

Adaptive Wind Field Estimation Using an Empirical Bayesian Approach

Daniel Kiehn^{*1}, Julius Schultz^{†2}, Nicolas Fezans^{‡1}, and Ulrich Römer^{§2}

¹German Aerospace Center (DLR), Institute of Flight Systems, Lilienthalplatz 7, 38108 Braunschweig, Germany.

²Technische Universität Braunschweig, Institute for Acoustics and Dynamics, Langer Kamp 19, 38106 Braunschweig, Germany

In lidar-based gust load alleviation, the wind profile ahead of the aircraft cannot be measured directly, but has to be reconstructed (estimated) based on the acquired line-of-sight measurements. Such wind reconstruction algorithms typically include regularization in order to adequately handle the noise within the data. This paper presents an empirical Bayesian approach to choose optimal regularization parameters for any given set of measurements. Using simulations of flight through turbulence, the Bayesian approach is compared with a former approach (based on engineering guess) and an *omniscient optimizer* which yields the best achievable results for a constant set of parameters by using the full knowledge of the wind field. The Bayesian approach is shown to outperform the engineering guess and performs close to the omniscient optimizer while purely relying on the lidar measurement data.

I. Introduction

Recent years have shown a renewed interest in the application of Doppler-lidar (light detection and ranging) wind sensors for feedforward gust load alleviation (GLA). Gusts and turbulence cause structural loads on aircraft which the aircraft structures must be designed to withstand. GLA controllers are dedicated flight control functions whose aim is to actively reduce these loads by anticipating or reacting to the encountered gusts and turbulence. The use of GLA controllers offers the potential to reduce structural loads on the aircraft, thereby improving safety of operations, passenger comfort, and allowing for lighter aircraft design [1–4]. It has been shown that the usage of feedforward GLA controllers based on remote lidar measurements increases the achievable performance of GLA functions, as these measurements allow anticipating the oncoming disturbance [5–8].

In the context of gust and turbulence load alleviation, the most relevant wind component is typically the vertical wind. However, Doppler lidars only measure the line-of-sight projection of the relative wind w.r.t. the aircraft, so the

^{*}Research Scientist, Department of Flight Dynamics and Simulation. Email: daniel.kiehn@dlr.de. Corresponding author.

[†]PhD Researcher, Faculty of Mechanical Engineering. Email: j.schultz@tu-braunschweig.de

[‡]Research Scientist, Department of Flight Dynamics and Simulation. Email: nicolas.fezans@dlr.de. AIAA Senior Member.

[§]Associate Professor, Faculty of Mechanical Engineering. Email: u.roemer@tu-braunschweig.de

^{*†} These authors contributed equally to this work

vertical wind ahead of the aircraft cannot be directly measured with a forward-looking sensor. Instead, it must be estimated (reconstructed) based on the line-of-sight measurements acquired with the Doppler lidar sensor. This task is performed by a wind reconstruction algorithm, which requires regularization to ensure a certain degree of smoothness of the resulting estimated wind field: all actions of the feedforward GLA controller are based on the estimated wind profiles, and the smoothness of these profiles reduces unnecessary control actions (and related actuator loads) due to measurement noise. The regularization parameters have a significant impact on the accuracy of the estimation result – which is the primary information that the feedforward GLA controller acts on – and hence significantly influences the performance of the GLA function. Oversmoothing causes higher frequency components of the wind field to not be captured by the wind reconstruction, but some degree of regularization is required to reduce the impact of measurement noise. Additionally, the optimal choice of regularization depends on the lidar sensor, the atmospheric conditions, and the wind conditions around the aircraft, as shown in a recent parameter study [9].

Previously, e.g., in [5] or [9], fixed regularization parameters (chosen based on engineering judgement) were used, which produced suboptimal results. It was found in sensitivity studies of the lidar system [9] that when the parameters of the system (e.g., laser power, measurement range, etc.) are varied, the optimal values for the regularization parameters also change. It is therefore crucial to find a suitable set of regularization parameters.

This paper addresses the challenge of choosing optimal regularization parameters for any given configuration. To this end, we reformulate the wind field estimation problem in a Bayesian framework. We demonstrate how the formerly used Tikhonov regularization can be transferred to a Bayesian setting by introducing Gaussian smoothness priors. In the Bayesian framework the optimal regularization parameters can then be found by optimizing the marginal likelihood. The proposed method is adaptive to lidar parameters, atmospheric conditions, and the wind field itself without requiring any explicit information about those variables, because it is fully reliant on the measurements. Additionally, it naturally provides uncertainty estimates of the reconstructed wind field. However, by introducing regularization, the uncertainty estimates can get overconfident, as observed in our study and discussed in [10].

The paper is structured in the following way: In Section II the wind estimation based on lidar measurements is explained, and the application of the Bayesian approach to wind field reconstruction is described in section III. Section IV introduces the error metric which is used to evaluate the performance of individual configurations in the frequency domain. In Section IV the analysis setup and the reference cases are introduced. The simulation results are discussed in Section VI, and the conclusions are presented in Section VII.

II. Wind estimation based on lidar measurements

A. Doppler wind lidars

The wind field ahead of the aircraft is retrieved from measurements of a Doppler wind lidar. Such sensors utilize the Doppler effect which describes a frequency shift of any observed wave caused by a relative motion between the emitter and receiver (i.e., observer). With such a system, the atmospheric disturbance (the wind field) can be detected and counteracted before it is encountered by the aircraft. The typical measurement distance for such systems is 50 to 300 m.

The lidar emits short laser pulses, which typically last for a few tens of nanoseconds. Only a specific area ahead of the aircraft is illuminated by the emitted light due to the low divergence of the laser beam. At each point of the laser beam, a tiny fraction of the pulse is scattered by the air molecules (and possibly some aerosols, if present). By comparing the frequency of the light scattered back to the lidar sensor with that of the emitted light, the relative velocity between the aerosols/molecules which scattered the light back can be determined: if the sensor and the aerosols/molecules are moving towards each other, a shift towards higher frequencies (so-called blue shift) is observed. Likewise, if they are moving away from each other, the spectrum is shifted towards lower frequencies (red shift).

So-called coherent heterodyne detection principles with a laser source in the infrared domain are generally the best choice if the presence of a sufficient number of aerosols can be assumed. For the alleviation of clear air turbulence at high altitudes however, the aerosol concentration is often insufficient. If a high availability of the remote wind measurement is desired (which a GLA function would require), a direct-detection principle with a laser source in the ultraviolet domain should be used, since it can work with Rayleigh-scattering from air molecules. More details on lidar-based remote wind measurement technologies and their respective capabilities can be found in the literature, for instance in [11, 12] and references therein.

With Doppler lidar sensors, only the relative wind component in the direction of the laser beam is measured, as illustrated in Fig. 1. The lidar sensor measures V_{LOS} , which is the projection of $V_{TAS}(M)$ at the measurement location M onto the sensor's line of sight (LOS). $V_{TAS}(M)$ is a combination of the inertial speed V_K of the aircraft and the local vertical wind speed $V_{wind}(M)$ at the measurement location [9]. If the wind is measured at locations ahead of the aircraft, then the direction of the laser beam (also called line-of-sight, LOS) is almost collinear with the flight path, and the measured relative wind speed is determined mostly by the true airspeed. The most relevant wind component for GLA is the vertical wind due to its direct effect on lift. Hence, this information has to be reconstructed from the relative wind measurements. The missing information (other velocity components as well as the wind at locations that were not directly measured) can be reconstructed by measuring the wind at locations with various vertical and lateral offsets with respect to the aircraft's flight path. Such offsets can be created by using multiple discrete line-of-sight directions, as was done in the AWIATOR programme (Aircraft Wing with Advanced Technology Operation) [11, 13],

or by pointing the sensor’s LOS off-center with a specific scan angle and rotating it around the aircraft’s longitudinal axis, as shown in Fig. 1. This way measurements are performed in different (i.e., non-collinear) directions, and the transversal components (lateral and vertical) of the wind can be estimated by analyzing the differences between the individual measurements. This method implicitly assumes that the wind is sufficiently homogeneous between the points where the measurements are made. This assumption is of course difficult to validate and strongly depends on the current atmospheric conditions encountered by the aircraft and the distance between the locations at which the considered measurements were performed. The closer the measurements are located, the more likely it is that this assumption is valid. Reducing the distance (laterally and vertically) between the measurements would however lead to reduced angles between the different line-of-sight directions, eventually leading to very small differences between the sensor line-of-sight velocities that are measured. This has a major drawback because these measurements cannot be perfect (noise, biases, etc.). In particular, the signal-to-noise ratio (signal being the difference due to the wind transversal components and noise being linked to the noise on each measurement) becomes very poor when calculating the difference between measurements taken under almost collinear directions. In other words, there will necessarily be a trade-off to be made between the validity of the homogeneity assumption (linked to the distance of the measurements to the flight path) and the signal-to-noise ratio for the reconstructed transversal wind components.




Fig. 1 Lidar measurement geometry. The ratio of the velocities is distorted for readability [9]. (Reproduced by copyright holder).

B. Wind field reconstruction

The aim of the wind reconstruction process is to determine the vertical wind velocities ahead of the aircraft based on the lidar line-of-sight measurements and some requirements about the smoothness of the wind field. The measurements are taken at the pulse repetition frequency of the lidar, which can be in the order of several kHz, and stored in a buffer together with metadata such as the coordinates, the LOS directions, and the expected standard deviation of the measurement.

