

Planar Orientation Control and Torque Maximization Using a Swarm With Global Inputs

Shiva Shahrokhi, Lillian Lin, and Aaron T. Becker

Abstract—This paper studies the torque applied by a large number of particles on a long aspect-ratio rod. The particles are all pushed in the same direction by a global signal. We calculate the force and torque generated by three canonical position distributions of a swarm: uniform, triangular, and normal. The model shows that for a pivoted rod the uniform distribution produces the maximum torque for small swarm standard deviations, but the normal distribution maximizes torque for large standard deviations. In simulation, we use these results to design PD controllers to orient rigid objects. We conclude showing experiments with up to 97 hardware robots to evaluate our theory in practice.

Note to Practitioners—Workspace clutter can prevent large steered particles from being able to manipulate objects and maneuver, while smaller particles can pass through this clutter. A small particle produces less force than a big particle, so to produce the same force, more are needed. Their small size makes onboard sensing and computation hard. Therefore they are often controlled by a shared control input. Manipulating objects with a swarm of particles actuated by a shared control input is a challenging task, but it is even more challenging when the object's final orientation needs to be set. Many applications including assembly and delivery require a specific orientation of the object.

Torque control with only one steered particle is easy: maximize torque by pushing on the object at a location as far from the pivot point as possible. However, a swarm of particles contributes force at different places on the object. This work studies how to maximize torque using a swarm of particles shaped in three canonical position distributions. The work is limited by assuming each particle touching the object transmits equal force, but hardware experiments validate the necessity to consider swarm distribution when applying torque. In future work, we will investigate how stochastic contacts between particles effects force and torque transmission, and examine control in 3D space.

Index Terms—Underactuated Robots, Torque Control, Swarm Control

I. INTRODUCTION

IN the future, large swarms of robots may be remotely guided through the human body, to cure disease, heal tissue, prevent infection, and assemble structures in parallel. For each application, large numbers of micro robots are required to deliver sufficient payloads, but the small size of these robots makes it difficult to perform onboard computation. Today, such robots are often controlled by a broadcast signal. Future implementations require control techniques that can reliably exploit large populations despite significant under-actuation.

This paper investigates maximizing torque applied by a large number of particles, hereafter called a *swarm*. The under-actuated swarm is steered by a shared signal such that the same force is applied to each particle. The robotic system is

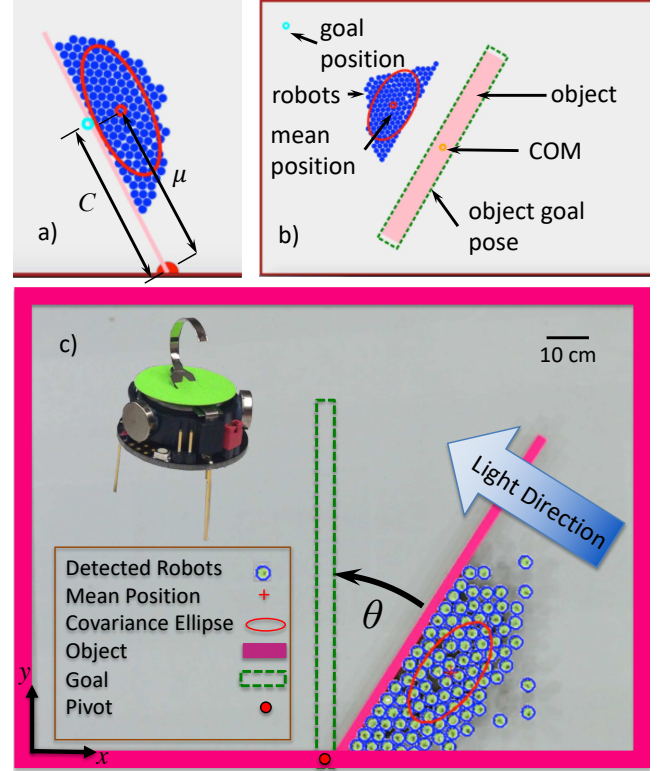


Fig. 1. Torque control of an object is essential for manipulation unless objects are homogeneous discs, especially when there are narrow passageways or the objects must be aligned, e.g. sensors and emitters. This paper calculates the optimal position for a swarm of particles to push to maximize torque production using a highly under-actuated system where all particles are controlled uniformly by the same input. (a) Simulation of particles exerting torque on a hinged “door”. (b) Orientation control of a free long rod. (c) Hardware robots applying torque to an object. See video attachment. A full resolution video is available at https://youtu.be/7Q5lu_ZFbxI.

comprised of the swarm, the shared control signal, and an external sensor that measures the swarm position. This paper examines analytically two representative aspects of planar swarm torque control: first, pushing a pivoted rod in Sec. IV-A, and second pushing a free body in Sec. IV-B. We show how implementing these results with PD controllers enables orientation control in simulation (Sec. V). We conclude with hardware experiments with centimeter-scale robots (Sec. VI).

The analysis in this paper is motivated by practical challenges when designing control techniques to orient objects with large aspect ratios in our preliminary conference paper [1]. The analytical results are new, and the hardware results have been extended.

With a single agent, torque control is straightforward: the agent simply maximizes the length of the moment arm to maximize torque. To make an agent push open a door, it should push on the edge furthest from the hinge. The optimal solution

for a swarm of particles is not straightforward because they cannot all push at one position.

Obtaining positions of each particle in the swarm is often hard, because remote sensing of tiny particles is difficult. However, though particles may be smaller than the minimum resolution of MRI, PET, and ultrasound, these sensing modalities can still return *aggregate* data in the form of an array of intensity values. From this data, some statistics of the swarm's position are easy to obtain such as mean position and variance [2]. This paper focuses on orientation control of a rectangular object by maximizing torque using statistics of the swarm's position distribution. Representative results are in Fig. 1.

