Experimental and numerical optimizations of an upwind mainsail trimming

Matthieu Sacher, Frédéric Hauville, Régis Duvigneau, Olivier Le Maître, Nicolas Aubin

THE 22nd CHESAPEAKE SAILING YACHT SYMPOSIUM
ANNAPOLIS, MARYLAND, MARCH 2016

Experimental and numerical optimizations of an upwind mainsail trimming

Matthieu Sacher[1], Naval Academy Research Institute – IRENav, France
Frédéric Hauville, Naval Academy Research Institute – IRENav, France
Régis Duvigneau, INRIA – Sophia Antipolis, France
Olivier Le Maître, LIMSI – CNRS, France
Nicolas Aubin, Naval Academy Research Institute – IRENav, France
Mathieu Durand, K-EPSILON, France

ABSTRACT

This paper investigates the use of meta-models for optimizing sails trimming. A Gaussian process is used to robustly approximate the dependence of the performance with the trimming parameters to be optimized. The Gaussian process construction uses a limited number of performance observations at carefully selected trimming points, potentially enabling the optimization of complex sail systems with multiple trimming parameters. We test the optimization procedure on the (two parameters) trimming of a scaled IMOCA mainsail in upwind conditions. To assess the robustness of the Gaussian process approach, in particular its sensitivity to error and noise in the performance estimation, we contrast the direct optimization of the physical system with the optimization of its numerical model. For the physical system, the optimization procedure was fed with wind tunnel measurements, while the numerical modeling relied on a fully non-linear Fluid-Structure Interaction solver. The results show a correct agreement of the optimized trimming parameters for the physical and numerical models, despite the inherent errors in the numerical model and the measurement uncertainties. In addition, the number of performance estimations was found to be affordable and comparable in the two cases, demonstrating the effectiveness of the approach.

INTRODUCTION

Researches on sailing yachts have fostered the development of advanced methods dedicated to the prediction and improvement of racing yacht performance. The performance is usually analyzed using so-called Velocity Prediction Programs (VPPs) (Oossanen, 1993), which solve equilibrium equations (balancing hull, appendages and sails loads) to determine several performance indicators, such as Boat Speed (BS) or Velocity Made Good (VMG). The different loads accounted by the VPPs can be based on empirical formulas, experimental data or numerical simulations (Hansen et al., 2003, Korpus, 2007). However, due to the complexity and multi-physic characters of the yachts dynamics, performance studies often consider the hydrodynamic (Huetz and Guillerm, 2014) and aerodynamic (Augier et al., 2012, Menotti et al., 2013, Trimarchi, 2012) aspects separately.

Here, we focus on the aerodynamics, optimizing the performance of a sail system, but the numerical procedure developed below can be used to perform hydrodynamic optimization or even fully coupled yacht performance optimization. Sail systems are subjected to very complex phenomena, such as nonlinear Fluid-Structure Interaction (FSI) effects and instabilities. Moreover, the modeling of real sailing conditions is still an open research problem due to the large uncertainties in wind and sea states. To our knowledge, the optimization of sails has thus been limited so far to idealized situations. For instance, sail shape optimizations (without accounting for the FSI problem) are performed in (Rousselon, 2008), while the trimming of two-dimensional sails is numerically considered in (Chapin et al., 2008). Regarding FSI in three spatial dimensions, the authors in (Ranzenbach et al., 2013) mention an optimization of the trimming of sails, but within an inviscid flow approximation and not much details are provided on the optimization procedure used.

The present work aims at pursuing these efforts toward the development of efficient numerical optimization procedures capable of dealing with complex sail systems, with realistic physical models (*e.g.* nonlinear FSI and turbulent flows) and a large number of optimization variables (*i.e.* trimming parameters). Denoting $\mathbf{x} \in \Omega$ the optimization variables, the optimization problem can be written as

$$\mathbf{x}_{\text{opt}} = \arg\min_{\mathbf{x}\in\Omega} -\mathcal{P}(\mathbf{x}),$$

where $\mathbf{x}_{\text{opt}} \in \Omega$ are the sought optimal parameters and $\mathcal{P} : \Omega \mapsto \mathbb{R}$ is a measure of performance. The main dif-

[1]matthieu.sacher@ecole-navale.fr

ficulty preventing the straightforward application of standard optimization procedures to sail systems is related to the cost of estimating of the performance $\mathcal{P}$ at tentative values $\mathbf{x}$ of the parameters. Indeed, the estimation of $\mathcal{P}(\mathbf{x})$ involves the resolution of a nonlinear FSI problem, which requires several convergence iterations between the nonlinear elastic and flow solvers. Also, adjoint based techniques are hardly amenable to FSI problems involving coupled nonlinear solvers; this fact precludes the use of efficient gradient-based iterative methods in favor of optimization algorithms such as the simplex based (Nelder and Mead, 1965) or evolutionary (Bäck and Schwefel, 1993, Hansen, 2006) methods. However, depending on the considered problem, these so called gradient-free algorithms are known to require a large number of evaluations of $\mathcal{P}(\mathbf{x})$, making applications to sail systems very costly as a single evaluation may routinely require several hours of CPU on modern parallel computers.

Based on these observations, we advocate the use of meta-model approaches to mitigate the large computational cost of optimizing the trimming parameters of sail systems. Specifically we rely below on Gaussian Process (GP) approximations for the mapping $\mathcal{P} : \Omega \mapsto \mathbb{R}$. This statistical approach uses a coarse set of performance evaluations at some parameters values $\mathbf{x} \in \Omega$ to infer a GP $\mathcal{G}(\mathbf{x}) \approx \mathcal{P}(\mathbf{x})$. One can then apply its favorite optimization procedure to $\mathcal{G}(\mathbf{x})$ to obtain the corresponding approximation of $\mathbf{x}_{\texttt{opt}}$. The surrogate-based optimization procedure is embedded in an iterative scheme, where new evaluations of the performance at carefully selected new points $\mathbf{x}$ are introduced in order to refine the GP approximation in regions of Ω susceptible to include the optimum. The GP approach is then expected to improve the optimization by a) requiring an overall lower number of performance evaluations, compared to direct gradient-free approaches, and b) enabling the use of efficient global optimization tools. In addition, the GP construction provides a natural way to estimate convergence on the approximation of $\mathcal{P}$ and then to characterize the accuracy on the retrieved optimum.

Another interest in considering an optimization based on GP meta-model is that it naturally accommodates for errors and noise in the performance evaluation. This specificity is exploited in the present work to perform the optimization of an actual physical sail system, consisting in a scaled IMOCA mainsail in upwind conditions. The objective is to find the optimal trimming of the sail, for a performance criterion combining the drive and side aerodynamic force coefficients. Here, the GP-based optimizer used values of $\mathcal{P}(\mathbf{x})$ measured in the wind tunnel of the Yacht Research Unit (Auckland), for the sequence of trimming points requested by the iterative optimization procedure. Because of the imperfections in the experimental apparatus and inherent noise in the measurements, the estimates of $\mathcal{P}(\mathbf{x})$ were subjected to significant errors, that would have compromised the convergence of descent methods (Saul'ev and Samoilova, 1975) without using a GP reconstruction.

In addition to evidence the robustness of GP based optimization, this experiment is used as a reference to assess the relevance of an optimization relying on numerical resolutions of the FSI problem to compute the performance. To this end, the experimental sail system was measured (dimensions, mechanical characteristics of mast and boom, ...), and wind tunnel inflow conditions recorded, to create a numerical model of the experiment. State of the art FSI solvers is then used for the resolution of the resulting numerical model at the sequence of trimming points requested by the GP-based optimizer. The numerical resolution involves a nonlinear structural solver with a mesh deformation utility (K-FSI tools) developed by K-EPSILON, coupled with the finite volume turbulent flow solver FINE$^{\mathrm{TM}}$/Marine from Numeca Software. The Unsteady Reynolds-Average Navier-Stokes Equations (URANS) turbulence model is used in these numerical experiments.

