

ŠPLETNO OGLAŠEVANJE

web advertising

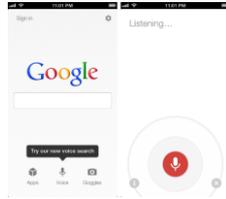
doc. dr. Matej Guid

Fakulteta za računalništvo in informatiko
Univerza v Ljubljani

marec 2022

ARTIFICIAL INTELLIGENCE MARKETING

- Smarter Search
- Ad Optimization
- Sentiment Analysis
- Customer Service
- AI in Images



ARTIFICIAL INTELLIGENCE & PREDICTIVE MARKETING

- reach customers on **the right device, channel at the right time**
- cross-device advertising and campaigning, based on patterns in **customer behavior**
- master the **customer journey** instead of marketing funnel
- understand **critical campaign indicators** and make data-driven decisions
- generate/select targeted marketing messages to **trigger an emotional response**



PROBLEM DESCRIPTION

- social media platforms optimize their own goals (not necessarily aligned with advertisers' goals!)
- micromanaging campaigns demands **a lot of human involvement** and strong analytical skills
- huge amount of employees' cumulative time spent (across 20+ countries)



Ad group name	Ads	Budget	Frequency	Reach	Impressions #	Clicks #	CTR #	#CPS #	Spent #	Offline conversion #	CPOC #
facebook 18-24, Desktop feed 23.02.2018 - 15.04.2018, completed	9	LIFETIME: EUR 80.01	2.41	126,356	304,869	13,383	4.3896 %	EUR 0.00	EUR 54.33	199	EUR 0.27
facebook 25-34, Desktop feed 23.02.2018 - 15.04.2018, completed	9	LIFETIME: EUR 139.55	1.73	53,363	92,708	3,253	3.5009 %	EUR 0.03	EUR 106.62	461	EUR 0.23
facebook 35-44, Desktop feed 23.02.2018 - 15.04.2018, completed	9	LIFETIME: EUR 152.98	1.66	31,052	1,391	2,6967 %	EUR 0.09	EUR 119.40	398	EUR 0.30	EUR 0.23
facebook 45-54, Desktop feed 23.02.2018 - 15.04.2018, completed	9	LIFETIME: EUR 177.70	1.73	20,617	35,756	1,306	3.6525 %	EUR 0.11	EUR 139.99	527	EUR 0.27
facebook 55+, Desktop feed 23.02.2018 - 15.04.2018, completed	9	LIFETIME: EUR 239.47	1.83	15,039	27,619	1,407	5.0943 %	EUR 0.13	EUR 181.39	609	EUR 0.30
facebook 18-24, Mobile feed 23.02.2018 - 15.04.2018, completed	9	LIFETIME: EUR 118.23	1.35	100,430	136,247	3,674	2.6956 %	EUR 0.02	EUR 88.36	365	EUR 0.24
facebook 25-34, Mobile feed 23.02.2018 - 15.04.2018, completed	9	LIFETIME: EUR 205.34	1.27	72,876	93,148	2,555	2.7429 %	EUR 0.06	EUR 164.63	852	EUR 0.19
facebook 35-44, Mobile feed 23.02.2018 - 15.04.2018, completed	9	LIFETIME: EUR 328.71	1.30								
facebook 45-54, Mobile feed 23.02.2018 - 15.04.2018, completed	9	LIFETIME: EUR 375.83	1.28								
facebook 55+, Mobile feed 23.02.2018 - 15.04.2018, completed	9	LIFETIME: EUR 342.74	1.23								



