


Article

Temporal fluctuations in chemotaxis gain implement a simulated-tempering strategy for efficient navigation in complex environments



Omer Karin, Uri Alon

uri.alon@weizmann.ac.il

Highlights
High chemotaxis gain impairs patch crossing, whereas low gain impairs peak localization

Stochastic switching between low and high gain periods provides a superior strategy

Stochastic strategy is analogous to simulated tempering sampling algorithm

Simulated tempering may be a widespread principle for efficient navigation

Karin & Alon, iScience 24, 102796
July 23, 2021 © 2021 The Authors.
<https://doi.org/10.1016/j.isci.2021.102796>



Article

Temporal fluctuations in chemotaxis gain implement a simulated-tempering strategy for efficient navigation in complex environments

Omer Karin^{1,2,3} and Uri Alon^{1,4,*}

SUMMARY

Bacterial chemotaxis is a major testing ground for systems biology, including the role of fluctuations and individual variation. Individual bacteria vary in their tumbling frequency and adaptation time. Recently, large cell-cell variation was also discovered in chemotaxis gain, which determines the sensitivity of the tumbling rate to attractant gradients. Variation in gain is puzzling, because low gain impairs chemotactic velocity. Here, we provide a functional explanation for gain variation by establishing a formal analogy between chemotaxis and algorithms for sampling probability distributions. We show that temporal fluctuations in gain implement simulated tempering, which allows sampling of attractant distributions with many local peaks. Periods of high gain allow bacteria to detect and climb gradients quickly, and periods of low gain allow them to move to new peaks. Gain fluctuations thus allow bacteria to thrive in complex environments, and more generally they may play an important functional role for organism navigation.

INTRODUCTION

Bacteria navigate up and down gradients of chemicals in a process called chemotaxis. Chemotaxis is achieved by transitions between swimming at a roughly uniform speed and constant direction (“runs”) and random reorientations (“tumbles”). Bacteria climb gradients of chemical ligands by modulating their tumbling frequency, so that tumbling rate decreases when the bacteria move up gradients of chemoattractants and increases when they go down the gradient (Figure 1A).

Chemotaxis is well-characterized in terms of its molecular signaling circuit. The circuit implements a non-linear integral feedback loop, which brings tumbling rates precisely back to the baseline after changes in input (exact adaptation) and allows the bacteria to respond to relative changes in ligand input (fold-change detection [FCD]) across a wide dynamic range (Shoval et al., 2010; Lazova et al., 2011). Analogous navigation systems also appear in eukaryotes (Polin et al., 2009; Arrieta et al., 2017) and simple animals (Larsch et al., 2015; Borba et al., 2020).

Several aspects of the chemotaxis phenotype are highly variable between individual isogenic bacteria (reviewed by Waite et al. [2018]; Colin and Sourjik [2017]). This variability results, at least in part, from stochastic fluctuations in the abundance of proteins in the chemotaxis signaling pathway (Spudich and Koshland, 1976; Levin et al., 1998) and from tunable post-translational modifications (Kamino et al., 2020) and is subject to genetic control.

A key question is whether phenotypic variability has an adaptive functional role for chemotaxis. Most of the research on the adaptive role of variability has focused on the baseline tumbling rate (Spudich and Koshland, 1976; Park et al., 2010; Matthäus et al., 2009; Frankel et al., 2014; Dufour et al., 2016; Keegstra et al., 2017). Some individual cells show high tumbling and others show low tumbling rate, and this behavior lasts for an entire cell generation. In each individual cell, the adaptation time to a step attractant change is inversely proportional to that cell's tumbling frequency (Spudich and Koshland, 1976). A recent study also demonstrated that some cells have discrete on/off temporal fluctuations in the chemotaxis signaling pathway, with a timescale of minutes (Keegstra et al., 2017). In an important theoretical study, Frankel et al. proposed that high tumbling rate is preferable when the navigation goal is near, whereas low tumbling rate is preferable when the goal is far (Frankel et al., 2014). This prediction was confirmed by subsequent experimental work (Dufour et al., 2016; Waite et al., 2016), which also suggested that heterogeneity can be

¹Department Molecular Cell Biology, Weizmann Institute of Science, Rehovot, Israel

²Wellcome Trust–Cancer Research UK Gurdon Institute, University of Cambridge, Cambridge, UK

³Department of Applied Mathematics and Theoretical Physics, Centre for Mathematical Sciences, University of Cambridge, Cambridge, UK

⁴Lead contact

*Correspondence: uri.alon@weizmann.ac.il
<https://doi.org/10.1016/j.isci.2021.102796>






Figure 1. Stochastic heterogeneity in pathway gain in bacterial chemotaxis

(A) Bacterial chemotaxis is the navigation process by which bacteria accumulate near peaks of chemoattractants and away from chemorepellants. Chemotaxis is performed by adjusting the tumbling rate according to the sensed concentration of the ligand, so that bacteria tumble less when they move up attractant gradients or down repellent gradients.

(B) The accumulation of bacteria around attractant peaks is determined by an inverse temperature parameter β , which is proportional to the chemotaxis gain, also called pathway sensitivity. Recent experiments showed that there is large and persistent non-genetic heterogeneity in gain, suggesting that there may be temporal fluctuations in β (Salek et al., 2019). These fluctuations may be short lasting or long lasting (e.g., due to cell-to-cell variation in receptor copy numbers). The functional significance of these fluctuations is currently unclear.

beneficial for amplifying the contribution of high-performing individuals at the tail of the phenotype distribution. Stochastic fluctuations in tumbling rate may also contribute to efficient exploration of the environment by generating Lévy walk search behavior (Tu and Grinstein, 2005; Matthäus et al., 2009).

Recently, another major source of cell-cell variability was discovered, whose functional role is still not clear. This is variability in the *pathway gain* (Salek et al., 2019). The gain, also called sensitivity, determines how strongly bacteria change their tumbling rate in response to a given gradient (Figure 1B) (Sourjik and Berg, 2002; Colin and Sourjik, 2017). Salek et al. (2019) found that gain varies between genetically identical individual cells by approximately an order of magnitude. Cell-cell variation in pathway gain is surprising, because large gain increases both chemotaxis efficiency and how tightly bacteria accumulate around attractant peaks (Salek et al., 2019). One may therefore expect that gain should be maximized. It has been proposed that cell-cell variability in pathway gain may be beneficial for bet-hedging by preventing the accumulation of too many bacteria in a small spatial region (Salek et al., 2019).

Here we consider the possibility that the observed cell-cell variability in gain is due to temporal fluctuations, which can last for less or more than the generation time of an individual cell.

We propose a functional role for temporal fluctuations in pathway gain using a formal analogy to sampling algorithms in physics and computer science. Gain fluctuations allow bacteria to efficiently navigate in complex environments where the ligand distribution has many local maxima and minima. Without gain fluctuations, bacteria would get stuck at local maxima. To show this, we establish a mapping between bacterial chemotaxis and a widely used method from statistical physics for sampling complex probability distributions known as *simulated tempering*, which uses dynamic changes in temperature to cross energy barriers (Swendsen and Wang, 1986; Marinari and Parisi, 1992; Hukushima and Nemoto, 1996; Earl and Deem, 2005; Lee et al., 2018). We show that stochastic fluctuations in pathway gain provide a biological implementation of simulated tempering and demonstrate that it allows colonization and growth in complex, patchy environments.

Our model makes minimal assumptions and can be readily generalized to FCD-based navigation systems in other organisms. We therefore conclude by discussing several possible generalizations of our results.

RESULTS

Model for bacterial chemotaxis as a sampling process

We begin by establishing a connection between chemotaxis and the sampling of a probability distribution, where the distribution is the distribution of chemoattractant in space. The sampling framework we use is

the Langevin Monte Carlo (LMC) approach. LMC is widely used for efficient sampling (Roberts and Tweedie, 1996; Neal, 2011; Girolami and Calderhead, 2011; Dalalyan, 2014) and for global optimization (Chiang et al., 1987; Gelfand and Mitter, 1991; Lee et al., 2018; Erdogdu et al., 2018; Ma et al., 2019; Chen et al., 2020). In this article, we exploit the analogy between chemotaxis and LMC to understand the efficiency of chemotaxis from an algorithmic perspective.

To establish the relation between bacterial chemotaxis and LMC sampling, we use the chemotaxis model of Tu et al. (2008). This model provides good agreement with experimentally measured responses of bacteria to temporally varying stimuli (Shimizu et al., 2010) and has become a standard model for chemotaxis (Menolascina et al., 2017; Salek et al., 2019; Alon, 2019). The model describes the fast (sub-second) inhibition of the tumbling rate $\lambda(t)$ by the attractant ligand input $L(t)$. It also describes the slower (seconds-minutes) adaptation to a step change in $L(t)$ that is due to receptor methylation, $m(t)$, a negative feedback process. The tumbling rate is controlled by a fast internal “receptor activity” variable $a(t)$:

$$a(t) = \frac{1}{1 + e^{N\left(\alpha(1-m) + \log\left(1 + \frac{L}{K_I}\right) - \log\left(1 + \frac{L}{K_A}\right)\right)}} \approx \frac{1}{1 + e^{N\left(A(1-m) + \log\left(\frac{L}{K_I}\right)\right)}} \quad (\text{Equation 1})$$

where K_I, K_A, α, N are constants and the approximation holds in the range of receptor sensing: $K_I \ll L \ll K_A$. The dynamics of the slow negative feedback process is given by:

$$\dot{m}(t) = \omega \left(1 - \frac{a}{a_0}\right) \quad (\text{Equation 2})$$

where a_0 is the steady-state level of a and ω sets the timescale of adaptation. The tumbling rate is given by:

$$\lambda(a) = \frac{1}{\tau} \left(\frac{a}{a_0}\right)^H \quad (\text{Equation 3})$$

where τ is steady-state run time, which is on the order of 1 s. Typical values for the constants $\alpha, N, a_0, K_I, K_A, \omega, H$ are provided in Table S1. The adaptation time of m , which is approximately inversely proportional to $(1 - a_0)N\alpha\omega$ (STAR Methods), is on the order of seconds to minutes.


