

Whole-Body Ergodic Exploration with a Manipulator Using Diffusion

Cem Bilaloglu, Tobias Löw, and Sylvain Calinon

Abstract—This paper presents a whole-body robot control method for exploring and probing a given region of interest. The ergodic control formalism behind such an exploration behavior consists of matching the time-averaged statistics of a robot trajectory with the spatial statistics of the target distribution. Most existing ergodic control approaches assume the robots/sensors as individual point agents moving in space. We introduce an approach that decomposes the whole-body of a robotic manipulator into multiple kinematically constrained agents. Then, we generate control actions by calculating a consensus among the agents. To do so, we use an ergodic control formulation called heat equation-driven area coverage (HEDAC) and solve the diffusion using the non-stationary heat equation. Our approach extends HEDAC to applications where robots have multiple sensors on the whole-body (such as tactile skin) and use all sensors to optimally explore the given region. We show that our approach increases the exploration performance in terms of ergodicity and scales well to real-world problems. We compare our method in kinematic simulations with the state-of-the-art and demonstrate the applicability of an online exploration task with a 7-axis Franka Emika robot. Additional material available at <https://sites.google.com/view/w-ee-d/>

Index Terms—Optimization and Optimal Control, Whole-Body Motion Planning and Control, Ergodic Exploration

I. INTRODUCTION

VARIOUS exploration tasks require physical interaction to collect information due to contact requirements or sensory occlusion. Existing work in tactile exploration ranges from object shape reconstruction with tactile skins [5], and probing for stiffness mapping [4], to exploring wrench space of an articulated object [13] and exploring end-effector poses for insertion tasks [19]. Although it is possible to formulate these tasks as an optimization of information measures [21], setting up the objective is challenging and computation is expensive. Fortunately, for a subset of problems, one can formulate the autonomous exploration as a coverage of a region (target distribution).

Manuscript received: June, 20, 2023; Revised September, 20, 2023; Accepted October, 18, 2023.

This paper was recommended for publication by Editor L. Pallottino upon evaluation of the Associate Editor and Reviewers' comments. This work was supported by the State Secretariat for Education, Research and Innovation in Switzerland for participation in the European Commission's Horizon Europe Program through the INTELLIMAN project (<https://intelliman-project.eu/>), HORIZON-CL4-Digital-Emerging Grant 101070136) and the SES-TOSENSE project (<http://sestosenso.eu/>, HORIZON-CL4-Digital-Emerging Grant 101070310).

The authors are with the Idiap Research Institute, CH-1920 Martigny, Switzerland and also with the EPFL, 1015 Lausanne, Switzerland cem.bilaloglu@idiap.ch; tobias.loew@idiap.ch; sylvain.calinon@idiap.ch

Digital Object Identifier (DOI): see top of this page.




Fig. 1. Whole-body exploration of a target distribution using the last three links of the robot manipulator. In kinematic simulation, the exploration target is given in red. In the real-world experiment, the robot explores the cube region in dashed lines to localize a target object (a tennis ball whose location is unknown). Blue, turquoise, and purple spheres are the virtual agents constrained to the 5-th, 6-th, and 7-th links, respectively. The green and yellow arrows show the net virtual force and torque acting on each link's center of mass calculated by our agent weighing strategy. We further weigh the net wrenches acting on the active link to obtain the consensus control action for the robot.

Existing work in spatial coverage focuses on multi-agent systems with high-range sensors mounted on drones and ground robots (camera, LiDAR, time of flight sensor arrays [6]) for increasing coverage speed. Applications of these methods explore vast regions such as crop fields [9] or bridges [10]. However, robotic manipulators are advantageous for exploring small but complex target distributions with physical interaction. Nevertheless, existing methods do not consider the link footprints and the kinematic chain composing a robot manipulator's whole-body for exploration and lack the potential performance gains of increased sensor coverage. We argue that, as the accessibility of anthropomorphic robotic hands [20], and arms equipped with joint force-torque sensors and tactile skins [2, 8] increases, so will the need for whole-body exploration. Therefore, in this letter, we propose a whole-body ergodic exploration method for robotic manipulators that we summarized in Figure 1.

II. RELATED WORK

The common challenge in autonomous exploration is the inherent curse of dimensionality. Indeed, a naive exploration

based on random actions quickly becomes infeasible and instead requires intelligent strategies leveraging prior information [3]. In settings with minimal uncertainty about the task priors, the best approach is to use coverage path planning (CPP) [7, 22] or informative path planning (IPP) if the information is unevenly distributed [18, 24]. Nevertheless, planning methods require specifying the horizon in advance and boil down to solving a trajectory optimization problem that is intractable for the most general case.

