Learning and Optimization for Next Generation Wireless Networks

Tara Javidi

S. Chiu, A. Lalitha, N. Ronquillo, O. Shayevitz, S. Shubhanshu, Y. Kaspi

Learning and Parameter Tuning for Next Generation Networks

- Next-generation wireless systems are increasingly complex

- Next-generation wireless systems are increasingly complex
- Each layer has an increasingly large number of parameters to be optimally tuned

- Next-generation wireless systems are increasingly complex
- Each layer has an increasingly large number of parameters to be optimally tuned
- Networks operate at an increasingly diverse settings

Learning and Optimization for Next Generation Wireless

- ☐ Next-generation wireless systems are increasingly complex
- ☐ Each layer has an increasingly large number of parameters to be optimally tuned
- ☐ Networks operate at an increasingly diverse settings
  - Performance relies on learning and parameter optimization
  - Example: network control's main task involves iterative enhancements of PHY parameters

- Unlike in legacy systems the overhead associated with this learning/optimization can be significant

- □ Unlike in legacy systems the overhead associated with this learning/optimization can be significant
- □ Parameter space is increasingly large and complex

- Unlike in legacy systems the overhead associated with this learning/optimization can be significant
- Parameter space is increasingly large and complex
  - Ultra Wideband spectrum sensing
  - Ultra narrow beam alignment for mmWave communication
  - Empirical network parameter tuning

- Unlike in legacy systems the overhead associated with this learning/optimization can be significant
- Parameter space is increasingly large and complex
  - Ultra Wideband spectrum sensing
  - Ultra narrow beam alignment for mmWave communication
  - Empirical network parameter tuning
- Our objective is to characterize/minimize the network overhead associated w learning/optimization

Spectrum Sensing and Initial Access

Spectrum Sensing: Problem Statement

- Spectrum with total bandwidth of B is available for transmission
- Primary users have dedicated sub-bands of bandwidth δ each

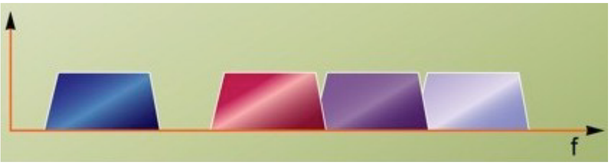

Spectrum Sensing: Problem Statement

- Spectrum with total bandwidth of B is available for transmission
- Primary users have dedicated sub-bands of bandwidth δ each

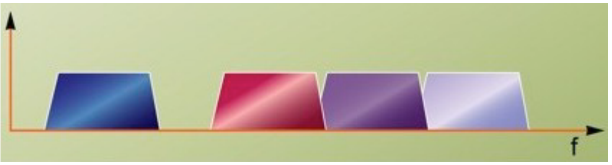

- Subset of subbands inspected sequentially by secondary user

Spectrum Sensing: Problem Statement

- Spectrum with total bandwidth of B is available for transmission
- Primary users have dedicated sub-bands of bandwidth δ each

- Subset of subbands inspected sequentially by secondary user

time	1	$\dots$	$\tau-1$	τ
sample	$A(1)$	$\dots$	$A(\tau-1)$	
observation	$Y(1)$	$\dots$	$Y(\tau-1)$	
declaration				$\hat{W}=d(Y^{\tau-1},x^{\tau-1})$
error				$\mathbf{1}_{\{\hat{W}\neq W\}}$

Spectrum Sensing: Problem Statement

- Spectrum with total bandwidth of B is available for transmission
- Primary users have dedicated sub-bands of bandwidth δ each

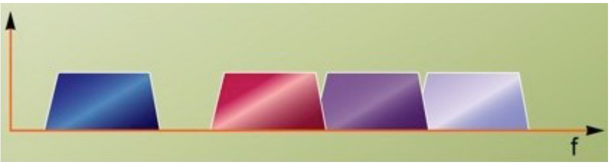

- Subset of subbands inspected sequentially by secondary user

time	1	$\dots$	$\tau - 1$	τ
sample	$A(1)$	$\dots$	$A(\tau - 1)$	
observation	$Y(1)$	$\dots$	$Y(\tau - 1)$	
declaration				$\hat{W} = d(Y^{\tau-1}, x^{\tau-1})$
error				$\mathbf{1}_{\{\hat{W} \neq W\}}$

- Inspection of a subset results in a signal plus noise measurement

Spectrum Sensing: Problem Statement

- Spectrum with total bandwidth of B is available for transmission
- Primary users have dedicated sub-bands of bandwidth δ each

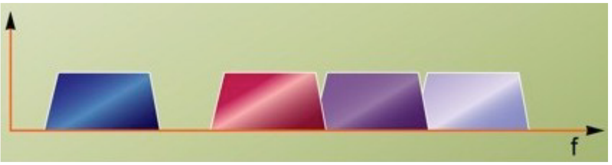

- Subset of subbands inspected sequentially by secondary user
- Inspection of a subset results in a signal plus noise measurement
  - Unit signal associated w the availability of band
  - Sensing noise/unit of spectrum $\approx$ 0-mean, σ^2-var Gaussian

Spectrum Sensing: Problem Statement

- Spectrum with total bandwidth of B is available for transmission
- Primary users have dedicated sub-bands of bandwidth δ each

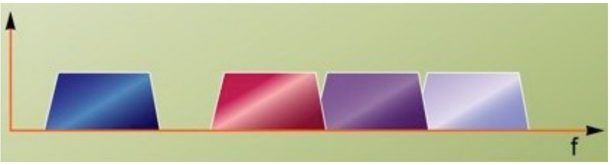

- Subset of subbands inspected sequentially by secondary user
- Inspection of a subset results in a signal plus noise measurement

$$Y^a = \mathbf{a}^T(\mathbf{W} + \mathbf{Z})$$

Spectrum Sensing: Problem Statement

- Spectrum with total bandwidth of B is available for transmission
- Primary users have dedicated sub-bands of bandwidth δ each

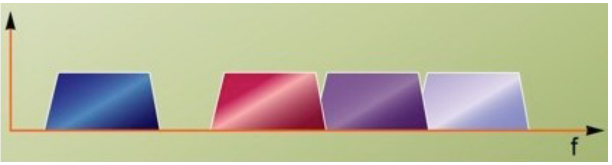

- Subset of subbands inspected sequentially by secondary user
- Inspection of a subset results in a signal plus noise measurement

$$Y^a = \mathbf{a}^T(\mathbf{W} + \mathbf{Z})$$

$$\mathbf{a} \in \mathcal{A}, \quad \mathbf{W} \in \{0,1\}^{\frac{B}{\delta}} \qquad ||W||_0 = K \qquad \mathbf{Z} \sim \mathcal{N}(0, \delta\sigma^2\mathbf{I})$$

Spectrum Sensing: Problem Statement

- Spectrum with total bandwidth of B is available for transmission
- Primary users have dedicated sub-bands of bandwidth δ each

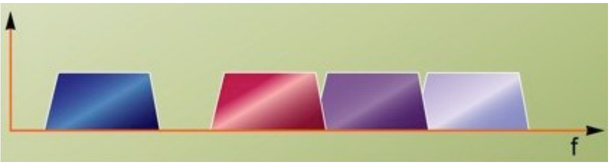

- Subset of subbands inspected sequentially by secondary user
- Inspection of a subset results in a signal plus noise measurement

$$Y^a = \mathbf{a}^T(\mathbf{W} + \mathbf{Z}) \qquad \mathbf{a}, \mathbf{W} \in \{0,1\}^{\frac{B}{\delta}} \quad ||W||_0 = K \quad \mathbf{N} \sim \mathcal{N}(0, B\sigma^2/\delta\mathbf{I})$$

- Minimize $\mathbb{E}\{\tau\}$ subject to Pe $\leq \epsilon$

Spectrum Sensing: Problem Statement

- Spectrum with total bandwidth of B is available for transmission
- Primary users have dedicated sub-bands of bandwidth δ each

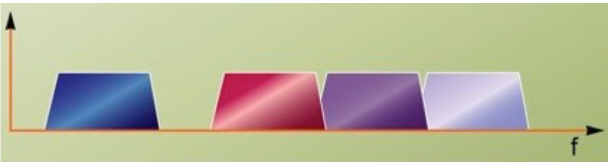

- Subset of subbands inspected sequentially by secondary user
- Inspection of a subset results in a signal plus noise measurement

$$Y^a = \mathbf{a}^T(\mathbf{W} + \mathbf{Z}) \qquad \mathbf{a}, \mathbf{W} \in \{0,1\}^{\frac{B}{\delta}} \quad ||W||_0 = K \quad \mathbf{N} \sim \mathcal{N}(0, B\sigma^2/\delta\mathbf{I})$$

- Minimize $\mathbb{E}\{\tau_\epsilon\}$

Spectrum Sensing: Problem Statement

- Directional transmission $B \subset 2\pi$ is available for transmission
- Angular resolution of $\delta \leq B$

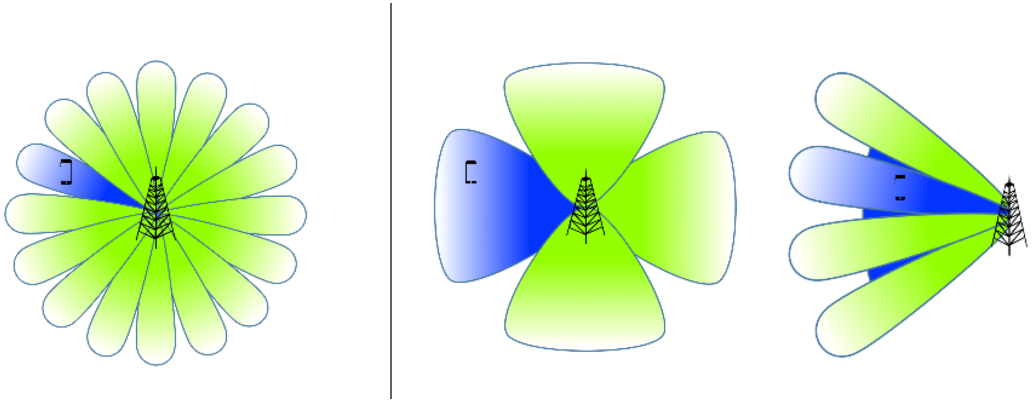

Spectrum Sensing: Problem Statement

- Directional transmission $B \subset 2\pi$ is available for transmission
- Angular resolution of $\delta \leq B$

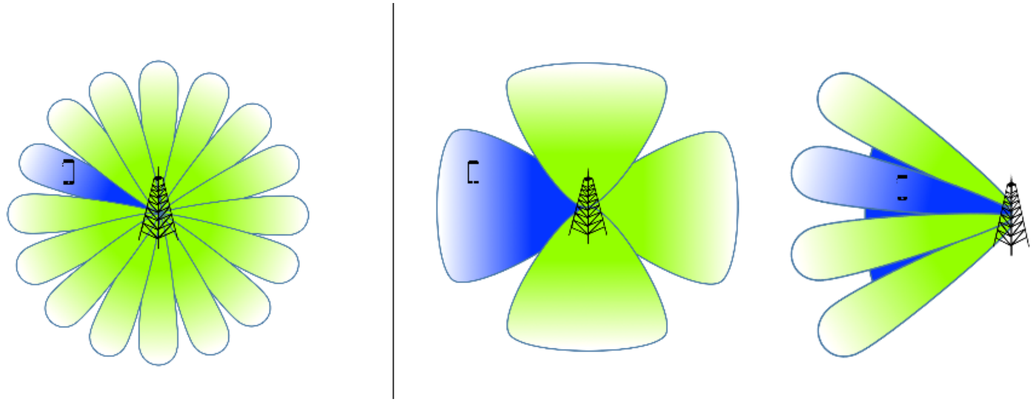

- Subsets of B are used sequentially by transmitter (receiver)

time	1	...	$\tau - 1$	τ
sample	$A(1)$	...	$A(\tau - 1)$	
observation	$Y(1)$	...	$Y(\tau - 1)$	
declaration				$\hat{W} = d(Y^{\tau-1}, x^{\tau-1})$
error				$\mathbf{1}_{\{\hat{W} \neq W\}}$

Spectrum Sensing: Problem Statement

- Directional transmission $B \subset 2\pi$ is available for transmission
- Angular resolution of $\delta \leq B$

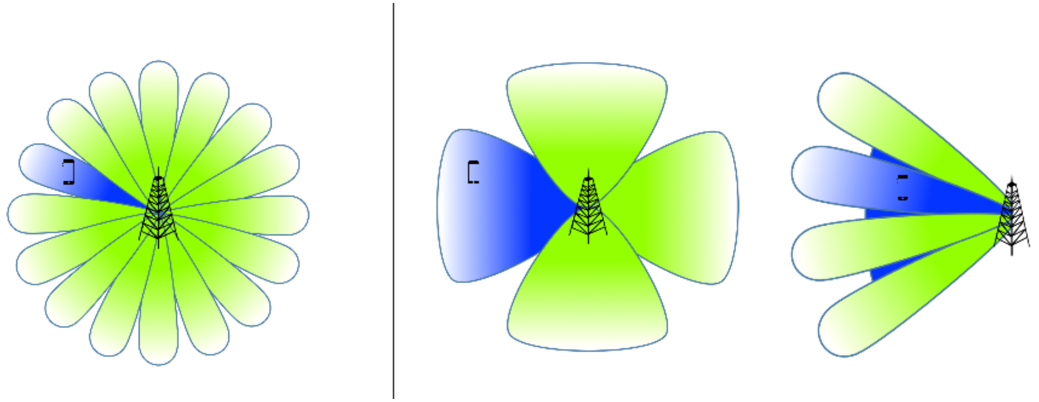

- Subsets of B are used sequentially by transmitter (receiver)
- Inspection of a subset results in a signal plus noise measurement

