CAUSAL INFERENCE WITHIN CLOSE ACTION-PERCEPTION LOOPS
SUPPLEMENTARY INFORMATION

We here describe in detail the normative model that we used in the main text to predict behavioral signatures of causal inference.

1. Task setup and generative model

A visual target appears on the screen for $T = N\delta t$ seconds, or N time steps of size δt. While, for intuition, one can think of these time steps as individual video frames, our formulation is independent of the choice of δt. The target appears at location z_1 and thereafter moves with velocity v. Therefore, its location in the nth time-step is given by

$$z_n = z_1 + (n-1)v\delta t = z_N + (n-N)v\delta t. \tag{1}$$

Across trials, we assume the target to be moving ($\gamma = 1$) with probability $p(\gamma = 1) = p_\gamma$, and to be stationary ($\gamma = 0$) otherwise. When moving, its velocity is drawn from $p(v|\gamma = 1) = \mathcal{N}\left(v|0, \sigma_0^2\right)$, a normal distribution with zero mean and variance σ_0^2. Overall, this leads to the prior over v to be given by

$$p(v) = p(v|\gamma = 0)p(\gamma = 0) + p(v|\gamma = 1)p(\gamma = 1) = (1 - p_\gamma)\delta(v) + p_\gamma \mathcal{N}\left(v|0, \sigma_0^2\right), \tag{2}$$

where $\delta(\cdot)$ is the Dirac delta function. This is our causal inference prior as it encapsulates the two different hypotheses (stationary vs. moving target, $\gamma = 0$ vs. $\gamma = 1$) for what caused the sensory percepts, together with their associated priors on the underlying latent target velocity. We assume a uniform prior over the target's initial location z_1 over a bounded range, whose form we make more precise later.

In each time step, the observer makes a noisy observation x_n of the target's true location z_n, distributed independently across time as

$$x_n|z_n \sim \mathcal{N}\left(x_n|z_n, \sigma^2/\delta t\right). \tag{3}$$

We here scale the observation's variance by $1/\delta t$, such that the overall amount of information that the observer receives per unit time remains invariant to the choice of δt.

Having observed $x_{1:N} \equiv x_1, \ldots, x_N$, the observer wants to infer whether the target is moving or not, that is $p\left(\gamma = 1|x_{1:N}\right)$. Furthermore, they want to infer the target's final location for a stationary target, $p\left(z_N|\gamma = 0, x_{1:N}\right)$, or the target's final location and velocity for a moving target, $p\left(z_N, v|\gamma = 1, x_{1:N}\right)$. In the next section we derive these quantities. Following this, we turn to the question of how the observer uses these quantities to act upon them by reporting whether the target is stationary or moving, and how they steer toward the target.

2. Inferring the target location and velocity

Here, we first start with assuming that the target is stationary to find $p\left(z_N|\gamma = 0, x_{1:N}\right)$. Then, we assume a moving target and compute $p\left(z_N, v|\gamma = 1, x_{1:N}\right)$. Lastly, we use the found expressions to derive $p\left(\gamma = 1|x_{1:N}\right)$.

2.1. A stationary target, $\gamma = 0$.

For the stationary case, we only need to find the posterior over the target's single location z as its velocity is fixed to zero. Then, it is easy to show that

$$p\left(z|x_{1:N}\right) \propto p(x_{1:N}|z) \propto \mathcal{N}\left(z|\bar{x}, \frac{\sigma^2}{T}\right), \tag{4}$$

where we have implicitly conditioned on $\gamma = 0$, have assumed a uniform prior over z over the relevant range of z's, and have defined

$$\bar{x} = \frac{1}{N}\sum_{n=1}^{N} x_n, \tag{5}$$

that is, the average observed location.

Causal inference also requires the marginal likelihood of $\boldsymbol{x}_{1:N}$, which is given by

$$\begin{aligned} p\left(x_{1:N}\right) &= \int p\left(x_{1:N}|z\right) p\left(z\right) \mathrm{d}z \\ &= p_{z0} \int_{z_{min}}^{z_{max}} \left[\prod_{n=1}^{N} \mathcal{N}\left(x_n|z, \sigma^2/\delta t\right)\right] \mathrm{d}z \\ &= p_{z0} \left(\frac{\delta t}{2\pi\sigma^2}\right)^{N/2} \int_{z_{min}}^{z_{max}} e^{-\frac{\delta t \sum_{n=1}^{N}(x_n - z)^2}{2\sigma^2}} \mathrm{d}z, \end{aligned} \tag{6}$$

where we assumed a uniform prior of probability $p(z) = p_{z0}$ over a wide z-range from z_{min} to z_{max}. The integral evaluates to

