters θ initially coming from Gamma or Log-normal distributions are now extended with the Poisson rate l . If $E \sim \text{Poisson RFS}(l, L_\theta)$, E is described as :

$$\begin{cases} |E| \sim \text{Poisson}(l) \\ e \text{ in } E \text{ are i.i.d. with likelihood distribution } L_\theta(e) \end{cases} \quad (11)$$

The Poisson cardinality distribution is described as :

$$\rho(n) = \frac{\exp(-l) l^n}{n!} \quad (12)$$

Given eq.9, the PDF of a Poisson RFS E is :

$$P(E|l, \theta) = \exp(-l) l^n \prod_{i=1}^n L_\theta(e_i) \quad (13)$$

- Bernoulli RFS

The Bernoulli distribution is a discrete probability distribution which describes the set of possible outcomes of any single experiment that asks a “yes-no” question. The number of LiDAR echoes received from weather particles can be modelled using a Bernoulli distribution, where each laser shot is a Bernoulli trial leading to a receiving echo or not. A Bernoulli RFS E is parametrized by the probability of success r and its cardinality distribution follows a Bernoulli distribution. E can be empty with probability $1 - r$ or a singleton with probability r . In the latter case, it contains the distance e of the returned echo following a likelihood distribution L_θ .

$$P(E|r, \psi) = \begin{cases} 1 - r & E = \emptyset \\ r \cdot L_\theta(e) & E = \{e\} \end{cases} \quad (14)$$

- Binomial RFS

By repetitively sending independent laser shots and receiving echoes (or not), we repeat Bernoulli trials. The binomial distribution is a discrete probability distribution which describes the number of successes in a sequence of independent Bernoulli trials, which in our case are the independent laser shots. So, the Binomial RFS is an extension of the Bernoulli RFS, where its cardinality distribution is described by the binomial distribution with

parameters m (number of binary experiments), n (number of successes) and r (the probability of success of each trial). It is defined by :

$$\rho(n) = \binom{m}{n} r^n (1 - r)^{m-n} \quad (15)$$

Let m be the number of laser shots fired towards a target and $E = \{e_1, e_2, \dots, e_n\}$ be a Binomial RFS containing our set of n returned echoes’ distances from frustum detections, the PDF of E is :

$$P(E|r, \theta) = \binom{m}{n} r^n (1 - r)^{m-n} \prod_{i=1}^n L_\theta(e_i) \quad (16)$$

In this case, the vector of parameters θ is extended with the success probability r .

3. Experimental setups

Two distinct experiments in laboratory conditions have been carried out. The goal is to assess degradation in the 3D-LiDAR data against DVE conditions and infer DVE properties using the distance distributions of frustum LiDAR echoes and the model described previously. Fog and smoke artificial conditions are produced in climatic chambers and, context sensors are used to measure the evolving degraded conditions and label the data. The next sections depict the experimental setups carried for both experiments with the description of the LiDAR sensor, conditions produced, targets, tests protocols and context sensors. In each test, the chamber is first saturated with the DVE particles and data is recorded until full dissipation.

3.1. LiDAR sensor

The LiDAR sensor used for this study is the Ouster OS1-128 spinning 3D-LiDAR. It uses 865 nm wavelength vertical cavity surface emitting laser (VCSEL) and SPAD detectors. Two columns of 128 emitters and detectors spinning around a vertical axis provide the 360° field of view with 128 point layers. The minimum range of the sensor is around 0.25 m.




Figure 6: Targets layout inside the CEREMA chamber. These targets remain static for all conducted tests in the room.

3.2. Artificial fog experiments

The fog experiments have been conducted in the CEREMA² laboratory. The climatic chamber is a 30×10 m rectangle facility able to produce two types of fog conditions using particle diameter modes centered at $1 \mu\text{m}$ and $10 \mu\text{m}$. No additional information is known about the fog particles and in the following, we will refer these two types of fog respectively as *small* fog and *big* fog.

The arrangement of the scene for this campaign is static and shown on figure 6. The different targets are listed on table 1 along with their distance from the LiDAR. *Back* target is a portion of the end of the chamber and is the target used similarly to a *free-sky* target. The other targets are Lambertian with calibrated reflectivity.

Target	<i>back</i>	<i>a1</i>	<i>b1,2,3</i>	<i>c1,2,3</i>
d (m)	29	23	17	11

Table 1: Targets distances in the CEREMA climatic chamber.

Once the climatic chamber is saturated with homogeneous water-fog particles, the measured visibility reaches its lowest value. Dissipations are then performed to acquire data while visibility increases until complete dissipation of fog. A calibrated transmissiometer sensor operated by the CEREMA team is used to measure visibility inside the chamber at 1 Hz, a resolution of 1 m with a minimum visibility value of 5 m.

3.3. Artificial smoke experiments

The smoke experiments have been conducted in a climatic chamber owned by ONERA, shown on figure 7. The facility is a 15×5 m rectangular room in which it is possible to produce artificial smoke conditions. The smoke is produced by a fog-oil smoke machine with properties described in table 2, as measured in [34].

Type of aerosol	oil-fog
Aerosol size distribution	Log-normal
Modal radius	$0.18 \pm 0.01 \mu\text{m}$
Complex refractive index	$1.508 + i10^{-5}$

Table 2: Microphysical properties of the artificial smoke particles generated during the smoke experiment.

Fig.7 displays the experimental setup with and without smoke. The smoke study uses 3 reflectance calibrated targets of 1m^2 each, namely *a1*, *a2* and *a3* for respectively the 80%, 50% and 10% reflectance values (from left to right on fig.7a). A target called *free-sky* is added, it represents the empty space defined by the rear doors at the end of the chamber. When the rear side of the chamber is open, the

laser beams shot by the sensor inside the *free-sky* frustum do not reach any solid targets.

The used methodology is the same as presented in 3.2. The *free-sky* target is used with the same purpose as the *back* target in the fog experiment as a higher level of frustum points is expected for beams shot in this direction. Concerning the monitoring of the smoke conditions, dissipations are performed. The chamber is first saturated with smoke particles while the doors are closed. Dissipations are then performed by opening the rear-doors while data are recorded until the smoke particles are completely dissipated. In total, 7 dissipations are performed with targets *a1*, *a2* and *a3* located at different distances from the OS1 LiDAR, as summarized in Table 3.



(a) Sensors and targets layout in the artificial fog experimental setup. At the rear-side of the chamber are disposed the *a1*, *a2* and *a3* targets in this order (decreasing reflectivity order). At the entrance, from left to right : the PSD sensor, the OS1-128 LiDAR, a Biral VPF700 scatterometer and the smoke machine.



(b) Example of smoke dissipation performed by opening the rear-doors. Note that the smoke is not uniformly distributed in the chamber during dissipation and evacuates mainly by the top of the chamber.

Figure 7: Experimental setup used during the artificial smoke dissipation experiments at ONERA.

Two context sensors providing visibility measurements have been used in this campaign. The first sensor is the Biral VPF700 scatterometer measuring visibility at 0.1 Hz within 10 m to 75 km with a resolution of 10 m. Its characteristics are not suited for measuring short visibility ranges generated during this campaign. The second sensor is the Lighthouse Handheld 3016 particle size diameter (PSD) counters. It can simultaneously measure the concentra-

²Centre d'Etudes et d'expertise sur les Risques, l'Environnement, la Mobilité et l'Aménagement", Clermont-Ferrand, France.