Spectrum Sensing: Problem Statement

- Directional transmission $B \subset 2\pi$ is available for transmission
- Angular resolution of $\delta \leq B$

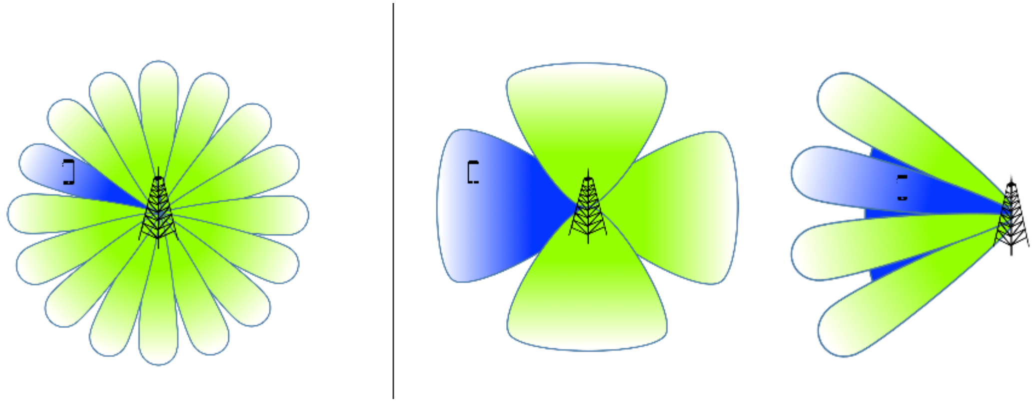

- Subsets of B are used sequentially by transmitter (receiver)
- Inspection of a subset results in a signal plus noise measurement

$$Y^a = \mathbf{a}^T(\mathbf{W} + \mathbf{Z})$$

$$\mathbf{a}, \mathbf{W} \in \{0,1\}^{\frac{B}{\delta}} \qquad ||W||_0 = K \qquad \mathbf{Z} \sim \mathcal{N}(0, \delta\sigma^2\mathbf{I})$$

Spectrum Sensing: Problem Statement

- Directional transmission $B \subset 2\pi$ is available for transmission
- Angular resolution of $\delta \leq B$

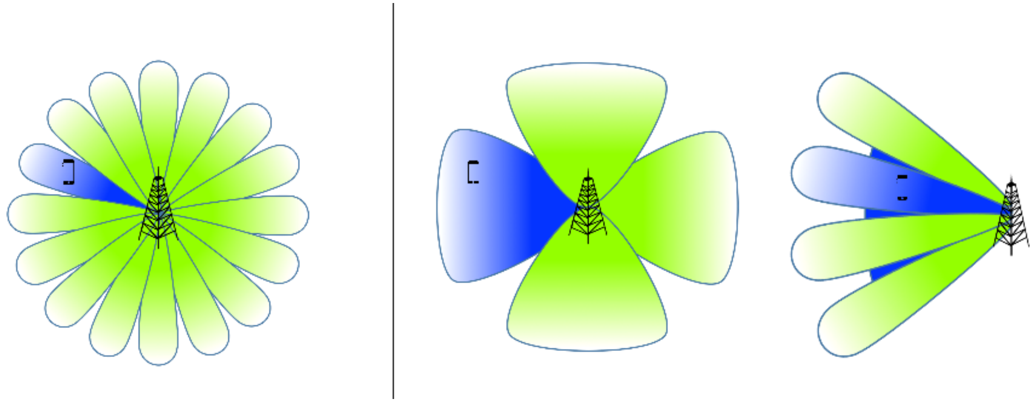

- Subsets of B are used sequentially by transmitter (receiver)
- Inspection of a subset results in a signal plus noise measurement

$$Y^a = \mathbf{a}^T(\mathbf{W} + \mathbf{Z})$$

$$\mathbf{a}, \mathbf{W} \in \{0,1\}^{\frac{B}{\delta}} \qquad ||W||_0 = K \qquad \mathbf{Z} \sim \mathcal{N}(0, \delta\sigma^2\mathbf{I})$$

Spectrum Sensing: Problem Statement

- Directional transmission $B \subset 2\pi$ is available for transmission
- Angular resolution of $\delta \leq B$

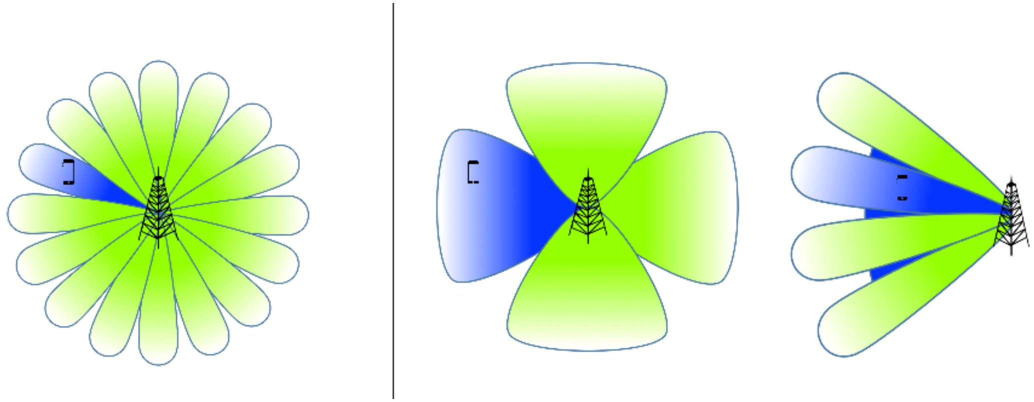

- Subsets of B are used sequentially by transmitter (receiver)
- Inspection of a subset results in a signal plus noise measurement

$$Y^a = \mathbf{a}^T(\mathbf{W} + \mathbf{Z})$$

$$\mathbf{a}, \mathbf{W} \in \{0,1\}^{\frac{B}{\delta}} \qquad ||W||_0 = K \qquad \mathbf{Z} \sim \mathcal{N}(0, \delta\sigma^2\mathbf{I})$$

- Minimize $\mathbb{E}\{\tau_\epsilon\}$

Measurement-Dependent Noisy Search

Measurement-Dependent Noisy Search

- Uknown parameter: $W \in \{0,1\}^{\frac{B}{\delta}}$, $||W||_0 = 1$
- Actions $A(t) \in \mathcal{A} \subset \{0,1\}^{\frac{B}{\delta}}$ chosen sequentially
- $Y(t) = A(t)(\mathbf{W} + \mathbf{Z})$

Measurement-Dependent Noisy Search

- □ Uknown parameter: $W \in \{0,1\}^{\frac{B}{\delta}}$, $||W||_0 = 1$
- □ Actions $A(t) \in \mathcal{A} \subset \{0,1\}^{\frac{B}{\delta}}$ chosen sequentially
- □ $Y(t) = A(t)(\mathbf{W} + \mathbf{Z}) = A(t)\mathbf{W} + \hat{Z}$
  - – Observation noise variance increases w $|A(t)|$

Measurement-Dependent Noisy Search

- ☐ Uknown parameter: $W \in \{0,1\}^{\frac{B}{\delta}}$, $||W||_0 = 1$
- ☐ Actions $A(t) \in \mathcal{A} \subset \{0,1\}^{\frac{B}{\delta}}$ chosen sequentially
- ☐ $Y(t) = A(t)(\mathbf{W} + \mathbf{Z}) = A(t)\mathbf{W} + \hat{Z}$
  - Observation noise variance increases w $|A(t)|$

time	1	...	$\tau - 1$	τ
sample	$A(1)$	...	$A(\tau-1)$	
observation	$Y(1)$	...	$Y(\tau-1)$	
declaration				$\hat{W} = d(Y^{\tau-1}, x^{\tau-1})$
error				$\mathbf{1}_{\{\hat{W} \neq W\}}$

Objective:

Find τ, $A(0), \ldots, A(\tau-1)$, and $d(\cdot)$ that minimize $\mathbb{E}[\tau]$ s.t. Pe $\leq \epsilon$

- Uknown parameter: $W \in \{0,1\}^{\frac{B}{\delta}}$, $||W||_0 = 1$
- Actions $A(t) \in \mathcal{A} \subset \{0,1\}^{\frac{B}{\delta}}$ chosen sequentially
- $Y(t) = A(t)(\mathbf{W} + \mathbf{Z}) = A(t)\mathbf{W} + \hat{Z}$
  - Observation noise variance increases w $|A(t)|$

time	1	$\dots$	$\tau - 1$	τ
sample	$A(1)$	$\dots$	$A(\tau - 1)$	
observation	$Y(1)$	$\dots$	$Y(\tau - 1)$	
declaration				$\hat{W} = d(Y^{\tau-1}, x^{\tau-1})$
error				$\mathbf{1}_{\{\hat{W} \neq W\}}$

Objective:
Find τ, $A(0), \dots, A(\tau - 1)$, and $d(\cdot)$ that minimize $\mathbb{E}[\tau]$ s.t. Pe $\leq \epsilon$

- Numerical solution via a dynamic programming equation

Simpler Questions of General Consequence

- □ Role of allowable actions set $\mathcal{A}$

Simpler Questions of General Consequence

- Role of allowable actions set $\mathcal{A}$
  - Designing $\mathcal{A}$ can significantly reduce the overhead

Simpler Questions of General Consequence

- Role of allowable actions set $\mathcal{A}$
  - Designing $\mathcal{A}$ can significantly reduce the overhead
    - Even though noise variance increases w $|a|$ linearly!

- Role of allowable actions set $\mathcal{A}$
  - Designing $\mathcal{A}$ can significantly reduce the overhead
    - Even though noise variance increases w $|a|$ linearly!
- Selecting $A(t)$ based on past observations (a feedback scheme) or off-line (non-adaptively)?

Simpler Questions of General Consequence

- Role of allowable actions set $\mathcal{A}$
  - Designing $\mathcal{A}$ can significantly reduce the overhead
    - Even though noise variance increases w $|a|$ linearly!
- Selecting $A(t)$ based on past observations (a feedback scheme) or off-line (non-adaptively)?
  - What is the adaptivity gain?
  - Feedback policies are computationally expensive

Role of Measurements

- □ Role of allowable actions set $\mathcal{A}$

Role of Measurements

- Role of allowable actions set $\mathcal{A}$
  - Advantages of group testing

Role of Measurements

- Role of allowable actions set $\mathcal{A}$
  - Advantages of group testing

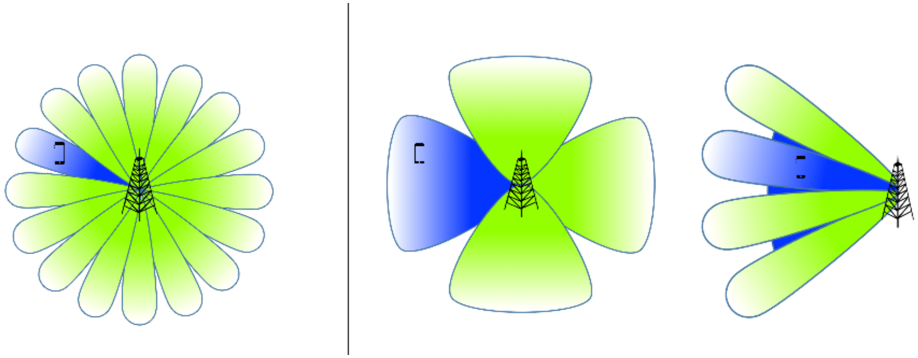

Role of Measurements

- Role of allowable actions set $\mathcal{A}$
  - Advantages of group testing

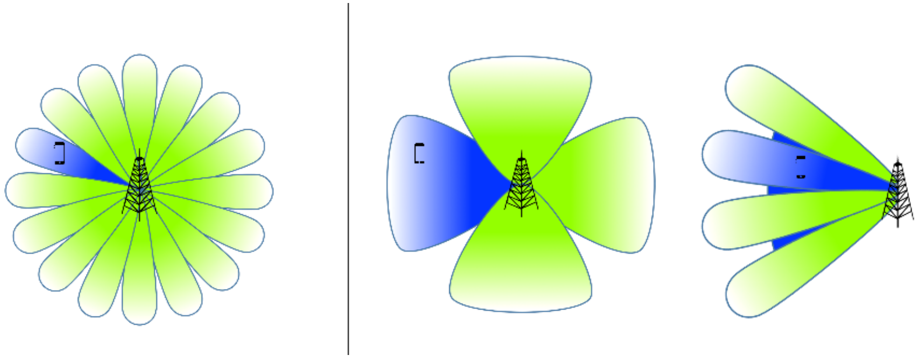

  - If $\mathcal{A}$ only singletons ($||A(t)|| = 1$) $\Rightarrow$ search time $\mathcal{O}(B/\delta)$
  - If $\mathcal{A}$ includes intervals, can be $\mathcal{O}\left(\log(B/\delta\epsilon)\right)$

