Clustering Items through Bandit Feedback[1]

Maximilian Graf* Victor Thuot † Nicolas Verzelen †

* Institut für Mathematik, Universität Potsdam, Potsdam, Germany

† INRAE, Mistea, Institut Agro, Univ Montpellier, Montpellier, France

June 3rd, MIPS week 2025, Montpellier

1. Maximilian Graf, Victor Thuot, and Nicolas Verzelen. *Clustering Items through Bandit Feedback: Finding the Right Feature out of Many*. Accepted at the 42nd International Conference on Machine Learning. 2025. arXiv: 2503.11209 [stat.ML].

Motivating example

- a set of forest patches
- a set of automatic biodiversity sensors

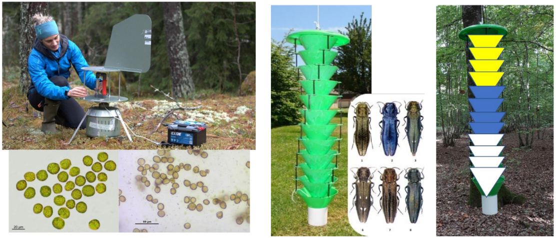

Figure: DNA sensor (left), optical sensor (right)[2]

- Objective: partition the forest patches by their biodiversity
- Limitations: cost, lack of specialists, unknown sensors

2. Christophe Bouget et al. "Bioc@pt : Capteurs automatiques de biodiversité en forêt". In: 2021.

Model

- n: number of forest patches
- d: number of sensors

Model

- n: number of forest patches
- d: number of sensors
- $M_{i,j}$: mean value of the j-th sensor on the i-th patch

Model

- n: number of forest patches
- d: number of sensors
- $M_{i,j}$: mean value of the j-th sensor on the i-th patch

$$M = \begin{bmatrix} M_{1,1} & \cdots & M_{1,j} & \cdots & M_{1,d} \\ \vdots & & \vdots & & \vdots \\ M_{i,1} & \cdots & M_{i,j} & \cdots & M_{i,d} \\ \vdots & & \vdots & & \vdots \\ M_{n,1} & \cdots & M_{n,j} & \cdots & M_{n,d} \end{bmatrix} \leftarrow M_{i,\cdot}$$

Model

- n: number of forest patches (items)
- d: number of sensors (features)
- $M_{i,j}$: mean value of the j-th sensor on the i-th patch

$$M = \begin{bmatrix} M_{1,1} & \cdots & M_{1,j} & \cdots & M_{1,d} \\ \vdots & & \vdots & & \vdots \\ M_{i,1} & \cdots & M_{i,j} & \cdots & M_{i,d} \\ \vdots & & \vdots & & \vdots \\ M_{n,1} & \cdots & M_{n,j} & \cdots & M_{n,d} \end{bmatrix} \leftarrow M_{i,\cdot}$$

Model

- n: number of forest patches (items)
- d: number of sensors (features)
- $M_{i,j}$: mean value of the j-th sensor on the i-th patch

$$M = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ \color{red}{0} & \color{red}{1} & \color{red}{0} & \color{red}{0.5} & \color{red}{0.05} & \color{red}{0} \\ 0 & 0 & 0 & 0 & 0 & 0 \\ \color{red}{0} & \color{red}{1} & \color{red}{0} & \color{red}{0.5} & \color{red}{0.05} & \color{red}{0} \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \leftarrow M_{i,\cdot} \in \{0, \color{red}{\Delta}\}$$

Model

- n: number of forest patches (items)
- d: number of sensors (features)
- $M_{i,j}$: mean value of the j-th sensor on the i-th patch

$$M = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0.5 & 0.05 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0.5 & 0.05 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \leftarrow M_{i,\cdot} \in \{0, \Delta\}$$

Main objective

Partition the patches into two groups of similar biodiversity

Model

- n: number of forest patches (items)
- d: number of sensors (features)
- $M_{i,j}$: mean value of the j-th sensor on the i-th patch

$$M = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0.5 & 0.05 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0.5 & 0.05 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \leftarrow M_{i,\cdot} \in \{0, \Delta\}$$

Main objective

Partition the patches into two groups of similar biodiversity

- Vanilla clustering problem: observe the entire matrix

Model

- n: number of forest patches (items)
- d: number of sensors (features)
- $M_{i,j}$: mean value of the j-th sensor on the i-th patch

$$M = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0.5 & 0.05 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0.5 & 0.05 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \leftarrow M_{i,\cdot} \in \{0, \Delta\}$$

Main objective

Partition the patches into two groups of similar biodiversity

- Vanilla clustering problem: observe the entire matrix
- Bandit clustering problem: construct the sampling protocol on the fly

Sampling protocol

Learning protocol

At each time step t,

- choose a patch I (*based on the past*)
- choose a sensor J (*based on the past*)
- observe $X_t = M_{I,J} +$ noise,

Sampling protocol

Learning protocol

At each time step t,

- choose a patch I (*based on the past*)
- choose a sensor J (*based on the past*)
- observe $X_t = M_{I,J} + \text{noise}$,

- **Objective 1**: recover the partition of the patches
- **Objective 2**: sequentially build the sampling protocol
  focus the budget on the most informative sensor

Sampling protocol (example)

$$
\begin{array}{c}
\downarrow \\
\begin{bmatrix}
\boxed{0} & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & 0 & 0.5 & 0.1 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & 0 & 0.5 & 0.1 & 0 \\
0 & 0 & 0 & 0 & 0 & 0
\end{bmatrix}
\end{array}
\begin{array}{l}
\leftarrow \\ \\ \\ \\ \\
\end{array}
$$

Time t	1	2	3	4
(patch,sensor)	(1,1)			
$X_t = \boxed{} +$ noise				

Sampling protocol (example)

$$\begin{bmatrix} \boxed{0} & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0.5 & 0.1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0.5 & 0.1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

