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

• Humans/animals use self-motion to **actively** obtain sensory data, in a **goal-directed**, **context-sensitive** manner

• Existing models: predict fixations based on **abstract statistical measures** e.g. saliency (Itti, 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**

• Validate model predictions with **human behavioral data**

• Devise simple but effective **approximate algorithms**

Bayesian Inference & Stochastic Control

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

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

$$\beta_0 = 1 - \beta_1$$

• 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

• Bayesian inference for belief update:

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

• One stopping action: current fixation location

• Optimize behaviorally defined goals: speed, accuracy, efficiency

• Cost for a one "trial": delay + # switches + error

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

• Optimal policy minimizes the expected cost (Bayes' risk):

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

• Bellman dynamic programming equation:

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

• Discretize belief state space, grid size used = 200

• Initialize V as stopping cost for each value of p_t, then iterate until convergence to find Q-factors (and value function) of the optimal policy

Model Behavior and Predictions

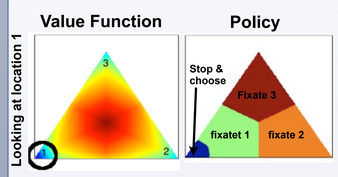

• rotationally symmetric policies for other locations

• environment (c,c_s,β): (0.1,0,0.9)

• A: (0.1,0,0.7)
B: (0.1,0.1,0.9)
C: (0.2,0,0.9),
D: (0.2,0.2,0.9)

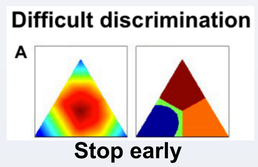

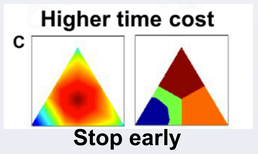

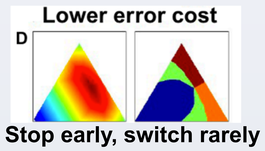

Human Behavioral Experiment: Visual Search

(Huang & Yu, *SfN*, 2010)

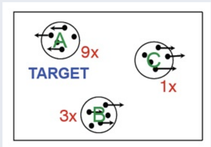

• Task: Find target (⟸) amongst distractors (⟹)

• Some locations more likely to contain target (1:3:9)

• Sequence of observations controlled by subject

• Gaze contingent display allows exact measurement of where subject obtains sensory input

• Reward policy penalizes for total search time, # switches, and incorrect response

Model-Data Comparison

↑ prior expectation ⇒ ↑ P(target selection)

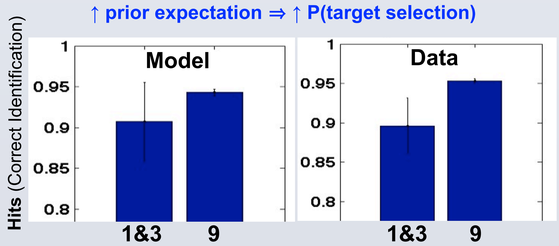

*False Alarms exhibit similar bias (not shown)

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

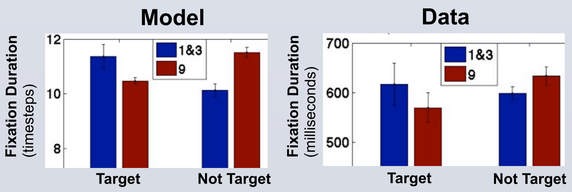

Relationship to Existing Probabilistic Models

• Greedy MAP (Najemnik & Geisler, 2005): maximize one-step look ahead probability of finding target

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

• Infomax (Butko & Movellan, 2010): minimize expected cumulative future entropy (uncertainty)

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

• These abstractive, statistical objectives are intuitively appealing, but hard to relate to behavioral goals

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

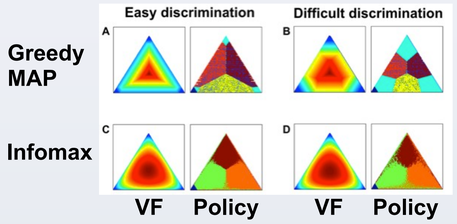

• Neither adapts to different contexts (behavioral goals)

• Greedy MAP: incurs high behavioral cost across settings

• Infomax: does well for no switch cost, unclear how to modify threshold for each setting, performance suffers otherwise

Generalization to More Complex Problems

• Exact solution: computational complex

• O(kn^{k-1}) for grid size n, # hypotheses k

• Grid solution: good/ improves with grid density (Lovejoy, 1991), poor scalability

• Are there simpler approximations that retain context-sensitivity?

Radial Basis Approximation

• Represent/learn the Q-factors/value function in a low dimensional space (radial basis function; RBF; hypothesized to underlie natural visual processing, Poggio & Girosi, 1990)

• At each step, choose m random samples at which to apply the dynamic programming equation

• Update the basis representation using these new values, performing least square, minimum norm regression

• Iterate until convergence

• O($k(mM+M^3)$) for m points at each step, M bases

• Complexity linear vs. exponential in k (# hypotheses, e.g. target locations) for exact solution

Non-parametric Approximation

• Hand-tuning parameters not ideal for complex problems

• One possible solution: Gaussian Process Regression (GPR) (Williams & Rasmussen, 1996), and GPR with Automatic Relevance Determination (ARD)

• Algorithm proceeds similarly, except for representing Q-factor/value function non-parametrically

• O(kN^3) for N sample points used in regression

Results (Approximate Policy)

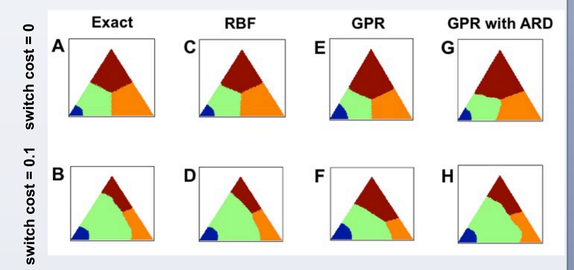

• RBF: # points = 1000, # bases = 49

• GPR: # points = 200, length scale = 1, signal strength = 1, noise strength = 0.1

• GPR with ARD: # points = 200, parameters learned

Conclusions & Discussion

• Active sensing modeled as optimal stochastic control

• Model and humans produce similar behavior (confirmation bias)

• Greedy MAP and Infomax lack context sensitivity (but may be augmented)

• Approximate algorithms (RBF and GPR): much simpler but retain context-sensitivity

• **Future work:**

• ↑# locations • Other methods to learn hyperparameters

• Other tasks • Augment existing models

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. C K I Williams and C E Rasmussen. Gaussian processes for regression. Advances in Neural Information Processing Systems 8, pages 514–20. MIT Press, Cambridge, MA, 1996.

6. W S Lovejoy. Computationally feasible bounds for partially observed Markov decision processes. Operations research 39.1: 162-175, 1991.

7. T Poggio and F Girosi. Networks for approximation and learning. Proceedings of the IEEE, 1990.