The paper is organized as follows. The first Section briefly reviews the construction of the Gaussian Process approximation and the selection of the new parameters in the iterative optimization procedure. We then detail the experimental set-up and the results of the corresponding optimization procedure in the second Section. The numerical modeling of the experiment and optimization results are reported in third Section. Finally, conclusions of this work and direction for future developments are provided in the fourth Section.

GP MODEL BASED OPTIMIZATION

In this Section we start by briefly summarizing the construction of a Gaussian Process to model a function from noisy measurements. More details on GP models can be found for instance in (Gibbs, 1997, Rasmussen and Williams, 2006). We then describe the GP model optimization procedure (Duvigneau and Chandrashekar, 2012), detailing the selection of successive optimal candidates.

Gaussian Process Model

Consider a dataset $\mathbf{X}_n = (\mathbf{x}_1 \cdots \mathbf{x}_n)^{\mathrm{T}}$ of n training inputs vectors $\mathbf{x}_i \in \Omega \subset \mathbb{R}^d$. Each element $\mathbf{x}_i \in \mathbf{X}_n$ is associated to an observation (or measurement) $y_i \in \mathbb{R}$ which is assumed to be dependent on a latent function $f(\mathbf{x})$ through

$$y_i = f(\mathbf{x}_i) + \varepsilon_i, \quad i = 1, \ldots, n, \tag{1}$$

where ε_i is a random measurement error (*i.e.* the measurement noise). In this work, the ε_i are assumed *independent* and to follow the same (centered) Gaussian distribution:

$$\varepsilon_i \sim \mathcal{N}\left(0, {\sigma_\epsilon}^2\right), \tag{2}$$

where $\mathcal{N}\left(\mu, \Sigma^2\right)$ denotes the Gaussian distribution with mean μ and variance Σ^2. Thus, ${\sigma_\epsilon}^2$ is referred to as noise variance. The objective is therefore to model the latent function $f(\mathbf{x})$ on the basis on the noisy observations y_i.

The latent function is considered as a realization of a zero-mean multivariate Gaussian process F, with unknown covariance function C_F, that is $F \sim \mathcal{N}(0, C_F)$, with

$$C_F(\mathbf{x}, \mathbf{x}') \doteq \mathbb{E}\{F(\mathbf{x}), F(\mathbf{x}')\}, \tag{3}$$

where $\mathbb{E}\{\cdot\}$ denotes the expectation operator.

The covariance function of F must be specified. In this work, we consider the Matérn class (Stein, 2012) of stationary covariance functions having for one-dimensional generator

$$M_\nu(r, l) = \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{\sqrt{2\nu} r}{l} \right)^\nu K_\nu \left(\frac{\sqrt{2\nu} r}{l} \right). \tag{4}$$

Here $r = |x - x'|$, ν and l are two positive parameters, and K_ν is the modified Bessel function of the second kind. We shall further restrict ourselves to covariances with $\nu \to \infty$, leading to the squared exponential covariance family with generator

$$M_\infty(r, l) = \exp\left(\frac{-r^2}{2l^2} \right). \tag{5}$$

The multidimensional counterpart is obtained by tensor product of the one-dimensional generator. The final expression of the covariance function for the GP approximation is

$$c(\mathbf{x}, \mathbf{x}'; \boldsymbol{\Theta}) = \theta_1 \prod_{i=1}^{d} \exp\left(\frac{-(x_i - x_i')^2}{2 l_i^2} \right) + \theta_2. \tag{6}$$

In the expression (6) of c, $\boldsymbol{\Theta} = \{\theta_1, \theta_2, l_1, l_2, \ldots, l_d\}$ is a vector of hyper-parameters. The first hyper-parameter, θ_1, scales the distance-dependent correlation, while θ_2 is an offset from zero. The other parameters l_i are the anisotropic correlation lengths associated to the d directions of Ω. From the parametrized covariance function $c(\mathbf{x}, \mathbf{x}'; \boldsymbol{\Theta})$ we derive the covariance matrix $\mathbf{C}(\boldsymbol{\Theta}) \in \mathbb{R}^{n \times n}$ of the observation points in $\mathbf{X}_n$. The covariance matrix $\mathbf{C}(\boldsymbol{\Theta})$ has for entries

$$C_{i,j}(\boldsymbol{\Theta}) \doteq c(\mathbf{x}_i, \mathbf{x}_j; \boldsymbol{\Theta}), \quad 1 \leq i, j, \leq n. \tag{7}$$

Given the n noisy observations y_i, collected into the vector $\mathbf{Y}_n = (y_1 \cdots y_n)^{\mathrm{T}}$, the predicted observation $y(\mathbf{x})$ at a new point $\mathbf{x} \in \Omega$ is given by the joint Gaussian distribution

$$\begin{pmatrix} \mathbf{Y}_n \\ y(\mathbf{x}) \end{pmatrix} \Bigg| \, \mathbf{X}_n, \boldsymbol{\Theta} \sim \mathcal{N}\left(\mathbf{0}, \begin{bmatrix} \mathbf{C} + {\sigma_\epsilon}^2 \mathbf{I} & \mathbf{k}(\mathbf{x}) \\ \mathbf{k}^{\mathrm{T}}(\mathbf{x}) & \kappa(\mathbf{x}) + {\sigma_\epsilon}^2 \end{bmatrix} \right). \tag{8}$$

In (8), the dependence of $\mathbf{C}$ on the hyper-parameters has been removed to simplify the notation, and

$$\kappa(\mathbf{x}) \doteq c(\mathbf{x}, \mathbf{x}; \boldsymbol{\Theta}), \quad \mathbf{k} \doteq (c(\mathbf{x}, \mathbf{x}_1; \boldsymbol{\Theta}) \cdots c(\mathbf{x}, \mathbf{x}_n; \boldsymbol{\Theta}))^{\mathrm{T}},$$

while $\mathbf{I}$ is the identity of $\mathbb{R}^n$. Using the conditional rules of a joint Gaussian distribution (Rasmussen and Williams, 2006, Von Mises, 1964), it comes

$$y(\mathbf{x}) | \, \mathbf{Y}_n, \mathbf{X}_n, \boldsymbol{\Theta}, {\sigma_\epsilon}^2 \sim \mathcal{N}\left(\hat{y}(\mathbf{x}), \hat{\sigma}_y^2(\mathbf{x}) \right). \tag{9}$$

The *best* prediction of $y(\mathbf{x})$ is the mean $\hat{y}(\mathbf{x})$ of the distribution; the prediction variance $\hat{\sigma}_y^2(\mathbf{x})$ quantifies the uncertainty in the prediction. The second order properties of the prediction $y(\mathbf{x})$ can be explicitly expressed as

$$\hat{y}(\mathbf{x}) = \mathbf{k}^{\mathrm{T}}(\mathbf{x}) \left(\mathbf{C}(\boldsymbol{\Theta}) + {\sigma_\epsilon}^2 \mathbf{I} \right)^{-1} \mathbf{Y}_n, \tag{10}$$

$$\hat{\sigma}_y^2(\mathbf{x}) = \kappa(\mathbf{x}) + {\sigma_\epsilon}^2 - \mathbf{k}^{\mathrm{T}}(\mathbf{x}) \left(\mathbf{C}(\boldsymbol{\Theta}) + {\sigma_\epsilon}^2 \mathbf{I} \right)^{-1} \mathbf{k}(\mathbf{x}). \tag{11}$$