facebook 18-24, Audience network 23.02.2018 - 15.04.2018, completed	12										
facebook 25-34, Audience network 23.02.2018 - 15.04.2018, completed	9										
facebook 35-44, Audience network 23.02.2018 - 15.04.2018, completed	9										
facebook 45-54, Audience network 23.02.2018 - 15.04.2018, completed	9	LIFETIME: EUR 373.97	1.35	41,193	55,714	2,093	3.7507 %	EUR 0.12	EUR 253.66	1,444	EUR 0.18
facebook 55+, Audience network 23.02.2018 - 15.04.2018, completed	9	LIFETIME: EUR 253.66	1.31	28,751	39,092	1,531	3.9164 %	EUR 0.14	EUR 209.51	700	EUR 0.30
facebook 18-24, Messenger 23.02.2018 - 15.04.2018, completed	9	LIFETIME: EUR 140.28	1.39	106,772	148,420	4,116	2.7732 %	EUR 0.03	EUR 107.22	509	EUR 0.21
facebook 25-34, Messenger 23.02.2018 - 15.04.2018, completed	9	LIFETIME: EUR 237.13	1.35	85,079	114,927	3,259	2.8357 %	EUR 0.06	EUR 192.54	1,229	EUR 0.16
facebook 35-44, Messenger 23.02.2018 - 15.04.2018, completed	9	LIFETIME: EUR 372.86	1.44	99,828	144,277	5,701	3.9514 %	EUR 0.05	EUR 312.18	2,655	EUR 0.12
facebook 45-54, Messenger 23.02.2018 - 15.04.2018, completed	9	LIFETIME: EUR 301.81	1.23	27,820	34,374	1,635	4.7565 %	EUR 0.15	EUR 252.17	1,119	EUR 0.23
facebook 55+, Messenger 23.02.2018 - 15.04.2018, completed	9	LIFETIME: EUR 292.70	1.31	17,266	22,675	1,394	6.1477 %	EUR 0.17	EUR 242.90	902	EUR 0.27

THE SOLUTION: ADPLATFORM'S AUTOPILOT

ADPLATFORM | BMW / BMW (CPC) Post Boost | 12:00 PM | SETTINGS

BMW (CPC) Post Boost

STATUS: active | BUDGET: 10,000 | DAILY BUDGET: 10,000 | START DATE: 23.08.2018 - 31.12.2018 | END DATE: 23 | TARGET OPTIONS: 3

AD GROUPS | COMMENTS | OPTIMIZATIONS | TARGETS | AUTOPILOT | CAMPAIGN RULES

AUTOMATED BUDGET ALLOCATION

NAME: BMW (CPC) Post Boost

AUTOMATION BUDGET

Turn On Off

BUDGET: 10,000 | DAILY BUDGET: 10,000

TRY TO SPEND BUDGET UNTIL: 01.09.2018 | 11:00 AM

Spend budget as evenly as possible
The algorithm will suggest a mix of bids to spend the budget evenly over the campaign period.

Spend as much as possible

OPTIMIZATION GOALS:

Spend out of the budget and more

Spend as much as possible

HOW DO YOU WANT TO AUTOMATE YOUR CAMPAIGN:

Automate all ad groups

Only subman selected ad groups in the campaign

AD GROUPS RUN BY YOU:

Name of this group (M / 2 / has interests)

Group name (don't want to be automated)

Add one more group (don't want to automate)

AUTOMATED AD GROUPS:

AD group name

Another group name

Group with many bids in it

Here's a group name full of bids

This is the lowest name reach value of all

CAMPAIGN RULES

ADD RULE:

Turn off when more than 1000 Impressions daily
Group: Liberate | Post: Hourly

10000 Impressions on Ad Campaign is greater or equal to 1,000

10000 Notify: notify@stuart@gmail.com

Double the budget if zero impressions
Group: 7 days | Post: Daily

10000 Impressions on Ad Campaign is 0

10000 Increase budget by 100% with no bid

10000 Notify: notify@stuart@gmail.com

Call 911

Hitpool AdPlatform Support | +360 1 478 91 80 | info@hitpool.com

- reach delivery and performance goals
- execute the budget pacing plan
- continuously allocate budget for the highest ROI with respect to advertiser's KPI's

ADPLATFORM | BMW / BMW (CPC) Post Boost | 12:00 PM | SETTINGS

BMW (CPC) Post Boost

STATUS: active | BUDGET: 10,000 | DAILY BUDGET: 10,000 | START DATE: 23.08.2018 - 31.12.2018 | END DATE: 23 | TARGET OPTIONS: 3

AD GROUPS | COMMENTS | OPTIMIZATIONS | TARGETS | AUTOPILOT | CAMPAIGN RULES

AUTOMATED BUDGET ALLOCATION

NAME: BMW (CPC) Post Boost

Smart budget allocation feature is AdPlatform's "autopilot". It was build to help you increase campaign's ROI.

The autopilot enables you to maintain control over the campaign:

Turn autopilot on or off anytime you want
Use autopilot on all groups or only on selected ones
Constrain the budget that will be run by the autopilot

How it works:

When you turn the autopilot on the system asks you to define campaign's lifetime budget, and the number of campaign days. At first, all ad groups will receive initial exploration budget in order to learn about how promising they are. The autopilot will then each day select the most promising groups and boost them, while decreasing budget for the less promising ones. The system continuously balances between exploration and exploitation. This means all groups get a chance, but best performing ones receive more budget.