The wind reconstruction algorithm takes all available lidar measurements which were taken within a certain range in front of and behind the aircraft and reconstructs the most likely wind field. When reconstructing the wind profile, no

assumption on a particular model structure and shape shall be assumed for the present gust(s) or turbulence. In this case, a *free-form model* structure as described in [5] can be used. The idea is to represent the gust/turbulence wind field by a mesh where a velocity vector is set for each node of the mesh. Any wind field can in principle be represented by such a mesh, as long as enough nodes are taken. For the application to gust load alleviation, small-scale wind variations are not relevant and there is no real benefit in using a very fine mesh [5]. The spatial extent of the mesh is defined by the lead and lag times $\tau_{\text{lead}}/\tau_{\text{lag}}$ and the inertial speed of the aircraft V_K .

The result of the wind reconstruction is a set of n estimated vertical wind speeds at equidistant nodes of the free-form wind field model, spanning from $-\tau_{\text{lag}}V_K$ to $\tau_{\text{lead}}V_K$ (assuming that τ_{lag} and τ_{lead} are both positive), as shown in Fig. 2. This wind profile is re-interpolated and provided as the input to the load alleviation controller, which generally runs at a much faster rate than the wind reconstruction algorithm. The re-interpolation is not discussed further in this paper, details can be found in [14]. The focus here is purely on the lidar measurements and the wind reconstruction process. For details about the controller design, the interested reader may refer to [14–16].

The discrete rate at which the wind estimation is run can generally be chosen freely, but due to the relatively high computational effort of the problem a rate of 10 Hz is currently used. The reconstructed wind profile is hence updated every 100 ms, as visualized in Fig. 3, where reconstructed wind profiles of a turbulent wind field are shown for three subsequent estimation windows made at the points x_i , x_{i+1} , and x_{i+2} , respectively.

We denote the ideal set of measurements that a perfect lidar sensor would provide at the measurement locations

$$\{y_i \mid i = 1, \dots, m\} \tag{1}$$

with m indicating the number of measurements currently stored in the buffer, as well as the set of measurements used during the wind reconstruction by

$$\{z_i \mid i = 1, \dots, m\} \tag{2}$$

The lidar sensor is assumed to have been calibrated beforehand (or continuously during operation, e.g., with an autocalibration feature). Measurements are assumed to be independent and subject to mean-free Gaussian measurement noise defined by a respective standard deviation

$$\{\sigma_i \mid i = 1, \dots, m\} \tag{3}$$

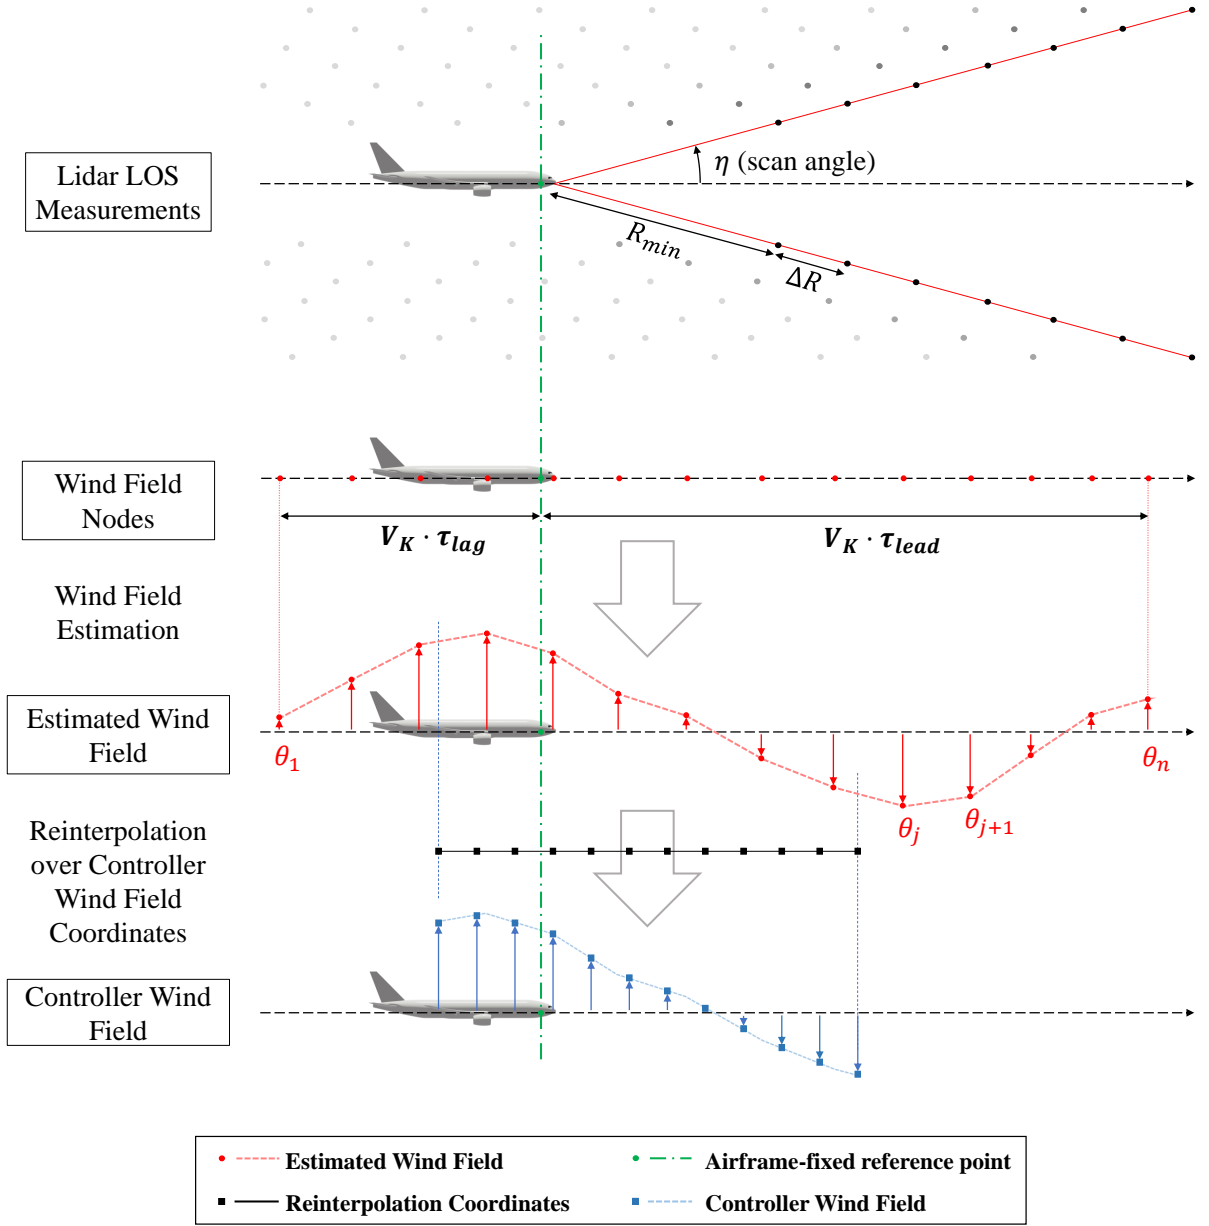


Fig. 2 Wind estimation process from lidar measurements to controller wind field (adapted from [15]). Reproduced by copyright holder.

such that the measurement equation reads

$$z_i = y_i + \epsilon_i, \quad i = 1, \dots, m \quad (4)$$

with $\epsilon_i \sim \mathcal{N}(0, \sigma_i)$.




Fig. 3 Three subsequent estimation windows in the presence of simulated von Kármán turbulence

Let $\theta = [\theta_1, \theta_2, \dots, \theta_n]$ be the vector of n wind field model parameters, which describe the vertical wind speeds of the free-form wind field model. For a given vector of wind field model parameters we can thus define the forward model outputs

$$\{\xi_i(\theta) \mid i = 1, \dots, m\} \quad (5)$$

which consist of the LOS wind speeds $\xi_i(\theta)$ at the same location and conditions as the measurements z_i based on the free-form model and the parameters θ . The forward model can then be compared against the measurements to estimate the wind field parameters. The relation of the measurements and the model parameter is given by

$$z_i = \underbrace{\xi_i(\theta)}_{y_i} + d_i + \epsilon_i \quad (6)$$

where d_i denotes a model discrepancy of the free-form wind field model. The model discrepancy mainly concerns the high-frequency components of the wind field, which are not important for gust load alleviation; hence, d_i is omitted in the following. Based on Eq. (6) we can define a probability density function (pdf) for the measurement data, given a

vector of wind field model parameters

$$p(z_i|\boldsymbol{\theta}) \sim \mathcal{N}(\xi_i(\boldsymbol{\theta}), \sigma_i). \quad (7)$$

Since the measurements z_i are independent, the pdf of the measurement vector \mathbf{z} is given by the product of the marginal pdfs as

$$p(\mathbf{z}|\boldsymbol{\theta}) = \prod_{i=1}^m p(z_i|\boldsymbol{\theta}) \propto \exp\left(-\frac{1}{2} \sum_{i=1}^m \frac{(z_i - \xi_i(\boldsymbol{\theta}))^2}{\sigma_i^2}\right). \quad (8)$$

The likelihood function is defined by $\boldsymbol{\theta} \mapsto p(\mathbf{z}|\boldsymbol{\theta})$ describing the likelihood of the data given model parameter. The maximum likelihood (ML) estimate $\hat{\boldsymbol{\theta}}$ is obtained by minimizing the negative logarithm of the likelihood, which leads to

$$\hat{\boldsymbol{\theta}} = \arg \min_{\boldsymbol{\theta}} \left(\boldsymbol{\theta} \mapsto \sum_{i=1}^m \frac{(z_i - \xi_i(\boldsymbol{\theta}))^2}{\sigma_i^2} \right). \quad (9)$$

In [5] and [9] a regularized ML approach is applied to incorporate smoothing requirements in the estimated wind field. A Gauss-Newton algorithm then solves the potentially non-linear least squares problem. In this work, we use a Bayesian approach which derives a posterior probability density for the model parameters. The model parameters describe the vertical wind speeds in the free-form wind field model and thus the posterior uncertainty can be used for uncertainty quantification of the wind field estimate. Additionally, the Bayesian approach offers concepts to automatically estimate the hyperparameter such as the weights for the smoothing penalties, purely based on the measurement data.

