

Figure S1, related to figure 1. Mouse performance improves with training. (A) Histograms of performance shown mouse-by-mouse (colormap) as well as summed by group (bar plots). (B) Histogram of the performance for all trials summarized in Figure 1C. (C) Average performance for each mouse across the sequence of training (same data as that which was averaged over all mice and shown in Figure 1). (D) Performance separated by location of the source relative to the air flow direction (downwind, location 1 vs. upwind, locations 2 and 3). Mice show consistently higher performance navigating towards the upwind sources early in training (9.86 ± 0.76 sec at upwind sources, 6.33 ± 0.68 sec at downwind source, $p=0.0027$, Mann-Whitney U-test). This preference is eliminated late in training (16.43 ± 0.67 sec at upwind sources, 15.25 ± 0.67 sec at downwind source, $p=0.16$; not significant, Mann-Whitney U-test). Error bars indicate standard error of the mean taken across the 6 mice.

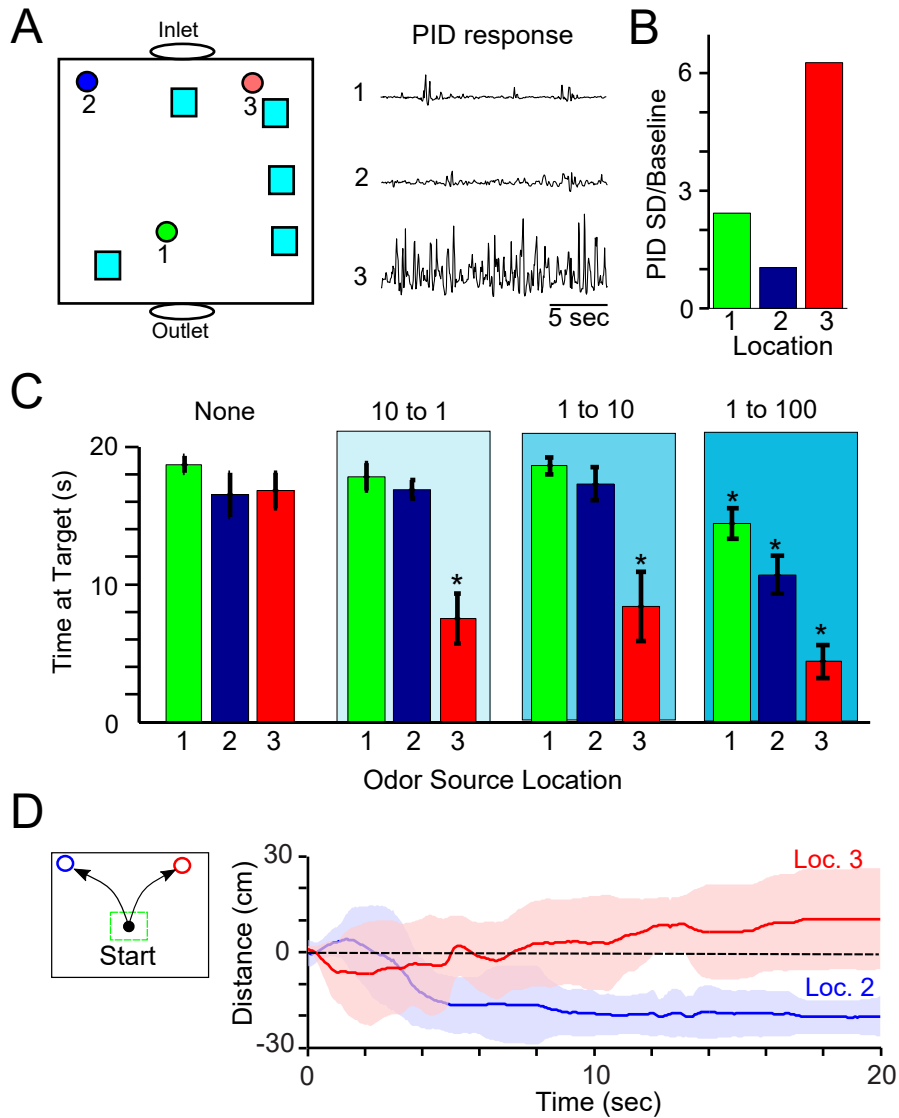


Figure S2, related to figure 1. Background sources of odor disrupt odor-based navigation. (A) Left: Diagram showing the location of fixed pieces of filter paper that carried a given concentration (see panel (C)) of the target odor, isoamyl acetate. These background sources of odor were placed near location 3 so as to maximally mask the odor signal in an area near this source. All mice used were late in training, during which time they sampled odor mostly near possible sources to stop at the right location. Right: PID recordings of background fluctuations taken near each source location. These recordings were taken without odor emanating from the source and thus reflect the strength of the background sources near each source location. Only location 3 shows extremely large fluctuations near the source, while the others do not. (B) Data quantified from the experiment shown in panel (A). (C) Bar graphs taken from sessions using 6 well-trained mice (corresponding to the “Late” category in the main text) showing time at a given odor/water source for four different ratios of target to background odor sources. Even with the target at a 10 fold higher concentration than the background (second panel from the left), the ability of mice to navigate to odor source 3 was significantly reduced (16.8 ± 1.3 sec without background, 7.5 ± 1.8 sec with background, $p < 0.009$). Only extremely high concentrations of background odor (100% isoamyl acetate) were capable of reducing navigation performance to sources 1 and 2 (18.7 ± 0.6 sec source 1 without background, 16.5 ± 1.6 sec source 2 without background, 14.4 ± 1.1 sec source 1 with background, 10.7 ± 1.4 source 2 with background, $p < 0.02$ for both comparisons). (D) Left: Lateral distance from the center port was measured for all paths that started in the middle of the arena. Right: Mice were capable of rapidly navigating to location 2 following odor release (blue, significant divergence from center by 5 seconds), while location 3 was much more challenging (no significant divergence from the center). Colored lines represent the mean and the colored areas represent the standard error of the mean for all trajectories that started in the center.