$$\int e^{-\frac{\delta t \sum_{n=1}^{N}(x_n - z)^2}{2\sigma^2}} \mathrm{d}z = -\sqrt{\frac{\pi\sigma^2}{2N\delta t}} e^{-\frac{\delta t \sum_{n=1}^{N}(x_n - \bar{x})^2}{2\sigma^2}} \mathrm{erf}\left(\frac{N(\bar{x} - z)}{\sqrt{2N\sigma^2/\delta t}}\right), \tag{7}$$

where the sum in the exponential equals $N\mathrm{var}(x)$, that is N times the empirical variance of $x_{1:N}$. Furthermore, the error function approaches -1 for large z_{max} and 1 for small z_{min}, such that its contribution in the definite integral approaches -2. Therefore, the final marginal likelihood is

$$p(x_{1:N}) = p_{z0}\left(\frac{\delta t}{2\pi\sigma^2}\right)^{N/2}\sqrt{\frac{2\pi\sigma^2}{T}} e^{-\frac{T\mathrm{var}(x)}{2\sigma^2}}. \tag{8}$$

We can find the same result by writing down Bayes' rule for $p(z|x_{1:N})$ and solving for $p(x_{1:N})$, which appears in the denominator.

2.2. **A moving target, $\gamma = 1$.** Using the previous identity, $z_n = z_N - (N-n)v\delta t$, leads to the likelihood of each x_n to be given by

$$p(x_n|z_N, v) = \mathcal{N}\left(x_n|z_N - (N-n)v\delta t, \sigma^2/\delta t\right). \tag{9}$$

Our aim is to find the joint posterior over z_N and v, which is given by the expression

$$\begin{aligned} p(z_N, v|x_{1:N}) &\propto \mathcal{N}\left(v|0, \sigma_0^2\right)\prod_{1:n}^{N}\mathcal{N}\left(x_n|z_N - (N-n)v\delta t, \sigma^2/\delta t\right) \\ &\propto e^{-\frac{1}{2}\left(\left(\frac{1}{\sigma_0^2} + \frac{T\delta t^2 - 3T^2\delta t + 2T^3}{6\sigma^2}\right)v^2 + \frac{T}{\sigma^2}z_N^2 - \frac{T^2 - T\delta t}{\sigma^2}z_N v - \frac{2T\bar{x}}{\sigma^2}z_N + \frac{2}{\sigma^2}x_{acc}v\right)}, \end{aligned} \tag{10}$$

where we implicitly conditioned on $\gamma = 1$, have used the same definition of $\bar{x}$ as further above, and have defined

$$x_{acc} = \delta t^2 \sum_{n=1}^{N}(N-n)x_n \tag{11}$$

To find the posterior moments we first take $\delta t \to 0$, removing all the δt-dependent terms. To describe the full posterior, we denote it by

$$p(z_N, v|x_{1:N}) = \mathcal{N}\left(\begin{pmatrix} z_N \\ v \end{pmatrix} | \begin{pmatrix} \mu_z \\ \mu_v \end{pmatrix}, \begin{pmatrix} \Sigma_{zz} & \Sigma_{zv} \\ \Sigma_{zv} & \Sigma_{vv} \end{pmatrix}\right). \tag{12}$$

A lengthy, but unspectacular, derivation reveals

$$\Sigma_{zz} = \frac{4\sigma^2\left(3\sigma^2 + T^3\sigma_0^2\right)}{T\left(12\sigma^2 + T^3\sigma_0^2\right)}, \tag{13}$$

$$\Sigma_{vv} = \frac{12\sigma^2\sigma_0^2}{12\sigma^2 + T^3\sigma_0^2}, \tag{14}$$

$$\Sigma_{zv} = \frac{6T\sigma^2\sigma_0^2}{12\sigma^2 + T^3\sigma_0^2}, \tag{15}$$

$$\mu_z = \frac{12\sigma^2\bar{x} + 4T^3\sigma_0^2\bar{x} - 6Tx_{acc}\sigma_0^2}{12\sigma^2 + T^3\sigma_0^2}, \tag{16}$$

$$\mu_v = \frac{6\sigma_0^2\left(T^2\bar{x} - 2x_{acc}\right)}{12\sigma^2 + T^3\sigma_0^2}. \tag{17}$$

Interestingly, in the $\sigma_0 \to \infty$ limit, the posterior variance Σ_{zz} scales as $1/T$, as before. The posterior variance Σ_{vv}, in contrast, drops more rapidly with $1/T^3$. Therefore, temporal integration of evidence provides qualitatively more velocity than location information. Intuitively, this is because any (x_i, x_j) pair can be used to infer velocity, whereas the location estimate relies on the across-x_i average.

