Active Sensing: An Optimal Inference Model

Sheeraz Ahmad[1] and Angela J. Yu[2]

University of California, San Diego

1. Computer Science and Engineering
2. Cognitive Science

Introduction

• Biological agents employ active sensing: goal-directed, context-dependent, self-motion to selectively attend to most informative aspects of the environment

• Need for goal oriented “active” models that allow for control of visual inputs.

• Existing models predict fixations based on abstract statistical measures e.g. saliency (Itt, Koch & Niebur, 1998), posterior max (Najemnik & Geisler, 2005), entropy (Butko & Movellan, 2010).

• Our approach: directly optimizing behavioral objectives such as accuracy, temporal delay and effort.

Visual Search Experiment

• Task: Find target (⇐) amongst distractors (⇒).

• Some locations privileged.

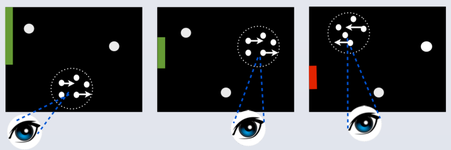

• Sequence of stimulus controlled by subject.

• Gaze contingent display prevents subjects from using covert attention.

• Reward policy:

$$\text{Score (\$)} = 100 - 12.5 \times T - 25 \times S + 50 \times C$$

base points, search cost, switch cost, accuracy

Model: Bayesian Inference & Control

• Maximize behaviorally relevant goals: speed, effort and accuracy.

• Define belief state: $\mathbf{p}_t := (p(s=1|\mathbf{x}_t;\boldsymbol{\lambda}_t),\ldots,p(s=k|\mathbf{x}_t;\boldsymbol{\lambda}_t))$

where s : target location

$\boldsymbol{\lambda}_t$: set of fixation locations till time 't'

$\mathbf{x}_t$: set of observations till time 't’

• Uses Bayesian inference for belief update:

$$\mathbf{p}_t^i \propto p(x_t|s=i;\lambda_t)p(s=i|\mathbf{x}_{t-1};\boldsymbol{\lambda}_{t-1}) = f_{s,\lambda_t}(x_t)\mathbf{p}_{t-1}^i$$

• Define the cost for a trial run:

$$l(\tau,\hat{s};\boldsymbol{\lambda}_\tau,s) = c\tau + c_s n_\tau + 1_{\{\hat{s}\neq s\}}$$

• Uses stochastic control theory to find the optimal policy that minimizes the expected cost:

$$L_\pi := \mathbb{E}[l(\tau,\boldsymbol{\lambda}_\tau,\hat{s})]_{\mathbf{x},s} = c\mathbb{E}[\tau] + c_s\mathbb{E}[n_s] + P(\hat{s}\neq s)$$

• Bellman optimality principle guarantees:

$$V(p_t,\lambda_t=j) = \min\begin{cases} ct + c_s n_s(t) + (1-p_t^j) & \lambda_{t+1}=0 \\ ct + c_s n_s(t) + c + \mathbb{E}[V(p_{t+1},\lambda_{t+1}=i)] & \lambda_{t+1}=i=j \\ ct + c_s n_s(t) + c + c_s + \mathbb{E}[V(p_{t+1},\lambda_{t+1}=i)] & \lambda_{t+1}=i\neq j \end{cases}$$

Results (Confirmation Bias)

Huang & Yu, 2010: ↑ **prior expectation ⇒ ↑ P(target selection).**

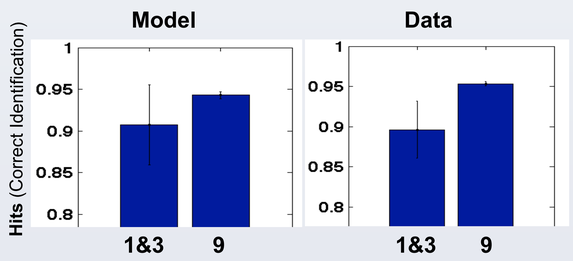

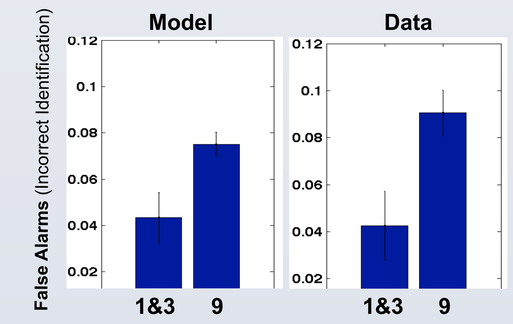

↑ prior expectation ⇒↓ time to confirm, ↑ time to disconfirm

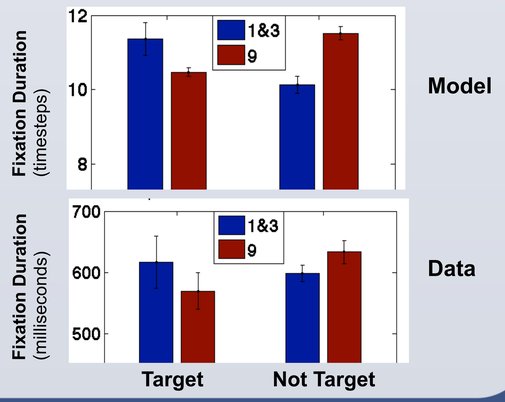

Optimal Model Details

• Observations are modeled as i.i.d. Bernoulli:

$$\begin{aligned} p(x_t=x|s=i;\lambda_t=j) &= 1_{\{i=j\}}Bern(x;\beta_1) + 1_{\{i\neq j\}}Bern(x;\beta_0) \\ &= 1_{\{i=j\}}\beta_1^x(1-\beta_1)^{1-x} + 1_{\{i\neq j\}}\beta_0^x(1-\beta_0)^{1-x} \end{aligned}$$

$$\beta_0 = 1-\beta_1$$

• One stopping action: current fixation location.

• Discretized belief state.

• Resulting value function for $\beta = 0.9,\ c = 0.1,\ c_s = 0$

• And the corresponding optimal policy:

Optimal Model Predictions

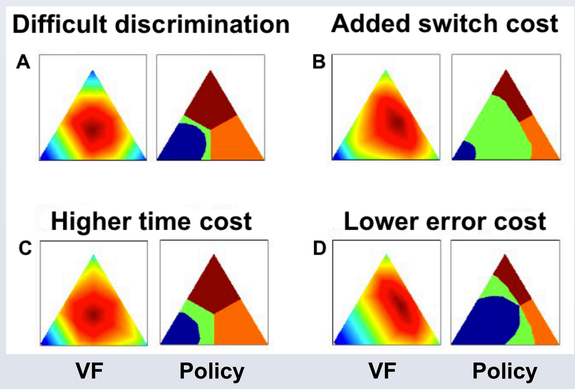

Relationship to Existing Probabilistic Models

• Greedy MAP (Najemnik & Geisler, 2005): Maximize the expected one step look ahead probability of the target.

$$\lambda'_{t+1}(\mathbf{p}_t) = \underset{i}{\text{argmax}}\ \mathbb{E}[\max_i \mathbf{p}_{t+1}^i]$$

• Infomax (Butko & Movellan, 2010): Minimize the expected long term entropy (uncertainty).

$$\lambda''_{t+1}(\mathbf{p}_t) = \underset{i}{\text{argmin}} \sum_{t'=t+1}^{T} \mathbb{E}[H(\mathbf{p}_{t'})|\lambda_{t+1}=i]$$

• Goals are not defined in terms of behavioral objectives but rather abstract quantities.

• Ad hoc stopping policy: stop when a certain level of certainty is reached.

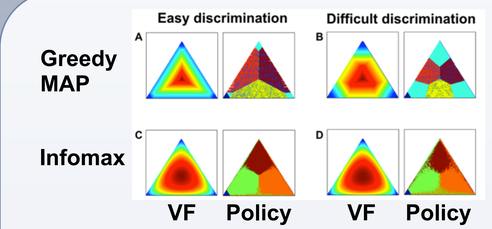

Radial Basis Approximation

• Solving for value function is computationally complex.

• Learn the function in a low dimensional basis space.

• At each step, choose a fixed number of random points, and update the value function via value iteration at those points using the current basis representation.

• Update the basis representation using these new values, and iterate till convergence.

Results (Approximate Policy)

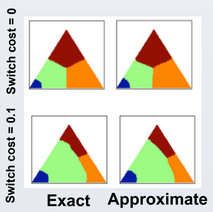

• We note that approximate policy is context sensitive.

• Complexity linear in number of hypotheses, as opposed to exponential for exact solution.

• RBF networks previously proposed to underlie visual processing (Poggio & Girosi, 1990).

Conclusion

• Policy optimizing behavioral objectives is intuitive and provides an account of fixation location as well as duration in a visual search task.

• The model also replicates confirmation bias.

• Heuristic policies fail to adapt to changing environment.

• Whereas, the radial-basis approximate policy is both computationally simple and context-sensitive.

Future work: ↑# locations, other visual tasks, Bayesian nonparametric approximations, competitive foraging.

References

1. L Itti, C Koch and E Niebur. A model of saliency based visual attention for rapid scene analysis. IEEE Transactions on Pattern Analysis and Machine 1998.
2. J Najemnik and W S Geisler. Optimal eye movement strategies in visual search. Nature, 434(7031):387-91, 2005.
3. N J Butko and J R Movellan. Infomax control of eye movements. IEEE Transactions on Autonomous Mental Development, 2(2):91-107, 2010.
4. H Huang and A J Yu. Society for Neural Science Annual Meeting, page online, 2010.
5. T Poggio and F Girosi. Networks for approximation and learning. Proceedings of the IEEE, 1990