

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

Giuseppe Silano<sup>1</sup> and Martin Saska<sup>2</sup>

**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 robust NMPC for a tilted multirotor that robustifies the tower-clearance constraint via online constraint tightening. First-order parametric state sensitivities provide a structured-uncertainty margin, while bounded gusts are handled by a stage-dependent additive margin. The formulation augments the nominal NMPC with sensitivity propagation and margin evaluation only, leaving the receding-horizon optimization structure unchanged. Monte-Carlo evaluation over 500 uncertainty realizations on a boundary-critical helical inspection trajectory shows that the proposed controller eliminates the clearance violations observed under nominal NMPC at the cost of a moderate increase in solve time.

## 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]–[3]. 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 [4], [5]. 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.

Robustifications of NMPC have been extensively studied in the process and robotics control communities. Min-max formulations explicitly optimize against worst-case disturbance realizations but are generically computationally prohibitive for fast aerial platforms [6]. Multi-stage and scenario-based NMPC mitigate this cost by branching the

prediction tree over a finite set of uncertainty realizations [7], but the combinatorial growth of scenarios remains demanding for the state dimensions and sample rates typical of multirotors. Tube-based NMPC, in either rigid or homothetic form, decouples a nominal trajectory from an ancillary feedback law that keeps the perturbed state inside a robust positive invariant set [8], [9]; for tilted multirotors, the joint presence of structured parametric mismatch (mass, inertia, thrust and drag effectiveness) and exogenous wind makes the construction of such invariant tubes both restrictive and conservative. Sensitivity-based constraint tightening, instead, propagates first-order parametric sensitivities along the prediction horizon and uses them to inflate the constraints on the fly [10], [11]. The proposed controller belongs to this last family, specialized to the cylindrical clearance constraint relevant for close-proximity tower inspection and complemented by a closed-form gust margin that does not enter the sensitivity propagation.

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.

## II. PROBLEM SETUP AND PREDICTION MODEL

This paper considers close-proximity helical tower inspection with guaranteed minimum clearance. The tower is modeled as a vertical cylinder defining both the inspection geometry and safety-critical clearance constraint, while the aerial platform uses 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  $\rho(\mathbf{p}) \triangleq \sqrt{x^2 + y^2}$  and  $d_T(\mathbf{p}) \triangleq \rho(\mathbf{p}) - R$ . Collision avoidance requires

$$d_T(\mathbf{p}(t)) \geq d_{\min}, \quad \forall t \geq 0, \quad (1)$$

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

$$d_T(\mathbf{p}_{\text{ref}}(t)) = d_{\text{ref}} > d_{\min}, \quad (2)$$

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

Partially funded by the Italian Electrical System research fund (decree n. 388, Nov. 6, 2024), GAČR project no. 26-22419S, CTU grant no. SGS26/077/OHK3/1T/13, and the European project “Robotics and Advanced Industrial Production” (no. CZ.02.01.01/00/22008/0004590). <sup>1</sup>G. Silano is with the Ricerca sul Sistema Energetico S.p.A., Milan, Italy, and the Czech Technical University in Prague, Czechia (giuseppe.silano@fel.cvut.cz). <sup>2</sup>M. Saska is with the Czech Technical University in Prague, Czechia (martin.saska@cvut.cz).

### B. Control-oriented GTMR model

The aerial platform is described by a control-oriented Generically-Tilted Multirotor (GTMR) model [12], [13] with input-rate actuation. Compared with collinear designs, the GTMR allocation provides non-zero lateral thrust components, which (i) partially decouple position and attitude control, allowing the vehicle to maintain a tower-facing orientation along the helical orbit without pitching toward the structure, and (ii) increase the admissible thrust direction set, improving rejection of lateral wind gusts during close-proximity flight. These features make the GTMR allocation particularly suitable for the safety-critical inspection scenario considered in this paper. Let  $N_p$  denote the number of rotor-propeller units, and define the system state as  $\mathbf{x} = [\mathbf{p}^\top \boldsymbol{\eta}^\top \mathbf{v}^\top \boldsymbol{\omega}^\top \boldsymbol{\xi}^\top]^\top$ , where  $\boldsymbol{\eta} = [\varphi \ \vartheta \ \psi]^\top \in \mathbb{R}^3$  is the roll-pitch-yaw Euler-angle attitude parametrization, used here as a local representation valid away from the gimbal-lock singularities encountered in close-proximity orbital flight,  $\mathbf{v} \in \mathbb{R}^3$  is the linear velocity expressed in the inertial frame,  $\boldsymbol{\omega} \in \mathbb{R}^3$  is the body angular velocity,  $\boldsymbol{\xi} \in \mathbb{R}^{N_p}$  collects the rotor-speed 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

