Olfactory Navigation in the Real World: Simple Local Search Strategies for Turbulent Environments

James B. Hengenius[a,*], Erin G. Connor[b], John P. Crimaldi[b], Nathaniel N. Urban[c], and G. Bard Ermentrout[d]

[a]Department of Epidemiology, University of Pittsburgh, Pittsburgh, PA 15260, USA

[b]Civil, Environmental and Architectural Engineering, University of Colorado at Boulder, Boulder, CO 80309, USA

[c]Department of Biological Sciences, Lehigh University, Bethlehem, PA 18015, USA

[d]Department of Mathematics, University of Pittsburgh, Pittsburgh, PA 15260, USA

[*]Corresponding author, james.hengenius@pitt.edu

December 9, 2020

Abstract

Olfaction informs animal navigation for foraging, social interaction, and threat evasion. However, turbulent flow on the spatial scales of most animal navigation leads to intermittent odor information and presents a challenge to simple gradient-ascent navigation. Here we present two strategies for iterative gradient estimation and navigation via olfactory cues in 2D space: tropotaxis, spatial concentration comparison (*i.e.*, instantaneous comparison between lateral olfactory sensors on a navigating animal) and klinotaxis, spatiotemporal concentration comparison (*i.e.*, comparison between two subsequent concentration samples as the animal moves through space). We then construct a hybrid model that uses klinotaxis but utilizes tropotactic information to guide its spatial sampling strategy. We find that for certain body geometries in which bilateral sensors are closely-spaced (*e.g.*, mammalian nares), klinotaxis outperforms tropotaxis; for widely-spaced sensors (*e.g.*, arthropod antennae), tropotaxis outperforms klinotaxis. We find that both navigation strategies perform well on smooth odor gradients and are robust against noisy gradients represented by stochastic odor models and

real turbulent flow data. In some parameter regimes, the hybrid model outperforms klinotaxis alone, but not tropotaxis.

Highlights:

- We simulate two olfactory navigation strategies: tropotaxis and klinotaxis
- Both strategies locate point sources, trails, and plumes; they can also follow trails
- Both strategies use local gradient estimates, but are robust against turbulence
- The geometry of olfactory sensors determines which strategy is more successful

Keywords: klinotaxis, tropotaxis, animal navigation, olfaction, computational modeling

1 Introduction

A spawning salmon can navigate up turbulent streams to find the location where it hatched [1]; a trained rescue dog can locate a lost hiker using a hours-old scent trail [2]; a moth can detect a potential mate over long distances by pheromones alone [3]. Olfaction is an evolutionarily ancient sense vital across *Animalia* for directing navigation[4]. It informs foraging, finding mates, and avoiding danger. Animals extract information about the amount and type of chemical odorants by sampling their fluid medium. However, fluid flow - and by extension the olfactory cues carried in the fluid - is turbulent on the spatial scales of most animal navigation [5, 6]. The fluctuating local gradients experienced by animals may even point away from the true chemical maximum. Turbulence leads to an intermittent and noisy odor signal that might prevent animals from using simple chemical gradient ascent (descent) algorithms to navigate toward (away from) odor sources.

Despite these challenges, animals routinely perform feats of olfactory navigation on scales from local foraging and predation (*e.g.,* foraging insects and rodents, hunting sharks)[7, 8, 9, 10, 11, 12] to long-range homing and migration (*e.g.,* spawning salmon and controversially, homing pigeons)[1, 13, 14]. Animals use multiple strategies to perform olfactory navigation tasks, and we propose that the strategies they adopt may depend on their body and olfactory sensor morphologies. Most animals with bilateral bodyplans have paired left-right olfactory sense organs. However, the relative distance between these two sensors varies across species [15]. Widely-spaced sensors allow for the resolution of two spatially distinct left-right odor concentration samples which may be used to estimate a component of the local odor gradient across the two sensors. However, the intermittent and spatially-correlated nature of turbulent odors may prevent closely-spaced sensors from accurately estimating the local odor gradient.

Perhaps the simplest chemosensory search strategy is bacterial chemotaxis. Bacterial sensors are very closely-spaced (receptors on the surface of single cells) and the chemoattractant distri-

bution is subject to microscopic fluctuations that may obscure its gradient[16]. The closeness of the sensors relative to the fluctuating environment rules out comparison of concentrations across sensors for gradient estimation; bacterial chemosensory input is essentially a point estimate in space. To overcome this limitation and navigate up macroscopic gradients, a bacterium uses the change in chemoattractant concentrations over time as a cue to switch between two behavioral states[17]: When concentrations are increasing, a bacterium prefers the "running" state and travels in straight trajectories. When concentrations are decreasing, it prefers the "tumbling" state and engages in a random walk. Runs allow bacteria to exploit the information that they are traveling up the gradient by continuing to travel in the same direction. Tumbling allows bacteria traveling down the gradient to randomly sample different directions until they reorient in the proper direction. In contrast to this stochastic search strategy, animal navigation occurs at macroscopic scales. Here, sensors may sample odors from widely-separated volumes of air (*e.g.,* between antennae or between laterally-separated nares). In these cases, sensors and their sample volumes may be sufficiently separated in space to estimate gradients between the two samples. Even when sensors are relatively closely-spaced, animals may move their sensors through space to make spatiotemporal comparisons (*e.g.,* an animal may inhale a sample volume, move its head, then inhale again). Here, we explore two animal strategies that navigate toward an odor source using local estimates of the odor gradient.

The most straightforward strategy we consider is tropotaxis, the simultaneous comparison of olfactory concentrations from two spatial locations. Dual left-right olfactory channels such as antennae or nares enable comparison of two concentrations at spatially separated positions. Arthropods with widely separated antennae relative to their body size (and the consequent ability to resolve spatially distinct odor concentrations) provide examples of such stereo-olfactory navigation [7, 8, 18, 19, 20], as do specialized vertebrates like the hammerhead shark[21]. This method allows animals to approximate the odor gradient between their sensors and orient left or right in the direction of higher concentration. Such a stereo strategy requires no memory of previously sampled odor concentrations.

For organisms that lack widely spaced left-right olfactory sensors, stereo gradient estimation may become unreliable when the odor gradient across the left-right sample volumes is small and the signal is noisy. To increase the effective distance between concentration samples and improve gradient estimation, animals moving through space can use klinotaxis, the comparison of subsequent odor samples taken at laterally separated locations[22]. Klinotaxis has been observed across *Animalia*, from nematodes and insect larvae to vertebrates such as sharks and mammals [23, 24, 25, 26]. Unlike tropotaxis, the intrinsic delay between samples requires that animals store at least one concentration value in memory for comparison with the subsequent concentration

sample.

Experimental evidence suggests that even some mammals, whose nares are not widely separated relative to their body size, can use stereo-olfactory information [27, 28, 29]. In a set of elegant experiments involving blind moles, Catania demonstrated that occlusion of one naris lead to turning over-correction toward the open naris; reversal of the left-right nares' signals (by means of nasal tubes) led to repulsion from the odor source. Catania concluded that these mammals used a hybrid strategy: klinotaxis allowed the animal to orient toward the source from a distance where the gradient between the nares was too small to be differentiated from noise. Near the source, where the gradient across the nares was larger, navigation was aided by this additional left-right stereo information.

It should be noted that this work primarily considers searches in two dimensions (*i.e.*, near surfaces) such as animals navigating along the ground or an aquatic bed. While organisms navigating 3D environments (airborne or aquatic) typically have also symmetric sensors, they may adopt more complex olfactory search behaviors than simple tropotaxis or klinotaxis. Notably, they may incorporate additional wind-directional (anemometric) information in their steering, taking advantage of the fact that in the presence of convective flow an odor source is likely to be upwind of odor perceived by an organism [30]. For example, moths "surge" upwind upon the detection of a chemoattractant pheromone, aligning opposite the direction of airflow. Upon loss of pheromone detection, they orient perpendicular to the flow direction and "cast" back and forth in an attempt to reacquire the signal. Flying insects such as *Drosophila* adopt similar strategies when navigating 2D environments [31]. We do not consider such strategies here.

Many species, including humans, are aided in olfactory navigation by internal cognitive maps of their odor environments [4, 32]. For example, rats are capable of learning the locations of odor sources via olfactory navigation and then revisiting them without continued reliance on olfactory cues[33, 11]. Navigational models such as infotaxis and entrotaxis couple odor detection events with spatial maps of odor source location likelihood [34, 35]. These models are most informative in odor regimes where encounters with the odorant are infrequent; they eliminate regions of the map where odor encounters are sparse by assigning these regions a low likelihood for the odor source location. In contrast, we consider odor regimes where the odorant is plentiful (*i.e.*, consistently above detection thresholds), and characterize navigation methods driven by local gradient estimation rather than mapping.

In this paper, we consider basic 2D olfactory search algorithms - tropotaxis and klinotaxis - in the absence of cognitive maps. We show that the optimal strategy depends on the navigator's body morphology, with widely-spaced bilateral sensors favoring tropotaxis and closely-spaced bilateral sensors favoring klinotaxis. A hybrid model, driven by klinotaxis with a tropotactic bias in

sampling, improves successful localization of odor sources, mimicking observations of Catania and modeling predictions of Liu *et al.*[27, 36]. These minimal algorithms are capable of navigating to odor point sources, along odor trails, and through odor plumes represented by smooth (*i.e.*, spatially continuous) time-invariant gradients and discontinuous, dynamic stochastic models. Finally, we test the navigators' performance in real plumes using planar laser-induced fluorescence (PLIF) data of acetone vapor as a surrogate for a chemoattractant.

2 Model and Methods

We evaluate the performance of navigating model agents ("navigators") using the tropotactic, klinotactic, and hybrid strategies across a series of increasingly realistic odorant landscapes. The navigators operate in discrete-time and continuous-space, where the time interval reflects the sampling interval of the navigator's olfactory sensors (*i.e.,* the time between "sniffs"). The landscapes examined include idealized smooth gradients representing odor point-sources, trails, and plumes; we further examine stochastic odor models representing turbulence; finally, we examine the navigators' performance on real plume data generated by planar laser-induced fluorescence of acetone vapor plumes.

2.1 The Navigators

2.1.1 The Tropotactic Navigator

The tropotactic navigator is a simple model driven by bilateral concentration comparison. It is inspired by earlier work by this group [37] and shares parallels with several extant agent-based models in the literature[38, 39, 40]. The navigator has a geometry consisting of a body and a pair of bilateral olfactory sensors (representative of nares or antennae) illustrated in **Figure 1a** and superimposed on an ant body in **Figure 1b**. The body coordinates are (x, y) and the forward heading is θ. The bilateral sensors are at a distance ℓ from the body and separated from the midline by an inter-sensor angle, γ. Thus, the positions of the left and right sensors are given as

$$\begin{aligned}(x_{\mathrm{L}}, y_{\mathrm{L}}) &= (x + \ell\cos(\theta+\gamma), y + \ell\sin(\theta+\gamma)\\ (x_{\mathrm{R}}, y_{\mathrm{R}}) &= (x + \ell\cos(\theta-\gamma), y + \ell\sin(\theta-\gamma)\end{aligned}. \tag{1}$$

The navigator moves through 2D space in discrete odor-sampling timesteps, Δt, according to

$$\begin{aligned}x_{t+\Delta t} &= x_t + V(x,y,t)\cos(\theta)\Delta t\\ y_{t+\Delta t} &= y_t + V(x,y,t)\sin(\theta)\Delta t\end{aligned}, \tag{2}$$

where V, the velocity, may be constant or depend on concentration (and therefore depend on

time and space). For simplicity, we adopt the approach of [39] and choose a constant velocity for our analysis. The heading θ is updated by a steering function F,

$$\theta_{t+\Delta t} = \theta_t + F(x, y, t) \quad , \tag{3}$$

which likewise depends on the dynamic concentration distributions (and therefore time and space). The tropotactic navigator samples the concentrations at its left and right sensor point locations, C_L and C_R, at each time step. This approach represents a simplification of the ant chemo-sensing model of Amorim *et al.*, which integrates odor concentration over an area which is "swept" by the ant antennae [40]. It uses this information to update its heading θ at each step. This leads to a simple steering function,

$$F(x, y, t) = \beta \frac{C_\mathrm{L} - C_\mathrm{R}}{C_\mathrm{L} + C_\mathrm{R}} \Delta t \tag{4}$$

which turns the heading θ to the left if C_L is greater than C_R and to the right if C_R is greater than C_L. The scalar value β controls the rate of turning per concentration sample. The term $(C_\mathrm{L} - C_\mathrm{R})/(C_\mathrm{L} + C_\mathrm{R})$ is known as the Michelson contrast and arises from Weber's law: the perceived change in a stimulus is proportional to the total stimulus. The neurological underpinnings of Weber's law have been explored in animal olfaction[41] and the Michelson contrast been incorporated in similar models of tropotactic navigation [38].