II. RELATED WORK

Unlike *caging* manipulation, where robots form a rigid arrangement around an object [3], [4], our swarm of particles is unable to grasp the blocks they push, and so manipulation requires *nonprehensile manipulation* techniques, e.g. [5]–[7].

Robotic manipulation by pushing has a long and successful history [5], [8]–[10]. Key developments introduced the notion of stable pushes and a friction cone. A *stable push* is a pushing operation by a robot with a flat-plate pushing element in which the object does not change orientation relative to the pushing robot [5]. The *friction cone* is the set of vector directions a robot in contact with an object can push that object with a stable push. Stable pushes can be used as primitives in a rapidly-expanding random tree to form motion plans. A key difference is that our swarm of particles do not form a rigid structure and instead tend to flow around the object, with similarities to fluidic trapping [11], [12].

Though some particles self-aggregate, e.g. ferrous particles tend to clump in a magnetic field, many do not. The magnetotactic bacteria of [13], [14] and protists of [15] are directed by the orientation of the magnetic field and do not suffer from magnetic aggregation. This paper does not consider self-aggregation.

Controlling the *shape*, or relative positions, of a swarm of particles is a key ability for a myriad of applications. Correspondingly, it has been studied from a control-theoretic perspective in both centralized, e.g. virtual leaders in [16], and decentralized approaches, e.g. decentralized control-Lyapunov function controllers in [17]. Most approaches assume a level of intelligence and autonomy in the individual robots that exceeds the capabilities of current micro- and nano-robots [13], [18], [19]. Instead, this paper focuses on a centralized technique that applies the same control input to each member of the swarm, as in [20].

III. TORQUE CONTROL

We derive inspiration from recent work on pulling with a swarm [21]. The contribution of this work is to map swarm distributions to torque production. To change the output torque τ of a single steered particle, we can choose the direction and magnitude of the force applied F , and the moment arm from the object's center of mass (COM) O to the point of contact P . We define a coordinate frame rooted at the COM, with

the x -axis parallel to the object's longest axis. The resulting torque is:

$$\tau = F \times (P - O). \quad (1)$$

We assume that the forces produced by individual particles add linearly. As [21] indicates, this is often a simplification of the true dynamics. The swarm version of (1) is the summation of the forces contributed by n individual particles:

$$\tau_{\text{total}} = \sum_{i=1}^n \rho_i F_i \times (P_i - O), \quad (2)$$

$$F_{\text{total}} = \sum_{i=1}^n \rho_i F_i. \quad (3)$$

Here F_i is the force that the i^{th} particle applies. Not all particles are in contact with the object. The indicator variable ρ_i is 1 if the particle is in direct contact with the object or touching a chain of particles where at least one particle is in contact with the object, otherwise $\rho_i = 0$. The moment arm is the particle's position P_i to the object's COM $O = [O_x, O_y]^T$. If all particles are identical and the control input is uniform, the force is equivalent for every particle and so F_i is a constant for all i .

This analysis assumes perfect transmission of force from each particle in contact with the object. The swarm, when steered toward an object, begins interacting with the object at different times. The number of particles touching the object as a function of time is difficult to predict and often impossible to directly measure. Stochastic effects make exact long-term prediction challenging. Even when it is possible to predict which particles will hit the object first, as particles interact with the object, the swarm's configuration changes. The challenge is not only limited to swarm-object interaction, but also to swarm-swarm interactions when the swarm self-collides or is split into multiple components. As a result, the instantaneous force the swarm will exert on the object is not easy to predict.

However, as explained by the central limit theorem, the time-averaged configuration and the time-averaged forces are predictable. This claim is validated by hardware experiments in Section VI.

IV. CALCULATING TORQUE FROM SWARM DISTRIBUTION

Assume a swarm of particles are pushing on a rigid object. Let the marginal distribution of the swarm along x have probability density $p(x)$, where x is defined as perpendicular to the object's long axis. This section considers three canonical probability distributions: uniform, triangular, and normal, all parameterized by mean μ and standard deviation σ . They are plotted in Fig. 2 and described by

$$p_u(x) = \begin{cases} \frac{1}{2\sqrt{3}\sigma}, & \text{for } \mu - \sqrt{3}\sigma \leq x \leq \mu + \sqrt{3}\sigma \\ 0, & \text{otherwise} \end{cases}, \quad (4)$$

$$p_t(x) = \begin{cases} \frac{x - \mu + \sqrt{6}\sigma}{6\sigma^2}, & \text{for } \mu - \sqrt{6}\sigma \leq x \leq \mu \\ \frac{-x + \mu + \sqrt{6}\sigma}{6\sigma^2}, & \text{for } \mu < x \leq \mu + \sqrt{6}\sigma \\ 0, & \text{otherwise} \end{cases}, \quad (5)$$

$$p_n(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-(x-\mu)^2/2\sigma^2}. \quad (6)$$