Role of Measurements

- Role of allowable actions set $\mathcal{A}$
  - Advantages of group testing

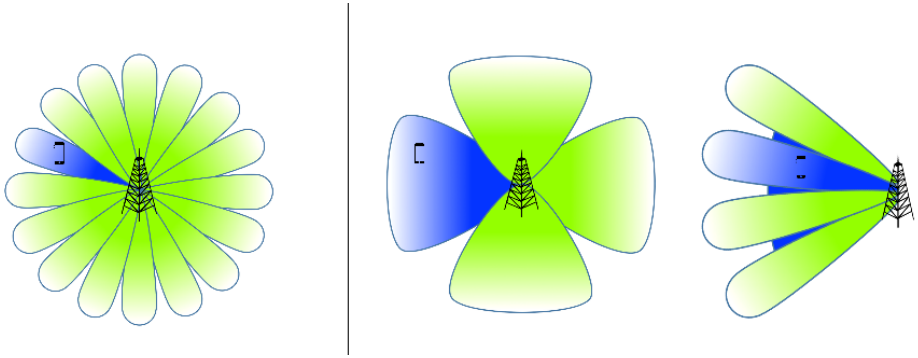

  - If $\mathcal{A}$ only singletons ($||A(t)|| = 1$) $\Rightarrow$ search time $\mathcal{O}(B/\delta)$
  - If $\mathcal{A}$ includes intervals, can be $\mathcal{O}\left(\log(B/\delta\epsilon)\right)$

Observation:

If $Y^a = \overbrace{\mathbf{1}_{\{\text{object in } a\}}}^{X} + Z,\ Z \sim \mathcal{N}(0, \sigma_z^2), \Rightarrow \mathbb{E}[\tau] \approx \frac{\log B/\delta\epsilon}{I(X, Y^a)}$

Adaptivity Gain

- Selecting $A(t)$ based on past observations (a feedback scheme) is computationally expensive

Adaptivity Gain

- Selecting $A(t)$ based on past observations (a feedback scheme) is computationally expensive
- Critical to quantify the **Adaptivity (feedback) gain** $\mathbb{E}\left[\tau_\epsilon\right]$:

$$\mathbb{E}\left[\tau_\epsilon^{na}\right] - \mathbb{E}\left[\tau_\epsilon^{*}\right]$$

Adaptivity Gain

- Selecting $A(t)$ based on past observations (a feedback scheme) is computationally expensive
- Critical to quantify the **Adaptivity (feedback) gain** $\mathbb{E}\left[\tau_\epsilon\right]$:

$$\mathbb{E}\left[\tau_\epsilon^{na}\right] - \mathbb{E}\left[\tau_\epsilon^{*}\right]$$

- Asymptotic analysis when B/δ grows

Adaptivity Gain

- Selecting $A(t)$ based on past observations (a feedback scheme) is computationally expensive
- Critical to quantify the **Adaptivity (feedback) gain** $\mathbb{E}\left[\tau_\epsilon\right]$:

$$\mathbb{E}\left[\tau_\epsilon^{na}\right] - \mathbb{E}\left[\tau_\epsilon^{*}\right]$$

- Asymptotic analysis when B/δ grows
  - Qualitative difference when B grows versus δ shrinks

Adaptivity Gain

- Selecting $A(t)$ based on past observations (a feedback scheme) is computationally expensive
- Critical to quantify the **Adaptivity (feedback) gain** $\mathbb{E}\left[\tau_\epsilon\right]$:

$$\mathbb{E}\left[\tau_\epsilon^{na}\right] - \mathbb{E}\left[\tau_\epsilon^{*}\right]$$

- Asymptotic analysis when B/δ grows
  - Qualitative difference when B grows versus δ shrinks
    - When B grows overall noise variance grows
    - Overall noise is constant even when $1/\delta$ grows

Adaptivity Gain

- Selecting $A(t)$ based on past observations (a feedback scheme) is computationally expensive

- Critical to quantify the **Adaptivity (feedback) gain** $\mathbb{E}\left[\tau_\epsilon\right]$:

$$\mathbb{E}\left[\tau_\epsilon^{na}\right] - \mathbb{E}\left[\tau_\epsilon^{*}\right]$$

- Asymptotic analysis when B/δ grows
  - Qualitative difference when B grows versus δ shrinks
    - When B grows overall noise variance grows
    - Overall noise is constant even when $1/\delta$ grows
  - Need for a fairly **tight non-asymptotic** analysis

Our Contributions: Main Take-aways (general K, $K = 1$)

- Searching with codebooks with feedback over a stateful channel

- Searching with codebooks with feedback over a stateful channel

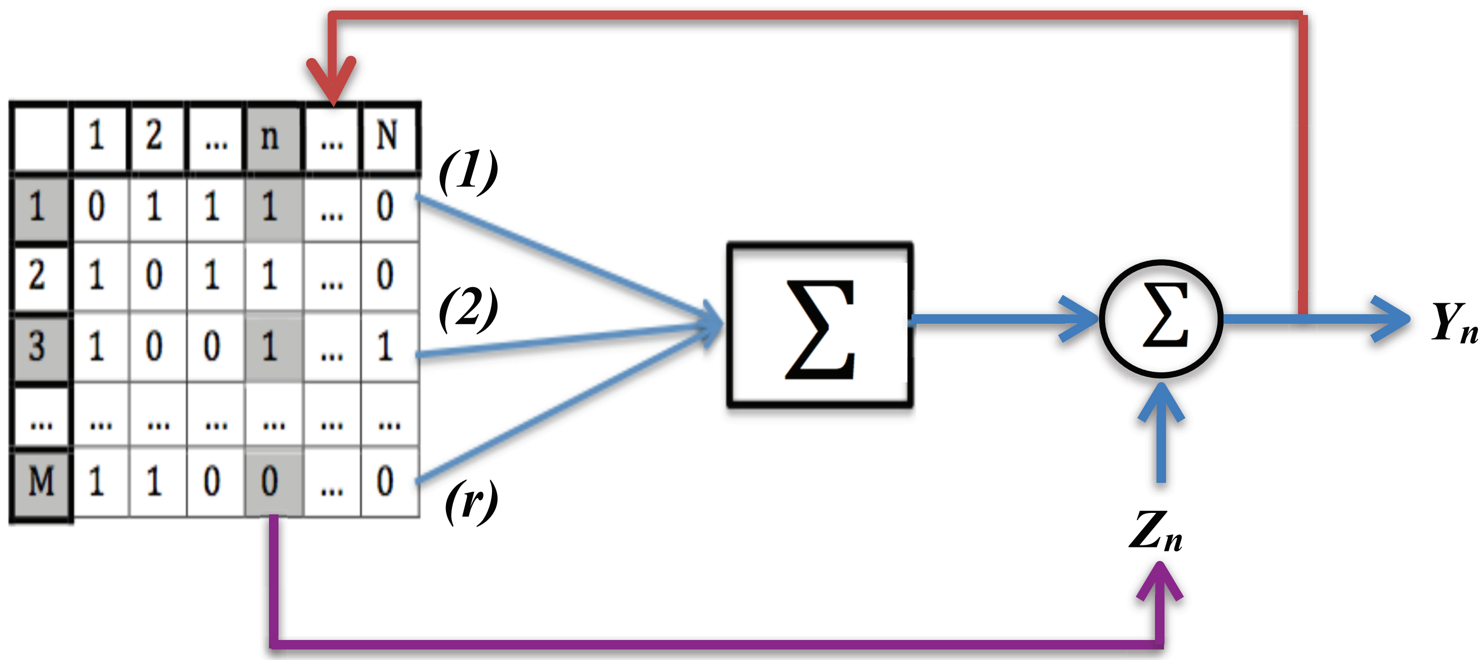

Our Contributions: Main Take-aways (general K, $K = 1$)

- Searching with codebooks with feedback over a stateful channel ($K = 1$)
  - Reduces the non-adaptive case to known IT problems
  - Adaptive strategy as a variant of feedback code

Our Contributions: Main Take-aways (general K, $K = 1$)

- Searching with codebooks with feedback over a stateful channel ($K = 1$)
  - Reduces the non-adaptive case to known IT problems
  - Adaptive strategy as a variant of feedback code
- Non-asymptotic achievability analysis for an adaptive scheme
  - Sorted Posterior Matching (SortPM) search strategy

Our Contributions: Main Take-aways (general K, $K = 1$)

- Searching with codebooks with feedback over a stateful channel ($K = 1$)
  - Reduces the non-adaptive case to known IT problems
  - Adaptive strategy as a variant of feedback code
- Non-asymptotic achievability analysis for an adaptive scheme
  - Sorted Posterior Matching (SortPM) search strategy
- Characterize daptivity gain with two distinct asymptotic regimes $B/\delta \to \infty$

Our Contributions: Main Take-aways (general K, $K = 1$)

- Searching with codebooks with feedback over a stateful channel ($K = 1$)
  - Reduces the non-adaptive case to known IT problems
  - Adaptive strategy as a variant of feedback code
- Non-asymptotic achievability analysis for an adaptive scheme
  - Sorted Posterior Matching (SortPM) search strategy
- Characterize daptivity gain with two distinct asymptotic regimes $B/\delta \to \infty$
  - Fixed search interval and increasing resolution (initial access)
  - Fixed resolution and increasing search (primary user detection)

- Searching with codebooks with feedback over a stateful channel ($K = 1$)
  - Reduces the non-adaptive case to known IT problems
  - Adaptive strategy as a variant of feedback code

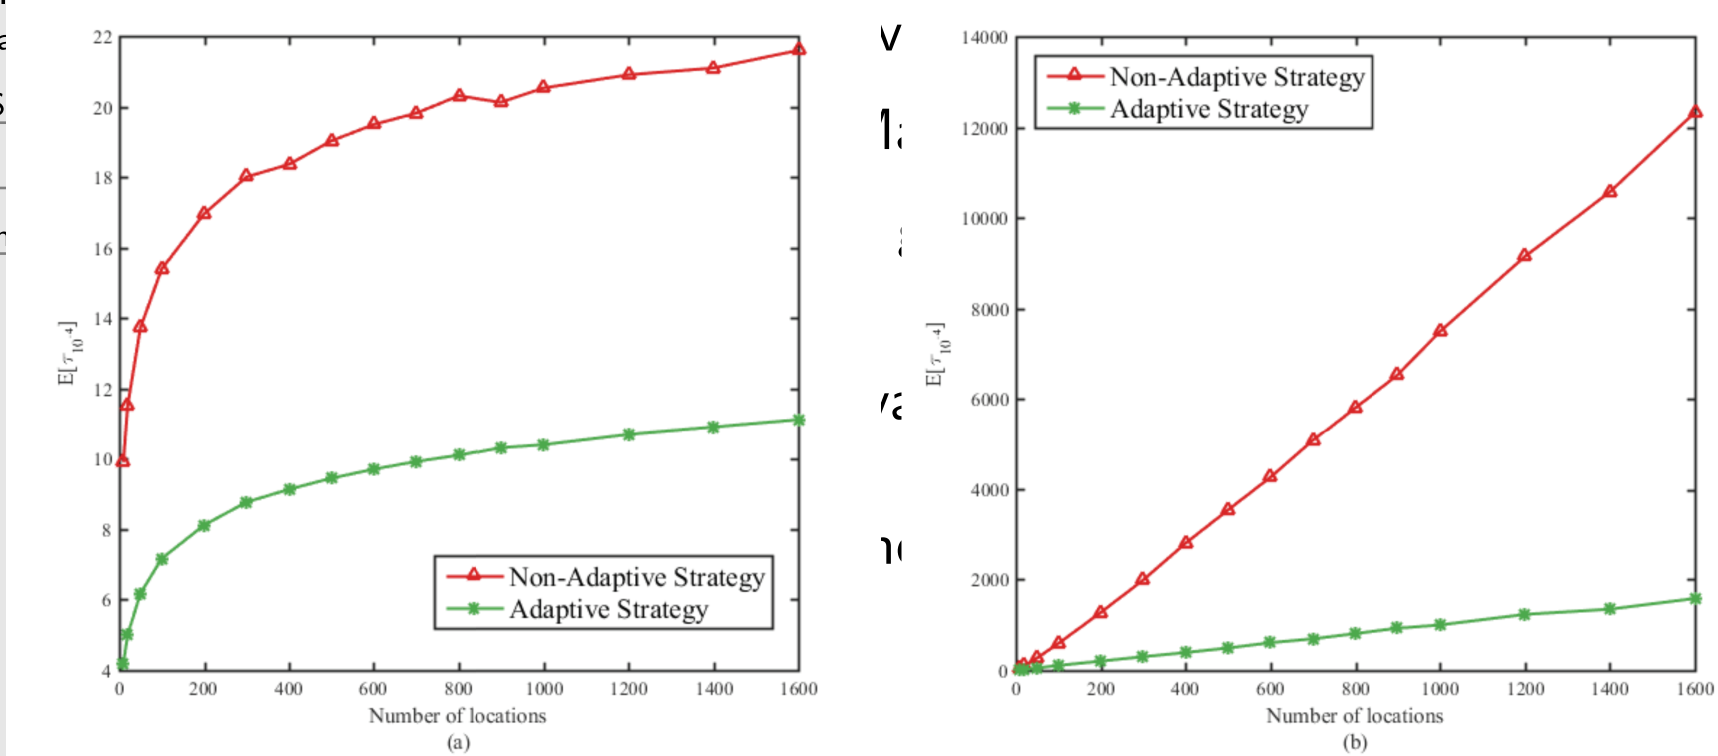

Analysis

Non-asymptotic Converse for Non-adaptive Search:

- Searching via coding over a stateful channel

Non-asymptotic Converse for Non-adaptive Search:

- ☐ Searching via coding over a stateful channel

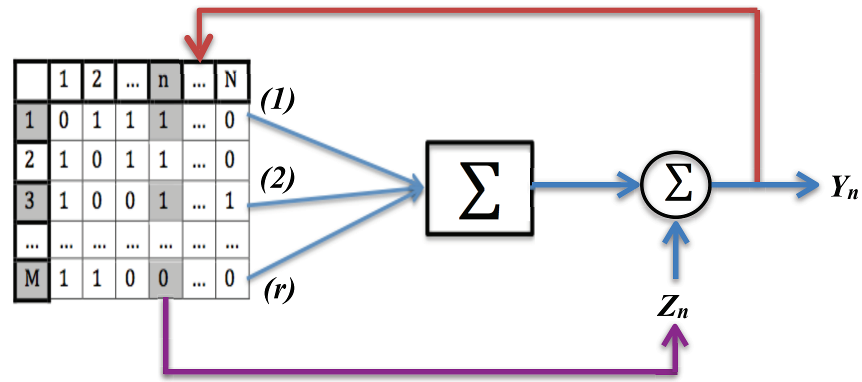

- Searching via coding over a stateful channel

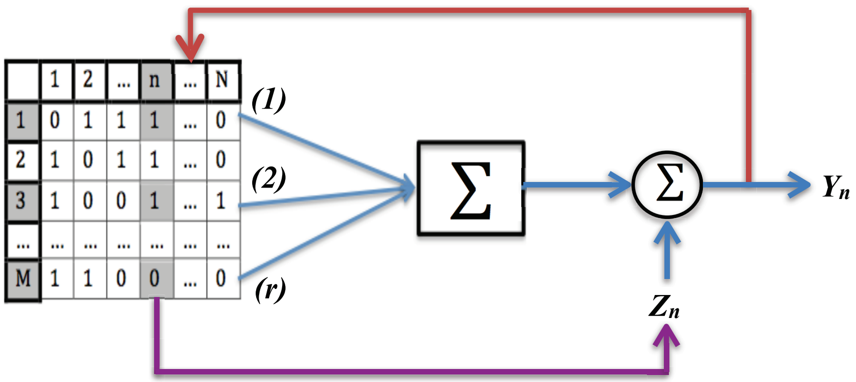

– Reduces non-adaptive case to known IT problem:

$$Y = X^q + Z^q, \qquad X^q \sim \text{Ber}(q),\ Z^q \sim \mathcal{N}(0, \tfrac{qB}{\delta}\sigma^2)$$

Non-asymptotic Converse for Non-adaptive Search:

- Searching via coding over a stateful channel

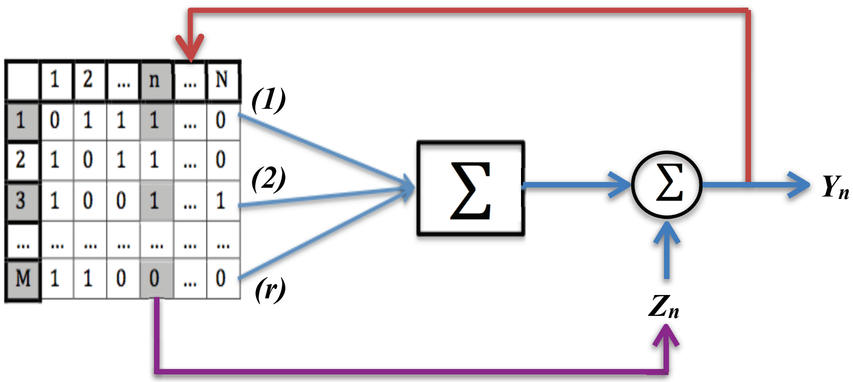

– Reduces non-adaptive case to known IT problem:

$$Y = X^q + Z^q, \qquad X^q \sim \text{Ber}(q),\ Z^q \sim \mathcal{N}(0, \tfrac{qB}{\delta}\sigma^2)$$

$$\mathbb{E}[\tau_\epsilon^{\text{NA}}] \geq \frac{(1-\epsilon)\log\frac{B}{\delta} - h(\epsilon)}{C_{\text{BPSK}}(q, \sigma\sqrt{qB/\delta})}$$

Non-asymptotic Converse for Non-adaptive Search:

- Searching via coding over a stateful channel

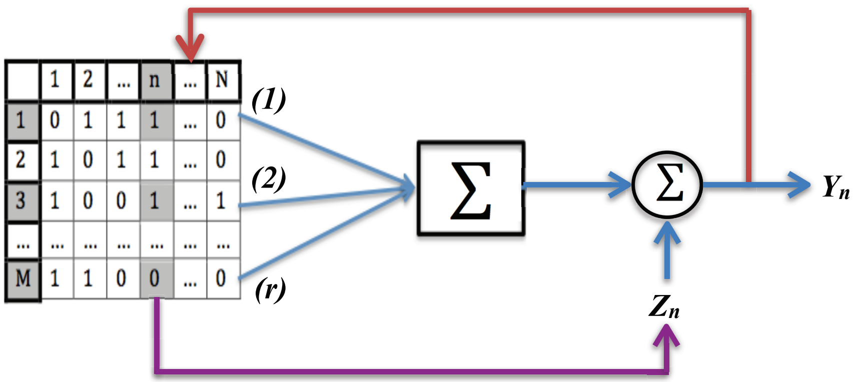

– Reduces non-adaptive case to known IT problem:

$$Y = X^q + Z^q, \qquad X^q \sim \text{Ber}(q),\ Z^q \sim \mathcal{N}(0, \tfrac{qB}{\delta}\sigma^2)$$

$$\mathbb{E}[\tau_\epsilon^{\text{NA}}] \geq \frac{(1-\epsilon)\log\frac{B}{\delta} - h(\epsilon)}{C_{\text{BPSK}}(q^*, \sigma\sqrt{q^*B/\delta})}$$

Non-asymptotic Converse for Non-adaptive Search:

- Searching via coding over a stateful channel

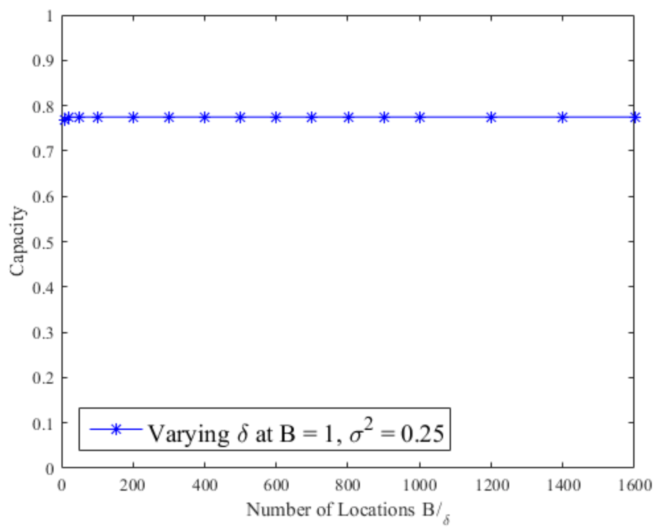

– Reduces non-adaptive case to known IT problem:

$$Y = X^q + Z^q, \qquad X^q \sim \mathrm{Ber}(q),\ Z^q \sim \mathcal{N}(0, \tfrac{qB}{\delta}\sigma^2)$$

$$\mathbb{E}[\tau_\epsilon^{\mathsf{NA}}] \geq \frac{(1-\epsilon)\log\frac{B}{\delta} - h(\epsilon)}{C_{\mathsf{BPSK}}(q^*, \sigma\sqrt{q^* B/\delta})}$$

Non-asymptotic Converse for Non-adaptive Search:

- Searching via coding over a stateful channel

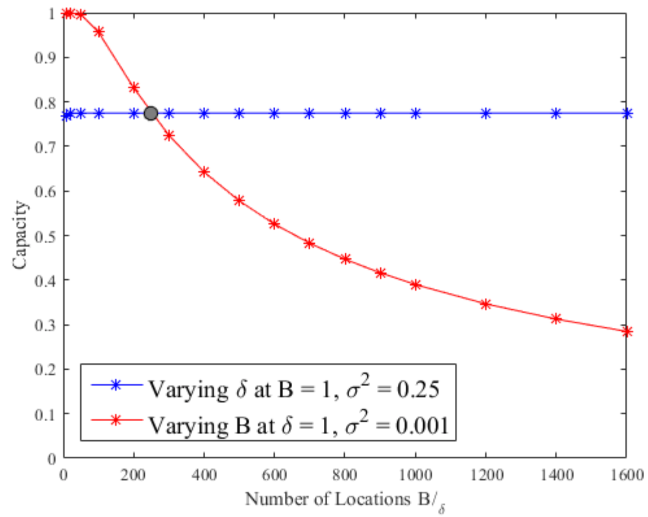

– Reduces non-adaptive case to known IT problem:

$$Y = X^q + Z^q, \qquad X^q \sim \mathrm{Ber}(q),\ Z^q \sim \mathcal{N}(0, \tfrac{qB}{\delta}\sigma^2)$$

$$\mathbb{E}[\tau_\epsilon^{\mathsf{NA}}] \geq \frac{(1-\epsilon)\log\frac{B}{\delta} - h(\epsilon)}{C_{\mathsf{BPSK}}(q^*, \sigma\sqrt{q^* B/\delta})}$$

Non-adaptive and Adaptive Search Strategies

Non-adaptive Strategy:

Non-adaptive and Adaptive Search Strategies

Non-adaptive Strategy:

Fix the number of samples $\tau = T$

Non-adaptive Strategy:
Fix the number of samples $\tau = T$

□ select T to be such that $\mathbb{E}\{P_e\} \leq \epsilon$

Non-adaptive Strategy:
Fix the number of samples $\tau = T$

- select T to be such that $\mathbb{E}\{P_e\} \leq \epsilon$
- for all $t \leq T$ query random set a such that $|a| = q^* B/\delta$ optimized

Sorted Posterior Matching (sortPM) Strategy:

Consider prior $\boldsymbol{\rho}(t) := (\mathbb{P}\{W = e_i | A(0:t-1), Y(0,t-1)\})$

- declares i as the target, if $\rho_i(t) \geq 1 - \epsilon, \; i \in \Omega$
- otherwise, queries the bins left of the median of the sorted prior
  - observe (noisy) Y
  - update the prior (posterior) via the Bayes' rule

Sorted Posterior Matching (sortPM) Strategy:

Consider prior $\boldsymbol{\rho}(t) := (\mathbb{P}\{W = e_i | A(0:t-1), Y(0,t-1)\})$

- declares i as the target, if $\rho_i(t) \geq 1-\epsilon, \;\; i \in \Omega$
- otherwise, queries the bins left of the median of the sorted prior
  - observe (noisy) Y
  - update the prior (posterior) via the Bayes' rule

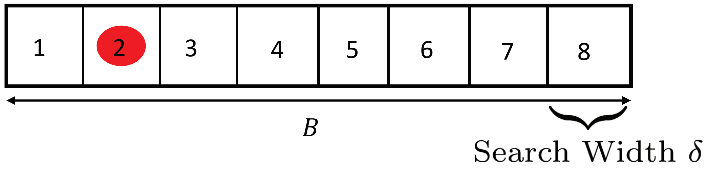

Sorted Posterior Matching (sortPM) Strategy:

Consider prior $\boldsymbol{\rho}(t) := (\mathbb{P}\{W = e_i | A(0:t-1), Y(0,t-1)\})$

- ☐ declares i as the target, if $\rho_i(t) \geq 1 - \epsilon, \;\; i \in \Omega$
- ☐ otherwise, queries the bins left of the median of the sorted prior
  - – observe (noisy) Y
  - – update the prior (posterior) via the Bayes' rule

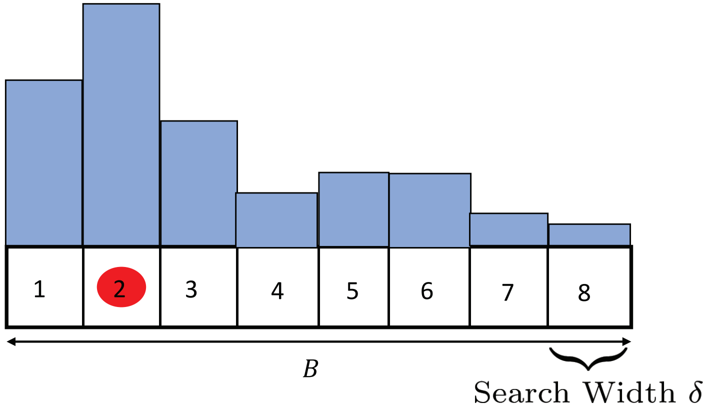

Sorted Posterior Matching (sortPM) Strategy:

Consider prior $\boldsymbol{\rho}(t) := (\mathbb{P}\{W = e_i | A(0:t-1), Y(0,t-1)\})$

- ☐ declares i as the target, if $\rho_i(t) \geq 1-\epsilon, \;\; i \in \Omega$
- ☐ otherwise, queries the bins left of the median of the sorted prior
  - – observe (noisy) Y
  - – update the prior (posterior) via the Bayes' rule