2.1.2 The Klinotactic Navigator

Contrasting the spatial concentration comparison of the tropotactic navigator, the klinotactic navigator relies on spatiotemporal concentration comparison. The navigator has a body geometry similar to that of the tropotactic navigator, though it has an additional internal degree of freedom allowing its head (sensors) to move laterally independent of the navigator's forward heading. Like the tropotactic navigator, body coordinates are (x, y), angular body heading is θ, and two sensors are located at length ℓ from the body. However, the sensors may move laterally relative to θ, with a lateral neck deflection angle ϕ bounded by $\pm\phi_\mathrm{max}/2$. The sensors remain separated by an inter-sensor angle, γ, with respect to ϕ rather than the body midline (see **Figure 1c**). The navigator geometry is shown superimposed on a rat silhouette in **Figure 1d**. Accounting for this addition degree of freedom in the body geometry, eqn. 1 is modified to give coordinates for the two sensors,

$$\begin{aligned} (x_\mathrm{L}, y_\mathrm{L}) &= (x + \ell\cos(\theta + \phi + \gamma), y + \ell\sin(\theta + \phi + \gamma)) \\ (x_\mathrm{R}, y_\mathrm{R}) &= (x + \ell\cos(\theta + \phi - \gamma), y + \ell\sin(\theta + \phi - \gamma)) \end{aligned} \tag{5}$$

This body geometry allows the navigator to perform klinotaxis, increasing the effective lateral

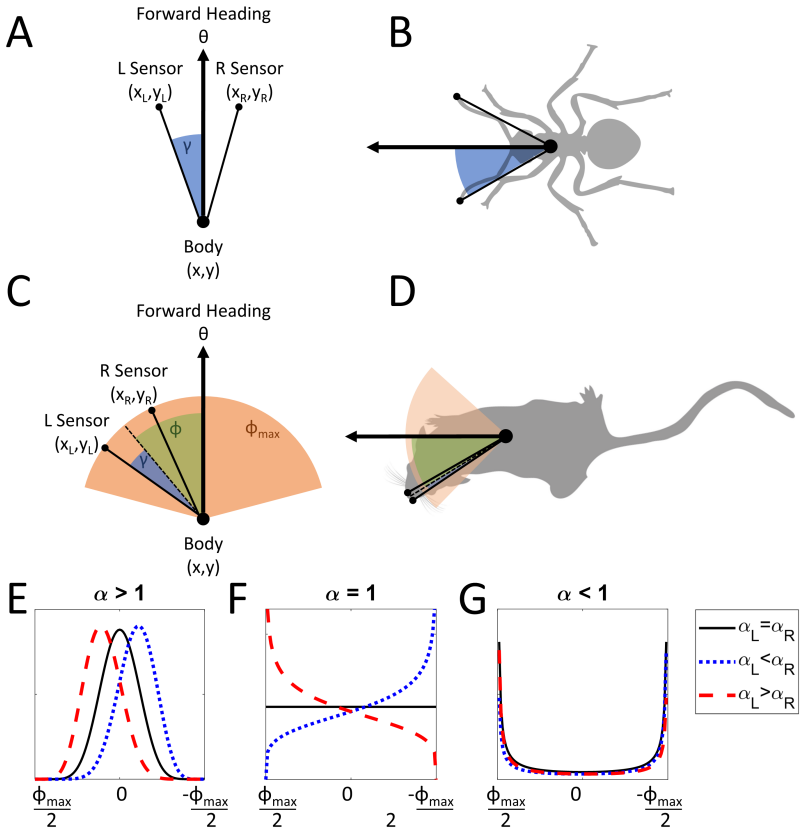

Figure 1: **Navigator geometries.** The tropotactic navigator (A) has a body (x, y) with forward heading θ and two olfactory sensors $(x_{\mathrm{L/R}}, y_{\mathrm{L/R}})$ separated from the body midline by angle γ (blue arc). An example of the tropotactic geometry is shown imposed over an insect bodyplan (B). The tropotactic navigator (C) has a body (x, y) with forward heading θ. It also has paired sensors $(x_{\mathrm{L/R}}, y_{\mathrm{L/R}})$ separated by an angle γ (blue arc), but these are allowed to move independent of the midline with a nose deflection angle ϕ (green arc). This deflection is bounded by a maximum range $\phi_{\max}$ (orange arc). An example of the tropotactic geometry is shown imposed over a mammal bodyplan (D). The nose deflection angle is drawn from a beta distribution (eqn. 8) that may be unimodal, uniform, or bimodal depending on the value of parameter α (black lines, E-G). The hybrid navigator uses stereo-olfactory information to guide its nose deflection sampling strategy, shifting sampling to the left when the left sensor detects a higher concentration (red dashed lines, E-G) and to the right when the right sensor detects a higher concentration (blue dotted lines, E-G). Note that the x-axes of E-G are reversed because a leftward turn is in the positive x-direction and *vice versa*.

separation between its sensor samples by sampling odor concentration, moving its sensors by altering neck angle ϕ, and resampling concentration every Δt time units. The concentration registered at time t, C_t, is the mean value of concentrations at the left and right nares,

$$C_t = (C_{\mathrm{L}} + C_{\mathrm{R}})_t/2. \tag{6}$$

The navigator still moves according to eqns. 2-3, but requires memory of its previously sampled concentration, $C_{t-\Delta t}$, obtained when its sensors were positioned at $\phi_{t-\Delta t}$. It also requires a new steering function, F, capable of comparing $C_{t-\Delta t}$ with the current concentration sample, C_t, when the sensors were are positioned at ϕ_t. Here, we adopt a simple all-or-nothing steering function for klinotaxis. Whenever the navigator moves its sensors (updates ϕ) and detects a higher concentration than the one it previously sampled (at $\phi_{t-\Delta t}$), it aligns its heading with its sensors according to eqn. 3:

$$F(x, y, t) = \begin{cases} \phi_t & \text{if } C_t > C_{t-\Delta t} \\ 0 & \text{otherwise} \end{cases} \tag{7}$$

Finally, the klinotactic navigator requires a mechanism for updating the position of its sensors relative to its body. We consider random sensor movement, where ϕ_t is a random variable drawn from a scaled beta distribution. The beta distribution pdf is defined on the interval $[0, 1]$ by shape parameters α_1 and α_2:

$$\mathrm{Beta}(\alpha_1, \alpha_2) = \frac{x^{\alpha_1}(1-x)^{\alpha_2 - 1}}{\mathrm{B}(\alpha_1, \alpha_2)}, \tag{8}$$

where $\mathrm{B}(\alpha_1, \alpha_2)$ is a normalization term,

$$\mathrm{B}(\alpha_1, \alpha_2) = \frac{\Gamma(\alpha_1)\Gamma(\alpha_2)}{\Gamma(\alpha_1 + \alpha_2)}, \tag{9}$$

and Γ is the gamma function. The distribution is scaled so that sensor deflection angles are chosen from $\pm\phi_{\max}/2$:

$$\phi_t \sim \phi_{\max}(\mathrm{Beta}(\alpha, \alpha) - \tfrac{1}{2}), \tag{10}$$

Here, $\phi_{\max}$ is the maximum angular deflection of the sensors (see **Figure 1c**) and α is the parameter of the beta distribution. Both α parameters are identical so the sensor deflection distribution is symmetric. This formula allows for control of the sensor sampling distribution; α values less than one lead to bimodal sampling toward the extremes of $\phi_{\max}$, α values greater than one lead to central unimodal sampling, and α values of one lead to uniform sampling (see **Figure 1e-g**, black

lines).

2.1.3 The Hybrid Navigator

Work by Catania suggests that animals with closely spaced nares (sensors) may still utilize concentration differences across bilateral sensors to localize odor sources[27]. Because the sensors are close together and the turbulent signal is noisy, this stereo information can only be reliably used when the concentration difference across the sensors is large.

The hybrid navigator uses the geometry of the klinotactic navigator (**Figure 1c**, eqn. 5). Like the klinotactic navigator, the hybrid navigator moves according to eqns. 2-3 and steers according to eqn. 7. However, the nose deflection distribution (eqn. 10) is modified to incorporate stereo information from the two sensors,

$$\phi_t \sim \phi_{\max}(\text{Beta}(\alpha_{\text{L}}, \alpha_{\text{R}}) - \tfrac{1}{2}). \tag{11}$$

Here, the two parameters of the beta distribution are allowed to vary such that the nose preferentially samples to the navigator's left side if the left sensor detects a higher concentration and samples to the right side if the right sensor detects a higher concentration. The two α parameters take the form,

$$\begin{aligned} \alpha_{\text{L}} &= \alpha + \alpha\tfrac{(C_{\text{L}} - C_{\text{R}})_{t-\Delta t}}{C_{\max}} \\ \alpha_{\text{R}} &= \alpha - \alpha\tfrac{(C_{\text{L}} - C_{\text{R}})_{t-\Delta t}}{C_{\max}} \end{aligned}. \tag{12}$$

When concentrations $C_{\text{L}} > C_{\text{R}}$ at time $t - \Delta t$, then $\alpha_{\text{L}} > \alpha_{\text{R}}$ at time t and the nose will preferentially sample to the navigator's left side (**Figure 1e-g**, blue lines). If $C_{\text{L}} < C_{\text{R}}$ at time $t - \Delta t$, then $\alpha_{\text{L}} < \alpha_{\text{R}}$ at time t and the nose will preferentially sample to the navigator's right side (**Figure 1e-g**, red lines). $C_{\max}$, the maximum detectable odor concentration, is introduced as a normalization so that small fluctuations in the turbulent environment do not lead the navigator astray. Animals may engage in more complicated dynamic thresholding rather than normalization by $C_{\max}$, and this is in line with Catania's suggestion that animals with closely-spaced nares can only detect large stereo differences in odor concentration where the odor gradient is steep.

For smooth gradient environments (see **Section 2.2.1**), $C_{\max}$ set to the concentration value at coordinates x, y = (1,0). These coordinates are chosen because navigators reaching a distance of 1 from the source are considered to have successfully located it. These coordinates therefore represent the highest concentration detectable by the navigator prior to successful identification of the source, regardless of the odor landscape. Note that these coordinates were chosen because the concentration is not defined for $x = 0$ in the plume (see eqn. 16).

For stochastic odor environments (see **Section 2.2.2**), $C_{\max}$ is set the the average Poisson

particle detection rate at the point $x, y = (1,0)$ near the source. Because the capture radius around the source is 1, these coordinates correspond to the highest average detection rate experienced by navigators prior to successfully locating the source.We choose these coordinates because the average particle detection rate is undefined at the source (see eqns. 17,20).

For PLIF plumes (see **Section 2.2.3**), $C_{\max}$ is set to one. This is because the plume data is normalized to a maximum value of one.

2.1.4 The Correlated Random Walk Agent

To confirm that the navigators perform better than chance, we constructed an agent that explores via a correlated random walk (CRW) to represent random searches. Like the navigators, it consists of a body at (x, y) and a heading θ. It moves according to eqn. 2 with a constant velocity $V = 1$. Unlike the navigators, whose steering functions are driven by odorant concentrations, the CRW agent heading is updated stochastically,

$$F(x, y, t) = \mathrm{N}(0, \sigma_{\mathrm{CRW}})\sqrt{\Delta t} \quad . \tag{13}$$

Here, σ_{CRW} is the standard deviation of the heading distribution (set to 0.5). This results in a random walk in which the heading at each time point is correlated with the subsequent time point.

2.2 The Odor Environments

In an idealized smooth and time-invariant odor landscape that increases monotonically approaching the odor source, gradient ascent would be an efficient strategy for locating the source. In contrast to this ideal case, the odor landscape is typically turbulent and dynamic rather than smooth and time-invariant; it is not guaranteed to increase monotonically approaching an odor source and may be highly intermittent. We model the odor environments in three increasing levels of realism and evaluate the performance of the navigators in each.

2.2.1 Smooth Odor Gradient Environments

We first tested the navigators' performance in smooth, time-invariant, monotonically increasing odor landscapes to verify that they could navigate using odor cues.

We define a continuously released odor point source located at the origin of a Cartesian plane as

$$C(x, y) = \exp -\sqrt{\frac{x^2 + y^2}{\sigma^2}} \tag{14}$$

where σ is the space constant determining the width of the odor distribution. We modify eqn.

14 by removing the dependence on the y direction to create an exponential odor trail along the y-axis:

$$C(x, y) = \exp -\sqrt{\frac{x^2}{\sigma^2}}. \tag{15}$$

Finally, we define an odor plume based on a 2D cross-section of a 3D plume [42],

$$C(x, y) = \frac{q}{2\pi kx} \exp \frac{-uy^2}{4kx}, \tag{16}$$

where q is the mass emission rate, u is the velocity of flow in the positive x direction of the Cartesian plane, and k is the effective diffusivity transverse to the flow. Note that in atmospheric dispersion modeling k is typically an empirically-fit function of x; we adopt Stockie's simplifying assumption here by treating k as a constant[42].

Examples of the smooth point source, trail, and plume may be seen in **Figure 2a-c**, respectively.

2.2.2 Stochastic Odor Environments

Unlike smooth gradients described above, real-world odor distributions are subject to turbulent advection on the length scales of most animal navigation. This leads to highly intermittent and noisy odor signals that fluctuate over time and space, complicating navigation [5]. To simply capture some of these complex dynamics without relying on computationally costly direct numerical simulation of the Navier-Stokes equations, we use the 3D Poisson detection event model derived in [34]. This model represents the stationary distribution of a convected odor plume, with stochastic events representing turbulent fluctuations.

It assumes that an odor source located at the origin is an emitter of detectable odor "particles." These particles are released at rate Q, with an average lifetime τ and an effective diffusivity constant D. The particles are advected away from the source in the positive x direction with a mean velocity U. For a sensor of radius a located at (x, y), the mean number of particle detection events is given as

$$\bar{R}(x, y) = \frac{aQ}{\sqrt{x^2 + y^2}} \exp \frac{-\sqrt{x^2 + y^2}}{\lambda} \exp \frac{xU}{2D} \tag{17}$$

$$\lambda = \sqrt{\frac{D\tau}{1 + \frac{U^2\tau}{4D}}}. \tag{18}$$

The particle encounters with the sensor are modeled as a Poisson process with mean rate $\bar{R}(x, y)$. Because the navigators rapidly sample the environment over small time intervals Δt during which the sensor movement is assumed to be negligible, the rate of detection over Δt is the product of $\bar{R}$