Distribution	The μ location to push that maximizes torque		Maximum possible torque	
Pivoted Uniform	$\mu_{p_{u_{\max}}} =$	$\begin{cases} 1 - \sqrt{3}\sigma & \text{for } \sigma < \frac{1}{2\sqrt{3}} \\ 1 - \sqrt{3}\sigma < \mu < \sqrt{3}\sigma & \text{for } \sigma \geq \frac{1}{2\sqrt{3}} \end{cases}$	$\tau_{p_{u_{\max}}} =$	$\begin{cases} 1 - \sqrt{3}\sigma & \text{for } \sigma < \frac{1}{2\sqrt{3}} \\ \frac{1}{4\sqrt{3}\sigma} & \text{for } \sigma \geq \frac{1}{2\sqrt{3}} \end{cases}$
Pivoted Triangular	$\mu_{p_{t_{\max}}} =$	$\begin{cases} \sqrt{12\sigma^2 + 1} - \sqrt{6}\sigma & \text{for } \sigma < \frac{1}{2\sqrt{3}} \\ \frac{\sqrt{2}}{2} & \text{for } \sigma \geq \frac{1}{2\sqrt{3}} \end{cases}$	$\tau_{p_{t_{\max}}} =$	$\begin{cases} \frac{(1+12\sigma^2)^{\frac{3}{2}} - 1}{18\sigma^2} - \sqrt{6}\sigma & \text{for } \sigma < \frac{1}{2\sqrt{3}} \\ \frac{\sqrt{2}-2+3\sqrt{6}\sigma}{36\sigma^2} & \text{for } \sigma \geq \frac{1}{2\sqrt{3}} \end{cases}$
Free Uniform	$\mu_{u_{f_{\max}}} =$	$\begin{cases} 1 - \sqrt{3}\sigma & \text{for } \sigma < \frac{1}{2\sqrt{3}} \\ \sqrt{3}\sigma & \text{for } \sigma \geq \frac{1}{2\sqrt{3}} \end{cases}$	$\tau_{u_{f_{\max}}} =$	$\begin{cases} 1 - \sqrt{3}\sigma & \text{for } \sigma < \frac{1}{2\sqrt{3}} \\ \frac{1}{4\sqrt{3}\sigma} & \text{for } \sigma \geq \frac{1}{2\sqrt{3}} \end{cases}$
Free Triangular	$\mu_{t_{f_{\max}}} =$	$\begin{cases} \sqrt{12\sigma^2 + 1} - \sqrt{6}\sigma & \text{for } \sigma < \frac{2}{\sqrt{6}} \\ 1 < \mu < \sqrt{12\sigma^2 + 1} - \sqrt{6}\sigma & \text{for } \sigma \geq \frac{2}{\sqrt{6}} \end{cases}$	$\tau_{t_{f_{\max}}} =$	$\begin{cases} \frac{(1+12\sigma^2)^{\frac{3}{2}} - 1}{18\sigma^2} - \sqrt{6}\sigma & \text{for } \sigma < \frac{2}{\sqrt{6}} \\ \frac{1}{9\sigma^2} & \text{for } \sigma \geq \frac{2}{\sqrt{6}} \end{cases}$

TABLE I
MAIN RESULTS FROM SECTION IV FOR MAXIMIZING TORQUE WITH THREE COMMON DISTRIBUTIONS.

The next section examines where to steer the mean of the probability distribution to maximize torque. We discuss two problems: pivoted object torque and free object torque. All the results are summarized in Table I, and the calculations are included in the Mathematica file in the multimedia attachment.

A. Pivoted object torque

In this problem, the torque applied to a rod of length 1 pivoted at 0 when $\theta = 0$ is

$$\tau_p = \int_0^1 x p(x) dx. \quad (7)$$

a) Defining torque for distributions

First, for simplicity of the following derivations, the lower bound l and upper bound u are defined for the uniform distribution as

$$l = \max(0, \mu - \sqrt{3}\sigma), u = \min(1, \mu + \sqrt{3}\sigma). \quad (8)$$

Torque by a uniformly distributed swarm is

$$\tau_{u_p} = \frac{u^2 - l^2}{4\sqrt{3}\sigma}. \quad (9)$$

To simplify derivations for the triangular distribution, we define the bounds of integration as

$$\begin{aligned} l_1 &= \max(0, \mu - \sqrt{6}\sigma), & l_2 &= \max(0, \mu), \\ u_1 &= \min(1, \mu), & u_2 &= \min(1, \mu + \sqrt{6}\sigma). \end{aligned}$$

Torque by a triangularly distributed swarm is

$$\tau_{t_p} = \frac{\tau_l + \tau_r}{36\sigma^2}, \quad (10)$$

where τ_l and τ_r are defined as

$$\begin{aligned} \tau_l &= l_1^2 \left(2l_1 - 3(\mu - \sqrt{6}\sigma) \right) + u_1^2 \left(2u_1 - 3(\mu - \sqrt{6}\sigma) \right), \\ \tau_r &= l_2^2 \left(2l_2 - 3(\mu + \sqrt{6}\sigma) \right) - u_2^2 \left(2u_2 - 3(\mu + \sqrt{6}\sigma) \right). \end{aligned}$$

Torque by a normally distributed swarm is

$$\begin{aligned} \tau_{n_p} &= \frac{\left(e^{\frac{\mu^2}{2\sigma^2}} - e^{\frac{(1-\mu)^2}{2\sigma^2}} \right) \sigma}{\sqrt{2\pi}} \\ &\quad + \frac{\mu}{2} \left(\operatorname{erf} \left(\frac{1-\mu}{\sqrt{2}\sigma} \right) + \operatorname{erf} \left(\frac{\mu}{\sqrt{2}\sigma} \right) \right). \end{aligned} \quad (11)$$

The $\operatorname{erf}(\cdot)$ is the error function, defined as

$$\operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt. \quad (12)$$

Torque is plotted as a function of μ for representative σ values in Fig. 3.

b) Maximum torque for distributions

For a uniformly distributed swarm, the mean position that maximizes torque on a pivoted object is found by setting the first derivative of (9) to zero and solving, and is

$$\mu_{u_{p_{\max}}} = \begin{cases} 1 - \sqrt{3}\sigma & \text{for } \sigma < \frac{1}{2\sqrt{3}} \\ 1 - \sqrt{3}\sigma < \mu < \sqrt{3}\sigma & \text{for } \sigma \geq \frac{1}{2\sqrt{3}} \end{cases}. \quad (13)$$

