NEURAL INFORMATION PROCESSING SYSTEMS

# Symmetric Linear Bandits with Hidden Symmetry

Nam Phuong Tran[1], The-Anh Ta[2],

Debmalya Mandal[1], Long Tran-Thanh[1]

University of Warwick[1], Data 61 Australia[2]

# Linear Bandits

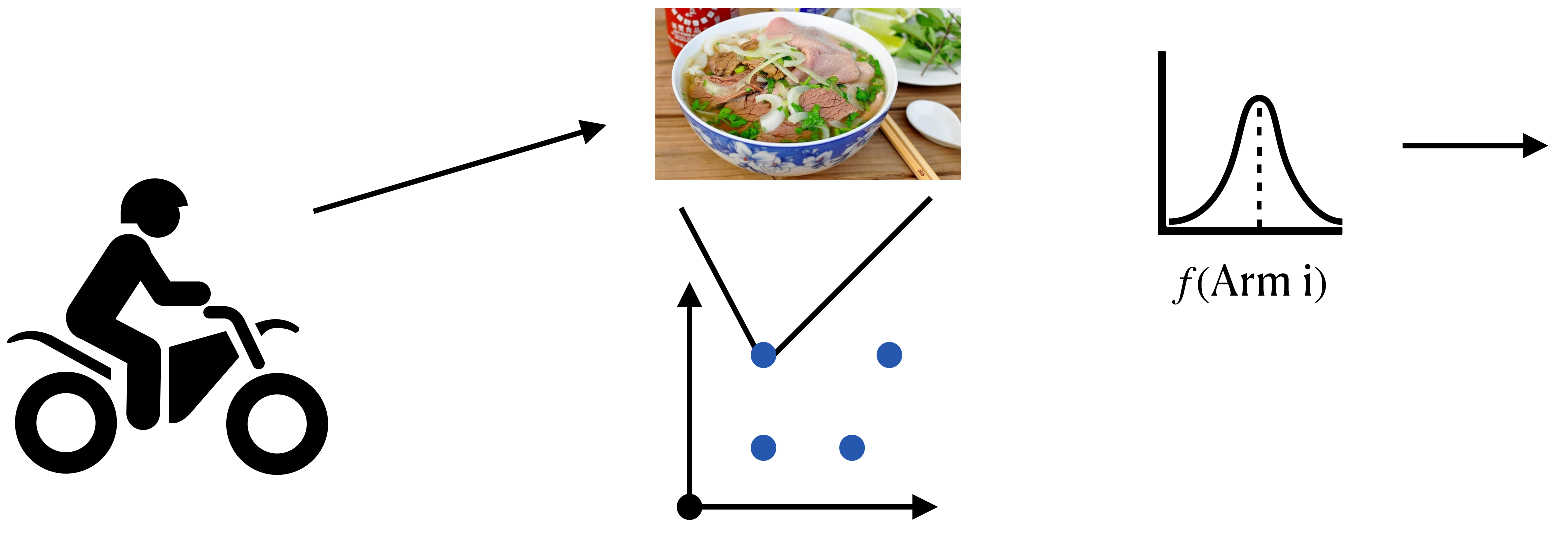


$$Y_t = f(X_t) + \eta_t$$

**Linear Bandit**: $y_t = f(x_t) + \eta_t, \quad f(x_t) = \left\langle x_t, \theta \right\rangle .\ \theta \in \mathbb{R}^d, \eta \sim N(0, \sigma).$

The goal is to minimise the expected regret $\mathbf{R}_T = \mathbb{E}\left[\sum_{t=1}^{T} f_\star - f(X_t)\right].$

Minimax regret: $\mathbf{R}_T = \Theta(d\sqrt{T})$

# Linear Bandits and The Curse of Dimensionality

**Linear Bandit**: $y_t = f(x_t) + \eta_t, \quad f(x_t) = \left\langle x_t, \theta \right\rangle.\ \theta \in \mathbb{R}^d, \eta \sim N(0, \sigma).$

- Minimax regret: $\mathbf{R}_T = \Theta(d\sqrt{T})$
- Curse of dimensionality: $d$ is large, vacuous regret.

**Inductive bias:** $\theta$ lies in low-dimensional structures.

- **Sparsity** $\|\theta\|_0 \leq s$**:** $\mathbf{R}_T = O(s \log(d)\sqrt{T})$
- Question: Is there any other low-dimensional structure that help overcome the curse of dimensionality?