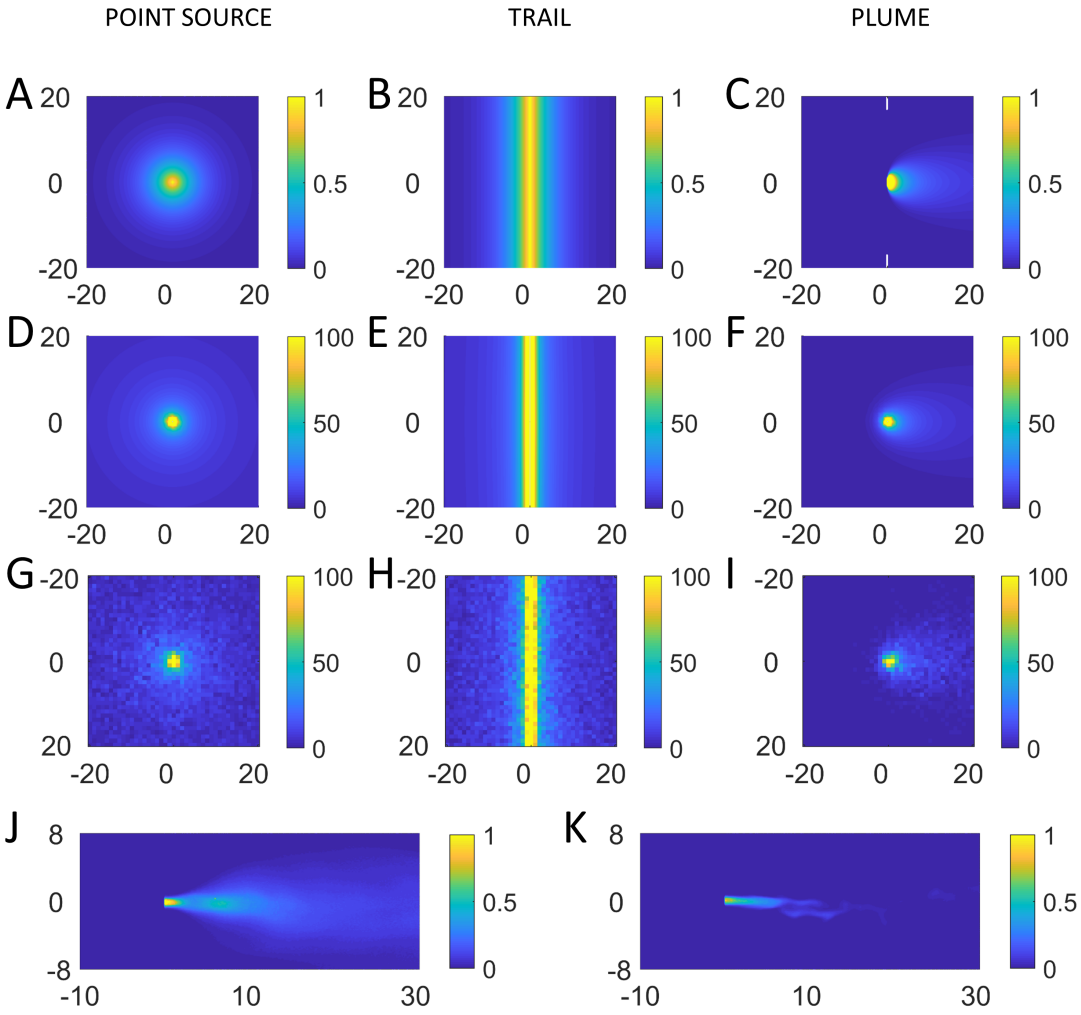

Figure 2: **Simulated and PLIF odor environments.** Smooth gradient odor point source, trail, and plume are shown in A-C, respectively. Average Poisson detection rates for the navigator sensors ($\bar{R}(x, y)\Delta t$, see eqn. 19) are shown for point source, trail, and plume in D-F, respectively. Single independent readings of the Poisson process sampled on a 40x40 spatial grid are shown for point source, trail, and plume in G-I, respectively. Example frames of the near-bed PLIF plume (J) and free-stream PLIF plume (K) show the plume broadening effect of the bed relative to the narrow free-stream plume.

and Δt. Samples at each sensor are drawn from

$$R(x,y) \sim \text{Pois}(\bar{R}(x,y)\Delta t). \tag{19}$$

The concentration of each sample is assumed to be proportional to the number of particle detection events, $C(x,y) \propto R(x,y)$, for the purposes of navigation in eqns. 4, 6, 7, and 12. Point source and plume navigator detection rates ($\bar{R}(x,y)\Delta t$) are shown in **Figure 2d,f**.

Eqns. 17-18 capture the shapes of the smooth gradient point source (eqn. 14, when U is set to zero) and smooth plume (eqn. 16, when U is greater than zero). To generate a Poisson model of a continuous trail, we set airflow U to zero and modify eqns. 17-18 to remove the concentration dependence on the y direction:

$$\bar{R}(x,y) = \frac{aQ}{\sqrt{x^2}} \exp -\sqrt{\frac{x^2}{D\tau}}. \tag{20}$$

As with the point source and plume, Poisson sampling for model navigation is performed according to eqn. 19 and concentrations are assumed to be proportional to number of detections. An illustration of the trail's average navigator detection rate ($\bar{R}(x,y)\Delta t$) is shown in **Figure 2e**. Independent Poisson samples on a 40-by-40 grid around the odor sources illustrates the characteristic fluctuations in the point source, trail, and plume models **Figure 2g-i**.

2.2.3 Planar Laser-Induced Fluorescence Plumes

Connor *et al.* used planar laser-induced fluorescence of neutrally-buoyant acetone vapor to characterize the behavior of airborne chemical plumes in open air (free-stream) and near-bed environments[5]. We use these chemical concentration distributions as surrogates for odor concentration to test the performance of the navigators in real turbulent flow conditions.

Plume data were obtained in a bench-top wind tunnel into which acetone vapor was released through a cylindrical tube. The odor was released through the source tube at a flow rate matching the mean flow rate in the wind tunnel (i.e., isokinetically). The near-bed and free-stream plume data differed in flow rate and release condition. For the near-bed case, the odor source was positioned 6 mm above the bed and had a flow rate of 10 cm/s. This location led to a more persistent plume and resulted in a broad distribution shown in **Figure 2j**. For the free-stream case, the odor source was positioned in the center of the wind tunnel (far from the walls/bed) with a greater flow rate of 20 cm/s. This location led to a more intermittent (rapidly fluctuating) plume structure and resulted in a narrow plume shown in **Figure 2k**. Both cases consisted of four minutes of video data (each) having a field of view of 30x16 cm. A camera mounted perpendicular to the laser light sheet collected images at 15 Hz.

All negative pixel values are set to zero and data is scaled so that the maximum pixel value in each plume is equal to one. Navigators sample the plume by reading the intensity value of the pixel nearest to each sensor's coordinates. To randomize the plume experienced by the navigators, the plume is initialized at a random time between zero and four minutes at the beginning of each navigator trial. The plume is then allowed to loop until the trial is complete.

2.3 Analysis Methods

All models were run and analyzed using MATLAB R2018a (The MathWorks, Inc.) on a 2.8 GHz quad-core computer with 32 GB RAM. We evaluated navigator performance according to their ability to find odor sources across a range of initial conditions and according to their ability to follow odor trails.

2.3.1 Odor Source Localization

The three navigation strategies were characterized by their search success rate across a range of initial positions, x_0, initial angles, θ_0, and odor environments. In these scenarios, a search was considered successful if the navigator moved within a capture radius of 1 distance unit of the odor source, whether it was a point source, trail, or plume source (green lines, **Figure 3a-d**). A search was considered a failure if it reached a distance greater than 20 from a point source or trail (blue dashed lines, **Figure 3a-b**), if it reached the edges of a 40x30 box surrounding a simulated plume (blue dashed lines, **Figure 3c**), or if it reached the edges of a 40x16 box surrounding the PLIF data plumes (blue dashed lines, **Figure 3d**). One hundred trials were simulated for each x_0, θ_0 combination (or x_0, y_0, θ_0 combination in the case of plumes) to estimate the average success rate for each set of initial conditions.

We examined the spatial distribution of success rates relative to odor sources and the average success rate for each navigator in each odor environment. Odor environment parameter values used in these odor localization tasks (illustrated in **Figure 2a-i**, eqns. 14-20) are listed in Table 1. Parameter values were chosen so that smooth gradient environments and stochastic environments were qualitatively similar in shape.

2.3.2 Trail Following Behavior

To evaluate the performance of the navigators on trails, we performed trials in which a navigator was started directly on the trail ($x_0, y_0 = 0, 0$) and oriented in a variety of initial headings ($\theta_0 =$ 0-$\pi/2$). We measured three features of navigators' trail-following performance.

First, we measured the total length traveled along the trail. This is defined as the total distance traveled in the y direction given that the navigator remains within a distance threshold of the trail

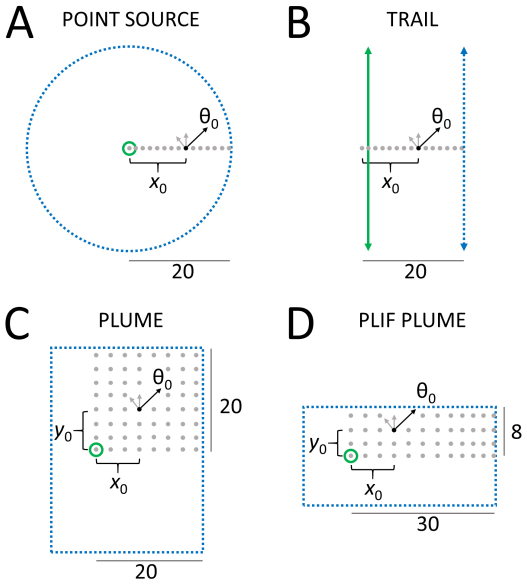

Figure 3: **Odor source localization schema.** Odor localization success rate trials for the point source (A) were conducted in a radial arena of radius 20; navigators were initialized at starting locations x_0 (dots) and starting headings θ_0 (arrows) ranging from distances 0 to 20 and 0-π, respectively. Trials were deemed successful if they reached a capture distance of 1 from the source (green line) and were considered failures if they reached a distance greater than 20 from the source (blue dashed line). This process was repeated for trails (B) in an arena of width 20 and infinite height. In plumes (C), navigators were initialized on a 20x20 spatial grid (x_0, y_0) with θ_0 values ranging from 0-2π. Total arena size for plumes was 30x40. In PLIF plumes (D), navigators were initialized on a 30x8 spatial grid (x_0, y_0) matching the dimensions of one half the recorded plume (30x16) and assuming symmetry about the x-axis; θ_0 values ranged from 0-2π. Total PLIF arena size was 40x16.

Table 1: **Odor Environment Parameters**

Parameter	Range	Dimension
σ^2	20	length^2
q	10000	mass·time^{-1}
k	1000	length^2· time^{-1}
u	1000	length· time^{-1}
a	1	length
Q	1000	particles·time^{-1}
τ	5000	time
U	10	length· time^{-1}
D	20	length^2· time^{-1}

($|x| < x_{thres}$, $x_{thres} = 2\sqrt{\sigma^2}$), so it excludes any distance traveled if the navigator leaves the trail.

Second, we measured the maximum excursion in the y direction. This is defined as the maximum value of $|y|$ given that the navigator remains within the distance threshold of the trail. Because the navigators may reverse direction while traveling along the trail, the maximum excursion may be less than the total length traveled along the trail.

Third, we measured the number of reversals the navigators make while traveling along the trail. These are defined in the change in sign of Δy given that the navigator remains within the distance threshold of the trail.

One thousand trials were simulated for each θ_0 to estimate average values for length traveled, maximum excursion, and reversal number.

3 Results

3.1 Tropotaxis and klinotaxis strategies find sources and follow trails on smooth odor gradients

We initially tested both navigation strategies in smooth gradient odor environments. These environments - consisting of point sources, trails, and plumes - are time-invariant and monotonically increase in concentration toward each odor source (**Figure 2a-c**). Under these conditions, gradient estimation and ascent should be a sufficient strategy to localize the odor sources. Animals are constrained by their physiology to make local estimates of gradient components between two points: either between two sensors in the case of tropotaxis or between two serially-sampled locations in the case of klinotaxis. While some animals are observed to periodically stop and sniff during olfactory searches (*e.g.*, mice and dogs), which may allow for more complete gradient estimation (*i.e.*, by utilizing multiple pairwise cocentration comparisons along different directions), animals often perform olfactory sampling while in continuous motion. As such, they must make navigational decisions with stereo or serial estimates of the local odor gradient. We find that even these estimates of the gradient are sufficient to guide tropotactic and klinotactic navigators toward

smooth gradient odor sources.

Three tropotactic example trajectories ($x_0, y_0 = 5, 0$; $\theta_0 = 0, \pi/2, \pi$; $\beta = 10, \gamma = \pi/4$, $\ell = 1$, $V = 1$, $\Delta t = 0.1$) are superimposed on the smooth gradient point source, trail, and plume in **Figure 4a-c**, respectively. Because there is no stochasticity in either the environment or navigator, these trajectories are deterministic. When initialized perpendicular to the odor gradient, the tropotactic navigator is quickly able to align with the gradient based on successive inter-sensor concentration comparisons (**Figure 4a-c**, yellow lines). When oriented directly parallel to the gradient and pointing toward the source, the navigator proceeds directly up the gradient (**Figure 4a-c**, black lines). It should be noted that the tropotactic navigator fails to arrive at the source when oriented parallel to the gradient and pointing away from the source (**Figure 4a-c**, white lines) In these cases, the inter-sensor concentration difference is zero and the steering function (eqn. 4) does not turn the navigator. In nature, animals may have additional strategies to prevent this outcome, such as initiating turning behavior when the concentration decreases over time (*e.g.*, the pirouette of *C. elegans*[23]). Natural sensor asymmetries and movements in animal sensor placement may also prevent this failure mode.

Three klinotactic example trajectories ($x_0, y_0 = 5, 0$; $\theta_0 = 0, \pi/2, \pi$; $\phi_{\max} = \pi/2, \alpha = 1, \gamma = \pi/16$, $\ell = 1$, $V = 1$, $\Delta t = 0.1$) are superimposed on the smooth gradient point source, trail, and plume in **Figure 4d-f**, respectively. The klinotactic navigator relies on a stochastic sampling mechanism: the "nose" is allowed to move independent of the forward heading according to eqn. 10. Because the navigator "follows its nose" according to the steering function (eqn. 7), the four sample klinotactic example trajectories shown are not deterministic. Although the klinotactic trajectories more tortuous than their tropotactic counterparts, they remain capable of localizing the source. The stochastic nature of their sampling mechanism also allows the klinotactic navigator to avoid the failure state of the tropotactic navigator (**Figure 4d-f**, white lines). Navigators initialized parallel to the gradient and oriented away from the source are capable of turning to successfully localize the source.