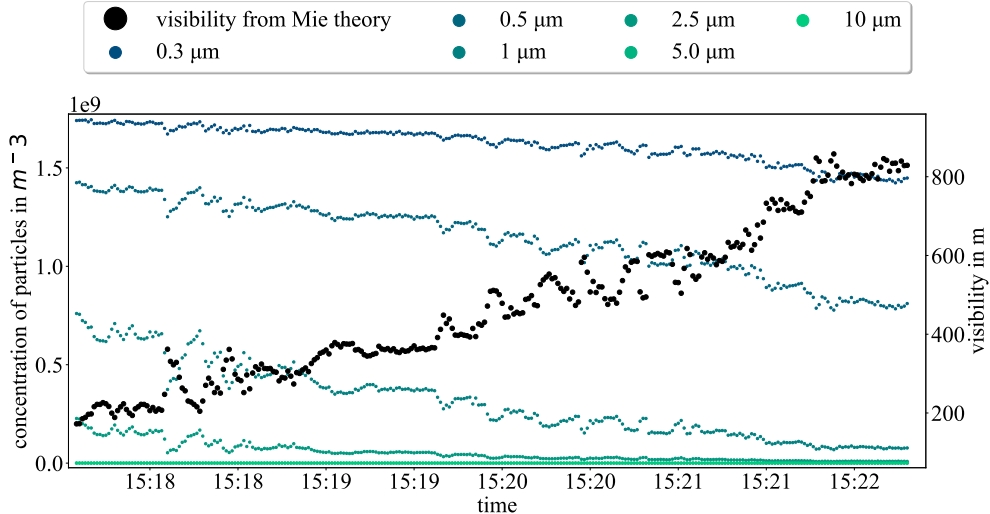


Figure 8: Example of conversion of PSD data into visibility using eqs.17 and 18 for one dissipation which gives an estimation of the evolution of visibility at 1 Hz recording frequency. Although the visibility values can seem to be too high, only its evolution is of interest for the proposed method.

Target	<i>a1</i>	<i>a2</i>	<i>a3</i>	<i>free-sky</i>
$d_1(m)$	12.4	13.3	13.4	∞
$d_2(m)$	10.9	11.1	11.0	
$d_3(m)$	9.4	9.4	9.3	
$d_4(m)$	7.1	7.2	7.2	
$d_5(m)$	5.5	5.5	5.5	
$d_6(m)$	4.1	3.8	3.7	
$d_7(m)$	1.5	1.4	1.3	

Table 3: Targets positions for the smoke experiments. Targets are moved manually before each dissipation. Note that, by definition, the *free-sky* target does not have any defined distance.

tion of particles according to size channels of 0.3, 0.5, 1.0, 2.5, 5.0 and 10 μm at a max frequency of 1 Hz. This study relies only on the PSD data to give context about the severity of the smoke conditions. Moreover, Mie theory about scattering allows to convert PSD counts into a visibility value [35]. Visibility is related to the atmospheric extinction coefficient by :

$$V = \frac{3.912}{\sigma} \quad (17)$$

Using Mie scattering formalism, a polydisperse distribution of particles gives σ_p as eq. 18 where r is the size of the particles counted, Q_{ext} is the scattering cross section and $n(r)$ is the continuous function representing the concentration of particles, given by the PSD sensor.

$$\sigma_p = \sum_r \pi r^2 Q_{ext} n(r) \quad (18)$$

In the case of visibilities below a few kilometers, the effect of the air molecules is considered negligible and the extinction coefficient is caused only by the contribution of atmospheric particles σ_p ($\sigma \sim \sigma_p$). An example is shown

on fig.8, where each PSD measurement during a dissipation of smoke is converted into a visibility value.

4. Experimental results

This last section showcases the results obtained from the two campaigns successively. Measurements highlighting the evolution of the weather conditions are given as well as statistical results on the point clouds data and, visibility classification results using the various inference models.




Figure 9: Visibility evolution measured every second during the dissipation in *small fog* and *big fog* conditions. The difference between these evolutions is explained by the difference of size distributions of the generated particles. Smaller and lighter particles tend to stay longer in the atmosphere.

4.1. Experimental results in artificial fog

The results shown below are drawn from the experiment described in section 3.2.

4.1.1. Evolution of visibility during fog dissipations

Fig.9 shows the evolution of the visibility measured by the transmissiometer during both *small* and *big fog* dissipations. For both tests, the trend of visibility is not linear

but increases as visibility rises and fog dissipates. As a result, the number of point clouds (recorded concurrently) corresponding to specific visibility values varies and can be low when visibility rises fast. This affects the classification results, as shown further. In addition, the *small* fog test shows a constant visibility of 5 m (minimum measurement of the transmissiometer) at the beginning of the dissipation as well as a slower rise than *big* fog. The visibility measured during the *big* fog test rises faster because bigger particles are heavier and fall faster.

4.1.2. Impacts on the point clouds

Following the test protocol and methodology described section 2.1 and using the *back* target, the corresponding detections in the point cloud as well as the associated frustum detections are extracted. The probability of detection (POD) of target or frustum points, is defined by the ratio between the number of detections and the number of laser beams shot towards the target.

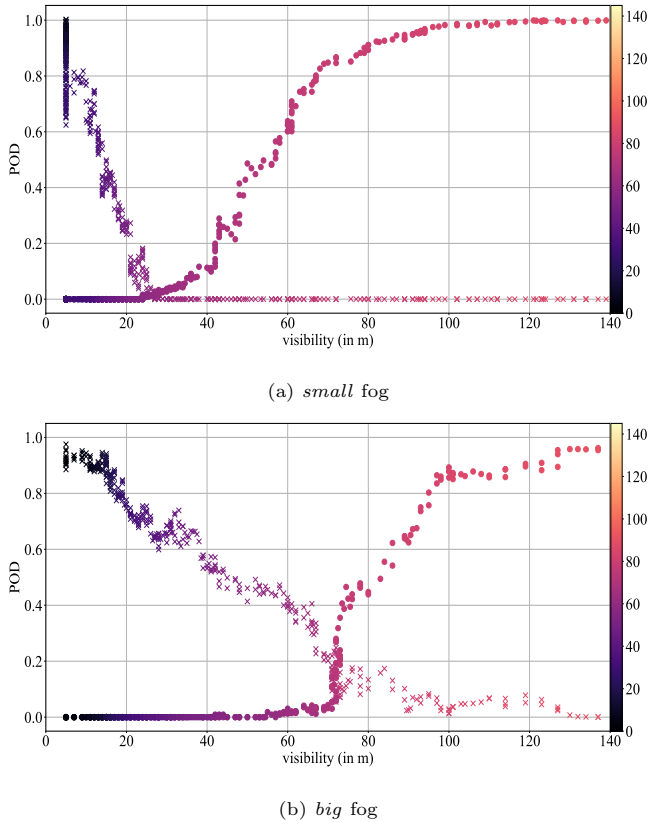


Figure 10: POD of *back* target (circles) and associated frustum (crosses) over visibility during both fog dissipations. Color refers to visibility in meter.

Figure 10 shows the target and frustum POD evolution for the *back* target, respectively for the *small* and *big* fog experiments. As fog dissipates and visibility rises, frustum POD decreases as target POD increases. While the POD of the *back* target is comparable to a sigmoid function for both tests, the frustum detections do not show a similar

behaviour. Thus, there is no clear trivial relation between the target POD and the frustum POD. Furthermore, the noise levels decrease as visibility rises but the *small* fog test shows a steeper slope, reaching close to 0 frustum detections at approximately 25 m of visibility. The constant visibility values of 5 m at the beginning of the *small* fog dissipation (Fig. 9) gives high variation in frustum POD (Fig. 10a) which forces us to consider only the end of the 5 m period for the labelled dataset. Besides, both tests show a short constant level of frustum POD values (0.9 for *big* fog and 0.8 for *small* fog) when visibility is very low.

These frustum noise points are used to perform the inference on visibility. Finally, it is noted that bigger fog particles require higher visibility values for the LiDAR to start detecting the target (around 70 m for *big* fog and 30 m for *small* fog). The next section concerns the distance distributions of noise detections inside the *back* frustum.

4.1.3. Distance distributions of the noise detections

An intuitive way to visualize the labelled dataset (\mathbf{Y} from 2.2) on which we perform the inference is displayed on the small-multiple graph of fig.11a and 11b, with respectively the results for *big* and *small* fog tests. Each sub-graph represents, in the form of a histogram, the distance distribution of frustum noise detections when the 3D-LiDAR aims at the *back* target according to specific values of visibility (here with classes of 5 m). It first shows that these detections are located at very close-range of the sensor, around 0.3 m. This close proximity is convenient for our application of inference because it is unlikely that a real object is located at this short distance. In operational conditions, these points could be filtered out and used for inference. Then, the mean distance of the echoes tend to shift to a shorter distance with increasing visibility (from almost 0.4 m to 0.3 m), especially for the *big* fog test. Finally and as expected from fig.10, the amplitude of the histograms decreases with visibility. The graphs also show how the Gamma and Log-normal distributions fit the frustum distance histograms using the parameters obtained in the learning phase for each visibility class.

