Balancing Suspense and Surprise: Timely Decision Making with Endogenous Information Acquisition

AHMED M. ALAA† AND MIHAELA VAN DER SCHAAR†,*, †UCLA, * UNIVERSITY OF OXFORD

OBJECTIVES

What is the optimal policy for Active sequential hypothesis testing with an unknown deadline?

- **Active**: observations are costly
- **Sequential**: observable data is a time series
- **Deadline**: opportunities to observe may end at an unknown, random time!

Many applications: Optimal policies for medical screening and diagnosis, cognitive alternative choice tasks, financial investment policies!
Our focus: theoretical analysis for the structure of the optimal policy.
Our framework: Bayesian sequential analysis.

RELATED WORK

- **Sequential hypothesis testing (Wald 1947, Shiryaev 1973)**
- **Quickest detection (Veeravalli 2001)**
- **Optimal experimentation (Smith 2001)**
- **Two-alternative forced choice (2AFC) task (Frazier & Yu 2006)**
- **Adaptive sensing (Candes 2013)**

Our model: Decide WHEN to sense a process, WHEN to stop sensing and WHICH hypothesis to accept under a deadline pressure.

MODEL

• **Time series:** an Observable stationary Markov process $X(t)$. **[Accessible via costly discrete samples]**
• **Hypotheses:** $X(t)$ is generated via the Markov kernel $\mathbb{P}_\theta, \theta \in \{0,1\}$ is latent. **[Binary hypothesis θ]**
• **Stopping time:** If $\theta = 1$, the process $X(t)$ is stopped at an $\mathcal{F}$-stopping time τ with a Markovian stopping rule. If $\theta = 0$, the process never stops. **[Hypothesis $\theta = 1$ models an adverse event]**
• **A decision-making policy π:** a stopping time T_π, a decision (estimate of θ) $\hat{\theta}_\pi$, and a partitioning $P^\pi_{T_\pi}$ of $X^\tau(t)$, with a **loss function:**

$$\ell(\pi;\Theta) \triangleq \Big(\underbrace{C_1\,\mathbf{1}_{\{\hat{\theta}_\pi=0,\theta=1\}}}_{\text{Type I error}} + \underbrace{C_o\,\mathbf{1}_{\{\hat{\theta}_\pi=1,\theta=0\}}}_{\text{Type II error}} + \underbrace{C_d\,T_\pi}_{\text{Delay}}\Big)\mathbf{1}_{\{T_\pi\leq\tau\}} + \underbrace{C_r\,\mathbf{1}_{\{T_\pi>\tau\}}}_{\text{Deadline missed}} + \underbrace{C_s N(P^\pi_{T_\pi\wedge\tau})}_{\text{Information}}$$

• **Risk of a policy π:** $R(\pi) \triangleq \mathbb{E}\left[\ell(\pi;\Theta)\right]$, the optimal policy is $\pi^* = \arg\inf_{\pi\in\Pi} R(\pi)$.

• **Optimal policy π^* is computed via the Bellman optimality condition.**
• **The decision-maker has two sources of information that govern her beliefs at time t:**
▷ Information conveyed in the realization $X(P_t)$.
▷ Information conveyed in the process survival up to time t.
• **The decision-maker decides either to stop observations and declare a hypothesis $\hat{\theta}$, or continue observations and acquire a new sample at a specific future date $t + \delta_t$.**