Understanding what the posterior means needs more work. In particular, let us define

$$\hat{v} = \frac{6}{T^3}\left(T^2\bar{x} - 2x_{acc}\right) = \frac{6}{T}\left(\frac{1}{N}\sum_{n=1}^{N}\left(\frac{2n}{N} - 1\right)x_n\right). \tag{18}$$

The term in (outer) parenthesis is a weighted sum of the x_n's. In this sum, x_1 is weighted by -1, and x_N by 1. Inbetween these extremes, the weights increase linearly from -1 to 1. Therefore, if we group equally-weighted terms (modulus the sign of the weight), the sum is a weighted combination of location differences, with the highest weight on $x_N - x_1$, less weight on $x_{N-1} - x_2$, and so on. To see how this can act as a velocity estimate, assume a noise-free $x_n = vn\delta t$. Substituting this in the above expression shows that $\hat{v} = v$ (this is where the $6/T$ pre-factor comes from). This justifies denoting it $\hat{v}$.

Substituting $\hat{v}$ into the posterior means results in the new expressions

$$\mu_z = \frac{12\sigma^2\bar{x} + T^3\sigma_0^2\left(\bar{x} + \frac{\hat{v}T}{2}\right)}{12\sigma^2 + T^3\sigma_0^2} = \bar{x} + \frac{T^3\sigma_0^2}{12\sigma^2 + T^3\sigma_0^2}\frac{T}{2}\hat{v}, \tag{19}$$

$$\mu_v = \frac{T^3\sigma_0^2}{12\sigma^2 + T^3\sigma_0^2}\hat{v} \tag{20}$$

This shows that the posterior mean μ_z start with $\bar{x}$ for small T, and then shifts toward $\bar{x} + \mu_v T/2$, which is the mean location plus half the estimated distance that the target moved, which is sensible. The posterior mean μ_v is initially biased toward zero, due to the prior, and later approaches $\hat{v}$. Overall, with $\sigma_0^2 \to 0$, the posterior approaches that for a stationary target, as desired.

To find the marginal likelihood, we solve Bayes's rule for the posterior for $p(x_{1:N})$, which yields

$$\begin{aligned} p\left(x_{1:N}\right) &= \frac{p\left(x_{1:N}|z_N, v\right)p\left(z_N\right)p\left(v\right)}{p\left(z_N, v|x_{1:n}\right)} \\ &= p_{z0}\left(\frac{\delta t}{2\pi\sigma^2}\right)^{N/2}\sqrt{\frac{2\pi|\mathbf{\Sigma}|}{\sigma_0^2}}e^{-\frac{T\left(\mathrm{var}(x)+\bar{x}^2\right)}{2\sigma^2}+\frac{1}{2}\boldsymbol{\mu}^T\mathbf{\Sigma}^{-1}\boldsymbol{\mu}}, \end{aligned} \tag{21}$$

where we have chosen $z_N = v = 0$ for the second equality (as the expression holds for any choice of z_N and v), and $\boldsymbol{\mu}$ and $\mathbf{\Sigma}$ denote the posterior mean and covariance. The remaining terms evaluate to

$$\boldsymbol{\mu}^T\mathbf{\Sigma}^{-1}\boldsymbol{\mu} = \frac{T\bar{x}^2}{\sigma^2} + \frac{T^6\sigma_0^2\hat{v}^2}{12\sigma^2\left(12\sigma^2 + T^3\sigma_0^2\right)}, \tag{22}$$

$$|\mathbf{\Sigma}| = \frac{12\sigma^4\sigma_0^2}{T\left(12\sigma^2 + T^3\sigma_0^2\right)}, \tag{23}$$

such that the marginal likelihood becomes

$$p\left(x_{1:N}\right) = p_{z0}\left(\frac{\delta t}{2\pi\sigma^2}\right)^{N/2}\sqrt{\frac{24\pi\sigma^4}{T\left(12\sigma^2 + T^3\sigma_0^2\right)}}e^{\frac{T^6\hat{v}^2\sigma_0^2}{24\sigma^2\left(12\sigma^2+T^3\sigma_0^2\right)}-\frac{T\mathrm{var}(x)}{2\sigma^2}} \tag{24}$$

As for the posterior, this marginal likelihood approaches that for a stationary target with $\sigma_0^2 \to 0$.