4.1.4. Learning and classification results

For both fog tests, classes of visibility start from 5 m, the minimum value of the transmissiometer and the classes range is set to 5 m which ensures enough data to train and test the model on each class. The last visibility class is limited by the number of frustum points available at higher visibility values, resulting in 100 m for *big* fog and 30 m for *small* fog.

Learning the optimal parameters of the likelihood distributions is then performed with this class configuration using the method described at the end of 2.2. The training results for the Gamma and Log-normal optimal distributions are depicted on fig.11a and 11b where for each class of visibility, Gamma and Log-normal plotted curves are parametrized with the trained parameters and correctly

fit the distance histograms. The dataset is split into a training-set (2/3 of the scans) and a test-set (remaining 1/3 of the scans). The model depicted in section 2 shows that two couples of combinations of likelihood distributions (Gamma or Log-normal) or RFS models (Poisson or Binomial) can be used to perform the inference.

Likelihood model	Gamma	Log-normal
No RFS model	1.8 m 6.4 m	1.8 m 6.3 m
Poisson RFS	1.5 m 5.4 m	1.4 m 4.9 m
Binomial RFS	1.7 m 6.0m	1.3 m 5.4m

Table 4: RMSE values for each model combination and fog test, **orange** for *big* fog and **blue** for *small* fog. Better RMSE values in the small fog cases can be explained by differences in the dataset. Given the resolution of the reference visibility sensor (1 m), no significant improvement is noted in using one RFS model over another when looking only at the RMSE.

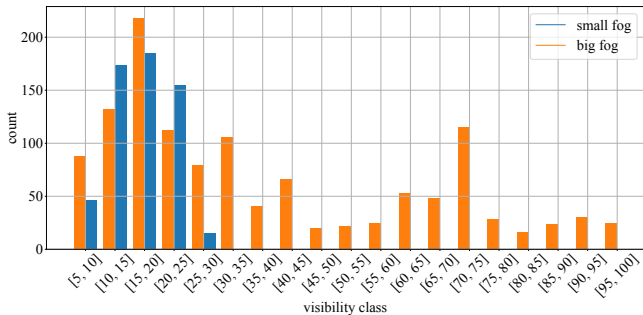


Figure 12: Number of LiDAR scans used for the test phase for each visibility class, representing 1/3 of the total dataset. We note that the number of scans per visibility class is not homogenous. Furthermore, the *small* fog experiment only generated frustum detections for lower visibility values (up to 30 m), thus limiting inference possibilities to those classes.

To quantify the performance of the inference model, we use confusion matrices and the Root Mean Square Error (RMSE) metric. Results from each configuration are given and summarized in table 4. Fig.13 and 14 give, in the form of confusion matrices, the classification results for each combination of likelihoods on the whole dataset and for the two types of fog. A confusion matrix is often used to give the results of a supervised learning algorithm where each row and column represents respectively the instance of a ground truth class and the instance of the inferred class. To compute a RMSE from the confusion matrix, the following equation is used :

$$RMSE_V = s \sqrt{\sum_{i=1}^N (V_i - \hat{V}_i)^2 / N} \quad (19)$$

where V_i is the ground truth visibility class id, \hat{V}_i is the

inferred class id, s is the visibility step of each class and N is the number of classes.

As shown on fig.12, the number of 3D-LiDAR scans within the test dataset representing 1/3 of the total dataset and thus the number of LiDAR points available for the classification differ for each class. The classification results are mitigated by this non-uniformity in the dataset. Considering this non-uniformity, the 1m resolution of the transmissiometer and the 5 m range of our visibility classes, our model remains robust and shows relatively good performances for both fog tests and with any combination of likelihood, with RMSE values from 4.9 m to 6.4 m. The Log-normal distribution seems to hold the best performances for modelling the distance distributions of the frustum echoes. There is no clear preference between the Poisson or Binomial RFS models but the capture of the cardinality induced by the RFS extension generally improves the overall classification results. For some classes, the model results in a null probability to infer the right class. These failures should be investigated closely but given the low amount of data available for this study and the inference of neighboring classes in these worst cases, the proposed methods still remain relevant.

The *small* fog test shows the best results in terms of RMSE but has less classes with relatively high number of available LiDAR scans (fig.12) compared to *big* fog. However, low number of available scans to test the model does not always cause low classification results. For example, the class [25, 30] for *small* fog has a very small number of scans but is perfectly classified by the model. This may be explained by clear differences in distance histograms from one class to another in Fig. 11b.


Although the classification results show relatively good results, it would be interesting to improve the model to learn how to differentiate between the two types of fog because *big* and *small* fogs cause quite different impacts on the point clouds. This is outside the scope of this paper but future works should include improving the inference model by adding a variable corresponding to the fog type.

4.2. Experimental results in artificial smoke


The results shown in this part come from the smoke experiments described in section 3.3. The context sensors (visibility and PSD) used for this campaign make the creation of the labelled dataset not as straightforward as for the fog experiments described above.

4.2.1. Context data for the smoke experiments

Measurements from the context sensors of the smoke experiments are available on fig.15. The graph shows simultaneously the concentrations of particles measured by the PSD sensor in color as well as the converted visibility value using the PSD data and Mie theory, in black. In addition, timestamps corresponding to the recording times of the LiDAR data are shown with the light purple vertical bands. These time intervals correspond to the dissipation



(a) *small fog*



(b) *big fog*

Figure 11: Distance histograms of frustum fog points detected with the *back* target for each visibility class in both fog conditions. Color refers to visibility in meter. Trained Gamma and Log-normal distributions fit the histograms.