The hyper-parameters $\boldsymbol{\Theta}$ and noise variance ${\sigma_\epsilon}^2$ are unknown *a priori* and need to be learned from the data. They can be determined by maximizing the log-marginal likelihood (Rasmussen and Williams, 2006) given by

$$\begin{aligned} \mathcal{L}(\boldsymbol{\Theta}, {\sigma_\epsilon}^2) = & -\frac{n}{2} \log(2\pi) - \frac{1}{2} \log \left| \mathbf{C}(\boldsymbol{\Theta}) + {\sigma_\epsilon}^2 \mathbf{I} \right| \\ & - \frac{1}{2} \mathbf{Y}_n^{\mathrm{T}} \left(\mathbf{C}(\boldsymbol{\Theta}) + {\sigma_\epsilon}^2 \mathbf{I} \right)^{-1} \mathbf{Y}_n. \end{aligned} \tag{12}$$

The optimal hyper-parameters and noise variance are then found by minimizing $\mathcal{L}$ with respect to its arguments. An evolution strategy algorithm (Hansen, 2006) is used for this purpose.

Once the optimal hyper-parameters are determined, the GP model can be used to predict values at new points using (10) and (11). The most computationally consuming part of the GP construction is the assembly of the (full) matrix $\left(\mathbf{C} + {\sigma_\epsilon}^2 \mathbf{I} \right)$, and the evaluation of its determinant and inverse required in the definition of the log-marginal likelihood and for new predictions. This can be done efficiently by LU decompositions.

Figure 1 illustrates for a one-dimensional function the effect of the observation noise σ_ϵ on the constructed GP model. The constructions use 6 observations points depicted with circles in the plots and the covariance hyper-parameters are determined by maximizing the log-marginal likelihood. However, in the first case, shown in Figure 1(a) a value $\sigma_\epsilon = 0$ is imposed, while in the second case in Figure 1(b) the noise level is also optimized. In addition to the observations, the two plots report the mean of the GP models with classical $\pm 3\hat{\sigma}$ uncertainty range. It is seen from Figure 1(a) that when the measurements are assumed to be noise-free ($\sigma_\epsilon = 0$), the resulting GP model is interpolating the data, *i.e.* the variance of the prediction is zero at the data points. However, the mean of the GP model exhibits significant oscillations such that over-fitting can be suspected. On the contrary, optimizing the noise level σ_ϵ results in a mean process free of spurious oscillations but that is no more interpolating, as it can be appreciated from Figure 1(b). The averaged distance of the best prediction to the observations is $\sim \sigma_\epsilon$.

Optimization strategy using Gaussian Process models

The GP model (build, previously) can be used to determine the next control parameters $\mathbf{x}_*$ to be included in the data

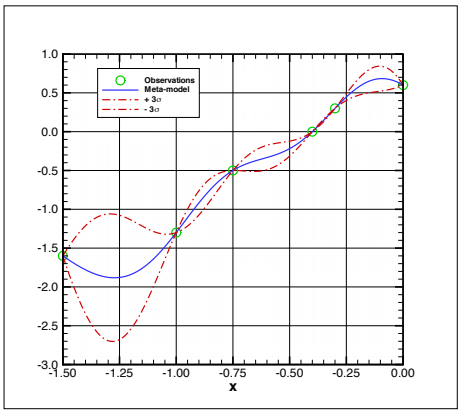

(a) Assuming $\sigma_\epsilon = 0$.

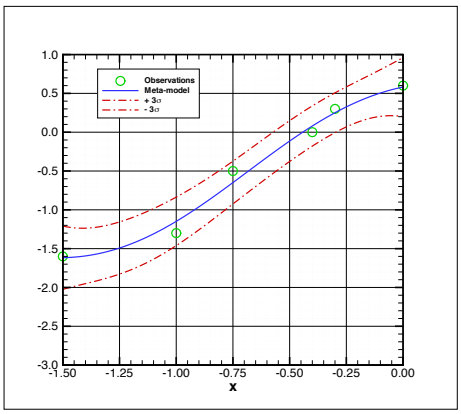

(b) Optimal $\sigma_\epsilon = 0.143$.

Figure 1: Effect of σ_ϵ on the GP model.

base. Deterministic optimization approaches would classically choose $\mathbf{x}_*$ as the best control parameters, that is the minimizer of the meta-model over Ω. However, GP models are random and not even bounded, so the definition of the best control parameters needs to be clarified in this context. This is classically achieved by introducing an appropriate (deterministic) merit function, and combining the expected prediction $\hat{y}$ and its variance $\hat{\sigma}_y^2$, requiring $\mathbf{x}_*$ to be the maximizer of this merit function. In fact, the merit function should balance a selection of $\mathbf{x}_*$ yielding the minimal expected prediction $\hat{y}$ (optimality) with a selection of $\mathbf{x}_*$ in areas of large variance $\hat{\sigma}_y^2$ to reduce the GP model uncertainty. A complete summary of various merit functions proposed in the literature is provided in (Jones, 2001). In this work, we use the Augmented Expected Improvement (AEI) merit function (Huang et al., 2006) which is an extension of the popular Expected Improvement (EI) (Jones et al., 1998) in the case of noisy estimations. The AEI merit function $AEI(\mathbf{x})$ estimates the expected increase in the performance, taking into account the noise in the observed values and penalizing areas where the variance $\hat{\sigma}_y^2$ is small. It writes as

$$AEI(\mathbf{x}) = EI(\mathbf{x}) \left(1 - \frac{\sigma_\epsilon}{\sqrt{\hat{\sigma}_y^2(\mathbf{x}) + {\sigma_\epsilon}^2}} \right), \qquad (13)$$

with the Expected Improvement defined by

$$\begin{aligned} EI(\mathbf{x}) &= \hat{\sigma}_y^2(\mathbf{x}) \left[u(\mathbf{x}) \Phi\left(u(\mathbf{x})\right) + \phi\left(u(\mathbf{x})\right) \right], \\ u(\mathbf{x}) &= \frac{\hat{y}(\mathbf{x}_{**}) - \hat{y}(\mathbf{x})}{\hat{\sigma}_y^2(\mathbf{x})}, \end{aligned} \qquad (14)$$

where Φ and $\phi = \Phi'$ denote respectively the cumulative distributions (Erf-function) and density of the standard Gaussian distribution, and $\mathbf{x}_{**}$ is the effective best solution:

$$x_{**} \doteq \arg\min_{\mathbf{x} \in \Omega} \left[\hat{y}(\mathbf{x}) + \hat{\sigma}(\mathbf{x}) \right]. \qquad (15)$$

When the optimum $\mathbf{x}_*$ of the AEI is determined, the corresponding performance $\mathcal{P}(\mathbf{x}_*)$ is evaluated and is included in the database. A new GP model can then be reconstructed using the extended data base, leading to a new maximizer $\mathbf{x}_*$ of the updated AEI, and so on. The iterations carry on until some convergence criterion is satisfied or the resources allocated to the optimization procedure have been exhausted (*e.g.* reaching a prescribed number of performance evaluations). Classical convergence criteria compare the distance between two successive iterates, in terms of optimal parameters $\mathbf{x}_*$ or/and performance prediction $\mathcal{P}(\mathbf{x}_*)$. Overall, we remark that each iteration requires the resolution of two optimization problems (one for the covariance parameters and one for the AEI) and one evaluation of the performance. In practice, all results reported in this work were obtained using the nonlinear non-convex black-box optimization library based on the Covariance Matrix Adaptation Evolution Strategy (Hansen, 2006, Hansen and Ostermeier, 2001).