C. Orders of Magnitude for the Application to Gust Load Alleviation

In order to generate the input signal for the feedforward controller, the estimated wind profile is reinterpolated at the controller sampling rate as described in [14, sect. 2.6]. It is important to choose the lead time τ_{lead} large enough to allow sufficient anticipation time for the controller, but also small enough to still ensure a certain density of measurements around the forward-most nodes of the estimation mesh. The required anticipation time is based on the flight speed and the wind component with the longest wavelength of interest. The longest gust for which the aircraft has to be certified according to CS25.341 has a length of 107 m, so at 250 m/s true airspeed a lead time of 0.5 s is enough to cover half of the longest gust of interest, which already permits to initiate the adequate commands sufficiently early. As illustrated in [7] on a fairly simple aeroelastic model, the achievable load alleviation performance increases very significantly with the previewed horizon at first but the curve flattens out fairly quickly. The orders of magnitudes can hardly be transferred from [7] to a fast jet aircraft (very different speeds, wind loading, number of flexible modes and load stations), but the overall trend remains valid and this behavior explains the choice made of a lead time of 0.5 s, which may appear rather short at first.

The minimum and maximum frequencies to be captured by the lidar and the wind reconstruction algorithm generally depend on the aircraft configuration. Generally, the frequency range of interest for feedforward GLA is approximately in the range of 0.5 Hz to 10 Hz. Below 0.5 Hz, the frequencies are low enough to be easily compensated by feedback control functions that are not using the lidar information. Above 10 Hz, the frequencies are typically too high to be effectively countered by currently available actuators and are usually not relevant in terms of load levels. Generally, the most relevant structural mode for GLA is the the first wing bending mode, which typically occurs between 1 Hz and 5 Hz [4, sect. 3.2].

III. Bayesian approach to wind field reconstruction

We use a Bayesian approach [17] to estimate the wind field parameters which builds on Bayes' law

$$p(\boldsymbol{\theta}|\mathbf{z}) = \frac{p(\mathbf{z}|\boldsymbol{\theta})p(\boldsymbol{\theta})}{p(\mathbf{z})}, \quad (10)$$

where $p(\mathbf{z}|\boldsymbol{\theta})$ denotes the likelihood introduced above, $p(\boldsymbol{\theta})$ the prior distribution and $p(\mathbf{z})$ a normalization constant, also called marginal likelihood. The maximum a posteriori estimate of the posterior is directly related to the regularized ML solution, which has been used in previous works for wind field reconstruction [6]. The reader is referred to VII for further details.

A. Smoothing priors

In this work we use Gaussian smoothness priors [18, 19] to penalize the first and second derivative of the wind field parameters in order to induce a certain degree of smoothness in the estimated wind field. A Gaussian prior has the general form

$$p(\boldsymbol{\theta}) = \frac{|\mathbf{Q}|^{1/2}}{(2\pi)^{n/2}} \exp\left(-\frac{1}{2}\boldsymbol{\theta}^T \mathbf{Q} \boldsymbol{\theta}\right) \quad (11)$$

with $|\mathbf{Q}|$ denoting the determinant of the precision (inverse covariance) matrix. The smoothing property is achieved by enforcing correlation between the wind field nodes. We will derive the precision matrix \mathbf{Q} from stochastic partial differential equations (SPDE) [20, 21].

An SPDE is defined as

$$\mathcal{G}u(x) = v(x), \quad x \in D, \quad (12)$$

where \mathcal{G} denotes a differential operator, u a stochastic process, and v a white noise process over a domain D . The

white noise process is uncorrelated and defines at every position x a normal random variable with zero mean and finite variance. The solution of the SPDE is a Gaussian Process (GP) with a covariance structure induced by the choice of \mathcal{G} . Indeed, there is a direct relationship between the operator \mathcal{G} and the precision matrix \mathbf{Q} of the Gaussian process u .

We state this relationship in the discretized version of Eq. (12):

$$\mathbf{\Gamma}\mathbf{u} = \mathbf{v}. \quad (13)$$

In this context $\mathbf{\Gamma}$ is a whitening matrix, transforming a correlated Gaussian random vector \mathbf{u} into a white noise random vector \mathbf{v} . The precision matrix \mathbf{Q} of the multivariate Gaussian random vector \mathbf{u} is then given as

$$\mathbf{Q} = \mathbf{\Gamma}^T \mathbf{\Gamma}. \quad (14)$$

In [20] it is shown that a certain choice of operator \mathcal{G} induces well known covariance structures, e.g., a Matérn covariance. Any type of correlated stochastic process will serve as a smoothing prior. We will use very fundamental SPDEs to derive the precision matrix of the smoothing priors. A straightforward choice is the Laplace operator, e.g., $\mathcal{G} = \frac{d^2}{dx^2}$ for the 1-D case (vertical wind speed w_W as a function of the aircraft position x). This operator requires the solution \mathbf{u} to be twice differentiable. Another possible choice is the Nabla operator, e.g., $\mathcal{G} = \frac{d}{dx}$ for the 1-D case.

A discretization of both operators with finite differences on a grid of size n in 1-D results in

$$\mathbf{\Gamma}_1 = \begin{bmatrix} -1 & 1 & 0 & \dots & 0 \\ 0 & -1 & 1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & -1 & 1 \end{bmatrix}_{(n-1 \times n)} \quad \mathbf{\Gamma}_2 = \begin{bmatrix} -1 & 2 & -1 & 0 & \dots & 0 \\ 0 & -1 & 2 & -1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & -1 & 2 & -1 \end{bmatrix}_{(n-2 \times n)} \quad (15)$$

where $\mathbf{\Gamma}_1$ is the discretized Nabla operator and $\mathbf{\Gamma}_2$ the discretized Laplace operator with omitted boundary conditions.

The precision matrices are thus given as

$$\mathbf{Q}_1 = \mathbf{\Gamma}_1^T \mathbf{\Gamma}_1, \quad (16)$$

$$\mathbf{Q}_2 = \mathbf{\Gamma}_2^T \mathbf{\Gamma}_2. \quad (17)$$

The weighted precision matrix combining a penalty for the first and second derivative eventually reads

$$\mathbf{Q}_{12} = \gamma_1 \mathbf{Q}_1 + \gamma_2 \mathbf{Q}_2. \quad (18)$$

Note that due to the omission of boundary conditions the precision matrix is singular and has no full column rank, which will be discussed below.

In its initial design, the wind reconstruction algorithm only included a regularization of the second derivative, but not the first [22]. It was found later that adding a small amount of regularization on the first derivative can be beneficial in some cases [23]. However, using regularization based on *only* the first derivative can be harmful, as it penalizes linear trends (e.g., wind gradients) just as much as a sawtooth profile with the same piecewise absolute gradient: $\gamma_1 \rightarrow \infty$ and $\gamma_2 = 0$ produces a straight line with no slope as it infinitely penalizes any gradient. Terms of higher than second order might theoretically be beneficial, but the effect will be very small compared to the first and second order terms. In this paper, we apply regularization only on the first and second derivatives, because we do not want to enhance the wind reconstruction algorithm itself, but exploit its full potential by optimizing the hyperparameters.

B. Posterior for linear models

We now recall the posterior distribution for the case that the wind field model depends linear on its parameters θ . Bayesian inference for a linear model is well-understood, see, e.g., [24], but nevertheless stated here for completeness. In this case the likelihood reads

$$p(\hat{\mathbf{z}}|\theta) \propto \exp\left(-\frac{1}{2}\|\hat{\mathbf{z}} - \mathbf{A}\theta\|^2\right) \quad (19)$$

with

$$\hat{z}_i = \frac{z_i}{\sigma_i} \quad \forall i = 1, \dots, m \quad \text{and} \quad (20)$$

$$A_{i,j} = \frac{\partial}{\partial \theta_j} \frac{\xi_i(\theta)}{\sigma_i} \quad \forall i = 1, \dots, m. \quad (21)$$

Applying Bayes' law and omitting the normalization constant, we obtain

$$p(\theta|\hat{\mathbf{z}}) \propto \exp\left(-\frac{1}{2}\|\hat{\mathbf{z}} - \mathbf{A}\theta\|^2 - \frac{1}{2}\theta^T \mathbf{Q}_{12}\theta\right). \quad (22)$$

Since the measurement noise and the prior are Gaussian distributions the posterior also follows a Gaussian distribution as

$$p(\boldsymbol{\theta}|\hat{\mathbf{z}}) \sim \mathcal{N}(\boldsymbol{\mu}_{\text{Post}}, \boldsymbol{\Sigma}_{\text{Post}}). \quad (23)$$

Analytical expression can be derived by comparing Eq. (22) to a normal distribution [18]. The mean vector and covariance matrix are then obtained as

$$\boldsymbol{\mu}_{\text{Post}} = (\mathbf{A}^T \mathbf{A} + \mathbf{Q}_{12})^{-1} \mathbf{A}^T \hat{\mathbf{z}} \quad (24)$$

$$\boldsymbol{\Sigma}_{\text{Post}} = (\mathbf{A}^T \mathbf{A} + \mathbf{Q}_{12})^{-1}. \quad (25)$$

For a Gaussian distribution the mean is also the mode of the distribution and thus the maximum a posteriori (MAP) estimate reads $\boldsymbol{\theta}_{\text{MAP}} = \boldsymbol{\mu}_{\text{Post}}$.