As soon as the width of the uniformly distributed swarm is longer than the rod ($\sigma > \frac{1}{2\sqrt{3}}$), there exists a range of optimal solutions: any μ such that the swarm covers the rod from 0 to 1. The resulting torque is

$$\tau_{u_{p_{\max}}} = \begin{cases} 1 - \sqrt{3}\sigma & \text{for } \sigma < \frac{1}{2\sqrt{3}} \\ \frac{1}{4\sqrt{3}\sigma} & \text{for } \sigma \geq \frac{1}{2\sqrt{3}} \end{cases}. \quad (14)$$

For a triangularly distributed swarm, the mean position that maximizes torque is found by the first derivative test on (10) and is

$$\mu_{t_{p_{\max}}} = \begin{cases} \sqrt{12\sigma^2 + 1} - \sqrt{6}\sigma & \text{for } \sigma < \frac{1}{2\sqrt{3}} \\ \frac{\sqrt{2}}{2} & \text{for } \sigma \geq \frac{1}{2\sqrt{3}} \end{cases}. \quad (15)$$

Thus the torque is

$$\tau_{t_{p_{\max}}} = \begin{cases} \frac{(1+12\sigma^2)^{\frac{3}{2}} - 1}{18\sigma^2} - \sqrt{6}\sigma & \text{for } \sigma < \frac{1}{2\sqrt{3}} \\ \frac{\sqrt{2}-2+3\sqrt{6}\sigma}{36\sigma^2} & \text{for } \sigma \geq \frac{1}{2\sqrt{3}} \end{cases}. \quad (16)$$

For all distributions, the μ that maximizes torque if $\sigma = 0$ is 1. The uniform distribution maximizes torque for $\sigma \lesssim 0.2403$, and the normal distribution maximizes torque for larger σ values. For the τ_p case, the optimal μ moves in the $-x$ direction, and for some distributions approaches a limit as σ increases. This limit is $\mu = \frac{\sqrt{2}}{2}$ for a triangularly distributed swarm. In a normal distribution, the first derivative of (11) is

$$\frac{d\tau}{d\mu} = \frac{1}{2} \left(\operatorname{erf} \left(\frac{1-\mu}{\sqrt{2}\sigma} \right) + \operatorname{erf} \left(\frac{\mu}{\sqrt{2}\sigma} \right) \right) - \frac{e^{-\frac{(\mu-1)^2}{2\sigma^2}}}{\sqrt{2\pi}\sigma}. \quad (17)$$




Fig. 2. Three distributions are examined in this work: uniform, triangular, and normal. Each is plotted above with $\mu = 1$ for representative σ values.




Fig. 3. Torque on a pivoted rod as a function of mean position μ . Mean position is the pushing location along a rod extending from 0 to 1. For all distributions, pushing at $\mu = 1$ is not optimal unless $\sigma = 0$ or the swarm is uniformly distributed with $\sigma > \frac{1}{2\sqrt{3}}$. Black lines show where μ maximizes torque in the limit as σ grows.




Fig. 4. Optimal location to push and maximum torque plots for a pivoted object of length 1, pivoted at 0. Generating code is in the attachment.

The limit as $\sigma \rightarrow \infty$ is $\frac{2-3\mu}{\sigma^3 6\sqrt{2\pi}}$ which goes to zero as the swarm variance increases. However, the zero of the numerator is $\mu = \frac{2}{3}$, which means as σ increases the best position to push the rod asymptotically reaches $\mu = \frac{2}{3}$.

The central limit theorem explains why systems where all the particles have independent noise form a normal distribution. However, optimizing torque for a normal distribution involves two error functions (erf) and the first derivative test has no closed-form solution. Instead, we have solved all the equations for the triangular distribution. The numerical studies illustrated in Fig. 4 b and c show the triangular and normal distributions perform similarly.

B. Free object torque

In this problem, the torque applied to a free rod of length 2 from $[-1, 0]$ to $[1, 0]$, is

$$\tau_f = \int_{-1}^1 x p(x) dx. \quad (18)$$

a) Defining torque to a free rod for distributions

To calculate free object torque for a uniformly distributed swarm, for simplicity of derivations, l_f and u_f are defined as

$$l_f = \max(-1, \mu - \sqrt{3}\sigma), \quad u_f = \min(1, \mu + \sqrt{3}\sigma).$$

The torque applied by a uniformly distributed swarm is

$$\tau_{u_f} = \frac{u_f^2 - l_f^2}{4\sqrt{3}\sigma}. \quad (19)$$

To simplify derivations for a triangularly distributed swarm, we define bounds of integration as:

$$\begin{aligned} l_{f1} &= \max(-1, \mu - \sqrt{6}\sigma), & l_{f2} &= \max(-1, \mu), \\ u_{f1} &= \min(1, \mu), & u_{f2} &= \min(1, \mu + \sqrt{6}\sigma). \end{aligned}$$

The torque applied by a triangularly distributed swarm is

$$\tau_{t_f} = \frac{\tau_{f1} + \tau_{f2}}{36\sigma^2}, \quad (20)$$




Fig. 5. Torque on a free rod as a function of mean position μ . Mean position is the pushing location along a rod extending from -1 to 1. For all distributions, pushing at $\mu = 1$ is not optimal unless $\sigma = 0$.