2.3. **Is the target stationary or moving?** To find the full posterior over the target's state, we use the causal inference target velocity prior, $p(v) = (1 - p_\gamma)\delta(v - 0) + p_\gamma\mathcal{N}\left(v|0, \sigma_0^2\right)$, with which the posterior becomes

$$p\left(z_N, v|x_{1:N}\right) = p\left(z_N|x_{1:N}, \gamma = 0\right)\delta\left(v - 0\right)p\left(\gamma = 0|x_{1:N}\right) + p\left(z_N, v|x_{1:N}, \gamma = 1\right)p\left(\gamma = 1|x_{1:N}\right). \tag{25}$$

In this mixture distribution, the first mixture component is that for the stationary target, and the second that for the moving one. These two components are weighted by the causal modeling posterior $p(\gamma|x_{1:N})$ which indicates the probability of the target being stationary or moving given the data. This probability can again be found by Bayes' rule, and results in

$$\begin{aligned} p\left(\gamma = 1|x_{1:N}\right) &= \frac{p\left(x_{1:N}|\gamma = 1\right)p(\gamma = 1)}{p\left(x_{1:N}|\gamma = 1\right)p(\gamma = 1) + p\left(x_{1:N}|\gamma = 0\right)p(\gamma = 0)} \\ &= \frac{\sqrt{\frac{12\sigma^2}{12\sigma^2+T^3\sigma_0^2}}e^{\frac{T^6\sigma_0^2\hat{v}^2}{24\sigma^2\left(12\sigma^2+T^3\sigma_0^2\right)}}p_\gamma}{\sqrt{\frac{12\sigma^2}{12\sigma^2+T^3\sigma_0^2}}e^{\frac{T^6\sigma_0^2\hat{v}^2}{24\sigma^2\left(12\sigma^2+T^3\sigma_0^2\right)}}p_\gamma + 1 - p_\gamma}, \end{aligned} \tag{26}$$

where the last expression results from substituting the marginal likelihoods for the stationary and moving target, and cancelling all the shared terms. This expression shows that, for a uniform prior, $p(\gamma = 0) = p(\gamma = 1) = 1/2$, the target is deemed more likely moving for larger velocity estimates $\hat{v}$. If this estimate is zero, that is, $\hat{v} = 0$, then, the more time has passed, the less likely is the target considered to be moving.

3. Acting upon the inferred target location and velocity

3.1. **Choosing stationary vs. moving.** Decision-makers would choose between a stationary and a moving target according to $p\left(\gamma|x_{1:N}\right)$. In particular, they would decide that the target is moving if $p(\gamma = 1|x_{1:N}) > 1/2$, that is, if

$$\frac{1}{2}\log\frac{12\sigma^2}{12\sigma^2 + T^3\sigma_0^2} + \frac{T^6\sigma_0^2\hat{v}^2}{24\sigma^2\left(12\sigma^2 + T^3\sigma_0^2\right)} > \log\frac{p(\gamma = 0)}{p(\gamma = 1)}. \tag{27}$$

3.1.1. *Empirical distribution of stationary reports.* As experimenters we cannot directly observe $\hat{v}$, such that we need to estimate it. Furthermore, it will fluctuate across trials, even if the same evidence is presented, making the decisions more noisy.

To build a model for $\hat{v}$ we note that $\hat{v}$ is a weighted sum of x_n's, and rely on our generative assumptions of x_n for a stationary and a moving target. For a stationary target, we have $x_n \sim \mathcal{N}\left(z, \sigma^2/\delta t\right)$. In this case, $\langle \hat{v} | \gamma = 0 \rangle = 0$, and its variance is given by

$$\text{var}\left(\hat{v}|\gamma = 0\right) = \left(\frac{6}{TN}\right)^2 \sum_{n=1}^{N} \left(\frac{2n}{N} - 1\right)^2 \frac{\sigma^2}{\delta t} = \frac{12\sigma^2}{T^3}. \tag{28}$$

This variance decreases rapidly with time, as more and more x_n's are used to estimate $\hat{v}$. This results in the required moments

$$\left\langle \hat{v}^2 | \gamma = 0 \right\rangle = \frac{12\sigma^2}{T^3}, \qquad \text{var}\left(\hat{v}^2 | \gamma = 0\right) = \left\langle \hat{v}^4 | \gamma = 0 \right\rangle - \left\langle \hat{v}^2 | \gamma = 0 \right\rangle^2 = \frac{288\sigma^4}{T^6}. \tag{29}$$