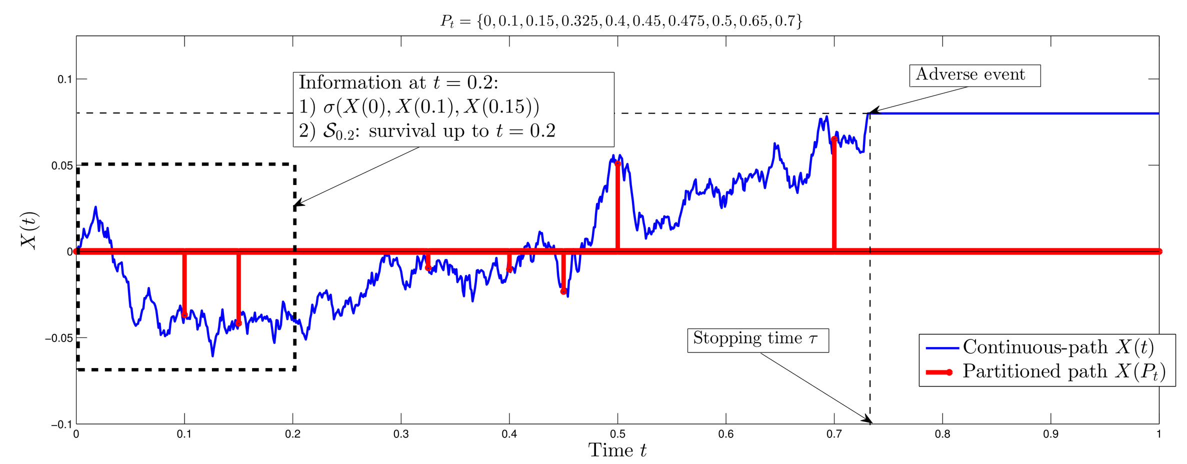

Figure 1: Exemplary sample path for $X(t)$ and π.

POSTERIOR BELIEFS

• The posterior belief $\mu_t = \mathbb{P}(\Theta = 1|\tilde{\mathcal{F}}_t)$ is given by

$$\mu_t = \begin{cases} 1, & \text{for } t \geq \tau \\ \left(1 + \frac{1-p}{p}\frac{d\tilde{\mathbb{P}}_o(P^\pi_t)}{d\tilde{\mathbb{P}}_1(P^\pi_t)}\right)^{-1}, & \text{for } 0 \leq t < \tau \end{cases}$$

where $\frac{d\tilde{\mathbb{P}}_o(P^\pi_t)}{d\tilde{\mathbb{P}}_1(P^\pi_t)}$ depends on the **likelihood ratio**

$$\frac{\mathbb{P}(X(P^\pi_t)|\theta = 1)}{\mathbb{P}(X(P^\pi_t)|\theta = 0)},$$

and the posterior **survival probability**

$$\mathbb{P}(\tau > t|\sigma(X(P^\pi_t), \theta = 1).$$

• **The process survival biases the belief process to the hypothesis $\theta = 0$: for every policy $\pi \in \Pi$, we have that**

$$\mathbb{E}[\mu_{t+\Delta t}|\mathcal{F}_t] \leq \mu_t,\ \forall \Delta t \in \mathbb{R}_+,$$

i.e. the belief process is a supermartingale.

SUSPENSE & SURPRISE

• The belief process μ_t governs the decision-maker's actions
• The decision-maker's actions shape her beliefs: information in $\sigma(X(P^\pi_t))$ depends on P^π_t.
• **Trade-off:** Less intense sampling reduces the cost but bears the risk of not declaring $\hat{\theta}$ before the stopping time τ.
• **Key quantities:**
▷ **Information gain** (surprise):

$$I_t(\Delta t) = (\mu_{t+\Delta t} - \mu_t)\,|\mathcal{F}_t$$

The amount of drift in the decision-maker's belief at time $t + \Delta t$ with respect to her belief at time t, given the information available up to time t.
▷ **Survival** (suspense):

$$S_t(\Delta t) = \mathbb{P}(\tau > t + \Delta t|\mathcal{F}_t, \theta = 1)$$

The probability that a process generated with $\theta = 1$ survives up to time $t+\Delta t$ given the information observed up to time t.