Control methods are robust to uncertainty compared to planning approaches and can be used when the terminal time for planning is unknown [17]. In [16], Mathew and Mezić proposed to use ergodicity—a concept that originated from statistical physics providing a metric that measures the difference between the target distribution and time-averaged statistics of agent trajectories [1] for control. This particular method called *spectral multiscale coverage* (SMC), minimizes the ergodic metric by matching the Fourier series weights of the target distribution and the reconstructed distribution from the agent trajectories (coverage). Although the original formulation uses the Dirac delta function as the agent’s footprint to simplify the computation of coverage, Ayvali *et al.* [3] proposed using KL-divergence for arbitrary agent footprints, which also alleviates the need for approximating the target distribution by a finite number of basis functions. However, the techniques based on KL-divergence are sampling-based planners, and they were not tested in online control settings. In [19], Shetty *et al.* proposed an online ergodic exploration technique for peg-in-hole tasks. They extended the SMC algorithm by a low-rank approximation called tensor train factorization to scale the computation of Fourier series weights in the 6-D pose space describing the end-effector location.

A recent alternative to SMC-based methods is *heat equation-driven area coverage* (HEDAC) [11]. HEDAC encodes the target distribution as a virtual heat source and calculates the potential field resulting from diffusion, modeled as heat conduction in a uniform medium. Virtual heat conduction provides a model to smooth the gradient field and propagate information about unexplored regions to the agents. This technique is based on the diffusion (heat) equation, an extensively studied partial differential equation (PDE) that can be solved in various domains such as mesh surfaces point clouds and by using explicit or implicit time stepping schemes [15]. Additionally, it is possible to introduce internal and external domain boundaries where no heat conduction is allowed to encode exploration with embedded obstacle avoidance behavior. For instance, Ivić *et al.* adopted a finite element method to solve the heat equation on a planar domain with obstacles modeled as internal boundaries [12] using the Neumann boundary conditions. They later extended this approach to a three-dimensional planning setting [10] because the re-planning frequency was too low for online control. Existing work in HEDAC focuses on multi-agent systems, with the only exception being drozBot [14], a robot manipulator drawing artistic portraits. DrozBot solves the non-stationary heat equation for a single explicit timestep to have an artistic effect similar to doodling. The method considers only a single point (the tip of the pen) for the planar coverage and not the




Fig. 2. The HEDAC [11] method computes the potential field $u(\mathbf{x}, t)$ that guides the agents for ergodic exploration. The time-averaged coverage of the agent(s) $c(\mathbf{x}, t)$ at time t is subtracted from the target distribution $p(\mathbf{x})$ and positive values corresponding to unexplored regions are squared and used as the virtual heat source $s(\mathbf{x}, t)$. The diffusion (heat) equation (2) is then used for diffusing the potential field and propagating information of unexplored regions to the agents.

whole-body of the manipulator.

In this work, on the other hand, we consider the robot’s link footprints and the kinematic chain for increasing sensor coverage. For that purpose, we decompose the whole-body into kinematically constrained virtual agents. The primary challenge here is to formulate a locally consistent exploration behavior in time and space, combining agents for global exploration with the robot. We achieve this by solving the non-stationary diffusion equation to slow the diffusion to increase local exploration. Then, we combine local exploration and introduce weighting strategies on the agent and link levels to simplify reaching a consensus among the agents. Consequently, we present the first whole-body ergodic exploration method and the first three-dimensional control implementation of the HEDAC approach. We summarize our contributions as follows:

- Increased sensor coverage by modeling the whole-body as a collection of virtual exploration agents constrained to the robot’s kinematic chain;
- Formulating a locally consistent exploration behavior in time and space, combining agents for global exploration;
- Controlling the robot with consensus among virtual agents and links by the introduced weighting strategies

III. WHOLE-BODY CONTROL USING KINEMATICALLY CONSTRAINED VIRTUAL AGENTS

A. Smooth Potential Field Resulting from Diffusion

We extend the state-of-the-art ergodic control technique HEDAC to obtain the potential field guiding the exploration behavior. Original implementation [11] solves stationary ($\dot{u}(\mathbf{x}, t) = 0$) diffusion with virtual source $s(\mathbf{x}, t)$, sink $a(\mathbf{x}, t)$ and convective loss $u(\mathbf{x}, t)$ terms

$$\alpha \cdot \Delta u(\mathbf{x}, t) = \beta \cdot u(\mathbf{x}, t) + \gamma \cdot a(\mathbf{x}, t) - s(\mathbf{x}, t). \quad (1)$$