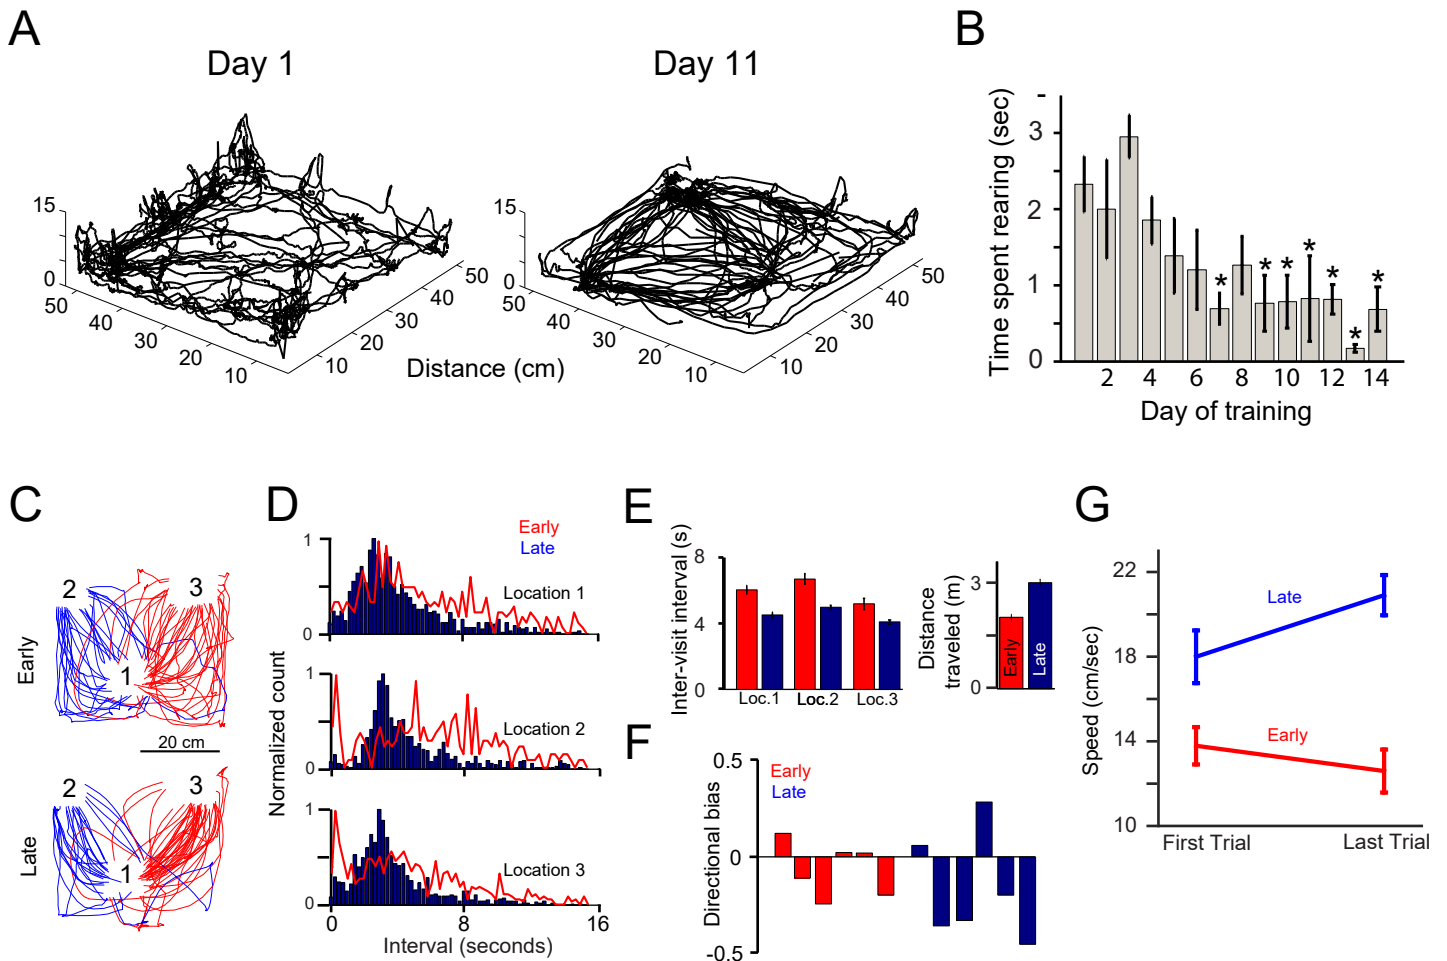


Figure S3, related to figure 2. Rearing by mice decreases with training as they develop foraging strategies.