We also performed systematic evaluation of tropotactic and klinotactic success rates as a function of initial position and heading (schema in **Figure 3a-c**). Using the parameter values illustrated in **Figure 4a-f**, tropotaxis and klinotaxis success rates were similar in the smooth gradient environment. Both strategies were capable of successfully finding the smooth gradient point sources and trails across nearly the entire range of initial conditions (see **Figure 4g-l**). The only observed failure state across initial conditions was when the tropotactic navigator was initialized pointing directly down the symmetric concentration gradient (**Figure 4a-c**, white lines corresponding to the $\theta_0 = 0$ regions in **Figure 4g-h**). Both strategies successfully located the source of a plume in a rectangular arena (**Figure 4i,l**). (Note that **Figure 4i,l** show success rates averaged over θ_0 for

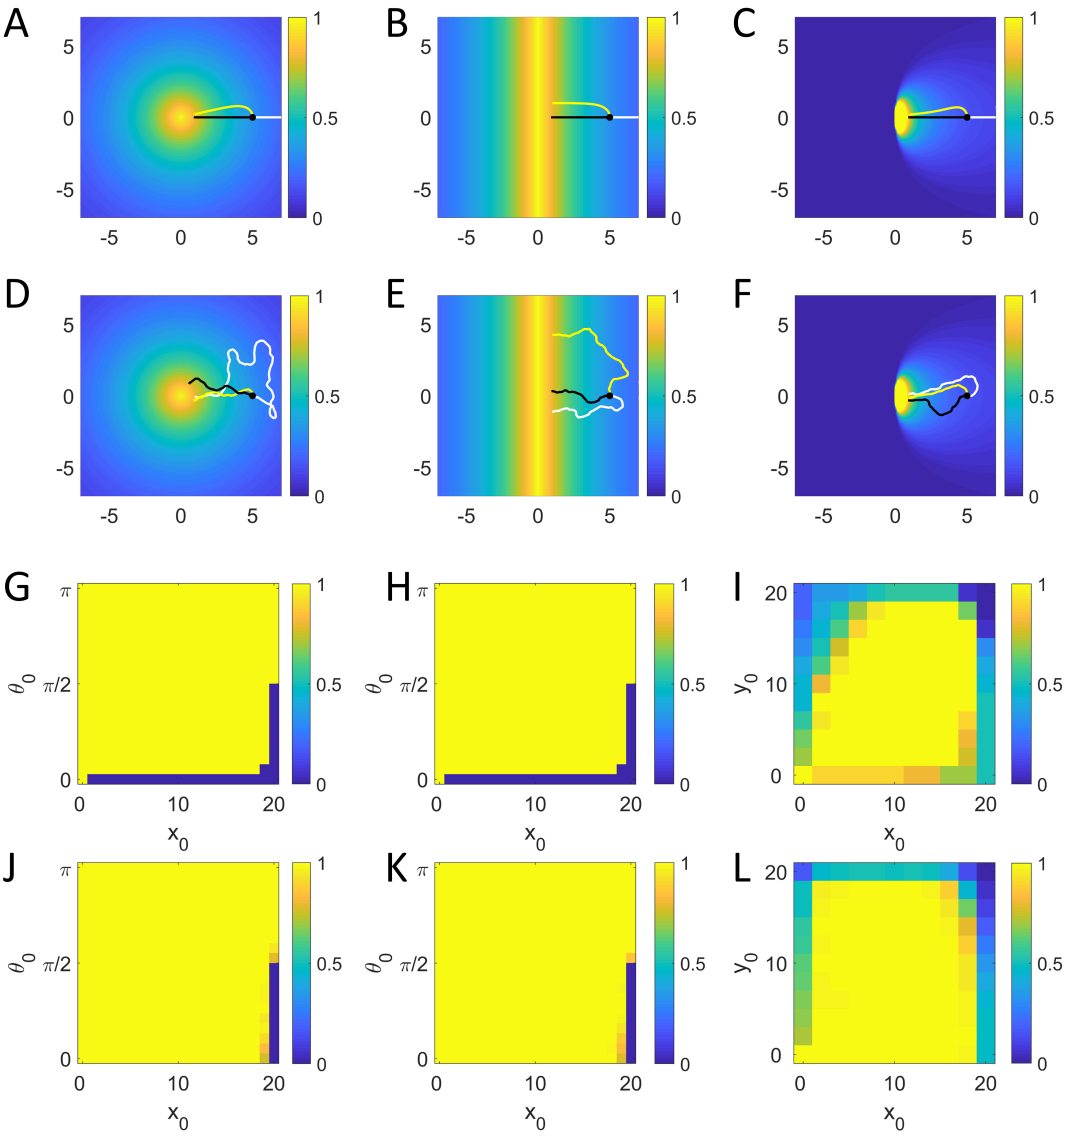

Figure 4: **Navigators successfully find sources using smooth gradients.** Tropotactic trajectories ($x_0, y_0 = 5, 0$; $\theta_0 = 0$ white line, $\theta_0 = \pi/2$ yellow line, $\theta_0 = \pi$ black line) are shown for the smooth gradient point source (A), trail (B), and plume (C). Klinotactic trajectories with the same initial conditions are shown for the point source (D), trail (E), and plume (F). Fraction of successful trials for tropotactic initial conditions are shown for the point source (G), trail (H), and plume (I). Fraction of successful trials for klinotactic initial conditions are shown for the point source (J), trail (K), and plume (L). Point source and trail success rate plots show success at initial distance x_0 and initial heading θ_0. Plume fraction of successful trial plots show spatial locations (x_0, y_0, see Figure 3c) and success rates averaged over θ_0 at each location. All fractional success rates are estimated from 100 trials per initial condition x_0, θ_0 (G,H,J,K) or x_0, y_0, θ_0 (I,L).

each initial position x_0, y_0.) Both navigators outperformed the CRW agent in all three types of environments, indicating that the algorithms performed better than chance (see **Supplemental Figure S1**.)

Finally, we evaluated navigators' ability to follow trails. The navigators were started on the trail ($x_0, y_0 = 0, 0$) at starting headings $\theta_0 = 0$-$\pi/2$ and allowed to run for 1000 time units using previously described parameter values. We examined each navigator's total length traveled along the trail in either direction, their maximum excursion along the trail in either direction, and the number of times they reversed direction while traveling along the trail. We observed that the stochastic nature of the klinotactic navigator led to more tortuous trajectories than those of the tropotactic navigator (see example trajectories in **Figure 5a**). Tropotactic and klinotactic navigators traveled similar total distances along the trail averaged over initial headings ranging from 0-$\pi/2$ (**Figure 5b**, 949.54 ± 0 vs 928.80 ± 3.20 distance units, respectively). The slight reduction in distance traveled by the klinotactic navigator can be attributed to the greater excursions in the x-direction exhibited in **Figure 5a**. The only observed failure state occurred in tropotaxis when the navigator's initial heading was perpendicular to the trail ($\theta_0 = 0$). In this case, the navigator failed to turn and acquire the trail because it was experiencing equal concentration values at both sensors (as in **Figure 4b**, white line). While both navigators traveled similar lengths along the trail, they differed in their greatest excursion along the trail (*i.e.*, their maximum $|y|$). Tropotaxis traveled an average of 949.54 ± 0 units away from the initial position (**Figure 5c**, red line) while klinotaxis traveled only 620.40 ± 229.35 units on average (**Figure 5c**, blue line). This difference can be attributed to the number of times each navigator reversed direction on the trail. The tropotactic navigator never reversed direction after it began following the trail, while the klinotactic navigator reversed direction an average of 3.7 ± 2.7 times per trial (**Figure 5d**). The high number of klinotactic reversals suggests that this minimal mechanism is capable of maintaining a navigator's position on a trail, but that additional mechanisms (*e.g.,* spatial memory) may be necessary to ensure consistent direction of travel along a trail.

3.2 Both strategies are robust against Poisson-distributed noise in odor concentration

Whereas the smooth gradient environments represent a time-average of the odor distributions, real odor environments are dynamic. The Poisson particle detection model derived by Vergassola *et al.* provides a computationally efficient means of simulating turbulent dynamic odor environments without relying on time-consuming direct numerical simulation of the Navier-Stokes equations [34]. This stochastic model simulates the number of detectable odor particles arriving at a sensor positioned in the vicinity of an odor source. One limitation of this model is that the resulting

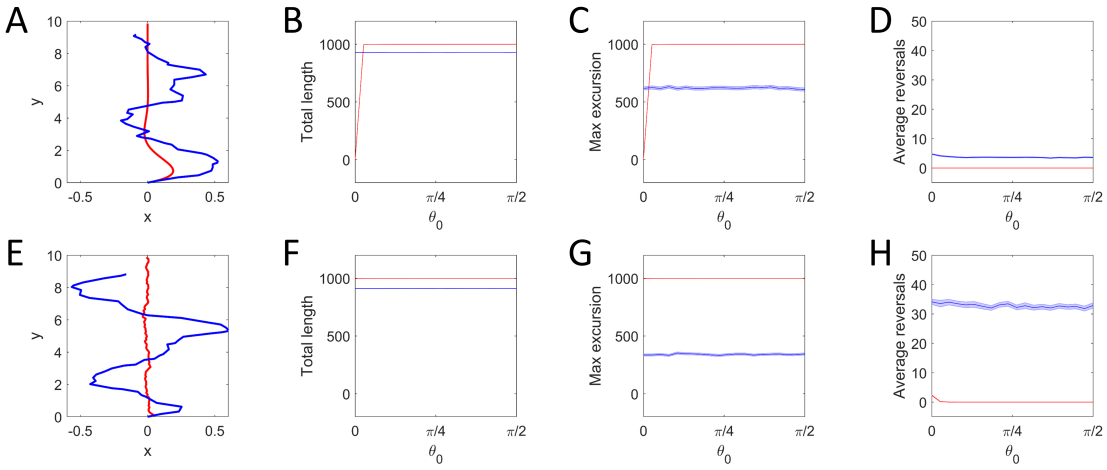

Figure 5: **Navigators follow trails in smooth gradient and stochastic odor environments.** Example tropotactic (red line) and klinotactic (blue line) trajectories are shown for initial conditions $x_0, y_0 = 0, 0$, $\theta_0 = \pi/3$ in the smooth gradient environment (A) and stochastic environment (E). Total distance traveled along the trail in either direction is shown for a range θ_0 in the smooth gradient environment (B) and stochastic environment (F). Maximum excursions along the trail (*i.e.,* how far the navigator traveled from its starting location) is shown for a range of θ_0 in the smooth gradient environment (C) and stochastic environment (G). The number of reversals in navigator direction on the trail are shown as a function of θ_0 for the smooth gradient environment (D) and stochastic environment (H). All values are averaged over 1000 trials for each initial angle θ_0. Shaded regions indicate 95% confidence intervals.

particle detection events lack the transient spatial correlations of true turbulent fluid flow. However, the loss of these correlations should increase intersensor variability in eqns. 3, 6, and 12, making the navigation task more challenging. As a consequence of the stochasticity in the environment, we find that overall navigator success is reduced relative to smooth gradient environments. However, both tropotactic and and klinotactic navigators remain able to localize sources and follow trails.

Sample trajectories for tropotactic and klinotactic navigators are shown in **Figure 6a-c** and **Figure 6d-f**, respectively. All simulations in **Figure 6** use the same parameter values as **Figures 4-5**. When oriented perpendicular to the gradient or facing up the gradient, both models found the source (**Figure 4a-c**, yellow and black lines). In smooth gradient environments, tropotaxis exhibited a failure state when oriented parallel to the gradient and away from the source (**Figure 4a-c**, white lines). The stochastic environment prevents this failure of tropotaxis by breaking the left-right concentration symmetry in eqn. 4 and allowing the navigator to turn toward the odor source (**Figure 6a-c**, white lines corresponding to the $\theta_0 = 0$ regions of **Figure 6g-i**).

Though the stochastic environment corrects this failure of tropotaxis, its overall effect is to decrease success rates of both navigators. We systematically evaluated navigator success rates across a range of initial conditions x_0, θ_0 (point sources and trails) and x_0, y_0, θ_0 (plumes). When seeking stochastic odor sources, both navigators show reduced success when the initial distance is large and the initial heading is pointing away from the source (**Figure 6g-l**, *cf.* **Figure 4g-l**). Notably,