We summarize the HEDAC algorithm which uses the potential field $u(\mathbf{x}, t)$ resulting from integrating (1) in Figure 2. We

modify HEDAC to formulate a locally consistent exploration behavior in time and space, combining agents for global exploration. For that purpose, unlike the original formulation, we solve non-stationary ($\dot{u}(\mathbf{x}, t) \neq 0$) diffusion to increase local exploration by slowing down the diffusion

$$\dot{u}(\mathbf{x}, t) = \alpha \cdot \Delta u(\mathbf{x}, t) - u(\mathbf{x}, t) + s(\mathbf{x}, t). \quad (2)$$

We omit the sink term because we want to combine agents therefore inter-agent collision is not a concern. Moreover, we use a single parameter (α) for tuning the exploration behavior as proposed in [10]. We numerically integrate (2) using an explicit time-stepping scheme

$$u(\mathbf{x}, t + 1) = u(\mathbf{x}, t) + \dot{u}(\mathbf{x}, t) \cdot \delta t, \quad (3)$$

where the maximum stable timestep δt is a function of the diffusion rate α and given by Courant–Friedrichs–Lewy condition

$$\delta t \left(\sum_{i=1}^{N_d} \frac{\alpha_i}{\delta x_i} \right) \leq 1, \quad (4)$$

with the spatial discretization δx and the number of spatial dimensions N_d . We introduce the parameter N_k , for the number of integration steps we take using the maximum stepsize δt computed by (4). Note that choosing high values for N_k would mean integrating (2) until $u(\mathbf{x}, t + 1) - u(\mathbf{x}, t) \approx 0$ and would be equivalent to stationary diffusion.

B. Kinematically Constrained Virtual Agents

We define *virtual agents* as atomic particles that compose a rigid body and interact with the potential field $u(\mathbf{x}, t)$. Depending on the task, virtual agents abstract a sensor/tool used for physical interaction during exploration. Additionally, we use the term *whole-body* if the sensor/tool spans multiple bodies on different links of the manipulator. By construction, we can locate each agent in the potential field using the forward kinematics function f_{kin} of the robot

$$\mathbf{x}_{i,j} = \mathbf{f}_{\text{kin}}(\mathbf{q}, i, j) \quad \forall i = 1, \dots, N_j \quad \forall j = 1, \dots, M, \quad (5)$$

where N_j is the number of virtual agents on the j -th link, and M is the number of links composing the whole-body. $\mathbf{x}_{i,j}$ is the position of the i -th agent on the j -th link and \mathbf{q} is the vector of joint variables.

C. Active Agents and Local Weighting

We call the virtual agents that we use for computing coverage but not contributing to the control action as *passive*. We use the term *passive* because these agents explore regions indirectly as a secondary effect of the robot's primary goal. Still considering their coverage relieves the need to revisit those locations later. On the other hand, *active* agents, contribute to the control command of the robot with their local exploration goal in addition to coverage.

According to our model, the potential field exerts a fictitious force on each active agent based on the gradient of the potential field, and we multiply this force by an agent weight $w_{i,j}$

$$\mathbf{f}_{i,j} = w_{i,j} \nabla u(\mathbf{x}_{i,j}(t), t). \quad (6)$$




Fig. 3. Comparison of uniform and local temperature weighting. The green square is the exploration target and small arrows show the temperature gradient. Blue dots and arrows show active agents and the force exerted on each agent after weighting.

The naive method of computing the agent weights is to use a uniform weighting strategy, thus assigning equal importance to every agent. However, this is suboptimal since the significance of agents differs depending on their position in the potential field and the current state of the potential field itself. The value of the potential field at a given point encodes how much this particular region is underexplored. Accordingly, we embed this information by using the local value sensed by the agent as its weight. Hence, agents on the frontier of exploration (the ones closer to the underexplored regions) will have a higher weight than those on the overexplored regions. Thus, we set the normalized weight of the i -th agent on the j -th link as

$$w_{i,j} = \frac{\tilde{w}_{i,j}}{\sum_{i=1}^{N_j} \tilde{w}_{i,j}} \quad \text{with} \quad \tilde{w}_{i,j} = u(\mathbf{x}_{i,j}(t), t), \quad (7)$$

and call it *local weighting strategy*. Note that, the local weight is a function of the potential field, i.e. both space and time, and is therefore computed online. We show the difference between local and uniform weighting strategies in Figure 3.

