Sensitivity-Based Tube NMPC for Close-Proximity Offshore Wind Turbine Inspection with a Tilted Multirotor

Giuseppe Silano[1] and Martin Saska[2]

***Abstract*— Close-proximity offshore wind turbine inspection requires strict clearance control around large cylindrical structures under wind and model mismatch. Nominal Nonlinear Model Predictive Control (NMPC) may violate safety constraints when mass, inertia, thrust effectiveness, drag, or wind conditions differ from nominal assumptions. We propose a sensitivity-based tube NMPC for a tilted multirotor that robustifies the tower-clearance constraint via online tightening. First-order parametric state sensitivities provide a structured-uncertainty margin, while bounded gusts are handled by a stage-dependent additive margin. The formulation preserves the computational structure of nominal NMPC and is designed for boundary-critical helical inspection trajectories.**

***Index Terms*— Offshore wind turbine inspection, close-proximity flight, robust NMPC, sensitivity-based tube tightening, tilted multirotor.**

I. Introduction

Aerial inspection can reduce offshore maintenance cost and human exposure, but reliable autonomy near wind-turbine towers remains difficult because of persistent wind, aerodynamic effects, and plant–model mismatch [1], [2]. Since inspection-quality sensing requires tracking close-proximity trajectories around the tower while enforcing a strict minimum clearance, safety preservation rather than trajectory tracking alone becomes the primary control objective.

Nonlinear Model Predictive Control (NMPC) is well suited to this setting because it handles nonlinear dynamics and constraints within a unified receding-horizon problem [3], [4]. However, nominal NMPC may be overly optimistic under uncertainty, yielding trajectories that are feasible in prediction but unsafe on the real system. This paper therefore proposes a practical robustification for a tilted multirotor based on sensitivity-driven tightening of the cylindrical clearance constraint under bounded structured uncertainty and bounded wind gusts.

More specifically, the contribution is threefold: i) a structured uncertainty model capturing mismatch in mass, inertia, propulsion effectiveness, drag, and persistent wind bias; ii) an online first-order parametric tightening term obtained from propagated state sensitivities; and iii) an additive stage-dependent gust margin in explicit closed form.

This work was partially funded by the research fund for the Italian Electrical System through decree n. 388 of November 6th, 2024, by the GAČR project no. 26-22419S, by the CTU grant no. SGS23/177/OHK3/3T/13, and by the European project "Robotics and Advanced Industrial Production" (reg. no. CZ.02.01.01/00/22008/0004590).

[1]Giuseppe Silano is with the Department of Power Generation Technologies and Materials, Ricerca sul Sistema Energetico S.p.A., Milan, Italy, and also with the Department of Cybernetics, Czech Technical University, Prague, Czechia (e-mail: giuseppe.silano@fel.cvut.cz).

[2]Martin Saska is with the Department of Cybernetics, Czech Technical University, Prague, Czechia (e-mail: martin.saska@cvut.cz).

II. Problem Setup and Prediction Model

The mission considered in this paper is close-proximity helical tower inspection with guaranteed minimum clearance. The tower is modeled as a vertical cylinder defining the inspection geometry and the safety-critical clearance constraint, while the aerial platform is described by a control-oriented tilted-multirotor prediction model with input-rate actuation.

A. Wind turbine inspection geometry

The tower is modeled as a vertical cylinder of radius R, aligned with the inertial z-axis, which captures the dominant obstacle during close-proximity tower inspection. Let $\mathbf{p} = [x\ y\ z]^\top$ denote the vehicle position, and define

$$\rho(\mathbf{p}) \triangleq \sqrt{x^2+y^2}, \qquad d_T(\mathbf{p}) \triangleq \rho(\mathbf{p}) - R. \tag{1}$$

Collision avoidance requires

$$d_T(\mathbf{p}(t)) \geq d_{\min}, \quad \forall t \geq 0, \tag{2}$$

where $d_{\min} > 0$ is the minimum admissible clearance. The inspection reference is a helical trajectory $\mathbf{p}_{\mathrm{ref}}$ on the preferred offset surface

$$d_T(\mathbf{p}_{\mathrm{ref}}(t)) = d_{\mathrm{ref}} > d_{\min}, \tag{3}$$

so that the vehicle orbits the tower while maintaining the desired sensing distance.