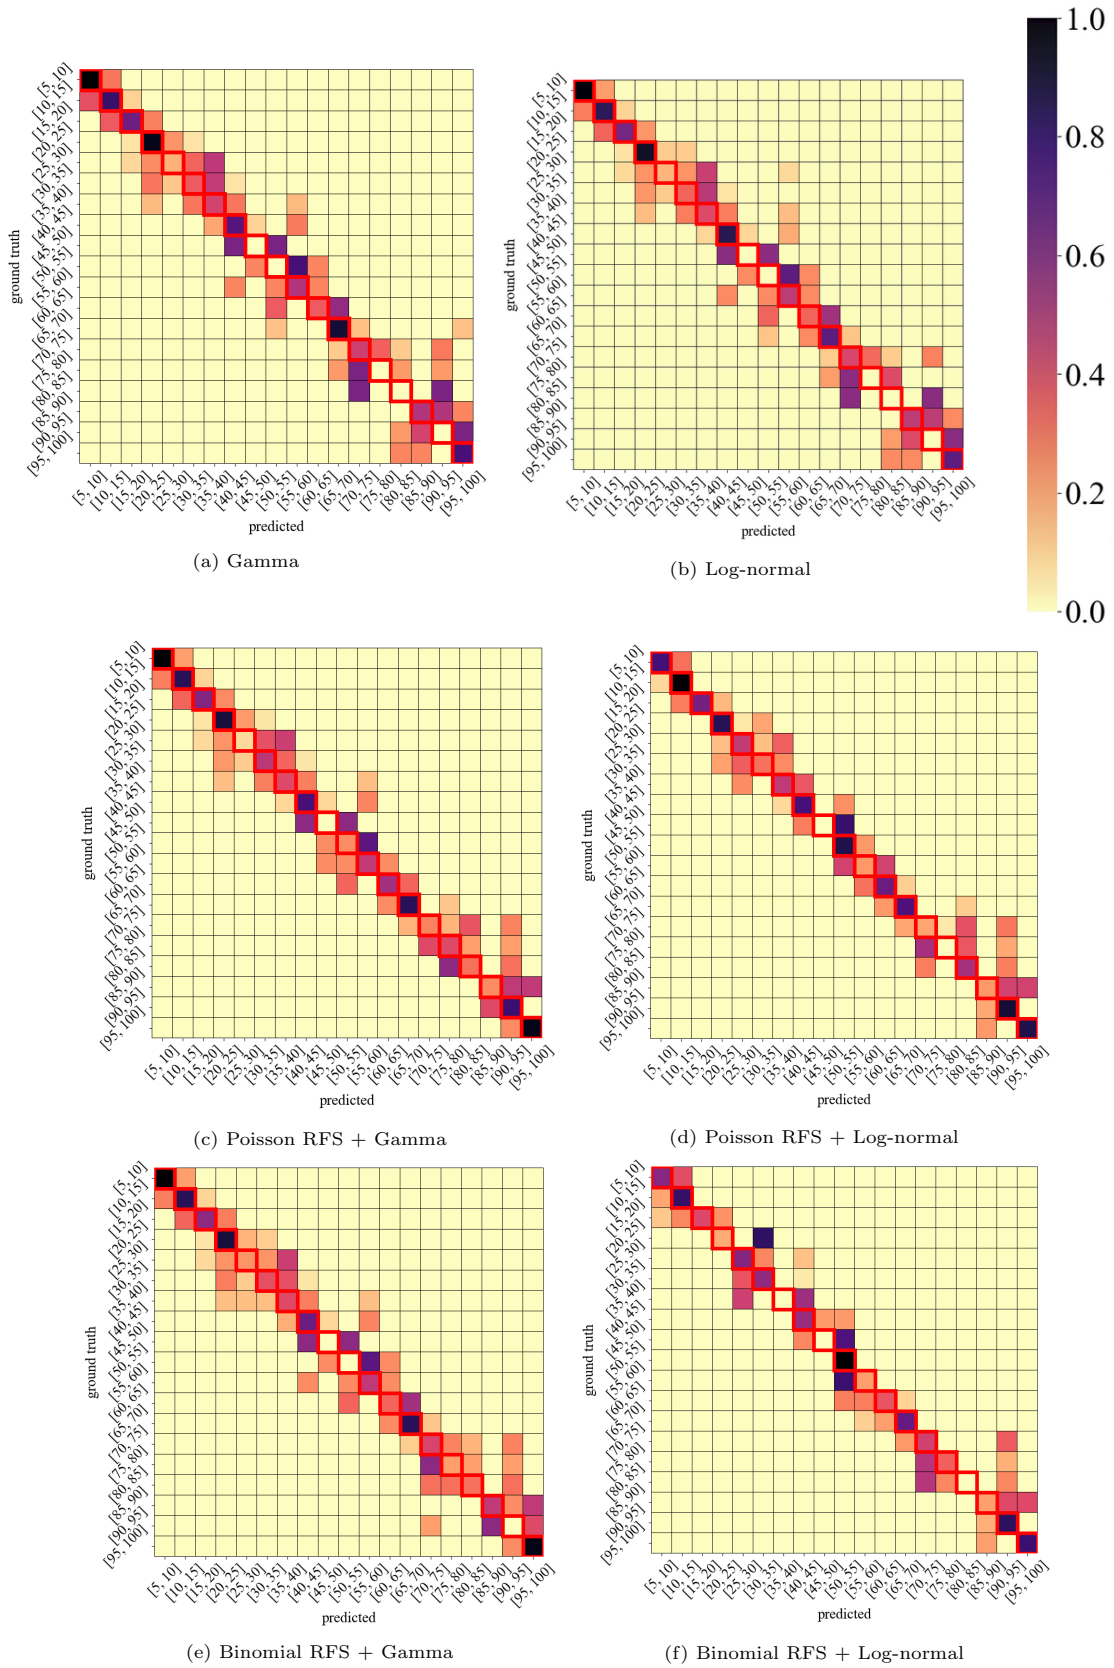


Figure 13: Confusion matrices from visibility classification using the *back* target in *big* fog conditions and the multiple combinations of likelihoods. Color represents the probability of true positive classification. Classes of 5 m visibility range are chosen for both training and testing as a first approximation of sufficient resolution for AV scenarios.