D. Active Links and Manipulability Weighting

Similarly to active agents, we call a rigid body composed of active agents *active link* and compute the net force and moment acting on it by all the agents

$$\mathbf{f}_{\text{net}_j} = \sum_{i=1}^{N_j} \mathbf{f}_{i,j}, \quad \mathbf{m}_{\text{net}_j} = \sum_{i=1}^{N_j} \mathbf{r}_{i,j} \times \mathbf{f}_{i,j}, \quad (8)$$

where $\mathbf{r}_{i,j}$ is the displacement vector connecting the active agent and the j -th link's center of mass. We concatenate force and moment into a net wrench acting on the j -th link of the manipulator. For the simplest kinematic control strategy, we set the desired twist of the link $\mathbf{v}_{\text{des}_j}$ equal to the net wrench acting on the j -th link

$$\mathbf{v}_{\text{des}_j} = \begin{bmatrix} \mathbf{f}_{\text{net}_j}^\top & \mathbf{m}_{\text{net}_j}^\top \end{bmatrix}^\top \quad (9)$$

corresponding to having identity inertia.

Here the twist commands that we generate correspond to consensus among the agents because the gradient field exerts similar forces to neighboring agents (local consistency) since we set the local cooling term used for collision avoidance in HEDAC to zero and slowed down diffusion by solving the non-stationary diffusion equation. Moreover, the gradient (force) field resulting from spatial diffusion is smooth, and we perform first local exploration (moving based on the force field) and

then global exploration (by the propagation of information as diffusion).

We weight the contribution of the active links to the control command because each link has different *manipulability* [23] and volume ν , resulting in distinct volumetric coverage rates. For that purpose, we compute the link weights using the scalar manipulability index μ

$$w_j = \nu_j \mu_j, \quad (10)$$

$$\text{with } \mu_j = \sqrt{\det(\mathbf{J}_j(\mathbf{q})\mathbf{J}_j(\mathbf{q})^\top)}, \quad (11)$$

where $\mathbf{J}_j(\mathbf{q})$ is the Jacobian of the j -th active link computed at its center of mass.

E. Consensus Control for Whole-Body

Our setting consists of multiple weight-prioritized tasks, encoded as task velocities with corresponding Jacobians, and we would like to perform them in the least square optimal sense by exploiting the redundancy of our robot

$$\dot{\mathbf{q}}_{\text{des}} = \arg \min_{\dot{\mathbf{q}}} (\bar{\mathbf{v}}_{\text{des}} - \bar{\mathbf{J}}\dot{\mathbf{q}})^\top \mathbf{W} (\bar{\mathbf{v}}_{\text{des}} - \bar{\mathbf{J}}\dot{\mathbf{q}}) \quad (12)$$

$$\text{with } \mathbf{W} = \text{blockdiag}(\mathbf{I}w_1, \mathbf{I}w_2, \dots, \mathbf{I}w_m) \quad (13)$$

where $\bar{\mathbf{J}}$ is the vertical concatenation of the Jacobians, and $\bar{\mathbf{v}}_{\text{des}}$ is the vertical concatenation of the task velocities of active links. We enforce task priorities by using the weighted pseudoinverse and computing the desired joint velocities as

$$\dot{\mathbf{q}}_{\text{des}} = (\bar{\mathbf{J}}^\top \mathbf{W} \bar{\mathbf{J}})^{-1} \bar{\mathbf{J}}^\top \mathbf{W} \bar{\mathbf{v}}_{\text{des}}. \quad (14)$$

Next, we use desired joint velocity to either kinematically simulate the robot $\mathbf{q}_{t+1} = \mathbf{q}_t + \dot{\mathbf{q}}_{\text{des}} \cdot \Delta t$ or as a desired joint velocity to an impedance controller. Then, we clamp the desired joint positions as the simplest strategy to comply with joint limits. We give the full procedure for robot control in Algorithm 1.

Algorithm 1: Whole-Body Exploration

Input: target distribution $p(\mathbf{x})$, forward kinematics function $\mathbf{f}_{\text{kin}}(\mathbf{q}, i, j)$, initial joint configuration \mathbf{q}_0 , number of active links M , number of active agents N_j in j -th link

Output: control commands $\mathbf{u}(t)$

initialize agents and bodies using using (5) and \mathbf{q}_0
compute δt using (4)

while true do

 compute $c(\mathbf{x}, t)$ and $s(\mathbf{x}, t)$ as in [11]

for $k \leftarrow 1$ to N_k **do**

 integrate (2) using (3)

for all M bodies **do**

for all N_j agents **do**

 compute $\mathbf{x}_{i,j}$ and $\mathbf{f}_{i,j}$ using (5) and (6)

 compute $\mathbf{v}_{\text{des},j}$ and w_j using (9) and (11)

 compute $\dot{\mathbf{q}}_{\text{des}}$ using (14)

IV. WHOLE-BODY EXPLORATION

We give the videos for all the simulated and real-world experiments on the website.