We recall that the construction of the Gaussian smoothness priors with omitted boundary conditions led to improper prior densities with singular precision matrix. However, the posterior density is still a proper density that defines a unique $\boldsymbol{\theta}_{\text{MAP}}$ if the matrices \mathbf{A} and \mathbf{Q}_{12} have a trivial common nullspace. We can interpret this fact as follows: in case of an improper prior, sampling from the prior distribution is not possible since the covariance matrix does not exist. However in the case of the posterior, the data in the likelihood and the prior together contain enough information to turn the posterior in a proper density with defined mean and covariance.

C. Hyperparameter learning

The choice of the weights for the smoothing priors (hyperparameters) has a strong influence on the quality of the reconstructed wind field. In the case of under-smoothing the wind field is too “oscillatory” and in the case of over-smoothing the reconstruction is smooth but has large deviation from the true wind field. Within the Bayesian approach there is a way to infer the regularization parameters from the data by maximizing the marginal likelihood [25]. This approach is referred to as empirical Bayesian approach [26]. The marginal likelihood is defined as

$$p(\hat{\mathbf{z}}) = \int p(\hat{\mathbf{z}}|\boldsymbol{\theta})p(\boldsymbol{\theta})d\boldsymbol{\theta}. \quad (26)$$

It is obtained by marginalizing the likelihood over the prior probability density. In the particular case of a linear model with Gaussian noise and Gaussian prior, there exists an analytical form of the marginal likelihood. The log marginal

likelihood then reads

$$\log p(\hat{\mathbf{z}}) = -\frac{1}{2} \log |\mathbf{A}\mathbf{Q}_{12}^{-1}\mathbf{A}^T + \mathbf{I}| - \frac{m}{2} \log 2\pi - \frac{1}{2} \hat{\mathbf{z}}^T (\mathbf{A}\mathbf{Q}_{12}^{-1}\mathbf{A}^T + \mathbf{I})^{-1} \hat{\mathbf{z}}. \quad (27)$$

The precision matrix \mathbf{Q}_{12} depends on the choice of the hyperparameter γ_1 and γ_2 . We highlight this dependence explicitly by writing

$$\begin{aligned} \log p(\hat{\mathbf{z}}|\gamma_1, \gamma_2) &= -\frac{1}{2} \log |\mathbf{A}\mathbf{Q}_{12}^{-1}(\gamma_1, \gamma_2)\mathbf{A}^T + \mathbf{I}| \\ &\quad - \frac{m}{2} \log 2\pi \\ &\quad - \frac{1}{2} \hat{\mathbf{z}}^T (\hat{\mathbf{A}}\mathbf{Q}_{12}^{-1}(\gamma_1, \gamma_2)\mathbf{A}^T + \mathbf{I})^{-1} \hat{\mathbf{z}}. \end{aligned} \quad (28)$$

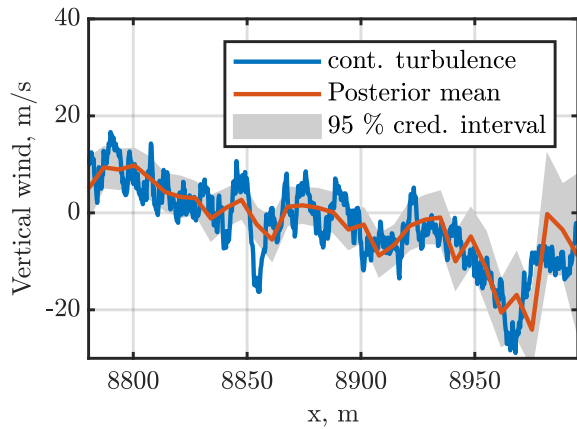
The marginal likelihood can now be optimized with respect to the hyperparameter, i.e.,

$$\gamma_1^*, \gamma_2^* = \arg \min_{\gamma_1, \gamma_2} \left(\gamma_1, \gamma_2 \mapsto -\log p(\hat{\mathbf{z}}|\gamma_1, \gamma_2) \right). \quad (29)$$

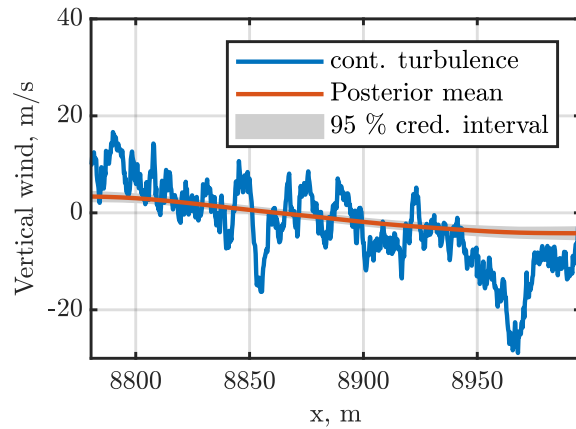
Using the optimized hyperparameters has often been interpreted as applying Occam’s razor, which states that “Models should not be more complex than necessary to explain the data.”. This is achieved by optimizing the marginal likelihood, since the first term is a penalty on complexity and the last term is a penalty on data-fit. The middle term is a constant and can be neglected in optimization. It is important to note that the calculation of the marginal likelihood requires the precision matrix \mathbf{Q}_{12} to be invertible. Thus, in case of a smoothing prior with no boundary conditions, a small noise term needs to be added on the diagonal to regularize the precision matrix.

Figure 4 illustrates the behavior of the regularization parameters and the wind field reconstruction for extreme choices of the regularization parameters as well as for the optimized values. Very small regularization parameters (Fig. 4a) lead to wind fields that attempt to follow the small scale variation of the turbulent wind, however, peaks are still missed and sometimes the estimation even goes in the wrong direction. On the other hand, too large regularization parameters (Fig. 4b) lead to a poor fit of the data, since the high frequency content is mostly lost. Figure 4d shows the result of the wind field reconstruction with optimized hyperparameters obtained by optimizing the marginal likelihood (Fig. 4c). The posterior mean yields a good representation of the wind field following a reasonable compromise between data fit and complexity.

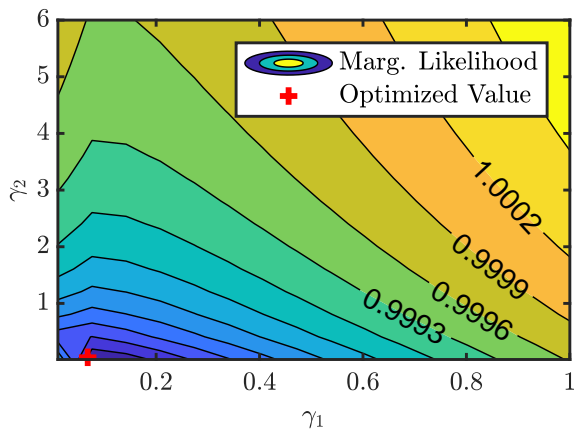
Additionally, Fig. 4 presents the uncertainty of the vertical wind field estimates indicated by the 95 % credible interval. The 95 % credible interval is defined as the central region of 95 % posterior probability. Credible intervals take into account the measurement noise (via the likelihood) and the regularization (via the prior). It is discussed in [10] that in




(a) Under-smoothing: $\gamma_1 = \gamma_2 = 1e-5$



(b) Over-smoothing: $\gamma_1 = \gamma_2 = 10$



(c) Contour of marginal likelihood



(d) Optimized smoothing: $\gamma_1 = 0.1, \gamma_2 = 0.06$

Fig. 4 Wind field reconstruction for one window of a simulated continuous turbulence.

case of an over-smoothing prior the credible interval gets overconfident and the coverage of the true data is low. In this case the prior assumptions overrule the data. On the other hand, for under-smoothing the credible interval purely reflects the uncertainty due to measurement noise. In our study, we recover these effects as seen in Fig. 4a and in Fig. 4b. The 95 % credible interval for the wind field prediction with optimized smoothing parameters represents a reasonable interval for likely free-form model wind field reconstructions. The width of the interval is influenced by the smoothing prior and the fact that the free-form model can not represent the high frequency content by construction. For feedforward gust load alleviation we are interested in a good mean estimate covering the important (low) frequency content of the wind field. Thus, in the remainder of the paper we purely focus on the posterior mean estimate.

IV. Evaluation metric

To automatically assess the performance of the lidar sensor and the wind reconstruction technique used, we apply the frequency-domain metric introduced in [9]. The relevant information on this metric and its physical interpretation is reminded in this section. The reader is referred to [9] for further details and explanations.

The proposed metric generally works by simulating a flight through broadband turbulence, and comparing the reconstructed wind field to the turbulent wind field which the aircraft actually encountered. Theoretically, this analysis could be performed with any turbulent wind spectrum, but we decide to use the broadband turbulence spectrum defined in the EASA certification specification CS 25.341(b) [27], also known as *von Kármán* spectrum. The von Kármán spectrum defines the spectral shape of the turbulent wind via the power spectral density (PSD) $\Phi_w(\Omega)$, which is a function of the spatial frequency Ω . The spatial frequency is related to the temporal frequency f via the true airspeed V_{TAS} with $\Omega = 2\pi f/V_{TAS}$ [28]. In this paper, we will use this relationship and express the power spectral density $\Phi_w(f)$ as a function of the temporal frequency for convenience.

