IEEE Access
Multidisciplinary ⋮ Rapid Review ⋮ Open Access Journal

Date of publication xxxx 00, 0000, date of current version xxxx 00, 0000.

Digital Object Identifier 10.1109/ACCESS.2024.0429000

Hybrid force–motion control with learning-augmented surface-normal estimation for contact-rich robotic swabbing

HARUN TUGAL[1], (Member, IEEE), WILLIAM HARWIN[1], LUCAS BROWN[1], KAIQIANG ZHANG[1] and ROBERT SKILTON[1], (Member, IEEE)

[1]Authors are with UK Atomic Energy Authority (UKAEA), Culham Science Centre, Abingdon, Oxfordshire OX14 3DB, United Kingdom.

Corresponding author: Harun Tugal (e-mail: harun.tugal@ukaea.uk).

This work has been carried out within the framework of the EUROfusion Consortium by the UK Atomic Energy Authority (UKAEA). The work was funded by the UKAEA/EPSRC Fusion Grant 2022/27 (grant number EP/W006839/1). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union, the European Commission, or the UK Atomic Energy Authority.

ABSTRACT This paper presents an operational-space framework for regulating contact forces on geometrically unknown surfaces by integrating online surface-normal estimation with hybrid motion–force control via task-space admittance. Cartesian motion is decomposed into normal and tangential components defined with respect to an estimated local surface frame. Surface normals are computed using a combination of force–velocity-based estimation with friction compensation, local quadratic surface fitting, and a regression-based prediction method.

The estimated normal is used to define the force-control direction and to align the end-effector orientation during contact. The approach is validated in simulation on a cylindrical surface and experimentally on a UR5 manipulator equipped with a wrist-mounted force/torque sensor. Results demonstrate sub-3° mean angular error under ideal conditions and consistent reductions in angular misalignment and RMSE when geometric and regression-based refinements are incorporated, while maintaining stable force tracking. The proposed framework enables robust surface interaction for contact-intensive tasks such as robotic swabbing.

INDEX TERMS Admittance control, Contact-rich manipulation, Gaussian process regression, Hybrid motion–force control, Quadratic surface fitting, Surface normal estimation, Robotic swabbing.

I. INTRODUCTION

Robotic manipulation in contact-rich environments requires tight coupling between geometric estimation and force regulation. Among the geometric quantities involved in physical interaction, the surface normal is particularly critical, as it directly determines tool orientation, the direction of applied contact forces, and ultimately task success. Accurate and timely estimation of surface normals enables robots to interact safely and effectively with surfaces of unknown or varying geometry in applications ranging from industrial surface processing and precision assembly to medical and service robotics.

In contact-intensive tasks such as polishing, grinding, deburring, and assembly, maintaining a prescribed normal force while following surface contours is essential for achieving consistent process quality (e.g., see Figure1). Purely position-based control strategies are ill-suited to these scenarios, as they cannot accommodate geometric uncertainty, surface compliance, or unmodeled contact dynamics. These limitations motivated the development of hybrid position/force control frameworks, which explicitly decompose the task space into motion-controlled and force-controlled directions [1]. By regulating motion along an unconstrained tangential direction while enforcing force constraints along a contact normal direction, hybrid controllers provide a solution for managing the dual objectives of geometric tracking and force regulation.

The effectiveness of hybrid position/force control is closely tied to the availability of reliable surface normal information. In practice, estimating surface normals during dynamic contact remains challenging due to sensor noise, highly variable frictional effects, tool compliance, and computational delays [2], [3]. Recent advances in force-based observers, adaptive estimation techniques, and learning-based approaches have shown promise in addressing these challenges, enabling on-

FIGURE 1: Sample retrieval trial using a robotic mobile manipulator equipped with a swabbing tool for radiological contamination monitoring (Source: Robotics & AI Collaboration/UKAEA)

line normal estimation without requiring explicit geometric models of the environment [4], [5]. These developments have expanded the applicability of hybrid control to unstructured and variable environments.

Operational space control provides a natural framework for integrating geometric estimation and force regulation, as it formulates control objectives directly in task-relevant Cartesian coordinates. Within this paradigm, admittance control is particularly attractive for contact tasks, as it maps measured interaction forces to commanded motion, thereby yielding compliant behavior that can be specified explicitly in task space [6]. Operational space admittance control is well suited to surface-following applications, where compliance must be defined relative to the surface normal and tangential directions. Moreover, adaptive admittance strategies can accommodate variations in contact stiffness and surface geometry encountered during task execution [7].

Despite substantial progress in surface normal estimation, hybrid position/force control, and operational space admittance control as largely independent research areas, relatively few studies have presented unified frameworks that integrate these components into a cohesive system with validated performance and stability considerations. Existing works typically focus on individual aspects—such as observer design for normal estimation, hybrid control structures for specific tasks, or admittance schemes for compliant interaction—without fully addressing their mutual coupling in real-time operation [8], [9]. As a result, important questions remain regarding robustness, stability, speed, and practical deployment in complex contact scenarios.

This article addresses this gap by presenting an integrated framework that combines online surface normal estimation with a hybrid position/force controller implemented in operational space using admittance control. The proposed approach enables robots to perform contact-rich tasks on surfaces of unknown or varying geometry by: (i) estimating surface normals online from force/torque sensing and kinematic information; (ii) decomposing the task into tangential position control and normal force control through a hybrid formulation; and (iii) implementing the force-controlled subspace via an operational space admittance controller that ensures compliant and stable interaction. The overall architecture is designed for real-time implementation and robust performance under geometric and dynamic uncertainties.

II. RELATED WORK

A. SURFACE NORMAL ESTIMATION

Force/torque sensing at the end-effector provides direct information about contact interactions and has been widely exploited for geometric estimation. Sloth et al. [2] proposed an unknown-input observer that estimates both the contact point and surface normal using force/torque measurements, explicitly accounting for tool flexibility. Their method was validated in simulation and experiments, demonstrating improved tracking performance when integrated with adaptive parallel position/force control. Similarly, Li et al. [4] demonstrated real-time normal force regulation in surface processing tasks using a hierarchical macro–mini robotic system, achieving robust performance without relying on prior geometric models of the workpiece.

Online normal estimation has also been incorporated directly into hybrid control frameworks. Dimeas and Doulgeri [1] introduced a progressive automation approach in which surface normals are estimated online to adapt the force direction of a hybrid controller for planar periodic tasks. Their framework enables rapid adaptation to rotated or misaligned surfaces and supports a transition from human-guided demonstrations to autonomous execution.