B. Control-oriented GTMR model

The aerial platform is described by a control-oriented Generically-Tilted Multirotor (GTMR) model [5], [6] with input-rate actuation. Let N_p denote the number of rotor–propeller units, and define the system state as

$$\mathbf{x} = \begin{bmatrix}\mathbf{p}^\top & \boldsymbol{\eta}^\top & \mathbf{v}^\top & \boldsymbol{\omega}^\top & \boldsymbol{\xi}^\top\end{bmatrix}^\top, \tag{4}$$

where $\boldsymbol{\eta}$ is a local attitude parametrization, $\mathbf{v} \in \mathbb{R}^3$ is the linear velocity, $\boldsymbol{\omega} \in \mathbb{R}^3$ is the body angular velocity, $\boldsymbol{\xi} \in \mathbb{R}^{N_p}$ collects the actuator states, and $\mathbf{u} = \dot{\boldsymbol{\xi}} \in \mathbb{R}^{N_p}$ is the actuator-rate input. The prediction model is written compactly as

$$\dot{\mathbf{x}} = f(\mathbf{x}, \mathbf{u}, \boldsymbol{\zeta}) + \mathbf{G}_w \mathbf{w}_g, \tag{5}$$

where $\boldsymbol{\zeta}$ collects the structured uncertainty, including the slowly varying wind bias, $\mathbf{w}_g$ is the gust disturbance, and $\mathbf{G}_w$ is a constant selection matrix mapping gusts to the translational acceleration channels of the state dynamics. The state and input are constrained to the admissible sets $\mathbf{x} \in \mathcal{X}$ and $\mathbf{u} \in \mathcal{U}$, which encode actuator magnitude and rate limits.

III. Uncertainty Description

The uncertainty considered in this paper is split into two parts: a slowly varying structured mismatch and a fast bounded gust component. This separation matches the robustification strategy developed later, where structured uncertainty is handled through sensitivity propagation and gusts through an additive margin.

A. Structured parametric uncertainty

The structured uncertainty vector is

$$\boldsymbol{\zeta} = \begin{bmatrix} \delta_m & \delta_{J_x} & \delta_{J_y} & \delta_{J_z} & \delta_T & \delta_D & w_{b,x} & w_{b,y} & w_{b,z} \end{bmatrix}^\top, \quad (6)$$

where δ_m, δ_{J_x}, δ_{J_y}, and δ_{J_z} denote relative mismatch in mass and principal inertia components, δ_T and δ_D represent relative mismatch in thrust effectiveness and drag, and $w_{b,x}, w_{b,y}, w_{b,z}$ are the components of the persistent wind bias. These quantities enter the prediction model as bounded perturbations of the corresponding nominal parameters, e.g., through multiplicative relations of the form $m = (1+\delta_m)m_0$. The admissible set is defined as

$$\mathcal{Z} = \{\boldsymbol{\zeta} : \underline{\boldsymbol{\zeta}} \le \boldsymbol{\zeta} \le \bar{\boldsymbol{\zeta}}\}, \quad (7)$$

where the inequalities are understood componentwise, and $\underline{\boldsymbol{\zeta}}$ and $\bar{\boldsymbol{\zeta}}$ denote the corresponding lower and upper bounds on the uncertainty vector.

B. Wind decomposition

The total wind action is decomposed as $\mathbf{w}(t) = \mathbf{w}_b + \mathbf{w}_g(t)$, where $\mathbf{w}_b$ is a slowly varying bias included in $\boldsymbol{\zeta}$, and $\mathbf{w}_g(t)$ is a fast bounded gust satisfying $\|\mathbf{w}_g(t)\| \le \bar{w}_g$ for all $t \ge 0$. This separation matches the robustification strategy adopted later, where structured uncertainty is handled through sensitivity propagation and gusts through an additive margin.