In the first step, a flight through continuous turbulence is simulated using MATLAB[®]/Simulink[®]. Even though the lidar simulation generally includes the coupled motion of the lidar and the aircraft (e.g., change of the lidar measurement direction due to change of the aircraft's attitude caused by turbulence or control surface deflections), this feature is disabled in order to remove all effects of aircraft and control system characteristics from the analysis and focus purely on the lidar performance. The movement of the lidar is hence unaffected by the surrounding turbulence; it is simulated as moving with constant inertial speed and constant pitch offset (representing the trim angle of the aircraft). During the simulation, the wind profile is continuously estimated at a rate of 10 Hz and stored with its corresponding coordinates in space. Since the estimation windows overlap spatially as shown in Sec. II.B, a time series can be assembled for each relative position (with respect to the aircraft) as follows. At each time $t \in [t_i, t_{i+1})$ between the estimation times t_i and t_{i+1} , the wind at the current position $x(t)$ can be obtained by interpolating the estimated wind profile between the nodes

surrounding that position:

$$x(t) = x_i + \Delta x = x_i + (t - t_i)V_K .$$

The obtained time series is piecewise continuous: it remains continuous as long as no new estimation is performed, but jumps may occur upon switching between two successively estimated wind profiles. A new estimation is performed every $\Delta t_{est} = 100$ ms. This process is shown exemplarily for three successive estimation windows in Fig. 5, which correspond to the estimation windows shown in Sec. II.B. The thick lines represent the segments which are stitched together to form the time series, and the thin lines represent the portions of the original estimation windows that are discarded when assembling the time series*. The result at the end of step 1 is a piecewise continuous time series of the estimated wind at the considered location. In the following, for conciseness, only the current position of the aircraft nose is considered, but the process is identical and the obtained results would be fairly similar for other positions.




Fig. 5 Stitching together subsequent estimation windows to obtain a piecewise continuous signal.

In the second step, the absolute error between the time signal obtained in step 1 and the turbulent wind field is determined by calculating the difference signal Δw_W . A segment of a reconstructed wind profile, the actual turbulent wind w_W , and the resulting difference signal Δw_W are shown in Fig. 6.

The third step consists of obtaining the relative spectral error $\varepsilon(f)$ which describes the ratio of the PSD of the error

*These portions are discarded for the present analysis, but do contain useful information for a preview-control-based gust load alleviation controller, see e.g. [14].




Fig. 6 Time signal produced to obtain the spectral error.

signal $\Phi_{\Delta w}$ over the PSD of the original turbulent wind, i.e.,

$$\varepsilon(f) = \frac{\Phi_{\Delta w}(f)}{\Phi_w(f)}. \quad (30)$$

Equation (30) describes the amount of conserved spectral power for a given frequency and, as shown in [14], represents the equivalent transfer function from the true vertical wind to the output of the wind reconstruction process. It includes the errors caused by the measurement volume[†], the lidar instrument errors, and the additional errors induced by the required regularization in the wind estimation algorithm. A relative error $\varepsilon(f)$ of zero indicates that the turbulent wind was perfectly reconstructed at this frequency, hence the spectral content of the error signal at this frequency is zero. On the other hand, if the relative error is 1, the spectral content of the difference signal at this frequency is the same as that of the original wind signal, so this frequency component was completely lost in the wind reconstruction process. Note that values $\varepsilon > 1$ are locally possible due to several reasons, e.g., because the interpolation process used to obtain the error signal Δw_W may introduce small spectral components at certain frequencies. The progression of $\varepsilon(f)$ for the example used in this section is shown in Fig. 7.

[†]Each measurement is not taken at an exact location but rather corresponds to an average line-of-sight relative wind velocity over a slender volume along the laser beam.

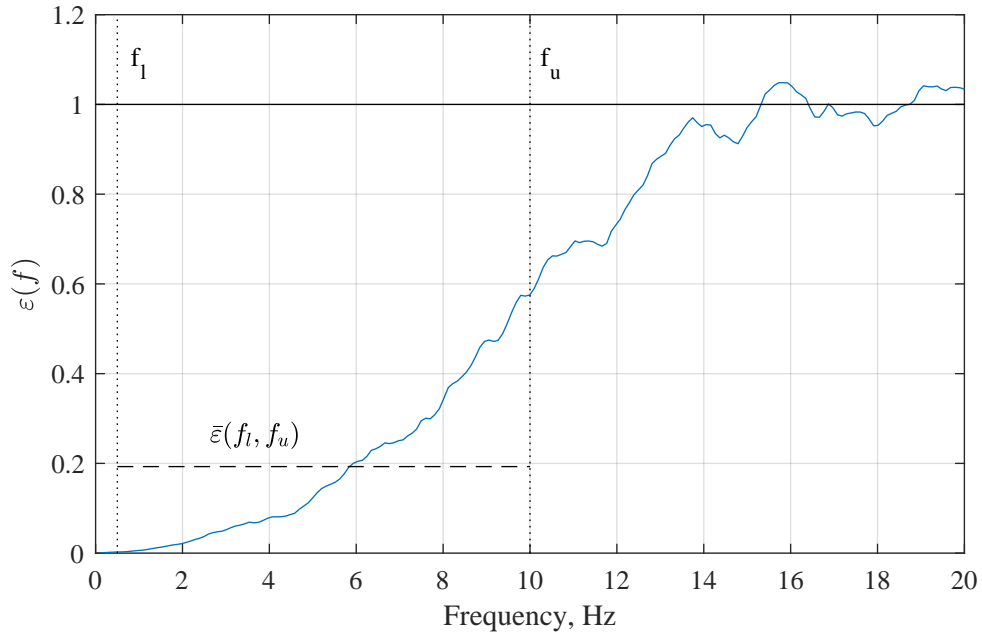


Fig. 7 Relative error $\varepsilon(f)$ and mean relative error $\bar{\varepsilon}$ between 0.5 and 10 Hz.

This formulation of the relative error characterizes the transfer function of the lidar sensor and the subsequent wind reconstruction process, i.e., the transfer function from the turbulent wind to the reconstructed wind profile. The knowledge of this transfer function can be used to tune a feedforward gust load alleviation controller, as shown in [14].

In order to obtain a single quantity which summarizes the performance of the lidar measurement and wind reconstruction process, the average of the relative error $\bar{\varepsilon}$ within a certain frequency band can be used:

$$\bar{\varepsilon} = \frac{1}{f_u - f_l} \int_{f_l}^{f_u} \varepsilon(f) df \quad (31)$$

where f_l and f_u are the lower and upper frequency bounds, respectively. This mean relative error is used as the primary indicator for the performance of lidar sensor and wind reconstruction algorithm configurations in this paper. For the presented analysis, $f_l = 0.5$ Hz and $f_u = 10$ Hz were chosen: below 0.5 Hz, the frequencies are low enough to be easily compensated by feedback control functions that are not using the lidar information. The lidar and wind reconstruction are expected to perform well at these frequencies, but their performance in this frequency range is not particularly relevant for the considered application and therefore not included in the chosen performance metric. Above 10 Hz, the frequencies are typically[‡] too high to be effectively countered by currently available actuators and are usually not relevant in terms of load levels.

[‡]Depending on the aircraft and its flight control systems.

V. Considered Test Case

A. Continuous Turbulence Scenario

In order to evaluate the empirical Bayesian approach outlined in Section III, three different methods to choose the regularization parameters are compared based on a simulated flight through continuous von Kármán turbulence. The chosen flight point represents high altitude cruise flight conditions of a typical business jet aircraft at an altitude of 40000 ft and a Mach number of 0.86 (true airspeed of approx. 254 m/s). The simulation duration is 1000 seconds with a wind estimation rate of 10 Hz, which yields 10000 individual estimation windows. The simulated wind field has a turbulence intensity (i.e., standard deviation) of 20 m/s and a spatial resolution of approx. 25 cm, which yields a Nyquist frequency of approx. 500 Hz.

The wind field was produced prior to the simulation using a random-phase method. The von Kármán turbulence spectrum has most of its spectral power concentrated at very low frequencies, so when running long simulations using this spectrum, the high-amplitude/low-wavelength components of the spectrum can cause high absolute wind speeds over several hundreds of meters. In order to avoid unrealistically high wind speeds, we eliminated all spectral components of more than 2500 m wavelength (i.e., all components below approx. 0.1 Hz at the current flight point). This way we avoid over-optimistic signal-to-noise ratios for the lidar, while the aeroelastic characteristics of the aircraft remain unaffected because this frequency region is far below the frequency of the short-period motion.

B. Compared Methods

Three methods to choose the regularization parameters are compared: first, fixed values $\gamma_1 = 0.128$ and $\gamma_2 = 0.22$ were used for all windows based on insights gained from the parameter study in [9]. This set of parameters will be referred to as *engineering guess* for the remainder of this paper.

In the second case, the values were individually determined using the *empirical Bayesian approach* described in Section III. For each estimation window, the marginal likelihood is optimized using a gradient based constraint non-linear optimizer [29]. Eventually, the obtained optimal hyperparameters are used for the wind reconstruction. This approach does not require the knowledge of the true wind field and thus could theoretically be applicable in an online setting. However, since each call of the marginal likelihood is computationally expensive, approximations of the optimum would likely be necessary for a real time application.

For the third case, a single set of optimal values γ_1 and γ_2 (same values used for all windows) is determined by minimizing the mean spectral error $\bar{\epsilon}$ between 0.5 Hz and 10 Hz as described in Eq. (31). Again, a gradient based constraint optimizer [29] was used for optimization. Lower bounds for γ_1 and γ_2 were specified in order to limit the jumps between subsequent estimation windows, as explained in subsection V.C. Unlike the first and second method, the

optimization needs the results of all estimation windows for the full simulation duration, and it requires total information about the actual wind field. Each call of the objective function thus requires re-estimating all estimation windows for a new set of weights. We call this approach *omniscient optimization*. The omniscient optimization is a purely theoretical case, as it requires full and exact information about the actual wind field, which cannot be acquired in-flight (making it unfeasible for online wind estimation). It only serves as a reference case for assessing the other two, as it provides the best achievable results (in the sense of the minimum mean spectral error) for a constant set of parameters γ_1 and γ_2 over all estimation windows.