Hybrid force–motion estimation strategies further exploit the coupling between commanded motion and measured forces. Nasiri and Wang [8] estimated surface normals using force measurements and velocity information while compensating for friction-induced biases common in manufacturing environments. Their approach was implemented on a seven-degree-of-freedom manipulator and demonstrated real-time performance in surface-following tasks with uncertain geometry.

Learning-based approaches have recently gained attention as flexible alternatives to purely model-based estimators. Xu et al. [5] proposed a neural network model that predicts local surface orientation to adapt end-effector alignment during polishing operations. Experimental results on curved components showed low angular estimation errors and improved surface quality. While learning-based methods offer adaptability to complex geometries, their integration with control architectures raises questions regarding interpretability, generalization, and stability.

Overall, real-time surface normal estimation remains challenging due to sensor noise, friction effects, tool compliance, and computational constraints [2], [8]. Recent trends emphasize hybrid approaches that combine physics-based observers with data-driven components to balance robustness and adaptability [5], [10].

B. HYBRID POSITION/FORCE CONTROL

Hybrid position/force control has long served as a foundational paradigm for robotic interaction with the environment. Classical formulations decompose the task space into position-controlled and force-controlled subspaces using selection matrices, enabling simultaneous regulation of motion and contact forces [11]. This structure aligns naturally with contact tasks in which tangential directions require precise motion control while normal directions demand force regulation.

Hybrid position/force control has been applied across a wide range of domains. In medical robotics, Jaroonsorn et al. [12] employed hybrid control to ensure accurate positioning and safe contact forces during robot-assisted transcranial magnetic stimulation. In manufacturing, hybrid force–motion frameworks with online geometric estimation enable robust surface-following and contact tracing under geometric uncertainty [8]. Assembly tasks, such as peg-in-hole insertion, benefit from force/motion tracking methods that estimate contact geometry to support precise alignment [13].

Despite these advances, challenges remain in guaranteeing stability during mode transitions, managing sensor noise and model uncertainty, and integrating hybrid control with higher-level planning and perception systems [2], [6].

Operational space (task space) control provides a framework for expressing robot dynamics and control objectives directly in task-relevant coordinates, simplifying the specification of desired behaviors for manipulation tasks.

C. OPERATIONAL SPACE AND ADMITTANCE CONTROL

Compliance in operational space is commonly achieved through impedance or admittance control. Impedance control maps motion errors to interaction forces, whereas admittance control maps measured forces to commanded motion [14], [15]. Admittance control is especially suitable for robots with high-bandwidth position or velocity control, as it allows compliant behavior to be implemented on top of an inner control loop.

Task-space admittance controllers enable explicit specification of compliance along normal and tangential directions, making them well suited for surface-following and force-maintenance tasks. Hierarchical formulations based on HQP integrate admittance behaviors with task prioritization and secondary objectives, such as joint limit avoidance or redundancy resolution [6]. Adaptive admittance strategies further adjust compliance parameters online to cope with varying contact dynamics and environmental uncertainty [7].

Stability remains a central concern in admittance control, particularly under varying contact stiffness, sensing delays, and uncertain geometry [16], [17]. Passivity-based design principles and conservative parameter selection are commonly employed to preserve stability margins [18]. Recent studies have explored energy-based and passivity-enforcing methods to ensure stable interaction across a wide range of environments [19].

D. INTEGRATION OF ESTIMATION AND CONTROL

Several recent works have demonstrated partial integration of surface normal estimation with hybrid or admittance control. Nasiri and Wang [8] combined real-time normal estimation with a hybrid force–motion controller on a redundant manipulator, validating the approach in manufacturing tasks. Xu et al. [5] integrated neural network-based normal prediction with adaptive force control for polishing curved surfaces, achieving improved surface quality through orientation alignment and force regulation. For flexible tools, Sloth et al. [2] estimated contact geometry to support adaptive parallel position/force control.

Although these studies demonstrate the feasibility of integrated estimation and control, most address only subsets of the overall problem. In particular, few approaches provide an end-to-end operational space admittance controller that incorporates online surface normal estimation, hybrid task decomposition, and explicit consideration of stability in a single validated framework. Addressing this gap is essential for reliable deployment of contact-rich robotic systems in unstructured and safety-critical environments.

III. NOTATION AND PRELIMINARIES

Throughout this paper, bold lower-case letters $\mathbf{x}$ denote column vectors, non-bold upper-case letters X denote matrices, and non-bold lower-case letters x denote scalar quantities. The transpose operator is denoted by $(\cdot)^\top$. The identity and zero matrices of appropriate dimensions are denoted by I and 0, respectively. The skew-symmetric operator mapping a vector $\mathbf{r} \in \mathbb{R}^3$ to its matrix form is denoted by $[\mathbf{r}]_\times$, such that $[\mathbf{r}]_\times \mathbf{a} = \mathbf{r} \times \mathbf{a}$.

Consider an n-degree-of-freedom serial manipulator with joint coordinates $\mathbf{q} \in \mathbb{R}^n$. The end-effector pose in operational space is defined as

$$\mathbf{x} = [\mathbf{p}^\top, \boldsymbol{\phi}^\top]^\top \in \mathbb{R}^6,$$

where $\mathbf{p} \in \mathbb{R}^3$ represents position and $\boldsymbol{\phi} \in \mathbb{R}^3$ represents orientation.

The differential kinematics are given by

$$\dot{\mathbf{x}} = J(\mathbf{q})\dot{\mathbf{q}},$$

where $J(\mathbf{q}) \in \mathbb{R}^{6\times n}$ is the geometric Jacobian.

The joint-space dynamics are described by

$$M(\mathbf{q})\ddot{\mathbf{q}} + C(\mathbf{q}, \dot{\mathbf{q}})\dot{\mathbf{q}} + \mathbf{g}(\mathbf{q}) = \boldsymbol{\tau} + \boldsymbol{\tau}_{\text{ext}},$$

where $M(\mathbf{q})$ is the inertia matrix, $C(\mathbf{q}, \dot{\mathbf{q}})$ captures Coriolis and centrifugal effects, $\mathbf{g}(\mathbf{q})$ is the gravity vector, $\boldsymbol{\tau}$ is the commanded joint torque, and $\boldsymbol{\tau}_{\text{ext}}$ represents torques induced by external contact.