EXPERIMENTAL OPTIMIZATION

This Section concerns the sail trimming optimization performed in the wind tunnel of the Yacht Research Unit (YRU) (Flay, 1996) at the University of Auckland.

Experimental setup

The sail model is inspired by an IMOCA 60-foot design mainsail at 1:13 scale. It was designed and produced by the sail-makers of INCIDENCES SAILS company. The sail has a surface area of 1 m^2, for an height of $h = 2$ m, and is supported by a rig consisting of a flexible circular section carbon mast (constant diameter 14 mm), clamped at its base and without spreader, backstay or forestay. The sail and rig are set in the open jet test section of the YRU wind tunnel, see Figure 2(a). The test section is 7.2 wide and 3.5 m high.

Three stepper motors and a control card control remotely the main sheet length (L_{sheet}) and main car position (L_{car})

as shown in Figure 2(b). In the following, L_{sheet} and L_{car} are the only trimming parameters to be optimized. Note that the remote system allows for changing these trimming parameters without switching off the wind tunnel flow and making a new tare of the measurement instruments. A precision of ± 2 mm on the imposed trimming parameters was estimated through repeated measurements.

(a) Mainsail in the YRU wind tunnel.

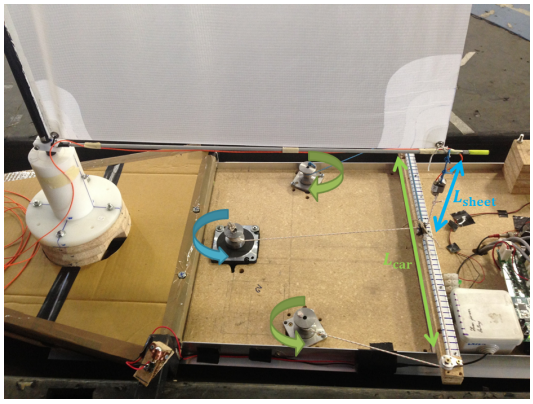

(b) Close view on the trimming system.

Figure 2: Experimental setup.

A six-components force balance, located under the floor of the wind tunnel, was used to measure the aerodynamic forces. The X-direction corresponds to the model longitudinal forward direction (*i.e.* thrust force direction), while the Y-direction is defined as the positive port-side (*i.e.* side force direction) and the Z-direction is the vertical. After careful calibration and testing, the balance precisions in the X, Y and Z directions were estimated to be ± 0.09 N, ± 0.11 N and ± 0.27 N respectively. An additional load sensor, with 5 daN range, was used to record the load in the sheet with a precision of ± 0.02 N. Flying shapes were also recorded with a V-SPARS acquisition system (Le Pelley and Modral, 2008), tracking the position of five dark red stripes across the sail (see Figure 2(a)).

The wind tunnel inflow velocity was measured and found to have an apparent wind speed (AWS) of 3.5 ± 0.15 m/s for an apparent wind angle (AWA) of 40 ± 2 deg. The corresponding Reynolds number, based on the reference chord length $c = S/h = 0.5$ m, is $\text{Re} = 1.2 \times 10^5$. A multi-hole pressure probe (Cobra Probe) was used to measure profiles of the flow velocity at several locations inside the wind tunnel. These measurements were repeated with and without the sail model in the test section to verify its effect on the flow field. Typical profiles are shown in Figure 3; it can be seen that the inflow has no twist.

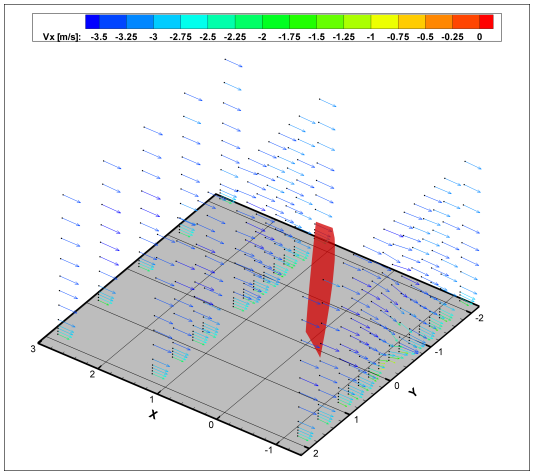

Figure 3: Velocity profiles with sail model.

The optimization problem is finally defined as the maximization with respect to $\mathbf{x} = (L_{\text{sheet}}, L_{\text{car}})$ of the performance $\mathcal{P}(\mathbf{x})$ taken as a composite function of the thrust C_X and side C_Y aerodynamic coefficients:

$$\mathcal{P}(\mathbf{x}) = C_X + 0.1 C_Y. \tag{16}$$

The aerodynamic coefficients are deduced from the aerodynamic forces through normalization by the reference force being $q_\infty S$, where q_∞ is the reference dynamic pressure measured in the wind tunnel (precision ± 1 Pa). In (16), the coefficient 0.1 penalizes the side force to account for the resulting hydrodynamic drag and leeway that would be detrimental to the performance. The optimization of the trimming parameters then follows the procedure illustrated in Figure 4. The primary loop, *"Sampling loop"*, generates an initial Latin Hypercube Sample (LHS) set of 10 trimming parameters (McKay et al., 2000). For each initial sample, the experimental model is remotely set to the corresponding trimming values of L_{sheet} and L_{car}. After the transient flow is over, the aerodynamic loads reported by the balance are averaged over an acquisition time of 30 s to smooth-out the remaining noisy fluctuations in the signals. When the (time-averaged) loads are collected for all the initial samples (blue loop in Figure 4), the GP based optimization is carried out on the initial data set (red block *"Trimming optimization"*). The optimization provides a new trimming point, $\mathbf{x}_*$, maximizing the AEI merit function.

In the *"Optimization loop"*, the experimental apparatus is remotely tuned to the new trimming point $\mathbf{x}_*$. The aerodynamic loads are then averaged over 30 s and the new estimate of $\mathcal{P}(\mathbf{x}_*)$ is included in the database. The Gaussian Process model is updated consequently, generating a new trimming points $\mathbf{x}_*$. These steps (red loop in Figure 4) are repeated until the convergence of the trimming parameters, which is considered achieved when the algorithm proposes two successive trimming points within a distance less than 1 mm (*i.e* the precision on the enforcement of the trimming parameters in the experimental apparatus).

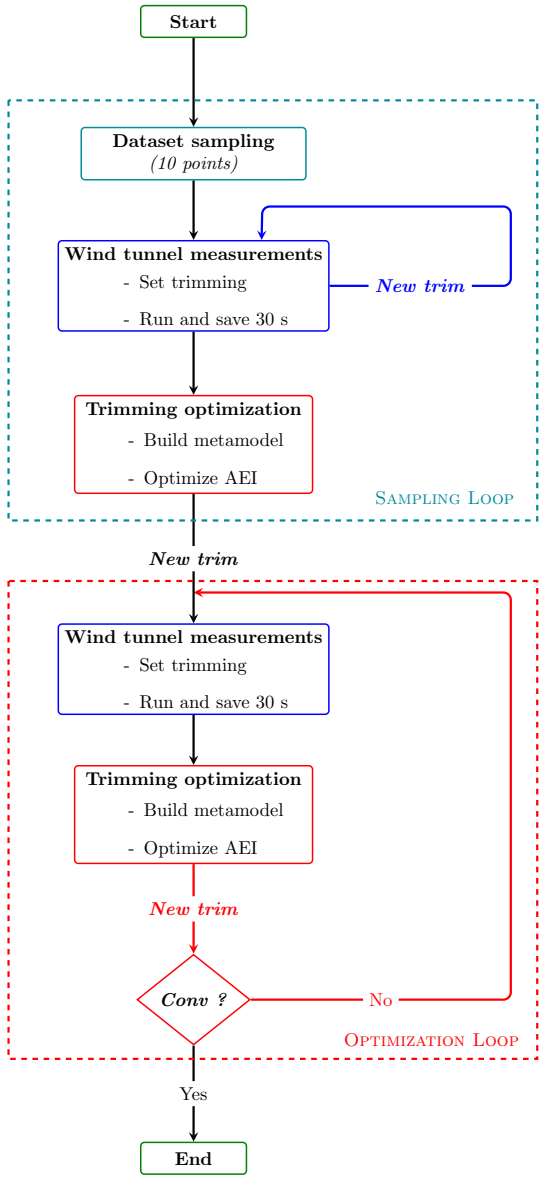

Figure 4: Trimming optimization procedure.

