

Finite-armed bandits

Case (a): \mathcal{A}_t has always the same number of vectors in it:
“finite-armed stochastic contextual bandit”.

Case (b): Also, \mathcal{A}_t does not change, or $\mathcal{A}_t = \{\mathbf{a}_1, \dots, \mathbf{a}_K\}$:
“finite-armed stochastic linear bandit”.

Case (c): If the vectors in \mathcal{A}_t are also orthogonal to each other:
“finite-armed stochastic bandit”.

Difference between cases (c) and (b):

- Case (c): Learn about mean of arm $i \Leftrightarrow$ Choose action i ;
- Case (b): Learn about mean of arm $i \Leftrightarrow$ Choose action j s.t.
 $\langle \mathbf{x}_j, \mathbf{x}_i \rangle \neq 0$.