STRUCTURE OF THE OPTIMAL POLICY

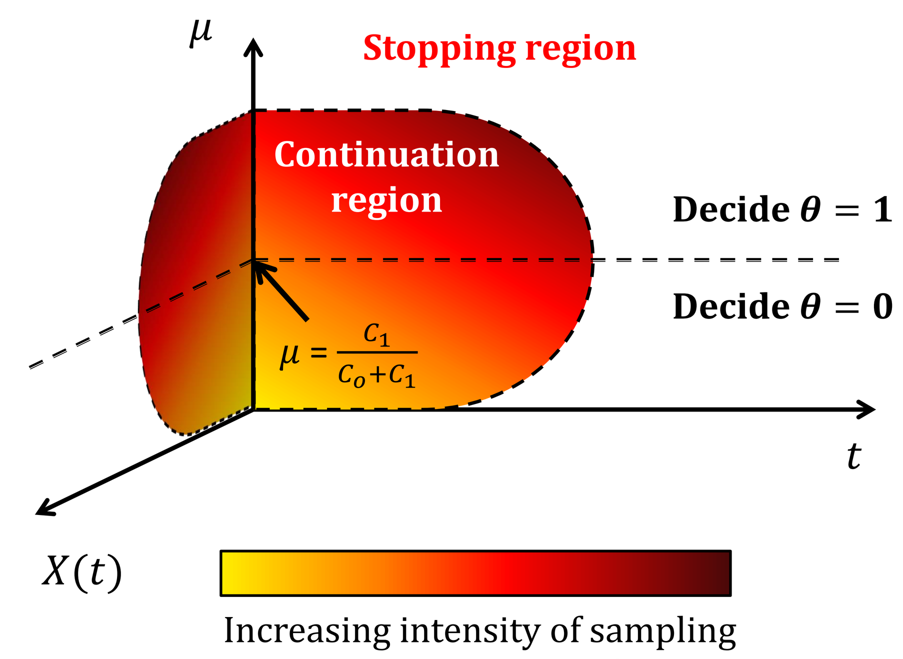

The optimal policy π^* generates a sequence of actions $(\hat{\theta}^{\pi^*}_t, \delta^{\pi^*}_t)$, with the following properties:

- **(Sufficient statistics)**
▷ The process $(t, \mu_t, \bar{X}(P^{\pi^*}_t))_{t\in\mathbb{R}_+}$ is a Markov sufficient statistic for π^*, where $\bar{X}(P^{\pi^*}_t)$ is the most recent sample in $P^{\pi^*}_t$.
- **(Continuation and stopping)**
▷ The policy π^* recommends *continuation* as long as the belief μ_t is in a time and context-dependent *continuation set* $\mathcal{C}(t, \bar{X}(P^{\pi^*}_t))$, where $\mathcal{C}(t', X) \subset \mathcal{C}(t, X), \forall t' > t$.
- **(Suspense and surprise)**
▷ Whenever $\mu_t \in \mathcal{C}(t, \bar{X}(P^{\pi^*}_t))$, the time for acquiring a new sample $\delta^{\pi^*}_t$ is set as follows $\delta^{\pi^*}_t = \arg\inf_{\delta\in\mathbb{R}_+} \mathbb{E}[f(I_t(\delta), S_t(\delta))]$, where $f(I_t(\delta), S_t(\delta))$ is decreasing in $I_t(\delta)$ and $S_t(\delta)$.
- **(Decisions)**
▷ The stopping time is given by $T_{\pi^*} = \inf\{t \in \mathbb{R}_+ : \mu_t \notin \mathcal{C}(t, \bar{X}(P^{\pi^*}_t))\}$, and the corresponding decision is $\hat{\theta}_{\pi^*} = \mathbf{1}_{\left\{\mu_t \geq \frac{C_1}{C_o+C_1}\right\}}$.

CONCLUSIONS

We characterize the optimal policy (continuation, sampling and stopping conditions) for active sequential hypothesis testing with an uncertain deadline.
The optimal policy chooses sampling times to balance information gain (surprise) and survival probability (suspense).