In operational space, the dynamics can be written as

$$\boldsymbol{\Lambda}(\mathbf{q})\ddot{\mathbf{x}} + \boldsymbol{\mu}(\mathbf{q}, \dot{\mathbf{q}}) + \mathbf{p}(\mathbf{q}) = \mathbf{F} + \mathbf{F}_{\text{ext}},$$

where $\boldsymbol{\Lambda} = \left(JM^{-1}J^{\top}\right)^{-1}$ is the operational-space inertia matrix, $\boldsymbol{\mu}$ and $\mathbf{p}$ denote the Coriolis/centrifugal and gravity terms mapped to task space, and $\mathbf{F}_{\text{ext}}$ is the external wrench at the end-effector.

During contact, a local surface frame $\{\mathcal{S}\}$ is defined at the contact point. Let $\mathbf{n} \in \mathbb{R}^3$ denote the unit surface normal and $T = [\mathbf{t}_1\ \mathbf{t}_2]$ an orthonormal basis spanning the tangent plane. Any Cartesian vector $\mathbf{v} \in \mathbb{R}^3$ can be decomposed as

$$\mathbf{v} = T\mathbf{v}_t + \mathbf{n}v_n,$$

where $\mathbf{v}_t$ and v_n are the tangential and normal components. This decomposition underpins the hybrid position/force control formulation.

IV. METHODOLOGY

This section presents the proposed framework, which integrates online surface normal estimation with hybrid position/force control implemented via operational space admittance control. The objective is to execute a surface-following task on an unknown or partially known surface while: (i) tracking a desired tangential trajectory, (ii) maintaining a specified normal contact force, and (iii) aligning the end-effector orientation with the local surface normal.

The surface normal $\mathbf{n}$ is not known a priori and must be estimated online from force/torque and kinematic measurements.

A. HYBRID POSITION/FORCE CONTROL VIA OPERATIONAL SPACE ADMITTANCE

A six-axis force/torque (FT) sensor is assumed to be mounted at the robot end-effector to measure interaction wrenches during contact. The measured wrench is used to generate compliant motion commands through a task-space admittance controller, enabling simultaneous regulation of contact forces and motion.

Let $\mathbf{f}_r, \mathbf{f}_e \in \mathbb{R}^6$ denote the reference and measured end-effector wrenches, respectively, and let $\mathbf{v}_f \in \mathbb{R}^6$ denote the desired task-space velocity generated by the admittance controller. The admittance dynamics are defined as

$$\mathbf{f}_r - \mathbf{f}_e = M_a\dot{\mathbf{v}}_f + D_a\mathbf{v}_f, \tag{1}$$

where $M_a, D_a \in \mathbb{R}^{6\times 6}$ are diagonal matrices representing the virtual inertia and damping of the admittance model. These parameters are independently tuned for translational and rotational subspaces to balance responsiveness and stability during contact.

In practical scenarios, the FT sensor is not colocated with the actual contact point due to tooling. Let $\mathbf{r}_{cs} \in \mathbb{R}^3$ denote the position vector from the sensor frame origin to the contact point, expressed in the sensor frame. The wrench at the contact point $\mathbf{f}_e$ is obtained from the sensor measurement $\mathbf{f}_s$ as

$$\mathbf{f}_e = \begin{bmatrix} \mathrm{I} & 0 \\ [\mathbf{r}_{cs}]_\times & \mathrm{I} \end{bmatrix} \mathbf{f}_s, \tag{2}$$

where $[\cdot]_\times$ denotes the skew-symmetric operator. Prior to this transformation, gravitational and inertial effects of the tool are compensated, and the sign of the measured wrench is inverted to reflect the force exerted by the robot on the environment.

At the contact point, the task objective is to maintain a constant normal contact force while minimizing parasitic torques that induce misalignment. Accordingly, the reference wrench is defined as $\mathbf{f}_r = \begin{bmatrix} 0 & 0 & f_z^r & 0 & 0 & 0 \end{bmatrix}^\top$, where f_z^r is the desired normal force expressed in the local task frame aligned with the estimated surface normal. Zero reference torques about the tangential axes enforce perpendicular alignment of the end-effector with respect to the surface.

To decouple motion and force objectives, the velocity command generated by the admittance controller is filtered through a task selection matrix $\Omega = \mathrm{diag}(0, 0, 1, 1, 1, 0)$, which suppresses undesired motion components and ensures that only the intended force-controlled and orientation-controlled directions influence the manipulator motion.

The desired task-space velocity is mapped to joint space using the manipulator Jacobian $J(\mathbf{q})$:

$$\dot{\mathbf{q}}_r = J^{\dagger}(q_m)\mathbf{v}_r, \tag{3}$$

where $J^{\dagger}$ denotes the generalized inverse of the Jacobian. Near kinematic singularities, direct inversion may lead to unstable joint velocities. To ensure numerical robustness and safe operation, a damped least-squares formulation is employed $J^{\dagger} = J^{\top}(JJ^{\top} + \lambda^2 I)^{-1}$, with damping factor $\lambda > 0$. This approach guarantees bounded joint velocities while preserving task execution accuracy [20].

B. SURFACE NORMAL ESTIMATION

Accurate estimation of the local surface normal is a key part of the depicted control architecture shown in Figure 2. The estimated normal defines the force-controlled direction v_n, determines the task decomposition inside the admittance controller, and generates the orientation correction required to align the tool with the contacted surface along with the admittance controller.

In the implemented framework, surface normal estimation is performed online and updated at each control cycle. Three complementary mechanisms are employed: (i) force–velocity decomposition, (ii) local quadratic surface fitting, and (iii) predictive Gaussian Process Regression (GPR).

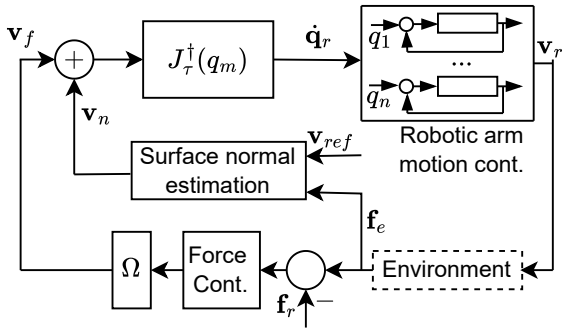

FIGURE 2: Block diagram representation of the implemented position–force control algorithm for a manipulator in contact with the environment, where task-space velocity commands are mapped to joint-space commands using a damped least-squares pseudo-inverse of the Jacobian.

a: Reactive force–velocity decomposition

During sliding contact, the measured interaction force is decomposed into tangential and normal components using the instantaneous tool velocity direction as in [8]. Let $\mathbf{f} \in \mathbb{R}^3$ denote the bias-compensated contact force at the contact point (obtained from $\mathbf{f}_e$), and let $\mathbf{v} \in \mathbb{R}^3$ denote the commanded translational velocity component of $\mathbf{v}_f$.