Fig. 6. Optimal location to push and maximum torque plots for a free object of length 2, located from $x = -1$ to 1.

where τ_{f_l} and τ_{f_r} are defined as:

$$\begin{aligned}\tau_{f_l} &= u_{f_1}^2 (2u_{f_1} - 3(\mu - \sqrt{6}\sigma)) - l_{f_1}^2 (2l_{f_1} - 3(\mu - \sqrt{6}\sigma)), \\ \tau_{f_r} &= -u_{f_2}^2 (2u_{f_2} - 3(\mu + \sqrt{6}\sigma)) + l_{f_2}^2 (2l_{f_2} - 3(\mu + \sqrt{6}\sigma)).\end{aligned}$$

The torque applied by a normally distributed swarm is

$$\begin{aligned}\tau_{n_f} &= \frac{\left(e^{-\frac{(\mu+1)^2}{2\sigma^2}} - e^{-\frac{(\mu-1)^2}{2\sigma^2}} \right) \sigma}{\sqrt{2\pi}} \\ &+ \frac{1}{2}\mu \left(\operatorname{erf}\left(\frac{1-\mu}{\sqrt{2}\sigma}\right) + \operatorname{erf}\left(\frac{1+\mu}{\sqrt{2}\sigma}\right) \right).\end{aligned}\quad (21)$$

Torque is plotted as a function of μ for representative σ values in Fig. 5.

b) Maximum torque for distributions

Maximum torque of a uniformly distributed swarm to a free object has the same upper solution as with a pivoted object:

$$\mu_{u_{f_{\max}}} = \begin{cases} 1 - \sqrt{3}\sigma & \text{for } \sigma < \frac{1}{2\sqrt{3}} \\ \sqrt{3}\sigma & \text{for } \sigma \geq \frac{1}{2\sqrt{3}} \end{cases}.\quad (22)$$

The major difference is the existence of only one σ value solution for torque on a free object. The left bound of the distribution must never be less than 0, or torque applied left of the center will cancel torque to the right of the center. The maximum torque produced is given by (14).

For a triangularly distributed swarm the optimal mean is

$$\mu_{t_{f_{\max}}} = \begin{cases} \sqrt{12\sigma^2 + 1} - \sqrt{6}\sigma & \text{for } \sigma < \frac{2}{\sqrt{6}} \\ 1 < \mu < \sqrt{12\sigma^2 + 1} - \sqrt{6}\sigma & \text{for } \sigma \geq \frac{2}{\sqrt{6}} \end{cases}.\quad (23)$$

As with the uniform distribution in (13) the optimal solution for triangular distributions with $\sigma > \frac{2}{\sqrt{6}}$ has a range of solutions. Setting $\mu = 1$ maximizes force for the optimal

torque. However, moving the mean right reduces the negative torque applied by particles to the left of 0, but this gain is canceled by a corresponding reduction in positive torque.

Therefore, the maximum torque is

$$\tau_{t_{f_{\max}}} = \begin{cases} \frac{(1+12\sigma^2)^{\frac{3}{2}} - 1}{18\sigma^2} - \sqrt{6}\sigma & \text{for } \sigma < \frac{2}{\sqrt{6}} \\ \frac{1}{9\sigma^2} & \text{for } \sigma \geq \frac{2}{\sqrt{6}} \end{cases}.\quad (24)$$

Again, for a normally distributed swarm, we have numerically found the maximum torque and optimal pushing location μ . The results are closely approximated by the upper bound solution for the triangular system. The results are plotted in Fig. 6. The normal distribution maximizes torque for $0.2405 \lesssim \sigma \lesssim 0.8793$, and the uniform distribution has the maximum torque for other values.

V. SIMULATION

This section examines three challenges for torque control of an object, arranged in increasing difficulty. Each task uses a PD controller that regulates the swarm's mean position, as in [22]. Although the controller only regulates the swarm's mean and not its distribution, simulations and hardware experiments confirm that as long as the swarm does not partially pass the object, the resulting distribution is approximately normal. The control input is the global force applied to each particle:

$$\begin{aligned}u_x &= K_p(G_x - \bar{x}) + K_d(0 - \bar{v}_x), \\ u_y &= K_p(G_y - \bar{y}) + K_d(0 - \bar{v}_y),\end{aligned}\quad (25)$$

where K_p is the proportional gain, and K_d is the derivative gain. The swarm's average position is $[\bar{x}, \bar{y}]^T$ and mean velocity is $[\bar{v}_x, \bar{v}_y]^T$. Each task uses a different algorithm to select the swarm's goal position $[G_x, G_y]^T$. The derivative gain K_d limits overshoot.

a) *Maximizing torque on a pivoted body:* An object with a pivot point can rotate, but not translate. A door is an example. If there was only one particle touching the object, that particle should push at the point which maximizes the moment arm, at the extreme end of the object furthest from the pivot point. The optimal pushing location provides the maximum force, because it maximizes the distance $\|P - O\|$ in (1). However, given a swarm of particles, maximizing $\|P - O\|$ is no longer the optimal solution. If the mean of the swarm hits the object at the extreme edge, half of the particles will miss the object and the swarm will be split. Because few particles remain, the force is significantly decreased and torque is not maximized. Unless a separate *regathering* step is added, particles that pass the rod do not contribute any force or torque. For the simulations in this subsection, the swarm mean and variance (μ, σ) are only calculated over the particles that have not passed the rod. The key parameter of interest for a pivoted rod of length L is C , the position along the rod where the mean of the swarm will push. The goal position in (25) is set to:

$$\begin{aligned} G_x &= O_x + C \cos(O_\theta), \\ G_y &= O_y + C \sin(O_\theta), \end{aligned} \quad (26)$$

where $[O_x, O_y]$ is the pivot point, and O_θ is the orientation of the object's major axis, measured from the world x -axis. Figure 7 illustrates how different values of C result in different rates of turning. These simulations tested $C = \tau_{t_{p_{\max}}}$ calculated by (15) and $C = \{1/4, 1/2, 2/3, 1\}L$. The fastest turning rates occurred with the dynamic goal location using $\tau_{t_{p_{\max}}}$. We have used $\tau_{t_{p_{\max}}}$ using the triangular distribution because the results for triangular and normal distributions are approximately the same.