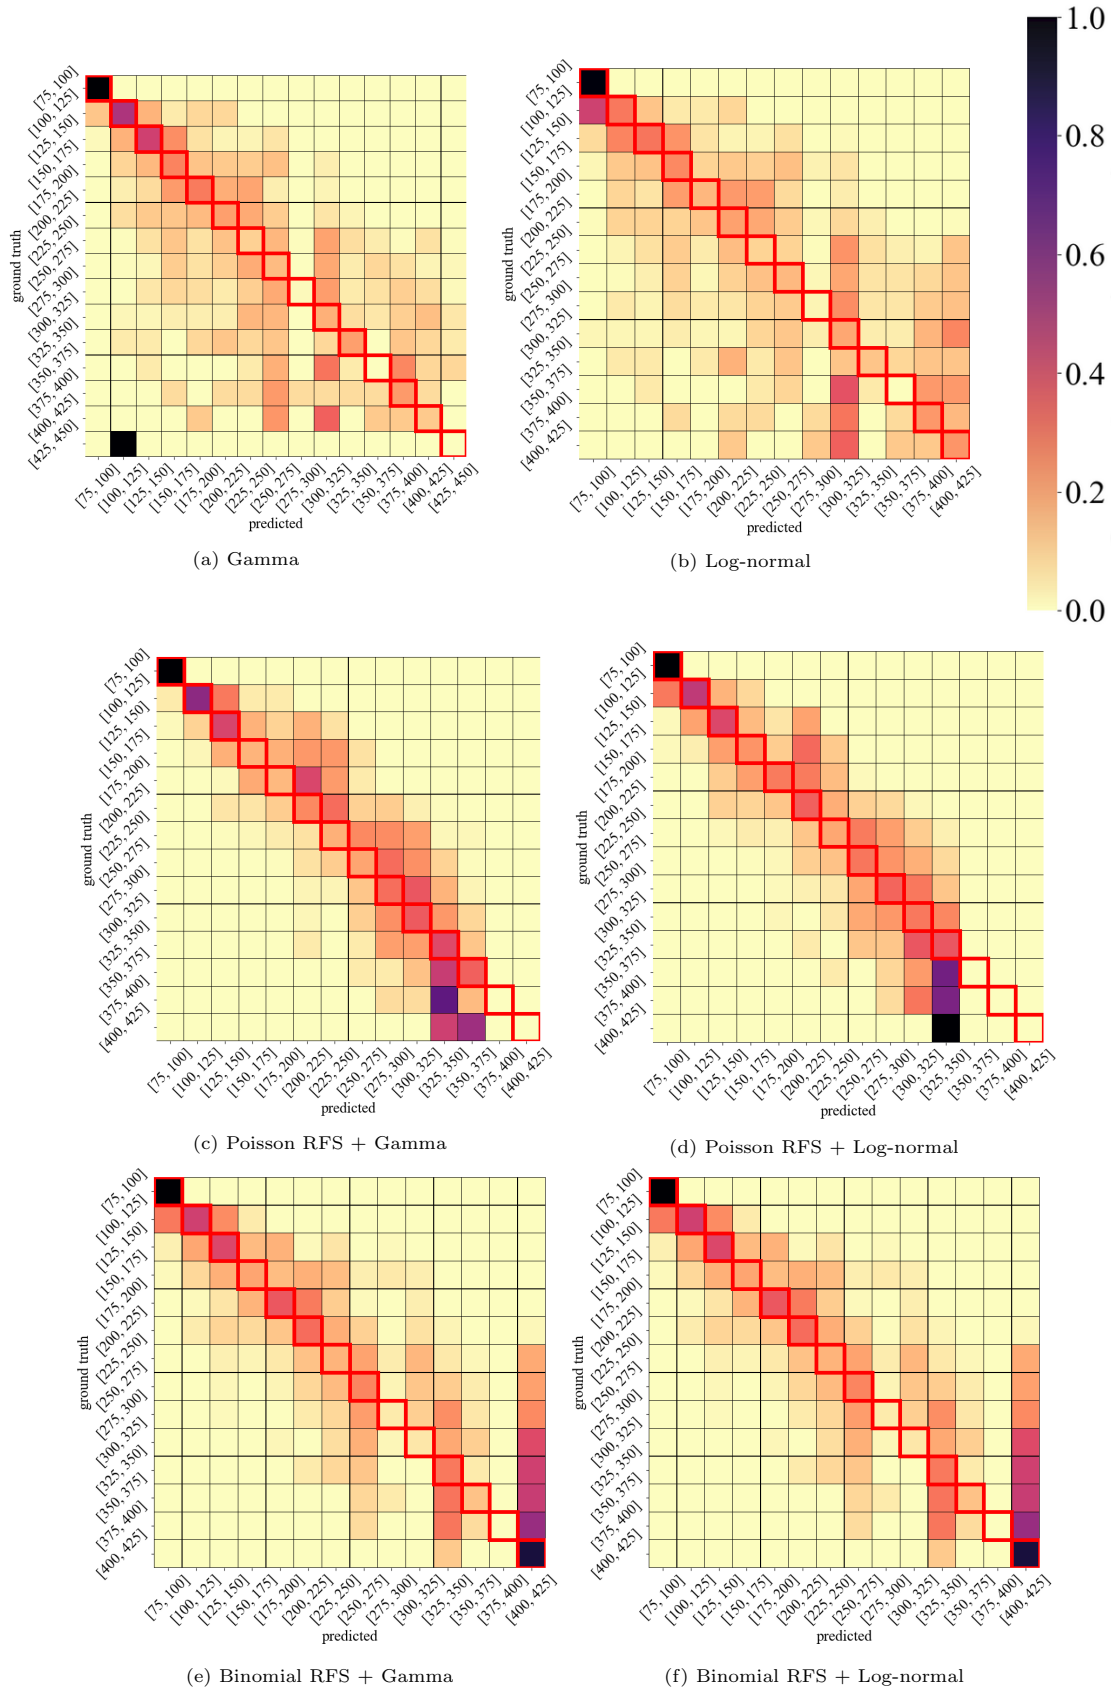


Figure 14: Confusion matrices from visibility classification using the *back* target in *small* fog conditions and the multiple combinations of likelihoods. Color represents the probability of true positive classification. Classes of 5 m visibility range are chosen for both training and testing as a first approximation of sufficient resolution for AV scenarios.

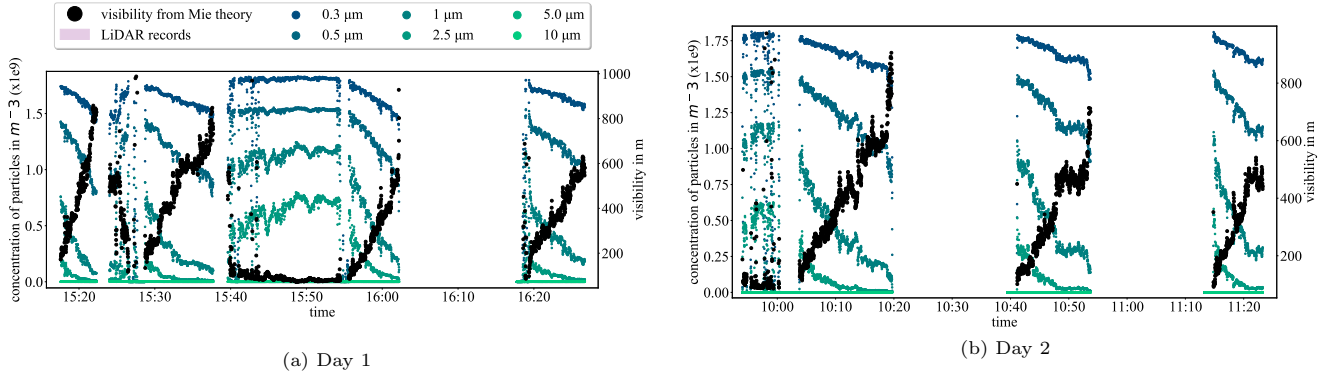


Figure 15: Measurements from PSD sensor (left y-axis) and conversion into visibility (right y-axis) with synchronized recording times of LiDAR point clouds for all performed dissipations. Although the dissipations are not fully controlled and changes in the dynamics of the visibility evolutions are observed, the particles distributions and computed Mie visibilities show similar behaviours.

durations and last approximately 5 minutes each. This is significantly shorter than for the fog experiments performed at CEREMA, which results in less LiDAR data to build the labelled dataset. This short duration originates from the test protocol performed and characteristics of the facility. Indeed, to obtain free-sky frustum LiDAR points as explained in section 2.1, the rear-doors of the chamber are opened at the beginning of the recordings. This leads to faster dissipation of the smoke particles, as opposed to the fog protocol where the chamber remains closed.

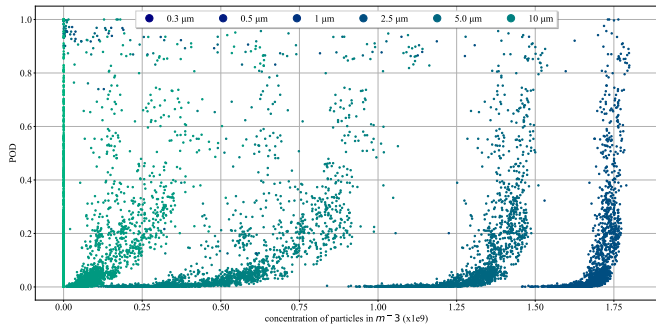


Figure 16: *free-sky* frustum POD according to each particle concentration PSD data. The PSD sensor is able to consider particles ranging from $0.3\ \mu\text{m}$ to $10\ \mu\text{m}$. This summarizes the overall particles distributions across recorded data in artificial smoke conditions with a direct relation to OS1-128 LiDAR noise detections (through the frustum POD).

Results are separated according to the test day. As fog dissipates, visibility rises but the concentrations of all particles between $0.3\ \mu\text{m}$ and $10\ \mu\text{m}$ decrease. Visibility computed with Mie theory results in values from approximately 70 m to a maximum of 980 m. These results are quite different than the ones from the fog experiments, with visibility values higher by one order of magnitude. And these high visibility values do not match with either visual observations or impacts on the point clouds. This should be investigated but it is not a limitation to evaluate our methodology nor the performances of the inference model so the following shows the results applied to these

Mie-converted visibility values.

4.2.2. Impacts on the point clouds

As said in 3.3, the smoke experiments rely on the PSD data to first obtain visibility values and then perform inference. As a result, the following shows the correspondence between PSD data and LiDAR *free-sky* frustum points with which we constitute the labelled dataset for the smoke experiments. Fig.16 displays, for all recorded dissipations, the resulting *free-sky* POD according to the concentrations measured by the PSD sensor for each particle size. Across all dissipations, a high concentration of particles leads to a high level of noise in the 3D-LiDAR point clouds and inversely. An exponential tendency is observable, with an increased slope with smaller particle size.

Then, the evolution of the *free-sky* POD according to the converted visibility is displayed on fig.17. This graph shows the results for each dissipation separately to highlight the differences between the dissipations. A common decreasing exponential tendency is observed but disruptions in the evacuation and homogeneity of the smoke particles already discussed cause different results. This has an effect on the inference model performances because the model is trained on data accumulated from every dissipations.

4.2.3. Distance distributions of the noise detections

The labelled dataset built from the smoke experiments is given on fig.19. Here, visibility values are much higher than with the fog tests and range from 100 m to 800 m approximately. We build classes of visibility with a size of 25 m and label the distance distributions of the *free-sky* frustum points according to them. Observations found during the fog experiments about the distance distributions of frustum echoes with the OS1-128 LiDAR are again noticed here. Detections located around 0.3 m have characteristic forms (potentially Gamma or Log-normal) and changes in the shapes, position and amplitudes depending on the evolution of the visibility conditions are observed.