The results and the comparison of the three different methods are presented in Section VI.

C. Parameter constraints

As explained in Section IV, the continuous output of the wind estimation is obtained by stitching partial estimated wind profiles from successive estimation windows together, where the spatial length Δx of every segment corresponds to the product of the aircraft's inertial speed V_K and the wind estimation update time, i.e., $\Delta x = V_K / f_{est} \approx 25.4$ m where $f_{est} = 10$ Hz is the refresh rate ("frequency") of the wind estimation process. Generally, actual continuity of the resulting combined signal is not ensured since individual windows can have different values at the start and end of the segment than the previous/subsequent segment, which can be seen by the small jumps between profiles in Fig. 5. We define the difference between two subsequent estimation windows at the boundary of their connection as

$$e_i = w_{est,i}(x_{bound,i}) - w_{est,i-1}(x_{bound,i}) \quad , \quad (32)$$

where $x_{bound,i}$ is the coordinate of the boundary between the i -th and $i + 1$ -th window, i.e.,

$$x_{bound,i} = \frac{i}{f_{est}} V_K. \quad (33)$$

Note that e_i is only defined for $i > 1$ since window $i = 0$ corresponds to $t = 0$, where no previous estimation window is available. The jumps described by e_i should be avoided to some degree since they introduce spectral components in the combined signal which have no physical meaning. However, they should not be removed via post-processing, but the regularization parameters should be chosen in a way that such discontinuities only occur to a low extent in the first place. In order to evaluate the impact of the regularization parameters on the extent of e_i , a full factorial variation of γ_1 and γ_2 is conducted and the mean of the absolute differences between subsequent window segments at the respective boundaries is determined for each combination

$$\bar{e} = \frac{1}{N_w - 1} \sum_{i=1}^{N_w-1} |e_i|, \quad (34)$$

where N_w is the number of wind estimation windows.

The progression of \bar{e} as a function of γ_1 and γ_2 is shown in Fig. 8. It is observable that the jumps tend to decrease quadratically with γ_1 and γ_2 . For a given target value of \bar{e} , the relationship of γ_1 and γ_2 turns out to be almost perfectly linear. This linear relationship allowed specifying a constraint for the *a-posteriori optimization* and the *Bayesian approach* in order to implicitly force an upper bound for \bar{e} . We decided to use $\bar{e} \leq 0.2 \text{ m/s}$ as a limit, which leads to the linear constraint

$$\gamma_2 \geq -0.291 \gamma_1 + 0.024. \quad (35)$$




Fig. 8 Mean absolute jumps \bar{e} as function of regularization parameters.

VI. Results

The performance of the three investigated methods for choosing the regularization parameters, quantified using the evaluation metric based on the spectral error, is shown in Fig. 9. The approach using fixed parameters clearly yields the worst results, as its spectral error is by far the highest, and about 25 % of the spectral content is lost already at 5 Hz. This clearly points to overregularization caused by a too high choice for γ_1 and γ_2 . Interestingly, the empirical Bayesian approach and the *a-posteriori optimization* method yield similar results in the frequency domain. The results using the *a-posteriori optimization* method are slightly better (when measured using the proposed evaluation metric), as can be seen by the vertical offset most pronounced at frequencies between 5 and 12 m/s. However, the only goal of this




Fig. 9 Comparison of the spectral error of the Bayesian approach against two references.

approach is to minimize the mean spectral error, whereas the empirical Bayesian approach does not explicitly consider the frequency-domain based evaluation metric at all – it is only implicitly affected due to the penalty on the complexity of the model as explained in subsection III.C. Even though these two approaches have different objective functions, both produce very similar results.

Figure 10 exemplarily shows the progression of the reconstructed time signals for a spatial length of 200 m. It can be seen that the *engineering guess* causes oversmoothing, as most of the peaks are not captured well. The empirical Bayesian approach and the omniscient optimization yield very similar results, but the omniscient case proves to be capable of following the true wind slightly better, especially at its peaks (e.g., at approx. 5505 m and 5615 m) due to the weaker regularization.

The distribution of regularization parameters of the three investigated approaches is shown in Fig. 11. The figure visualizes the joint distribution as well as the marginal distributions in form of histograms. Additionally, a zoom at the region of the linear constraint defined by Eq. (35) is plotted. In case of the omniscient optimization, the optimum is found for $\gamma_1 = 4e^{-7}$ and $\gamma_2 = 0.024$, whereas the median values obtained using the Bayesian approach are $\gamma_1 = 0.05$ and $\gamma_2 = 0.025$ (mean: $\gamma_1 = 0.066$, $\gamma_2 = 0.071$).

A possible explanation for the variation of parameters found by the Bayesian approach was sought in the inevitable local variation of spectral content between windows, but a deeper investigation[§] revealed that there is no obvious correlation

[§]This analysis is omitted in this paper for the sake of brevity.

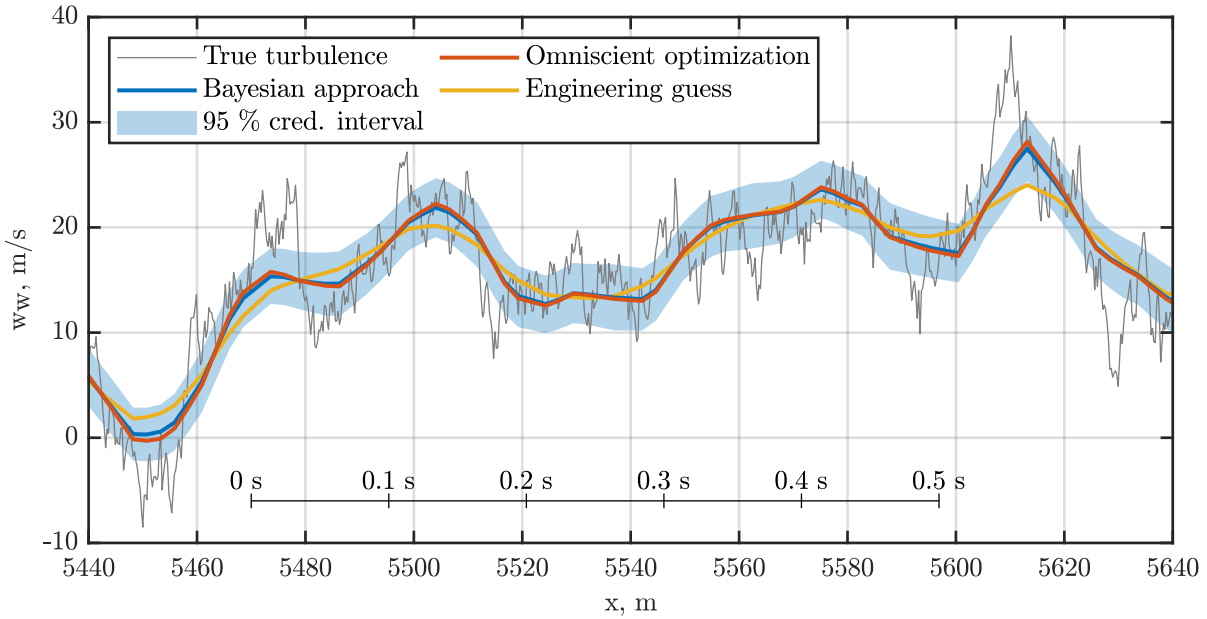


Fig. 10 Reconstructed wind profiles for the three tested methods (example window).

between the spectral content of the wind in the considered time window and the obtained values for γ_1 and γ_2 . This could be explained by the fact that the sensitivity of the reconstructed wind profiles in the time domain is relatively low for such low values of γ_1 and γ_2 , i.e., the parameter values are so small that even quite high relative changes do not affect the shape of the reconstructed signal (and hence the transfer function) in any significant way. Additionally, both parameters have somewhat similar effects: in the extreme case of $\gamma_1 = 0$ and $\gamma_2 \rightarrow \infty$, the resulting profile will be a straight line with any slope, but in the case of $\gamma_1 \rightarrow \infty$ and $\gamma_2 = 0$, the resulting profile is a straight line with zero slope.

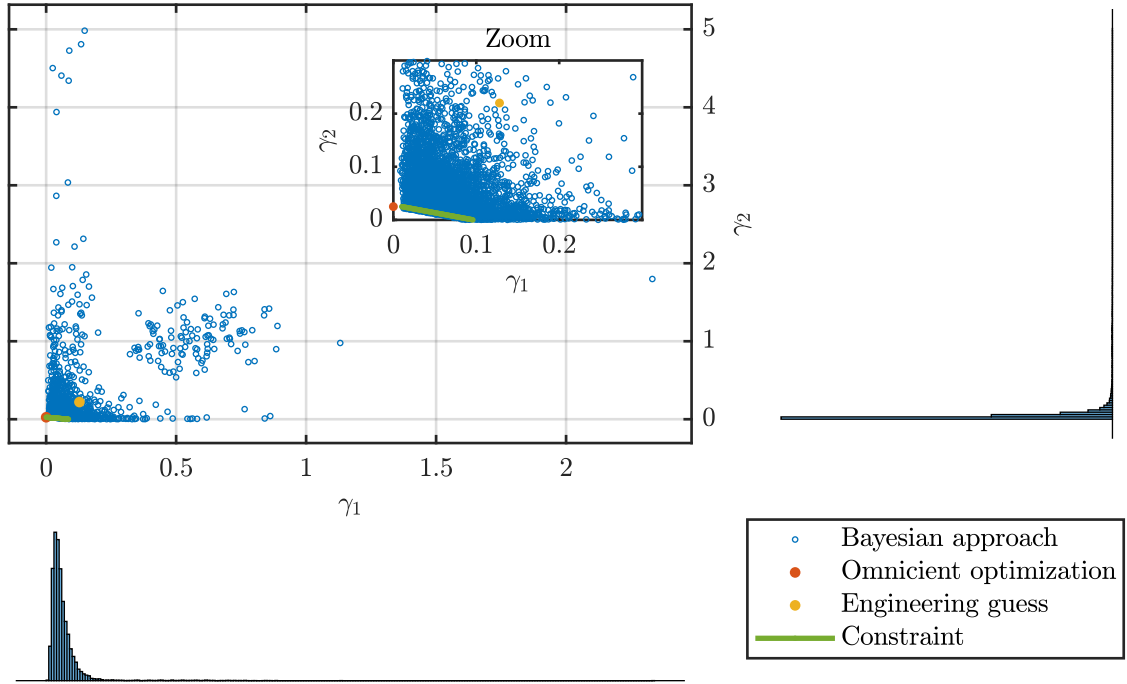


Fig. 11 Distribution and fixed values of the regularization parameters.