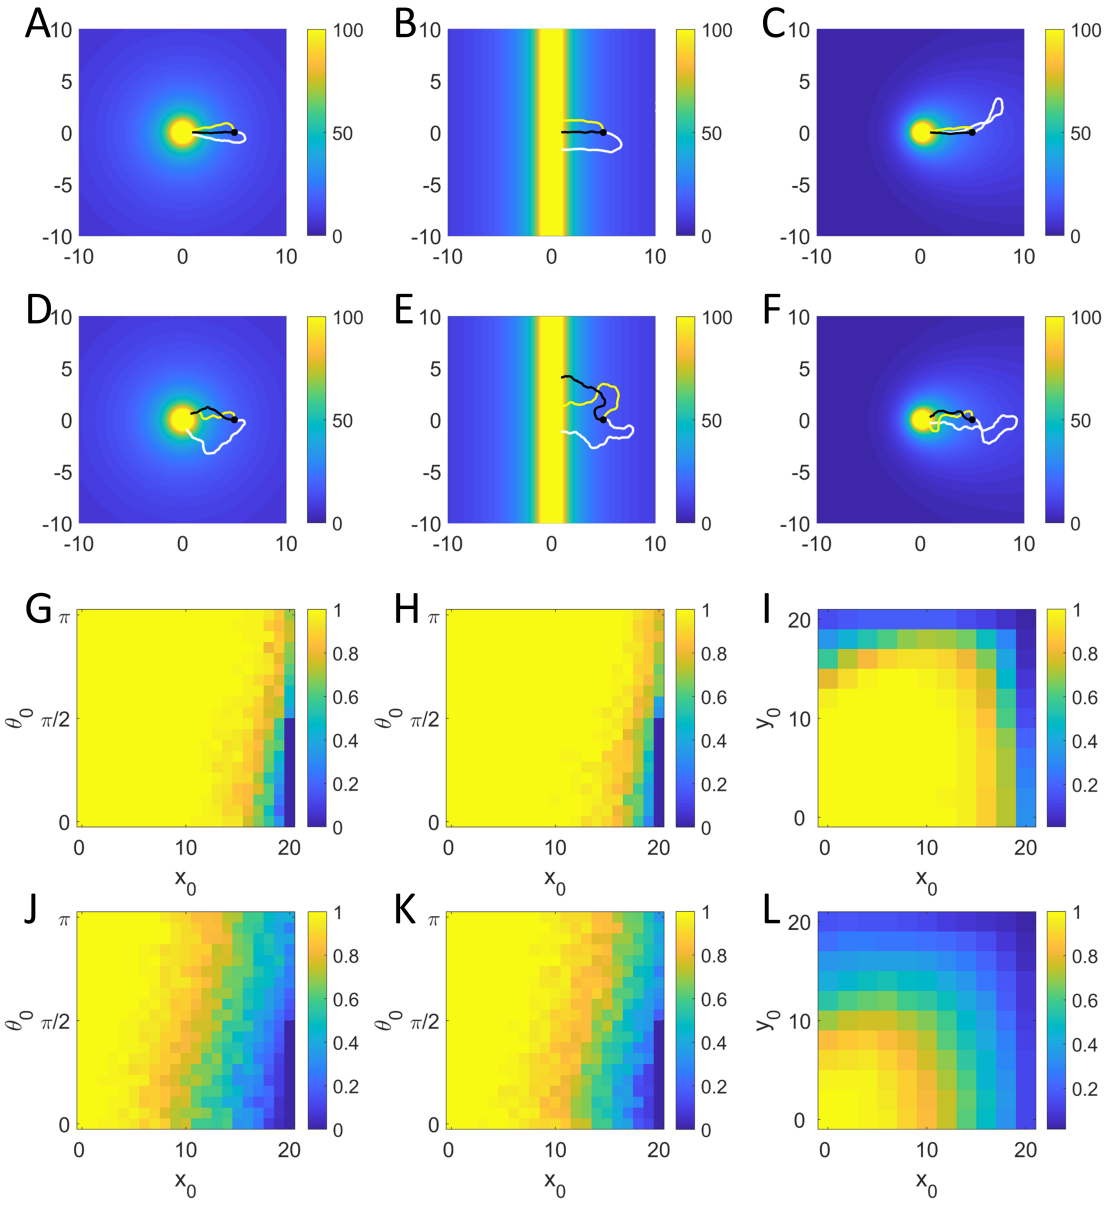

Figure 6: **Navigators successfully travel up gradients despite turbulent noise.** Tropotactic trajectories ($x_0, y_0 = 5, 0$; $\theta_0 = 0$ white line, $\theta_0 = \pi/2$ yellow line, $\theta_0 = \pi$ black line) are shown for the stochastic point source (A), trail (B), and plume (C). Klinotactic trajectories with the same initial conditions are shown for the point source (D), trail (E), and plume (F). Fraction of successful trials for tropotactic initial conditions are shown for the point source (G), trail (H), and plume (I). The fraction of successful trials for klinotactic initial conditions are shown for the point source (J), trail (K), and plume (L). Point source and trail success rate plots show success at initial distance x_0 and initial heading θ_0. Plume fraction of successful trial plots show spatial locations (x_0, y_0, see Figure 3c) and success rates averaged over θ_0 at each location. All fractional success rates are estimated from 100 trials per initial condition x_0, θ_0 (G,H,J,K) or x_0, y_0, θ_0 (I,L).

tropotaxis is more robust against stochastic odor inputs than klinotaxis. Tropotactic success rates averaged over initial conditions are reduced only slightly for each stochastic environment relative to the smooth gradient environment: point source stochastic vs. smooth success 0.92 vs. 0.93, trail 0.93 vs. 0.93, plume 0.76 vs. 0.80. When compared to tropotaxis, klinotactic success rates averaged over initial conditions show notable reductions in stochastic environments relative to smooth gradient environments: point source stochastic vs. smooth success 0.73 vs. 0.97, trail 0.79 vs. 0.97, plume 0.50 vs. 0.84. The greater impact of stochasticity on klinotaxis compared to tropotaxis may be attributed to differences in the steering functions. The tropotaxis navigator turns at a rate proportional to the concentration difference across its two sensors (eqn 4); if environmental noise leads it to incorrectly turn away from the source, it will only turn by an angle proportional to the small concentration fluctuations that induced the incorrect course change. The klinotactic navigator is governed by a heading update strategy (eqn. 7) that only exploits increases in concentration (*i.e.*, it only turns when a newly sampled concentration is higher than the previous). In noisy conditions it may therefore turn in the incorrect direction (away from the true maximum) by a ϕ angle that is not proportional to the concentration difference because it does not utilize multiple simultaneous samples. Even in stochastic conditions, both navigators outperformed the CRW agent in point source, trail, and plume environments (see **Supplemental Figure S1**.)

We evaluated the navigators' trail following performance in stochastic environments. Example trajectories for both navigators are qualitative similar in smooth gradient and stochastic environments (**Figure 5a,e**). Tropotactic performance is not substantially changed by the stochastic environment. Averaged over initial headings $\theta_0 = 0\text{-}\pi/2$, the tropotactic navigator travels a total length of 997.88 ± 0.07 distance units along the trail. This is a slight improvement over smooth gradient tropotactic performance (949.54 ± 0 distance units) because noise in the environment corrects the failure state at $\theta_0 = 0$ (**Figure 5b,f**, red lines). The maximum excursion of tropotaxis along the trail is also similar in both environments: 997.86 ± 0.09 and 949.54 ± 0 distance units for stochastic and smooth gradients, respectively (**Figure 5c,g**, red lines). This is accounted for by the low number of reversals made by the tropotactic navigator (≈ 0) for both stochastic and smooth gradient environments (**Figure 5d,h**, red lines). The klinotactic navigator also travels a similar total distance along the trail in smooth gradient and stochastic environments (928.80 ± 3.20 and 911.68 ± 8.33 distance units, respectively; **Figure 5b,f**, blue lines). However, the klinotactic navigator's maximum excursion along the trail is considerably longer in the smooth gradient environment than the stochastic environment (620.40 ± 229.35 and 340.24 ± 146.34 distance units, respectively; **Figure 5c,g**, black lines). This difference can be attributed to the increase in klinotactic trail reversals per trial in the stochastic environment (32.8 ± 14.2) compared to the smooth gradient (3.7 ± 2.7). The high number of reversals made by the klinotactic navigator suggest that

klinotaxis alone may not be sufficient for navigating along trails in turbulent odor environments; the use of additional sensory modalities (*e.g.,* visual landmarks) or cognitive maps may be required to navigate along a trail without reversing direction. Alternatively, animals may use a concentration gradient along the trail to determine the direction of travel. Real trails are deposited and dissipate over time, creating a gradient along the trail in the direction of travel of the trail-laying agent. Dogs have been shown to preferentially follow trails up their odor gradient, following the agent laying down the trail [43].

3.3 Closely-spaced sensors favor klinotaxis and widely-spaced sensors favor tropotaxis in stochastic odor environments

With some exceptions, vertebrates have closely-spaced nares relative to their body size [15]. Many vertebrates also have flexible necks that may facilitate the movement of the head (and nose) in klinotactic searches [44]. Conversely, mature arthropods may have antennae that are widely-spaced relative to the body size of the animal and branched to increase sensory surface area [45]. We wished to determine how inter-sensor distance (represented in our navigators by the inter-sensor angle γ) affected the success rates of tropotaxis and klinotaxis. To test this, we performed a parameter sweep of γ from 0 (representing two sensors directly ahead of the navigator) to $\pi/2$ (representing two sensors perpendicular to the body of the navigator). All other parameters for the tropotactic navigator ($\beta = 10, \ell = 1, V = 1$) and klinotactic navigator ($\phi_{\max} = \pi/2, \alpha = 1, \ell = 1, V = 1$) were held constant. For each value of γ, we evaluated both navigators' success across initial conditions x_0, θ_0 (point source and trail) or x_0, y_0, θ_0 (plume) as shown in **Figure 3a-c**. Results for each γ value are reported as the average success across all initial conditions.

In smooth gradient odor environments, we find klinotaxis outperforms tropotaxis at the task of locating sources in the parameter regime where γ is narrower than $\pi/2$ in point source and trails environments and when γ is narrower than $\pi/4$ in the plume environment (**Figure 7a-c**, blue lines). For angles equal to $\pi/2$, klinotactic performance is substantially reduced. This is because the klinotactic navigator is using an average concentration value across its two sensors (eqn. 6) and the klinotactic navigation strategy relies on this average represents a point estimate in front of the navigator. (It "follows its nose," and therefore the nose cannot be sampling at or behind the body.) When γ is $\pi/2$, the two sensors are located directly lateral to the body; the point estimate is at the location of the body rather than in front of it. When γ is greater than $\pi/2$ (not shown), the sensors and the point estimate falls behind the body, disrupting navigation. Unlike klinotaxis, tropotaxis relies on the difference between its two sensors to guide its heading update (eqn. 4). It turns the fastest when the relative difference between its concentration samples is the greatest, so average success is maximized when the distance between sensors is maximized at $\pi/2$ (**Figure**

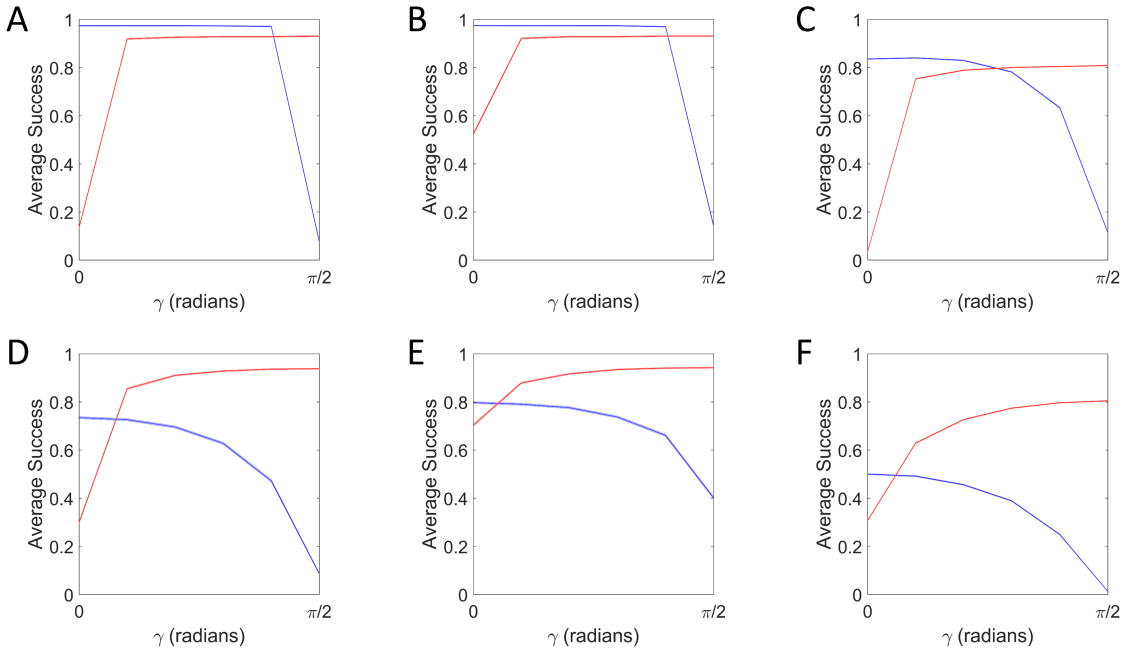

Figure 7: **Klinotaxis and tropotaxis have different optimal inter-sensor angles.** Success rates averaged over initial conditions are plotted versus inter-sensor angle γ for the smooth gradient point source (A), trail (B), and plume (C). Success rates averaged over initial conditions are plotted versus inter-sensor angle γ for the stochastic point source (D), trail (E), and plume (F). Blue lines represent klinotactic average success rates and red lines represent tropotactic average success rates. Shaded regions denote 95% Clopper-Pearson confidence intervals[46].

7a-c, red lines). Tropotaxis does not depend on sampling in front of the body; so long as there is some separation between the sensors (γ values greater than 0), the tropotactic navigator can locate sources. Note that success in **Figure 7b,e** is defined as reaching the trail, not following it. As expected based on findings by Amorim *et al.*[40] and Riman *et al.*[37] in their tropotactic models of trail following, tropotaxis fails at trail following when $\gamma \geq \pi/2$. This observation, coupled with the rarity of insects with antennae angles greater than $\pi/2$, suggests that intersensor angles greater than $\pi/2$ may not be advantageous in ecologically valid settings.

In stochastic odor environments, we find that klinotactic performance is globally reduced relative to smooth gradient environments across the range of γ values (**Figure 7d-f**, blue lines). This is in line with observations in **Section 3.2** of reduced klinotactic performance under stochastic conditions. Also consistent with observations in **Section 3.2** is the robustness of tropotaxis to stochastic environments. The success of the tropotactic navigator is not substantially reduced in stochastic environments across the range of γ values (**Figure 7d-f**, red lines).

Despite the decrease in klinotactic success in the stochastic environment, it continues to outperform tropotaxis for small values of γ. This reflects the opposing functions of paired sensors in klinotaxis and tropotaxis. In klinotaxis, the navigator averages concentrations across its two sensors to obtain an estimate of the concentration at the midpoint between the sensors. When sensors

are closely-spaced (small γ), the resulting average consists of two independent samples acquired near the midpoint of the sensors, providing a more accurate estimate by improving signal-to-noise. Conversely, widely-spaced (large γ) sensors sample further away from the midpoint of the nares, introducing error in the presence of noise. In tropotaxis, the navigator uses the difference of concentration samples across its two sensors instead of the average. Rather than favoring closely-spaced sensors (small γ), tropotaxis performs optimally when the distance between sensors (and therefore the difference between left-right concentration values) is large (large γ).

The closely-spaced sensor (small γ) parameter regime in one in which many vertebrates appear to operate [15]. In Liu *et al.*[36], the γ value for a parameterized model of mouse olfactory navigation was estimated to be approximately 0.02 radians, suggesting that the closely-spaced nares of mice favor klinotactic strategies.

3.4 Stereo information improves klinotaxis performance by increasing the signal-to-noise ratio

Louis *et al.* demonstrated that the accuracy of larval chemotaxis in the fruit fly *Drosophila melanogaster* was reduced in animals with either the left or right olfactory sensor genetically inactivated [47]. The larvae use a klinotactic strategy to navigate and improve the signal-to-noise ratio of their concentration estimates by using two bilateral sensors (as in eqn. 6). To determine if our klinotactic navigator was taking advantage of improved signal-to-noise by averaging over its two sensors, we compared success rates using both unoccluded sensors and with one occluded sensor ($C_{\mathrm{R}} := 0$ for all samples). The klinotactic navigator parameters were set to $\gamma = \pi/16, \phi_{\max} = \pi/2, \alpha = 1, \ell = 1, V = 1$. We evaluated both the occluded- and unoccluded-sensor klinotactic navigators' success across initial conditions x_0, θ_0 (point source and trail) or x_0, y_0, θ_0 (plume) as shown in **Figure 3a-c**.