A projection operator along the velocity direction is defined as

$$\Omega_v = \mathbf{v}(\mathbf{v}^\top \mathbf{v})^{-1}\mathbf{v}^\top .$$

The tangential force component is

$$\mathbf{f}_t = (I - \Omega_v)\mathbf{f}.$$

Assuming Coulomb friction, the estimated friction force is

$$\mathbf{f}_\tau = -\hat{\mu}\|\mathbf{f}_t\|\frac{\mathbf{v}}{\|\mathbf{v}\|}.$$

The reactive surface normal estimate is then

$$\hat{\mathbf{n}}_s = \frac{\mathbf{f} - \mathbf{f}_\tau}{\|\mathbf{f} - \mathbf{f}_\tau\|}.$$

This estimate is subsequently transformed to the base frame as $\hat{\mathbf{n}}_b = R_b^e\,\hat{\mathbf{n}}_s$, with the rotation matrix from the end-effector frame to the base frame (R_b^e). The base-frame representation is used for geometric processing and learning.

The friction coefficient $\hat{\mu}$ is updated online using a weighted moving-average filter. Although this reactive estimator provides fast adaptation, it remains sensitive to measurement noise and stick–slip phenomena.

b: Local quadratic surface fitting

To improve geometric robustness, recent contact positions $\{\mathbf{p}_i\}_{i=1}^N$ (expressed in the base frame via forward kinematics) are stored in a sliding window. A second-order surface model

$$z = ax^2 + by^2 + cxy + dx + ey + f\ ,$$

is fitted using least-squares regression. The surface gradient evaluated at the most recent contact point (x_0, y_0) is

$$\nabla z = \begin{bmatrix} 2ax_0 + cy_0 + d \\ 2by_0 + cx_0 + e \\ -1 \end{bmatrix},$$

and the geometric normal is

$$\hat{\mathbf{n}}_q = \frac{\nabla z}{\|\nabla z\|}\,.$$

Quadratic fitting captures local curvature and provides improved robustness over purely force-based estimates, particularly on smooth curved geometries.

c: Predictive Gaussian Process normal estimation

To incorporate spatial memory and enable anticipatory correction, contact samples $(\mathbf{p}_i, \hat{\mathbf{n}}_i)$ are stored in a bounded buffer and used to construct a local Gaussian Process model with squared-exponential kernel

$$k(\mathbf{x}, \mathbf{x}_i) = \sigma_f \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}_i\|^2}{2\ell^2}\right).$$

Instead of querying only at the current contact location, the model predicts the normal at a short look-ahead point

$$\mathbf{x}_{\text{ahead}} = \mathbf{x} + \delta\frac{\mathbf{v}}{\|\mathbf{v}\|},$$

where δ is proportional to tangential speed and bounded for stability. The predicted normal is

$$\hat{\mathbf{n}}_p = \frac{\sum_i k(\mathbf{x}_{\text{ahead}}, \mathbf{x}_i)\hat{\mathbf{n}}_i}{\sum_i k(\mathbf{x}_{\text{ahead}}, \mathbf{x}_i)}.$$

This predictive mechanism reduces orientation lag and mitigates contact transients when traversing compound curvature.

d: Normal fusion

When a valid prediction exists, the reactive estimate and GPR prediction are blended:

$$\hat{\mathbf{n}} = (1 - \alpha)\hat{\mathbf{n}}_b + \alpha\hat{\mathbf{n}}_p, \qquad \alpha \in (0, 1),$$

and then normalized. In free space, the estimate defaults to the current tool z-axis to avoid spurious corrections.

e: Orientation alignment with normal estimation

Let $\mathbf{z}_{ee}$ denote the current tool z-axis expressed in the base frame. The orientation correction is computed using the geometric cross-product error

$$\boldsymbol{\omega}_c = -k_R\,\mathbf{z}_{ee} \times \hat{\mathbf{n}}, \tag{4}$$

which drives $\mathbf{z}_{ee}$ toward alignment with the estimated surface normal.

The tight coupling between reactive force decomposition, quadratic geometric fitting, predictive GPR estimation, and orientation feedback enables stable interaction with unknown and varying surface geometries while maintaining compliant force control.

C. SIMULATION AND EXPERIMENTAL SETUP

The proposed control architecture and surface-normal estimation framework were evaluated in both simulation and physical experiments. Simulation was performed in Gazebo using a kinematic and dynamic UR5 model, a simulated wrist-mounted force/torque (FT) sensor, and a cylindrical pipe $0.5\,\mathrm{m}$ diameter representing the contact surface. Ground-truth geometric information, including the analytical surface normal of the pipe, enabled quantitative assessment of the normal estimation accuracy and verification of the hybrid motion–force control performance prior to hardware deployment.

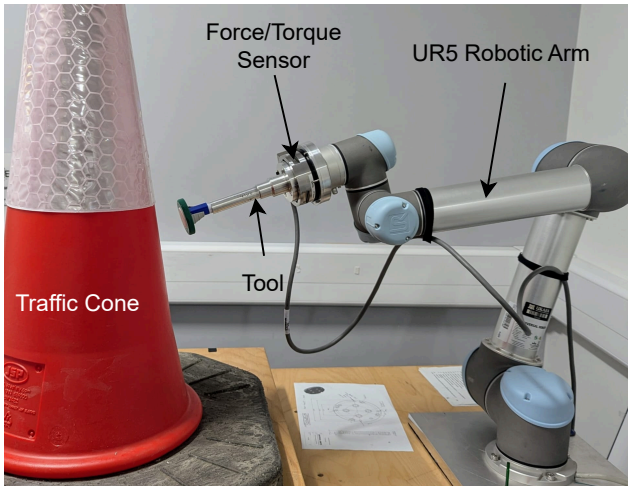

FIGURE 3: Experimental setup consisting of a UR5 robotic arm and ATI FT sensor. A traffic cone is used as the environment that the robot arm is interacting with.