A. Effects of the parameter constraint

As explained in section V.C, a linear constraint was applied for the a-posteriori optimization and the empirical Bayesian approach to ensure a certain degree of smoothness of the time signal which is produced by combining the individual estimation windows to evaluate the spectral error. The upper bound for the “jumps” between subsequent wind profiles, which was used to derive the constraint (cf. Eq. (34)), can be chosen somewhat arbitrarily, so it is important to understand its impact on the optimization results.

Figure 12 shows the cost function of the a-posteriori optimization (i.e., the mean spectral error \bar{e} between 0.5 and 10 Hz) as a function of γ_1 and γ_2 . The white lines are isolines for the mean absolute jumps \bar{e} (cf. Sec. V.C). The red dot marks the unconstrained minimum, and the red cross marks the constrained minimum for $\bar{e} \leq 0.2$ m/s. Without constraints on the mean absolute jumps \bar{e} , the optimal point of the cost function is found at $\gamma_1 = 0$ and $\gamma_2 \approx 0.002$ (red dot in this figure). At this point, the mean absolute jump value \bar{e} is $\bar{e}_{opt} \approx 0.41$ m/s. Hence, setting an upper bound on \bar{e} lower than this value will lead to a constrained optimum that differs from the unconstrained one. Based on the respective orientations of (a) the gradient of the cost function and (b) the constraints on \bar{e} , the constrained optima will always lie on the ($\gamma_1 = 0, \gamma_2 \geq 0.002$) half-line. The optimum under the constraint $\bar{e} \leq 0.2$ m/s lies approximately at $(0, 0.024)$, as

already found numerically, and is marked with a red cross in this figure.

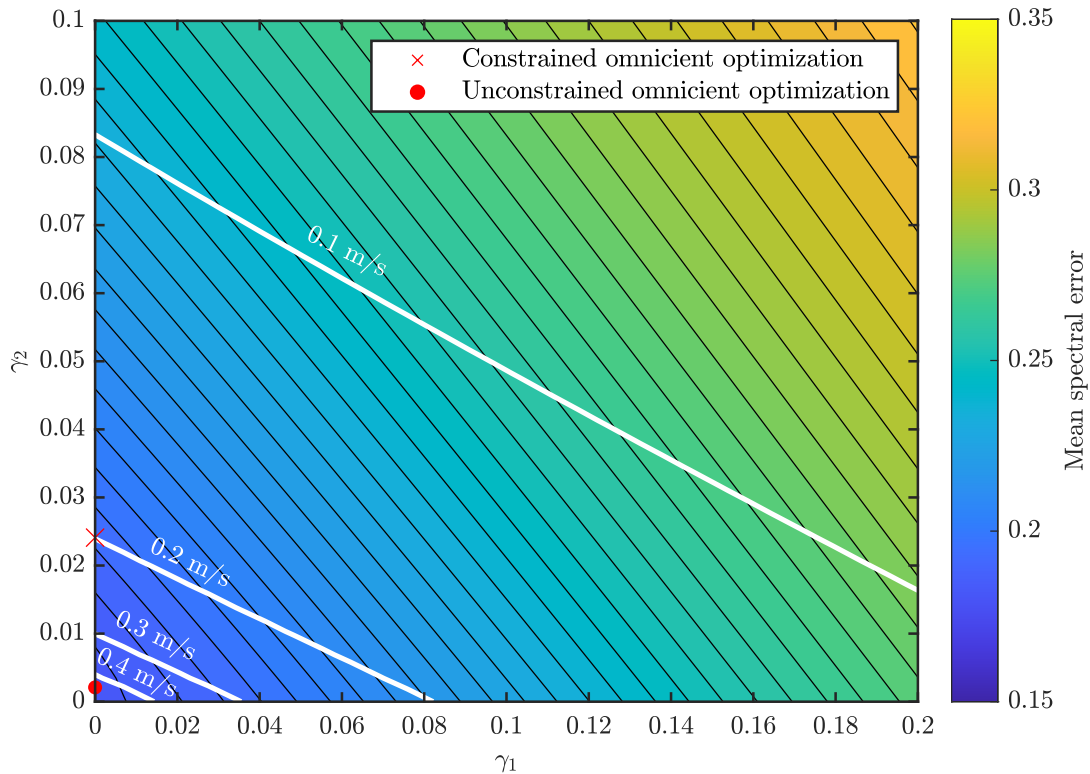


Fig. 12 Mean spectral error $\bar{\varepsilon}$ as a function of γ_1 and γ_2 .

VII. CONCLUSIONS

This paper presents a methodology to adaptively choose optimal regularization parameters based on an empirical Bayesian approach. It is shown that the developed method, which aims at optimizing the marginal likelihood without requiring any information about the wind field itself, is generally able to produce similar results as the a-posteriori method, which aims at reducing the error in the frequency domain and requires the full wind field information. It can be seen that the reconstructed wind profiles in the time domain (and hence their representations in the frequency domain) are not too sensitive w.r.t. variations of γ_1 and γ_2 , as long as these remain in the low-value range. This is supported by the fact that 97.5% of the optimized regularization parameters are in the range $\gamma_1, \gamma_2 \in [0.0, 0.5]$. The results using the a-posteriori optimization method are slightly better, i.e., on average 3.3% better conservation of spectral power, which is expected as it uses this metric explicitly in its cost function. An additional insight of the study is that the “engineering guess” for γ_1 and γ_2 was still too conservative, since the optimum obtained by the a-posteriori method yields much lower values and in terms of the error metric on average 15% better conservation of spectral power could be achieved.

A drawback of the Bayesian method is that the characterization of the transfer function of the lidar measurement and wind reconstruction process cannot be determined a priori, and hence cannot be directly used for controller tuning. In case of the omniscient optimizer, where an optimized set of regularization parameter values is found for all windows, the transfer function is well known. The parameter estimation with the omniscient optimizer could be made semi-adaptive by generating sets of optimal parameter values over a large number of points within an envelope (with varying flight conditions and atmospheric parameters) and using a selection logic similar to gain scheduling. This however, would still not be truly adaptive. The Bayesian method, on the other hand, is able to find optimal parameter values for any given atmospheric conditions and thus paves the way for adaptive wind field estimation for different atmospheric conditions.

Acknowledgements

We would like to acknowledge the funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy—EXC 2163/1—Sustainable and Energy Efficient Aviation – Project-ID 390881007.

Appendix A: Connection of the Bayesian approach to the regularized maximum likelihood solution

This paper presents a Bayesian approach to solve the inverse problem for wind field reconstruction, whereas previous works (e.g., [5, 9]) applied an approach based on the maximum likelihood solution with Tikhonov regularization. We now show the connection between these two approaches. The goal is to show that the maximum a posteriori (MAP) solution of the Bayesian approach given by Eq. (22) is equivalent to the maximum likelihood approach with Tikhonov regularization.

The MAP solution is defined as

$$\boldsymbol{\theta}_{MAP} = \arg \max_{\boldsymbol{\theta}} p(\boldsymbol{\theta}|\hat{\mathbf{z}}) \quad (36)$$

This can be achieved by minimizing the negative logarithmic posterior and hence

$$\boldsymbol{\theta}_{MAP} = \arg \min_{\boldsymbol{\theta}} \|\hat{\mathbf{z}} - \mathbf{A}\boldsymbol{\theta}\|^2 + \boldsymbol{\theta}^T \mathbf{Q}_{12} \boldsymbol{\theta} \quad (37)$$

$$= \arg \min_{\boldsymbol{\theta}} \|\hat{\mathbf{z}} - \mathbf{A}\boldsymbol{\theta}\|^2 + \gamma_1 \boldsymbol{\theta}^T \mathbf{Q}_1 \boldsymbol{\theta} + \gamma_2 \boldsymbol{\theta}^T \mathbf{Q}_2 \boldsymbol{\theta} \quad (38)$$

From the construction of the priors we know that

$$\boldsymbol{\theta}^T \mathbf{Q}_1 \boldsymbol{\theta} = \boldsymbol{\theta}^T \boldsymbol{\Gamma}_1^T \boldsymbol{\Gamma}_1 \boldsymbol{\theta} = (\boldsymbol{\Gamma}_1 \boldsymbol{\theta})^T \boldsymbol{\Gamma}_1 \boldsymbol{\theta} = \|\boldsymbol{\Gamma}_1 \boldsymbol{\theta}\|^2 \quad (39)$$

$$\boldsymbol{\theta}^T \mathbf{Q}_2 \boldsymbol{\theta} = \boldsymbol{\theta}^T \boldsymbol{\Gamma}_2^T \boldsymbol{\Gamma}_2 \boldsymbol{\theta} = (\boldsymbol{\Gamma}_2 \boldsymbol{\theta})^T \boldsymbol{\Gamma}_2 \boldsymbol{\theta} = \|\boldsymbol{\Gamma}_2 \boldsymbol{\theta}\|^2 \quad (40)$$

and thus we can also write Eq. (38) as:

$$\boldsymbol{\theta}_{MAP} = \arg \min_{\boldsymbol{\theta}} \|\hat{\mathbf{z}} - \mathbf{A}\boldsymbol{\theta}\|^2 + \gamma_1 \|\boldsymbol{\Gamma}_1 \boldsymbol{\theta}\|^2 + \gamma_2 \|\boldsymbol{\Gamma}_2 \boldsymbol{\theta}\|^2 \quad (41)$$

Equation (41) is exactly the expression for the maximum likelihood estimate with the Tikhonov regularization matrices $\boldsymbol{\Gamma}_1$, $\boldsymbol{\Gamma}_2$ and the corresponding weighting coefficients γ_1 and γ_2 .