# Hint from supervised learning

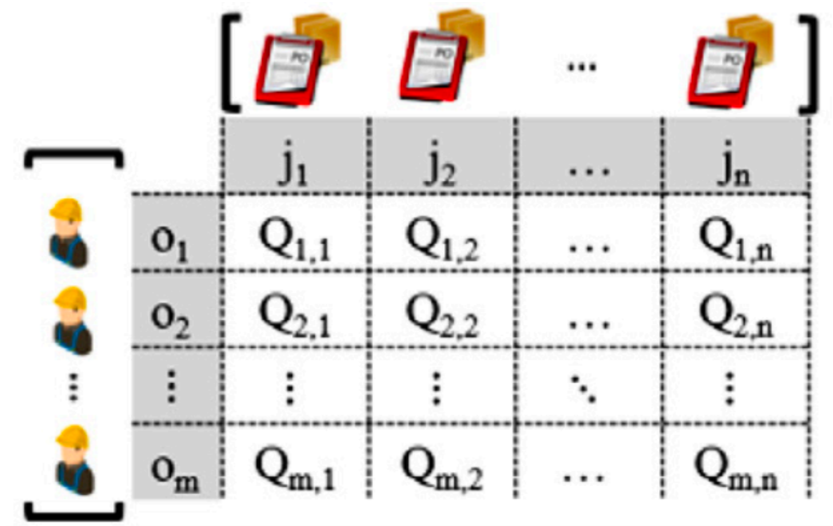


**What is symmetry?** $\underbrace{f(g \cdot x) = f(x)}_{\text{invariance}}$ .

**Break the curse of dimensionality** with *partial* knowledge of the symmetry structure?

**Optimal regret bound**: What is the optimal regret bound with the presence of symmetry.

# Symmetric Linear Bandit

**Linear Bandit**: $y_t = f(x_t) + \eta_t, \quad f(x_t) = \left\langle x_t, \theta \right\rangle$ .

- $x \in \mathcal{X} \subset \mathbb{R}^d; \theta \in \mathbb{R}^d$, $\eta \sim N(0, \sigma)$, where $d \geq T$ .
- **Symmetry**: $\mathcal{G} \leq \mathcal{S}_d$ be a subgroup of symmetric group of set $[d] = \{1,\ldots,d\}$.
- **Group action**: For $g \in \mathcal{G}$, let $g \cdot x$ permute the coordinate of $x$.
- **Invariant mean reward**: $\forall x \in \mathcal{X}, g \in \mathcal{G}; \ \langle x, \theta \rangle = \langle g \cdot x, \theta \rangle$.
- **Hidden symmetry**: $\mathcal{G}$ is only known to be a subgroup of $\mathcal{S}_d$.

# Impossibility Result

**Proposition 1 (Impossibility result)**: Assume, $\mathcal{X}$ is unit cube, and $f$ is invariant w.r.t. action of $\mathcal{G}$, s.t., $\dim(\mathrm{Fix}_{\mathcal{G}}) = 2$, any bandit algorithm suffer regret $\Omega(d\sqrt{T})$.

**Extra assumptions needed!!!**

**Assumption 2(Sub-exponential number of partitions).** The partition corresponding to $\mathcal{G}$ belongs to a small subclass of partitions $\mathcal{Q}_{d,\leq d_0} \subset \mathcal{P}_{d,\leq d_0}$. Moreover, for any $k \leq d_0$, there exist a constant $c > 0$ such that $|\mathcal{Q}_{d,k}| \leq O(d^{cd_0})$.