A. Simulated Experiments

We performed kinematic simulations to measure the exploration performance. We used the normalized ergodicity over the target distribution as the exploration metric

$$\varepsilon = \frac{\|\max(e(\mathbf{x}, t), 0)\|_2}{\int_{\Omega} p(\mathbf{x}) d\mathbf{x}}, \quad (15)$$

where $p(\mathbf{x})$ is the target distribution, $e(\mathbf{x}, t)$ is the residual given by $p(\mathbf{x}) - c(\mathbf{x}, t)$ as given in Figure 2 and calculated as in [11]. This metric shows how good the time-averaged statistics of agent trajectories match the target distribution, and lower values indicate higher performance.

1) *Planar Experiments:* We first present the results of planar simulations showing the exploration performance qualitatively using images. We tested four different configurations: (i) *SMC*, (ii) *passive*, (iii) *active/stationary* using (1), and (iv) *active/non-stationary* ($N_k = 1$) using (2). SMC and passive configurations use a single active agent at the tip of the last link that we control by the state-of-the-art ergodic control methods SMC and HEDAC, respectively, where the passive configuration also considers the coverage of the last link indirectly for the control. We used active agents sampled on the last link equally spaced by the unit distance for the active configurations. We computed the metric (15) using the coverage of the last link for all of these configurations. We used a 75×75 grid for discretizing the exploration domain and 20 basis functions for the SMC method. We plotted the trajectories and the target distribution in Figure 4. Next, we repeated the same experiment for two different target distributions starting from 100 uniformly sampled initial joint configurations. We chose one diffuse and one fine-detailed target distribution to test the scenarios advantaging stationary/non-stationary diffusion, respectively. Although the exploration behavior would continue indefinitely, we stopped the simulation after 1000 timesteps. We calculated the mean and standard deviation of the normalized ergodic metric and plotted the results in Figure 5.

2) *Three Dimensional Experiments:* In the first 3D experiment, we measured the effect of using multiple active links on the exploration performance. We assumed no prior for the exploration target and used a cube discretized on a $50 \times 50 \times 50$ grid where each point corresponds to 1 cm as the target distribution. We placed the target in front of a 7-axis Franka Emika robot and sampled active agents on links 5, 6, and 7 using Poisson-disk sampling. Note that we used the same agent configuration for computing the coverage and the metric using (15), but we changed the number of active links used for computing robot control commands. Similarly to the planar experiments, we used 30 pre-sampled initial joint configurations and plotted the results in Figure 6. Secondly, we compared the exploration performance of the proposed method to that of using a search pattern. In this task, the robot explored the same target region using its whole-body




Fig. 4. Planar exploration using different configurations. The black shape is the target distribution, and the colored lines are agent trajectories, where the blue dashed lines correspond to the agent’s path at the tip of the last link. We show the configuration of the planar manipulator at equally spaced timesteps. Red is the start configuration, and the transparency decreases as time increases.

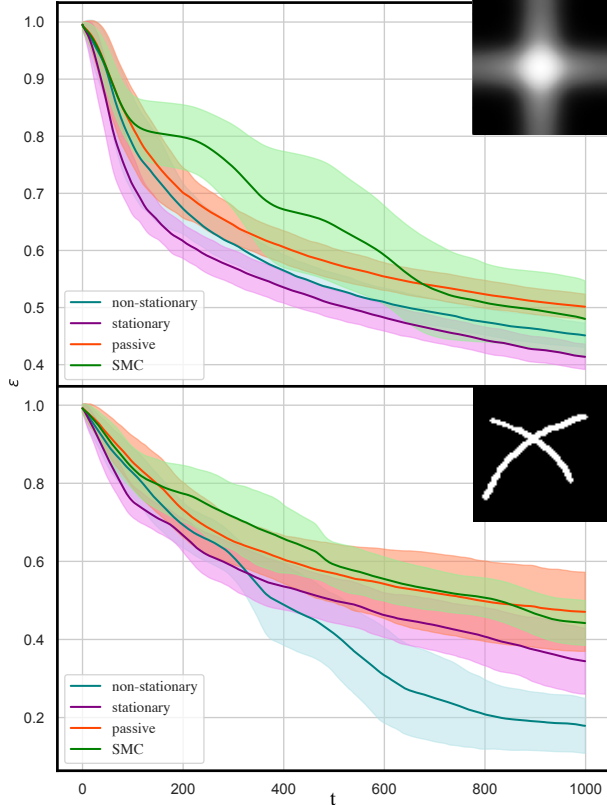


Fig. 5. Coverage performance given by the normalized ergodic metric for different virtual agent configurations. Target distributions for the coverage task are given on the top right.

composed of links 5, 6, and 7 until one of the links touched the sphere target placed at an unknown position. We repeated the experiment 30 times for randomly sampled target sphere positions but used fixed initial configurations corresponding to the end-effector poses 1, 2 given in Figure 7b. We recorded the time until the first contact and plotted the results in Figure 7a.