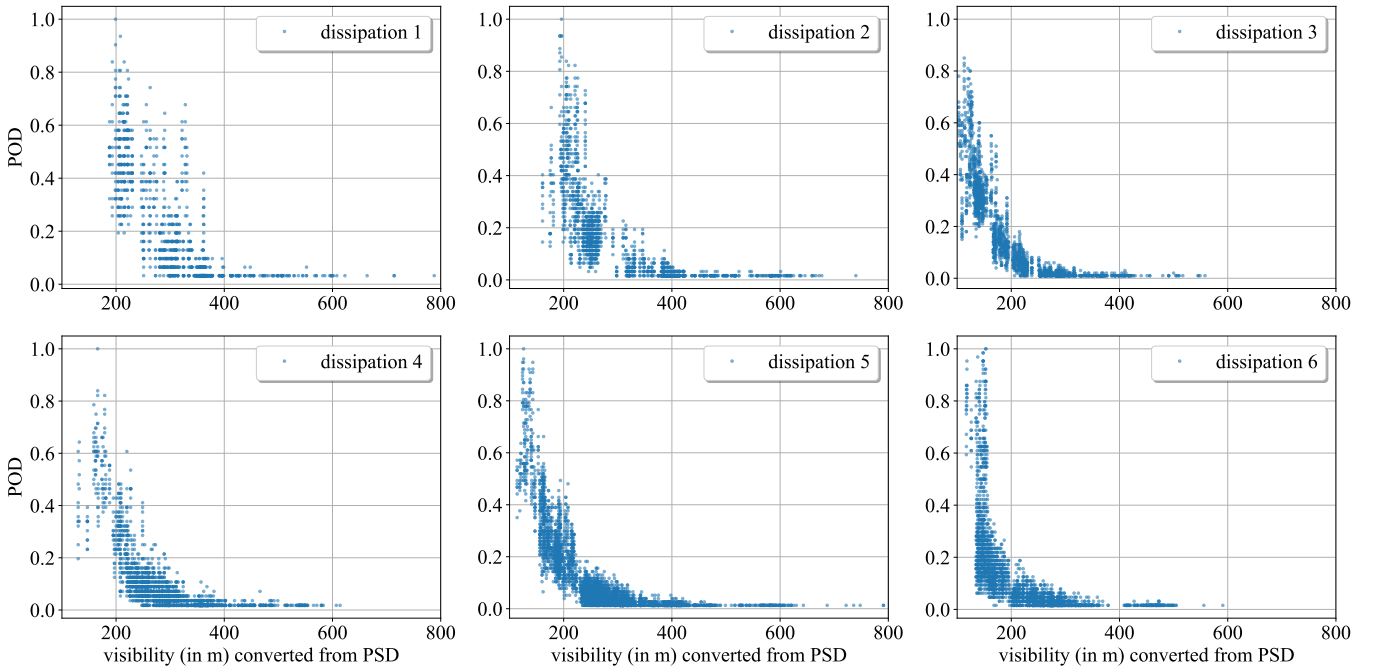


Figure 17: *Free-sky* frustum POD according to visibility converted from the PSD data for each dissipation performed.

Though, the distance distributions of the smoke experiments are more chaotic than the ones of the fog of fig.11, which confirms the erratic nature of the smoke dissipations as discussed previously. Fig.19 also shows the results of the Gamma and Log-normal distributions trained on each visibility class of the smoke tests.

4.2.4. Classification results

Classes of visibility are set from 75 m (minimum value of the converted visibility) to 425 m (too few points are available above) by steps of 25 m. This configuration is chosen as it holds the best results notably because it ensures enough data available to train and test the model on each class.

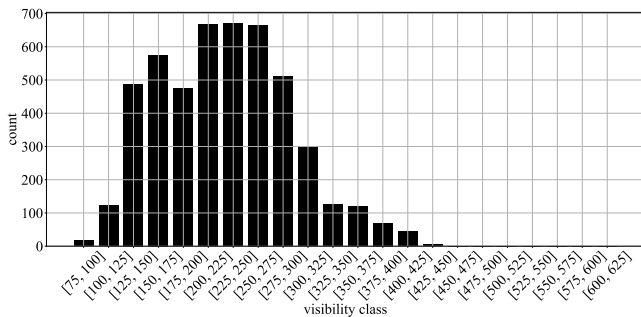


Figure 20: Number of LiDAR scans used for the test phase for each visibility class in the smoke dataset. This represents the distribution of 1/3 of the overall dataset, remaining 2/3 being used for training purposes. Note that lower and higher visibility classes have fewer data for both testing and training purposes.

To evaluate the classification results, the dataset is split into a training set (2/3 of the scans) and a test set (1/3 of the scans). Results from each configuration are given and summarized in table 5 and fig.18. Similarly to the fog tests, the number of available LiDAR scans for the smoke test (available fig.20) differs for each visibility class and this mitigates the classification results. In addition, smoke dissipations have erratic behaviors which lead to non homogeneous results when confronting the *free-sky* POD data with visibility.

Likelihood model	Gamma	Log-normal
No RFS model	85 m	69 m
Poisson RFS	39 m	39 m
Binomial RFS	51 m	42 m

Table 5: RMSE values for each model combination in smoke conditions. We first note that the RFS formalism improves inference results. Furthermore, the RMSE metric suggests that the Poisson RFS is better suited to infer visibility in smoke conditions than the Binomial RFS.

The classification results, highlighted on the confusion matrices of fig.18 and RMSE (table 5) show lower performances than for the fog tests. While the fog tests have RMSE values around the class size of 5 m, the RMSE values of the smoke tests range from 1.5 to 3.4 times the class size of 25 m. This is explained by the accumulation of data from non-uniform dissipations and the erratic behaviors as described before. Nevertheless, we remain confident that the model can be efficient and robust with our approach