b) *Orientation of an object:* These simulations used a uniform density rectangle as the object. This object was $30\times$ larger than the particles. Using the pure torque control discussed in the previous paragraph, the orientation of the object can be controlled by applying force. The rectangular object is not pivoted, so it moves in addition to rotating. The swarm may split into multiple components if some particles move past the object, so we use the hysteresis variance control from [22] to gather the swarm when its variance grows larger than a threshold value. This variance control steers the swarm mean toward a nearby region with boundaries that form a corner. If the variance is less than the variance threshold, the controller chooses a goal position to regulate the orientation of the object:

$$\begin{aligned} G_x &= O_x + K_{\text{orient}} \text{AngErr}(O_\theta, G_\theta) \cos(O_\theta), \\ G_y &= O_y + K_{\text{orient}} \text{AngErr}(O_\theta, G_\theta) \sin(O_\theta). \end{aligned} \quad (27)$$

Here K_{orient} is a positive gain on the control input, G_θ is the goal orientation, $[O_x, O_y]$ is the object's center of mass, and $\text{AngErr}(A, B) = \text{ArcTan}(\cos(A) - \cos(B), \sin(A) - \sin(B))$.

Fig. 8 illustrates this controller with different starting positions. When the plot traces are constant the swarm is no longer pushing the object and instead is being regathered in a corner of the workspace.

c) *Straight translation while regulating object orientation:* When the total force is applied perpendicular to the

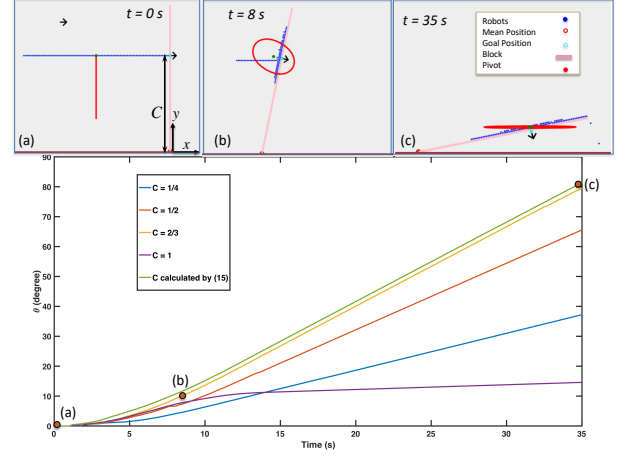


Fig. 7. Simulation results from a swarm applying force to a pivoted rod (hinged door). The swarm mean is steered toward a point C units along the object from the pivot point. Simulation used 100 particles of diameter 0.08 m with a dynamic standard deviation and an object length of 12 m.

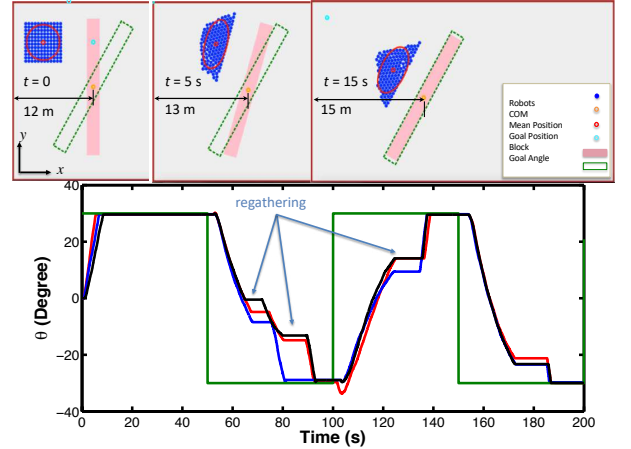


Fig. 8. Plot demonstrating orientation control of a rectangular object. The green line is the goal orientation. Other lines show results with three different initial average y -coordinates of the swarm. When the plot traces are constant the swarm is no longer pushing the object and instead is being regathered in a corner of the workspace until the variance is below a desired threshold.

object and in line with the center of mass, according to (1) there will be no torque. The following goal position for the mean position of the swarm regulates the object's orientation using Δ_θ for proportional feedback to determine where to apply force where $\Delta_\theta = \text{AngErr}(G_\theta, O_\theta)$ is the difference between the goal angle and the current object angle. K_τ is a constant and is tuned manually to 10. (O_x, O_y) is the position of the object's COM.

$$\begin{aligned} G_x &= O_x \\ G_y &= K_\tau \Delta_\theta + O_y \end{aligned} \quad (28)$$

Fig. 9 shows how Δ_θ converges to zero with different initial configurations of the swarm. When the swarm is above or below the object the swarm applies a torque to the object.

VI. EXPERIMENTS

To demonstrate the analytical results from Section IV hold, we performed experiments using centimeter-scale hardware robots called *Kilobots*. These allow us to emulate a variety

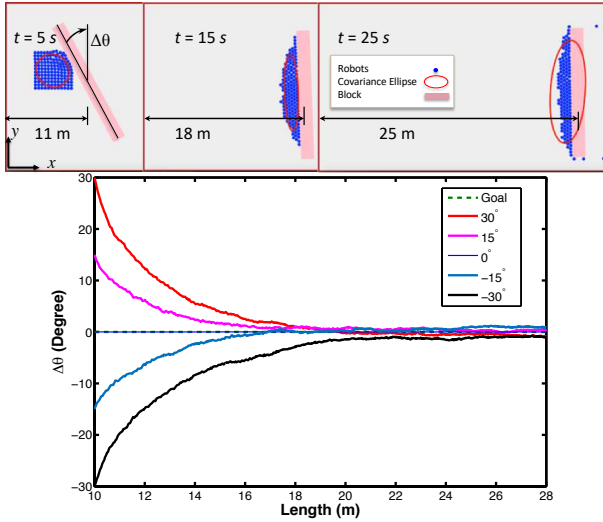


Fig. 9. In this task, the swarm pushed the object in the $+x$ direction while trying to regulate the orientation to $G_\theta = 0^\circ$. The swarm can push the object without changing its orientation only if it pushes along a line intersecting the COM of the object. A feedback control law regulates the object's orientation.

of dynamics, while enabling a high degree of control over robot function, the environment, and data collection.