Experimental validation was performed using a Universal Robots UR5 manipulator instrumented with an ATI Gamma force/torque (FT) sensor and an ATI FT 9105-NETB interface unit. The sensor was rigidly mounted to the end-effector via a custom-designed adaptor to enable controlled interaction with curved surfaces. A traffic cone of $0.75\,\mathrm{m}$ height served as the test object, see Figure 3. The robot was operated through its native controller at $125\,\mathrm{Hz}$, while all control, estimation, and data acquisition modules were implemented in ROS to ensure real-time execution and synchronized sensing. Experiments were conducted under fixed and repeatable initial conditions to facilitate consistent evaluation and direct comparison with the simulation results.

V. RESULTS

The experimental and simulation results are evaluated with respect to the closed-loop architecture defined in (1)–(4) and depicted in Figure 2. The admittance dynamics in (1) regulate the interaction behavior by mapping the wrench error to a desired task-space velocity, while the wrench transformation in (2) ensures that control actions are computed at the actual contact point. The resulting task-space command is projected into joint space through the generalized inverse kinematics in (3). Concurrently, orientation alignment is enforced via the geometric cross-product correction in (4), which drives the tool z-axis toward the estimated surface normal. The parameters used in the control architecture is given in Table 1.

TABLE 1: Controller parameters. Adt. C. denotes the admittance controller and Ort. C. denotes the orientation controller. Subscripts $_p$ and $_o$ denote the translational and rotational components, respectively.

Adt. C.	$m_p = 1\,\mathrm{kg}$	$m_o = 0.04\,\mathrm{kg}$	$d_p = 150\,\mathrm{Ns/m}$	$d_o = 1.5\,\mathrm{Ns/m}$
Ort. C.	$k_R = 0.08$			

A. SIMULATION RESULTS

The proposed estimation and control framework was validated on a UR5 platform during pipe- and cone-swabbing tasks under the setup described in Section IV-C. A raster (square sweeping) trajectory was executed in the Cartesian operational space along the x–y directions, generated via piecewise-constant velocity segments. Specifically, a linear velocity of 0.02, m/s was applied over 10, s intervals per axis, interleaved with pause phases, resulting in a periodic back-and-forth motion in x combined with stepwise transitions in y, forming a full square scan pattern. The trajectory was sampled at 125, Hz over a total duration of 30, s, and positions were obtained via discrete-time integration of the velocity profiles. The normal force reference was set to $f_z^{\mathrm{ref}} = 3\,\mathrm{N}$ for all experiments. Performance is assessed by measuring the angular difference between the estimated normal $\hat{\mathbf{n}}$ and the negative tool-frame z-axis (i.e., $-n_{\mathrm{tool}}$), which represents the desired contact orientation during perpendicular swabbing.

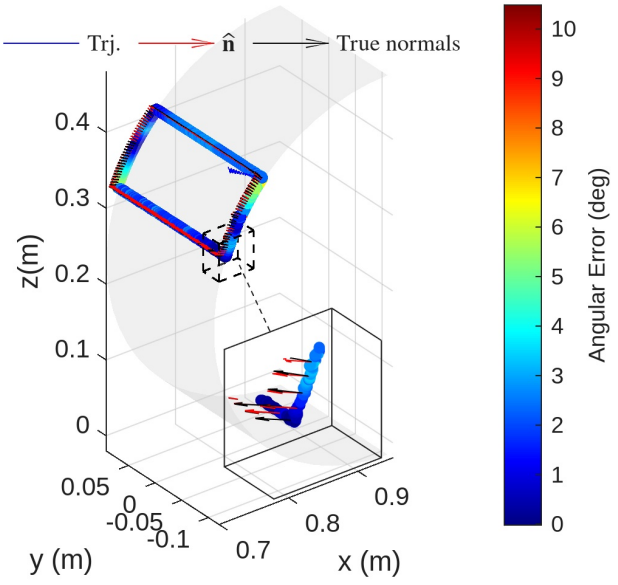

FIGURE 4: Robot trajectory, estimated surface normals obtained via GPR, ground-truth normals, and the angular error between the estimated and ground-truth normals during a pipe-swabbing simulation in Gazebo.

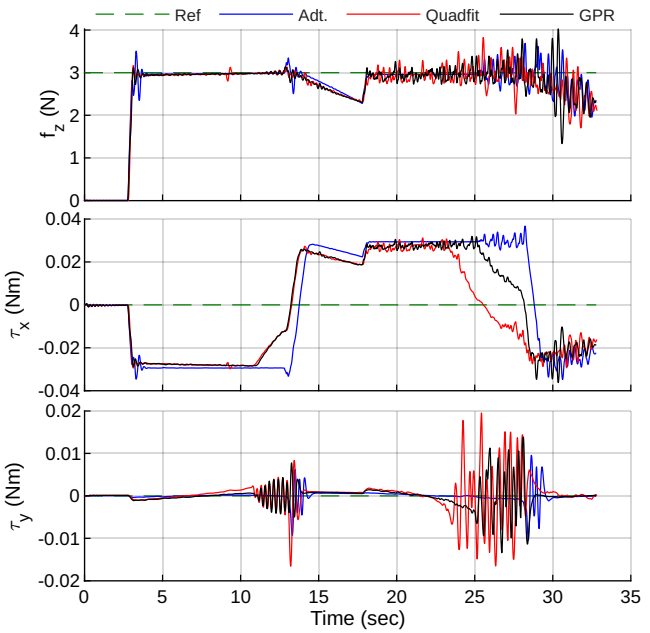

FIGURE 5: Simulated force in z and torques in x and y directions at tool frame; force tracking achieved once the contact is established and the tool remains perpendicular to the surface.

Figure 4 illustrates the robot trajectory, the estimated surface normals, the ground-truth normals of the cylindrical pipe, and the angular error between them during interaction. Since the simulated pipe geometry is analytically defined, the exact surface normal at each contact point is available, enabling direct quantitative comparison. The angular error between the estimated normal $\hat{\mathbf{n}}$ and the ground-truth normal $\mathbf{n}$ is computed as $\theta_{\text{error}} = \cos^{-1}\left(\hat{\mathbf{n}}^\top \mathbf{n}\right)$, expressed in degrees. Furthermore, the interaction forces and torques during the simulated swabbing task are shown in Figure 5. The force in the z direction converges to the reference value, while the torques about the x and y axes remain near zero, indicating that the estimated normals successfully maintain tool alignment with the pipe surface.