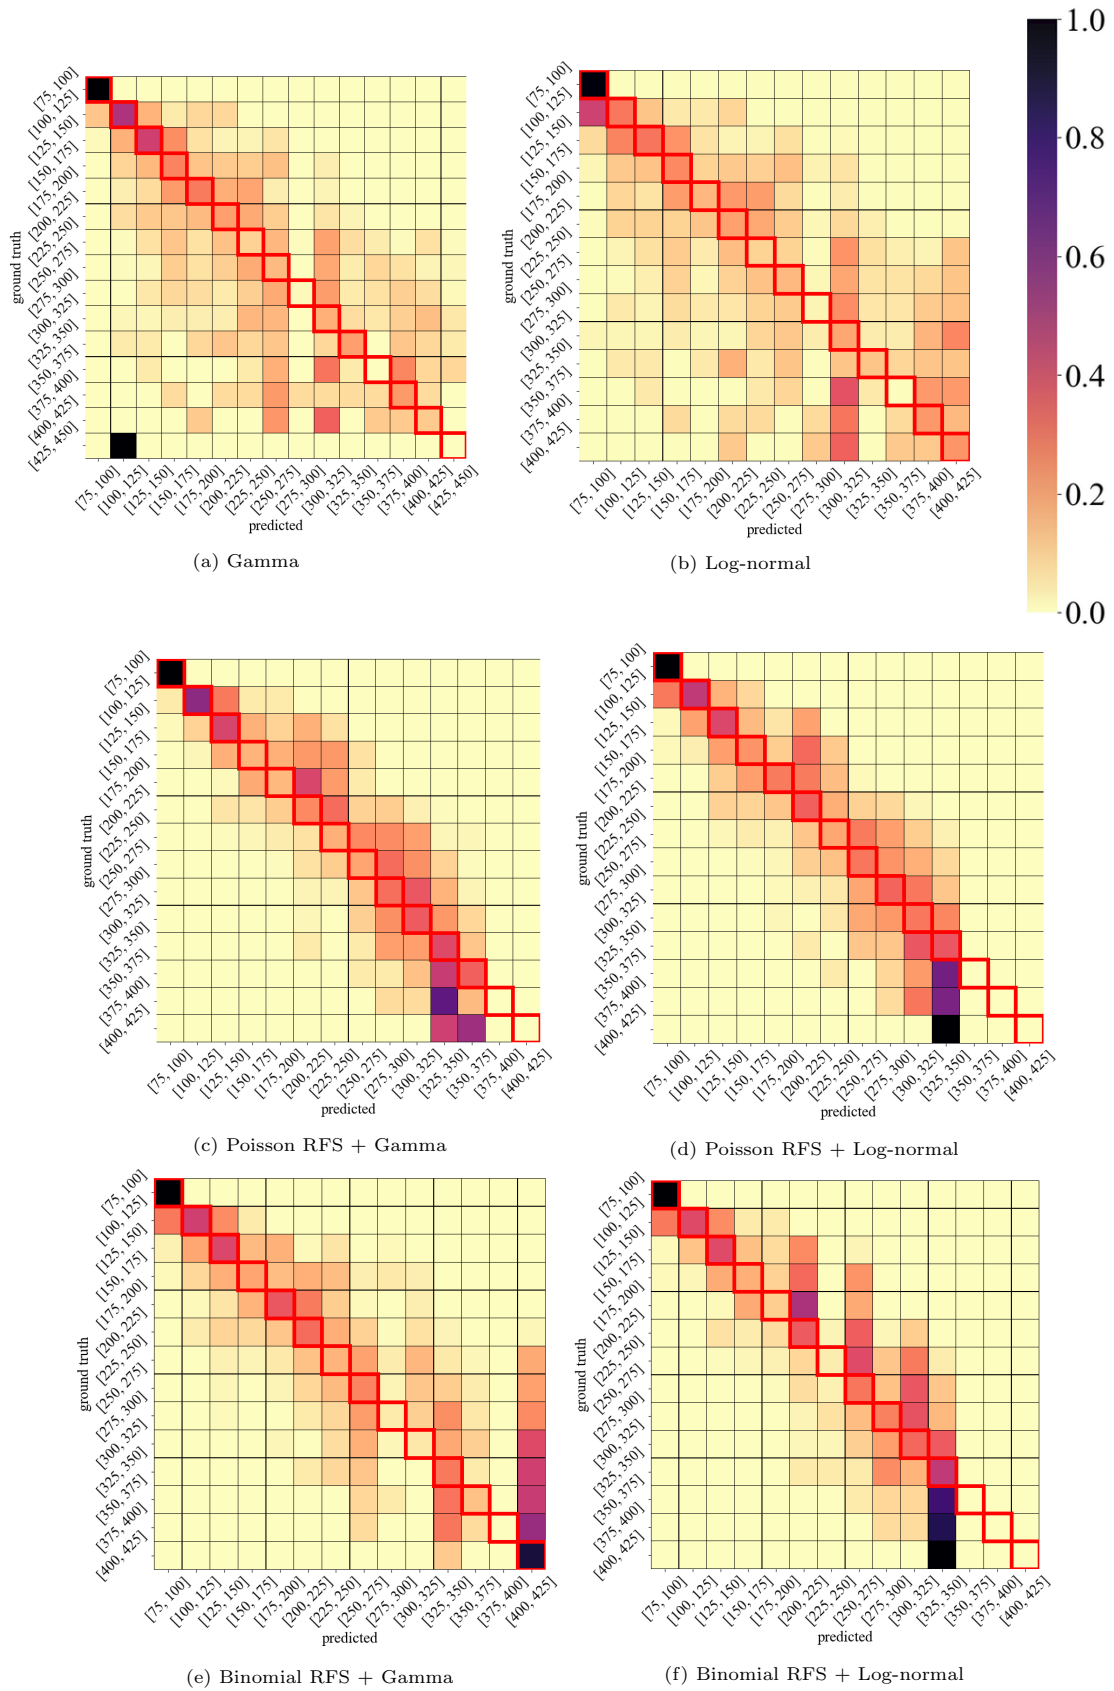


Figure 18: Confusion matrices from visibility classification using the *free-sky* target in smoke conditions and the multiple combinations of likelihoods. Color represents the probability of true positive classification.

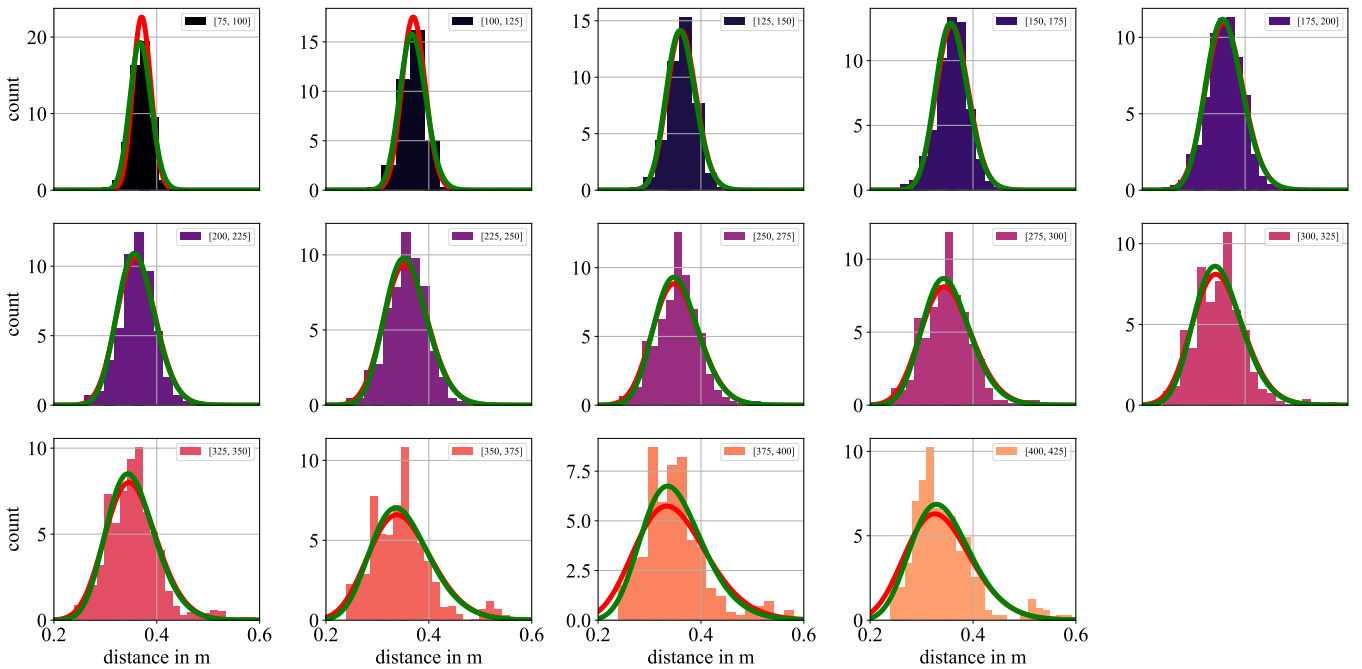


Figure 19: Distance histograms of *free-sky* frustum smoke points detected for each visibility class, color refers to visibility in meter. Trained **Gamma** and **Log-normal** distributions are also plotted for each class in red and green.

if the dissipations produced are improved. The RFS extension still enhances the overall inference results but here, the Poisson RFS model holds the best results with a RMSE value of 39 m for both Gamma and Log-normal likelihood distributions.

5. Conclusion

This paper presents a methodology to infer optical visibility values from automotive 3D-LiDAR point clouds. The goal is to improve the capacity of an autonomous vehicle to retrieve information about the weather conditions with embedded perception sensors. Particles of fog and smoke produced in artificial conditions create false positive echoes in the LiDAR point cloud data. Using the Ouster OS1-128 LiDAR, the distance distribution of these echoes have characteristic shapes and evolve with changing visibility values. In addition these points are located at very close range which is convenient to perform the weather condition inference because there is hardly any chance that a real object is located at such small distances.

A fully probabilistic Bayesian inference framework is described with the aim of inferring discrete visibility classes from these LiDAR points where Gamma and Log-normal distributions model the shapes of the distance distributions, and Poisson and Binomial RFS models capture the cardinality of the sets of echoes. The approach uses a *free-sky* target with which the field-of-view does not contain any objects because it increases the odds of producing false positive echoes, and thus increases the inference performance. This reproduces outdoor conditions where parts

of the field of view comprise the sky – alternatively, the model could use known targets with appropriate learned hyper-parameters (the paper shows an example with a low reflectivity target). The generalization of the model to targets of multiple distance and reflectivity is a crucial point of improvement for the deployment of this approach.

Labelled datasets are created using ground truth visibility sensors or PSD data converted into visibility. The creation of such labelled dataset is not straightforward because the dissipation of the fog or smoke particles are not totally predictable and can create non homogeneous results. Nevertheless, the model remains robust with relatively good RMSE results on visibility classification, especially for the fog experiments which are performed in a more controlled environment. Gamma and Log-normal distributions are suitable to model the distance distributions of false positive echoes in DVE, and the RFS formalism improves the classification results by statistically capturing the evolution of the number of noise points.