Time t	1	2	3	4
(patch,sensor)	(1,1)			
$X_t = \boxed{} +$ noise	0.1			

Sampling protocol (example)

$$
\begin{array}{cc}
\begin{array}{cccccc} \downarrow & & & & & \end{array} & \\
\begin{bmatrix}
0 & 0 & 0 & 0 & 0 & 0 \\
\boxed{0} & 1 & 0 & 0.5 & 0.1 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & 0 & 0.5 & 0.1 & 0 \\
0 & 0 & 0 & 0 & 0 & 0
\end{bmatrix} & \begin{array}{c} \\ \leftarrow \\ \\ \\ \\ \end{array}
\end{array}
$$

Time t	1	2	3	4
(patch,sensor)	(1,1)	(2,1)		
$X_t = \boxed{} + \text{noise}$	0.1			

Sampling protocol (example)

$$
\begin{array}{cc}
\begin{array}{cccccc} \downarrow & & & & & \end{array} & \\
\begin{bmatrix}
0 & 0 & 0 & 0 & 0 & 0 \\
\boxed{0} & 1 & 0 & 0.5 & 0.1 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & 0 & 0.5 & 0.1 & 0 \\
0 & 0 & 0 & 0 & 0 & 0
\end{bmatrix} & \begin{array}{c} \\ \leftarrow \\ \\ \\ \\ \end{array}
\end{array}
$$

Time t	1	2	3	4
(patch,sensor)	(1,1)	(2,1)		
$X_t = \boxed{} + \text{noise}$	0.1	-0.05		

Sampling protocol (example)

$$
\begin{array}{cc}
\begin{array}{cccccc} & & & \downarrow & & \end{array} & \\
\begin{bmatrix}
0 & 0 & 0 & \boxed{0} & 0 & 0 \\
0 & 1 & 0 & 0.5 & 0.1 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & 0 & 0.5 & 0.1 & 0 \\
0 & 0 & 0 & 0 & 0 & 0
\end{bmatrix} & \begin{array}{c} \leftarrow \\ \\ \\ \\ \\ \end{array}
\end{array}
$$

Time t	1	2	3	4
(patch,sensor)	(1,1)	(2,1)	(1,4)	
$X_t = \boxed{} +$ noise	0.1	-0.05		

Sampling protocol (example)

$$
\begin{array}{cc}
\begin{array}{cccccc} & & & \downarrow & & \end{array} & \\
\begin{bmatrix}
0 & 0 & 0 & \boxed{0} & 0 & 0 \\
0 & 1 & 0 & 0.5 & 0.1 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & 0 & 0.5 & 0.1 & 0 \\
0 & 0 & 0 & 0 & 0 & 0
\end{bmatrix} & \begin{array}{c} \leftarrow \\ \\ \\ \\ \\ \end{array}
\end{array}
$$

Time t	1	2	3	4
(patch,sensor)	(1,1)	(2,1)	(1,4)	
$X_t = \boxed{} + \text{noise}$	0.1	-0.05	0.1	

Sampling protocol (example)

$$
\begin{array}{cc}
\begin{array}{cccccc} & & & \downarrow & & \end{array} & \\
\begin{bmatrix}
0 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & 0 & 0.5 & 0.1 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & 0 & \boxed{0.5} & 0.1 & 0 \\
0 & 0 & 0 & 0 & 0 & 0
\end{bmatrix} & \begin{array}{c} \\ \\ \\ \leftarrow \\ \\ \end{array}
\end{array}
$$

Time t	1	2	3	4
(patch,sensor)	(1,1)	(2,1)	(1,4)	(4,4)
$X_t = \boxed{} + \text{noise}$	0.1	-0.05	0.1	

Sampling protocol (example)

$$
\begin{array}{cc}
\begin{array}{cccccc} & & & \downarrow & & \end{array} & \\
\begin{bmatrix}
0 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & 0 & 0.5 & 0.1 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & 0 & \boxed{0.5} & 0.1 & 0 \\
0 & 0 & 0 & 0 & 0 & 0
\end{bmatrix} & \begin{array}{c} \\ \\ \\ \leftarrow \\ \\ \end{array}
\end{array}
$$

Time t	1	2	3	4
(patch,sensor)	(1,1)	(2,1)	(1,4)	(4,4)
$X_t = \boxed{} + \text{noise}$	0.1	-0.05	0.1	0.4

Sampling protocol (example)

$$\begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0.5 & 0.1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0.5 & 0.1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

Time t	1	2	3	4	... until T
(patch,sensor)	(1,1)	(2,1)	(1,4)	(4,4)	
$X_t = \boxed{} + \text{noise}$	0.1	-0.05	0.1	0.4	

$\rightarrow$ At time T, output a partition of the patches

Results

1. **Algorithmic solution:** new online clustering algorithms
   - identification of the most discriminative sensor
   - classification of the patches

Results

1. **Algorithmic solution:** new online clustering algorithms
   - identification of the most discriminative sensor
   - classification of the patches
2. **Optimality:** provably optimal sampling complexity
   - bound on the budget of our algorithm
   - impossibility result for the lower bound

Results

1. **Algorithmic solution:** new online clustering algorithms
   - identification of the most discriminative sensor
   - classification of the patches
2. **Optimality:** provably optimal sampling complexity
   - bound on the budget of our algorithm
   - impossibility result for the lower bound
3. **Reference:**

Maximilian Graf, Victor Thuot, and Nicolas Verzelen. *Clustering Items through Bandit Feedback: Finding the Right Feature out of Many*. Accepted at the 42nd International Conference on Machine Learning. 2025. arXiv: 2503.11209 [stat.ML]