TABLE 2: Angular difference (degrees) between the estimated surface normal and the ground-truth surface normal in simulation.

Controller	Mean	Median	STD	RMSE
Adt. C.	2.713	2.483	3.958	4.798
Adt. C.+ Quadratic	2.368	2.173	3.957	4.61
Adt. C.+ GPR	2.411	2.190	3.929	4.609

Table 2 summarizes the statistical performance of the different estimation strategies integrated with the admittance controller. The baseline admittance controller without geometric refinement (Adt. C.) exhibits a mean angular error of $2.713°$ and an RMSE of $4.798°$. Incorporating local quadratic surface fitting (Adt. C.+Quadratic) reduces the mean error to $2.368°$ and the RMSE to $4.61°$. Similarly, the Gaussian Process Regression approach (Adt. C.+GPR) achieves a mean error of $2.411°$ and an RMSE of $4.609°$.

The results indicate that both learning-based (GPR) and analytical (quadratic fitting) refinement strategies improve normal estimation accuracy compared to the baseline controller. Since the quadratic model assumes a height-field representation $z = f(x, y)$, estimation accuracy can exhibit mild orientation dependence when the surface normal approaches the horizontal plane of the robot base frame. In the present experiments with a vertically oriented cone, the observed errors did not show a strong systematic dependence on surface direction, suggesting that the refinement methods remain effective over the tested range of surface orientations. The reduction in mean angular error, although moderate, demonstrates improved consistency in surface alignment during motion. The comparable RMSE and standard deviation values across methods suggest that the dominant error components arise primarily during transient contact phases rather than steady-state interaction.

Overall, the simulation results confirm that the proposed estimation framework achieves sub-$3°$ mean angular accuracy under ideal sensing conditions, validating the correctness of the formulation before experimental deployment.

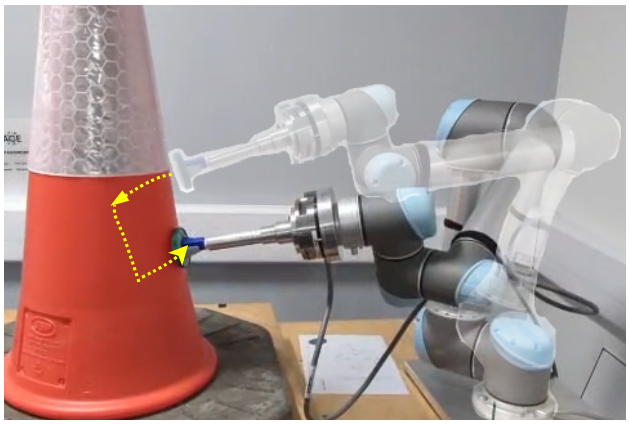

FIGURE 6: Example raster (square sweeping) trajectory executed by the UR5 robotic arm on a conical surface. The solid configuration shows the current robot pose, while the grayed configuration indicates a previous pose along the trajectory.

B. EXPERIMENTAL RESULTS

The proposed estimation and control framework was validated on a UR5 platform during pipe- and cone-swabbing tasks under the setup described in Section IV-C with above mentioned trajectory, , see Figure 6. Here the performance is assessed by measuring the angular difference between the estimated normal $\hat{\mathbf{n}}$ and the negative tool-frame z-axis (i.e., $-n_{\text{tool}}$), which represents the desired contact orientation during perpendicular swabbing.

Figure 7 depicts a representative cone-swabbing trial, showing the end-effector trajectory, the estimated surface normals,

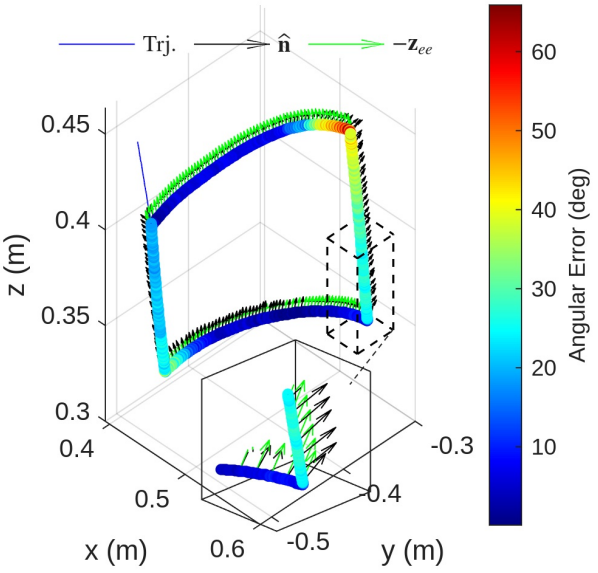

FIGURE 7: Robot trajectory, estimated surface normals obtained via GPR, the tool normals, and the angular error between the estimated and the negative tool normals ($-n_{tool}$) during cone-swabbing experiment.

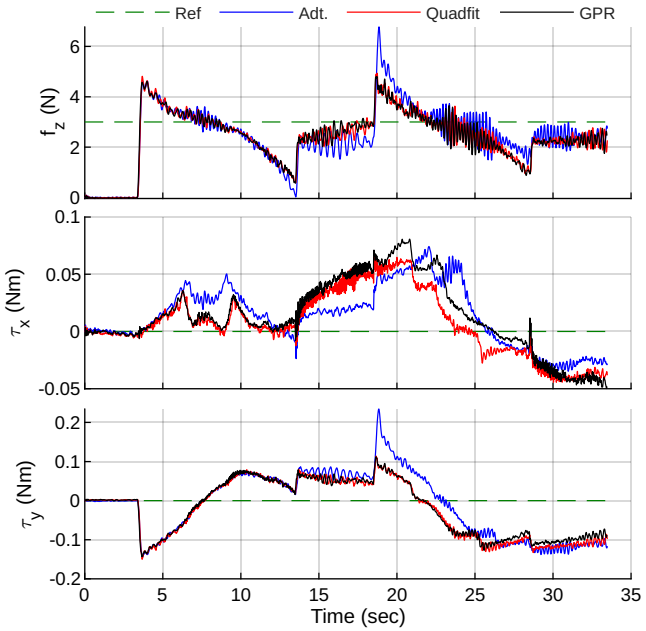

FIGURE 8: Measured force in z and torques in x and y directions at tool frame; force tracking achieved once the contact is established and the tool closely remains perpendicular to the surface.

and the angular error relative to $-n_{\text{tool}}$. Consistent with the intended behavior, the estimator aligns the tool approximately perpendicular to the local surface as the robot progresses along the object, with larger errors appearing predominantly during transitions, turns, and intermittent contact reconfigurations. Complementary to the geometric view, Figure 8 reports the measured interaction wrench at the tool frame. The normal force (f_z) converges to the reference value after contact establishment, while the torques about x and y remain near zero during steady sliding, indicating that the estimated normals contribute to maintaining perpendicular tool alignment.