B. Real-world Experiment

For the real-world experiments, we used a 7-axis Franka Emika robot with the same object localization task performed in kinematic simulation. We placed a tennis ball inside the target distribution as the target object at an unknown location. We first recorded the trajectory resulting from the exploration behavior and then tracked the trajectory with a stiff impedance




Fig. 6. Coverage performance given by the normalized ergodic metric for different virtual agent configurations. The inset boxplot shows the performance in detail for each configuration at 500-th timestep. During the experiments, we set $N_k = 3$ as an empirically found moderate value and attained a control frequency of 35, 18, and 13 Hz for active agents on the last one, two, and three links respectively using a laptop processor.




Fig. 7. Coverage performance given by the timesteps required until the first contact of the whole-body with the target object (green sphere). a) The horizontal axis of the boxplot corresponds to our method and a planned search pattern starting from the initial end-effector poses 1 and 2. b) We show the initial poses 1, 2 and the direction of the search pattern is indicated by the ordering of the poses.




Fig. 8. Real-world experiment of the robot exploring the cube in dashed lines using its last three links, until either b) link 5, c) link 6 or d) link 7 contacts the target object.

controller to ensure that the robot safely contacted the tennis ball and did not hit the second robot or the stick. We ran the experiment until one of the links made contact with the target object, which we detected by using the joint torque sensors of the robot. We provide the experiment setup in Figure 8.

V. DISCUSSION

We investigated the effect of using active/passive agents and stationary/non-stationary diffusion on the exploration trajectory in Figure 4. We see SMC and passive configurations do not, and stationary diffusion configuration can not efficiently align with the target. On the other hand, in the non-stationary case, agents align with the target distribution for most of their trajectories (only get misaligned when rotating for re-alignment). This is expected since non-stationary diffusion increases local exploration and agent coordination, resulting in better alignment behavior. Next, we tested the time evolution of the ergodic metric in Figure 5. On top of the qualitative results of Figure 4, we can quantitatively argue active agents have higher performance than the state-of-the-art SMC method and the passive configuration. However, the performance difference between stationary/non-stationary diffusion depends on exploration time and the target distribution. This is because stationary diffusion ($N_k \gg 1$) smooths the details and coarsely explores diffuse targets but can not explore distal and fine-detailed regions. In contrast, non-stationary diffusion with $N_k = 1$ performs significantly better in distal and fine-grained regions but slightly worse for a diffuse target. For $N_k = (1, 10]$, we obtained results that interpolate the stationary/non-stationary diffusion results given in Figure 5. Accordingly, for the number of integration steps N_k , we recommend using a single step ($N_k = 1$) for concentrated targets, high values ($N_k \approx 10$) for diffuse targets, and low to medium values ($N_k \approx 3$) for generic cases.

In the first 3-D experiment, we see that using multiple links increases the exploration performance as depicted in Figure

6. However, this gain diminishes as we move to the links with less manipulability. Links closer to the base, including link 5, have negligible weights when computing the robot control command due to the normalization. Consequently, they practically become passive, contributing to coverage but not to the robot control commands. This partly explains why we observe marginal gains after link 6. Nevertheless, we think the primary reason is that the target distribution is small compared to the link size. This decreases the benefit of increasing the sensor footprint by using more links. Unfortunately, we can not make the target distribution larger due to the manipulator's joint limits and fixed base. Nonetheless, if the link and target size ratio were lower, the performance would increase with the sensor footprint. Two such examples are using a mobile manipulator where we can increase the target size or using the whole-body of a robot hand connected to an arm where we can decrease the link size relative to the target distribution.

In Figure 7, we compare the performance of the proposed method to using a planned search pattern. On the path from 1 to 2, links 5, 6, and 7 initially stay inside the target distribution, whereas on the path from 2 to 1, all links explore only at the end of the path. Accordingly, these two paths correspond to the best and worst-case scenarios for exploration. The results show that our approach is more robust to the initial robot configuration and can effectively leverage the whole-body as it performs better than the planned path of 1-2. Although here we used a uniform distribution in the form of a cube, ergodic exploration methods such as ours generalize to arbitrary target distributions, whereas pattern-based or coverage, informative path-planning approaches are limited to uniform distributions or simple geometries.