If we denote the $\hat{v}^2$-related term in the above decision criterion, Eq. (27), by α, this α has moments

$$\langle \alpha | \gamma = 0 \rangle = \frac{T^3\sigma_0^2}{2\left(12\sigma^2 + T^3\sigma_0^2\right)}, \qquad \sqrt{\text{var}\left(\alpha|\gamma = 0\right)} = \frac{T^3\sigma_0^2}{\sqrt{2}\left(12\sigma^2 + T^3\sigma_0^2\right)}, \tag{30}$$

leading to the signal-to-noise ratio $\langle \alpha | \gamma = 0 \rangle / \sqrt{\text{var}(\alpha|\gamma = 0)} = 1/\sqrt{2}$. More relevant, under the assumption that α is Gaussian whose parameters are fully determined by mean and variance, the probability that the decision criterion is met becomes

$$p\left(\text{choose } \gamma = 1 | \gamma = 0\right) = \Phi\left(\frac{1}{\sqrt{2}} + \frac{\sqrt{2}\left(12\sigma^2 + T^3\sigma_0^2\right)}{T^3\sigma_0^2}\left(\log\frac{p(\gamma = 1)}{p(\gamma = 0)} + \frac{1}{2}\log\frac{12\sigma^2}{12\sigma^2 + T^3\sigma_0^2}\right)\right) \tag{31}$$

For $T \to 0$ or $\sigma_0^2 \to 0$ this probability is dominated by the prior and lead to a choice of $\gamma = 1$ if $p(\gamma = 1) > p(\gamma = 0)$. It shrinks with increasing T, as more evidence results in higher certainty that the target is not moving. Increasing the observation noise σ^2 has two counteracting effects. First, it increases the pre-factor to the inner-most brackets, thus boosting the prior. Second it results in a weaker drop of the last term in brackets with time, indicating that more evidence will be required to discard the possibility that the target is moving.

For a moving target, $x_n \sim \mathcal{N}\left(z_N - (N - n)v\delta t, \sigma^2/\delta t\right)$. This yields $\hat{v}$ to be Gaussian, with moments

$$p\left(\hat{v}|v\right) = \mathcal{N}\left(\hat{v}|v, \frac{12\sigma^2}{T^3}\right), \tag{32}$$

which only differs in the non-zero mean from the $\hat{v}$ for the stationary case. With the above, the moments of $\hat{v}$ are given by

$$\left\langle \hat{v}^2 | \gamma = 1 \right\rangle = \frac{12\sigma^2}{T^3}\left(1 + \tilde{v}^2\right), \qquad \text{var}\left(\hat{v}^2 | \gamma = 1\right) = \frac{288\sigma^4}{T^6}\left(1 + 2\tilde{v}^2\right), \tag{33}$$

where we have defined the time-rescaled velocity $\tilde{v}^2 = T^3v^2/(12\sigma^2)$. The previously defined α then has moments

$$\langle \alpha | \gamma = 1 \rangle = \frac{T^3\sigma_0^2\left(1 + \tilde{v}^2\right)}{2\left(12\sigma^2 + T^3\sigma_0^2\right)}, \qquad \sqrt{\text{var}\left(\alpha|\gamma = 1\right)} = \frac{T^3\sigma_0^2\sqrt{1 + 2\tilde{v}^2}}{\sqrt{2}\left(12\sigma^2 + T^3\sigma_0^2\right)}, \tag{34}$$

leading to the signal-to-noise ratio $\left(1 + \tilde{v}^2\right)/\sqrt{2 + 4\tilde{v}^2}$ which approaches the linear function $|\tilde{v}|/2$ for larger $\tilde{v}^2$. Plugging these moments into the decision criteria and assuming Gaussianity leads to the choice probability

$$p\left(\text{choose } \gamma = 1 | \gamma = 1\right) = \Phi\left(\frac{1 + \tilde{v}^2}{\sqrt{2 + 4\tilde{v}^2}} + \frac{\sqrt{2}\left(12\sigma^2 + T^3\sigma_0^2\right)}{T^3\sigma_0^2\sqrt{1 + 2\tilde{v}^2}}\left(\log\frac{p(\gamma = 1)}{p(\gamma = 0)} + \frac{1}{2}\log\frac{12\sigma^2}{12\sigma^2 + T^3\sigma_0^2}\right)\right). \tag{35}$$

For $\tilde{v}^2 \to 0$, this probability becomes equivalent to the one for $\gamma = 0$. The larger $\tilde{v}^2$, the stronger the influence of the first term, and the weaker the influence of the remaining terms. In particular, the larger $\tilde{v}^2$, the higher the probability of choosing $\gamma = 1$.