The model has several important features. It has the FCD property, where tumble responses to a time-varying input $L(t)$ depend only on changes of $L(t)$ relative to its baseline (Shoval et al., 2010; Shimizu et al., 2010; Adler and Alon, 2018). More generally, $a(t)$ is proportional to the logarithmic derivative of the low-frequency signals in the input, with a frequency threshold that is inversely proportional to the adaptation time. These features, which are common in sensory circuits (Adler and Alon, 2018), have been experimentally demonstrated (Lazova et al., 2011). Thus, a bacteria running in a static ligand field L in direction \vec{v} relative to the gradient will have $a(t) \propto v \vec{v} \cdot \nabla \log[L](x)$, where v is the swimming speed of the bacteria.

The run-and-tumble movement of *E. coli* is nearly isotropic at time and length scales much larger than $\tau, v\tau$ (Berg, 1993). This movement is thus well approximated by a diffusion process with diffusion constant $D = \frac{v^2 d - 1}{z_\theta + \tau^{-1}}$, where d is the dimension and z_θ is rotational diffusivity, which is small compared with the tumbling frequency τ^{-1} . The temporal evolution of the bacterial population can therefore be modeled by a Fokker-Planck equation that depends on the aforementioned signaling dynamics, as shown by several studies that also provide excellent correspondence to empirical distributions with plausible parameters (Si et al., 2012; Dufour et al., 2014; Menolascina et al., 2017). Here we analyze the equivalent Langevin equation, which describes the stochastic time evolution of the location of an individual bacterium:

$$dx = \chi \nabla \log L(x) dt + \sqrt{2D} dW \quad (\text{Equation 4})$$

where W is a 1-dimensional (1D) Wiener process and the advection parameter χ is proportional to the swimming speed squared: $\chi = [v^2 \tau^{-1} / (z_\theta + \tau^{-1})^2] H N (1 - a_0)$ (Si et al., 2012). Note again that Equation 4 applies in the range where sensing is logarithmic $K_I \ll L \ll K_A$.

The advection parameter χ is also called the chemotactic sensitivity coefficient (Salek et al., 2019). We will study the 1D case ($d = 1$), but our conclusions generalize to 2 and 3 dimensions. We will also neglect the fact that the run direction after tumble may be correlated with the previous run direction, which increases diffusivity but does not affect our conclusions.



$F(x)$, experiment ..
 $L(x)^\beta$, $\beta=14$
 $L(x)^{\xi\beta}$, $\langle \xi \rangle=14, \xi \sim \text{Gamma}$

Figure 2. Inference of the inverse temperature parameter β from chemotaxis in linear gradients

Kalinin et al. (2009) performed experiments where a bacteria (*E. coli*) navigated linear gradients $L(x)$ of the chemoattractant MeAsp in a microfluidic device (insets). The accumulation profiles $F(x)$ (blue circles) correspond well to an invariant probability distribution proportional to $L(x)^\beta$ where $\beta = 14$ (black lines). This correspondence holds even if we account for the heterogeneity in β observed in Salek et al., by taking the invariant probability distribution to be proportional to $L(x)^{\xi\beta}$, $\xi \sim \text{Gamma}[2.25, 35]$ (red lines).

Equation 4 is equivalent to a continuous LMC process that samples an invariant probability distribution $P(x)$: at long times, the bacterial density converges to $P(x)$. This distribution is proportional to a power β of the ligand distribution (Lee et al., 2018):

$$P(x) \propto e^{\beta \log L(x)} = L(x)^\beta \quad (\text{Equation 5})$$

where:

$$\beta = \frac{\chi}{D} = \frac{\tau^{-1}}{z_\theta + \tau^{-1}} HN(1 - a_0) \approx HN(1 - a_0) \quad (\text{Equation 6})$$

Equation 5 is due to the log-sensing property of chemotaxis, which results in an advection term that is proportional to the logarithmic gradient in Equation 4. The logarithmic gradient is equivalent to the advection term of LMC. A search strategy that does not respond to the logarithmic gradient will therefore not be equivalent to LMC.


The parameter β , which in statistical physics is proportional to the inverse temperature, determines the degree to which the invariant distribution is concentrated around the peaks of the attractant profile $L(x)$. For our purposes, it is important to note that the “inverse temperature” β is proportional to the chemotaxis gain, as well as to the baseline receptor activity a_0 (Equation 6). As there are temporal fluctuations in a_0 (Keegstra et al., 2017) and potentially also in chemotaxis gain, there may be large temporal fluctuations in β . Remarkably, β does not depend sensitively on swimming speed v or average run duration τ , which may be environment dependent; it depends only on intracellular signaling parameters.

When the inverse temperature parameter β is very small ($\beta \ll 1$), the bacterial distribution $P(X)$ is nearly uniform on its support; when it is very large ($\beta \gg 1$), $P(X)$ is concentrated around the global maximum of attractant $L(x)$; and when $\beta = 1$, the process samples the attractant distribution precisely, $P(x) \sim L(x)$.

What is the typical value of β for *E. coli* chemotaxis? Microfluidics experiments in a linear gradient in the sensing range (Kalinin et al., 2009) show that bacteria converge to a population distribution of $P(X) \propto L(x)^\beta$ where $\beta \sim 14$ (Figure 2, black lines). The analytical estimate from Equation 4 with the parameters provided in Si et al. (2012) also gives a high value of $\beta \sim 27$. The discrepancy may be partially due to a being away from the adapted state a_0 during gradient climbing (see Hu and Tu [2014] for details). Both estimates, however, suggest that *E. coli* tends strongly toward optimization, in the sense that cells converge to a distribution that is tightly concentrated around the attractant peak.

The experiments of Salek et al. suggest that there is large cell-cell variation in the pathway gain HN (and hence in β) (Salek et al., 2019). Accounting for this variation does not hinder the fit to the empirical distributions (Figure 2, red lines) and can explain the long tail observed in the steep linear gradient (Figure 2, right panel). However, the variance in pathway gain measured by Salek et al. suggests large cell-cell variation in β (i.e., a 95% range of $2 < \beta < 37$ for the estimate $\beta \approx 14$). This raises the question of what is the role of the large variability of β .

A Unimodal attractant distribution



B Bimodal attractant distribution




Figure 3. Stochastic tempering allows efficient movement between attractant peaks in multi-modal attractant distributions

(A and B) We simulated the chemotaxis Langevin processes in either unimodal (A) or bimodal (B) attractant distributions. The processes had a high gain, parametrized by an “inverse temperature” parameter β ($\beta = 30$, blue, left panel); a low gain ($\beta = 2$, orange, middle panel); or a tempering strategy with stochastic switches between high and low gain ($\beta = 30 - \beta = 2$). Parameters for all processes are provided in the STAR Methods section. Although the high- β process localizes well around the unimodal peak, it fails to cross between peaks in the bimodal case. The low- β process, on the other hand, crosses well between the peaks but localizes poorly around each peak. The tempering strategy can both localize well around each peak (in periods here β is high) and cross well between the peaks in periods where β is low.

A constant gain would prevent efficient chemotaxis

Understanding the adaptive importance of a trait requires us to understand how this trait contributes to the ability of the organism to survive and reproduce. For bacterial chemotaxis, an effective chemotaxis strategy is considered to be one that has rapid and tight accumulation of bacterial populations around peaks of attractants (or away from repellants) (Clark and Grant, 2005; Celani and Vergassola, 2010). In the simple case of a single-peaked Gaussian patch of attractant (Figure 3A), this is clearly optimized by maximizing pathway gain (and therefore maximizing “inverse temperature” β), because this maximizes the velocity in the direction of the gradient, as well as the tightness of the invariant distribution around the peak.

Such single-peaked patches are unlikely, however, to fully represent the natural environments in which *E. coli* live. We will therefore consider a more general environment that contains multiple attractant peaks, namely, a multimodal attractant distribution (Figure 3B). In such complex environments, the bacteria face two important tasks. The first is to navigate toward a maximum of attractant, and the second is to colonize many attractant peaks. As we will show using the analogy with sampling by LMC, these tasks are incompatible with simply maximizing or minimizing β .

The fundamental problem of sampling in complex multimodal environments is that of metastability—becoming stuck at one peak for long times. Consider an environment $L(x)$ that consists of two Gaussians of unit width with means at $x = 0, x = 2\mu$ (this corresponds, for example, to patches of dissolved organic matter). If a sampling process that samples $L(x)^\beta$ starts around the first peak $x = 0$, when will it visit the second peak? This transit time will be limited by the probability density around the midpoint (or valley) between the peaks $x = \mu$, which is proportional to $e^{-\frac{1}{2}\beta\mu^2}$. Therefore, the time it takes to cross between modes increases exponentially with inverse temperature β and with the distance μ . The feasibility of exploring multiple peaks is thus extremely sensitive to the choice of β and its relation with the statistical properties of the environment. Similar considerations apply for linear gradients from steady sources that decay like $L(x) = (1/x)$, where the probability density at distance μ goes like $\mu^{-\beta}$ (Figure S1). One possibility to avoid metastability is to set β low enough to allow transitions between the peaks. This, however, may be an unfavorable solution, because it comes at the cost of worse localization near the peaks: the bacterial distribution is spread out and nearly uniform.

To summarize, a chemotactic strategy with a single (or narrowly distributed) β faces an inauspicious trade-off. Choosing a large inverse temperature β results in tight accumulation around the attractant peak, but

crossing between patches becomes intractable; a low β may allow efficient crossing between patches, but the bacteria will not accumulate around the peaks of the attractant distribution.

Tempering allows efficient navigation in complex environments

To address this problem, we propose that bacteria employ a tempering strategy with stochastic temporal changes in inverse temperature β . This provides a physiological function to the stochastic variation in gain, which proportionally affects β . Our inspiration is the simulated tempering approach to LMC, in which stochastic switches between temperatures allow sampling of complex, multi-modal distributions (Marinari and Parisi, 1992; Lyubartsev et al., 1992). When multiple instances are run in parallel, the method is known as parallel tempering (Hansmann, 1997; Earl and Deem, 2005). Simulated tempering is also related to the simulated annealing method for global optimization (Kirkpatrick et al., 1983; Van Laarhoven and Aarts, 1987), in which temperature is changed over time according to a defined schedule.