(A) Three dimensional trajectories of mice while odor is present taken from day 1 of training (left) and day 11 of training (right). Initially mice spent more time in an elevated position with head extended upward into the air (rearing, as shown by large upward deviations in body position). This behavior was largely eliminated following training on the odor search task. (B) Bar graph showing rearing as a function of training session. Initially mice spent 2-3 seconds per 20 second odor period rearing. With increased training on the task mice spent much less time rearing (typically less than a second). Training sessions indicated with an asterisk show significantly less rearing compared to the first two sessions ($p < 0.05$, corrected for multiple comparisons). (C) Representative foraging trajectories from 6 mice starting from the center source and moving towards source 2 (blue) and source 3 (red) during the inter-trial time period when odor is not present in the arena. With training, the paths taken during foraging become less variable. (D) Histograms of the time between visits to each location (inter-visit interval) during foraging. Late in training mice develop foraging strategies that involve regular visits to each possible source location. Peaks near zero early in training indicate an immediate return to check the same location twice in a row. (E) Quantification across 6 mice of behavioral variables during foraging. Left panel: mice visit each possible source more often late in training (Group averages across source locations: Early: 5.96 ± 0.23 seconds; Late: 4.52 ± 0.1 seconds, $p < 0.00003$). Right panel: mice travel a longer distance during foraging late in training (Early: 2.02 ± 0.084 meters, Late: 3 ± 0.084 meters). (F) Each mouse develops a directional preference (defined as the probability of a clockwise (positive values) vs. counter-clockwise (negative values) movement along their foraging path. A value of 1 is entirely clockwise while 0 indicates no preference. The magnitude of the directional preference for most mice increases from early to late training (5 of 6 mice show an increase; $p = 0.06$ across all 6 mice, Mann-Whitney U-test). (G) Self-generated odor trails could play a role in guidance late in training. Across 27 sessions for mice late in training a trend was seen towards an increase in speed from the first to last trial in each session (first trial: 17.99 ± 1.25 cm/sec, last trial $20.91 \pm .95$ cm/sec, $p = 0.0693$, Mann-Whitney U-test; blue lines). Since the arena was cleaned between sessions (eliminating surface-deposited odor cues), this result is consistent with usage of built-up self-generated odor trails to more efficiently guide movement. This was in marked contrast to the same mice early in training, where across 24 sessions a slight decrease in speed from the first to last trial occurred (first trial: $13.79 \pm .88$ cm/sec, last trial $12.6 \pm .1$ cm/sec, $p = 0.359$; red lines).


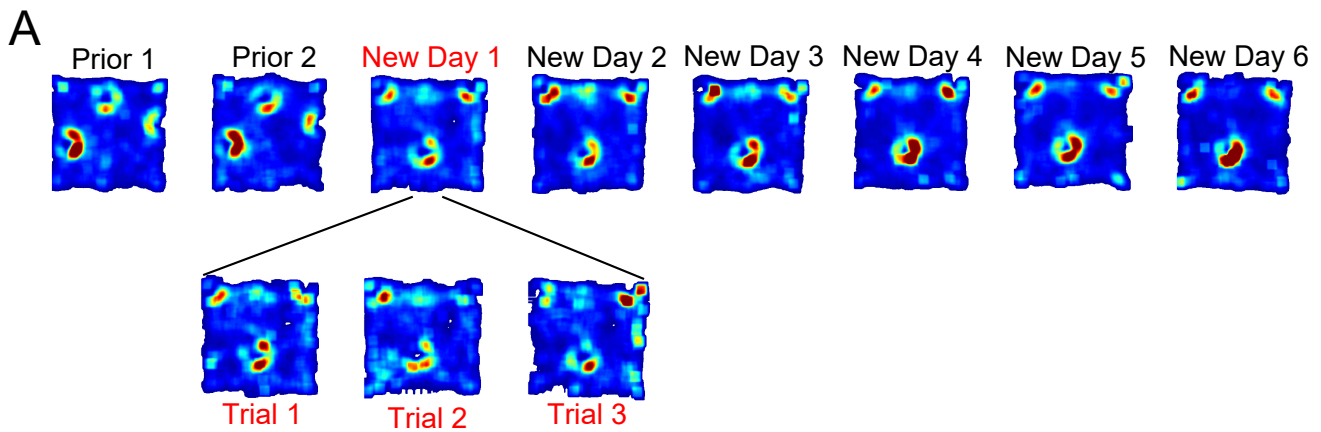


Figure S4, related to figure 2. Mice continue to search in areas where rewards were previously located. (A) Heat maps showing mouse behavior during odor search as path density. Arena configurations marked “prior” indicate the prior configuration of the odor sources (note high density areas which are the locations of the odor sources) and “New” indicate sessions with a new configuration of sources which is then maintained for the rest of the sessions. The top row represents data from the first 3 trials for each source location for the indicated day of training. The bottom row of 3 shows mouse behavior during the first 3 trials on the first day after the odor sources were switched. (B) Quantification of search path density in areas that contained the previous odor sources. Mice did not immediately stop visiting the previous odor source locations, but rather gradually eliminated visits to these possible source sites over a period of days.