A non-approximate approach to computing $p\left(\text{choose } \gamma = 1 | \gamma = 1\right)$ is to re-write the decision criterion for $p\left(\gamma = 1 | x_{1:N}\right) > 1/2$ as

$$|\hat{v}| > \beta \equiv \frac{\sigma}{T^3\sigma_0}\sqrt{24\left(12\sigma^2 + T^3\sigma_0^2\right)\left(\log\frac{p(\gamma = 0)}{p(\gamma = 1)} - \frac{1}{2}\log\frac{12\sigma^2}{12\sigma^2 + T^3\sigma_0^2}\right)}. \tag{36}$$

This criterion is only valid if the term in square-roots is non-negative, which is guaranteed as long as

$$\frac{p(\gamma = 0)}{p(\gamma = 1)} > \sqrt{\frac{12\sigma^2}{12\sigma^2 + T^3\sigma_0^2}}. \tag{37}$$

This is always satisfied if $p(\gamma = 0) > p(\gamma = 1)$. In general, it requires little evidence about the target's motion (i.e., small T and σ_0^2 and large σ^2), and a relatively strong prior toward the target not moving. If the above condition is violated, the decision criterion becomes $|\hat{v}| \geq 0$, which is always satisfied. That is, under these circumstances, the target will always be considered moving.

Assuming there is a non-zero chance of the target being stationary, then the probability of the decision-maker reporting a stationary target depends on the perceived $\hat{v}$ which is Gaussian in the true v (see above). Then, as $\beta \geq 0$, $|\hat{v}| > \beta$ is satisfied if either $\hat{v} > \beta$ or $-\hat{v} > \beta$. As these two options are mutually exclusive, their joint probability sums and is given by

$$p(\text{choose } \gamma = 1|v) = \Phi\left(\frac{v-\beta}{\sqrt{12\sigma^2/T^3}}\right) + \Phi\left(\frac{-v-\beta}{\sqrt{12\sigma^2/T^3}}\right). \tag{38}$$

Simulations confirmed that these expressions match simulated choices.

4. Stage II: intercepting the target

We assume agents travel at constant velocity, such that we only need to determine travel direction and stopping time for the best target interception. The objective is to minimize the expected cost, $\langle c(\boldsymbol{z}_a(t), \boldsymbol{z}_o(t))\rangle$, where $\boldsymbol{z}_a$ is the agent's location, $\boldsymbol{z}_o$ is the target's location, and the expectation is over the uncertainty involving both. We will assume a simple cost function $c(\boldsymbol{z}_a, \boldsymbol{z}_o) = -\delta(\boldsymbol{z}_a - \boldsymbol{z}_o)$, which is minimized if $\boldsymbol{z}_a = \boldsymbol{z}_o$.

The task itself is two-dimensional: the target appears at a certain distance and can move only laterally. We will assume that the depth is known and denoted $z_{o,d}$ (o for object, and d for depth). The agent moves at a constant velocity v_a at angle θ ($\theta = 0$ is straight-ahead) for some time t. Then, what needs to be determined for optimal interception is the agent's stopping time t^* and the angle θ^* that minimizes the expected cost,

$$t^*, \theta^* = \underset{t,\theta}{\text{argmin}} \langle c(\boldsymbol{z}_a(t,\theta), \boldsymbol{z}_o(t))\rangle = \underset{t,\theta}{\text{argmax}}\, p(\boldsymbol{z}_a(t,\theta) = \boldsymbol{z}_o(t)), \tag{39}$$

where the second equality follows from the delta cost-function.

4.1. **Agent motion model.** We will use the self-motion estimation model from Lakshminarasimhan et al. (2018), where they assume a Weber-like variance scaling of the self-location estimate,

$$\boldsymbol{z}_a(t,\theta) = \begin{pmatrix} z_{a,l} \\ z_{a,d} \end{pmatrix}(t,\theta) \sim \mathcal{N}\left(\begin{pmatrix} v_a t \sin\theta \\ v_a t \cos\theta \end{pmatrix}, \begin{pmatrix} k^2 (v_a t \sin\theta)^{2\lambda} & 0 \\ 0 & k^2 (v_a t \cos\theta)^{2\lambda} \end{pmatrix}\right), \tag{40}$$

with parameters k (k^2 is the variance scaling factor) and λ (determines the sub/supra-linearity of the variance scaling), and where we have assumed $\boldsymbol{z}_a(0,\theta) = (0,0)^T$.