For simplicity, we analyze stochastic switches between two inverse temperatures: $\beta_{hot}, \beta_{cold}$, with transition rates: $\eta_1 : \beta_{cold} \rightarrow \beta_{hot}, \eta_2 : \beta_{hot} \rightarrow \beta_{cold}$, similar to the binary fluctuations observed by Keegstra et al. (2017) (our results generalize to stochastic switches between many temperatures). Bacteria spend on average a fraction $\varepsilon = (\eta_1 / (\eta_1 + \eta_2))$ of the time at β_{hot} and $(1 - \varepsilon)$ of the time at β_{cold} . Intuitively, this strategy can be efficient because crossing valleys between peaks becomes feasible (because a bacterium spends a fraction ε of the time at β_{hot} , where it can overcome potential barriers), and it can also localize efficiently around each peak when at β_{cold} .

To quantitatively test the importance of tempering, we will analyze two scenarios: colonization of multiple patches and escape from an unfavorable patch.

For the first scenario, which we call the *serial patch model*, we will consider an environment that consists of a 1D array of Gaussians with standard deviation σ and with maxima placed at $\mu_n = 2n\mu, n \in \mathbb{Z}$: $L_n(x) = e^{-\frac{\beta}{2\sigma^2}(x-2n\mu)^2}$ (Figure 4A). The midway points between patches are at distance μ from each peak, which is large compared with the standard deviation $\mu/\sigma \gg 1$ (i.e., the patches are separated). We assume that the performance of an individual bacteria (or its lineage), f , depends on the product of how many patches the bacteria visited f_v and its performance over each patch f_p , that is, $f = f_v f_p$. The migration of bacteria in this model is limited by metastability at the midway points. The crossing rate between patches can therefore be estimated by the Kramers approximation (Kramers, 1940):

$$\tau_{escape}^{-1} \approx \frac{\sqrt{-U''(0)U''(\mu)}}{2\pi} e^{-\frac{U(\mu)-U(0)}{D}} \approx \frac{\beta D \sqrt{\mu^2 - \sigma^2}}{2\pi\sigma^3} 2^\beta e^{-\frac{\mu^2}{2\sigma^2}\beta} \quad (\text{Equation 7})$$

where U is the potential function $U = -\chi \log \left(e^{-\frac{\chi}{2}(x-2\mu)^2} + e^{-\frac{\chi}{2}x^2} \right)$ and $\chi = \beta D$. In Figure 4B we plot the value of τ_{escape}^{-1} for various values of μ, σ and β . One can see that the crossing rate τ_{escape}^{-1} is very sensitive to the choice of β (see also Figure S2 for estimation of τ_{escape} for a wide range of μ, σ and β).

As τ_{escape}^{-1} is the rate-limiting step for the colonization of new patches, we can model the colonization process as random walk where the bacteria have probability $p = \tau_{escape}^{-1}$ to move to each of the adjacent patches. The number of patches visited up to time t (divided by \sqrt{t}) will be on the order of the mean displacement:

$$f_v \approx \sqrt{2p} = \sqrt{\frac{\beta D \sqrt{\mu^2 - \sigma^2}}{\pi} 2^\beta e^{-\frac{\mu^2}{4\sigma^2}\beta}} \quad (\text{Equation 8})$$

Note that aforementioned equation can be generalized for patches with various separation distances (due to different variance or mean position) by taking the geometric mean of the values of τ_{escape}^{-1} .

For the performance in each patch, we consider a performance metric that depends on localization around the peak: the overlap integral between the location of the bacteria $P(x)$ and the attractant peak of $L(x)$ ($\delta \ll 1$):

$$f_p = \frac{\int_{-\delta}^{\delta} P(x) dx}{\int_{-\delta}^{\delta} L(x) dx} = \frac{\int_{-\delta}^{\delta} \sqrt{\frac{\beta}{2\pi\sigma^2}} e^{-\frac{\beta}{2\sigma^2}x^2} dx}{\int_{-\delta}^{\delta} \sqrt{\frac{1}{2\pi\sigma^2}} e^{-\frac{1}{2\sigma^2}x^2} dx} \approx \sqrt{\beta} \quad (\text{Equation 9})$$

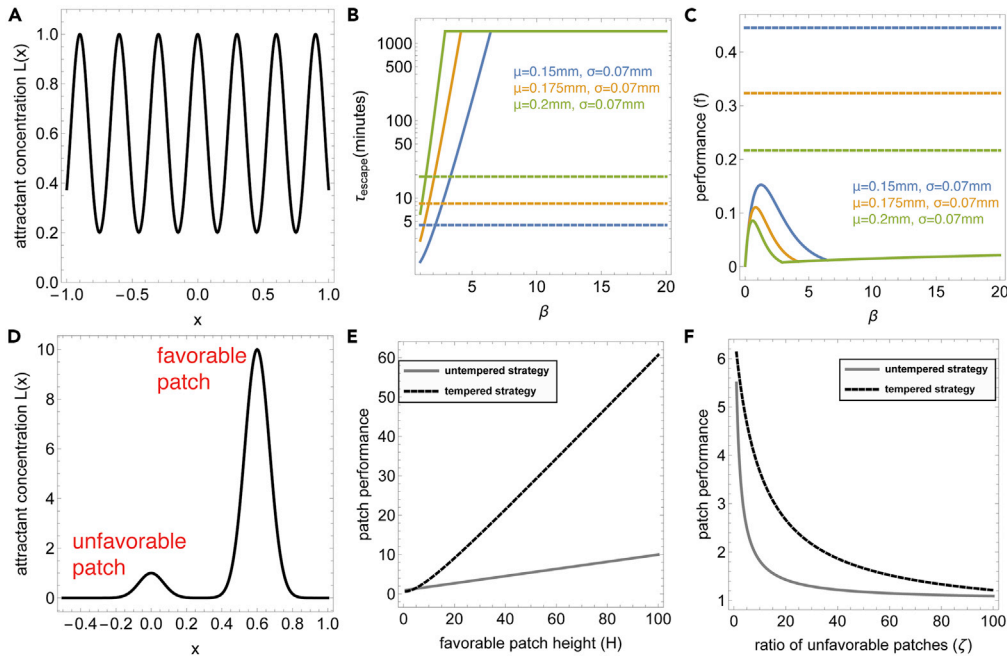


Figure 4. Stochastic tempering provides a superior strategy for patch colonization in multi-modal environments

(A) The serial patch model is an infinite mixture of Gaussian attractant patches with standard deviation σ and maxima placed at $\dots -4\mu, -2\mu, 0, 2\mu, 4\mu, \dots$. Bacteria start at $x = 0$.
 (B) Time to cross over to an adjacent patch τ_{escape} increases exponentially with β for the untempered strategy (solid lines, plotted for various μ). For the tempered strategy crossing time is determined by β_{hot} (dashed lines, plotted for $\beta_{\text{hot}} = 1$, $\beta_{\text{cold}} = 20$, $\varepsilon = \frac{1}{3}$, $D = 0.0009 \text{ mm}^2/\text{sec}$). We assume a patch lifetime of $\tau_{\text{patch}} = 1 \text{ day}$.
 (C) Performance f for the untempered strategy as a function of β (solid lines) decays to zero at low and high values of β due to the trade-off between colonization and localization. The untempered strategy outperforms it by decoupling colonization and localization (dashed lines).
 (D) The asymmetric patch model is an infinite mixture of small unfavorable Gaussian patches (height $h = 1$) and large favorable Gaussian patches (height $h = H \gg 1$). For panels (E and F) we denote by ζ the ratio between unfavorable and favorable patches.
 (E) Expected value of patch performance (normalized by $\sqrt{\beta_{\text{cold}}}$ or $\sqrt{\beta}$) for tempered (black) and untempered (gray) strategies for various values of H setting $\zeta = 10$ (other parameters are as in other simulations).
 (F) Normalized expected value of patch performance as a function of ζ , setting $H = 10$.

We can also consider other performance metrics—the important aspect is that β improves localization around the peak, and thus improves f_p . The overall performance is therefore:

$$f = f_p f_v = \beta \sqrt{\frac{D\sqrt{\mu^2 - \sigma^2}}{\pi}} 2^{\frac{\beta}{2}} e^{-\frac{\mu^2}{4\sigma^2}\beta} \quad (\text{Equation 10})$$

The performance rises and then falls as a function of β (Figure 4C, solid lines).

The equation for f reveals the fundamental trade-off between colonization and localization—as β is important for localization but detrimental for colonization, it is impossible to optimize both f_p and f_v at the same time. Moreover, the optimal β is highly sensitive to μ (which is contingent on the environment), and a choice of β that is too low or too high results in poor performance.

Now consider a tempered strategy with stochastic switches between $\beta_{\text{hot}}, \beta_{\text{cold}}$. We make the conservative assumption that the time the bacteria spends at β_{cold} makes no contribution to escaping from patches, and thus:

$$\tau_{\text{escape}}^{-1}(\mu) \approx \varepsilon \frac{\beta_{\text{hot}} D \mu}{2\pi\sigma^3} 2^{\frac{\beta_{\text{hot}}}{2}} e^{-\frac{\mu^2}{2\sigma^2}\beta_{\text{hot}}} \quad (\text{Equation 11})$$

This holds as long as the bacteria spends enough time in β_{hot} to escape (i.e., η_2 is smaller than $\tau_{\text{escape}}^{-1}$). Thus:

$$f_v \approx \sqrt{\varepsilon \frac{\beta_{\text{hot}} D \sqrt{\mu^2 - \sigma^2}}{\pi} 2^{\frac{\beta_{\text{hot}}}{2}} e^{-\frac{\mu^2}{4\sigma^2} \beta_{\text{hot}}}} \quad (\text{Equation 12})$$

We also make the conservative assumption that the time the bacteria spends at β_{hot} makes no contribution to patch performance, and thus patch performance is:

$$f_p \approx (1 - \varepsilon) \sqrt{\beta_{\text{cold}}} \quad (\text{Equation 13})$$

The overall performance is

$$f = f_p f_v = (1 - \varepsilon) \sqrt{\varepsilon} \sqrt{\beta_{\text{cold}}} \sqrt{\varepsilon \frac{\beta_{\text{hot}} D \sqrt{\mu^2 - \sigma^2}}{\pi} 2^{\frac{\beta_{\text{hot}}}{2}} e^{-\frac{\mu^2}{4\sigma^2} \beta_{\text{hot}}}} \quad (\text{Equation 14})$$

We can see a clear difference between the performance of the tempered strategy (Equation 14) and the untempered strategy (Equation 10)—the tempered strategy allows one to optimize separately f_p (by increasing β_{cold}) and f_v (by decreasing β_{hot}) at a fixed cost of the pre-factor $(1 - \varepsilon)\sqrt{\varepsilon}$. Therefore, choosing a low β_{hot} and a high β_{cold} (independent of μ) can easily outperform the untempered strategy (Figure 4C, dashed lines). Tempering therefore allows for efficient balancing of colonization and patch utilization.