Future work could include the production of a better quality dataset (especially for smoke). The inference model could also be trained and tested on real outdoor acquisitions, with for example the dataset proposed in [36]. Finally, the temporal evolution of visibility could be taken into account to improve the classification over time, with for example a Bayesian histogram filter.

References

- [1] Y. Li, J. Ibanez-Guzman, Lidar for Autonomous Driving: The principles, challenges, and trends for automotive lidar and per-

- ception systems, *IEEE Signal Processing Magazine* (Jul. 2020). doi:10.1109/MSP.2020.2973615.
- [2] F. Di Stefano, S. Chiappini, A. Gorreja, M. Balestra, R. Pierdicca, Mobile 3D scan LiDAR: a literature review, *Geomatics, Natural Hazards and Risk* (Jan. 2021). doi:10.1080/19475705.2021.1964617.
 - [3] R. H. Rasshofer, M. Spies, H. Spies, Influences of weather phenomena on automotive laser radar systems, *Advances in Radio Science* (Jul. 2011). doi:10.5194/ars-9-49-2011.
 - [4] A. M. Wallace, A. Halimi, G. S. Buller, Full waveform LiDAR for adverse weather conditions, *IEEE Transactions on Vehicular Technology* (2020) 1–doi:10.1109/TVT.2020.2989148.
 - [5] F. Christnacher, S. Schertzer, N. Metzger, E. Bacher, M. Laurenzis, R. Habermacher, Influence of gating and of the gate shape on the penetration capacity of range-gated active imaging in scattering environments, *Optics Express* 23 (26) (2015) 32897. doi:10.1364/OE.23.032897.
 - [6] M. Pfennigbauer, C. Wolf, J. Weinkopf, A. Ullrich, Online waveform processing for demanding target situations, Baltimore, Maryland, USA, 2014, p. 90800J. doi:10.1117/12.2052994.
 - [7] M. Ijaz, Z. Ghassemlooy, H. Le Minh, S. Rajbhandari, J. Perez, Analysis of fog and smoke attenuation in a free space optical communication link under controlled laboratory conditions, in: 2012 International Workshop on Optical Wireless Communications (IWOW), IEEE, Pisa, Italy, 2012, pp. 1–3. doi:10.1109/IWOW.2012.6349680.
 - [8] K. Montalban, C. Reymann, D. Atchuthan, P.-E. Dupouy, N. Riviere, S. Lacroix, A Quantitative Analysis of Point Clouds from Automotive Lidars Exposed to Artificial Rain and Fog, *Atmosphere* 12 (6) (2021) 738. doi:10.3390/atmos12060738.
 - [9] R. Heinzler, F. Piewak, P. Schindler, W. Stork, CNN-based Lidar Point Cloud De-Noising in Adverse Weather, *IEEE Robotics and Automation Letters* (Apr. 2020). doi:10.1109/LRA.2020.2972865.
 - [10] C. Courcelle, D. Baril, F. Pomerleau, J. Laconte, On the Importance of Quantifying Visibility for Autonomous Vehicles under Extreme Precipitation, Tech. rep. (Sep. 2022).
 - [11] Z. Jiajun, D. Dmitri, F. Dave, Methods and systems for detecting weather conditions including fog using vehicle onboard sensors.
 - [12] R. K. Heinzler, LiDAR-based Weather Detection: Automotive LiDAR Sensors in Adverse Weather Conditions, Ph.D. thesis.
 - [13] J. R. Vargas Rivero, T. Gerbich, V. Teiluf, B. Buschardt, J. Chen, Weather Classification Using an Automotive LIDAR Sensor Based on Detections on Asphalt and Atmosphere, *Sensors* (Aug. 2020). doi:10.3390/s20154306.
 - [14] G. Sebastian, T. Vattem, L. Lukic, C. Burgy, T. Schumann, RangeWeatherNet for LiDAR-only weather and road condition classification, *IEEE*, 2021. doi:10.1109/IV48863.2021.9575320.
 - [15] R. Karlsson, D. R. Wong, K. Kawabata, S. Thompson, N. Sakai, Probabilistic Rainfall Estimation from Automotive Lidar (Apr. 2021).
 - [16] R.-C. Miclea, C. Dughir, F. Alexa, F. Sandru, I. Silea, Laser and LIDAR in a System for Visibility Distance Estimation in Fog Conditions (Nov. 2020). doi:10.3390/s20216322.
 - [17] D. Blagojevic, Weather effects on short-range LiDAR and their classification (2022).
 - [18] Y. Li, P. Duthon, M. Colomb, J. Ibanez-Guzman, What happens for a ToF LiDAR in fog? (May 2020).
 - [19] T. Yang, Y. Li, Y. Ruicheck, Z. Yan, LaNoising: A Data-driven Approach for 903nm ToF LiDAR Performance Modeling under Fog, *IEEE*, 2020. doi:10.1109/IROS45743.2020.9341178.
 - [20] International Civil Aviation Organization, Council, Meteorological service for international air navigation: international standards and recommended practices : annex 3 to the Convention on International Civil Aviation., International Civil Aviation Organization, Montr??al, 2010, oCLC: 657037416.
 - [21] N. Rivière, P.-E. Dupouy, Perception lidar 3D pour l'aide à la conduite autonome, *Photoniques* (Sep. 2022). doi:10.1051/photon/202211528.
 - [22] G. Satat, M. Tancik, R. Raskar, Towards photography through realistic fog, *IEEE*, 2018. doi:10.1109/ICCPHOT.2018.8368463.
 - [23] D. Koller, N. Friedman, Probabilistic graphical models: principles and techniques, 2009.
 - [24] C. M. Bishop, Pattern recognition and machine learning, Information science and statistics, New York, 2006.
 - [25] K. P. Murphy, Machine learning: a probabilistic perspective, 2012.
 - [26] S. Thrun, Probabilistic robotics, *Communications of the ACM*-doi:10.1145/504729.504754.
 - [27] S. Weinzierl, Introduction to Monte Carlo methods, Tech. rep. (Jun. 2000).
 - [28] S. Ronnback, A. Wernersson, On filtering of laser range data in snowfall, *IEEE*, 2008. doi:10.1109/IS.2008.4670551.
 - [29] S. Michaud, J.-F. Lalonde, P. Giguere, Towards Characterizing the Behavior of LiDARs in Snowy Conditions (2015).
 - [30] D. J. Daley, D. Vere-Jones, An introduction to the theory of point processes, 2003.
 - [31] B.-N. Vo, B.-T. Vo, D. Phung, Labeled Random Finite Sets and the Bayes Multi-Target Tracking Filter, *IEEE Transactions on Signal Processing* 62 (24) (2014) 6554–6567. doi:10.1109/TSP.2014.2364014.
 - [32] R. P. S. Mahler, Advances in statistical multisource-multitarget information fusion, 2014.
 - [33] B. Ristic, B.-T. Vo, B.-N. Vo, A. Farina, A Tutorial on Bernoulli Filters: Theory, Implementation and Applications, *IEEE Transactions on Signal Processing* (Jul. 2013). doi:10.1109/TSP.2013.2257765.
 - [34] R. Ceolato, A. E. Bedoya-Velásquez, V. Mouysset, Short-Range Elastic Backscatter Micro-Lidar for Quantitative Aerosol Profiling with High Range and Temporal Resolution, *Remote Sensing* (Oct. 2020). doi:10.3390/rs12203286.
 - [35] E. M. Patterson, D. A. Gillette, G. W. Grams, The Relation Between Visibility and the Size-Number Distribution of Airborne Soil Particles, *Journal of Applied Meteorology* (May 1976). doi:10.1175/1520-0450(1976)015<0470:TRBVAT>2.0.CO;2.
 - [36] C. Linnhoff, K. Hofrichter, L. Elster, P. Rosenberger, H. Winner, Measuring the Influence of Environmental Conditions on Automotive Lidar Sensors, *Sensors* (Jul. 2022). doi:10.3390/s22145266.