IV. Nominal NMPC Formulation

With sampling time T_s, horizon N, and nominal uncertainty realization $\boldsymbol{\zeta}_0 = \mathbf{0}$, the discrete-time model is

$$\mathbf{x}_{i+1} = F(\mathbf{x}_i, \mathbf{u}_i, \boldsymbol{\zeta}_0). \quad (8)$$

The nominal NMPC problem, used as the baseline for the subsequent robust tightening, is

$$\min_{\{\mathbf{x}_i, \mathbf{u}_i\}} \sum_{i=0}^{N-1} \ell(\mathbf{x}_i, \mathbf{u}_i, \mathbf{r}_i) + \ell_f(\mathbf{x}_N, \mathbf{r}_N) \quad (9a)$$

$$\text{s.t. } \mathbf{x}_{i+1} = F(\mathbf{x}_i, \mathbf{u}_i, \boldsymbol{\zeta}_0),\ i = 0, \dots, N-1, \quad (9b)$$

$$\mathbf{x}_i \in \mathcal{X},\ \mathbf{u}_i \in \mathcal{U},\ i = 0, \dots, N-1, \quad (9c)$$

$$d_T(\mathbf{p}_i) \ge d_{\min},\ i = 0, \dots, N, \quad (9d)$$

$$\mathbf{x}_0 = \mathbf{x}(t_k). \quad (9e)$$

The stage and terminal costs are chosen as

$$\ell = \|h(\mathbf{x}_i) - \mathbf{r}_i\|^2_{\mathbf{Q}} + \|\mathbf{u}_i\|^2_{\mathbf{Q}_u}, \quad (10)$$

$$\ell_f = \|h_f(\mathbf{x}_N) - \mathbf{r}_N\|^2_{\mathbf{Q}_f}, \quad (11)$$

where $\mathbf{r}_i$ and $\mathbf{r}_N$ are obtained by sampling the helical inspection trajectory, and $h(\cdot)$ and $h_f(\cdot)$ select the tracked components of the state. The matrices $\mathbf{Q}$, $\mathbf{Q}_u$, and $\mathbf{Q}_f$ weight stage-output tracking, input effort, and terminal tracking, respectively.

V. Sensitivity-Based Tube Tightening

The proposed robustification augments the nominal clearance constraint with two additional margins: a parametric margin accounting for structured model mismatch and a gust margin accounting for fast bounded wind disturbances. Their combined effect is to enforce a conservative clearance condition along the prediction horizon while preserving a tractable optimization structure.

A. Parametric sensitivity dynamics

To quantify how structured uncertainty affects the predicted motion, the proposed method propagates first-order state sensitivities along the horizon. To this end, define the state sensitivity matrix as

$$\boldsymbol{\Pi} \triangleq \frac{\partial \mathbf{x}}{\partial \boldsymbol{\zeta}} \in \mathbb{R}^{(12+N_p)\times 9}, \quad (12)$$

which measures the sensitivity of the extended state to perturbations of the structured uncertainty vector. Its continuous-time evolution is

$$\dot{\boldsymbol{\Pi}} = \frac{\partial f}{\partial \mathbf{x}}\boldsymbol{\Pi} + \frac{\partial f}{\partial \boldsymbol{\zeta}}, \qquad \boldsymbol{\Pi}(t_k) = \mathbf{0}, \quad (13)$$

where the initialization $\boldsymbol{\Pi}(t_k) = \mathbf{0}$ reflects the fact that the sensitivity is reset at the beginning of each prediction horizon. The corresponding discrete-time propagation is

$$\boldsymbol{\Pi}_{i+1} = F_\Pi(\mathbf{x}_i, \mathbf{u}_i, \boldsymbol{\Pi}_i, \boldsymbol{\zeta}_0), \quad (14)$$

obtained consistently with the nominal discrete-time model.