The second scenario is the *asymmetric patch model*, which explores the problem of global optimization. Consider two Gaussian patches with standard deviation σ and with maxima placed at $0, 2\mu$, and with very different heights $h = 1, h = H$, so $H \gg 1$ (Figure 4D). We assume that the performance increases proportionally to patch height (which corresponds to attractant concentration). We can compute the transition rate from the unfavorable patch ($h = 1$) to the favorable patch ($h = H$):

$$\tau_{\text{escape}}^{-1} \approx \frac{\beta D \sqrt{h\mu^2 - \frac{1}{2}(1+h)^2 \sigma^2}}{(1+h)\pi\sigma^3} (1+h)^\beta e^{-\frac{\mu^2}{2\sigma^2} \beta} \quad (\text{Equation 15})$$

The exponential dependence of the escape rate on β again means that for well-separated patches (large μ/σ) it is unfeasible to leave the unfavorable patch when β is large. As the bacteria can in principle get trapped in either of the patches, the expected value for the performance is $f_p \approx \frac{1}{2}(H + 1)\sqrt{\beta}$.

For the tempered strategy, on the other hand, the bacteria can cross between the patches in the β_{hot} state and localize effectively in the β_{cold} state. From Equation 5, the relative time it spends at the favorable patch is $1/(1 + H^{-\beta_{\text{hot}}})$ (we again assume that patch crossing occurs only in the β_{hot} state), so the performance becomes proportional: $f_p = (1 - \varepsilon) \left(\frac{H}{1 + H^{-\beta_{\text{hot}}}} + \frac{1}{1 + H^{\beta_{\text{hot}}}} \right) \sqrt{\beta_{\text{cold}}}$. Taking, for example, $\beta = \beta_{\text{cold}}, \beta_{\text{hot}} = 1$, and $H \gg 1$, we find $f_p \approx (1 - \varepsilon) H \sqrt{\beta_{\text{cold}}}$. Thus, for two patches, the tempered strategy is better when the cost of tempering ε is smaller than $1/2$.

What about a more general case, where there are only a few favorable patches $n_{h=H}$ among many unfavorable patches $n_{h=1}$ (i.e. $\frac{n_{h=1}}{n_{h=H}} = \zeta \gg 1$)? The untempered strategy has a performance of $\frac{f_p}{\sqrt{\beta_{\text{cold}}}} = \frac{1}{\zeta^{-1} + 1} + \frac{1}{\zeta + 1} H \rightarrow 1$, because it is unlikely that the starting patch will be a favorable one. On the other hand, for the tempered strategy, we have $\frac{f_p}{\sqrt{\beta_{\text{cold}}}} = (1 - \varepsilon) \left(\frac{1}{\zeta^{-1} + 1} + \frac{H}{\zeta H^{-1} + 1} \right) = (1 - \varepsilon) \frac{H^2 + \zeta}{H + \zeta}$. This performance depends on the ratio of H/ζ —a very favorable patch ($H \gg \zeta \gg 1$) can yield high performance $\frac{f_p}{\sqrt{\beta_{\text{cold}}}} = (1 - \varepsilon) H$ even when there are many unfavorable patches in the environment. Tempering therefore allows for efficient utilization of favorable patches in the presence of other (potentially many) unfavorable patches (Figures 4E and 4F).

In conclusion, these models suggest that the tempered strategy can yield superior performance in the presence of multiple attractant patches.

This quantitative theory was developed for static attractant patches. In reality, patches have a finite lifetime τ_{patch} due to their consumption, diffusion/advection, or environmental effects. The timescale τ_{patch} is variable—for marine sources, for example, it can range from seconds to minutes for patches of dissolved

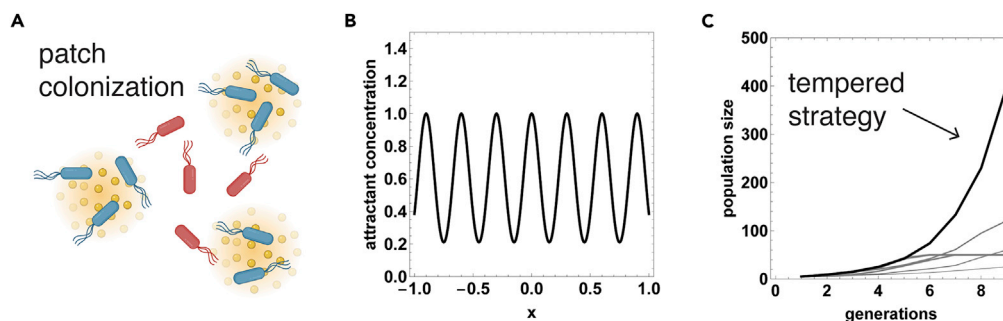


Figure 5. Tempering improves growth and colonization of multi-modal environments in population simulations
 (A) A simple model of patch colonization, where bacteria migrate and colonize attractant patches. Each simulated bacteria performs chemotaxis. At the end of each generation-interval, the bacteria replicates if local attractant concentration is higher than I_{rep} , and if the patch is occupied by less than K bacteria. Here $I_{rep} = 0.95$ and $K = 50$. Red bacteria represent the β_{hot} state, whereas blue bacteria represent the β_{cold} state.
 (B) Simulations for a 1D case where the attractant patches are Gaussians with height = 1 and standard deviation $\sigma = 0.07$ mm (see Figure S3 for other parameter choices). $G = 21$ patches are placed at $\mu = 0.15$ mm distance from each other.
 (C) Constant- β (constant gain) strategies (gray lines, thicker lines correspond to higher β , namely, $\beta = 2, 3, 4, 5, 9, 10, 1, 17.5, 30$) either fail to localize around peaks (low β) or suffer from metastability and fail to colonize new patches (high β). This results in lower overall population growth. The simulated tempering strategy (black line) localizes well around peaks and can colonize new patches, yielding larger population growth.

organic matter to hours to days for more persistent sources such as marine aggregates or phytoplankton (Kjørboe et al., 2002; Stocker et al., 2008). Patch lifetime can be incorporated into the model. It sets a maximal value for τ_{escape} : $\tau_{escape} < \tau_{patch}$. Given this value, it is possible to determine whether the cost of tempering (the prefactor of Equation 14) indeed outweighs the cost of using an untempered strategy, with the escape rate now capped by τ_{patch} . In Figures 4B and 4C we set $\tau_{patch} = 1$ day, and in Figure S3 we plotted τ_{escape} and f taking $\tau_{patch} = 10$ min and $\tau_{patch} = 1$ week.

Tempering provides efficient growth and colonization in complex environments

In addition to the theoretical analysis of tempering in a patchy environment, we also simulated a minimal model for bacterial population dynamics (Figure 5A). We assume that the environment consists of G equal-height Gaussian patches with peaks at $0, 2\mu, 4\mu, \dots, 2(G-1)\mu$, and standard deviation σ (Figure 5B). The bacteria navigate and grow in this environment, such that after every generation-interval (t_{gen}) the bacteria replicate if they sense an attractant level higher than a threshold I_{rep} (accumulation near peak). Bacteria start at $x = 0$. To encourage colonization of new patches, we also assume that each patch has a finite carrying capacity K . For simplicity we again assume a static attractant profile; the carrying capacity may then be due to the nutrient depletion, while the attractant is a separate navigational cue that is present at large quantities.

Constant- β strategies perform poorly in such environments (Figure 5C). If β is high, the bacteria localize effectively near the first local maximum of the attractant. They therefore initially grow rapidly; however, the finite carrying capacity K of the patch prevents further growth, and, due to metastability, they only very slowly colonize other patches. This results in limited population expansion. For the bacteria to colonize other patches, β needs to be lowered, but this prevents effective localization near the peak and so also compromises growth.

Bacteria with a tempering strategy that temporally switch between high and low β , on the other hand, can effectively grow and colonize (Figure 5C). During the time the bacteria spend at high β , they effectively accumulate at a local peak. Spending time at low β helps the bacteria cross valleys, thus colonizing new patches. This results in superior overall growth rate. Importantly, the strategy is robust to the geometry and size of the attractant peaks—spending some time at a low enough β prevents metastability for diverse environments, even if the average β is very high. As we noted earlier, it is only important that the bacteria spend long enough in β_{hot} to allow for patch crossing. As we showed analytically, the choice of $\beta_{hot}, \beta_{cold}$ does not need to be tuned for a particular environment, as the simple strategy of switching between high and low β either outperforms or performs similarly to constant- β strategies for diverse environments (Figure S4).

Theory on LMC sampling can provide insight into chemotaxis strategies

So far, we have shown that bacterial chemotaxis is analogous to LMC-based sampling and used this analogy to demonstrate the importance of stochastic fluctuations in pathway gain for efficient navigation in complex environments. We demonstrated this for an environment that consists of multiple Gaussian attractant patches. We conclude by developing the analogy between patch colonization and sampling more formally, which allows us to generalize our results.