We showcased the method's applicability with a manipulator in a real-world experiment and presented the first 3D control implementation of the HEDAC approach. The control frequency of the method primarily depends on (i) integrating the diffusion using (3) and (ii) the number of active agents and links. Note that our choice of solving non-stationary diffusion with (2) and using explicit time-stepping results in higher computational efficiency for small N_k . In contrast, for the stationary diffusion equation (1) using implicit time-stepping reduces to sparse matrix multiplication with pre-inverted system matrix and is more efficient than the iterative solution. To reduce the effect of the number of active links and agents on computation, we decomposed links to multiple agents instead of using a single agent with the shape of the link. Because coverage computation using a single agent with the link shape can not be parallelized and is computationally expensive due to including more zero entries in the convolution kernel. In the end, integrating the diffusion equation is the main bottleneck in our implementation, and increasing the number of agents has minimal effect. Still, considering that the exploration performance gain is minimal and the link weight is almost zero for link 5, we think it is better to consider the last two links for the whole-body exploration if the robot base is fixed.

An alternative use case for the presented method is using it as a long-horizon planner anticipating the joint limits instead of a myopic controller. Although this would improve the

performance, it would be minimal because, thanks to the smooth potential field resulting from diffusion, it is unlikely to get stuck in a local minima due to joint limits. Still, we perceive the limited workspace due to the joint limits and the fixed base as the primary limitation to using whole-body exploration. Using a mobile manipulator would drastically increase whole-body exploration's performance because the link size target size ratio can be lower, and we would benefit from even less manipulable but large links. However, a limitation of the current approach in such a scenario would be the computational complexity of solving diffusion in a larger domain. Fortunately, unsupervised learning techniques for approximating diffusion, such as tensor train decomposition and physics-informed neural networks, seem promising.

Finally, real-world experiments showed we can perform whole-body tactile exploration using off-the-shelf torque-controlled robots without additional sensors. This result provides insight into alternative platforms equipped with joint torque sensing and searching for contacts, such as legged robots and multi-fingered robot hands. For instance, legged robots can use our method when exploring ground contacts, and robot hands can use it for grasping where each finger moves until contacting the object of interest. Although we can sense contact by joint torques, localizing it on the robot surface is challenging and requires whole-body tactile sensors. Accordingly, in the future, we plan to use the method with recursive target distribution updates and tactile sensors to reconstruct the physical properties of the environment that can only be measured through contact, such as deformability or friction.

VI. CONCLUSION

In this letter, we presented a robot control method for efficiently exploring a target distribution using a robotic manipulator's whole-body. Unlike existing approaches for multi-robot systems using independent agents, we used kinematically constrained agents on the links of a robotic manipulator to increase coverage. We introduced active agents and links together with weighting strategies considering the shape and kinematic chain of the whole-body for exploration. We combined the agents composing the whole-body for global exploration by formulating a locally consistent exploration behavior in time and space using non-stationary diffusion. Lastly, we measured the performance of our method in terms of ergodicity in kinematic simulations and demonstrated its applicability in physical scenarios using the 7-axis Franka Emika in an object localization task.

