

Bandit Models for Recommendations

Maryam Tavakol


Computer Science 8, Artificial Intelligence

July 2019

Recommender Systems

Recommender Systems

- e-commerce



by Sony
★★★★☆ 508 customer reviews | 319 answered questions

Suggested price: €379.00
Our Price: **€280.00** FREE Delivery. Delivery Details
You Save: €99.00 (26%)
Prices for items sold by Amazon are inclusive of German VAT.

9 new from €280.00 17 used from €260.40

Colour: **Schwarz**

€280.00 €288.00

Style Name: **Kopfhörer**

- Aktives Noise Cancelling: Optimierung der Geräuschminimierung an Bord
- Quick Attention Modus: blitzschnell die Umgebung durch intuitives Handling
- Kompatibel mit Amazon Alexa. Perfekte Passform im Around-Ear Design
- Optimierbare Klangeinstellungen mit der Headphones Connect App
- Lange Akkulaufzeit: bis zu 30 Std. abhängig vom verwendeten Modus.

Gerät kaufen und 3 Monate unbegrenzt Musik etc.

Customers who viewed this item also viewed



←

- Sony Wireless High Resolution Headphones (Noise Cancelling, Bluetooth, NFC,...**
★★★★☆ 347
€356.39 ✓prime
- Sony WF-1000XM3 True Wireless Noise Cancelling Kopfhörer (bis zu 32h Akkulaufzeit, stabile Bluetooth Verbindung,...**
€249.00 ✓prime
- Sony MDR-1000X Wireless High Resolution Headphones Noise Cancelling, Sense...**
★★★★☆ 308
€319.90
- Bose 700 Headphones**
€399.00 ✓prime

Recommender Systems

- e-commerce
- entertainment

The screenshot shows an Amazon product page for Sony headphones. At the top, the product name "by Sony" is followed by a star rating and "508 customer reviews | 319 answered questions". The price information includes "Suggested price: €579.00", "Our Price: €280.00 FREE Delivery. Delivery Details", and "You Save: €99.00 (26%)". Below this, it states "Prices for items sold by Amazon are inclusive of German tax" and "9 new from €280.00 | 17 used from €260.40". The color is listed as "Schwarz".

A prominent recommendation banner is displayed, titled "Because you watched Black Mirror". It features three Netflix titles: "BLACK MIRROR BANDERSNATCH", "3%", and "THE SOCIETY".

Below the banner, four headphones are listed with their respective details:

Product Name	Price	Rating
Sony Wireless High Resolution Headphones (Noise Cancelling, Bluetooth, NFC,...	€356.39 ✓prime	★★★★☆ 347
Sony WF-1000XM3 True Wireless Noise Cancelling Kopfhörer (bis zu 32h Akkulaufzeit, stabile Bluetooth Verbindung,...	€249.00 ✓prime	★★★★☆ 308
Sony MDR-1000X Wireless High Resolution Headphones Noise Cancelling, Sense...	€319.90	★★★★☆ 308
Bose 700 Headphones	€399.00 ✓prime	