Below, we present the results of the trimming optimization, contrasting two cases. In the first case, the measurement noise is set to $\sigma_\epsilon = 0$ (noise-free situation) so the constructed GP models are interpolating the experimental data. In the second case, the actual measurement noise is determined experimentally and it is subsequently used in the GP model construction.

Experimental optimization with $\sigma_\epsilon = 0$

The results of the experimental trimming optimization in the case where σ_ϵ is set to zero is now presented.

The GP model for the experimental $\mathcal{P}(\mathbf{x})$ function, after 36 iterations of the optimization algorithm, is reported in Figure 5. Specifically, Figure 5(a) depicts the color contours of the GP model mean as a function of the trimming parameters L_{car} and L_{sheet}, while Figure 5(b) shows the standard deviation of the GP model. The black dots are the data points where the performance was experimentally estimated. Regarding the location of the observation points, we notice a large dispersion, highlighting the lack of convergence in the successive tentative optimal trimming candidates $\mathbf{x}_*$. In fact, the algorithm has explored the parameter domain Ω without discovering a particular sub-domain of Ω likely to contain the global optimum. This can also be appreciated from the mean field of the GP model, in Figure 5(a), which, although smooth, presents at least 2 local minima. The presence of multiple local minima is in fact spurious and induced by the interpolating nature of the GP model for $\sigma_\epsilon = 0$: the model is fitting the experimental noise. This can also be appreciated from the standard deviation field reported in Figure 5(b), which is zero at the observation points, denoting an inappropriate level of confidence in the GP model approximation of $\mathcal{P}(\mathbf{x})$ at these locations. Further, departing from the observation points, the variance of the GP model prediction quickly increases (observe, in particular, the standard deviation field in the neighborhood of isolated data points) and becomes large. As a result of the over-confidence in the model at measured points and high variance (low confidence) in unexplored areas, the optimization process is led by the AEI merit function to propose new candidates in relatively less populated areas.

Figure 6 depicts the measured values of $\mathcal{P}(\mathbf{x}_*)$ at the successive tentative optima as selected along the iterations of the algorithm (the first 10 iterations correspond to the initial Latin Hypercube Sampling of Ω, and are not actual iterations of the algorithm). The plot shows that the measured performance is not converging and it sustains large fluctuations having the same magnitude as for the initial random sample: the complete absence of an improvement trend in the successive measurement of $\mathcal{P}(\mathbf{x}_*)$ is characteristic of the failure of the present approach. This unsuccessful test highlights the negative effect of not accounting for the noise in the estimates of $\mathcal{P}(\mathbf{x})$: it prevents the GP model to discover trends in the actual performance function from the noisy observations.

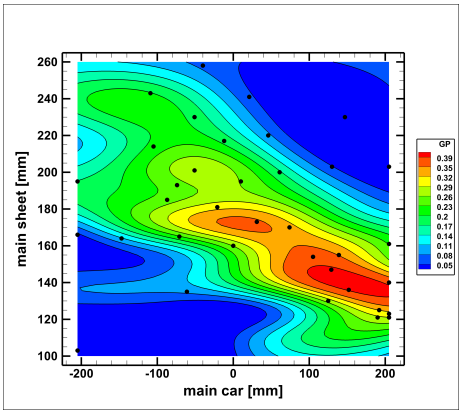

(a) GP model expectation.

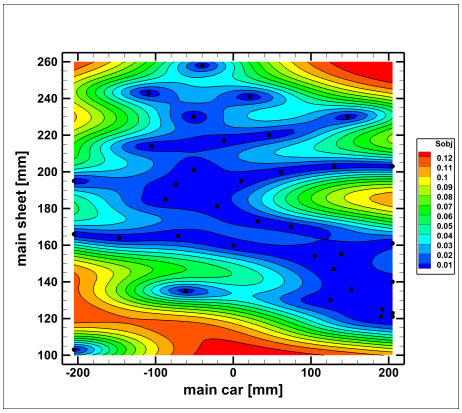

(b) GP model standard deviation.

Figure 5: GP model of the experimental performance $\mathcal{P}(\mathbf{x})$ ($\sigma_\epsilon = 0$).

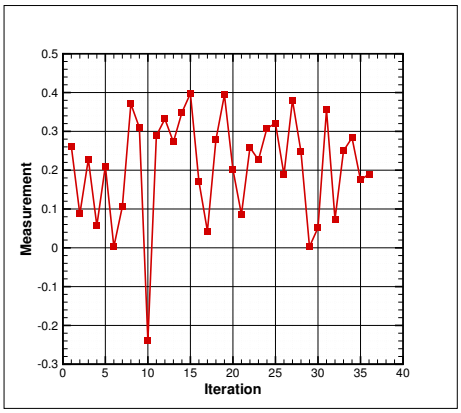

Figure 6: Experimental measurements of the performance for the sequence of proposed optima. Case of $\sigma_\epsilon = 0$.

Experimental optimization with noise

In a second experiment we set ${\sigma_\epsilon}^2 = 0.027^2$. This value is not determined as part of the optimization of the log-marginal likelihood in (12). Instead, σ_ϵ is directly estimated from the experimental apparatus, using repeated measurements at the same trimmings. The same 10 previous initial LHS points are used to determined the first GP model, with the prescribed value of σ_ϵ. Then the optimization proceeds and a different sequence of proposed optima is generated as the GP models differ from the previous experiment.

In particular, the optimization now converges in 33 iterations, the last two proposed optima being in sufficiently close distance (less than 1 mm). The convergence of the sequence of proposed optima can be seens in Figure 7, where the mean and standard deviation of the GP model of $\mathcal{P}(\mathbf{x})$ at convergence are plotted. In contrast to the case with $\sigma_\epsilon = 0$, the clustering of the successive proposed optima is clearly visible. Also, the mean of the GP model in Figure 7(a) remains smooth and now exhibits a single well-defined global maximum. The converged optimal trimming is found for $L_{\text{sheet}} = 133$ mm and $L_{\text{car}} = 138$ mm corresponding to a predicted performance $\mathcal{P}(\mathbf{x}_{\text{opt}}) = 0.397$. The standard deviation of the GP model, depicted in Figure 7(b), is seen to be minimal in the neighborhood of the optimum, though assuming values $\gtrsim \sigma_\epsilon$. Other regions of Ω far from the optimum are not explored by the optimization process, although the variance can be large.

The convergence of the optimization procedure can also be appreciated from the plot of Figure 8 which should be contrasted with the results shown Figure 6. It shows a clear improvement of the measured performance (after the first 10 random points). In fact after iteration 25, the remaining fluctuations in the measured performances can be essentially attributed to the measurement noise. These remaining fluctuations, with amplitude $\sim \sigma_\epsilon$, are much less significant compared to the previous case.