The Kilobot, from [23], [24], is a low-cost robot designed for testing collective algorithms with large numbers of robots. It is available as open-source hardware or can be purchased ready-made [25]. Each robot is approximately 3 cm in diameter, 3 cm tall, and uses two vibration motors to move on a flat surface at speeds up to 1 cm/s. Each robot has one ambient light sensor that is used to implement *phototaxis*, moving towards a light source.

These experiments used up to $n = 97$ Kilobots, a glass-covered $1.5 \text{ m} \times 1.2 \text{ m}$ whiteboard as the workspace, and 30 W and 50 W LED floodlights arranged 1.5 m above the plane of the table in a 6 m square centered on the workspace. Above the table, an overhead machine vision system tracked the position of the swarm. For the uniform distribution experiments, we designed an object to rigidly hold the swarm, and programmed the robots to go straight.

A. Pushing a pivoted object

Figure 10 shows snapshots of experiments using a pivoted object with uniformly distributed swarms that have the same size and standard deviation, σ , but different mean position, μ . The object was designed such that the Kilobots keep the starting uniform distribution with $\sigma = 0.05L$, where $L = 1.42 \text{ m}$. Kilobots were programmed to go straight so that they always apply force perpendicular to the object. The results of different mean positions are shown in Fig. 11 for four trials at $\mu = [0.5, 0.66, 0.8]L$. Mean positions near the optimal solution of (13) increase torque.

The experiments from Section V.a were manually demonstrated using this physical swarm. Fig. 12 illustrates an experiment showing pure torque control with a swarm of robots. In this figure a large-aspect-ratio rectangle ($91 \times 2 \text{ cm}$, colored pink in the image) was hinged to one side of the table. Like a door, this object could only be moved around this hinge pivot. Two trials were performed. In each trial the swarm

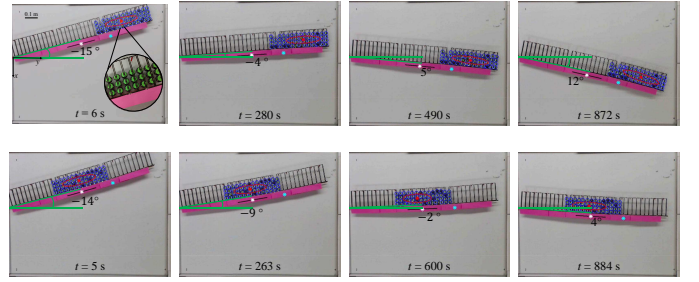


Fig. 10. Snapshots showing the effect of pushing a pivoted rectangular object at different distances from the pivot point. The top row shows frames with a mean position $\mu = 3/4$ of the object length. The bottom row show frames with swarm mean position $\mu = 1/2$.




Fig. 11. Pivoted object orientation using uniformly distributed Kilobots with $\sigma = 0.05L = 0.071 \text{ m}$, but different mean positions. Mean positions are normalized along a rod of length $L = 1.42 \text{ m}$.

was initialized in the lower right side of the table, and then commanded to push the object with the mean position of the swarm directed toward a point distance C from the pivot point. Data was recorded for 150 seconds. In the first trial, $C = L$, so the robots were commanded to push the door at the extreme edge of the door from the pivot. In the second trial $C = 1/2L$, so the swarm pushed the object in the center of the rectangle. As discussed in Section V, the robots spread when commanded to push the object at the extreme end, and half of the robots flowed past the end of the rectangle without engaging the rectangle. This illustrates a key difference between robotic swarms and a single pusher robot. The swarm exerts the most torque when (2) is maximized. (2) is maximized when the majority of the swarm engages the object. For this reason in Fig. 12, the trial in the second row of screenshots moves the door further than the trial in the first row.




Fig. 12. Snapshots showing the effect of pushing a pivoted rectangular object at different distances from the pivot point. The top row of snapshots illustrate the swarm pushing at the end of the object. The bottom row illustrates that when the swarm pushes at the middle of the object, the force provided by the swarm remains constant.




Fig. 13. Snapshots showing orientation control of a free object using 97 hardware robots that all receive the same control input. Light direction is the global control input and the robots are programmed to move toward the brightest light in the environment.

B. Orientation control of a free object

Figure 13 shows snapshots of orientation control of a free rod using 97 Kilobots. Kilobots were programmed to move toward the brightest light in the room. Therefore, the light is the global input. The goal orientation of the object toggled each time the goal orientation was achieved, switching between 15° and -15° . The mean position of the robots was steered to a workspace corner whenever the goal orientation was achieved, or the variance grew larger than a threshold amount. The cyan arrows around the object show the normalized artificial potential field forces used to steer the mean position of the swarm away from narrow ends of the object. These are used, as in [2], during the regathering process if the mean position of the swarm distribution passes the object. The swarm, using control law (27), achieved three goal orientations in 150 s.

VII. CONCLUSION AND FUTURE WORK

This paper presented an analysis of torque applied by a swarm to a long rod. Three distributions were studied: uniform, triangular, and normal. Though particles are often normally distributed, the closed-form solution for optimizing torque from normal distributions involves error functions. We illustrated numerically that the triangular distribution well-approximates the normal distribution, and gave analytic solutions with uniform and triangular distributions. We designed PD controllers using this analysis, and verified the analysis with simulation and hardware experiments. Future efforts should examine the effects of Brownian noise, pose control for multiple-part assembly, trajectory prediction, manipulation in crowded workspaces, and extend this analysis to 3D.