When navigating smooth gradient point source and trail environments, no notable difference in the performance of occluded (red dashed lines) versus unoccluded (blue lines) klinotactic navigators is evident (**Figure 8a-b**). In the absence of stochastic fluctuations in the environment, there is not increase in signal-to-noise - and thus no advantage - conferred by using two sensors rather than one. Slight differences in performance are noted in the smooth gradient plume (**Figure 8c**). This is due to the asymmetric boundary conditions of the plume (**Figure 3c**). Occluding the right sensor introduces a slight leftward-bias in klinotactic steering decisions. When boundary conditions are symmetric, such as in the point source and the trail (**Figure 3a-b**), the effects are negligible. In the case of the plume, they introduce small changes to the success rate.

In contrast to the smooth gradient case, the unoccluded (blue lines) klinotactic navigators outperform occluded (red dashed lines) navigators in the stochastic environments (**Figure 8d-f**). As

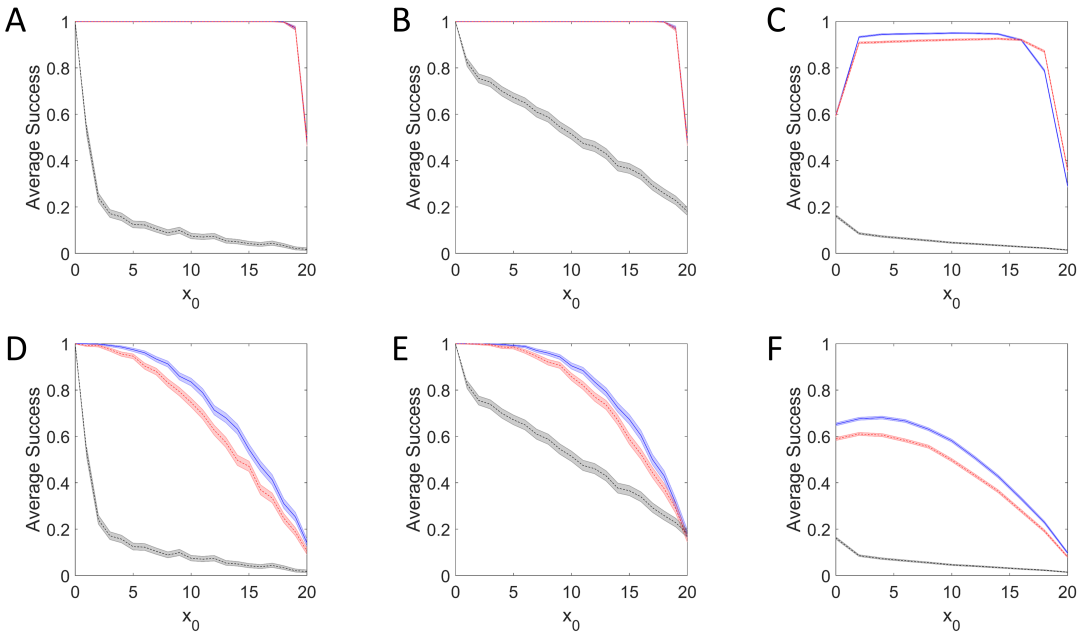

Figure 8: **Nares occlusion impairs klinotactic navigator performance.** The average success rate versus distance for klinotactic navigators with unoccluded (blue line) and one occluded (red dashed line) sensor are shown for smooth gradient point source (A), trail (B), and plume (C). The average success rate versus distance for unoccluded (blue line) and occluded (red dashed line) klinotactic navigators are shown for stochastic point source (D), trail (E), and plume (F). Success rates for the correlated random walk agent (black dashed line) are shown for point source (A,D), trail (B,E), and plume (C,F) environments. Point source and trail success rates (A,B,D,E) are averaged over θ_0; plume success rates (C,F) are averaged over θ_0 and y_0. Shaded regions denote 95% Clopper-Pearson confidence intervals.

discussed in **Section 3.2**, the klinotactic steering function (eqn. 7) exhibits reduced performance in stochastic environments. Averaging over two sensors therefore provides an advantage by increasing the signal-to-noise and reducing the probability of incorrectly turning away from the odor source. This suggests that at least one use of bilateral sensors is to improve the signal-to-noise ratio in stochastic odor environments. In the next section, we consider an additional use: improving klinotaxis with stereo information.

3.5 Stereo information improves the hybrid strategy's performance by providing left-right directional cues

Catania's observations of the blind eastern American mole (*Scalopus aquaticus*) led him to propose a hybrid strategy for olfactory searches [27]. In this strategy, moles use klinotaxis to orient on an odor source from a distance. Because moles' nares (sensors) are close together and the odor gradient is relatively flat far from the source, moles cannot use tropotactic cues to orient at long distances. However, when moles are close to the source and the gradient is steep, moles can resolve left-right nares differences and use them to inform nose movements. We represent this strategy in our hybrid navigator by introducing a left-right bias in nose deflection (eqns. 11-12) when the concentration difference across the sensors is large. These left-right biases are illustrated in **Figure 1e-g**. Parameter screens suggest that the hybrid strategy significantly outperforms klinotaxis when the beta distribution parameter α is greater than one. This parameter regime represents a unimodal nose sampling distribution that is consistent with observations of mice (Urban unpublished data) and may generalize to other mammalian species that use klinotaxis. Accordingly, we alter our klinotactic parameter set so that $\alpha = 10$. All other parameters remain unchanged ($\gamma = \pi/16, \phi_{\max} = \pi/2, \ell = 1, V = 1$).

It should be noted that the klinotactic navigator performance is lower in this unimodal sampling ($\alpha = 10$) parameter regime compared to the previously analyzed uniform sampling ($\alpha = 1$) regime, even in smooth gradient environments (**Figure 9a-c**, blue lines, *cf.* **Figure 8a-c**, blue lines). This is because the unimodal nose deflection distribution limits the turning angle of the klinotactic navigator, increasing its turning radius and course correction time. As such, it is more likely to encounter a failure boundary before orienting toward the odor source. In the unimodal sampling regime, the klinotactic navigator fares even worse in stochastic environments, where environmental noise may lead to erroneous steering away from the source (**Figure 9d-f**, blue lines, *cf.* **Figure 8d-f**, blue lines). Even though performance is reduced when $\alpha = 10$, the navigator still outperforms the CRW agent (**Figure 9d-f**, black lines).

Compared to klinotactic performance alone in the $\alpha = 10$ parameter regime, employing the hybrid strategy improved success rates in some - but not all - odor environments. In smooth

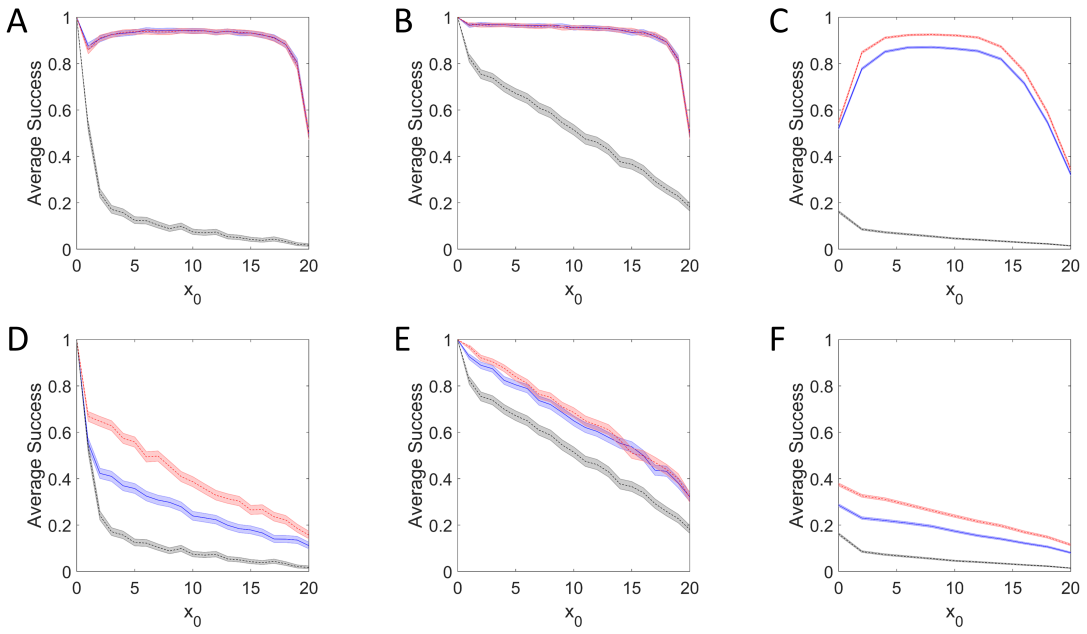

Figure 9: **A hybrid strategy outperforms klinotaxis alone.** The average success rate versus distance for klinotactic (blue line) and hybrid (red dashed line) navigators are shown for smooth gradient point source (A), trail (B), and plume (C). The average success rate versus distance for klinotactic (blue line) and hybrid (red dashed line) navigators are shown for stochastic point source (D), trail (E), and plume (F). Success rates for the correlated random walk agent (black dashed line) are shown for point source (A,D), trail (B,E), and plume (C,F) environments. Point source and trail success rates (A,B,D,E) are averaged over θ_0; plume success rates (C,F) are averaged over θ_0 and y_0. Shaded regions denote 95% Clopper-Pearson confidence intervals.

gradient environments, only plume navigation was improved by the incorporation of stereo information (**Figure 9a-c**, red lines). This may be attributed to the asymmetric boundaries and initial conditions of navigators in the plume (**Figure 3c**); across navigator starting positions, the hybrid strategy introduces a turning bias toward the plume midline and away from the boundaries at $y = \pm 20$. In contrast, the point source and trail environments are symmetric in the y-direction and do not benefit in this regard.

When odor is stochastic, the hybrid strategy outperforms klinotaxis in all three environments, although only significantly in point sources and plumes (**Figure 9d-f**, red lines). As discussed in **Section 3.2**, the nature of the klinotactic steering function can lead the navigator off-target in noisy environments. This effect has a large impact where the standard deviation of the particle detection events are high (near the source). The hybrid strategy appears to correct for this effect, preventing "near misses" when the navigator is near the source. It does this by biasing the nose deflection in the direction of the source and minimizing the chance that the navigator incorrectly turns away from the source. This also explains why the hybrid strategy does not significantly improve performance in the trail environment. Unlike point source and plume capture areas, which are circles of fixed circumference (**Figure 3a,c**, green lines), the capture boundary for trails is infinite (**Figure 3b**, green line). Klinotaxis alone is sufficient to navigate up the gradient, and the infinite nature of the boundary precludes near-misses that would be possible with finite capture boundaries.

3.6 Occluding or reversing stereo inputs disrupts the performance of the hybrid strategy

To demonstrate that moles use stereo information in their olfactory searches, Catania performed experiments in which he blocked one naris (sensor) and used nasal tubes to reverse the left-right olfactory inputs [27]. In the experiment with the occluded naris, he observed a bias in nose movement toward the side of the open naris (see Catania Figure 2). In the nares reversal experiment, he observed repulsion of the nose from the odor source (see Catania Figure 4). Together, he suggested that these observations were consistent with a hybrid search strategy using klinotaxis at a distance and incorporating stereo information near the source. We performed equivalent virtual experiments with our hybrid navigator to determine if it behaved similarly to the moles. Using the parameter values from **Section 3.5** ($\gamma = \pi/16, \phi_{\max} = \pi/2, \alpha = 10, \ell = 1, V = 1$) we occluded one naris (setting $C_{\mathrm{L}} := 0$ or $C_{\mathrm{R}} := 0$) or reversed nares inputs (switching α_{L} and α_{R}). To mimic Catania's experiment, we started the navigators near the point source at $x_0, y_0 = 0, 10$ with a heading pointing toward the source, $\theta_0 = 3\pi/2$. Simulations were allowed to run until the found the source or reached the failure boundary. The stochastic point source was used to simulate the

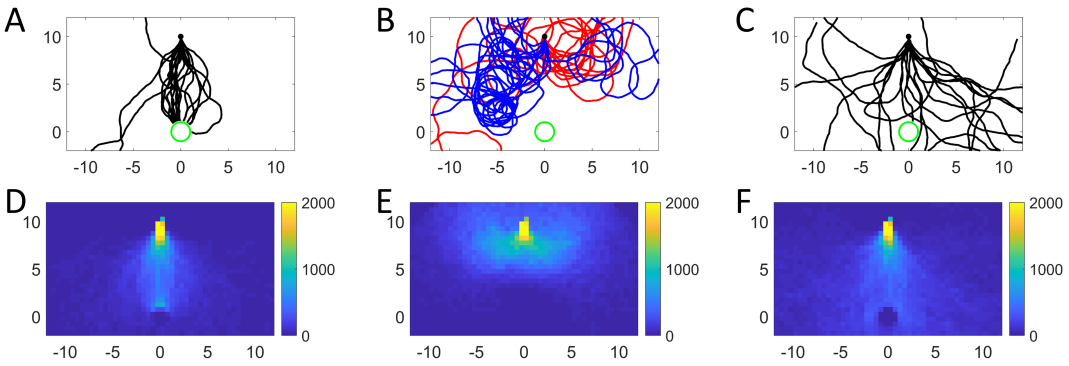

Figure 10: **Occlusion and reversal of hybrid sensors mimics observations in mammals.** Twenty trajectories (black lines) of the hybrid navigator are shown in (A); ten trajectories of the hybrid navigator with the right sensor occluded (red lines) and ten trajectories with the left sensor occluded (blue lines) are shown in (B); twenty trajectories of the hybrid navigator with reversed sensors (black) are shown in (C). All trajectories are initialized at $x_0, y_0, \theta_0 = 10, 0, 3\pi/2$ (black dots). The point source is located at the origin with a capture radius of 1 (green circles). Occupancy histograms for the hybrid navigator (D), occluded nares (E), and reversed nares (F) are constructed from one thousand trajectories each.

odor environment.

When neither naris was occluded, the hybrid navigator was capable of acquiring the source (20 sample trajectories in **Figure 10a**, occupancy histogram in **d**). In line with mole observations, occlusion of the left naris (**Figure 10b**, blue lines) or right naris (**Figure 10b**, red lines) led to biased nose motion in the direction of the open naris (occupancy histogram in **Figure 10e**). The navigator trajectories also showed increased tortuosity qualitatively consistent with Catania's observations. Finally, reversal of the left-right olfactory inputs led to repulsion from the odor source (20 sample trajectories in **Figure 10c**, occupancy histogram in **f**), qualitatively consistent with crossed-nares mole data. Taken together, this suggests that the hybrid strategy is consistent with behavior observed in moles. It may be a strategy that generalizes to other species such as mice as suggested by Liu *et al.*[36].