$$\dot{\mathbf{x}} = f(\mathbf{x}, \mathbf{u}, \boldsymbol{\zeta}) + \mathbf{G}_w \mathbf{w}_g, \quad (3)$$

where  $\boldsymbol{\zeta}$  collects the structured uncertainty, including the slowly varying wind-induced force bias,  $\mathbf{w}_g$  is the bounded fast-varying gust force acting on the airframe, 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.

The structured uncertainty vector is

$$\boldsymbol{\zeta} = [\delta_m \ \delta_{J_x} \ \delta_{J_y} \ \delta_{J_z} \ \delta_T \ \delta_D \ w_{b,x} \ w_{b,y} \ w_{b,z}]^\top, \quad (4)$$

where  $\delta_m$ ,  $\delta_{J_x}$ ,  $\delta_{J_y}$ , and  $\delta_{J_z}$  denote relative mismatch in mass and principal inertia components,  $\delta_T$  and  $\delta_D$  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  $\mathcal{Z} = \{\boldsymbol{\zeta} \in \mathbb{R}^{n_\zeta} : \underline{\boldsymbol{\zeta}} \leq \boldsymbol{\zeta} \leq \bar{\boldsymbol{\zeta}}\}$ , where the inequalities are understood componentwise, and  $\underline{\boldsymbol{\zeta}}, \bar{\boldsymbol{\zeta}} \in \mathbb{R}^{n_\zeta}$ , with  $n_\zeta = 9$ , denote the lower and upper bounds on the uncertainty vector.

The aerodynamic force induced by wind is decomposed as  $\mathbf{w}(t) = \mathbf{w}_b + \mathbf{w}_g(t)$ , where  $\mathbf{w}_b$  is a slowly varying bias force included in  $\boldsymbol{\zeta}$ , and  $\mathbf{w}_g(t)$  is a fast-varying bounded gust force satisfying  $\|\mathbf{w}_g(t)\| \leq \bar{w}_g$  for all  $t \geq 0$ .

### IV. NOMINAL NMPC FORMULATION

With sampling time  $T_s$ , horizon  $N$ , and nominal uncertainty realization  $\boldsymbol{\zeta}_0 = \mathbf{0}$ , the discrete-time model reads  $\mathbf{x}_{i+1} = F(\mathbf{x}_i, \mathbf{u}_i, \boldsymbol{\zeta}_0)$ . The nominal NMPC problem, used as the baseline for the subsequent robust tightening, is

$$\min_{\{\mathbf{x}_i, \mathbf{u}_i\}} \sum_{i=0}^{N-1} \ell(\mathbf{x}_i, \mathbf{u}_i, \mathbf{r}_i) + \ell_f(\mathbf{x}_N, \mathbf{r}_N) \quad (5a)$$

$$\text{s.t. } \mathbf{x}_{i+1} = F(\mathbf{x}_i, \mathbf{u}_i, \boldsymbol{\zeta}_0), \quad i = 0, \dots, N-1, \quad (5b)$$

$$\mathbf{x}_i \in \mathcal{X}, \quad \mathbf{u}_i \in \mathcal{U}, \quad i = 0, \dots, N-1, \quad (5c)$$

$$d_T(\mathbf{p}_i) \geq d_{\min}, \quad i = 0, \dots, N, \quad (5d)$$

$$\mathbf{x}_0 = \mathbf{x}(t_k). \quad (5e)$$

The stage and terminal costs are chosen as

$$\ell = \|\mathbf{h}(\mathbf{x}_i) - \mathbf{r}_i\|_{\mathbf{Q}}^2 + \|\mathbf{u}_i\|_{\mathbf{Q}_u}^2, \quad \ell_f = \|\mathbf{h}_f(\mathbf{x}_N) - \mathbf{r}_N\|_{\mathbf{Q}_f}^2. \quad (6)$$

The reference  $\mathbf{r}_i$  samples the helical trajectory and includes desired position, velocity, and acceleration. The output map  $\mathbf{h}(\mathbf{x}) = [\mathbf{p}^\top \mathbf{v}^\top \dot{\mathbf{v}}^\top(\mathbf{x}, \mathbf{u})]^\top$  selects the tracked channels, with  $\dot{\mathbf{v}}$  evaluated along (3);  $\mathbf{h}_f(\mathbf{x}) = \mathbf{p}$ . The block-diagonal weight  $\mathbf{Q} = \text{blkdiag}(\mathbf{Q}_p, \mathbf{Q}_v, \mathbf{Q}_a)$  uses the values in Table I;  $\mathbf{Q}_u$  regularizes the actuator-rate input and  $\mathbf{Q}_f$  the terminal position error.