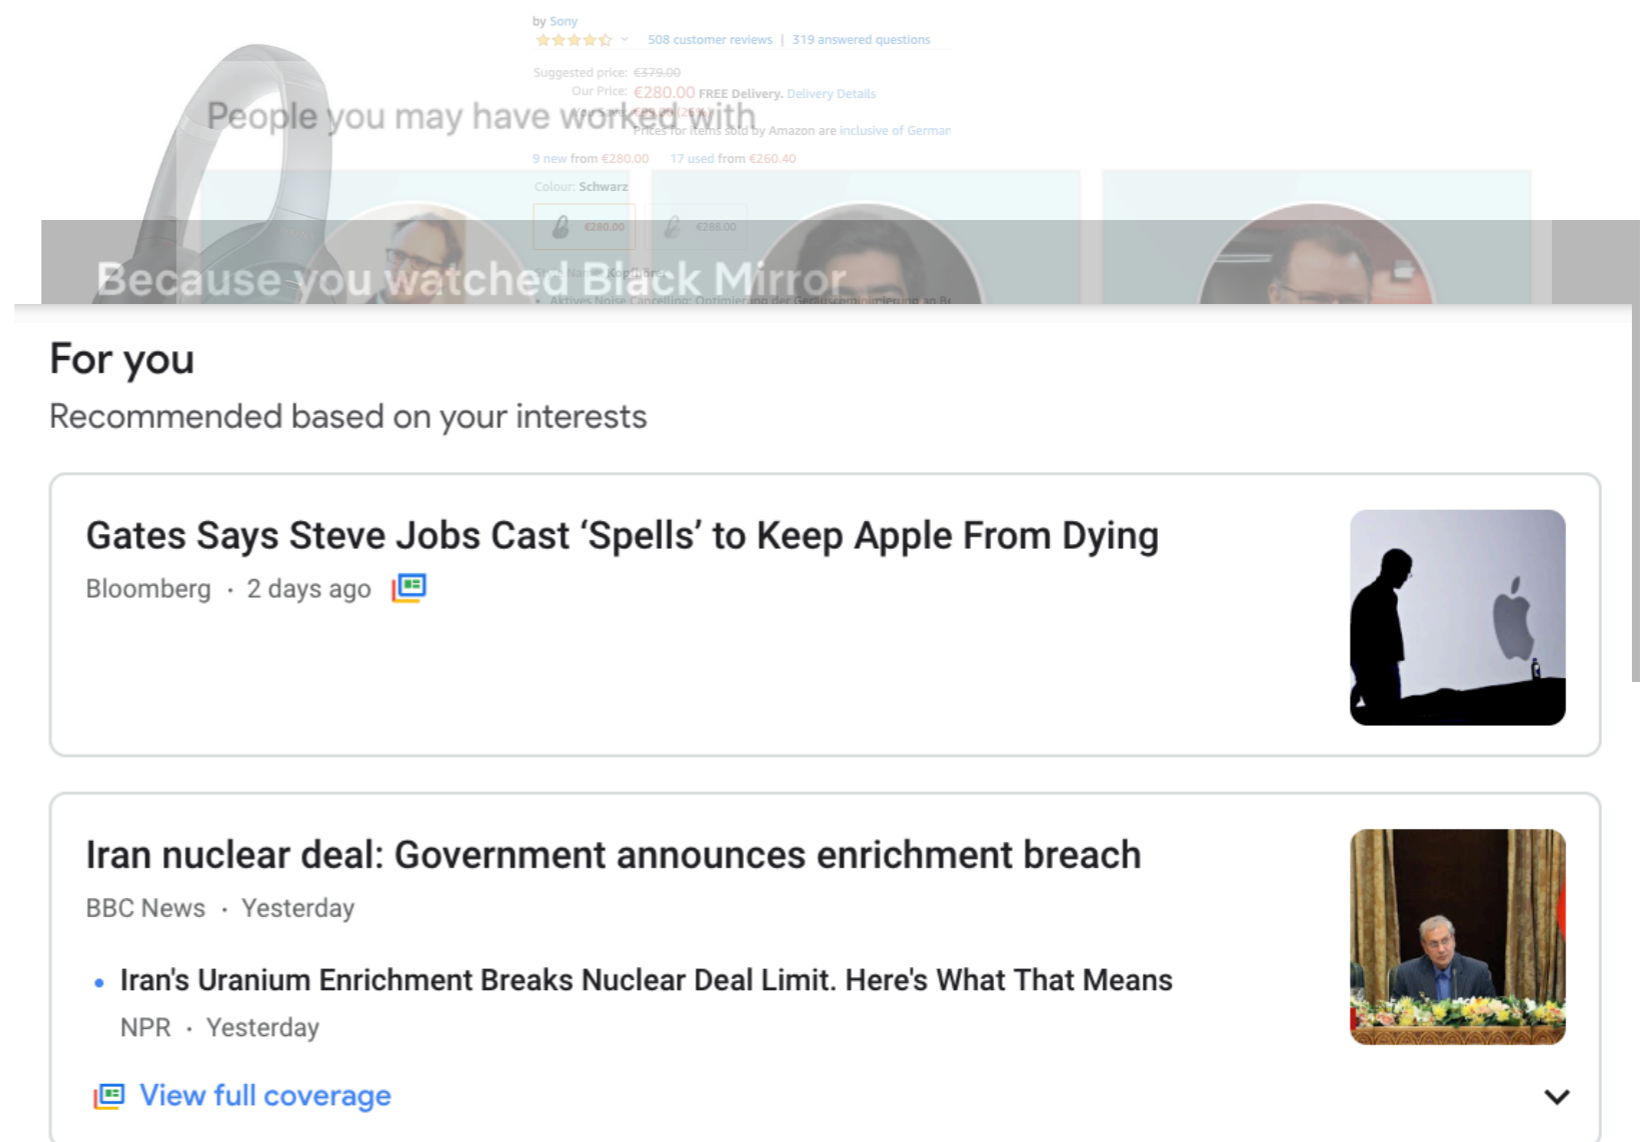
Recommender Systems

- e-commerce
- entertainment
- social networks

The screenshot displays a social network interface. At the top, it says 'by Sony' with a 4.5-star rating, '508 customer reviews', and '319 answered questions'. Below this is the section 'People you may have worked with', which features three profile cards. Each card includes a circular profile picture, a redacted name, a job title, and a 'Connect' button. The first card is for a 'Co-founder & CEO at Lifem' associated with 'Criteo'. The second is for a 'PhD Student at Joseph L. Rotman School of...' with '27 mutual connections'. The third is for a 'Professor at University of Tehran & Adjunct...' with '28 mutual connections'. Below the profiles are four product recommendations: 'Sony Wireless High Resolution Headphones' (€356.39), 'Sony WF-1000XM3 True Wireless Noise Cancelling Kopfhörer' (€249.00), 'Sony MDR-1000X Wireless High Resolution Headphones' (€319.90), and 'Bose 700 Headphones' (€399.00).