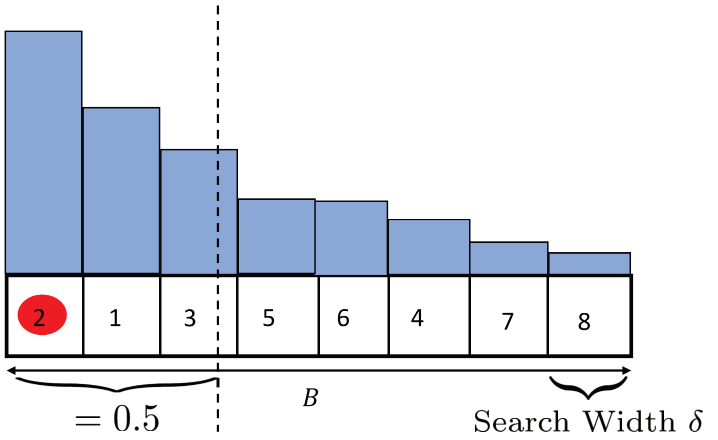

Sorted Posterior Matching (sortPM) Strategy:

Consider prior $\boldsymbol{\rho}(t) := (\mathbb{P}\{W = e_i | A(0:t-1), Y(0,t-1)\})$

- declares i as the target, if $\rho_i(t) \geq 1-\epsilon, \;\; i \in \Omega$
- otherwise, queries the bins left of the median of the sorted prior
  - observe (noisy) Y
  - update the prior (posterior) via the Bayes' rule

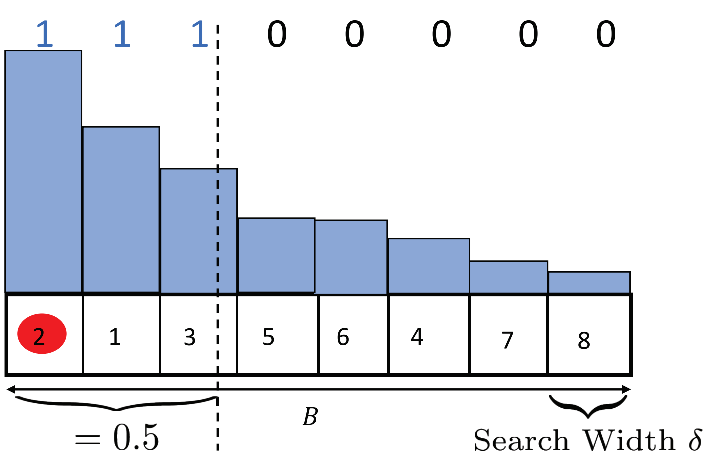

Sorted Posterior Matching (sortPM) Strategy:

Consider prior $\rho(t) := (\mathbb{P}\{W = e_i | A(0:t-1), Y(0,t-1)\})$

- ☐ declares i as the target, if $\rho_i(t) \geq 1 - \epsilon, \;\; i \in \Omega$
- ☐ otherwise, queries the bins left of the median of the sorted prior

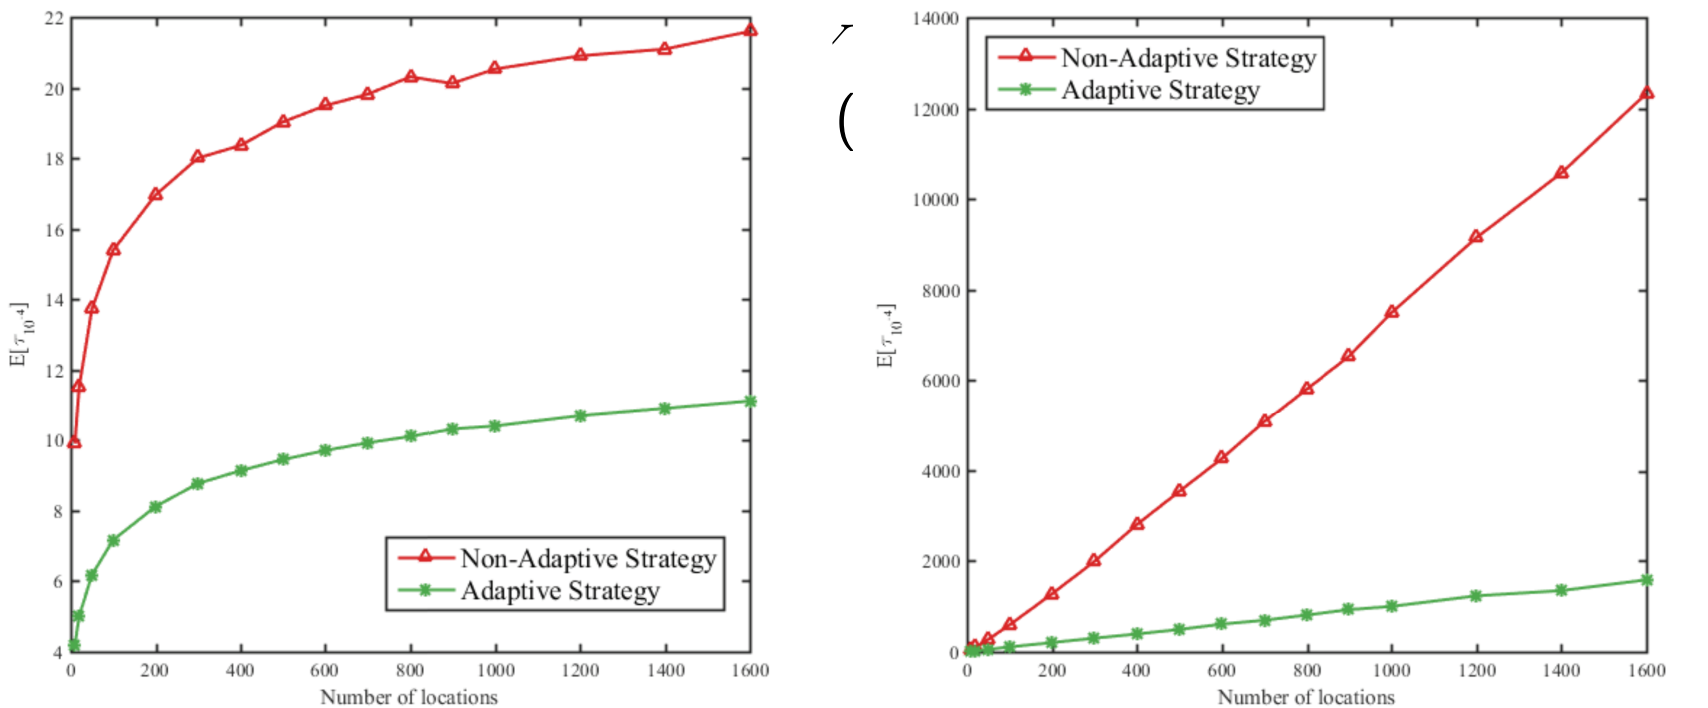

(a) (b)

SortPM: Upper Bound

Theorem. *[Lalitha, Ronquillo and J. 17] Under SortPM, we have*

$$\mathbb{E}[\tau_{SPM}] \leq \min_{\alpha} \frac{\log B/\delta\epsilon + \max\{\log\log B/\delta, \log\log \frac{1}{\epsilon}\}}{1 - h(Q((\sigma^2 \alpha B/\delta)^{-1/2}))} + K(\alpha).$$

where

$$h(p) = p \log \frac{1}{p} + (1-p) \log \frac{1}{1-p},$$

$K(\cdot)$ is non-increasing function

SortPM: Upper Bound

Theorem. *[Lalitha, Ronquillo and J. 17] Under SortPM, we have*

$$\mathbb{E}[\tau_{SPM}] \leq \min_{\alpha} \frac{\log B/\delta\epsilon + \max\{\log\log B/\delta, \log\log \frac{1}{\epsilon}\}}{1 - h(Q((\sigma^2 \alpha B/\delta)^{-1/2}))} + K(\alpha).$$

where

$$h(p) = p \log \frac{1}{p} + (1-p) \log \frac{1}{1-p},$$

$K(\cdot)$ is non-increasing function

✓ Analysis is based on a Lyapunov drift

SortPM: Upper Bound

Corollary. *[Lalitha, Ronqullio and J. 17] Relying on hard-detected output symbols, the asymptotic adaptivity gain for $B/\delta \to \infty$ is:*

$$\lim_{\delta\to 0} \frac{\tau_{opt}^{NA} - \mathbb{E}[\tau_{opt}^{A}]}{\log \frac{B}{\delta}} = \frac{1}{C_{BPSK}(q^*, B\sigma^2)} - 1.$$

$$\lim_{B\to\infty} \frac{\tau_{opt}^{NA} - \mathbb{E}[\tau_{opt}^{A}]}{\frac{B}{\delta}\log \frac{B}{\delta}} \geq \frac{\sigma^2\delta}{\log e}.$$

SortPM: Upper Bound

Corollary. *[Lalitha, Ronqullio and J. 17] Relying on hard-detected output symbols, the asymptotic adaptivity gain for $B/\delta \to \infty$ is:*

$$\lim_{\delta \to 0} \frac{\tau_{opt}^{NA} - \mathbb{E}[\tau_{opt}^{A}]}{\log \frac{B}{\delta}} = \frac{1}{C_{BPSK}(q^*, B\sigma^2)} - 1.$$

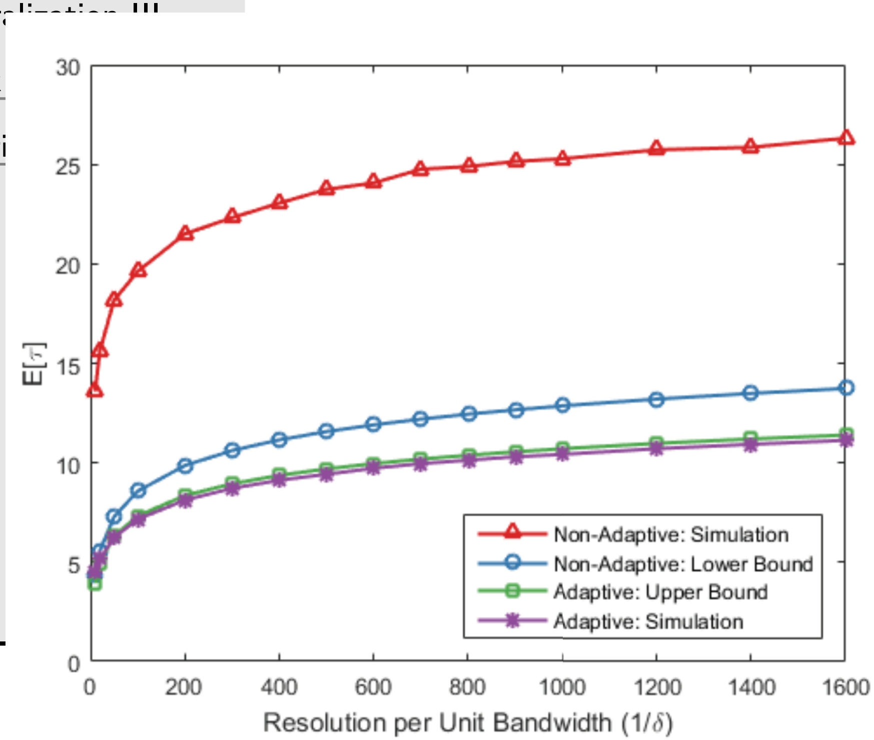

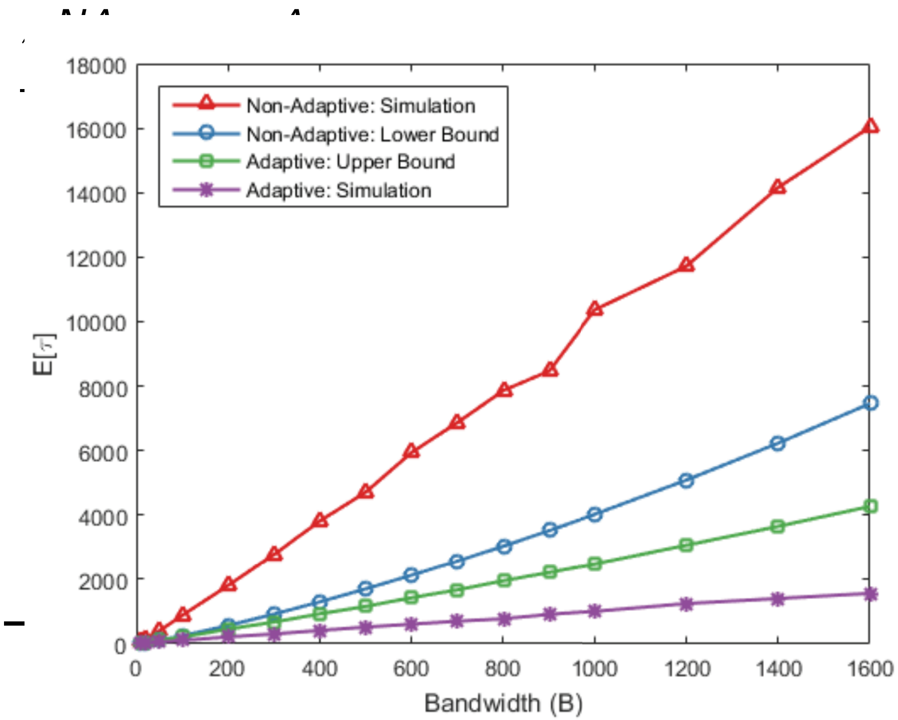

Prior Work: Measurement Independent Noise

- Generalized binary search [Burnashev and Zigangirov '74]
- Channel coding over DMC with feedback [Burnashev '75], [Yamamato and Itoh '79], ... [Naghshvar, Wigger and J '13]
- Posterior matching [Shayevitz and Feder '11]
- Bisection search with noisy responses [Horstein '63], [Waeber, Frazier, Henderson '13]