VIII. ACKNOWLEDGEMENTS

This material is based upon work supported by the National Science Foundation under IIS-1553063, CNS-1646566, and IIS-1619278.

REFERENCES

- [1] S. Shahrokhi and A. T. Becker, "Object manipulation and position control using a swarm with global inputs," in *IEEE International Conference on Automation Science and Engineering (CASE)*, 2016, pp. 561–566.
- [2] S. Shahrokhi, L. Lin, C. Ertel, M. Wan, and A. T. Becker, "Steering a swarm of particles using global inputs and swarm statistics," *IEEE Transactions on Robotics*, vol. 34, no. 1, pp. 207–219, 2018.
- [3] A. Sudsang, F. Rothganger, and J. Ponce, "Motion planning for disc-shaped robots pushing a polygonal object in the plane," *IEEE Trans. Robot. Autom.*, vol. 18, no. 4, pp. 550–562, Aug. 2002.
- [4] J. Fink, N. Michael, and V. Kumar, "Composition of vector fields for multi-robot manipulation via caging," in *Robotics Science and Systems*, Atlanta, GA, Jun. 2007.
- [5] K. Lynch, "Locally controllable manipulation by stable pushing," *IEEE Trans. Robot. Autom.*, vol. 15, no. 2, pp. 318–327, Apr. 1999.
- [6] F. Ruggiero, V. Lippiello, and B. Siciliano, "Nonprehensile dynamic manipulation: A survey," *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 1711–1718, July 2018.
- [7] J. Alonso-Mora, S. Baker, and D. Rus, "Multi-robot formation control and object transport in dynamic environments via constrained optimization," *The International Journal of Robotics Research*, vol. 36, no. 9, pp. 1000–1021, 2017.
- [8] K. M. Lynch, "Nonprehensile robotic manipulation: Controllability and planning," Ph.D. dissertation, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, Mar. 1996.
- [9] S. Akella, W. H. Huang, K. M. Lynch, and M. T. Mason, "Parts feeding on a conveyor with a one joint robot," *Algorithmica*, vol. 26, no. 3, pp. 313–344, 2000.
- [10] J. D. Bernheisel and K. M. Lynch, "Stable transport of assemblies by pushing," *IEEE Transactions on Robotics*, vol. 22, no. 4, pp. 740–750, Aug. 2006.
- [11] M. D. Armani, S. V. Chaudhary, R. Probst, and B. Shapiro, "Using feedback control of microflows to independently steer multiple particles," *Journal of Microelectromechanical systems*, vol. 15, no. 4, 2006.
- [12] A. Becker, R. Sandheinrich, and T. Bretl, "Automated manipulation of spherical objects in three dimensions using a gimbaled air jet," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct. 2009, pp. 781–786.
- [13] S. Martel, "Magnetotactic bacteria for the manipulation and transport of micro-and nanometer-sized objects," *Micro-and Nanomanipulation Tools*, vol. 13, 2015.
- [14] D. Lohin, C. Tremblay, M. Mohammadi, and S. Martel, "Exploiting the responses of magnetotactic bacteria robotic agents to enhance displacement control and swarm formation for drug delivery platforms," *The International Journal of Robotics Research*, vol. 36, no. 11, pp. 1195–1210, 2017.
- [15] Y. Ou, D. H. Kim, P. Kim, M. J. Kim, and A. A. Julius, "Motion control of tetrahymena pyriformis cells with artificial magnetotaxis: model predictive control (MPC) approach," in *IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2012, pp. 2492–2497.
- [16] M. Egerstedt and X. Hu, "Formation constrained multi-agent control," *IEEE Trans. Robotics Automat.*, vol. 17, pp. 947–951, 2001.
- [17] M. A. Hsieh, V. Kumar, and L. Chaimowicz, "Decentralized controllers for shape generation with robotic swarms," *Robotica*, vol. 26, no. 05, pp. 691–701, 2008.
- [18] X. Yan, Q. Zhou, J. Yu, T. Xu, Y. Deng, T. Tang, Q. Feng, L. Bian, Y. Zhang, A. Ferreira, and L. Zhang, "Magnetite nanostructured porous hollow helical microswimmers for targeted delivery," *Advanced Functional Materials*, vol. 25, no. 33, pp. 5333–5342, 2015.
- [19] J. Zhang, O. Onaizah, K. Middleton, L. You, and E. Diller, "Reliable grasping of three-dimensional untethered mobile magnetic microgripper for autonomous pick-and-place," *IEEE Robotics and Automation Letters*, vol. 2, no. 2, pp. 835–840, 2017.
- [20] A. Becker, G. Habibi, J. Werfel, M. Rubenstein, and J. McLurkin, "Massive uniform manipulation: Controlling large populations of simple robots with a common input signal," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Nov. 2013, pp. 520–527.
- [21] D. L. Christensen, S. A. Suresh, K. Hahn, and M. R. Cutkosky, "Let's all pull together: Principles for sharing large loads in microrobot teams," *IEEE Robotics and Automation Letters*, vol. 1, no. 2, pp. 1089–1096, 2016.
- [22] S. Shahrokhi and A. T. Becker, "Stochastic swarm control with global inputs," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2015, pp. 421–427.
- [23] M. Rubenstein, C. Ahler, and R. Nagpal, "Kilobot: A low cost scalable robot system for collective behaviors," in *IEEE International Conference on Robotics and Automation (ICRA)*, May 2012, pp. 3293–3298.
- [24] M. Rubenstein, A. Cornejo, and R. Nagpal, "Programmable self-assembly in a thousand-robot swarm," *Science*, vol. 345, no. 6198, pp. 795–799, 2014.
- [25] K-Team, "Kilobot, www.k-team.com," 2015.