TABLE 3: Angular difference (degrees) between the estimated surface normal and the negative tool normal ($-n_{tool}$) in experiments.

Controller	Mean	Median	STD	RMSE
Adt. C.	19.348	11.939	17.910	26.364
Adt. C.+ Quadratic	16.860	10.787	13.656	21.695
Adt. C.+ GPR	15.731	9.200	12.900	20.343

A quantitative summary of the angular discrepancy with respect to $-n_{\text{tool}}$ across controllers is provided in Table 3. The baseline admittance controller (Adt. C.) yields a mean angular error of $19.35°$ and an RMSE of $26.36°$. Incorporating the local quadratic surface-fitting refinement (Adt. C.+Quadratic) reduces the mean error to $16.86°$ and the RMSE to $21.70°$. The Gaussian Process Regression refinement (Adt. C.+GPR) further improves performance, achieving a mean error of $15.73°$ and an RMSE of $20.34°$. The reduction in both mean and RMSE indicates that model-based and learning-based refinements consistently enhance the surface-normal estimates relative to the admittance-only baseline.

Compared to simulation, the experimental errors are larger in both mean and dispersion, which is expected due to practical factors including sensor noise and bias in the FT measurement, unmodeled structural compliance at the tool–sensor–flange interface, residual kinematic/calibration errors, surface irregularities (e.g., seams and material discontinuities), and contact nonlinearities (e.g., friction and stick–slip). Notably, the improvement trends across methods are consistent with simulation results: both Quadratic and GPR refinements reduce the steady-state misalignment and attenuate the worst-case transients observable in the torque channels. These observations support the transferability of the estimation framework from idealized settings to real hardware.

Overall, the experiments demonstrate that the proposed hybrid motion–force controller, augmented with geometric refinement, achieves stable contact maintenance, reference force tracking, and improved surface alignment during swabbing on curved objects. While absolute angular accuracy is limited by sensing and interaction uncertainties, the systematic reduction in mean and RMSE with both refinements validates the practical value of embedding local surface modeling (Quadratic) or data-driven regression (GPR) within the estimation loop.

VI. CONCLUSION

This paper presented an integrated hybrid motion–force control framework with online surface-normal estimation for contact-intensive swabbing on curved geometries. Two refinement strategies—local quadratic surface fitting and Gaussian Process Regression (GPR)—were incorporated within an admittance control scheme to enhance normal estimation during physical interaction.

Simulation results demonstrated sub-3° mean angular error under ideal sensing conditions, validating the formulation and confirming that both refinement methods reduce mean error and RMSE relative to an admittance-only baseline, particularly during sustained contact. Experimental validation on a UR5 platform with an ATI wrist-mounted force/torque sensor showed consistent trends on real curved objects. Although absolute errors increased due to sensing noise, compliance, calibration imperfections, and contact nonlinearities, both refinement approaches improved alignment performance, with GPR yielding the lowest mean and RMSE angular discrepancies. Force tracking remained stable, and torque components stayed near zero during steady sliding, indicating effective perpendicular tool alignment.

Overall, embedding geometric or data-driven surface modeling within the estimation loop improves alignment accuracy without compromising contact stability, supporting practical deployment in swabbing and related contact-rich tasks. The gap between simulation and hardware performance motivates improvements in force/torque sensing (e.g., bias and temperature compensation and in-situ calibration), as well as better modeling of compliance and friction and adaptive tuning of refinement parameters. Future work will also address multi-contact and free-form surfaces with online coverage planning, and investigate fusion of vision or depth sensing with force feedback to enhance robustness and reduce angular error in unstructured environments. In scenarios where the robot performs repeated passes over the same region along similar trajectories, learning-based strategies such as iterative learning control (ILC) could further improve surface-normal prediction and alignment accuracy by exploiting information accumulated across successive executions.

REFERENCES

[1] F. Dimeas and Z. Doulgeri, "Progressive automation of periodic tasks on planar surfaces of unknown pose with hybrid force/position control," *IEEE International Conference on Intelligent Robots and Systems*, pp. 5246–5252, 2020.

[2] C. Sloth, A. Kramberger, E. L. Diget, and I. Iturrate, "Towards contact point and surface normal estimation for control of flexible tool," in *IEEE American Control Conference (ACC)*, 2022, pp. 500–505.

[3] I. Iturrate, E. L. Diget, and C. Sloth, "Surface following using direct adaptive admittance control," *IEEE/SICE International Symposium on System Integration (SII)*, pp. 1454–1459, 2025.

[4] J. Li, Y. Guan, H. Chen, B. Wang, T. Zhang, J. Hong, and D. Wang, "Real-time normal contact force control for robotic surface processing of workpieces without a priori geometric model," *The International Journal of Advanced Manufacturing Technology*, vol. 119, no. 3, pp. 2537–2551, 2022.

[5] J. Xu, H. Huang, H. Long, and S. Lei, "The adaptive trajectory of the normal force vector in the polishing of curved surface component robots," *Advanced Intelligent Systems*, vol. 7, no. 11, p. 2401044, 2025.

[6] M. Selvaggio, J. Cacace, C. Pacchierotti, F. Ruggiero, and P. R. Giordano, "A shared-control teleoperation architecture for nonprehensile object transportation," *IEEE Transactions on Robotics*, vol. 38, no. 1, pp. 569–583, 2022.

[7] F. Ferraguti, C. Talignani Landi, L. Sabattini, M. Bonfè, C. Fantuzzi, and C. Secchi, "A variable admittance control strategy for stable physical human–robot interaction," *The International Journal of Robotics Research*, vol. 38, no. 6, pp. 747–765, 2019.