### V. SENSITIVITY-BASED TUBE TIGHTENING

The proposed robustification augments the nominal clearance constraint with a parametric margin for structured mismatch and a gust margin for fast wind disturbances, enforcing a conservative clearance condition along the 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

$$\boldsymbol{\Pi} \triangleq \frac{\partial \mathbf{x}}{\partial \boldsymbol{\zeta}} \in \mathbb{R}^{(12+N_p) \times 9}, \quad (7)$$

which measures the sensitivity of the extended state to perturbations of the structured uncertainty vector. Its continuous-time evolution, evaluated along the nominal predicted trajectory  $(\mathbf{x}_i, \mathbf{u}_i, \boldsymbol{\zeta}_0)$ , is

$$\dot{\boldsymbol{\Pi}} = \frac{\partial f}{\partial \mathbf{x}} \boldsymbol{\Pi} + \frac{\partial f}{\partial \boldsymbol{\zeta}}, \quad \boldsymbol{\Pi}(t_k) = \mathbf{0}, \quad (8)$$

where the initialization  $\boldsymbol{\Pi}(t_k) = \mathbf{0}$  reflects the reset of the sensitivity at the beginning of each prediction horizon. The corresponding discrete-time propagation  $\boldsymbol{\Pi}_{i+1} = F_{\boldsymbol{\Pi}}(\mathbf{x}_i, \mathbf{u}_i, \boldsymbol{\Pi}_i, \boldsymbol{\zeta}_0)$  is obtained consistently with the nominal discrete-time model.

The sensitivity propagation introduces  $n_\zeta n_x = 9(12+N_p)$  additional scalar states per stage with respect to the nominal NMPC, evaluated at negligible marginal cost via the same automatic differentiation toolchain.<sup>1</sup>

<sup>1</sup>For the GTMR considered in this paper ( $N_p = 6$ ), this amounts to 162 states per stage and 3240 states over the horizon, translating to the approximately twofold solve-time increase reported in Section VII.

### B. Clearance constraint in standard form

To derive the tightening term, the clearance requirement is rewritten as a scalar inequality. Specifically, define  $y(\mathbf{x}) \triangleq d_{\min} - d_T(\mathbf{p}) \leq 0$ , so that satisfying  $y(\mathbf{x}) \leq 0$  is equivalent to enforcing the minimum-clearance condition. Since the constraint depends only on the position components of the state, its Jacobian with respect to (w.r.t.) the system state reads

$$J_{yx}(\mathbf{x}) = \frac{\partial y}{\partial \mathbf{x}} = [-x/\rho \quad -y/\rho \quad 0 \quad \mathbf{0}_{1 \times (9+N_p)}], \quad (9)$$

where the trailing zero block reflects the absence of direct dependence on the remaining state components. The gradient  $\partial y/\partial p$  is well defined throughout the inspection mission since the helical reference enforces  $\rho(\mathbf{p}) \geq R + d_{\min} > 0$  at all times. Therefore, the first-order sensitivity of the scalar clearance constraint w.r.t. the structured uncertainty reads  $\Pi_y \triangleq \partial y/\partial \zeta = J_{yx}(\mathbf{x})\Pi$ , 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,

$$\mathbf{w}_\zeta \triangleq \zeta - \zeta_0, \quad (10)$$

which lies in the componentwise-bounded set

$$\mathcal{W}_\zeta \triangleq \{\mathbf{w}_\zeta \in \mathbb{R}^{n_\zeta} : |w_{\zeta,j}| \leq \bar{w}_{\zeta,j}, j = 1, \dots, n_\zeta\}, \quad (11)$$