For an attractant ligand distribution $L(x)$ we can define the target probability distribution of the bacteria as $l_\beta(x) \propto L(x)^\beta$, which is the invariant distribution of Equation 4. Let $\rho_t(x)$ be the distribution of locations sampled by the bacteria by time t . We are interested in how fast $\rho_t(x)$ converges to $l_\beta(x)$. For this we need to measure how different $\rho_t(x)$ is from $l_\beta(x)$. One way to do this is to use the Kullback-Leibler (KL) divergence of $\rho_t(x)$ with respect to $l_\beta(x)$:

$$\mathcal{H}_{l_\beta}(\rho_t) = \int \rho_t(x) \log \frac{\rho_t(x)}{l_\beta(x)} dx \quad (\text{Equation 16})$$

KL divergence is a measure of *relative entropy*. It is non-negative, and minimized $\mathcal{H}_{l_\beta}(\rho_t) = 0$ at $\rho_t = l_\beta$. There is extensive mathematical literature on the convergence of the KL divergence for Langevin processes (Vilani, 2008). An important result is that Langevin dynamics can be viewed as optimization (or steepest descent) of $\mathcal{H}_{l_\beta}(\rho_t)$ in the space of probability distributions (Jordan et al., 1998; Wibisono, 2018), with dynamics given by:

$$\frac{d\mathcal{H}_{l_\beta}(\rho_t)}{dt} = -J_\beta(\rho_t) \quad (\text{Equation 17})$$

where J_β is the relative Fisher information:

$$J_\beta(\rho_t) = \int \rho_t(x) \|\nabla \log \frac{\rho_t(x)}{l_\beta(x)}\|^2 dx \quad (\text{Equation 18})$$

As can be seen from the aforementioned equations, convergence can be rapid if we can bound the KL divergence with the Fisher information:

$$\mathcal{H}_{l_\beta}(\rho) \leq \frac{1}{2A} J_\beta(\rho) \quad (\text{Equation 19})$$

which yields:

$$\mathcal{H}_{l_\beta}(\rho_t) \leq e^{-2At} \mathcal{H}_{l_\beta}(\rho_0) \quad (\text{Equation 20})$$

That is, the convergence of the KL divergence is exponentially fast. Equation 19 is equivalent to a condition known as the *Logarithmic Sobolev Inequality (LSI)* (Gross, 1975; Bakry and Émery, 1985), with a Sobolev constant A . LSI holds for strongly log-concave distributions (and in particular for Gaussians)— l_β satisfies Equation 19 if it is A -log concave. LSI is also equivalent to a Gaussian-like concentration of l_β (Ledoux, 1999; Gozlan et al., 2011). We can therefore conclude that chemotaxis (either untempered or tempered) can sample such Gaussian-like attractant patches rapidly.

For multi-modal distributions, such as mixtures of Gaussians, this rapid convergence does not hold, as we have seen the time to cross between modes increases exponentially with their distance for an untempered strategy. However, it is possible to prove that a tempered strategy converges rapidly, at least on a mixture of Gaussians with equal variance (Madras and Randall, 2002; Woodard et al., 2009; Lee et al., 2018). The proof, which is due to Lee et al. (2018), is based on a decomposition theorem by Madras and Randall (2002); informally, it states that if the space can be partitioned into subsets in such a way that the LMC process converges rapidly within each subset, and transitions between subsets occur rapidly enough, then the LMC process converges rapidly overall (STAR Methods). For an environment that consists of a collection of Gaussian attractant patches of equal variance, it is possible to define such a partition (Ge et al., 2017; Lee et al., 2018). If the high-temperature process crosses rapidly between the partition sets (that is, β_{hot} is chosen low enough), then the entire process can converge rapidly.

DISCUSSION

We studied the functional significance of variation in pathway gain in bacterial chemotaxis. Gain variation is puzzling, because individuals with low pathway gain climb gradients poorly and accumulate less tightly

around attractant peaks, and thus selection should push towards maximal gain. We considered the possibility that variation is due to (potentially long-lasting) temporal fluctuations. We suggest that such fluctuations in gain are crucial for efficient chemotaxis in complex environments with multiple attractant maxima and minima. This is because a strategy with a single (or tightly distributed) pathway gain either cannot escape local peaks, if gain is high, or localizes poorly around peaks, if gain is low. A variable-gain strategy overcomes both of these problems. This strategy is a biological implementation of the simulated-tempering algorithm used in statistical physics and computer science for sampling complex distributions.

To analyze the function of gain variation, we made a formal analogy between bacterial chemotaxis and the LMC algorithm, a widely used method for sampling probability distributions. Standard LMC often fails to efficiently sample from multi-modal distributions. It gets stuck on local maxima, because the time to cross valleys depends exponentially on the distance between the peaks, and can easily become intractable. Simulated tempering LMC employs a tempering strategy that consists of stochastic switches in the temperature parameter. The tempering strategy makes the crossing of valleys computationally tractable, while sacrificing efficient localization near the peaks for only a fraction of the time. We show that tempering is analogous to stochastic fluctuations in the chemotaxis pathway gain. This allows rapid colonization of new patches by bacteria, while maintaining efficient localization within each patch.

The model presented here is based on the mean-field approximation that averages over the internal state of the bacteria (Si et al., 2012). Although this approximation works well in some parameter regimes (such as shallow gradients), it cannot account for large deviations in internal state that bacteria experience due to positive feedback between run duration and sensing. This was addressed recently in an important theoretical paper (Long et al., 2017). Long et al. showed that the positive feedback between motion and sensing can result in an almost switch-like dynamics, with very long runs when climbing gradients and persistent tumbling when descending the gradient. These dynamics become dominant when the positive feedback timescale $\tau_E = (L/NHv)$ (L is the length scale of the gradient) becomes short compared with adaptation time. This effect causes large variability in the internal state of the bacteria (including receptor methylation), which depends on the behavior of the bacteria, and it improves gradient climbing efficiency and accumulation around peaks. Tempering is therefore still required to allow for efficient crossing between peaks.

The present tempering role for temporal fluctuations in gain adds to previous concepts on the functional benefits of phenotypic heterogeneity. Two well-studied concepts are bet-hedging (Kussell and Leibler, 2005; Wolf et al., 2005; Xue and Leibler, 2017; Martín et al., 2019) and division of labor (Ackermann, 2015; Adler et al., 2019). Bet-hedging is defined as diversification that is beneficial for buffering against uncertain changes in the environment. For instance, diverse behavioral strategies help animal populations to overcome different invasion stages and conditions (Sih et al., 2012; Wolf and Weissing, 2012; Carere and Gherardi, 2013; Forkosh et al., 2019). The tempering strategy studied here is beneficial even when the environment is static. In addition, unlike division of labor, which is a population-level property, the tempering strategy can be beneficial at the level of the individual organism, allowing efficient navigation.

It is interesting to consider how tempering is affected by another important source of heterogeneity in bacteria, namely, cell-cell variation in the adaptation time (Spudich and Koshland, 1976). In addition to the functional roles discussed in the introduction, fluctuations in adaptation time can also complement tempering for efficient navigation in complex environments. This is because the bacterial chemotaxis signaling circuit acts as a low-pass filter on the logarithmic derivative of the input (which controls the accumulation around peaks), with a filtering time-window that is proportional to the adaptation time (Tu et al., 2008). Slower adaptation therefore makes the bacteria “see” smoother inputs, which can also accelerate sampling (Ma et al., 2019), again at the cost of poorer localization around local peaks.

In addition to tempering, other behavioral strategies can also help bacteria escape from local attractant peaks. One such strategy is to increase the variation in run duration. It has been proposed that *E. coli* may perform an approximation of a Lévy walk, where the run durations are drawn from a power-law distribution instead of an exponential distribution, due to slow fluctuations in signaling molecules (Tu and Grinstein, 2005; Matthäus et al., 2009). This is supported by a recent study that characterized *E. coli* run-length distributions using 3D tracking (Huo et al., 2021). Mathematically, it is possible to account for the power-law run distribution in the sampling equation by using fractional differentiation, as shown in a recent article by Şimşekli (2017). Lévy walks are effective for searching for randomly distributed, sparse targets

(Viswanathan et al., 1999); however they may be suboptimal when the target is close, or in the presence of bias (Palyulin et al., 2014). In principle, a Lévy walk can occasionally generate very long runs that can cross between peaks, while most of the time staying close to the peak. However, there is an important difference between the tempering and Lévy walk strategies. The tempering strategy allows the bacteria to cross between peaks, but bacteria remain attracted to areas of high attractant concentration; in contrast, in the Lévy walk the long runs are random and can send the bacteria far from attractant patches, where the gradient may no longer be detectable. The tempering strategy may therefore be preferable when attractant peaks are dense, whereas the Lévy walk strategy may be preferable when they are sparse.

Although the LMC analogy was developed for bacterial navigation, its underlying assumptions are general and can extend to other organisms that combine FCD input sensing with stochastic navigation. As the benefit from tempering (balancing peak localization with exploration of new peaks) can in principle apply to a wide range of navigation systems, we would expect tempering to be a widespread strategy. One potential example for this is the chemotaxis of *Dictyostelium*, which is based on FCD sensing (Janetopoulos et al., 2004; Kamino and Kondo, 2016; Kamino et al., 2017) and stochastic navigation (Amselem et al., 2012). It will be interesting to test whether this system, which exhibits stochastic fluctuations in signaling components (Arai et al., 2010), also performs tempering. Another potential example is the dopamine system in vertebrates, which controls movement and is based on FCD sensing of expected reward (Tobler et al., 2005; Karin and Alon, 2021). The effective inverse temperature for the dopamine system is $\beta \approx 1$, and it is therefore closer to sampling of rewards rather than optimization, in line with experimental observations on choice behavior across vertebrates (McDowell, 2013). In this system, tempering can be implemented by fluctuations in dopamine gain, which may be due to changes in the level of other neuromodulators such as endogenous opioids (Smith et al., 2011).

The analogy between LMC-based sampling and chemotaxis can help address additional fundamental questions in biological navigation. One important outstanding question is why stochastic navigation is so prevalent across different organisms. The run-and-tumble navigation of bacteria is markedly less efficient in climbing gradients than direct reorientation according to the spatial gradient, because run-and-tumble movement is nearly isotropic and uncorrelated with the gradient (Berg, 1993). It is thought that bacteria use run-and-tumble because they cannot measure spatial gradients along their body, as shown by Berg and Purcell (1977) and Bialek and Setayeshgar (2005) (but see also Thar and Kühl [2003]). However, stochastic navigation is also employed by organisms that can sense spatial gradients along their body, such as *Chlamydomonas* and *C. elegans* (Polin et al., 2009; Luo et al., 2014; Pierce-Shimomura et al., 1999). This raises the question of whether stochastic navigation may be preferable under some conditions to direct spatial-gradient reorientation.

The analogy to sampling and optimization approaches from physics and computer science can provide insight into this question. Stochastic navigation is analogous to LMC-based sampling, whereas spatial gradient reorientation is analogous to gradient-descent-based optimization. Gradient descent is commonly used for convex optimization; however, it is notoriously ineffective for global optimization and for optimization in non-convex settings, when compared with Langevin diffusion (Ma et al., 2019; Raginsky et al., 2017; Xu et al., 2018). Comparing trade-offs between sampling and optimization may therefore help understand the navigational strategies employed by organisms.