3.7 Both strategies successfully navigate real odor plumes

The navigators perform better than chance in /bluetime-invariant smooth gradient environments and stochastic environments. However, neither of these environments captures the full complexity of true turbulent odor environments. To evaluate the performance of the navigators in true turbulent environments, we use planar laser induced fluorescence data from Connor *et al.* as a surrogate for odorant concentration data[5]. We tested the navigators in two environmental regimes: near-bed flow at 10 cm/s (representative of an animal navigating near a surface such as the ground) and free-stream flow at 20 cm/s (representative of an animal navigating in open space). Model success was evaluated over a range of initial conditions x_0, y_0, θ_0 (as shown in **Figure 3d**) in each

environmental regime.

Three example trajectories ($x_0, y_0 = 10, 0; \theta_0 = 0, \pi/2, \pi$) for the tropotactic navigator in the near-bed and free-stream plumes are shown in **Figure 11a-b**. Here, parameter values are the same as in **Figures 4** and **6** ($\beta = 10, \gamma = \pi/4$, $\ell = 1$ cm, $V = 1$ cm/s). Note that the formerly unspecified units of length and velocity are now scaled to the PLIF data scale of centimeters and seconds. The tropotactic navigator is capable of locating the source regardless of initial heading. Interestingly, the smooth gradient of the near-bed plume (**Figure 2j**) leads to smaller differences between tropotactic sensors and lower rates of turning. This is illustrated by the long excursion of the tropotactic navigator oriented away from the source before it successfully turns to the source (**Figure 11a**, white line). Conversely, the intermittent regions of high and low concentration in the free-stream plume (**Figure 2k**) lead to larger concentration differences across the sensors. This facilitates faster turning (**Figure 11b**, white line), but overall more tortuous trajectories (*e.g.*, **Figure 11b**, yellow line). Systematic evaluation of tropotactic success rates across initial conditions x_0, y_0, θ_0 reveals that the tropotactic navigator performs near 100% across the near-bed plume (**Figure 11e**). Performance is comparable near the midline of the free-stream plume, but the narrower and intermittent plume structure reduces tropotactic performance at the periphery of the plume (**Figure 11f**).

Three example trajectories ($x_0, y_0 = 10, 0; \theta_0 = 0, \pi/2, \pi$) for the klinotactic navigator in the near-bed and free-stream plumes are shown in **Figure 11c-d**. Here, parameter values are the same as in **Figures 4** and **6** ($\phi_{\max} = \pi/2, \alpha = 1, \gamma = \pi/16$, $\ell = 1$ cm, $V = 1$ cm/s). As noted in **Sections 3.1-3.2**, the steering mechanism of the klinotactic navigator leads to more tortuous trajectories when compared to tropotaxis. Regardless, klinotaxis is capable of navigating up the broad plume of the near-bed plume to locate the source (**Figure 11c**). It is less successful in the free-stream plume, where large regions of the plume periphery lack concentration gradients that can be used for reliable navigation (**Figure 11d**). Once the klinotactic navigator leaves the narrow region located at the plume midline, it embarks on meandering excursions away from the plume (*e.g.*, the white trajectory in **Figure 11d**). It can only successfully locate the source if it stays in the plume, or, if it previously left the plume, reacquires the plume. These examples are consistent with systematic evaluation of navigator performance across initial conditions x_0, y_0, θ_0. The broad plume of the near-bed case facilitates successful klinotactic navigation wherever there is a consistent gradient (**Figure 11g**). The narrow plume of the free-stream case leads to a substantial global reduction in performance relative to the near-bed case (**Figure 11h**).

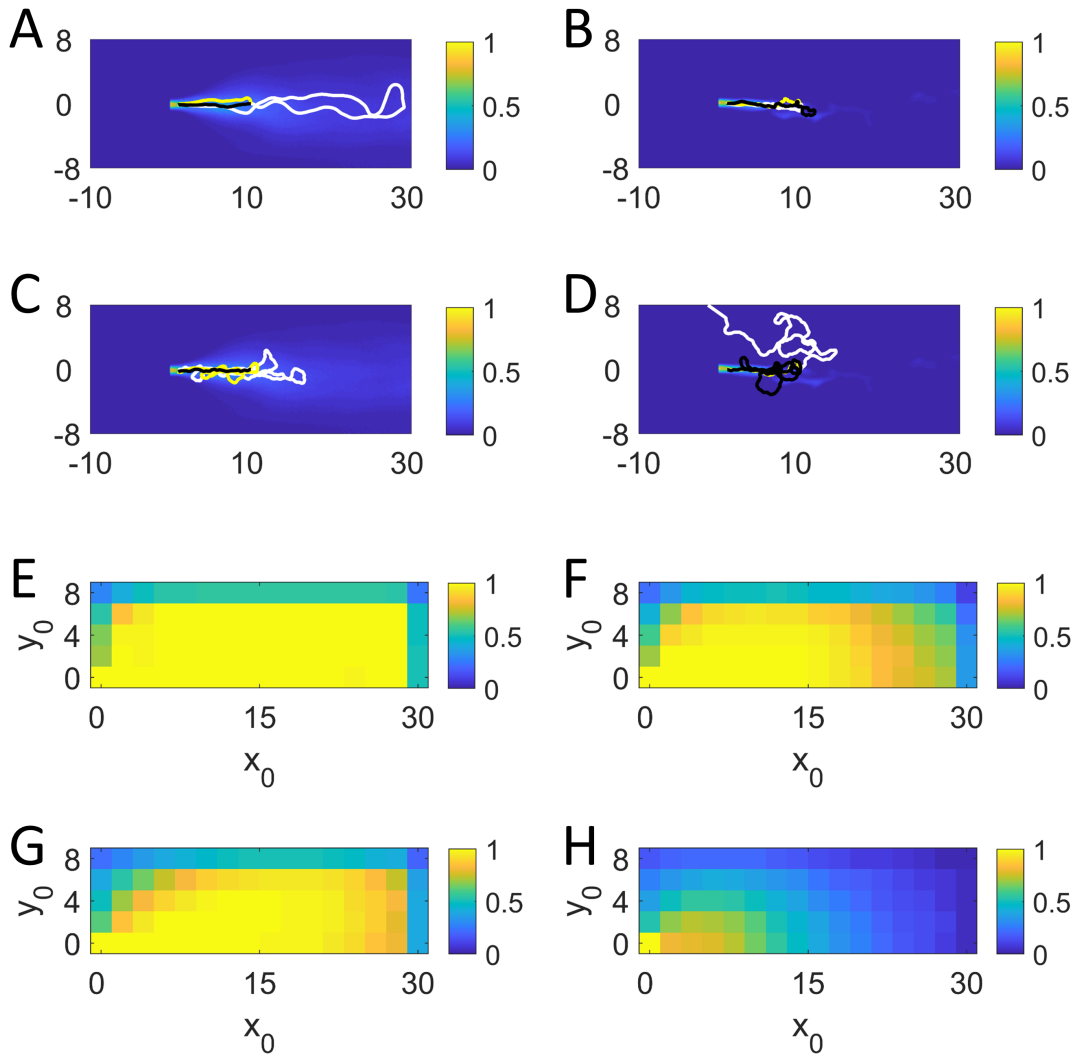

Figure 11: **Navigator success depends on intermittency of turbulent PLIF plumes.** Tropotactic trajectories ($x_0, y_0 = 10, 0$; $\theta_0 = 0$ white line, $\theta_0 = \pi/2$ yellow line, $\theta_0 = \pi$ black line) are shown for the near-bed plume (A) and free-stream plume (B). Klinotactic trajectories with the same initial conditions are shown for the near-bed plume (C) and free-stream plume (D). Fraction of successful trials for tropotactic initial conditions are shown for the near-bed (E) and free-stream (F) plumes. Fraction of successful trials for klinotactic initial conditions are shown for the near-bed (G) and free-stream (H) plumes. Plume fraction of successful trial plots show spatial locations (x_0, y_0, see Figure 3d) and success rates averaged over θ_0 at each location. All fractional success rates are estimated from 100 trials per initial condition x_0, y_0, θ_0. Note that for these simulations, $\ell = 1$cm, $V = 1$cm/s, and $\Delta t = 0.1$s.

4 Discussion

Olfaction is a sensory modality that guides navigation across the animal kingdom. Animals' olfactory environments often exhibit turbulent flow, leading to odor distributions which are dynamic, noisy, and intermittent. Intuitively, such landscapes can be disruptive to navigation strategies that rely on gradient estimation. Despite this apparent challenge, animals have been observed to use spatial and spatiotemporal concentration comparisons to navigate through their turbulent environments. One key challenge of studying these navigation strategies in behaving animals is simultaneously characterizing the dynamic odor environment in which they are occurring.

Modeling allows us to control features of the odor distribution and observe how different olfactory navigation strategies behave under increasingly realistic turbulent conditions. In this study, we develop general models of two gradient-driven navigation algorithms used by animals: tropotaxis and klinotaxis. We examine the performance of each strategy in smooth gradient, stochastic, and turbulent odor environments. Tropotaxis and klinotaxis performed well in smooth gradient environments where concentration comparisons were noise-free - both had near-100% success rates across the point source, trail, and plume domains where they were tested. A stochastic odor environment adversely impacted performance of both strategies, though klinotaxis, with its spatiotemporal concentration-based steering, fared worse than tropotaxis. Notably, both strategies performed significantly better than chance (represented by correlated random walks) in the presence of odor noise, indicating that these gradient-driven strategies remain beneficial even in stochastic odor environments. Finally, we tested the strategies in actual turbulent flow conditions using PLIF plume data. Both strategies performed well in the near-bed plume, indicating that both strategies are appropriate for navigating along flat surfaces such as the ground or seabed. Klinotaxis alone performed poorly (though still better than chance) in the free-stream plume, indicating that it may be inefficient for tracking air- or water-borne odorants far from surfaces.

Modeling also allows us to study the effect of olfactory sensor geometry on each strategy's performance. Different animals have vastly different sensor morphologies: arthropods have chemoreceptors on widely-spaced antennae, while mammals have nares that are typically located close together (*e.g.,* the ant and rat in **Figure 1b,d**). How do these different geometries affect the performance of different olfactory navigation strategies? We find that widely-spaced sensors favor tropotaxis while closely-spaced sensors favor klinotaxis. Widely-spaced sensors can sample two spatially-separated points on an odor gradient; in noisy environments, they are more likely to detect the underlying signal (*i.e.*, the direction of the gradient). Conversely, when sensors are closely-spaced in the same gradient the concentration difference between them will be smaller; detection of the signal becomes difficult in noisy environments. In this sense, klinotaxis may be viewed as solution to the "problem" of closely-spaced sensors operating in noisy environments:

rather than comparing concentrations across sensors, animals move their sensors through space to increase the effective distance between concentration comparisons and increase the likelihood of detecting the direction of the gradient in turbulent environments. While closely-spaced sensors may have trouble resolving the direction of a source where the gradient is relatively flat, the gradient is typically steeper near the source. Here, even closely-spaced sensors may be able to reliably detect the direction of the gradient. Animals that rely on klinotaxis for navigation far from a source may still make use of this additional stereo information when they approach a source. Catania provided a qualitative description of this form of hybrid strategy based on findings in moles. Here, we developed a general hybrid model of klinotaxis utilizing stereo information and found it it qualitatively recapitulated Catania's observations. This suggests that klinotactic mammals with closely-spaced nares may still utilize limited tropotactic information to guide navigation near odor sources.

Our modeling approach allowed us to observe the effects of odor environment and body geometry on olfactory navigation performance, but it is important to note the limitations of this study. Our tropotactic and klinotactic models represent abstractions of each behavior and are not closely representative of any particular species' behavior. For simplicity, we fix the velocity of navigators though some animals have been observed to alter their velocity profiles upon approach to an odor source[36]. In the case of klinotaxis, the stochastic nose deflection (eqn. 10) is a simplification of animal movement that may not always be physiologically realistic. (*I.e.*, the nose may sometimes "jump" between two spatially-distant positions while sampling the environment.) When evaluating our navigator performance on different types of odor environments, we chose to use the Poisson-distributed odor model presented in Vergassola *et al.*[34] to represent turbulence. While this model allows for rapid simulation of noisy odor environments, it was designed to represent the dilute limit when odor detection events are infrequent. As such, it lacks the spatial correlation structure that would be present in direct numerical simulations of turbulent environments. Finally, our navigators operate without any kind of cognitive map of their environment. Unlike models such as infotaxis or entrotaxis, which use maps to operate in dilute conditions, our minimal navigation strategies are only useful in odor environments where the odor is consistently above navigator's detection thresholds.

Several future avenues of research may address current limitations of this study. Navigator velocities may be modulated as a function of concentration to optimize performance. Decreasing the forward velocity relative to the turning angle as a navigator approaches an odor source allows for sharper turns and may lead to improved success rates. This deceleration has been observed in mice [36] and explored theoretically in models of ant pheromone tracking[40]. More realistic models of lateral nose movement than the beta-distributed motion described here, such as acceleration-based approaches that prevent "jumps" in nose position[48] or Ornstein-Uhlenbeck processes[36],

may yield better agreement with experimental data from various species. We chose simple steering functions for ease of analysis, but a survey of more complex proportional turning schemes may improve performance, especially when compared to the all-or-nothing nature of the klinotactic navigator. While organisms may use different strategies for 3D navigation, a natural extension of this work is consideration of navigation in three dimensions with and without anemometric steering inputs. We currently lack full 3D turbulent plume data for use as inputs in 3D navigation but that may be addressed by the following approach. Direct numerical simulation (DNS) of the Navier-Stokes equations may be used in place of the simple Poisson-distributed odor hit model, potentially improving agreement between navigators using inputs with simulated turbulence versus true turbulent data in 2D and 3D and allowing for exploration of diverse flow conditions. As noted, animals may "stop and sniff" to accumulate data; this may be used in future models that employ higher order spatial derivatives to model the odor environment. Navigating animals have been shown to switch strategies between active olfactory search and reliance on cognitive maps[11]; incorporation of simple maps of the odor landscape could allow these navigator models to generalize to dilute odor environments in addition to the supra-threshold regime in which they currently operate.