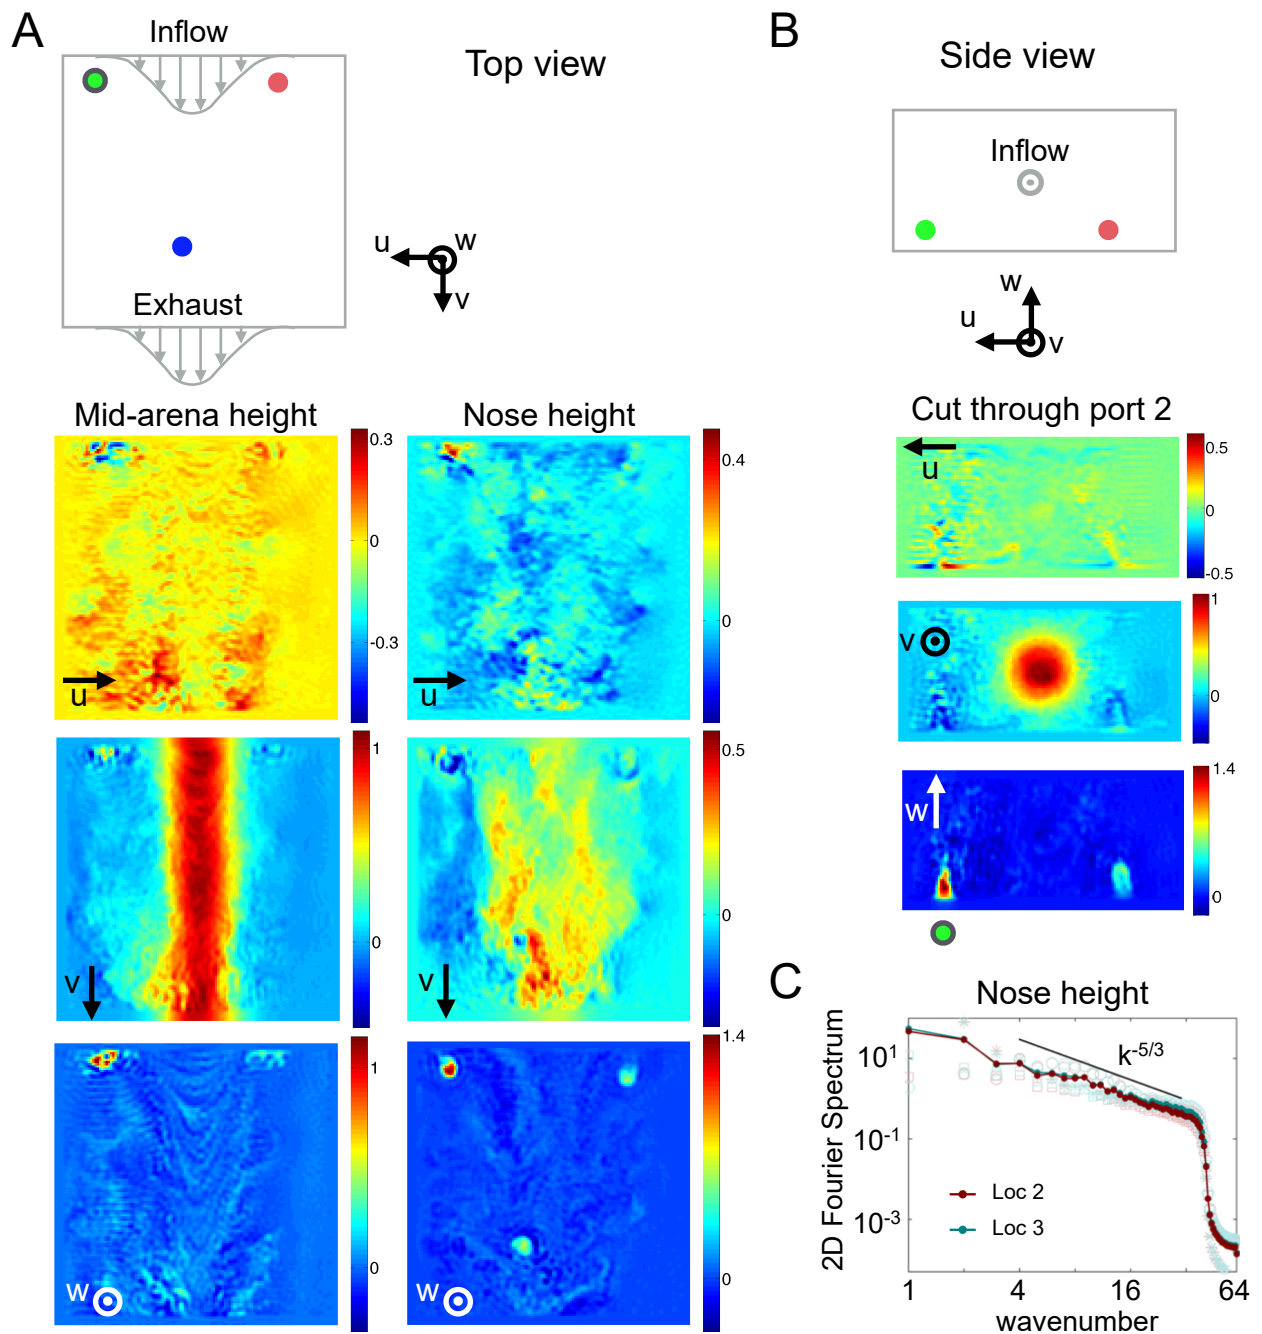


Figure S5, related to figure 3. Visualization of airflow in the arena. (A) Top view sketch of the arena (top). Snapshots of the three components of the velocity field in a horizontal cut at mouse nose (right) and mid-arena height (left). (B) Side view sketch of the arena (top). Snapshots of the three components of the velocity field in a vertical cut through Location 2 (bottom). (C) Two-dimensional spectrum of velocity fluctuations at mouse nose height as a function of wavenumber k displays a power law behavior, with an exponent close to $-5/3$.