B. Clearance constraint in standard form

To derive the tightening term, the clearance requirement is rewritten as a scalar inequality. Specifically, define

$$y(\mathbf{x}) \triangleq d_{\min} - d_T(\mathbf{p}) \le 0, \quad (15)$$

so that satisfying $y(\mathbf{x}) \le 0$ is equivalent to enforcing the minimum-clearance condition. Since the constraint depends only on the position components of the state, its gradient with respect to (w.r.t.) $\mathbf{p}$ is

$$\frac{\partial y}{\partial \mathbf{p}} = -\begin{bmatrix} x/\rho & y/\rho & 0 \end{bmatrix}, \quad (16)$$

and the corresponding Jacobian w.r.t. the system state is

$$J_{yx}(\mathbf{x}) = \frac{\partial y}{\partial \mathbf{x}} = \begin{bmatrix} \frac{\partial y}{\partial \mathbf{p}} & \mathbf{0}_{1\times(9+N_p)} \end{bmatrix}, \quad (17)$$

where the zero block reflects the absence of direct dependence on the remaining state components. Therefore, the first-order sensitivity of the scalar clearance constraint w.r.t. the structured uncertainty is

$$\boldsymbol{\Pi}_y \triangleq \frac{\partial y}{\partial \boldsymbol{\zeta}} = J_{yx}(\mathbf{x})\boldsymbol{\Pi}. \quad (18)$$

This quantity measures how uncertainty perturbs the clearance constraint along the prediction horizon and is the key ingredient for the parametric tightening term.

C. Parametric tightening margin

To account for bounded structured uncertainty, consider the deviation of the uncertainty vector from its nominal value,

$$\mathbf{w}_\zeta \triangleq \boldsymbol{\zeta} - \boldsymbol{\zeta}_0, \tag{19}$$

and the corresponding componentwise-bounded set

$$\mathcal{W}_\zeta \triangleq \{\mathbf{w}_\zeta : \ |\mathbf{w}_{\zeta,j}| \le \bar{w}_{\zeta,j}\}. \tag{20}$$

Using the first-order approximation of the constraint variation, the parametric tightening margin at prediction stage i is defined as

$$\alpha_{p,i} \triangleq \|\boldsymbol{\Pi}_{y,i}\|_\infty \, \|\bar{\mathbf{w}}_\zeta\|_1 + \varepsilon_s, \tag{21}$$

where $\boldsymbol{\Pi}_{y,i}$ denotes the sensitivity of the scalar clearance constraint at stage i, $\bar{\mathbf{w}}_\zeta$ collects the componentwise bounds of the structured uncertainty, and $\varepsilon_s > 0$ is a small slack term introduced for numerical robustness.

D. Wind-gust margin

The gust component is not included in the sensitivity propagation and is instead handled through an additive stage-dependent margin. Assuming that the fast wind disturbance satisfies $\|\mathbf{w}_g(t)\| \le \bar{w}_g$ for all $t \ge 0$. Since $m \ge m_{\min} \triangleq (1-\bar{\delta}_m)m_0$, the induced translational acceleration perturbation satisfies

$$\|\Delta\dot{\mathbf{v}}\| \le \bar{a}_g, \qquad \bar{a}_g \triangleq \frac{\bar{w}_g}{m_{\min}}. \tag{22}$$

In addition, assuming identical initial states for the nominal and perturbed trajectories at the beginning of the horizon, the corresponding position deviation over $\tau = iT_s$ is bounded by

$$\|\Delta\mathbf{p}_g(\tau)\| \le \tfrac{1}{2}\bar{a}_g\tau^2, \tag{23}$$

which yields the stage-dependent gust margin

$$\alpha_{g,i} \triangleq \tfrac{1}{2}\bar{a}_g(iT_s)^2. \tag{24}$$

This term bounds the clearance loss induced by unresolved gusts and grows naturally with the prediction step.