Limitations of the study

In this study we propose a function for temporal fluctuations in pathway gain—they allow bacteria to effectively balance localization around peaks with navigation to new peaks. It will be important to test this proposal experimentally. In particular it will be important to experimentally quantify the navigation behavior of bacteria in multi-modal environments and how it is influenced by the pathway gain of individual bacteria.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- RESOURCE AVAILABILITY
 - Lead contact
 - Materials availability

- Data and code availability
- EXPERIMENTAL MODEL AND SUBJECT DETAILS
- METHOD DETAILS
 - Stochastic simulations of chemotaxis
 - Heterogeneity in [MaterSection] from bacterial T-maze experiments
 - Adaptation time of tumbling rate
 - Mixing time of chemotaxis strategies
- QUANTIFICATION AND STATISTICAL ANALYSIS
 - Estimating the inverse temperature [MaterSection] from chemotaxis in linear gradients

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at <https://doi.org/10.1016/j.isci.2021.102796>.

ACKNOWLEDGMENTS

We thank Roman Stocker and Francesco Carrara for reading the manuscript and for their insightful feedback. Figures were created with [BioRender.com](https://www.biorender.com). This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No 856487). U.A. is the incumbent of the Abisch-Frenkel Professorial Chair, and O.K. is a JSMF postdoctoral fellow.

AUTHOR CONTRIBUTIONS

O.K. and U.A. conceptualized and wrote the paper.

DECLARATION OF INTEREST

The authors declare no conflict of interest.

Received: January 12, 2021

Revised: March 29, 2021

Accepted: June 24, 2021

Published: July 23, 2021

REFERENCES

- Ackermann, M. (2015). A functional perspective on phenotypic heterogeneity in microorganisms. *Nat. Rev. Microbiol.* *13*, 497–508.
- Adler, M., and Alon, U. (2018). Fold-change detection in biological systems. *Curr. Opin. Syst. Biol.* *8*, 81–89.
- Adler, M., Kohanim, Y.K., Tendler, A., Mayo, A., and Alon, U. (2019). Continuum of gene-expression profiles provides spatial division of labor within a differentiated cell type. *Cell Syst.* *8*, 43–52.
- Alon, U. (2019). *An Introduction to Systems Biology: Design Principles of Biological Circuits* (CRC press).
- Amselem, G., Theves, M., Bae, A., Bodenschatz, E., and Beta, C. (2012). A stochastic description of Dictyostelium chemotaxis. *PLoS One* *7*, e37213.
- Arai, Y., Shibata, T., Matsuoka, S., Sato, M.J., Yanagida, T., and Ueda, M. (2010). Self-organization of the phosphatidylinositol lipids signaling system for random cell migration. *Proc. Natl. Acad. Sci.* *107*, 12399–12404.
- Arrieta, J., Barreira, A., Chioccioli, M., Polin, M., and Tuval, I. (2017). Phototaxis beyond turning: persistent accumulation and response acclimation of the microalga *Chlamydomonas reinhardtii*. *Sci. Rep.* *7*, 3447.
- Bakry, D., and Émery, M. (1985). Diffusions hypercontractives. *Séminaire de Probabilités XIX 1983/84* (Springer), pp. 177–206.
- Berg, H.C. (1993). *Random Walks in Biology*, Expanded ed. (Princeton University Press).
- Berg, H.C., and Purcell, E.M. (1977). Physics of chemoreception. *Biophys. J.* *20*, 193–219.
- Bialek, W., and Setayeshgar, S. (2005). Physical limits to biochemical signaling. *Proc. Natl. Acad. Sci.* *102*, 10040–10045.
- Borba, C., Schwennicke, S., Kourakis, M.J., and Smith, W.C. (2020). Deep origins of chordate midbrain visual processing centers. *bioRxiv*.
- Carere, C., and Gherardi, F. (2013). Animal personalities matter for biological invasions. *Trends Ecol. Evol.* *1*, 5–6.
- Celani, A., and Vergassola, M. (2010). Bacterial strategies for chemotaxis response. *Proc. Natl. Acad. Sci.* *107*, 1391–1396.
- Chen, Y., Chen, J., Dong, J., Peng, J., and Wang, Z. (2020). Accelerating Nonconvex Learning via Replica Exchange Langevin Diffusion (ArXiv), p. 200701990, *Cs Math Stat*.
- Chiang, T.-S., Hwang, C.-R., and Sheu, S.J. (1987). Diffusion for global optimization in \mathbb{R}^n . *SIAM J. Control Optim.* *25*, 737–753.
- Clark, D.A., and Grant, L.C. (2005). The bacterial chemotactic response reflects a compromise between transient and steady-state behavior. *Proc. Natl. Acad. Sci.* *102*, 9150–9155.
- Colin, R., and Sourjik, V. (2017). Emergent properties of bacterial chemotaxis pathway. *Curr. Opin. Microbiol.* *39*, 24–33.
- Dalalyan, A.S. (2014). Theoretical Guarantees for Approximate Sampling from Smooth and Log-Concave Densities (ArXiv Prepr), *ArXiv14127392*.
- Dufour, Y.S., Fu, X., Hernandez-Nunez, L., and Emonet, T. (2014). Limits of feedback control in bacterial chemotaxis. *PLoS Comput. Biol.* *10*, e1003694.
- Dufour, Y.S., Gillet, S., Frankel, N.W., Weibel, D.B., and Emonet, T. (2016). Direct correlation between motile behavior and protein abundance in single cells. *PLoS Comput. Biol.* *12*, e1005041.