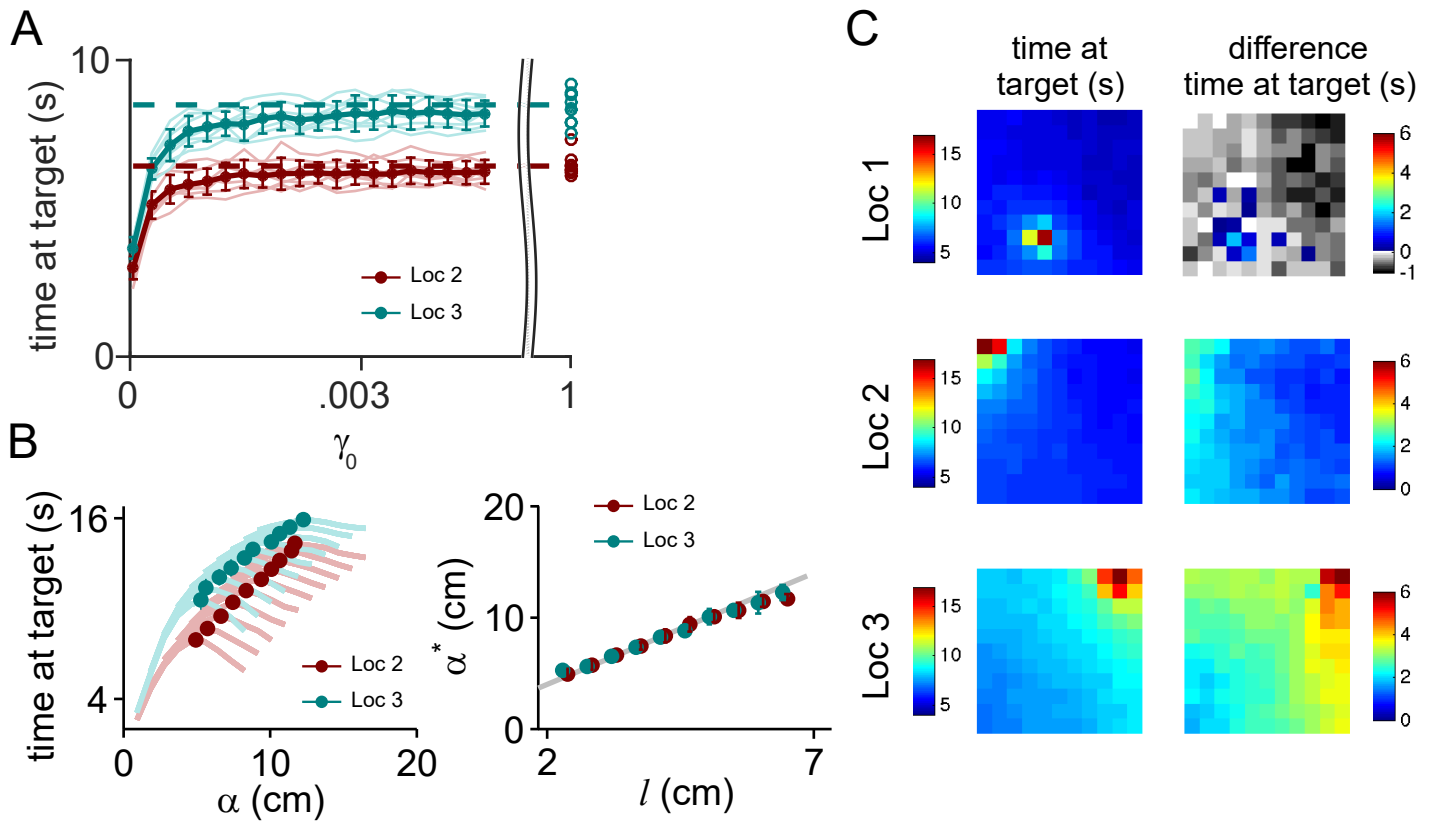


Figure S6, related to figure 4. Optimization gradient climb algorithm. (A) Time spent at target as a function of γ_0 . Performance of the gradient ascent algorithm is a monotonically growing function of γ_0 (see Materials and Methods). Solid bright lines: all realizations; darker lines and full dots: averages across realizations for Location 2 (red) and 3 (cyan). Open circles correspond to $\gamma_0 = 1$ for each realization, the point has been displaced in x for better visualization. Performance asymptotes to the optimal choice of an adaptive $\gamma = \alpha/|\nabla F|$, so that the step length $|dx_n| = \alpha$ is constant (horizontal dashed red and cyan lines, for location 2 and 3 respectively). (B) Left: Performance of the algorithm as a function of α for various values of the tolerance l . Bright red and bright cyan lines: time at target, averaged across all realizations for Location 2 and 3 respectively. Each curve corresponds to a value of the tolerance l that goes from 5 to 15 grid cells (2 to 7 cm) in 10 equally spaced intervals. Dark full red and cyan dots: optimal α^* , for each l for Locations 2 and 3 respectively. Right: optimal α^* as a function of the tolerance l . Red and cyan dots: optimal value of α averaged over realization, \pm standard deviation for Locations 2 and 3 respectively. Grey line: $\alpha^* = 2l$. The optimal α^* , depends linearly on l , representing the precision required for source localization. This is consistent with standard results of optimization using stochastic subgradient ascent methods for convex functions. (C) Left: performance of gradient ascent algorithm as a function of initial location of the searcher for Location 1, 2 and 3 (top to bottom). Location 3 generates the largest plume and is hence the easiest to track. Location 1 generates an extremely small plume, and the algorithm only performs well when starting in a 5 cm grid of locations around the source. Right: difference of time at target of gradient ascent algorithm and straight lines. Gradient ascent always outperforms straight lines toward the upstream sources (middle, bottom); whereas for the downstream source, it only outperforms the straight lines in the immediate vicinity of the source (top). Note that the color scale for Location 1 starts from negative values (Grey scale).

SUPPLEMENTAL EXPERIMENTAL PROCEDURES

Behavioral apparatus

Arena. The navigation arena was a fully enclosed acrylic box measuring 0.6 meters per side and 0.3 meters tall (see Figure 1A), with openings only for air inflow and exhaust. Odor source locations were intentionally placed asymmetrically within the arena in order to create distinct odor flow and mouse navigation trajectories that would be more readily analyzed with statistical measures. The only light was provided by deep red LEDs (wavelength 660 nm).