E. Tightened constraint

The final robustified clearance condition combines the parametric and gust margins at each prediction stage. In standard form, the tightened constraint is

$$y(\mathbf{x}_i) + \alpha_{p,i} + \alpha_{g,i} \le 0, \tag{25}$$

or, equivalently, in clearance form,

$$d_T(\mathbf{p}_i) \ge d_{\min} + \alpha_{p,i} + \alpha_{g,i}. \tag{26}$$

Thus, the nominal predicted trajectory must maintain an additional safety margin that compensates for both structured parametric uncertainty and unresolved wind-gust effects.

VI. Robustified NMPC Problem

The proposed controller augments the nominal NMPC formulation in Section IV with sensitivity propagation and replaces the nominal clearance constraint with its tightened counterpart. The resulting finite-horizon problem is

$$\min_{\{\mathbf{x}_i,\mathbf{u}_i,\boldsymbol{\Pi}_i\}} \ \sum_{i=0}^{N-1} \ell(\mathbf{x}_i,\mathbf{u}_i,\mathbf{r}_i) + \ell_f(\mathbf{x}_N,\mathbf{r}_N) \tag{27a}$$

$$\text{s.t. } \mathbf{x}_{i+1} = F(\mathbf{x}_i,\mathbf{u}_i,\boldsymbol{\zeta}_0), \ i=0,\dots,N-1, \tag{27b}$$

$$\boldsymbol{\Pi}_{i+1} = F_\Pi(\mathbf{x}_i,\mathbf{u}_i,\boldsymbol{\Pi}_i,\boldsymbol{\zeta}_0), \ i=0,\dots,N-1, \tag{27c}$$

$$\mathbf{x}_i \in \mathcal{X}, \ \mathbf{u}_i \in \mathcal{U}, \ i=0,\dots,N-1, \tag{27d}$$

$$y(\mathbf{x}_i) + \alpha_{p,i} + \alpha_{g,i} \le 0, \ i=0,\dots,N, \tag{27e}$$

$$\mathbf{x}_0 = \mathbf{x}(t_k), \quad \boldsymbol{\Pi}_0 = \mathbf{0}. \tag{27f}$$

Here, $\boldsymbol{\Pi}_i$ propagates the first-order sensitivity to structured uncertainty, while $\alpha_{p,i}$ and $\alpha_{g,i}$ denote the corresponding parametric and gust tightening margins at prediction stage i. Compared with nominal NMPC, the additional online computations are limited to sensitivity propagation and margin evaluation, while the overall receding-horizon structure is preserved.

VII. Simulation Results

This section presents the numerical assessment of the proposed robustified NMPC in a safety-critical tower-inspection scenario. The mission consists of tracking a helical reference trajectory around a cylindrical tower while enforcing the minimum-clearance requirement $d_T(\mathbf{p}) \ge d_{\min}$. The numerical values used in the simulations, including platform parameters, inspection geometry, symmetric uncertainty bounds, and controller weights, are summarized in Table I. The desired inspection distance d_{ref} is selected close to $d_{\min}$, so that moderate uncertainty can erode the clearance margin if it is not explicitly accounted for. Both the nominal and the robustified NMPC are implemented with the same prediction horizon, discretization, weighting matrices, actuator constraints, and solver settings to ensure a fair comparison. The continuous-time dynamics are discretized by a fixed-step fourth-order Runge–Kutta scheme within a multiple-shooting formulation, and the resulting nonlinear programs are solved in MATLAB using MATMPC [7] in real-time iteration mode with qpOASES [8].

The uncertainty realization includes structured mismatch in mass, inertia, thrust effectiveness, drag, and persistent wind bias, together with a bounded gust component. To keep the numerical assessment compact, two types of evidence are reported: i) a representative closed-loop trajectory comparison between nominal and robustified NMPC, and ii) a Monte-Carlo summary over 500 uncertainty realizations.

Figure 1 shows a representative inspection run. In the nominal case, the controller tracks the helical reference more aggressively, but under uncertainty the realized trajectory moves closer to the tower. By contrast, the proposed robustified formulation preserves a larger distance from the tower by anticipating the effect of model mismatch and gust disturbances through the tightening margins.