- Earl, D.J., and Deem, M.W. (2005). Parallel tempering: theory, applications, and new perspectives. *Phys. Chem. Chem. Phys.* **7**, 3910–3916.
- Erdogdu, M.A., Mackey, L., and Shamir, O. (2018). Global non-convex optimization with discretized diffusions. In *Advances in Neural Information Processing Systems 31*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, eds. (Curran Associates, Inc), pp. 9671–9680.
- Forkosh, O., Karamihalev, S., Roeh, S., Alon, U., Anpilov, S., Touma, C., Nussbaumer, M., Flachskamm, C., Kaplick, P.M., and Shemesh, Y. (2019). Identity domains capture individual differences from across the behavioral repertoire. *Nat. Neurosci.* **22**, 2023–2028.
- Frankel, N.W., Pontius, W., Dufour, Y.S., Long, J., Hernandez-Nunez, L., and Emonet, T. (2014). Adaptability of non-genetic diversity in bacterial chemotaxis. *Elife* **3**, e03526.
- Ge, R., Lee, H., and Risteski, A. (2017). Beyond log-concavity: provable guarantees for sampling multi-modal distributions using simulated tempering Langevin Monte Carlo (ArXiv), p. 1711002736, *Cs Math Stat*.
- Gelfand, S.B., and Mitter, S.K. (1991). Recursive stochastic algorithms for global optimization in \mathbb{R}^d . *SIAM J. Control Optim.* **29**, 999–1018.
- Girolami, M., and Calderhead, B. (2011). Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **73**, 123–214.
- Gozlan, N., Roberto, C., and Samson, P.-M. (2011). From concentration to logarithmic Sobolev and Poincaré inequalities. *J. Funct. Anal.* **260**, 1491–1522.
- Gross, L. (1975). Logarithmic sobolev inequalities. *Am. J. Math.* **97**, 1061–1083.
- Hansmann, U.H. (1997). Parallel tempering algorithm for conformational studies of biological molecules. *Chem. Phys. Lett.* **281**, 140–150.
- Hu, B., and Tu, Y. (2014). Behaviors and strategies of bacterial navigation in chemical and nonchemical gradients. *PLoS Comput. Biol.* **10**, e1003672.
- Hukushima, K., and Nemoto, K. (1996). Exchange Monte Carlo method and application to spin glass simulations. *J. Phys. Soc. Jpn.* **65**, 1604–1608.
- Huo, H., He, R., Zhang, R., and Yuan, J. (2021). Swimming *Escherichia coli* explore the environment by Lévy walk. *Appl. Environ. Microbiol.* **87**, e02429–20.
- Janetopoulos, C., Ma, L., Devreotes, P.N., and Iglesias, P.A. (2004). Chemoattractant-induced phosphatidylinositol 3, 4, 5-trisphosphate accumulation is spatially amplified and adapts, independent of the actin cytoskeleton. *Proc. Natl. Acad. Sci.* **101**, 8951–8956.
- Jordan, R., Kinderlehrer, D., and Otto, F. (1998). The variational formulation of the Fokker–Planck equation. *SIAM J. Math. Anal.* **29**, 1–17.
- Kalinin, Y.V., Jiang, L., Tu, Y., and Wu, M. (2009). Logarithmic sensing in *Escherichia coli* bacterial chemotaxis. *Biophys. J.* **96**, 2439–2448.
- Kamino, K., Keestra, J.M., Long, J., Emonet, T., and Shimizu, T.S. (2020). Adaptive tuning of cell sensory diversity without changes in gene expression. *Sci. Adv.* **6**, eabc1087.
- Kamino, K., and Kondo, Y. (2016). Rescaling of spatio-temporal sensing in eukaryotic chemotaxis. *PLoS One* **11**, e0164674.
- Kamino, K., Kondo, Y., Nakajima, A., Honda-Kitahara, M., Kaneko, K., and Sawai, S. (2017). Fold-change detection and scale invariance of cell–cell signaling in social amoeba. *Proc. Natl. Acad. Sci.* **114**, E4149–E4157.
- Karin, O., and Alon, U. (2021). The dopamine circuit as a reward-taxis navigation system. *bioRxiv*.
- Keestra, J.M., Kamino, K., Anquez, F., Lazova, M.D., Emonet, T., and Shimizu, T.S. (2017). Phenotypic diversity and temporal variability in a bacterial signaling network revealed by single-cell FRET. *Elife* **6**, e27455.
- Kjørboe, T., Grossart, H.-P., Ploug, H., and Tang, K. (2002). Mechanisms and rates of bacterial colonization of sinking aggregates. *Appl. Environ. Microbiol.* **68**, 3996–4006.
- Kirkpatrick, S., Gelatt, C.D., and Vecchi, M.P. (1983). Optimization by simulated annealing. *science* **220**, 671–680.
- Kramers, H.A. (1940). Brownian motion in a field of force and the diffusion model of chemical reactions. *Physica* **7**, 284–304.
- Kussell, E., and Leibler, S. (2005). Phenotypic diversity, population growth, and information in fluctuating environments. *Science* **309**, 2075–2078.
- Larsch, J., Flavell, S.W., Liu, Q., Gordus, A., Albrecht, D.R., and Bargmann, C.I. (2015). A circuit for gradient climbing in *C. elegans* chemotaxis. *Cell Rep.* **12**, 1748–1760.
- Lazova, M.D., Ahmed, T., Bellomo, D., Stocker, R., and Shimizu, T.S. (2011). Response rescaling in bacterial chemotaxis. *Proc. Natl. Acad. Sci.* **108**, 13870–13875.
- Ledoux, M. (1999). Concentration of measure and logarithmic Sobolev inequalities. In *Seminaire de probabilités XXXIII* (Springer), pp. 120–216.
- Lee, H., Risteski, A., and Ge, R. (2018). Beyond log-concavity: provable guarantees for sampling multi-modal distributions using simulated tempering Langevin Monte Carlo. In *Advances in Neural Information Processing Systems 31*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, eds. (Curran Associates, Inc), pp. 7847–7856.
- Levin, D.A., and Peres, Y. (2017). Markov Chains and Mixing Times (American Mathematical Soc).
- Levin, M.D., Morton-Firth, C.J., Abouhamad, W.N., Bourret, R.B., and Bray, D. (1998). Origins of individual swimming behavior in bacteria. *Biophys. J.* **74**, 175–181.
- Long, J., Zucker, S.W., and Emonet, T. (2017). Feedback between motion and sensation provides nonlinear boost in run-and-tumble navigation. *PLoS Comput. Biol.* **13**, e1005429.
- Luo, L., Cook, N., Venkatachalam, V., Martinez-Velazquez, L.A., Zhang, X., Calvo, A.C., Hawk, J., MacInnis, B.L., Frank, M., Ng, J.H.R., et al. (2014). Bidirectional thermotaxis in *Caenorhabditis elegans* is mediated by distinct sensorimotor strategies driven by the AFD thermosensory neurons. *Proc. Natl. Acad. Sci. U S A* **111**, 2776–2781.
- Lyubartsev, A.P., Martsinovski, A.A., Shevkunov, S.V., and Vorontsov-Velyaminov, P.N. (1992). New approach to Monte Carlo calculation of the free energy: method of expanded ensembles. *J. Chem. Phys.* **96**, 1776–1783.
- Ma, Y.-A., Chen, Y., Jin, C., Flammarion, N., and Jordan, M.I. (2019). Sampling can be faster than optimization. *Proc. Natl. Acad. Sci.* **116**, 20881–20885.
- Madras, N., and Randall, D. (2002). Markov chain decomposition for convergence rate analysis. *Ann. Appl. Probab.* **12**, 581–606.
- Marinari, E., and Parisi, G. (1992). Simulated tempering: a new Monte Carlo scheme. *Europhys Lett.* **19**, 451–458.
- Martin, P.V., Muñoz, M.A., and Pigolotti, S. (2019). Bet-hedging strategies in expanding populations. *PLoS Comput. Biol.* **15**, e1006529.
- Matthäus, F., Jagodic, M., and Dobnikar, J. (2009). *E. coli* superdiffusion and chemotaxis-search strategy, precision, and motility. *Biophys. J.* **97**, 946–957.
- McDowell, J.J. (2013). On the theoretical and empirical status of the matching law and matching theory. *Psychol. Bull.* **139**, 1000.
- Menolascina, F., Rusconi, R., Fernandez, V.I., Smruga, S., Aminzare, Z., Sontag, E.D., and Stocker, R. (2017). Logarithmic sensing in *Bacillus subtilis* aerotaxis. *NPJ Syst. Biol. Appl.* **3**, 16036.
- Neal, R.M. (2011). MCMC using Hamiltonian dynamics. *Handb Markov Chain Monte Carlo* **2**, 2.
- Palyulin, V.V., Chechkin, A.V., and Metzler, R. (2014). Lévy flights do not always optimize random blind search for sparse targets. *Proc. Natl. Acad. Sci.* **111**, 2931–2936.
- Park, H., Pontius, W., Guet, C.C., Marko, J.F., Emonet, T., and Cluzel, P. (2010). Interdependence of behavioural variability and response to small stimuli in bacteria. *Nature* **468**, 819–823.
- Pierce-Shimomura, J.T., Morse, T.M., and Lockery, S.R. (1999). The fundamental role of pirouettes in *Caenorhabditis elegans* chemotaxis. *J. Neurosci.* **19**, 9557–9569.
- Polin, M., Tuval, I., Drescher, K., Gollub, J.P., and Goldstein, R.E. (2009). *Chlamydomonas* swims with two “gears” in a eukaryotic version of run-and-tumble locomotion. *Science* **325**, 487–490.
- Raginsky, M., Rakhlin, A., and Telgarsky, M. (2017). Non-convex Learning via Stochastic Gradient Langevin Dynamics: A Nonasymptotic Analysis (ArXiv), p. 170203849, *Cs Math Stat*.

- Roberts, G.O., and Tweedie, R.L. (1996). Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli* 2, 341–363.
- Salek, M.M., Carrara, F., Fernandez, V., Guasto, J.S., and Stocker, R. (2019). Bacterial chemotaxis in a microfluidic T-maze reveals strong phenotypic heterogeneity in chemotactic sensitivity. *Nat. Commun.* 10, 1877.
- Shimizu, T.S., Tu, Y., and Berg, H.C. (2010). A modular gradient-sensing network for chemotaxis in *Escherichia coli* revealed by responses to time-varying stimuli. *Mol. Syst. Biol.* 6, 382.
- Shoval, O., Goentoro, L., Hart, Y., Mayo, A., Sontag, E., and Alon, U. (2010). Fold-change detection and scalar symmetry of sensory input fields. *Proc. Natl. Acad. Sci.* 107, 15995–16000.
- Si, G., Wu, T., Ouyang, Q., and Tu, Y. (2012). Pathway-based mean-field model for *Escherichia coli* chemotaxis. *Phys. Rev. Lett.* 109, 048101.
- Sih, A., Cote, J., Evans, M., Fogarty, S., and Pruitt, J. (2012). Ecological implications of behavioural syndromes. *Ecol. Lett.* 15, 278–289.
- Şimşekli, U. (2017). Fractional Langevin Monte Carlo: exploring lévy driven stochastic differential equations for Markov chain Monte Carlo. In *International Conference on Machine Learning (PMLR)*, pp. 3200–3209.
- Smith, K.S., Berridge, K.C., and Aldridge, J.W. (2011). Disentangling pleasure from incentive salience and learning signals in brain reward circuitry. *Proc. Natl. Acad. Sci.* 108, E255–E264.
- Sourjik, V., and Berg, H.C. (2002). Receptor sensitivity in bacterial chemotaxis. *Proc. Natl. Acad. Sci.* 99, 123–127.
- Spudich, J.L., and Koshland, D.E. (1976). Non-genetic individuality: chance in the single cell. *Nature* 262, 467–471.
- Stocker, R., Seymour, J.R., Samadani, A., Hunt, D.E., and Polz, M.F. (2008). Rapid chemotactic response enables marine bacteria to exploit ephemeral microscale nutrient patches. *Proc. Natl. Acad. Sci.* 105, 4209–4214.
- Swendsen, null, and Wang, null (1986). Replica Monte Carlo simulation of spin glasses. *Phys. Rev. Lett.* 57, 2607–2609.
- Thar, R., and Köhl, M. (2003). Bacteria are not too small for spatial sensing of chemical gradients: an experimental evidence. *Proc. Natl. Acad. Sci.* 100, 5748–5753.
- Tobler, P.N., Fiorillo, C.D., and Schultz, W. (2005). Adaptive coding of reward value by dopamine neurons. *Science* 307, 1642–1645.
- Tu, Y., and Grinstein, G. (2005). How white noise generates power-law switching in bacterial flagellar motors. *Phys. Rev. Lett.* 94, 208101.
- Tu, Y., Shimizu, T.S., and Berg, H.C. (2008). Modeling the chemotactic response of *Escherichia coli* to time-varying stimuli. *Proc. Natl. Acad. Sci.* 105, 14855–14860.
- Van Laarhoven, P.J., and Aarts, E.H. (1987). Simulated annealing. In *Simulated Annealing: Theory and Applications (Springer)*, pp. 7–15.
- Villani, C. (2008). *Optimal Transport: Old and New (Springer Science & Business Media)*.
- Viswanathan, G.M., Buldyrev, S.V., Havlin, S., Da Luz, M.G.E., Raposo, E.P., and Stanley, H.E. (1999). Optimizing the success of random searches. *Nature* 401, 911–914.
- Waite, A.J., Frankel, N.W., Dufour, Y.S., Johnston, J.F., Long, J., and Emonet, T. (2016). Non-genetic diversity modulates population performance. *Mol. Syst. Biol.* 12, 895.
- Waite, A.J., Frankel, N.W., and Emonet, T. (2018). Behavioral variability and phenotypic diversity in bacterial chemotaxis. *Annu. Rev. Biophys.* 47, 595–616.
- Wibisono, A. (2018). Sampling as optimization in the space of measures: the Langevin dynamics as a composite optimization problem. In *Conference on Learning Theory (PMLR)*, pp. 2093–3027.
- Wolf, D.M., Vazirani, V.V., and Arkin, A.P. (2005). Diversity in times of adversity: probabilistic strategies in microbial survival games. *J. Theor. Biol.* 234, 227–253.
- Wolf, M., and Weissing, F.J. (2012). Animal personalities: consequences for ecology and evolution. *Trends Ecol. Evol.* 27, 452–461.
- Woodard, D.B., Schmidler, S.C., and Huber, M. (2009). Conditions for rapid mixing of parallel and simulated tempering on multimodal distributions. *Ann. Appl. Probab.* 19, 617–640.
- Xu, P., Chen, J., Zou, D., and Gu, Q. (2018). Global convergence of Langevin dynamics based algorithms for nonconvex optimization. In *Advances in Neural Information Processing Systems* 31, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, eds. (Curran Associates, Inc), pp. 3122–3133.
- Xue, B., and Leibler, S. (2017). Bet hedging against demographic fluctuations. *Phys. Rev. Lett.* 119, 108103.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Software and algorithms		
Mathematica	https://www.wolfram.com/mathematica/	Version 12.1.1.0

RESOURCE AVAILABILITY

Lead contact

Further information and requests should be directed to and will be fulfilled by the lead contact, Uri Alon (uri.alon@weizmann.ac.il).