References

- [1] Regan, C. D., and Jutte, C. V., "Survey of Applications of Active Control Technology for Gust Alleviation and New Challenges for Lighter-weight Aircraft," Tech. rep., NASA Dryden Flight Research Center, Edwards, CA, USA, April 2012. NASA/TM-2012-216008.
- [2] Livne, E., "Future of Airplane Aeroelasticity," *Journal of Aircraft*, Vol. 40, No. 6, 2003, pp. 1066–1092. <https://doi.org/10.2514/2.7218>.
- [3] Wildschek, A., Haniš, T., and Stroscher, F., " L_∞ -Optimal Feedforward Gust Load Alleviation Design for a Large Blended Wing Body Airliner," *Progress in Flight Dynamics, Guidance, Navigation, Control, Fault Detection, and Avionics*, Vol. 6, EDP Sciences, 2013, pp. 707–728. <https://doi.org/10.1051/eucass/201306707>.
- [4] Handoyo, V., "Contribution to Load Alleviation in Aircraft Pre-design and its Influence on Structural Mass and Fatigue," Ph.D. thesis, Technische Universität Berlin, 2021. URL <https://elib.dlr.de/139558/>.
- [5] Fezans, N., Joos, H.-D., and Deiler, C., "Gust load alleviation for a long-range aircraft with and without anticipation," *CEAS Aeronautical Journal*, Vol. 10, No. 4, 2019, pp. 1033–1057. <https://doi.org/10.1007/s13272-019-00362-9>.
- [6] Fezans, N., Vrancken, P., Linsmayer, P., Wallace, C., and Deiler, C., "Designing and Maturing Doppler Lidar Sensors for Gust Load Alleviation: Progress Made Since AWIATOR," *Aerospace Europe Conference 2020, Bordeaux, France*, 2020. URL <https://elib.dlr.de/134227/>.
- [7] Khalil, A. K. A., and Fezans, N., "Gust load alleviation for flexible aircraft using discrete-time H-infinity preview control," *The Aeronautical Journal*, Vol. 125, No. 1284, 2021, pp. 341–364. <https://doi.org/10.1017/aer.2020.85>.
- [8] Fournier, H., Massioni, P., Tu Pham, M., Bako, L., Vernay, R., and Colombo, M., "Robust Gust Load Alleviation of Flexible Aircraft Equipped with Lidar," *Journal of Guidance, Control, and Dynamics*, Vol. 45, No. 1, 2022, pp. 58–72. <https://doi.org/10.2514/1.G006084>.
- [9] Kiehn, D., Fezans, N., Vrancken, P., and Deiler, C., "Parameter Analysis of a Doppler Lidar Sensor for Gust Detection and Load Alleviation," *International Forum on Aeroelasticity and Structural Dynamics (IFASD)*, 13.-15.06.2022, Madrid, Spain, 2022. URL <https://elib.dlr.de/187627/>.
- [10] Knapik, B. T., Van Der Vaart, A. W., and van Zanten, J. H., "Bayesian inverse problems with Gaussian priors," *The Annals of Statistics*, Vol. 39, No. 5, 2011, pp. 2626 – 2657. <https://doi.org/10.1214/11-AOS920>.
- [11] Schmitt, N. P., Rehm, W., Pistner, T., Diehl, H., Navé, P., Jenaro-Rabadañ, G., Mirand, P., and Reymond, M., "Forward Looking Clear Air Turbulence Measurement with the AWIATOR LIDAR Sensor," *Proceedings of the 1st CEAS European Air and Space Conference*, Berlin, Germany, 2007, pp. 179–184. CEAS-2007-204.
- [12] Vrancken, P., "Airborne Remote Detection of Turbulence with Forward-Pointing LIDAR," *Aviation Turbulence – Processes, Detection, Prediction*, edited by R. Sharman and T. Lane, Springer, 2016, Chap. 22, pp. 443–464. <https://doi.org/10.1007/978-3-319-23630-8>.

- [13] Schmitt, N. P., Rehm, W., Pistner, T., Zeller, P., Diehl, H., and Navé, P., “The AWIATOR airborne LIDAR turbulence sensor,” *Aerospace Science and Technology*, Vol. 11, No. 7-8, 2007, pp. 546–552. <https://doi.org/10.1016/j.ast.2007.03.006>.
- [14] Cavaliere, D., Fezans, N., and Kiehn, D., “Method to Account for Estimator-Induced Previewed Information Losses – Application to Synthesis of Lidar-Based Gust Load Alleviation Functions,” *6th CEAS Conference on Guidance, Navigation and Control (EuroGNC), 3–5 May 2022, Berlin, Germany, 2022*. CEAS-GNC-2022-063.
- [15] Cavaliere, D., Fezans, N., Kiehn, D., Quero, D., and Vrancken, P., “Gust Load Control Design Challenge Including Lidar Wind Measurements and Based on the Common Research Model,” *AIAA SciTech 2022 Forum, San Diego, USA, 2022*. <https://doi.org/10.2514/6.2022-1934>, URL <https://arc.aiaa.org/doi/abs/10.2514/6.2022-1934>.
- [16] Fezans, N., Wallace, C., Kiehn, D., Cavaliere, D., and Vrancken, P., “Lidar-Based Gust Load Alleviation - Results Obtained on the Clean Sky 2 Load Alleviation Benchmark,” *International Forum on Aeroelasticity and Structural Dynamics (IFASD), 13.-15.06.2022, Madrid, Spain, 2022*. URL <https://elib.dlr.de/187462/>.
- [17] Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., and Rubin, D. B., *Bayesian data analysis*, CRC press, 2013. <https://doi.org/10.1201/b16018>.
- [18] Kaipio, J., and Somersalo, E., *Statistical and computational inverse problems*, Vol. 160, Springer Science & Business Media, 2006. <https://doi.org/10.1007/b138659>.
- [19] Bardsley, J. M., *Computational Uncertainty Quantification for Inverse Problems*, Society for Industrial and Applied Mathematics, 2018. <https://doi.org/10.1137/1.9781611975383>.
- [20] Lindgren, F., Rue, H., and Lindström, J., “An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, Vol. 73, No. 4, 2011, pp. 423–498. <https://doi.org/10.1111/j.1467-9868.2011.00777.x>.
- [21] Miller, D. L., Glennie, R., and Seaton, A. E., “Understanding the stochastic partial differential equation approach to smoothing,” *Journal of Agricultural, Biological and Environmental Statistics*, Vol. 25, No. 1, 2020, pp. 1–16. <https://doi.org/10.1007/s13253-019-00377-z>.
- [22] Fezans, N., Schwithal, J., and Fischenberg, D., “In-flight Remote Sensing and Characterization of Gusts, Turbulence and Wake Vortices,” *Deutscher Luft- und Raumfahrtkongress 2015, Rostock, Germany, 22nd – 24th September 2015, 2015*.
- [23] Fezans, N., Schwithal, J., and Fischenberg, D., “In-flight Remote Sensing and Characterization of Gusts, Turbulence and Wake Vortices Using a Doppler LIDAR,” *CEAS Aeronautical Journal*, Vol. 8, No. 2, 2017, pp. 313–333. <https://doi.org/10.1007/s13272-017-0240-9>.
- [24] Lindley, D. V., and Smith, A. F., “Bayes estimates for the linear model,” *Journal of the Royal Statistical Society: Series B (Methodological)*, Vol. 34, No. 1, 1972, pp. 1–18. URL <https://www.jstor.org/stable/2985048>.

- [25] Williams, C. K., and Rasmussen, C. E., *Gaussian processes for machine learning*, Vol. 2, MIT Press Cambridge, MA, 2006. https://doi.org/10.1007/978-3-540-28650-9_4.
- [26] Teckentrup, A. L., “Convergence of Gaussian process regression with estimated hyper-parameters and applications in Bayesian inverse problems,” *SIAM/ASA Journal on Uncertainty Quantification*, Vol. 8, No. 4, 2020, pp. 1310–1337. <https://doi.org/10.1137/19M1284816>.
- [27] European Union Aviation Safety Agency (EASA), *Certification Specifications and Acceptable Means of Compliance for Large Aeroplanes (CS-25), Amendment 27*, 2023.
- [28] Hoblit, F. M., *Gust Loads on Aircraft: Concepts and Applications*, AIAA Education Series, Washington, D.C., 1988. <https://doi.org/10.2514/4.861888>.
- [29] Nocedal, J., Öztoprak, F., and Waltz, R. A., “An interior point method for nonlinear programming with infeasibility detection capabilities,” *Optimization Methods and Software*, Vol. 29, No. 4, 2014, pp. 837–854. <https://doi.org/10.1080/10556788.2013.858156>.