# Sub-exponential collection of models

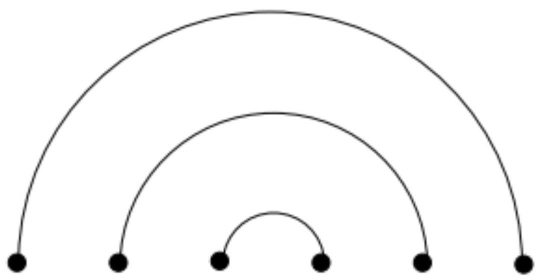

Non-crossing Partitions

Non-nesting Partitions

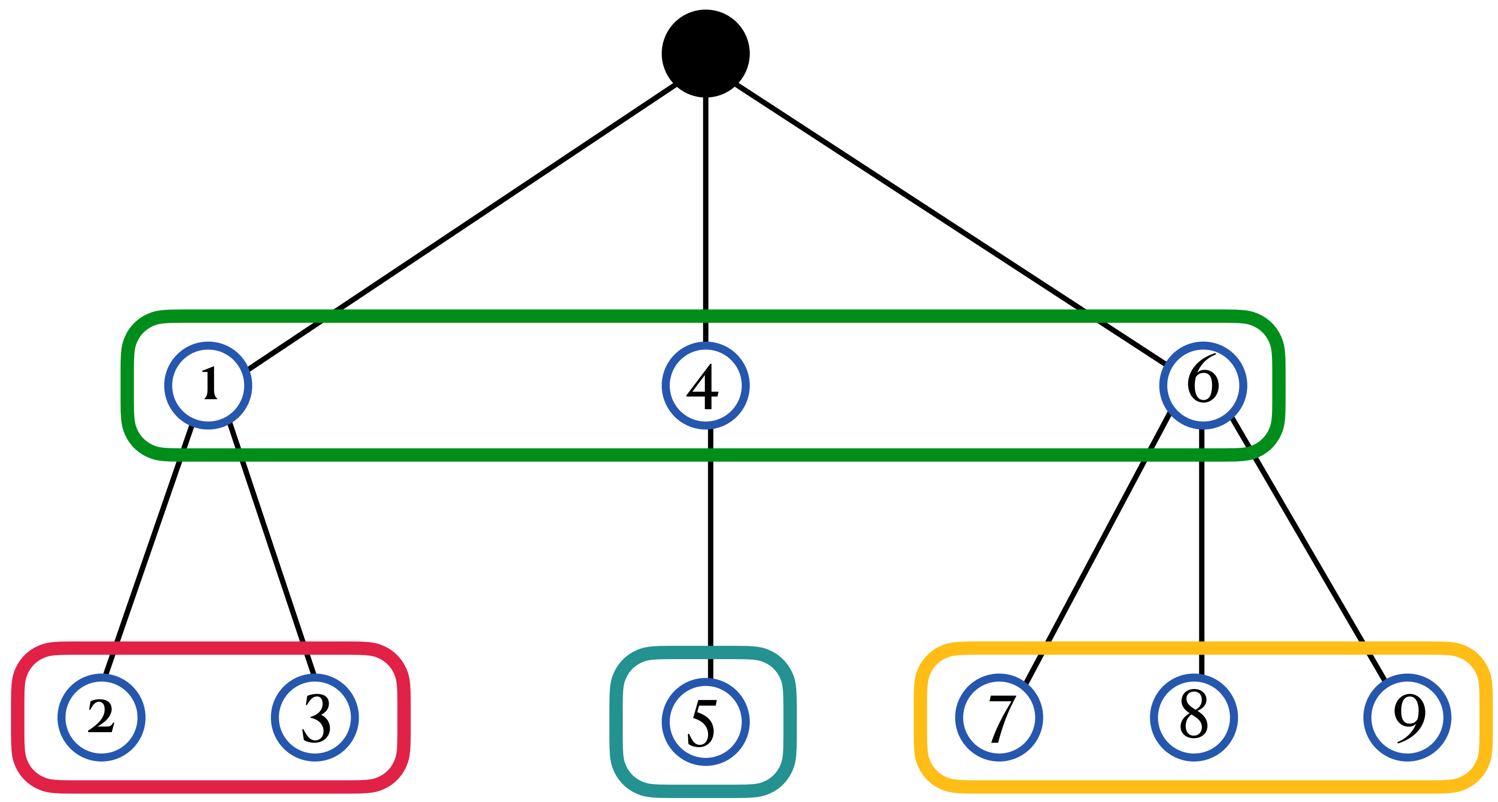


Partitions respect hierarchy

**Assumption 2(Sub-exponential number of partitions).** The partition corresponding to $\mathcal{G}$ belongs to a small subclass of partitions $\mathcal{Q}_{d,\leq d_0} \subset \mathcal{P}_{d,\leq d_0}$. Moreover, for any $k \leq d_0$, there exist a constant $c > 0$ such that $|\mathcal{Q}_{d,k}| \leq O(d^{cd_0})$.

# A Lower Bound of Symmetric Linear Bandits

*Recall the equivalence between interval partition and sparsity!*

**Regret Lower Bound:** There exist symmetric linear bandit instances in which Assumption 1 holds, such that any bandit algorithm must suffer regret

$$\mathbf{R}_T = \Omega\left(\min\left(C_{\min}^{-1/3}(\mathscr{X})d_0^{2/3}T^{2/3}, \sqrt{dT}\right)\right)$$

# Regret Upper Bound

**Theorem 7 (Regret upper bound):** with the choice of $t_1 \cong C_{\min}(\mathcal{X})^{-1/3} K_x^{1/3} d_0^{1/3} T^{2/3}$, the regret of Algorithm 1 is upper bounded as

$$\mathbf{R}_T = O\left(C_{\min}^{-1/3}(\mathcal{X}) d_0^{2/3} T^{2/3} \log(dT)^{1/3}\right)$$

**Theorem 1 (Regret upper bound):** With well-separateness assumption, Algorithm returns $\widehat{m} \ni \theta_\star$ with probability at least $1 - 1/T$, and it regret is upper bounded as

$$\mathbf{R}_T = O\left(d_0 \sqrt{T \log(d_0 T)}\right)$$

# Conclusion

- We show impossibility result $\Omega(d\sqrt{T})$ for general symmetry.

- With cardinality constraint, it possible to achieve regret $\tilde{\Theta}(d_0^{2/3}T^{2/3})$, and $\tilde{\Theta}(d_0\sqrt{T})$.