Generalizations

- General noise model: $Y(t) = A(t)\mathbf{W} + \hat{Z}$, $\hat{Z} = f(\mathbf{Z}, A(t))$

Generalizations

- General noise model: $Y(t) = A(t)\mathbf{W} + \hat{Z}$, $\hat{Z} = f(\mathbf{Z}, A(t))$

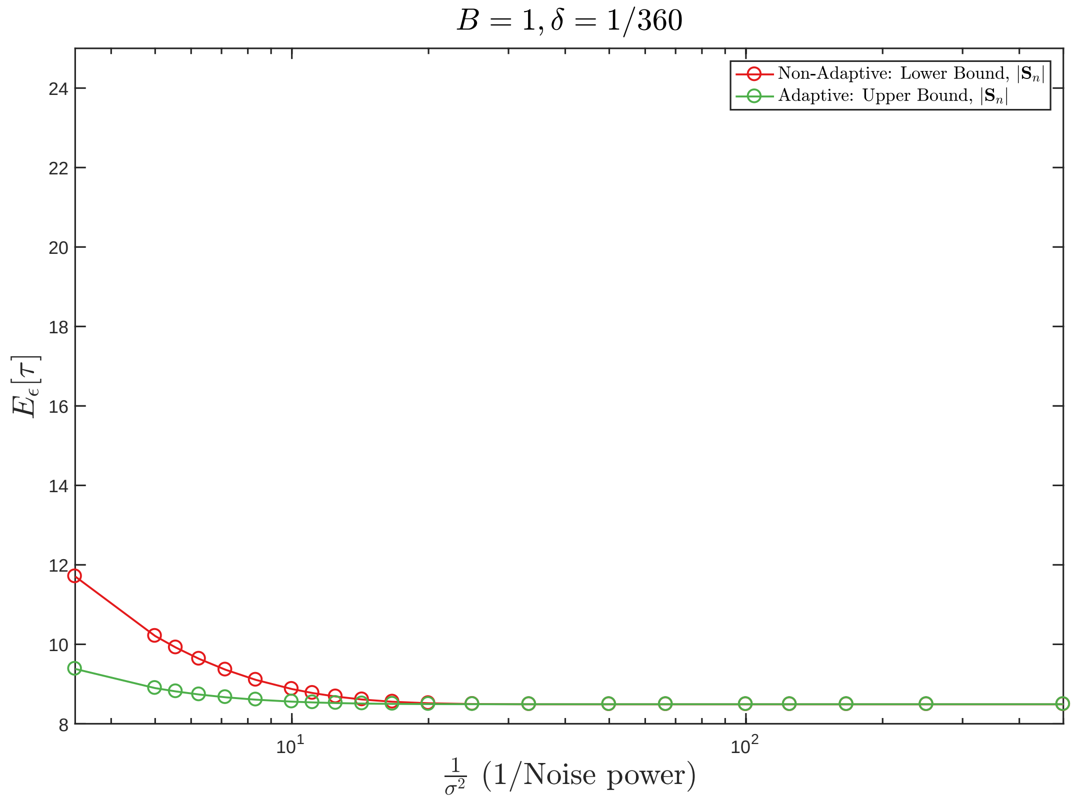

- General noise model: $Y(t) = A(t)\mathbf{W} + \hat{Z}$, $\hat{Z} = f(\mathbf{Z}, A(t))$

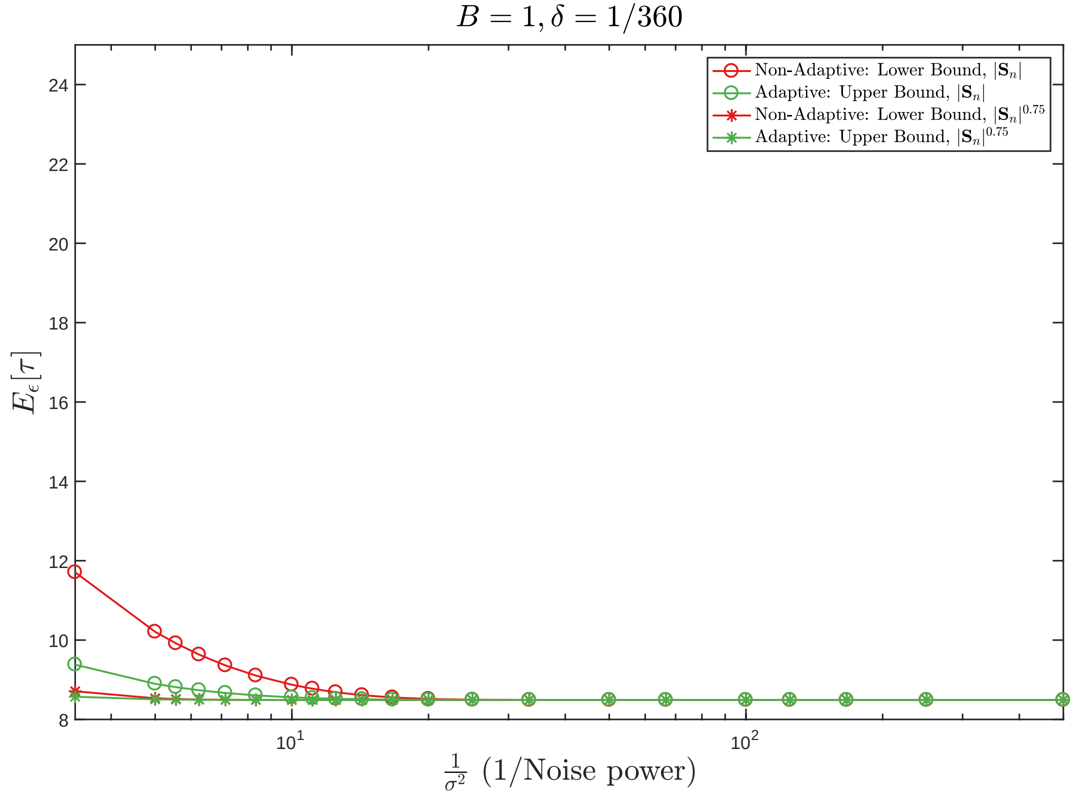

- ☐ General noise model: $Y(t) = A(t)\mathbf{W} + \hat{Z}$, $\hat{Z} = f(\mathbf{Z}, A(t))$

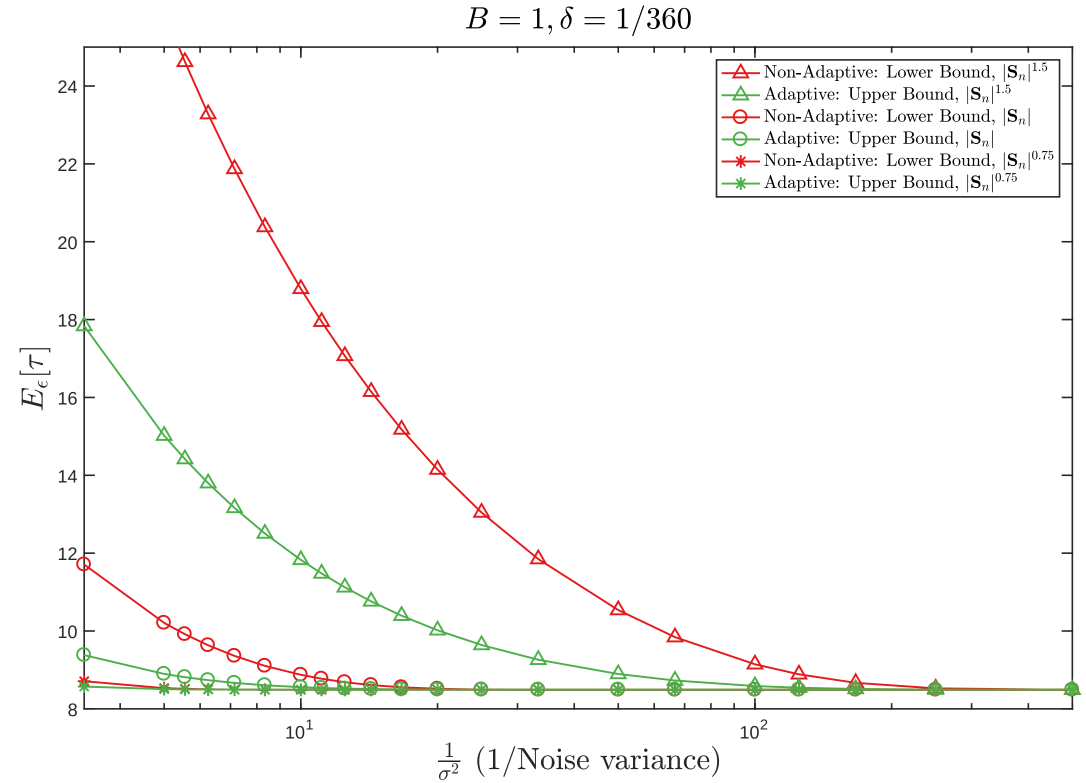

Generalizations

- General noise model: $Y(t) = A(t)\mathbf{W} + \hat{Z}$, $\hat{Z} = f(\mathbf{Z}, A(t))$
- Fixed (hierarchical) beam patterns

Generalizations

- General noise model: $Y(t) = A(t)\mathbf{W} + \hat{Z}$, $\hat{Z} = f(\mathbf{Z}, A(t))$
- Fixed (hierarchical) beam patterns

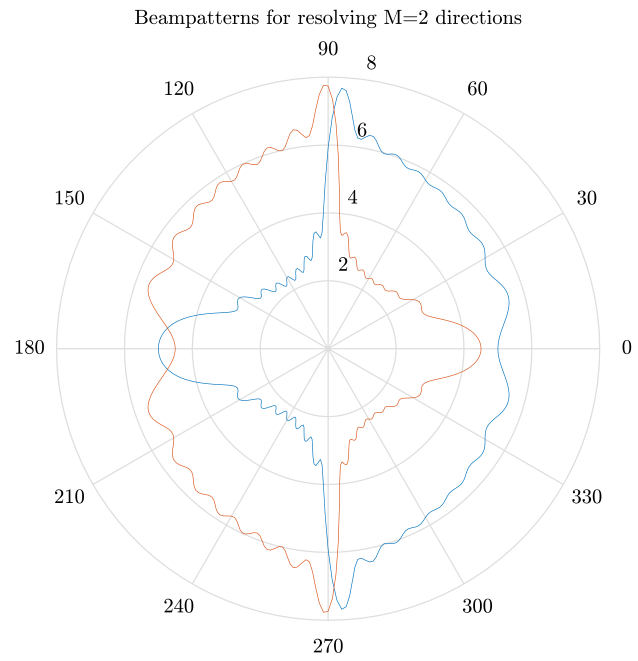

Generalizations

- ☐ General noise model: $Y(t) = A(t)\mathbf{W} + \hat{Z},\ \hat{Z} = f(\mathbf{Z}, A(t))$
- ☐ Fixed (hierarchical) beam patterns

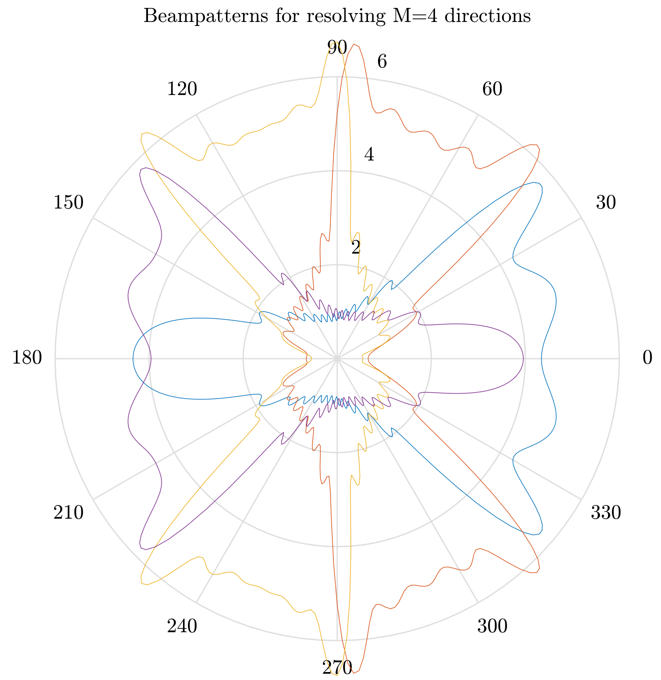

Generalizations

- ☐ General noise model: $Y(t) = A(t)\mathbf{W} + \hat{Z}$, $\hat{Z} = f(\mathbf{Z}, A(t))$