TABLE I: Simulation parameters.

Parameter	Symbol	Value	Unit
GTMR model			
Mass	m	2.57	kg
Principal inertia	$\mathbf{J}$	$\mathrm{diag}(0.11,\ 0.11,\ 0.19)$	$\mathrm{kg\,m^2}$
Arm length	d	0.39	m
Rotor force coefficient	c_f	11.8×10^{-4}	$\mathrm{N/Hz^2}$
Rotor torque coefficient	c_t	2.5×10^{-5}	$\mathrm{N\,m/Hz^2}$
Inspection geometry			
Tower radius	R	1.00	m
Minimum clearance	$d_{\min}$	0.20	m
Desired offset	d_{ref}	0.35	m
Altitude range	z	$[0.40,\ 6.00]$	m
Uncertainty bounds			
Mass mismatch	$\bar{\delta}_m$	0.10	–
Inertia mismatch	$\bar{\boldsymbol{\delta}}_J$	$[0.15,\ 0.15,\ 0.20]$	–
Propulsion mismatch	$\bar{\delta}_T$	0.10	–
Drag mismatch	$\bar{\delta}_D$	0.20	–
Wind-bias bounds	$\bar{\mathbf{w}}_b$	$[0.8,\ 0.8,\ 0.3]^\top$	N
Gust-force bound	$\bar{w}_g$	0.6	N
Numerical slack	ε_s	10^{-6}	–
NMPC weights			
Position tracking	$\mathbf{Q}_p$	$\mathrm{diag}(50,\ 50,\ 80)$	–
Velocity tracking	$\mathbf{Q}_v$	$\mathrm{diag}(5,\ 5,\ 8)$	–
Acceleration tracking	$\mathbf{Q}_a$	$\mathrm{diag}(1,\ 1,\ 2)$	–
Terminal position tracking	$\mathbf{Q}_f$	$\mathrm{diag}(80,\ 80,\ 120)$	–
Simulation setup			
Prediction horizon	N	30	–
Sampling time	T_s	0.01	s
Simulation time	T_{sim}	50	s
Prediction window	NT_s	0.30	s

TABLE II: Comparison between nominal and robustified NMPC over 500 uncertainty realizations.

Controller	Min. clearance	Violated samples	Avg. solve time
Nominal NMPC	0.182 [m]	230	9.4 [ms]
Robustified NMPC	0.364 [m]	0	18.7 [ms]

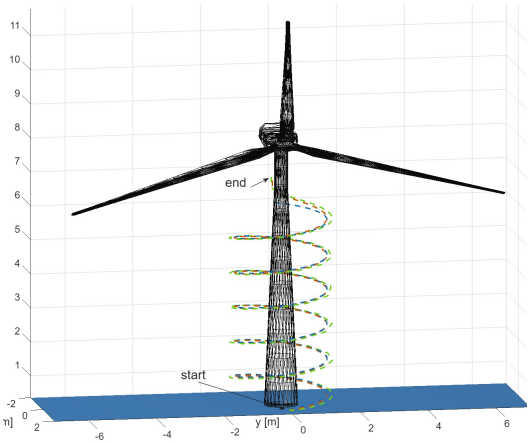

Fig. 1: Representative closed-loop trajectories under the same uncertainty realization. The blue curve denotes the helical inspection reference, the red curve the trajectory generated by the nominal NMPC, and the green curve the trajectory generated by the robustified NMPC. The robustified controller preserves a larger distance from the tower.

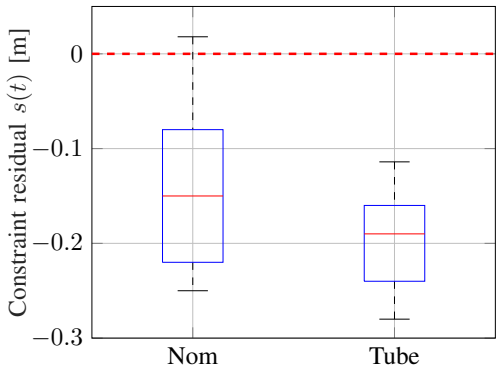

Fig. 2: Boxplots of the signed clearance residuals $s(t) = d_{\min} - d_T(\mathbf{p}(t))$ over 500 Monte-Carlo trials.