4.2. **Target motion model.** As the target depth is known, we only need to estimate its initial lateral location and eventual velocity. To do so, we use the target location/velocity estimates derived further above, which provide the joint estimate

$$p(z_{o,l}(0), v_{o,l}|X) = p(z_{o,l}(0)|X, \gamma = 0)\, p(\gamma = 0|X) + p(z_{o,l}(0), v_{o,l}|X, \gamma = 1)\, p(\gamma = 1|X), \tag{41}$$

where $X \equiv x_{1:N}$ are all target observations up to the target offset, $z_{o,l}(0)$ is the inferred lateral position at target offset (measured relative to agent), $v_{o,l}$ is the target's lateral velocity, and $\gamma \in \{0,1\}$ denotes the target being stationary/moving. To compute the posterior over $z_{o,l}(t)$ for $\gamma = 1$, we use $z_{o,l}(t) = z_{o,l}(0) + v_{o,l}$ to get, as before,

$$\mu_{oz,l}(t) = \langle z_{o,l}(t)|X, \gamma = 1\rangle = \mu_z + \mu_v t = \frac{12\sigma^2 \bar{x} + T^3\sigma_0^2\left(\bar{x} + \left(\frac{T}{2} + t\right)\hat{v}\right)}{12\sigma^2 + T^3\sigma_0^2}, \tag{42}$$

$$\sigma^2_{oz,l}(t) = \text{var}(z_{o,l}(t)|X, \gamma = 1) = \Sigma_{zz} + t^2\Sigma_{vv} + 2t\Sigma_{zv} = \frac{4\sigma^2\left(3\sigma^2 + \sigma_0^2\left(T^3 + 3T^2 t + 3Tt^2\right)\right)}{T\left(12\sigma^2 + T^3\sigma_0^2\right)}, \tag{43}$$

where μ_z, μ_v, and the Σ's are the posterior moments of $p(z_{o,l}(0), v_{o,l}|X, \gamma = 1)$, Eq. (12). Overall, this leads to the posterior to be given by

$$p(\boldsymbol{z}_o(t)|X) = \delta(z_{o,d}(t) - z_{o,d})\left(\mathcal{N}\left(z_{o,l}(t)|\bar{x}, \frac{\sigma^2}{T}\right) p(\gamma = 0|X) + \mathcal{N}\left(z_{o,l}(t)|\mu_{oz,l}(t), \sigma^2_{oz,l}(t)\right) p(\gamma = 1|X)\right). \tag{44}$$

4.3. **Optimal steering angle and stopping time.** To find the optimal steering angle θ^* and stopping time t^*, we use the independence of the components of $\boldsymbol{z}_a(t)$ and $\boldsymbol{z}_o(t)$ to find

$$z_{a,l}(t) - z_{o,l}(t)|X \sim \mathcal{N}\left(v_a t \sin\theta - \bar{x}, k^2 (v_a t \sin\theta)^{2\lambda} + \frac{\sigma^2}{T}\right) p(\gamma = 0|X) \tag{45}$$
$$+ \mathcal{N}\left(v_a t \sin\theta - \mu_{oz,l}(t), k^2 (v_a t \sin\theta)^{2\lambda} + \sigma^2_{oz,l}(t)\right) p(\gamma = 1|X),$$

$$z_{a,d}(t) - z_{o,l}(t)|X \sim \mathcal{N}\left(v_a t \cos\theta - z_{o,d}, k^2 (v_a t \cos\theta)^{2\lambda}\right) \tag{46}$$

The optimal θ^* and t^* maximize the probability of both differences being zero. This leads to the expression

$$t^*, \theta^* = \underset{t,\theta}{\operatorname{argmax}} \mathcal{N}\left(0|v_a t\cos\theta - z_{o,d}, k^2\left(v_a t\cos\theta\right)^{2\lambda}\right) \times \left(\mathcal{N}\left(0|v_a t\sin\theta - \bar{x}, k^2\left(v_a t\sin\theta\right)^{2\lambda} + \frac{\sigma^2}{T}\right) p\left(\gamma = 0|X\right) + \mathcal{N}\left(0|v_a t\sin\theta - \mu_{oz,l}(t), k^2\left(v_a t\sin\theta\right)^{2\lambda} + \sigma^2_{oz,l}(t)\right) p\left(\gamma = 1|X\right)\right). \quad (47)$$

This optimization is complex and doesn't have an analytical solution. Therefore, we need to use numerical optimization to find θ^* and t^*.