In summary, the global optimal trimming parameters are determined despite the noise in the measurements, thanks to the non-interpolating nature of the GP approximation which smooths out the noise. Further, the optimum is found in few iterations only, owing to the AEI merit function which is able to disregard non-interesting areas of Ω, even if they carry large prediction variance.

NUMERICAL OPTIMIZATION

A numerical model of the wind tunnel and sail model has been created to reproduce the previous experimental optimization problem. The objective is to assess the capabilities of the optimization method, when applied to a coupled FSI software, and compare the resulting optimum with the experimental one.

Numerical model

We briefly present the structural and fluid solvers used for the resolution of the FSI problem. Steady solution of the FSI problem are sought by means of a quasi-steady approach.

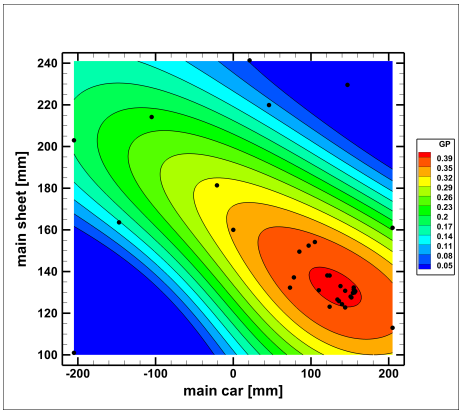

(a) GP model expectation.

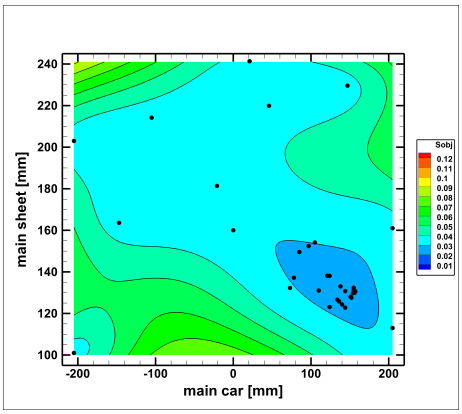

(b) GP model standard deviation.

Figure 7: GP model of the experimental performance $\mathcal{P}(\mathbf{x})$ ($\sigma_\epsilon = 0.027$).

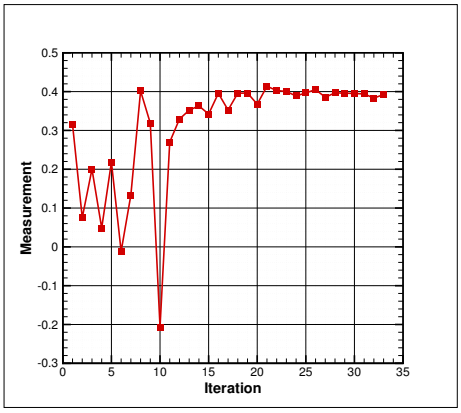

Figure 8: Experimental measurements of the performance for the sequence of proposed optima. Case of $\sigma_\epsilon = 0.027$.

For the structural model of the sail we rely on the ARA software developed by K-EPSILON. The code ARA considers different structural elements (*e.g.* Timoshenko beams, cables and Constant Strain Triangles (CST) membrane elements of various types) for the static or dynamic simulation of sail boat rigs in large displacement regime (Augier, 2012). The structural model for the simulations presented hereafter is illustrated in Figure 9; it uses dimensions and mechanical characteristics (for the mast, boom, and sail fabrics) measured on the experimental model.

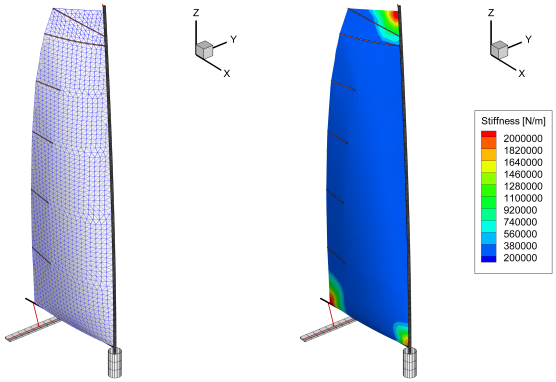

(a) Structural mesh. (b) Mainsail stiffness.

Figure 9: Numerical model of the sail.

The solver ARA is coupled to the ISIS-CFD software (from FINE$^{\text{TM}}$/Marine) which solves the Navier-Stokes equations in the flow domain. ISIS-CFD is based on finite volume methods accommodating both structured and unstructured meshes; it also proposes several turbulence models and boundary condition. For the present computations, we consider a parallelepiped computational domain, with spatial extension $7.5h$, $12h$ and $1.8h$ in the X, Y and Z directions respectively. These dimensions were selected on the basis of previous numerical experiments (Viola et al., 2013). The boundary conditions, applied on the different faces of the computational domain, are schematically illustrated in Figure 10.

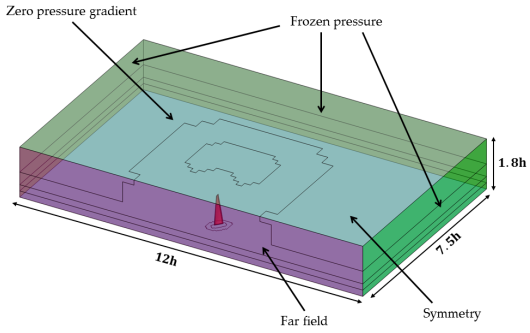

Figure 10: Boundary conditions for the flow solver.

The fluid domain is meshed using HEXPRESS$^{\text{TM}}$, a semi-automated mesh generator. Note that the mast is not meshed in the fluid solver. Regarding the turbulence model,

the SST $k-\omega$ model (Menter et al., 2003) was selected with wall function boundary conditions (Kalitzin et al., 2005). This choice requires a sufficiently fine mesh over the sail surface and at the bottom of domain (sea level) to correctly capture the vertical profile of the inflow velocity. The later is estimated from the experimental profiles previously reported in Figure 3, and its shown in Figure 11.

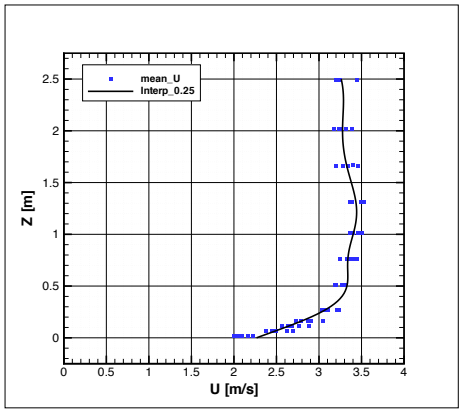

Figure 11: Inflow velocity profile.

For the FSI simulations, the mesh of the fluid domain has to be deformed to adapt to the changes in the boundary geometries. A mesh deformation propagation method, proceeding from the deformable boundaries toward the inside of the fluid domain, was developed at K-EPSILON for this purpose (Durand et al., 2014). The FSI problem is then solved coupling ARA and ISIS-CFD solvers with a quasi-monolithic algorithm (Durand, 2012), which is an implicit coupling procedure adapted to a partitioned solver. Briefly, the resolution of the structural problem is nested inside the iterations on the nonlinear steady flow solution. This approach preserves the convergence and stability properties of the monolithic approach. More details on the solvers and the coupling algorithms can be found in (Durand, 2012, Roux et al., 2002, 2008).