Olfactometer. Odors were presented within through narrow outlets close to the ground using a custom-built olfactometer. Odors were diluted by volume in solvent (1 ml of diethyl phthalate) and stored in glass tubes. Air was passed through these tubes and then fed into the arena through 1/8 inch diameter flexible Teflon (Tygon formulation 2375) tubes. In order to eliminate auditory and airflow cues due to changes in airflow with odor onset, a constant flow of clean air was passed through the tube at each odor source location and the output of the olfactometer was added to this flow at a single selected odor source location on each trial; this flow was also included in all fluid dynamic simulations (see Fluid Dynamic Simulation section below). To minimize auditory cues, all olfactometer valves were located together 1 meter away from the arena. There were no valves at the odor sources in the arena.

Behavioral monitoring system. Mice were monitored using a dual camera system (The Imaging Source DMK 72BUC02). Acquisition was triggered simultaneously in both cameras through a TTL pulse generated by custom-written software in labVIEW (National Instruments). The center of mass location for the mouse was recorded at 16 Hz. To generate coordinates for the mouse, the center of mass pixel coordinates were fed through a distortion model for each camera and lens system, and then combined to perform stereo triangulation (Camera calibration toolbox for MATLAB; http://www.vision.caltech.edu/bouguetj/calib_doc/). Where appropriate, analysis was simplified by considering only 2 dimensions. Arrival at the water spout was determined by a photo-diode placed immediately under the water spout.

Fluid dynamics simulations

Fresh air flows into the arena from a circular inlet 10 cm in diameter, located as shown in the sketches in Figure S5A. Odor and air flow out of the arena through a symmetric circular outlet at the other end of the arena. A small volume of air flows into the arena vertically from three tubes, located close to the ground, in the locations showed in the sketches in Figure S5A. Only one of the three tubes contains the odor. We used a pseudo spectral code to integrate the Navier-Stokes equations and we enforced the boundary conditions (at inlet, outlet, ports and walls) using a Brinkman penalization term [S1].

$$\begin{aligned} \partial_t \mathbf{u} + \mathbf{u} \cdot \nabla \mathbf{u} &= \nu \Delta \mathbf{u} - \nabla \frac{p}{\rho} - \frac{\mathbf{u} - \mathbf{u}_{bc}}{\sigma} \chi \\ \nabla \cdot \mathbf{u} &= 0 \end{aligned}$$

where ν is the kinematic viscosity of air, χ is 0 in the bulk and 1 at the boundaries and \mathbf{u}_{bc} is the boundary condition. For small σ , the penalized velocity field \mathbf{u} converges to the solution of the Navier-Stokes equation with the correct boundary conditions. We chose σ equal to the time step.