To complement the single-run illustration, Table II reports the main numerical results obtained over a Monte-Carlo campaign with 500 independent uncertainty realizations.In each trial, the uncertain parameters are sampled uniformly within the prescribed bounds. Here, "violated samples" denotes the total number of closed-loop time instants, aggregated over the full campaign, at which the minimum-clearance constraint is violated. To further characterize the distribution of these safety-critical events, Figure 2 shows the signed clearance residual $s(t) = d_{\min} - d_T(\mathbf{p}(t))$, for which $s(t) \leq 0$ indicates constraint satisfaction and $s(t) = 0$ corresponds to constraint activation. The nominal NMPC distribution is centered near -0.15 m, reflecting the intentional offset between the desired inspection distance d_{ref} and the safety threshold $d_{\min}$. However, under the combined effect of wind gusts and structured model mismatch, the nominal controller does not always compensate for the resulting deviations, leading to 230 violation instances over the full campaign, namely samples for which $s(t) > 0$. In contrast, the proposed robustified NMPC shifts the residual distribution further into the negative, safe region. By propagating parametric sensitivities and incorporating the stage-dependent gust margin $\alpha_{g,i}$, the controller proactively allocates additional clearance in high-sensitivity portions of the helical path. As a result, all simulated trials satisfy the minimum-clearance requirement $d_{\min} = 0.20$ m, yielding a 100% success rate at the cost of a moderate increase in computational effort, as also reflected in Table II.

VIII. Conclusions

This paper presented a preliminary sensitivity-based robustification of NMPC for close-proximity tower inspection with a tilted multirotor. The method augments the nominal formulation with a parametric tightening term and a stage-dependent gust margin, yielding a conservative clearance constraint under bounded uncertainty. Initial numerical results indicate improved clearance preservation relative to nominal NMPC. Future work will focus on broader validation, reduced conservatism, and field experiments in a mock-up inspection scenario.

REFERENCES

[1] Y. Liu *et al.*, "Review of robot-based damage assessment for offshore wind turbines," *Renewable and Sustainable Energy Reviews*, vol. 158, 2022, art. no. 112187.

[2] O. Moolan-Feroze *et al.*, "Improving Drone Localisation Around Wind Turbines Using Monocular Model-Based Tracking," in *IEEE International Conference on Robotics and Automation*, 2019, pp. 7713–7719.

[3] D. Bicego *et al.*, "Nonlinear Model Predictive Control with Enhanced Actuator Model for Multi-Rotor Aerial Vehicles with Generic Designs," *Journal of Intelligent & Robotic Systems*, vol. 100, pp. 1213–1247, 2020.

[4] R. Findeisen *et al.*, "State and Output Feedback Nonlinear Model Predictive Control: An Overview," *European Journal of Control*, vol. 9, no. 2–3, pp. 190–206, 2003.

[5] M. Hamandi *et al.*, "Design of multirotor aerial vehicles: A taxonomy based on input allocation," *The International Journal of Robotics Research*, vol. 40, no. 8-9, pp. 1015–1044, 2021.

[6] M. Ryll *et al.*, "6D interaction control with aerial robots: The flying end-effector paradigm," *The International Journal of Robotics Research*, vol. 38, no. 9, pp. 1045–1062, 2019.

[7] Y. Chen *et al.*, "MATMPC - A MATLAB Based Toolbox for Real-time Nonlinear Model Predictive Control," in *European Control Conference*, 2019, pp. 3365–3370.

[8] H. Ferreau *et al.*, "qpOASES: A parametric active-set algorithm for quadratic programming," *Mathematical Programming Computation*, vol. 6, no. 4, pp. 327–363, 2014.