[8] E. Nasiri and L. Wang, "Hybrid force motion control with estimated surface normal for manufacturing applications," in *IEEE 21st International Conference on Ubiquitous Robots (UR)*, 2024, pp. 125–132.

[9] S. Scherzinger, A. Roennau, and R. Dillmann, "Contact skill imitation learning for robot-independent assembly rogramming," in *IEEE International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 4309–4316.

[10] C. Y. Lin, C. C. Tran, S. H. Shah, and A. R. Ahmad, "Real-time robot pose correction on curved surface employing 6-axis force/torque sensor," *IEEE Access*, vol. 10, no. August, pp. 90 149–90 162, 2022.

[11] H. Tugal, K. Cetin, Y. Petillot, M. Dunnigan, and M. S. Erden, "Contact-based object inspection with mobile manipulators at near-optimal base locations," *Robotics and Autonomous Systems*, vol. 161, p. 104345, 2023.

[12] P. Jaroonsorn, P. Neranon, P. Smithmaitrie, and C. Dechwayukul, "Robot-assisted transcranial magnetic stimulation using hybrid position/force control," *Advanced Robotics*, vol. 34, no. 24, pp. 1559–1570, 2020.

[13] K. Van Wyk, M. Culleton, J. Falco, and K. Kelly, "Comparative peg-in-hole testing of a force-based manipulation controlled robotic hand," *IEEE Transactions on Robotics*, vol. 34, no. 2, pp. 542–549, 2018.

[14] N. Hogan, "Impedance control: An approach to manipulation," *IEEE American Control Conference*, vol. 1, no. March, pp. 304–313, 1984.

[15] O. Khatib, "A unified approach for motion and force control of robot manipulators: the operational space formulation," *IEEE Journal on Robotics and Automation*, vol. 3, no. 1, pp. 43–53, 1987.

[16] L. Pezeshki, H. Sadeghian, X. Chen, M. Keshmiri, S. Haddadin, and A. Mohebbi, "Assist-as-needed framework for robotic rehabilitation: adaptive admittance control with passivity-based safety features," *IEEE Transactions on Medical Robotics and Bionics*, pp. 1–12, 2025.

[17] Y. Huang and L. Yang, "Power adaptation-enabled admittance control for stable and safe actuated interaction in unmodeled environment," *IEEE Robotics and Automation Letters*, vol. 10, no. 3, pp. 2798–2805, 2025.

[18] D. Zeng, Y. Wang, Y. Jiang, J. Jiang, L. Wang, E. Ge, L. Huang, and H. Zhang, "Energy-aware adaptive force-admittance control for robotic manipulators: application in high-precision aircraft assembly," *IEEE Transactions on Industrial Electronics*, vol. 73, no. 2, pp. 2713–2724, 2025.

[19] F. Benzi and C. Secchi, "Unified power and admittance adaptation for safe and effective physical interaction with unmodelled dynamic environments," *IEEE Robotics and Automation Letters*, vol. 8, no. 12, pp. 8279–8286, 2023.

[20] R. S. Jamisola, D. N. Oetomo, M. H. Ang, O. Khatib, T. M. Lim, and S. Y. Lim, "Compliant motion using a mobile manipulator: An operational space formulation approach to aircraft canopy polishing," *Advanced Robotics*, vol. 19, no. 5, pp. 613–634, 2005.

HARUN TUGAL (M'21) received the B.Sc. degree in Electrical and Electronics Engineering from Karadeniz Technical University, Trabzon, Turkey, in 2011, and the M.Sc. degree in Advanced Control and Systems Engineering and the Ph.D. degree in Electrical and Electronics Engineering from the University of Manchester, Manchester, U.K., in 2013 and 2017, respectively. He was a Research Associate with the Institute of Sensors, Signals and Systems and the Edinburgh Centre for Robotics at Heriot-Watt University, Edinburgh, U.K., till 2021. He is currently a Principal Researcher at the U.K. Atomic Energy Authority (UKAEA). His research interests include adaptive, robust, and nonlinear control, human–robot interaction, bilateral telerobotics, haptic rendering, underwater robotics, and robotic systems for challenging environments. He is a member of the Automatic Control Engineering (ACE) Network.

WILLIAM HARWIN is an emeritus professor of human and interactive systems at the University of Reading and is a lead technologist at the UKAEA. He has a PhD in engineering from the University of Cambridge. His research interests are on understanding the mathematics and engineering of interacting systems, in particular understanding humans as a dynamic system who interact with real and simulated environments especially via telerobotics, exoskeletons or haptic devices. Particular projects include haptics in skills training and education, ubiquitous wearable sensors for healthcare assessments, control of contact forces in collaborative robotic systems, machine mediated stroke rehabilitation, human induced instability in haptic environments, and multi-finger bi-manual haptics.

Harwin is a senior member of the IEEE and a member of the IET, the IMA and the IPEM.

LUCAS BROWN received a MEng Degree in Electrical and Electronic Engineering from Imperial College in 2017. He is currently working at UKAEA, working on Control Systems for both Remote Maintenance and Fusion devices. He is particularly interested in System Identification and Model Predictive Control.

KAIQIANG ZHANG (M'19) received his B.Eng. degree in Automatic Control from the Xi'an Jiaotong University in 2012, his M.Sc. (dist.) degree in Advanced Engineering Robotics and Ph.D. degree in Control from the University of Bristol, in 2013 and 2019, respectively. Zhang is currently a Principal Researcher at the U.K. Atomic Energy Authority, where he is the Programme Area Manager of Robotics Research. His research interests include robust control, adaptive control, nonlinear modelling, safety guaranteed, control and the associated robotic applications. Zhang is an Honorary Professor with University of Nottingham, an executive member of IFAC Industry Committee, and a member of IEEE Nanotechnology committee.

ROBERT SKILTON received the B.Sc. and M.Sc. degrees in Computer Science and Cybernetics from the University of Reading, Reading, U.K., in 2006 and 2011, respectively, and the Ph.D. degree from the University of Surrey, Guildford, U.K., in 2023, where his research focused on generative artificial intelligence for visual inspection.

He is currently a Robotics Fellow and Head of Robotics Research and Technology at the U.K. Atomic Energy Authority (UKAEA), where he leads research on robotics and artificial intelligence for fusion energy applications and contributes to programmes in nuclear decommissioning and other high-integrity systems. His research interests include robot perception, control, autonomy, and the integration of artificial intelligence into real-world robotic systems operating in challenging environments.