For good initial guesses for this optimization, we consider the moving and stationary target case in isolation. For a stationary target, the Gaussian peaks at $v_a t\sin\theta = \bar{x}$ and $v_a t\cos\theta = z_{o,d}$, which leads to parameters

$$\theta^*_{\gamma=0} = \tan^{-1}\frac{\bar{x}}{z_{o,d}}, \qquad t^*_{\gamma=0} = \frac{\sqrt{\bar{x}^2 + z^2_{o,d}}}{v_a}. \quad (48)$$

For a moving target, the Gaussian peaks at $v_a t\sin\theta = \mu_{oz,l}(t) = \mu_z + \mu_v t$ and $v_a t\cos\theta = z_{o,d}$, resulting in

$$\theta^*_{\gamma=1} = \tan^{-1}\frac{\mu_z + \mu_v t^*_{\gamma=1}}{z_{o,d}}, \qquad t^*_{\gamma=1} = \frac{\mu_v\mu_z \pm \sqrt{v_a^2 z^2_{o,d} + v_a^2\mu_z^2 - \mu_v^2 z^2_{o,d}}}{v_a^2 - \mu_v^2}, \quad (49)$$

where we provide two solutions to $t^*_{\gamma=1}$ which result from solving a quadratic equation, and where μ_z and μ_v are the posterior moments of $p\left(z_{o,l}(0), v_{o,l}|X, \gamma = 1\right)$. To identify a unique $t^*_{\gamma=1}$ we assume that $v_a^2 > \mu_z^2$, such that the agent is guaranteed to be able to catch up with the target. Furthermore, we require $t^*_{\gamma=1} > 0$ which is guaranteed to be violated if $\mu_z^2\mu_v^2 < v_a^2 z^2_{o,d} + v_a^2\mu_z^2 - \mu_v^2 z^2_{o,d}$, or, equally, $\left(z^2_{o,d}/\mu_z^2 + 1\right)\left(v_a^2 - \mu_z^2\right) > 0$. The latter holds if both $v_a^2 - \mu_z^2 > 0$ (guaranteed by assumption) and $z^2_{o,d}/\mu_z^2 + 1 > 0$, or $z^2_{o,d} > -\mu_z^2$. As the right-hand side of the latter is always negative, the last inequality is always true, confirming that the only solution to $t^*_{\gamma=1}$ is the one with a sum (rather than difference) in the numerator. Substituting the expressions of μ_z and μ_v in terms of $\bar{x}$ and $\hat{v}$ does not lead to any appreciable simplifications.

Together, these two solutions allow us to approximate the optimal heading directions and stopping times by

$$\theta^* \approx \theta^*_{\gamma=0} p\left(\gamma = 0|X\right) + \theta^*_{\gamma=1} p\left(\gamma = 1|X\right), \quad (50)$$

$$t^* \approx t^*_{\gamma=0} p\left(\gamma = 0|X\right) + t^*_{\gamma=1} p\left(\gamma = 1|X\right). \quad (51)$$

We use these approximations to initiate a gradient-descent procedure to find the correct t^* and θ^*.

4.4. **Experimentally observable optimal stopping, and perceived velocity estimate.** As for the stationarity reports, the experimenter does not observe $\bar{x}$ and $\hat{v}$ and needs to marginalize over them. They are distributed as

$$\bar{x}|v, z_N\, \mathcal{N}\left(z_N - \frac{vT}{2}, \frac{\sigma^2}{T}\right), \qquad \hat{v}|v \sim \mathcal{N}\left(v, \frac{12\sigma^2}{T^3}\right). \quad (52)$$

We estimate the associated statistics of θ^* and t^*, by a Monte Carlo approximation. That is, we draw multiple $\bar{x}$ and $\hat{v}$, compute the associated θ^* and t^*, and use those to compute the posterior distributions over θ^* and t^*.

To assess the decision maker's estimate of the target's velocity, we will infer the assumed velocity from the decision maker's stopping location. The actual stopping location could feature a mismatch between the decision maker's depth and that of the target. By our initial assumption that the decision maker knows the target's depth, this mismatch arises from the decision maker's actual location while actively steering, and thus doesn't refect the decision maker's estimate. Thus, instead of using the agent's final location as a measure for the assumed velocity, we will instead use the point $z_{o,d}\tan\theta^*$ and time $\hat{t} = z_{o,d}/(v_a\cos\theta^*)$ at which the agent crossed the target's path (i.e. reaches depth $z_{0,d}$).

Therefore, we won't take the depth mismatch into account, and instead only focus on lateral location, using

$$z_{o,l}(0) + \hat{v}_o\hat{t} \approx v_a\hat{t}\sin\theta^*, \quad (53)$$

where the left-hand side is the target's location at time $\hat{t}$ (assuming velocity $\hat{v}_o$), and the right-hand side is the decision maker's location at the same time. Re-expressing the above in terms of $\hat{v}_o$ results in

$$\hat{v}_o \approx v_a\left(\sin\theta^* - \frac{z_{o,l}(0)}{z_{o,d}}\cos\theta^*\right). \quad (54)$$