A convergence analysis was conducted in order to select spatial discretization capturing correctly the physics of the FSI problem, while maintaining a reasonable computational cost permitting the optimization of the trimming parameters. In particular, different fluid meshes with up to 4.3 million finite volumes were considered. Eventually, a discretization of the fluid domain with roughly 1.5 million finite volumes and a sail discretization with 2 700 membrane elements was selected for the computations presented below. The computations were carried out on a 64 CPUs cluster; an averaged computational time of 5 h was reported for solving individual FSI problems. From the FSI solution, the associated aerodynamic forces acting on the sail are computed in the same reference frame as in the experimental setup, and the performance in (16) is finally returned to the optimizer. Except for the determination of $\mathcal{P}(\mathbf{x})$, the flowchart of the numerical optimization procedure is identical to that of the experiment in Figure 4. However, for the numerical optimization, σ_ϵ is directly inferred from the data when minimizing (12). In addition, the parameter domain for the numerical case has been increased to encompass higher values of L_{car}.

Numerical optimization results

Figure 12 depicts the mean value of the GP model based on the numerical evaluation of the performance. The GP model is reported at the end of the optimization procedure, which has converged in 34 iterations. The black dots correspond again to the sequence of optimization points $\mathbf{x}_*$ in the data set. We first remark the smoothness of the mean GP model which exhibits a single global optimum, as for the experimental case (considering measurement noise). In fact, the inferred $\sigma_\epsilon = 0.022$ has a value close to the experimental one. Further, in the range $L_{\text{car}} \in [-210, 150]$ mm, the terminal GP model is seen to be in good agreement with its experimental counterpart shown in Figure 7(a). However, the valley containing the numerical minimum is larger and flatter, compared to the experimental one, and the numerical optimum appears at a value of a L_{car} value larger than for the experimental case. The variance of the GP model prediction, σ_y^2, exhibits a structure similar in shape and magnitude to the experimental case in Figure 7(b) (not shown).

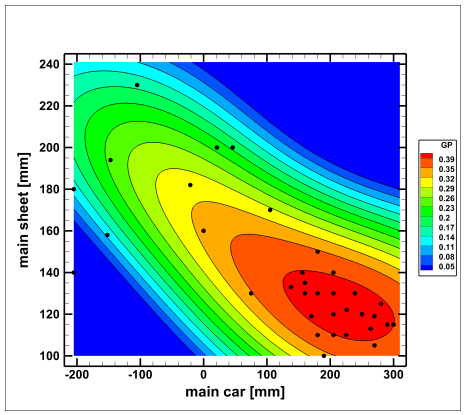

Figure 12: Mean GP model of the numerical performance (inferred $\sigma_\epsilon = 0.022$).

To understand the differences between the numerical and experimental optima, we first compare in Figure 13 and Figure 14 the flyings shapes for two trimming parameters $(L_{\text{sheet}}, L_{\text{car}})$ equal to $(160, 0)$ and $(133, 138)$ respectively (lengths in mm). For the case with centered car, $L_{\text{car}} = 0$, a small wrinkle is visible in the experimental flying shape (see Figure 13(a)). At the optimal experimental trimming point $(L_{\text{sheet}}, L_{\text{car}}) = (133, 138)$, shown in Figure 14(a), the wrinkle in the experimental flying shape is even more pronounced. This is in contrast with the corresponding numerical flying shapes, shown in Figures 13(b) and 14(b)

respectively, which present no such wrinkle. Modeling errors and experimental uncertainties are deemed responsible for this difference. In particular, the absence of wrinkle in the numerical solution could be mostly due to an incorrect prescription of the tensions in the two full battens of the sail. Another important source of discrepancy between the flying shapes are the boundary conditions of the numerical wind tunnel and effects of confinement, which were shown to have a significant impact on the computed aerodynamic forces (Viola et al., 2013).

(a) Experimental.

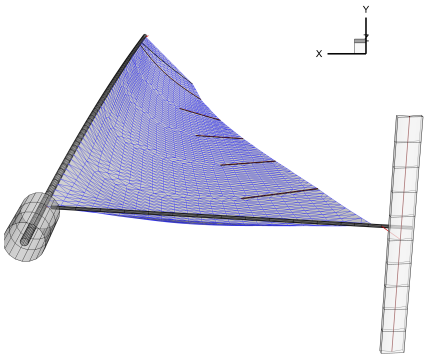

(b) Numerical.

Figure 13: Comparison of experimental and numerical flying shapes for trimming parameters $L_{\text{sheet}} = 160$ and $L_{\text{car}} = 0$.

To complete the comparison of the experimental and numerical optimizations, we report in Table 1 the computed location of the two optima and the corresponding best prediction of the performance. We observe that the experimental and numerical optimal trimming are significantly different for the optimal L_{car}. This difference can be explained by the numerical performance function that is particularly flat along the L_{car} direction around the optimal point: variation of L_{car} around the optimum weakly affects the predicted performance. This is consistent with the predicted performances reported in the last column of Table 1, that are in close agreement despite the differing trimming parameters.

A more detailed investigation of the measured and computed fluid forces at the optima, reported in Table 2, reveals

(a) Experimental.

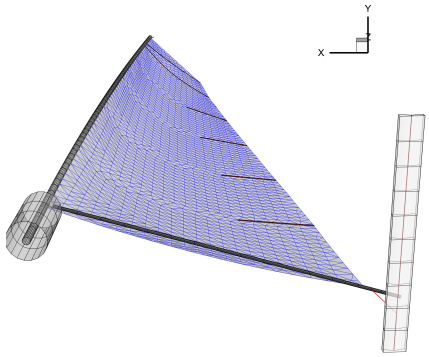

(b) Numerical.

Figure 14: Comparison of experimental and numerical flying shapes for trimming parameters $L_{\text{sheet}} = 133$ and $L_{\text{car}} = 138$.

	L_{sheet} [mm]	L_{car} [mm]	Pred. $\mathcal{P}(\mathbf{x}_{\text{opt}})$
Exp.	133	138	0.397 ± 0.027
Num.	122	226	0.413 ± 0.024

Table 1: Comparison of the experimental and numerical optima.

that the thrust coefficients are in fact equal up to the third significant digit. On the contrary, the side force coefficients and the sheet tensions exhibit larger discrepancies. The differences in the sheet tensions can be directly attributed to the different optimal L_{car}. The higher disagreement in the side force coefficients is not surprising. In our experience the side forces are very sensitive to model errors.

	C_X	C_Y	T [N]
Exp.	0.497 ± 0.012	-1.026 ± 0.015	14.9 ± 0.02
Num.	0.497	-0.803	20.1

Table 2: Comparison of aerodynamic coefficients and sheet tension at the experimental and numerical optima.

CONCLUSION AND DISCUSSION

We have proposed to use a Gaussian Process model to enable the optimization of the trimming parameters of a complex nonlinear sail systems. The approach has been first tested on the trimming of an experimental model sail in the Yacht Research Unit wind tunnel. The experiments have validated the approach and have shown its robustness against noisy estimates of the performance. For this two parameters problem, the experimental optimal point was found within few iterations of the algorithm. These tests have validated the proposed optimization method and have demonstrated its robustness against experimental variabilities, as well as the key role of the noise parameter.

A detailed numerical model of the wind tunnel experiment has been established, considering the full Fluid-Structure Interactions problems. The numerical model is based on a turbulent flow solver coupled with a nonlinear elastic solver, using a mesh deformation method. The numerical and experimental optima were found to be consistent, given all the modeling and measurement errors. In particular, it is found that the predicted performances and fluid forces are in much better agreement than the optimal parameters. In fact, it can be reasonably claimed that the differences are consistent with the current predictive capabilities of state of the art FSI solvers, and that the GP approximation does not introduce noticeable errors in the optimization procedure. Specifically, we have shown that, from a limited set of computations, the GP model is able to reconstruct accurately the numerical estimate of the performance. Therefore, for the present tests, only a better numerical modeling of the experimental set-up would help reducing the observed discrepancies. Possible avenues in this direction are: improved battens modeling, better boundary conditions for the flow, accounting for mast/flow interaction, more advanced turbulence model, ... At a more fundamental level, the question of the treatment of experimental uncertainties remains critical. On this aspect, we are considering sensitivity and uncertainty quantification (Le Maître and Knio, 2010) studies to account for the experimental uncertainties. In particular, the results presented in this work point to the need for an appropriate characterization of the optima and the assessment of their robustness to uncertainties and modeling errors. Future works will develop these aspects along with the deployment of the GP-based optimization method to problems involving large numbers of trimming parameters (full yacht rig).