where  $n_\zeta = 9$  and  $\bar{w}_{\zeta,j} > 0$  denotes the componentwise upper bound on the  $j$ -th uncertainty channel. Using the first-order approximation  $\Delta y_i \approx \Pi_{y,i} \mathbf{w}_\zeta$  and exploiting the box structure of  $\mathcal{W}_\zeta$ , the worst-case clearance perturbation at stage  $i$  admits a closed form as the support function of  $\mathcal{W}_\zeta$  evaluated at  $\Pi_{y,i}$ , attained by sign-aligning each component of  $\mathbf{w}_\zeta$  with the corresponding sensitivity. The parametric tightening margin is therefore defined as

$$\alpha_{p,i} \triangleq \sup_{\mathbf{w}_\zeta \in \mathcal{W}_\zeta} |\Pi_{y,i} \mathbf{w}_\zeta| + \varepsilon_s = \sum_{j=1}^{n_\zeta} |\Pi_{y,i,j}| \bar{w}_{\zeta,j} + \varepsilon_s, \quad (12)$$

where  $\Pi_{y,i,j}$  denotes the  $j$ -th component of  $\Pi_{y,i}$  and  $\varepsilon_s > 0$  is a small slack term introduced for numerical robustness. Equation (12) is tighter than the Hölder upper bound  $\|\Pi_{y,i}\|_\infty \|\bar{\mathbf{w}}_\zeta\|_1$  whenever the per-channel sensitivities differ in magnitude, which is typical here because  $\zeta$  aggregates parameters of heterogeneous physical scales (relative mass, inertia, propulsion, and drag mismatches, and wind-bias forces). Computing (12) requires only a weighted  $\ell_1$  norm of  $\Pi_{y,i}$  and therefore preserves the per-stage cost of the nominal NMPC.

### D. Wind-gust margin

The gust component is handled through an additive stage-dependent margin instead of being included in the sensitivity propagation. Since  $m \geq m_{\min} \triangleq (1 - \bar{\delta}_m)m_0$ , the bound  $\|\mathbf{w}_g(t)\| \leq \bar{w}_g$  implies  $\|\Delta \dot{\mathbf{v}}\| \leq \bar{a}_g \triangleq \bar{w}_g/m_{\min}$ . Assuming identical initial states for the nominal and perturbed trajectories at the start of the horizon and a worst-case open-loop

double integration of the gust acceleration over the prediction window—neglecting any in-horizon corrective action—the position deviation over  $\tau = iT_s$  satisfies  $\|\Delta \mathbf{p}_g(\tau)\| \leq \frac{1}{2} \bar{a}_g \tau^2$ , which yields the stage-dependent gust margin

$$\alpha_{g,i} \triangleq \frac{1}{2} \bar{a}_g (iT_s)^2. \quad (13)$$

Combining the parametric and gust margins yields the tightened clearance condition

$$d_T(\mathbf{p}_i) \geq d_{\min} + \alpha_{p,i} + \alpha_{g,i}, \quad (14)$$

or equivalently  $y(\mathbf{x}_i) + \alpha_{p,i} + \alpha_{g,i} \leq 0$ .

## 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

$$\min_{\{\mathbf{x}_i, \mathbf{u}_i, \Pi_i\}} \sum_{i=0}^{N-1} \ell(\mathbf{x}_i, \mathbf{u}_i, \mathbf{r}_i) + \ell_f(\mathbf{x}_N, \mathbf{r}_N) \quad (15a)$$

$$\text{s.t. } \mathbf{x}_{i+1} = F(\mathbf{x}_i, \mathbf{u}_i, \zeta_0), \quad i = 0, \dots, N-1, \quad (15b)$$

$$\Pi_{i+1} = F_\Pi(\mathbf{x}_i, \mathbf{u}_i, \Pi_i, \zeta_0), \quad i = 0, \dots, N-1, \quad (15c)$$

$$\mathbf{x}_i \in \mathcal{X}, \quad \mathbf{u}_i \in \mathcal{U}, \quad i = 0, \dots, N-1, \quad (15d)$$

$$y(\mathbf{x}_i) + \alpha_{p,i} + \alpha_{g,i} \leq 0, \quad i = 0, \dots, N, \quad (15e)$$

$$\mathbf{x}_0 = \mathbf{x}(t_k), \quad \Pi_0 = \mathbf{0}. \quad (15f)$$

Here,  $\Pi_i$  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  $i$ .

## 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  $d_T(\mathbf{p}) \geq d_{\min}$ . 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  $d_{\text{ref}}$  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 [14] in real-time iteration mode with qpOASES [15].

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.

TABLE I: Simulation parameters.