$$\mathbf{u}_{bc} = \begin{cases} 0 & \text{all walls} \\ v_{io} \hat{y} & \text{inlet, outlet} \\ w_p \hat{z} & \text{ports} \end{cases}$$

with the coordinates as in the sketches in Figure S5. To model the vents at the inlet/outlet and avoid Gibbs oscillations, v_{io} is a Gaussian velocity profile decreasing rapidly outside of the circular inlet. We added a sinusoidal perturbation with vanishing mean, and modulated amplitude 30% of the local speed, wavelength equal to the width of the inlet, and a random phase varying at each time step. We imposed the same boundary condition at the inlet and outlet, as required by periodicity of the pseudo spectral method. The maximum velocity at the inlet is 1m/s, with 30% fluctuations imposed at the boundary, matching the measures obtained through an anemometer (data not shown). Additionally, a small volume of air is emitted vertically at each port location to ensure odor release, with w_p modeled as a Gaussian profile rapidly decreasing at distances larger than 1cm (with 40 cm^3/s from the port releasing the odor and 15 cm^3/s from the other ports, matching the fluxes used in the experiments). These small fluxes were modeled with a Brinkman penalization term at the ports, conveniently located few grid cells above ground, at the x-y

locations shown in the schematics in Figure S5. This mimics the real situation, where the tubes bringing odor in the arena are slightly above ground, and also avoids a non-periodic boundary condition, which cannot be handled with a Fourier pseudo spectral method. The odor concentration near the source is large, thus the details of the flow right below the port are irrelevant to the search algorithm. The resolution is chosen to resolve the smallest turbulent fluctuations. At the stationary state the smallest turbulent structures have a typical size $\eta \sim L/Re^{3/4} \sim 0.9\text{cm}$ (23) where $L \sim 60\text{cm}$ is the size of the arena and Re is the Reynolds number based on the turbulent fluctuations $v_L \sim 4\text{cm/s}$ (about 4% of the inlet velocity) and inlet diameter. Accordingly, we choose $dx = 0.4\text{cm}$. The time step was chosen according to the numerical stability criterion $dt < C dx/U_{max} = 2\text{ms}$ where we chose empirically the Courant constant $C = 0.5$.

For each realization of the velocity field, the odor was simulated by tracking the trajectory of a total of 2.8×10^7 tracer particles, released continuously at a rate of 1400 particles per ms at the active port location, and transported passively by the flow \mathbf{v} . The position of each particle \mathbf{X} evolves in time according to:

$$\frac{d\mathbf{X}}{dt} = \mathbf{v}(\mathbf{X}(t), t) + \sqrt{2D}\eta$$

where D is the diffusivity of the odor in air, and $\mathbf{v}(\mathbf{X}(t), t)$ is the velocity of air, computed at the location of the particle. The concentration of particles is proportional to the concentration c of the odor, advected by the flow according to the equation of passive scalar transport:

$$\partial_t c + \mathbf{v} \cdot \nabla c = D\Delta c$$

with a boundary condition $c = c_0$ at the port location. At mouse nose height, the concentration right above the odor source is about 1000 particles, so that our resolution for the odor field is $c_0/1000$. Because these are forward rather than backward Lagrangian trajectories [S2], we evaluated the error empirically, by comparing the odor field reconstructed by independent sets of N particles, as a function of N . We used N such that the mean error fell below 5% to 9%. The grid spacing is about 0.4 cm , so that one grid cell represents about $64\ \mu\text{L}$, comparable with the volume of air typically inhaled by the animals. Brownian diffusion in the Lagrangian trajectories reproduces molecular diffusion of the odor (see e.g. [S2] and references therein) and is mostly relevant for particles starting initially closer than the Kolmogorov length scale.

Figure S5A and B represent snapshots of one instance of the air flow. We show the three components of the velocity along 2 horizontal cuts, at mouse nose (Figure S5A - right) and at vent height (Figure S5A - left) and 1 vertical cut at the position of the port emitting the odor (Figure S5B). The air stream in the box is rapidly fluctuating. The Reynolds number of the simulated air flow, based on the inlet diameter and maximum velocity is 6600. The main circular jet is thus in the mixing transition regime [S3], where the flow generates increasing interfacial area between odor rich regions and pure air, which favors mixing. In this regime, we observed the development of power law spectra for the velocity fluctuations, with respect to the time averaged mean flow. The 2D isotropic spectrum of the velocity at mouse nose height displays an exponent close to the dimensional prediction $-5/3$ (see Figure S5C), consistent with previous observations [S4]. The average velocity fluctuations are about 4% of the maximum velocity at the inlet, so that the Reynolds number based on the velocity fluctuations is about 260.