ACKNOWLEDGEMENTS

The authors wish to acknowledge the people of the Yacht Research Unit for their welcoming and help during the experiments in the wind tunnel. Contributions of INCIDENCES SAILS and NUMECA companies are also acknowledged. This work was partially funded by the European Union's Seventh Program for research, technological development and demonstration under grant agreement No PIRSES-GA-2012-318924, and from the Royal Society of New Zealand for the UK-France-NZ collaboration project SAILING FLUIDS (see `www.sailingfluids.org`). This work was also supported by the "Laboratoire d'Excellence" LabexMER (ANR-10-LABX-19) and co-funded by a grant from the French government under the program "Investissements d'Avenir".

REFERENCES

B. Augier. *Etudes expérimentales de l'interaction fluide-structure sur surface souple: application aux voiles de bateaux.* PhD Thesis, Universite de Bretagne Occidentale, 2012.

B. Augier, P. Bot, F. Hauville, and M. Durand. Experimental validation of unsteady models for fluid structure interaction: Application to yacht sails and rigs. *Journal of Wind Engineering and Industrial Aerodynamics*, 101:53–66, 2012.

T. Bäck and H. P. Schwefel. An overview of evolutionary algorithms for parameter optimization. *Evolutionary computation*, 1(1):1–23, 1993.

V. G. Chapin, R. Neyhousser, G. Dulliand, and P. Chassaing. Design optimization of interacting sails through viscous cfd. In *INNOVSail, Innovation in high performance sailing Yacht*, Lorient, 2008.

M. Durand. *Interaction fluide-structure souple et légère, application aux voiliers*. PhD Thesis, Ecole Centrale de Nantes, 2012.

M. Durand, A. Leroyer, C. Lothodé, F. Hauville, M. Visonneau, R. Floch, and L. Guillaume. Fsi investigation on stability of downwind sails with an automatic dynamic trimming. *Ocean Engineering*, 90:129–139, 2014.

R. Duvigneau and P. Chandrashekar. Kriging-based optimization applied to flow control. *International Journal for Numerical Methods in Fluids*, 69(11):1701–1714, 2012.

R. G. J. Flay. A twisted flow wind tunnel for testing yacht sails. *Journal of Wind Engineering and Industrial Aerodynamics*, 63(13):171 – 182, 1996. ISSN 0167-6105. Special issue on sail aerodynamics.

M. N. Gibbs. *Bayesian Gaussian Processes for Classification and Regression*. PhD thesis, University of Cambridge, 1997.

H. Hansen, P. S. Jackson, and K. Hochkirch. Real-time velocity prediction program for wind tunnel testing of sailing yachts. *Proc. The Modern Yacht, Southampton, UK*, 2003.

N. Hansen. The cma evolution strategy: a comparing review. In *Towards a new evolutionary computation*, pages 75–102. Springer, 2006.

N. Hansen and A. Ostermeier. Completely derandomized self-adaptation in evolution strategies. *Evolutionary Computation*, 9(2):159–195, 2001.

D. Huang, T. T. Allen, W. I. Notz, and N. Zeng. Global optimization of stochastic black-box systems via sequential kriging meta-models. *Journal of global optimization*, 34 (3):441–466, 2006.

L. Huetz and P. E. Guillerm. Database building and statistical methods to predict sailing yacht hydrodynamics. *Ocean Engineering*, 90:21–33, 2014.

D. R. Jones. A taxonomy of global optimization methods based on response surfaces. *Journal of global optimization*, 21(4):345–383, 2001.

D. R. Jones, M. Schonlau, and W. J. Welch. Efficient global optimization of expensive black-box functions. *Journal of Global optimization*, 13(4):455–492, 1998.

G. Kalitzin, G. Medic, G. Iaccarino, and P. Durbin. Near-wall behavior of rans turbulence models and implications for wall functions. *Journal of Computational Physics*, 204(1):265–291, 2005.

R. Korpus. Performance prediction without empiricism: A rans-based vpp and design optimization capability. In *The 18th Chesapeake Sailing Yacht Symposium, SNAME*, 2007.

O. P. Le Maître and O. M. Knio. Spectral methods for uncertainty quantification. *Scientific Computation. Springer, New York*, 2010.

D. J. Le Pelley and O. Modral. V-spars: A combined sail and rig shape recognition system using imaging techniques. In *Proc. 3rd High Performance Yacht Design Conference Auckland, New Zealand, Dec*, pages 2–4, 2008.

M. D. McKay, R. J. Beckman, and W. J. Conover. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 42(1):55–61, 2000.

W. Menotti, M. Durand, D. Gross, Y. Roux, D. Glehen, and L. Dorez. An unsteady fsi investigation into the cause of the dismating of the volvo 70 groupama 4. In *INNOVSail, Innovation in high performance sailing Yacht*, page 197, Lorient, 2013.

F. R. Menter, M. Kuntz, and R. Langtry. Ten years of industrial experience with the sst turbulence model. *Turbulence, heat and mass transfer*, 4:625–632, 2003.

J. Nelder and R. Mead. A simplex method for function minimization. *Computer Journal*, 7(4):208–313, 1965.

P. V. Oossanen. Predicting the speed of sailing yachts. *SNAME*, 101:337–397, June 1993.

R. Ranzenbach, D. Armitage, and A. Carrau. Mainsail planform optimization for irc 52 using fluid structure interaction. In *The 21st Chesapeake Sailing Yacht Symposium, SNAME*, 2013.

C. E. Rasmussen and C. K. I. Williams. *Gaussian processes for machine learning*. MIT Press, 2006.

N. Rousselon. Optimization for sail design. In *modeFRONTIER Conference*, 2008.

Y. Roux, S. Huberson, F. Hauville, J. P. Boin, M. Guilbaud, and B. Malick. Yacht performance prediction : Towards a numerical vpp. In *1st High Performance Yacht Design Conference Auckland, 4-6 December*, Auckland, 2002.

Y. Roux, M. Durand, A. Leroyer, P. Queutey, M. Visonneau, J. Raymond, J. M. Finot, F. Hauville, and A. Purwanto. Strongly coupled vpp and cfd ranse code for sailing yacht performance prediction. In *3rd High Performance Yacht Design Conference Auckland, 2-4 December*, pages 215–225, Auckland, 2008.

V. K. Saul'ev and I. I. Samoilova. Approximation methods for the unconstrained optimization of functions of several variables. *Journal of Soviet Mathematics*, 4(6):681–705, 1975.

M. L. Stein. *Interpolation of spatial data: some theory for kriging*. Springer Science & Business Media, 2012.

D. Trimarchi. *Analysis of downwind sail structures using non-linear shell finite elements*. PhD Thesis, University of Southampton, 2012.

I. M. Viola, P. Bot, and M. Riotte. Upwind sail aerodynamics: A rans numerical investigation validated with wind tunnel pressure measurements. *International Journal of Heat and Fluid Flow*, 39:90–101, 2013.

R. Von Mises. *Mathematical Theory of Probability and Statistics*. Academic press, 1964.