Parameter	Symbol	Value	Unit
<b>GTMR model</b>			
Number of rotors	$N_p$	6	–
Rotor tilt angle	$\alpha$	20	deg
Mass	$m$	2.57	kg
Principal inertia	$\mathbf{J}$	diag(0.11, 0.11, 0.19)	kg m <sup>2</sup>
Arm length	$d$	0.39	m
Rotor force coefficient	$c_f$	$11.8 \times 10^{-4}$	N/Hz <sup>2</sup>
Rotor torque coefficient	$c_t$	$2.5 \times 10^{-5}$	N m/Hz <sup>2</sup>
<b>Inspection geometry</b>			
Tower radius	$R$	1.00	m
Minimum clearance	$d_{\min}$	0.20	m
Desired offset	$d_{\text{ref}}$	0.35	m
<b>Uncertainty bounds</b>			
Mass mismatch	$\bar{\delta}_m$	0.10	–
Inertia mismatch	$\bar{\delta}_J$	[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
<b>NMPC weights</b>			
Position tracking	$\mathbf{Q}_p$	diag(50, 50, 80)	–
Velocity tracking	$\mathbf{Q}_v$	diag(5, 5, 8)	–
Acceleration tracking	$\mathbf{Q}_a$	diag(1, 1, 2)	–
Terminal position tracking	$\mathbf{Q}_f$	diag(80, 80, 120)	–
<b>Simulation setup</b>			
Prediction horizon	$N$	20	–
Sampling time	$T_s$	0.025	s
Simulation time	$T_{\text{sim}}$	50	s

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

Controller	Min. clearance	Trials with violation	Avg. solve time
Nominal NMPC	0.182 [m]	230/500	9.4 [ms]
Robustified NMPC	0.364 [m]	0/500	18.7 [ms]

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 II reports the Monte-Carlo results over 500 realizations with parameters sampled uniformly within the prescribed bounds; “trials with violation” counts realizations in which at least one closed-loop instant violates  $d_T(\mathbf{p}(t)) \geq d_{\min}$ . Figure 2 shows the signed clearance residual  $s(t) = d_{\min} - d_T(\mathbf{p}(t))$ . The nominal NMPC distribution is centered near  $-0.15$  m, reflecting the offset between  $d_{\text{ref}} = 0.35$  m and  $d_{\min} = 0.20$  m, but the combined effect of wind gusts and structured mismatch yields violations in 230 of 500 trials. The robustified NMPC shifts the distribution further into the safe region by allocating additional clearance in high-sensitivity portions of the orbit, with median clearance ( $\approx 0.40$  m) exceeding  $d_{\text{ref}}$ : the controller prioritizes the tightened safety constraint over reference tracking under uncertainty, as intended. This yields a 100% success rate at the cost of an approximately twofold solve-time increase. The reported MATLAB/MATMPC timings characterize rel-

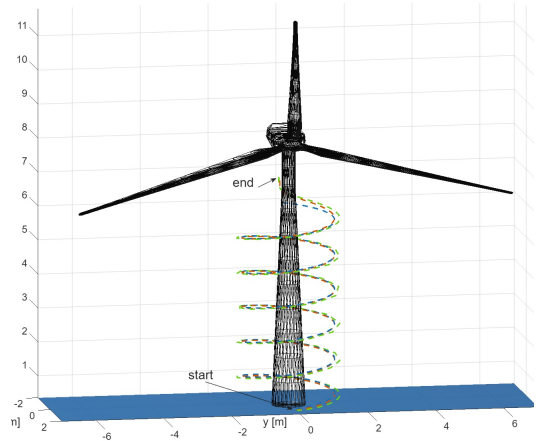


Fig. 1: Representative closed-loop trajectories under a wind bias  $\mathbf{w}_b = [0.6, 0.4, 0.0]^\top$  N: helical reference (blue), nominal NMPC (red), and robustified NMPC (green). The robustified controller anticipates the bias-induced drift and preserves a larger clearance from the tower.