In summary, we show that gradient-turning klinotactic and tropotactic olfactory search strategies are capable of locating odor sources and following odor trails. In addition, both strategies are robust against stochasticity and turbulent fluctuations in odor inputs. We further show that the optimal strategy to use depends on the navigator's body geometry, with closely-spaced sensors favoring klinotaxis and widely-spaced sensors favoring tropotaxis. This finding suggests that many vertebrates (with closely-spaced nares relative to body size) will utilize klinotaxis, while arthropods (with widely-spaced antennae relative to body size) will favor tropotaxis. A hybrid model combining klinotaxis with limited tropotactic information has been observed in some mammals, but has not previously been characterized analytically. Here, we show that such a hybrid strategy can outperform klinotaxis alone and may be used by even those vertebrates with closely-spaced nares to improve olfactory search performance. Future work will be needed to adapt these very general models if one wishes to predict species-specific olfactory search behaviors.

Acknowledgements

We thank Andrew Papale, Nour Riman, and Annie Liu for helpful discussions while developing models and revising the manuscript.

Funding

This work was supported by the National Science Foundation (NSF PHY 1555862, NSF PHY 1555916).

Declaration of Competing Interest

The authors declare no conflicts of interest.

CRediT Attribution

JBH: Conceptualization, Methodology, Software, Validation, Visualization, Writing - Original Draft, Writing - review & editing **EGC:** Conceptualization, Data Curation, Writing - Review & Editing **JPC:** Conceptualization, Data Curation, Writing - Review & Editing, Funding acquisition **NNU:** Conceptualization, Writing - Review & Editing, Funding acquisition **GBE:** Conceptualization, Methodology, Writing - Review & Editing, Supervision, Funding acquisition

References

[1] A. H. Dittman and T. P. Quinn, "Homing in pacific salmon: Mechanisms and ecological basis," *Journal of Experimental Biology*, vol. 199, no. 1, pp. 83–91, 1996.

[2] V. Fenton, "The use of dogs in search, rescue and recovery," *Journal of Wilderness Medicine*, vol. 3, pp. 292–300, Aug. 1992.

[3] N. J. Vickers, "Winging it: Moth Flight behavior and responses of Olfactory neurons are shaped by pheromone plume dynamics," *Chemical Senses*, vol. 31, pp. 155–166, Feb. 2006.

[4] L. F. Jacobs, "From chemotaxis to the cognitive map: The function of olfaction.," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 109 Suppl, pp. 10693–700, June 2012.

[5] E. G. Connor, M. K. McHugh, and J. P. Crimaldi, "Quantification of airborne odor plumes using planar laser-induced fluorescence," *Experiments in Fluids*, vol. 59, p. 137, Sept. 2018.

[6] J. P. Crimaldi, M. B. Wiley, and J. R. Koseff, "The relationship between mean and instantaneous structure in turbulent passive scalar plumes," *Journal of Turbulence*, vol. 3, p. N14, Jan. 2002.

[7] H. Martin, "Osmotropotaxis in the Honey-Bee," *Nature*, vol. 208, pp. 59 EP –, Oct. 1965.

[8] A. Borst and M. Heisenberg, "Osmotropotaxis in Drosophila melanogaster," *Journal of Comparative Physiology*, vol. 147, no. 4, pp. 479–484, 1982.

[9] K. Steck, M. Knaden, and B. S. Hansson, "Do desert ants smell the scenery in stereo?," *Animal Behaviour*, vol. 79, no. 4, pp. 939–945, 2010.

[10] S. B. Vander Wall, M. J. Beck, J. S. Briggs, J. K. Roth, T. C. Thayer, J. L. Hollander, and J. M. Armstrong, "Interspecific variation in the olfactory abilities of granivorous rodents," *Journal of Mammalogy*, vol. 84, pp. 487–496, May 2003.

[11] D. H. Gire, V. Kapoor, A. Arrighi-Allisan, A. Seminara, and V. N. Murthy, "Mice develop efficient strategies for foraging and navigation using complex natural stimuli," *Current Biology*, vol. 26, pp. 1261–1273, May 2016.

[12] J. M. Gardiner and J. Atema, "Sharks need the lateral line to locate odor sources: Rheotaxis and eddy chemotaxis.," *Journal of Experimental Biology*, vol. 210, pp. 1925–34, June 2007.

[13] F. Papi, "Olfactory navigation in birds," *Experientia*, vol. 46, pp. 352–363, Apr. 1990.

[14] H. G. Wallraff, "Avian olfactory navigation: Its empirical foundation and conceptual state," *Animal Behaviour*, vol. 67, pp. 189–204, Feb. 2004.

[15] D. M. Stoddart, "External nares and olfactory perception," *Experientia*, vol. 35, pp. 1456–1457, Nov. 1979.

[16] R. Stocker, "Marine microbes see a sea of gradients," *Science*, vol. 338, pp. 628–633, Nov. 2012.

[17] H. C. Berg, "Chemotaxis in Bacteria," *Annual Review of Biophysics and Bioengineering*, vol. 4, pp. 119–136, June 1975.

[18] P. B. Reeder and B. W. Ache, "Chemotaxis in the Florida spiny lobster, Panulirus argus," *Animal Behaviour*, vol. 28, no. 3, pp. 831–839, 1980.

[19] V. Calenbuhr and J. L. Deneubourg, "A model for osmotropotactic orientation (I)," *Journal of Theoretical Biology*, vol. 158, pp. 359–393, Oct. 1992.

[20] V. Calenbuhr, L. Chrétien, J. L. Deneubourg, and C. Detrain, "A model for osmotropotactic orientation (II)," *Journal of Theoretical Biology*, vol. 158, no. 3, pp. 395–407, 1992.

[21] P. B. Johnsen and J. H. Teeter, "Behavioral responses of bonnethead sharks (Sphyrna tiburo) to controlled olfactory stimulation," *Marine Behaviour and Physiology*, vol. 11, pp. 283–291, July 1985.

[22] D. Martinez, "Klinotaxis as a basic form of navigation," *Frontiers in Behavioral Neuroscience*, vol. 8, p. 275, Aug. 2014.

[23] Y. Iino and K. Yoshida, "Parallel use of two behavioral mechanisms for chemotaxis in Caenorhabditis elegans.," *Journal of Neuroscience*, vol. 29, pp. 5370–80, Apr. 2009.

[24] A. Gomez-Marin, G. J. Stephens, and M. Louis, "Active sampling and decision making in Drosophila chemotaxis.," *Nature communications*, vol. 2, p. 441, Aug. 2011.

[25] R. Mathewson and E. Hodgson, "Klinotaxis and rheotaxis in orientation of sharks toward chemical stimuli," *Comparative Biochemistry and Physiology Part A: Physiology*, vol. 42, pp. 79–84, May 1972.

[26] P. W. Jones and N. N. Urban, "Mice follow odor trails using stereo olfactory cues and rapid sniff to sniff comparisons," *bioRxiv*, p. 293746, Apr. 2018.

[27] K. C. Catania, "Stereo and serial sniffing guide navigation to an odour source in a mammal," *Nature Communications*, vol. 4, p. 1441, Dec. 2013.

[28] R. Rajan, J. P. Clement, and U. S. Bhalla, "Rats smell in stereo," *Science*, vol. 311, no. 5761, p. 666, 2006.

[29] Y. Wu, K. Chen, Y. Ye, T. Zhang, and W. Zhou, "Humans navigate with stereo olfaction," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 117, pp. 16065–16071, June 2020.

[30] N. J. Vickers, "Mechanisms of animal navigation in odor plumes," *The Biological Bulletin*, vol. 198, pp. 203–212, Apr. 2000.

[31] E. Álvarez-Salvado, A. M. Licata, E. G. Connor, M. K. McHugh, B. M. King, N. Stavropoulos, J. D. Victor, J. P. Crimaldi, and K. I. Nagel, "Elementary sensory-motor transformations underlying olfactory navigation in walking fruit-flies," *eLife*, vol. 7, aug 2018.

[32] X. Bao, E. Gjorgieva, L. K. Shanahan, J. D. Howard, T. Kahnt, and J. A. Gottfried, "Grid-like neural representations support olfactory navigation of a two-dimensional odor space," *Neuron*, Apr. 2019.

[33] U. Bhattacharyya and U. S. Bhalla, "Robust and rapid air-borne odor tracking without casting," *eNeuro*, vol. 2, no. 6, 2015.

[34] M. Vergassola, E. Villermaux, and B. I. Shraiman, "'Infotaxis' as a strategy for searching without gradients," *Nature*, vol. 445, pp. 406–409, Jan. 2007.

[35] M. Hutchinson, H. Oh, and W.-H. Chen, "Entrotaxis as a strategy for autonomous search and source reconstruction in turbulent conditions," *Information Fusion*, vol. 42, pp. 179–189, July 2018.

[36] A. Liu, A. E. Papale, J. Hengenius, K. Patel, B. Ermentrout, and N. N. Urban, "Mouse navigation strategies for odor source localization," *Frontiers in Neuroscience*, vol. 14, p. 218, Mar. 2020.

[37] N. Riman, J. D. Victor, S. D. Boie, and B. Ermentrout, "The dynamics of bilateral olfactory search and navigation," *arXiv*, Aug. 2020.

[38] A. Perna, B. Granovskiy, S. Garnier, S. C. Nicolis, M. Labédan, G. Theraulaz, V. Fourcassié, and D. J. Sumpter, "Individual rules for trail pattern formation in Argentine ants (linepithema humile)," *PLoS Computational Biology*, vol. 8, p. e1002592, July 2012.

[39] L. Barberis and F. Peruani, "Large-scale patterns in a minimal cognitive flocking model: incidental leaders, nematic patterns, and aggregates," *Physical Review Letters*, vol. 117, p. 248001, Dec. 2016.

[40] P. Amorim, T. Goudon, and F. Peruani, "An ant navigation model based on Weber's law," *Journal of Mathematical Biology*, vol. 78, pp. 943–984, Mar. 2019.

[41] S. Gorur-Shandilya, M. Demir, J. Long, D. A. Clark, and T. Emonet, "Olfactory receptor neurons use gain control and complementary kinetics to encode intermittent odorant stimuli," *eLife*, vol. 6, June 2017.

[42] J. Stockie, "The mathematics of atmospheric dispersion modeling," *SIAM Review*, vol. 53, no. 2, pp. 349–372, 2011.

[43] A. Thesen, J. B. Steen, and K. B. Døving, "Behaviour of dogs during olfactory tracking," *Journal of Experimental Biology*, vol. 180, no. 1, pp. 247–251, 1993.

[44] D. M. Stoddart, "Vertebrate olfaction," *Endeavour*, vol. 5, pp. 9–14, Jan. 1981.

[45] D. Schneider, "Insect antennae," *Annual Review of Entomology*, vol. 9, pp. 103–122, Jan. 1964.

[46] C. J. Clopper and E. S. Pearson, "The use of confidence or fiducial limits illustrated in the case of the binomial," *Biometrika*, vol. 26, p. 404, Dec. 1934.

[47] M. Louis, T. Huber, R. Benton, T. P. Sakmar, and L. B. Vosshall, "Bilateral olfactory sensory input enhances chemotaxis behavior," *Nature Neuroscience*, vol. 11, pp. 187–199, Feb. 2008.

[48] A. G. Khan, M. Sarangi, and U. S. Bhalla, "Rats track odour trails accurately using a multilayered strategy with near-optimal sampling," *Nature Communications*, vol. 3, p. 703, Jan. 2012.

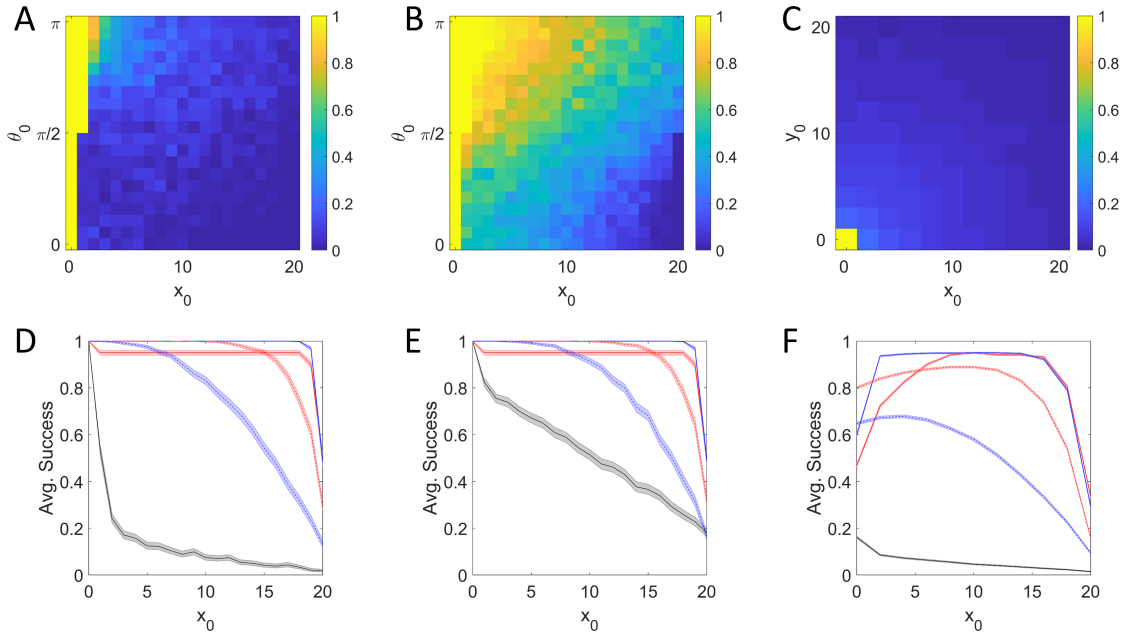

Figure S1: **Tropotaxis and klinotaxis outperform the CRW.** The CRW navigator was run across a grid of initial conditions as described in Figure 3. The fraction of successful trials for CRW at each initial condition is shown for the point source (A), trail (B), and plume (C). Average success is also displayed as a function of distance from the point source (D), trail (E), and plume (F). Solid red and blue lines respectively represent tropotaxis and klinotaxis in smooth gradient environments. Dashed red and blue lines respectively represent tropotaxis and klinotaxis in stochastic environments. Black lines represent the CRW. Point source and trail success rates (D,E) are averaged over θ_0; plume success rates (F) are averaged over θ_0 and y_0. Shaded regions denote 95% Clopper-Pearson confidence intervals.