- ☐ Fixed (hierarchical) beam patterns

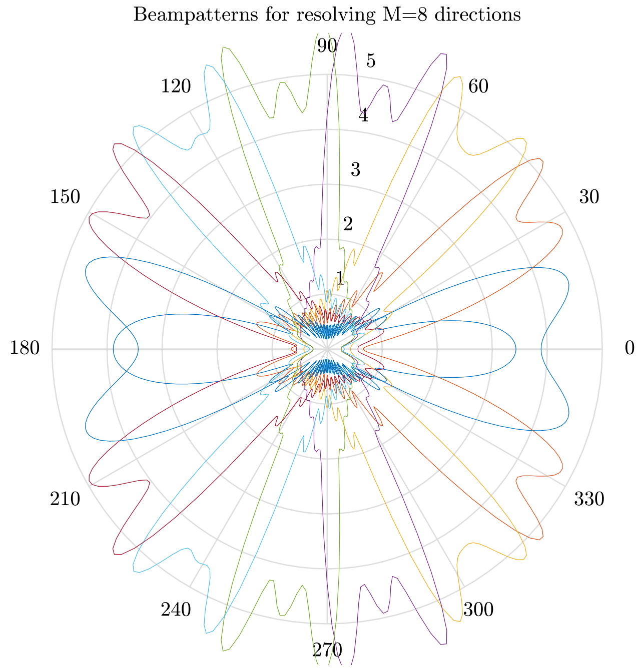

Generalizations and On-going Work

- □ Search for multiple target $(K > 1)$
- □ Dynamic case: $W(t)$
- □ Beyond Gaussian

Generalizations and On-going Work

- ☐ Search for multiple target $(K > 1)$
  - ✓ Noisy sequential group testing [Atia and Saligrama '12]; Mapped to an OR MAC [Kaspi, Shayevitz, J '15]
- ☐ Dynamic case: $W(t)$
- ☐ Beyond Gaussian

Generalizations and On-going Work

- ☐ Search for multiple target $(K > 1)$
  - ✓ Noisy sequential group testing [Atia and Saligrama '12]; Mapped to an OR MAC [Kaspi, Shayevitz, J '15]
  - ✓ Factor of $\frac{1}{K}$ in rate, where K bounds (is) the number of targets
- ☐ Dynamic case: $W(t)$
- ☐ Beyond Gaussian

Generalizations and On-going Work

- ☐ Search for multiple target $(K > 1)$
  - ✓ Noisy sequential group testing [Atia and Saligrama '12]; Mapped to an OR MAC [Kaspi, Shayevitz, J '15]
  - ✓ Factor of $\frac{1}{K}$ in rate, where K bounds (is) the number of targets
  - ‡ Case of an adder channel
- ☐ Dynamic case: $W(t)$
- ☐ Beyond Gaussian

- ☐ Search for multiple target ($K > 1$)
  - ✓ Noisy sequential group testing [Atia and Saligrama '12]; Mapped to an OR MAC [Kaspi, Shayevitz, J '15]
  - ✓ Factor of $\frac{1}{K}$ in rate, where K bounds (is) the number of targets
  - ‡ Case of an adder channel
- ☐ Dynamic case: $W(t)$
  - ✓ Results generalizes to unknown but constant speed (cut rate by half)
- ☐ Beyond Gaussian

Generalizations and On-going Work

- □ Search for multiple target ($K > 1$)
  - ✓ Noisy sequential group testing [Atia and Saligrama '12]; Mapped to an OR MAC [Kaspi, Shayevitz, J '15]
  - ✓ Factor of $\frac{1}{K}$ in rate, where K bounds (is) the number of targets
  - ‡ Case of an adder channel
- □ Dynamic case: $W(t)$
  - ✓ Results generalizes to unknown but constant speed (cut rate by half)
- □ Beyond Gaussian
  - ✓ Similar results for the binary symmetric noise (hard decoding) [Kaspi, Shayevitz, J '14]

Network performance function of network parameter $f : \mathcal{X} \to \mathbb{R}$.

Assumptions:

- ☐ $\mathcal{X}$ is the set of network parameters and protocols
- ☐ $f(x)$ is the network performance; $f(x_1)$ and $f(x_2)$ "correlated"
- ☐ f observed w noise: $y = f(x) + \eta(x)$, η non-presistent noise

Goal: Design a sequential strategy of selecting n query points $x_1, \ldots, x_n$ to identify a global optimizer of f.

- ☐ Performance measures:
  - Simple regret: $\mathcal{S}_n = f(x^*) - f(x_n^*)$
  - Cumulative regret: $\mathcal{R}_n = \sum_{t=1}^{n} f(x^*) - f(x_t)$

Network performance function of network parameter $f : \mathcal{X} \to \mathbb{R}$.

Assumptions:

- $\mathcal{X}$ is the set of network parameters and protocols
- $f(x)$ is the network performance; $f(x_1)$ and $f(x_2)$ "correlated"
- f observed w noise: $y = f(x) + \eta(x)$, η non-presistent noise

Goal: Design a sequential strategy of selecting n query points $x_1, \ldots, x_n$ to identify a global optimizer of f.

- Performance measures:
  - Simple regret: $\mathcal{S}_n = f(x^*) - f(x_n^*)$
  - Cumulative regret: $\mathcal{R}_n = \sum_{t=1}^{n} f(x^*) - f(x_t)$

Empirical Network Parameter tuning

Network performance function of network parameter $f : \mathcal{X} \to \mathbb{R}$.

Assumptions:

- $\mathcal{X}$ is the set of network parameters and protocols
- $f(x)$ is the network performance; $f(x_1)$ and $f(x_2)$ "correlated"
- f observed w noise: $y = f(x) + \eta(x)$, η non-presistent noise

<u>Goal:</u> Design a sequential strategy of selecting n query points $x_1, \dots, x_n$ to identify a global optimizer of f.

- Performance measures:
  - Simple regret: $\mathcal{S}_n = f(x^*) - f(x_n^*)$
  - Cumulative regret: $\mathcal{R}_n = \sum_{t=1}^{n} f(x^*) - f(x_t)$

Empirical Network Parameter tuning

Network performance function of network parameter $f : \mathcal{X} \to \mathbb{R}$.

Assumptio

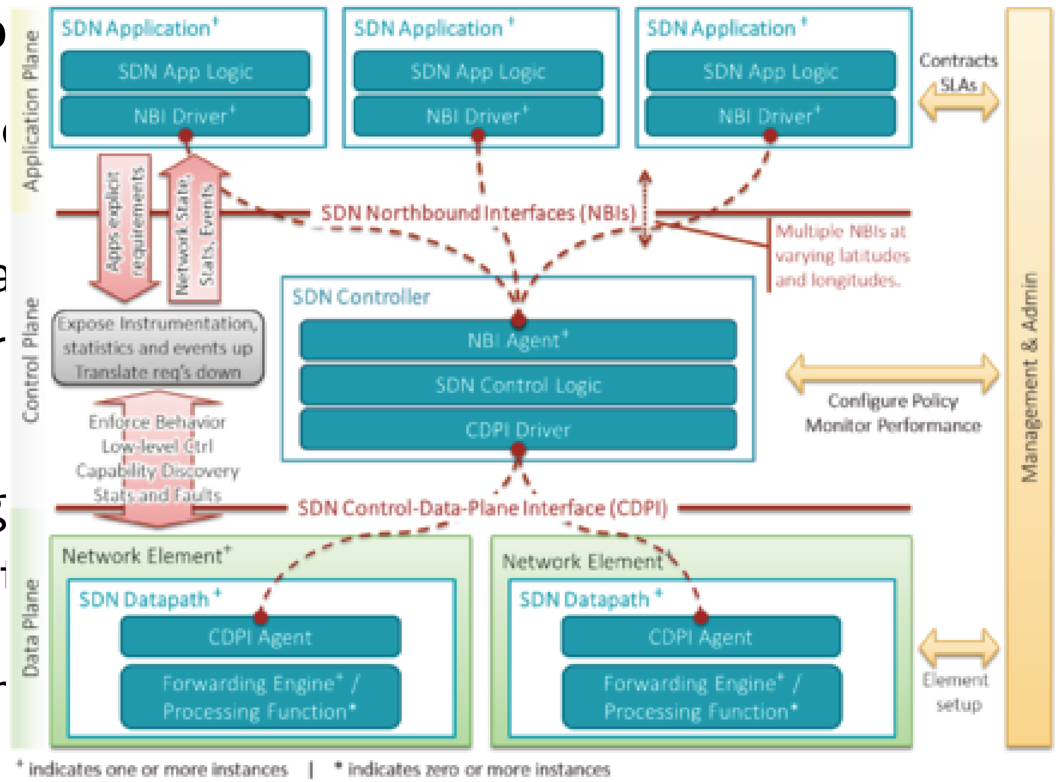

- □ $\mathcal{X}$ is th
- □ $f(x)$ is
  ”correla
- □ f obser ... tent noise

Goal: Desig ... points
$x_1, \ldots, x_n$

- □ Perform
  - Simple regret: $S_n = f(x^*) - f(x_n)$
  - Cumulative regret: $\mathcal{R}_n = \sum_{t=1}^{n} f(x^*) - f(x_t)$

Empirical Network Parameter tuning

Network performance function of network parameter $f : \mathcal{X} \to \mathbb{R}$.

Assumptions:

- $\mathcal{X}$ is the set of network parameters and protocols
- $f(x)$ is the network performance; $f(x_1)$ and $f(x_2)$ "correlated"
- f observed w noise: $y = f(x) + \eta(x)$, η non-presistent noise

Goal: Design a sequential strategy of selecting n query points $x_1, \ldots, x_n$ to identify a global optimizer of f.

- Performance measures:
  - Simple regret: $\mathcal{S}_n = f(x^*) - f(x_n^*)$
  - Cumulative regret: $\mathcal{R}_n = \sum_{t=1}^{n} f(x^*) - f(x_t)$ [bandit]

Questions?

Experiment Design: Single-shot

Design of Experiments [Blackwell '51]

- M mutually exclusive hypotheses: $H_i \Leftrightarrow \{\theta = i\}$, $i = 1, 2, \ldots, M$

- Prior $\boldsymbol{\rho}(0) = [\rho_1(0), \ldots, \rho_M(0)]$, $\rho_i(0) = P(\theta = i)$

Design of Experiments [Blackwell '51]

- M mutually exclusive hypotheses: $H_i \Leftrightarrow \{\theta = i\}$, $i = 1, 2, \ldots, M$
- Prior $\boldsymbol{\rho}(0) = [\rho_1(0), \ldots, \rho_M(0)]$, $\rho_i(0) = P(\theta = i)$
- Experiments $\mathcal{A}$ are available

Design of Experiments [Blackwell '51]

- M mutually exclusive hypotheses: $H_i \Leftrightarrow \{\theta = i\}$, $i = 1, 2, \ldots, M$

- Prior $\boldsymbol{\rho}(0) = [\rho_1(0), \ldots, \rho_M(0)]$, $\rho_i(0) = P(\theta = i)$

- Experiments $\mathcal{A}$ are available

- $Z|_{\{\theta=i, A=a\}} \sim q_i^a(\cdot)$: observation density given $a \in \mathcal{A}$ and H_i

Design of Experiments [Blackwell '51]

- M mutually exclusive hypotheses: $H_i \Leftrightarrow \{\theta = i\}$, $i = 1, 2, \ldots, M$
- Prior $\boldsymbol{\rho}(0) = [\rho_1(0), \ldots, \rho_M(0)]$, $\rho_i(0) = P(\theta = i)$
- Experiments $\mathcal{A}$ are available
- $Z|_{\{\theta=i, A=a\}} \sim q_i^a(\cdot)$: observation density given $a \in \mathcal{A}$ and H_i

Objective:
What is the best experiment $A = a$ to identify θ?

Learning from a Single Experiment

- Consider a single experiment $a \in \mathcal{A}$

Learning from a Single Experiment

- Consider a single experiment $a \in \mathcal{A}$
  - Prior $\theta \sim \boldsymbol{\rho}$
  - Noisy observations subject to $\{q_i^a(\cdot)\}_{i,a}$

Learning from a Single Experiment

- Consider a single experiment $a \in \mathcal{A}$
  - Prior $\theta \sim \boldsymbol{\rho}$
  - Noisy observations subject to $\{q_i^a(\cdot)\}_{i,a}$
- What should a be?

Learning from a Single Experiment

- Consider a single experiment $a \in \mathcal{A}$
  - Prior $\theta \sim \boldsymbol{\rho}$
  - Noisy observations subject to $\{q_i^a(\cdot)\}_{i,a}$
- What should a be? Compare experiment a with a'?