Materials availability

The study did not generate any new materials.

Data and code availability

- All data generated are available from the authors upon request.
- Code for the patch colonization simulations (Figures 5 and S4) is available at https://github.com/omerka-weizmann/chemotaxis_tempering.
- No additional information.

EXPERIMENTAL MODEL AND SUBJECT DETAILS

The study did not use experimental models.

METHOD DETAILS

Stochastic simulations of chemotaxis

We simulated the trajectories of bacteria using the Langevin process $dx = \chi \nabla \log L(x) dt + \sigma dW$, with $\sigma = \sqrt{2} \cdot 0.03 \text{ mm}^{1.5} \text{ sec}^{-0.5}$ (corresponding to a run duration of $\tau = 1 \text{ sec}$ and a swimming speed of $v = 0.03 \text{ mm/sec}$, and $D = v^2 \tau$). The advection parameter χ was adjusted according to β , where $\beta = 2\chi/\sigma^2$. All simulations were performed using the ItoProcess procedure of Mathematica, with a step size of 0.05sec. For tempering between β_{hot} , β_{cold} , we assumed stochastic changes with transition rates: $\eta_1 : \beta_{\text{cold}} \rightarrow \beta_{\text{hot}}$, $\eta_2 : \beta_{\text{hot}} \rightarrow \beta_{\text{cold}}$, so transition that times were drawn from exponential distributions with means $\tau_1 = (1/\eta_1)$, $\tau_2 = (1/\eta_2)$. Since stochastic fluctuations may occur on a timescale of minutes to hours, we chose for the simulations $\tau_1 = 6 \text{ min}$ and $\tau_2 = 24 \text{ min}$. The generation interval is $t_{\text{gen}} = 60 \text{ min}$.

Heterogeneity in β from bacterial T-maze experiments

A recent study (Salek et al., 2019) quantified non-genetic variability in pathway gain. While the fitted distribution is a product of several underlying components, we find that it is approximately a Gamma distribution with a shape parameter of 2.25 and a scale parameter of 35. This distribution has a coefficient of variation of approximately 2/3. We accounted for this heterogeneity in Figure 2 by computing the probability distribution $P(x) = \frac{1}{Z} \int_0^{\infty} \text{Prob}[\xi] \frac{L(x)^{\beta \xi}}{\int_{-\infty}^{\infty} L(x)^{\beta \xi} dx} d\xi$, where Z is a normalization factor, ξ is a random variable that is distributed $\xi \sim \text{Gamma}[2.25, 35]$, and β is chosen such that $\langle \beta \cdot \xi \rangle = 14$.

Adaptation time of tumbling rate

To derive the adaptation time for the tumbling rate we linearize Equations 1 and 2 around $a = a_0$ and $m = m_0 = 1 + \frac{1}{\alpha N} \log \frac{a_0}{1-a_0} + \frac{1}{\alpha} \log \frac{1}{K}$. Since tumbling rate adjusts on a timescale that is more rapid than methylation, we take its quasi-steady-state (given by Equation 1), so the equation for m becomes:

$$\dot{m}(t) \approx -(1 - a_0) N \alpha \omega (m - m_0)$$

From which we can see that the adaptation rate is inversely proportional to $(1 - a_0)N\alpha\omega$.

Mixing time of chemotaxis strategies

In order to analyze whether a given chemotaxis strategy is able to cross between peaks, we use a formal analogy between chemotaxis and the LMC algorithm. The LMC algorithm samples an invariant probability distribution that is proportional to the attractant distribution raised to the power β , where β is the inverse temperature of the Langevin process. Whether the chemotaxis strategy is able to cross between modes and localize around attractant peaks is then directly related to the convergence rate of the analogous sampling algorithm to its invariant distribution.

To quantify this convergence rate of the sampling process, we recall that we defined the KL divergence of the sampled distribution ρ_t relative to the invariant distribution l_β :

$$\mathcal{H}_{l_\beta}(\rho_t) = \int \rho_t(x) \log \frac{\rho_t(x)}{l_\beta(x)} dx$$

And that $\mathcal{H}_{l_\beta}(\rho_t) \rightarrow 0$ as $t \rightarrow \infty$. We define the *mixing time* as the time it takes for $\mathcal{H}_{l_\beta}(\rho_t)$ to reach a fraction γ of its original value, e.g. the half-way convergence point for $\gamma = 1/2$. In our case, $l_\beta(x)$ is given by a power β of the attractant ligand distribution. As we have shown in the main text, for an A log-concave attractant distribution such as a Gaussian $l_\beta(x) = e^{-\frac{\beta}{2\sigma^2}x^2}$ (which is log concave with $A = \beta/\sigma^2$), the convergence rate can be bound by:

$$\mathcal{H}_{l_\beta}(\rho_t) \leq e^{-2\alpha t} \mathcal{H}_{l_\beta}(\rho_0)$$

And therefore, the mixing time is polynomial in the length-scale σ .

For the case of two or more Gaussians with means placed at distance 2μ from each other however, the limiting timescale for the mixing time is given by the escape time from each Gaussian,

$$\tau_{\text{escape}} \approx \left(\frac{\beta D \sqrt{\mu^2 - \sigma^2}}{2\pi\sigma^2} 2^\beta \right)^{-1} e^{\frac{\mu^2}{2\sigma^2} \beta}, \text{ which is exponential in the length scale.}$$

What about a tempered strategy? Here we will outline the proof of (Lee et al., 2018) that a tempered strategy over Gaussians of similar variance has polynomial mixing time. For simplicity, we analyze the discrete time case, as the results extend in a straightforward manner to the continuous time case. Following the definitions in (Madras and Randall, 2002; Lee et al., 2018), a Markov Process \mathcal{M} is given by a probability transition kernel $\mathcal{P}(x, dy)$ over a measure space Ω (e.g. $\Omega = \mathbb{R}^n$), that is reversible with respect to a probability density π . The asymptotic convergence rate is related to a quantity that is called the *spectral gap* of the process. The spectral gap is defined as:

$$\text{Gap}(\mathcal{M}) = \inf_f \frac{1}{2} \iint |f(x) - f(y)|^2 \pi(dx) \mathcal{P}(x, dy)$$

Where f is a non-constant function with variance 1 and an expected value of 0 w.r.t π . Note that the quantity that is minimized is the *Dirichlet form* of f .

The spectral gap is interesting because it determines the mixing time (Levin and Peres, 2017). The mixing time is inversely related to the spectral gap, so algorithmic improvement is achieved by increasing the spectral gap.

In our case, we will be interested in giving lower bounds on the spectral gap for Langevin diffusion processes on Gaussian patches, using a decomposability theorem, due to Randal and Madras (Madras and Randall, 2002).

For the decomposability theorem, we assume that we are given a partition $\Phi = \{A_j, j \in 1, \dots, m\}$ of the state space Ω of the Markov chain $\Omega = \cup_{1 \leq j \leq m} A_j$. Let $\mathcal{P}|_A$ denote the restriction of \mathcal{P} to A , so transitions occur according to \mathcal{P} but are rejected if they leave A . Additionally, let $\overline{\mathcal{M}}$ denote the projected Markov chain with transition rates corresponding to average probability flows between the sets of Φ , i.e. the transition

probability between i, j is $\frac{1}{\pi(\mathcal{A}_i)} \int_{\mathcal{A}_i} \int_{\mathcal{A}_j} \mathcal{P}(x, dy) \pi(x) dx$. Then the decomposability theorem states the following lower bound on the spectral gap:

$$\text{GAP}(\mathcal{M}) \geq \frac{1}{2} \text{GAP}(\overline{\mathcal{M}}) \min_{1 \leq j \leq m} \text{GAP}(\mathcal{M}|_{\mathcal{A}_j})$$

The product is composed of two parts: the spectral gap of the projected Markov chain $\text{GAP}(\overline{\mathcal{M}})$, and the smallest spectral gap when the Markov process is restricted to each component of the partition. More recently, Lee et al. demonstrated a similar decomposition theorem, which concerned decomposing the stationary distribution instead of decomposing the state space (Lee et al., 2018). For a Langevin-diffusion process over a mixture of Gaussians of similar variance, the stationary distribution is similar to a mixture of Gaussians (raised by a power β), over each of which the Langevin process mixes rapidly (Bakry and Émery, 1985), so the partition components can simply be chosen to be the mixture components. In fact, this can be generalized to any mixture of log-concave distributions (Lee et al., 2018).

Let us now consider a tempering strategy as defined in the manuscript, which consists of stochastic switches between two inverse temperatures $\beta_{cold}, \beta_{hot}$ at rates $\eta_1 : \beta_{cold} \rightarrow \beta_{hot}, \eta_2 : \beta_{hot} \rightarrow \beta_{cold}$. If β_{hot} is chosen low enough then all the mixture components will become sufficiently close to allow for rapid transitions between the components at β_{hot} . It is therefore possible for the process to cross between the components at β_{cold} by transitioning to β_{hot} and then back to β_{cold} .

QUANTIFICATION AND STATISTICAL ANALYSIS

Estimating the inverse temperature β from chemotaxis in linear gradients

To quantify the inverse temperature parameter for *E. coli* chemotaxis, we used the experiments by (Kalinin et al., 2009) that measured population distributions $F(x)$ for *E. coli* in linear gradients after reaching dynamic equilibrium (Figure 5, panel E1, see Kalinin et al. for experimental details). In order to determine β , we estimated the following error: the absolute difference between the probability distribution proportional to $P(x) = L(x)^\beta$ (where $L(x)$ is ligand concentration) and $F(x)$ for each of the three gradients. We set β to minimize the average error for the three distributions. A similar fit can be obtained by taking the intercept of the linear fit of $\log F(x)$ against $\log L(x)$.