Source Localization Algorithm

Optimization of the stepsize. Because the odor gradient bears information about the location of the source, we focused on a generic gradient ascent algorithm:

$$\mathbf{x}_{t+1} = \mathbf{x}_t - \gamma_t \nabla F(\mathbf{x}_t, t) \quad (2)$$

where \mathbf{x}_t is the position of the searcher at time t ; γ_t is the step size that may depend on time; $F = [c(\mathbf{x}(t), t) - c^*]^2$ is the cost function to be minimized; $c(\mathbf{x}, t)$ is the two dimensional horizontal snapshot of odor concentration at time t , computed at mouse nose height and c^* is the value of the concentration at the source; the gradient ∇F is calculated at the current position of the searcher. We empirically optimized performance of the stochastic gradient ascent algorithm in eq (2). We assumed constant $\gamma_t = \gamma_0$ and search for its optimal value. Because the increments represent the distance travelled by the mouse in between subsequent sniffs, we enforced a maximum increment $|d\mathbf{x}_n| = |\mathbf{x}_{n+1} - \mathbf{x}_n| < \alpha$. This is equivalent to $\gamma_t = \min(\gamma_0, \alpha/|\nabla F|)$. We find that performance is a monotonically increasing function of γ_0 . The optimal choice corresponds to $\gamma_0 \rightarrow \infty$ which amounts to always travel at maximum speed, with $\gamma_t = \alpha/|\nabla F|$ (Figure S6A). Note that this is a

standard choice in stochastic subgradient methods [S5]. The algorithm thus becomes eq (1). Note that minimization of the cost function F in eq (2) corresponds to maximization of the odor concentration $c(\mathbf{x}(t), t)$, hence gradient “descent” of F is equivalent to gradient “ascent” of $c(\mathbf{x}(t), t)$.

We next varied the values for α and found an optimal value that scales linearly with the precision that we require for source localization, $\alpha \sim 2\ell$ (see Figure S6B). Note that, although convergence is not guaranteed, because F is not convex, this result is consistent with standard results obtained in optimization for convex functions [S6]. Convergence of standard subgradient methods for minimization of convex functions is ensured in a region of size $\ell = \alpha/2$ near the real minimum. Larger and smaller values of α are then sub-optimal because they either overshoot the target or reach it with unnecessary precision. If the signal can be decomposed in a convex average plus noise, similar convergence results can be proved in expectation, although the rates of convergence depend on the specific details of the signal. The results shown in Figure S6 indicate that this behavior is robust and holds in our case as well. This is somewhat surprising, as the fluctuations display non-trivial space and time correlations stemming from the physics of turbulent transport and the time averaged signal is not convex as smoothing out a turbulent plume requires an extremely long time averaging [S6].

Memory. We next asked whether it is advantageous to average over subsequent snapshots, to increase the signal to noise ratio. To this end we use a filtered signal with memory T_{memory} :

$$\bar{c} = \frac{1}{T_{\text{memory}}} \int_{t-T_{\text{memory}}}^t c(\mathbf{x}, t') dt'$$

Performance improves with averaging up to $T_{\text{memory}} \approx 0.2$ to $0.3T_{\text{task}}$ beyond which the performance measured as time to reach target saturates (Figure 4E left), and performance measured as success rate decreases (not shown). Because mice would have to physically pause to average the signal at any given location, we repeat the analysis allowing only 1 position update for each T_{memory} time interval, required to average. In this case, performance is monotonically decreasing with memory: it is never beneficial to pause and average the signal (Figure 4E right). This is true even in the case where the searcher pauses only when it gets the signal (empty dots

Figure 4E right). We conclude that, although computationally there is an optimal averaging window, memory is irrelevant for the algorithm analyzed in this paper. Mice are capable of more complex behavior, modulating their sniffing rates and likely the inhaled volume. A more complete analysis of averaging and sniffing behavior during olfactory tasks is left for future studies.

Direct approaches. We quantified the probability of direct approach toward Location 2 and 3, and we found that an average probability of direct approach toward Location 2 and 3 of 64% and 78% respectively. These values compare well with the experimental rates of direct approach for animals that have limited experience with source locations Locations 2 and 3, 74 and 81% respectively. Also note that because the arena is confined, the searcher does eventually find the source even when purely performing straight lines trajectories. However, the performance of the straight-lines behavior is poorer than the gradient ascent algorithm (dashed lines in Figure 4A) and its chance of direct approach is close to 50% (Figure 4B, bottom), inconsistent with mouse behavior under conditions in which they follow the plume (Figure 2C).

SUPPLEMENTAL REFERENCES

- S1. Kevlahan, N.K.-R., and Ghidaglia, J.-M. (2001). Computation of turbulent flow past an array of cylinders using a spectral method with Brinkman penalization. *European Journal of Mechanics - B/Fluids* 20, 333–350.
- S2. Celani, A., Cencini, M., Mazzino, A., and Vergassola, M. (2004). Active and passive fields face to face. *New J. Phys.* 6, 72.
- S3. Dimotakis, P.E. (2000). The Mixing Transition in Turbulent Flows. *J. Fluid Mech* 409, 6998.
- S4. Dimotakis, P.E. (2005). Turbulent Mixing. *Annual Review of Fluid Mechanics* 37, 329–356.
- S5. Shor, N.Z. (1985). *Minimization Methods for Non-Differentiable Functions and Applications* (Berlin ; New York: Springer-Verlag).
- S6. Celani, A., Villiermaux, E., and Vergassola, M. (2014). Odor Landscapes in Turbulent Environments. *Phys. Rev. X* 4, 041015.