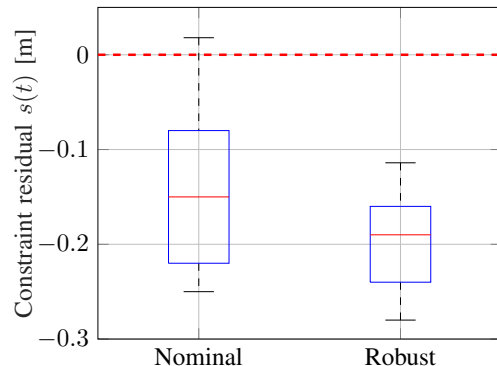


Fig. 2: Signed clearance residuals  $s(t) = d_{\min} - d_T(\mathbf{p}(t))$  over 500 Monte-Carlo trials;  $s(t) \leq 0$  indicates constraint satisfaction. The robustified controller keeps the entire distribution strictly below the safety threshold (dashed line).

ative overhead rather than embedded real-time feasibility; deployment will rely on a C-code-generated NMPC (e.g., via acados [16]) and is left to future work.

## 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, higher-order sensitivity terms to handle large parametric deviations, and field experiments in a mock-up inspection scenario.

## REFERENCES

- [1] Y. Liu, M. A. Hajj, R. L. W. van de Werken, and C. G. Soares, "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, K. Karachalios, D. N. Nikolaidis, and A. Calway, "Improving Drone Localisation Around Wind Turbines Using Monocular Model-Based Tracking," in *IEEE International Conference on Robotics and Automation*, 2019, pp. 7713–7719.
- [3] M. Car, L. Markovic, A. Ivanovic, M. Orsag, and S. Bogdan, "Autonomous Wind-Turbine Blade Inspection Using LiDAR-Equipped Unmanned Aerial Vehicle," *IEEE Access*, vol. 8, pp. 131 380–131 387, 2020.
- [4] D. Bicego, J. Mazzetto, M. Farina, R. Carli, and A. Franchi, "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.
- [5] R. Findeisen, F. Allgöwer, and B. A. Foss, "State and Output Feedback Nonlinear Model Predictive Control: An Overview," *European Journal of Control*, vol. 9, no. 2–3, pp. 190–206, 2003.
- [6] D. Q. Mayne, "Model predictive control: Recent developments and future promise," *Automatica*, vol. 50, no. 12, pp. 2967–2986, 2014.
- [7] S. Lucia, T. Finkler, and S. Engell, "Multi-stage nonlinear model predictive control applied to a semi-batch polymerization reactor under uncertainty," *Journal of Process Control*, vol. 23, no. 9, pp. 1306–1319, 2013.
- [8] P. Robuffo Giordano, Q. Delamare, and A. Franchi, "Trajectory Generation for Minimum Closed-Loop State Sensitivity," in *IEEE International Conference on Robotics and Automation*, 2018, pp. 286–293.
- [9] T. Belvedere, M. Cognetti, G. Oriolo, and P. Robuffo Giordano, "Sensitivity-Aware Model Predictive Control for Robots With Parametric Uncertainty," *IEEE Transactions on Robotics*, vol. 41, pp. 3039–3058, 2025.
- [10] A. Zanelli, J. Frey, F. Messerer, and M. Diehl, "Zero-order robust nonlinear model predictive control with ellipsoidal uncertainty sets," *IFAC-PapersOnLine*, vol. 54, no. 6, pp. 50–57, 2021.
- [11] B. Houska, M. E. Villanueva, and B. Chachuat, "Stable set-valued integration of nonlinear dynamic systems using affine set-parameterizations," *SIAM Journal on Numerical Analysis*, vol. 53, no. 5, pp. 2307–2328, 2015.
- [12] M. Hamandi, F. Usai, Q. Sablé, N. Staub, M. Tognon, and A. Franchi, "Design of multicopter aerial vehicles: A taxonomy based on input allocation," *The International Journal of Robotics Research*, vol. 40, no. 8-9, pp. 1015–1044, 2021.
- [13] 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.
- [14] Y. Chen, M. Bruschetta, E. Picotti, and A. Beghi, "MATMPC - A MATLAB Based Toolbox for Real-time Nonlinear Model Predictive Control," in *European Control Conference*, 2019, pp. 3365–3370.
- [15] H. Ferreau, C. Kirches, A. Potschka, H. Bock, and M. Diehl, "qpOASES: A parametric active-set algorithm for quadratic programming," *Mathematical Programming Computation*, vol. 6, no. 4, pp. 327–363, 2014.
- [16] R. Verschuere, G. Frison, D. Kouzoupis, J. Frey, N. van Duijkeren, A. Zanelli, B. Novoselnik, T. Albin, R. Quirynen, and M. Diehl, "acados - a modular open-source framework for fast embedded optimal control," *Mathematical Programming Computation*, vol. 14, pp. 147–183, 2021.