Recommender Systems

- e-commerce
- entertainment
- social networks
- online content




by Sony
★★★★☆ 508 customer reviews | 319 answered questions
Suggested price: €579.00
Our Price: €280.00 FREE Delivery. Delivery Details
Prices for items sold by Amazon are inclusive of German VAT.


9 new from €280.00 17 used from €260.40
Colour: Schwarz
€280.00 €280.00

People you may have worked with

Because you watched Black Mirror


For you
Recommended based on your interests


Gates Says Steve Jobs Cast 'Spells' to Keep Apple From Dying
Bloomberg · 2 days ago 



Iran nuclear deal: Government announces enrichment breach
BBC News · Yesterday

- **Iran's Uranium Enrichment Breaks Nuclear Deal Limit. Here's What That Means**
NPR · Yesterday

 [View full coverage](#)



Multi-armed Bandits

Multi-armed Bandits

- Which one to play?



Multi-armed Bandits

- Which one to play?



10\$



0\$



-5\$

...



1\$

Multi-armed Bandits

- Which one to play?



10\$

-8\$



0\$

4\$



-5\$

0\$

...



1\$

1\$

Multi-armed Bandits

- Which one to play?



- strategy: to better balance exploration vs. exploitation

Bandits in Recommendation

Bandits in Recommendation

- A user views the following item:



Bandits in Recommendation

- A user views the following item:



- Now what to recommend?



...



Bandits in Recommendation

- A user views the following item:



- Now what to recommend?



...



Bandits in Recommendation

- A user views the following item:



- Now what to recommend?



what about the context?

Context	Actions		
	1	2	3
(Male, 50, ...)	1.0	0.2	0.0
(Female, 18, ...)	1.0	0.0	1.0
(Female, 48, ...)	0.5	0.1	0.7
⋮		⋮	

Contextual Bandits

- In contextual bandit, the system receives an additional information (=context) at every step
 - the reward is represented as a function of context
- sequential procedure:

- repeat:
 1. learner presented with **context**
 2. learner chooses an **action**
 3. learner observes **reward** (but **only** for chosen action)
- **goal**: learn to choose actions to maximize rewards

LinUCB for News Recommendation

LinUCB for News Recommendation

- The payoff function is defined by a linear model

$$\mathbb{E}[r_{t,a} | \mathbf{x}_t] = \mathbf{x}_t^\top \mathbf{w}$$

LinUCB for News Recommendation

- The payoff function is defined by a linear model

$$\mathbb{E}[r_{t,a} | \mathbf{x}_t] = \mathbf{x}_t^\top \mathbf{w}$$

to be learned

LinUCB for News Recommendation

- The payoff function is defined by a linear model

$$\mathbb{E}[r_{t,a} | \mathbf{x}_t] = \mathbf{x}_t^\top \mathbf{w}$$

to be learned

- arm selection strategy:
 - the arm with the highest **mean reward** + **confidence interval**

LinUCB for News Recommendation

- The payoff function is defined by a linear model

$$\mathbb{E}[r_{t,a} | \mathbf{x}_t] = \mathbf{x}_t^\top \mathbf{w}$$

to be learned

- arm selection strategy:
 - the arm with the highest **mean reward** + **confidence interval**
 exploitation exploration

LinUCB for News Recommendation

- The payoff function is defined by a linear model

$$\mathbb{E}[r_{t,a} | \mathbf{x}_t] = \mathbf{x}_t^\top \mathbf{w}$$

to be learned

- arm selection strategy:
 - the arm with the highest **mean reward** + **confidence interval**

exploitation

exploration

Upper Confidence Bound (UCB)

Thanks for Your Attention

- Read more...
- contextual bandits in recommendation:
 - L. Li, W. Chu, J. Langford, and R. E. Schapire. A contextual-bandit approach to personalized news article recommendation. International Conference on the World Wide Web, 2010
 - M. Tavakol and U. Brefeld. A unified contextual bandit framework for long-and short- term recommendations. In Proceedings of the European Conference on Machine Learning, 2017.
- off-line evaluation of recommendations using bandits:
 - Y. Su, L. Wang, M. Santacatterina, and T. Joachims. CAB: Continuous Adaptive Blending for Policy Evaluation and Learning. International Conference on Machine Learning, 2019.

Questions?