REFERENCES

- [1] I. Abraham, A. Prabhakar, and T. D. Murphey. "An Ergodic Measure for Active Learning From Equilibrium". In: *IEEE Transactions on Automation Science and Engineering* 18.3 (July 2021), pp. 917–931.
- [2] A. Albin, F. Grella, P. Maiolino, and G. Cannata. "Exploiting Distributed Tactile Sensors to Drive a Robot Arm Through Obstacles". In: *IEEE Robotics and Automation Letters* 6.3 (July 2021), pp. 4361–4368.
- [3] E. Ayvali, H. Salman, and H. Choset. "Ergodic Coverage in Constrained Environments Using Stochastic Trajectory Optimization". In: *Proc. IEEE/RSJ Intl Conf. on Intelligent Robots and Systems (IROS)*. Sept. 2017, pp. 5204–5210.
- [4] E. Ayvali, R. A. Srivatsan, L. Wang, R. Roy, N. Simaan, and H. Choset. "Using Bayesian Optimization to Guide Probing of a Flexible Environment for Simultaneous Registration and Stiffness Mapping". In: *Proc. IEEE Intl Conf. on Robotics and Automation (ICRA)*. May 2016, pp. 931–936.
- [5] M. Bauza, O. Canal, and A. Rodriguez. "Tactile Mapping and Localization from High-Resolution Tactile Imprints". In: *Proc. IEEE Intl Conf. on Robotics and Automation (ICRA)*. May 2019, pp. 3811–3817.
- [6] C. Bilaloğlu, M. Şahin, F. Arvin, E. Şahin, and A. E. Turgut. "A Novel Time-of-Flight Range and Bearing Sensor System for Micro Air Vehicle Swarms". In: *Swarm Intelligence*. Lecture Notes in Computer Science. 2022, pp. 248–256.
- [7] A. Bircher, M. Kamel, K. Alexis, H. Oleynikova, and R. Siegwart. "Receding Horizon Path Planning for 3D Exploration and Surface Inspection". In: *Autonomous Robots* 42.2 (Feb. 2018), pp. 291–306.
- [8] X. Fan, D. Lee, L. Jackel, R. Howard, D. Lee, and V. Isler. "Enabling Low-Cost Full Surface Tactile Skin for Human Robot Interaction". In: *IEEE Robotics and Automation Letters* 7.2 (Apr. 2022), pp. 1800–1807.
- [9] S. Ivić, A. Andrejčuk, and S. Družeta. "Autonomous Control for Multi-Agent Non-Uniform Spraying". In: *Applied Soft Computing* 80 (July 2019), pp. 742–760.
- [10] S. Ivić, B. Crnković, L. Grbčić, and L. Matleković. "Multi-UAV Trajectory Planning for 3D Visual Inspection of Complex Structures". In: *Automation in Construction* 147 (Mar. 2023), p. 104709.
- [11] S. Ivić, B. Crnković, and I. Mezić. "Ergodicity-Based Cooperative Multiagent Area Coverage via a Potential Field". In: *IEEE Transactions on Cybernetics* 47.8 (Aug. 2017), pp. 1983–1993.
- [12] S. Ivić, A. Sikirić, and B. Crnković. "Constrained Multi-Agent Ergodic Area Surveying Control Based on Finite Element Approximation of the Potential Field". In: *Engineering Applications of Artificial Intelligence* 116 (Nov. 2022), p. 105441.
- [13] T. Löw, C. Bilaloglu, and S. Calinon. "Efficient Representation for Learning Object-Centric Articulation Models Using Geometric Algebra". In: *Proc. IEEE Intl Conf. on Robotics and Automation (ICRA)*. Workshop on Effective Representations, Abstractions, and Priors for Robot Learning (RAP4Robots).
- [14] T. Löw, J. Maceiras, and S. Calinon. "drozBot: Using Ergodic Control to Draw Portraits". In: *IEEE Robotics and Automation Letters* 7.4 (Oct. 2022), pp. 11728–11734.
- [15] C. B. Macdonald and S. J. Ruuth. "The Implicit Closest Point Method for the Numerical Solution of Partial Differential Equations on Surfaces". In: *SIAM Journal on Scientific Computing* 31.6 (Jan. 2010), pp. 4330–4350.
- [16] G. Mathew and I. Mezić. "Spectral Multiscale Coverage: A Uniform Coverage Algorithm for Mobile Sensor Networks". In: *Proceedings of the 48th IEEE Conference on Decision and Control (CDC)*. Dec. 2009, pp. 7872–7877.
- [17] G. Mathew and I. Mezić. "Metrics for Ergodicity and Design of Ergodic Dynamics for Multi-Agent Systems". In: *Physica D: Nonlinear Phenomena* 240.4 (Feb. 2011), pp. 432–442.
- [18] M. Popović, T. Vidal-Calleja, G. Hitz, J. J. Chung, I. Sa, R. Siegwart, and J. Nieto. "An Informative Path Planning Framework for UAV-based Terrain Monitoring". In: *Autonomous Robots* 44.6 (July 2020), pp. 889–911.
- [19] S. Shetty, J. Silvério, and S. Calinon. "Ergodic Exploration Using Tensor Train: Applications in Insertion Tasks". In: *IEEE Transactions on Robotics* 38.2 (Apr. 2022), pp. 906–921.
- [20] O. Shorthose, A. Albin, L. He, and P. Maiolino. "Design of a 3D-Printed Soft Robotic Hand With Integrated Distributed Tactile Sensing". In: *IEEE Robotics and Automation Letters* 7.2 (Apr. 2022), pp. 3945–3952.
- [21] A. T. Taylor, T. A. Berrueta, and T. D. Murphey. "Active Learning in Robotics: A Review of Control Principles". In: *Mechatronics* 77 (Aug. 2021), p. 102576.
- [22] L. Yoder and S. Scherer. "Autonomous Exploration for Infrastructure Modeling with a Micro Aerial Vehicle". In: *Field and Service Robotics: Results of the 10th International Conference*. 2016, pp. 427–440.
- [23] T. Yoshikawa. "Manipulability of Robotic Mechanisms". In: *The International Journal of Robotics Research* 4.2 (June 1985), pp. 3–9.
- [24] H. Zhu, J. J. Chung, N. R. Lawrence, R. Siegwart, and J. Alonso-Mora. "Online Informative Path Planning for Active Information Gathering of a 3D Surface". In: *Proc. IEEE Intl Conf. on Robotics and Automation (ICRA)*. May 2021, pp. 1488–1494.