Learning from a Single Experiment

- Consider a single experiment $a \in \mathcal{A}$
  - Prior $\theta \sim \boldsymbol{\rho}$
  - Noisy observations subject to $\{q_i^a(\cdot)\}_{i,a}$
- What should a be? Compare experiment a with a'?
  - Stochastically degraded case [Blackwell '53], [Stein '53]

Learning from a Single Experiment

- Consider a single experiment $a \in \mathcal{A}$
  - Prior $\theta \sim \boldsymbol{\rho}$
  - Noisy observations subject to $\{q_i^a(\cdot)\}_{i,a}$
- What should a be? Compare experiment a with a'?
  - Stochastically degraded case [Blackwell '53], [Stein '53]
- Given experiment a:

Learning from a Single Experiment

- Consider a single experiment $a \in \mathcal{A}$
  - Prior $\theta \sim \boldsymbol{\rho}$
  - Noisy observations subject to $\{q_i^a(\cdot)\}_{i,a}$
- What should a be? Compare experiment a with a'?
  - Stochastically degraded case [Blackwell '53], [Stein '53]
- Given experiment a:
  - True hypothesis $\theta = i$ with probability ρ_i

Learning from a Single Experiment

- Consider a single experiment $a \in \mathcal{A}$
  - Prior $\theta \sim \boldsymbol{\rho}$
  - Noisy observations subject to $\{q_i^a(\cdot)\}_{i,a}$
- What should a be? Compare experiment a with a'?
  - Stochastically degraded case [Blackwell '53], [Stein '53]
- Given experiment a:
  - True hypothesis $\theta = i$ with probability ρ_i
  - Output distribution $Z^a \sim \sum_{i=1}^{M} \rho_i q_i^a(\cdot)$

Learning from a Single Experiment

- Consider a single experiment $a \in \mathcal{A}$
  - Prior $\theta \sim \boldsymbol{\rho}$
  - Noisy observations subject to $\{q_i^a(\cdot)\}_{i,a}$
- What should a be? Compare experiment a with a'?
  - Stochastically degraded case [Blackwell '53], [Stein '53]
- Given experiment a:
  - True hypothesis $\theta = i$ with probability ρ_i
  - Output distribution $Z^a \sim \sum_{i=1}^M \rho_i q_i^a(\cdot)$
  - Posterior upon observation $\theta | Z^a \sim \Phi^a\underbrace{(\boldsymbol{\rho}, Z^a}_{\text{Bayes operator}})$

Learning from a Single Experiment

- Consider a single experiment $a \in \mathcal{A}$
  - Prior $\theta \sim \boldsymbol{\rho}$
  - Noisy observations subject to $\{q_i^a(\cdot)\}_{i,a}$
- What should a be? Compare experiment a with a'?
  - Stochastically degraded case [Blackwell '53], [Stein '53]
- Given experiment a:
  - True hypothesis $\theta = i$ with probability ρ_i
  - Output distribution $Z^a \sim \sum_{i=1}^{M} \rho_i q_i^a(\cdot)$
  - Posterior upon observation $\theta | Z^a \sim \underbrace{\Phi^a(\boldsymbol{\rho}, Z^a)}_{\text{Bayes operator}}$
  - How does $\Phi^a(\cdot,\cdot)$ compare with $\Phi^{a'}(\cdot,\cdot)$

Divergence-based Selection

- Define a "symmetrized divergence" among $q_1^a, q_2^a, \ldots, q_M^a$

Divergence-based Selection

- Define a "symmetrized divergence" among $q_1^a, q_2^a, \ldots, q_M^a$
- Best action must maximize the divergence

Divergence-based Selection

- Define a "symmetrized divergence" among $q_1^a, q_2^a, \ldots, q_M^a$
- Best action must maximize the divergence
  - maximize discrimination among $H_1, H_2, \ldots, H_M$

Divergence-based Selection

- Define a "symmetrized divergence" among $q_1^a, q_2^a, \ldots, q_M^a$
- Best action must maximize the divergence
  - maximize discrimination among $H_1, H_2, \ldots, H_M$

Information Utility Heuristics:

- Measure of uncertainty V [DeGroot 1962]
- Information utility associated with V

$$\mathcal{IU}(a, \boldsymbol{\rho}, V) = V(\boldsymbol{\rho}) - \mathbb{E}[V(\Phi^a(\boldsymbol{\rho}, Z))]$$

Divergence-based Selection

- Define a "symmetrized divergence" among $q_1^a, q_2^a, \dots, q_M^a$
- Best action must maximize the divergence
  - maximize discrimination among $H_1, H_2, \dots, H_M$

Information Utility Heuristics:

- Measure of uncertainty V [DeGroot 1962]
- Information utility associated with V

$$\mathcal{IU}(a, \boldsymbol{\rho}, V) = V(\boldsymbol{\rho}) - \mathbb{E}[V(\underbrace{\Phi^a(\boldsymbol{\rho}, Z)}_{\text{Bayes operator}})]$$

Divergence-based Selection

- Define a "symmetrized divergence" among $q_1^a, q_2^a, \ldots, q_M^a$
- Best action must maximize the divergence
  - maximize discrimination among $H_1, H_2, \ldots, H_M$

Information Utility Heuristics:

- Measure of uncertainty V [DeGroot 1962]
- Information utility associated with V

$$\mathcal{IU}(a, \boldsymbol{\rho}, V) = V(\boldsymbol{\rho}) - \mathbb{E}[V(\Phi^a \underbrace{(\boldsymbol{\rho}, Z)}_{\text{Bayes operator}})]$$

- Most informative action $\arg\max_a \mathcal{IU}(a, \boldsymbol{\rho}, V)$

Divergence-based Selection

- Define a "symmetrized divergence" among $q_1^a, q_2^a, \ldots, q_M^a$
- Best action must maximize the divergence
  - maximize discrimination among $H_1, H_2, \ldots, H_M$

Information Utility Heuristics:

- Measure of uncertainty V [DeGroot 1962]
- Information utility associated with V

$$\mathcal{IU}(a, \boldsymbol{\rho}, V) = V(\boldsymbol{\rho}) - \mathbb{E}[V(\underbrace{\Phi^a(\boldsymbol{\rho}, Z)}_{\text{Bayes operator}})]$$

- Most informative action $\arg\max_a \mathcal{IU}(a, \boldsymbol{\rho}, V)$

Noisy search reduces to maximizing the $\mathcal{IU}(a, \boldsymbol{\rho}, V^*)$

Recall

- The *Entropy* of $p(\cdot)$ on space $\mathcal{Z}$:

$$H(p) = \sum_{\mathcal{Z}} p(z) \log \frac{1}{p(z)}$$

- The *Kullback-Leibler (KL) divergence* between $p(\cdot)$ and $q(\cdot)$:

$$D(p||q) = \sum_{\mathcal{Z}} p(z) \log \frac{p(z)}{q(z)}$$

Mutual Information as Information Utility

A widely used heuristic

$$\pi_I(\boldsymbol{\rho}) = \arg\max_a I(\theta; Z^a), \qquad \text{where } Z^a \sim q^a_{\boldsymbol{\rho}} = \sum_{i=1}^{M} \rho_i q_i^a$$

[Chaloner Verdinelli 1995], [Lindley 1956], [MacKay 1992], [Paninski 2005], [Branson 2010], [Butko Movellan 2009], [Fleuret 2004], [Williams et al. 2007]

A widely used heuristic

$$\pi_I(\boldsymbol{\rho}) = \arg\max_a I(\theta; Z^a), \qquad \text{where } Z^a \sim q^a_{\boldsymbol{\rho}} = \sum_{i=1}^{M} \rho_i q_i^a$$

$$I(\theta; Z^a) = H(\boldsymbol{\rho}) - \mathbb{E}(H(\Phi^a(\boldsymbol{\rho}, Z^a)))$$

$$= \mathcal{IU}(a, \boldsymbol{\rho}, H)$$

Mutual Information as Information Utility

A widely used heuristic

$$\pi_I(\boldsymbol{\rho}) = \arg\max_a I(\theta; Z^a), \qquad \text{where } Z^a \sim q^a_{\boldsymbol{\rho}} = \sum_{i=1}^{M} \rho_i q^a_i$$

$$I(\theta; Z^a) = H(\boldsymbol{\rho}) - \mathbb{E}(H(\Phi^a(\boldsymbol{\rho}, Z^a)))$$

$$= \mathcal{IU}(a, \boldsymbol{\rho}, H)$$

Also

$$I(\theta; Z^a) = \sum_{i=1}^{M} \rho_i D(q^a_i || q^a_{\boldsymbol{\rho}})$$

A widely used heuristic

$$\pi_I(\boldsymbol{\rho}) = \arg\max_a I(\theta; Z^a), \qquad \text{where } Z^a \sim q^a_{\boldsymbol{\rho}} = \sum_{i=1}^{M} \rho_i q_i^a$$

$$I(\theta; Z^a) = H(\boldsymbol{\rho}) - \mathbb{E}(H(\Phi^a(\boldsymbol{\rho}, Z^a)))$$

$$= \mathcal{IU}(a, \boldsymbol{\rho}, H)$$

Also

$$I(\theta; Z^a) = \textstyle\sum_{i=1}^{M} \rho_i D(q_i^a || q^a_{\boldsymbol{\rho}})$$

Jensen-Shannon divergence [Lin 1991]

Generalizing L divergence: $D_L(f,g) = \frac{1}{2} D(f || \frac{f+g}{2}) + \frac{1}{2} D(g || \frac{f+g}{2})$

Mutual Information as Information Utility

A widely used heuristic

$$\pi_I(\boldsymbol{\rho}) = \arg\max_a I(\theta; Z^a), \qquad \text{where } Z^a \sim q^a_{\boldsymbol{\rho}} = \sum_{i=1}^{M} \rho_i q^a_i$$

$$I(\theta; Z^a) = H(\boldsymbol{\rho}) - \mathbb{E}(H(\Phi^a(\boldsymbol{\rho}, Z^a)))$$

$$= \mathcal{IU}(a, \boldsymbol{\rho}, H)$$

Also

$$I(\theta; Z^a) = \sum_{i=1}^{M} \rho_i D(q^a_i || q^a_{\boldsymbol{\rho}})$$

A widely used heuristic

$$\pi_I(\boldsymbol{\rho}) = \arg\max_a I(\theta; Z^a), \qquad \text{where } Z^a \sim q^a_{\boldsymbol{\rho}} = \sum_{i=1}^{M} \rho_i q^a_i$$

$$I(\theta; Z^a) = H(\boldsymbol{\rho}) - \mathbb{E}(H(\Phi^a(\boldsymbol{\rho}, Z^a)))$$

$$= \mathcal{IU}(a, \boldsymbol{\rho}, H)$$

Also

$$I(\theta; Z^a) = \sum_{i=1}^{M} \rho_i D(q^a_i || q^a_{\boldsymbol{\rho}})$$

As $\rho_i \to 1$, $D(q^a_i||q^a_{\boldsymbol{\rho}}) \to D(q^a_i||q^a_i) = 0$ for any experiment a

A New Symmetrized Divergence

Extrinsic Jensen-Shannon Divergence [Naghshvar, J. ISIT'12]

Extrinsic Jensen-Shannon Divergence [Naghshvar, J. ISIT'12]

The *Extrinsic Jensen-Shannon (EJS) divergence* among densities $q_1, q_2, \ldots, q_M$ with respect to $\boldsymbol{\rho} = [\rho_1, \rho_2, \ldots, \rho_M]$ is defined as

$$EJS(\boldsymbol{\rho}; q_1, q_2, \ldots, q_M) = \sum_{i=1}^{M} \rho_i D(q_i || \sum_{k \neq i} \frac{\rho_k}{1-\rho_i} q_k).$$

A New Symmetrized Divergence

Extrinsic Jensen-Shannon Divergence [Naghshvar, J. ISIT'12]

The *Extrinsic Jensen-Shannon (EJS) divergence* among densities $q_1, q_2, \ldots, q_M$ with respect to $\boldsymbol{\rho} = [\rho_1, \rho_2, \ldots, \rho_M]$ is defined as

$$EJS(\boldsymbol{\rho}; q_1, q_2, \ldots, q_M) = \sum_{i=1}^{M} \rho_i D(q_i || \sum_{k \neq i} \frac{\rho_k}{1-\rho_i} q_k).$$

Bayesian generalization of J-divergence [Jefferys 73]

$$D_J(f,g) = \frac{1}{2} D(f||g) + \frac{1}{2} D(g||f)$$

Extrinsic Jensen-Shannon Divergence [Naghshvar, J. ISIT'12]

The *Extrinsic Jensen-Shannon (EJS) divergence* among densities $q_1, q_2, \ldots, q_M$ with respect to $\boldsymbol{\rho} = [\rho_1, \rho_2, \ldots, \rho_M]$ is defined as

$$EJS(\boldsymbol{\rho}; q_1, q_2, \ldots, q_M) = \sum_{i=1}^{M} \rho_i D(q_i || \sum_{k \neq i} \frac{\rho_k}{1-\rho_i} q_k).$$

Bayesian generalization of J-divergence [Jefferys 73]

$$D_J(f, g) = \frac{1}{2} D(f||g) + \frac{1}{2} D(g||f)$$

Proposition

EJS is the information utility associated with the average likelihood function $U(\boldsymbol{\rho}) = \sum_{i=1}^{M} \rho_i \log \frac{1-\rho_i}{\rho_i}$, i.e.

$$EJS(\boldsymbol{\rho}; q_1^a, \ldots, q_M^a) = \mathcal{IU}(a, \boldsymbol{\rho}, U)$$

An Upper Bound on Expected Number of Searches

Theorem (Naghshvar et. al. 13)**.** *Suppose there is* $C > 0$ *s.t. when* a *is selected according to SortPM and* $|a| \leq \alpha B/\delta$*, for all* $\boldsymbol{\rho}$*,* $EJS(\boldsymbol{\rho}, a) \geq C$*. Then*

$$\mathbb{E}[\tau^*] \leq \mathbb{E}[\tau_{SortPM}] \leq \frac{\log M + \max\{\log\log M, \log\frac{1}{\delta}\} + 4\Delta}{C} + K(\alpha).$$

Theorem (Naghshvar et. al. 13)**.** *Suppose there is $C > 0$ s.t. when a is selected according to SortPM and $|a| \leq \alpha B/\delta$, for all $\boldsymbol{\rho}$, $EJS(\boldsymbol{\rho}, a) \geq C$. Then*

$$\mathbb{E}[\tau^*] \leq \mathbb{E}[\tau_{SortPM}] \leq \frac{\log M + \max\{\log\log M, \log\frac{1}{\delta}\} + 4\Delta}{C} + K(\alpha).$$

Lemma. *Fix $\alpha \in (0,1)$. Using hard-decoded observation sequence $\Rightarrow C(\alpha) = 1 - h\left(Q\left(\left(\sigma^2\alpha B/\delta\right)^{-1/2}\right)\right)$.*