Benefits:

Reach delivery and performance goals
Execute the budget pacing plan
Continuously allocate budget for the highest ROI with respect to advertiser's KPI's

To test how well it works, we suggest you do the following:

1. Turn on auto pilot
2. Assign a portion of the budget to manually managed groups
3. Compare the results

Hitpool AdPlatform Support | +360 1 478 91 80 | info@hitpool.com

AUTOPILOT CAMPAIGN SETTINGS

STATUS: active
SPENT: EUR 3,938.83
FACEBOOK AD ACCOUNT: ANIMO FORTI SIA EUR
ADS: 223
PEOPLE TARGETED: 309,000,000
OBJECTIVE: Conversions
AUTOMATION: Autopilot

Group list | Statistics | Automation | Segments | Creatives

AUTOPILOT

AUTOMATED BUDGET ALLOCATION
Automatically allocate EUR 5,000.00 for 51 days (until 15.04.2018)

OPTIMIZATION GOAL
Spend out all the budget and maximize Offsite conversions

AUTOPILOT

SETTINGS

Complete campaign budget: 5,000.00 EUR (min. 5,000.00)
Spent: EUR 3,938.83

Spend until 15.04.2018 Europe/Riga

⚠ Once you turn the autopilot on, the end dates of all groups will be set to 15.04.2018

Adjust period 1 h

OPTIMIZATION GOAL

Spend out all the budget and maximize Offsite conversions

Spend as much as possible with the max bid 0.05 EUR and optimize

Cancel Save settings and turn Autopilot

- Offsite conversions
- Clicks
- Link clicks
- Offsite conversions
- Video views
- Post engagements

CASE STUDY: THE CLIENT & THE CHALLENGE

- nurturing skin care company 
- twofold goal:
 - test two different kinds of products
 - maximize the number of conversions



 **Hemptouch**
Sponzorirano · €

Všeckaj stran

Being a teenager is all about you. AND ACNE. You can get clear skin by this summer.



Hemp cannabinoids
Brand new approach to acne treatment

WWW.HEMPTOUCH.COM [Nakup](#)



 **Hemptouch**
Sponzorirano · €

Všeckaj stran

Sometimes your body needs a boost. Try cannabinoids. Fights insomnia, autoimmune disease, cardiovascular disease, weak immune system.

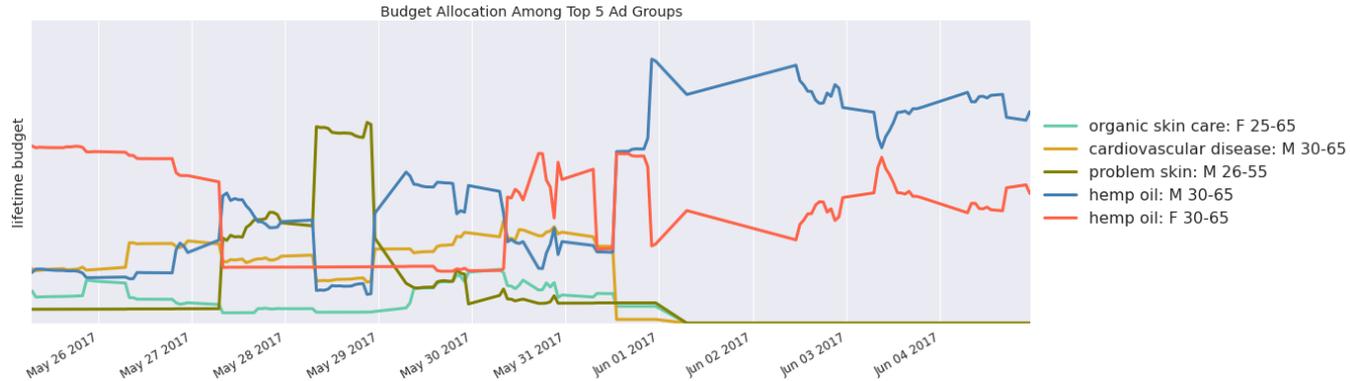


Give strength to your endocannabinoid system.
100% natural hemp CBD extract.

WWW.HEMPTOUCH.COM [Nakup](#)

CASE STUDY: AUTOMATED BUDGET ALLOCATION

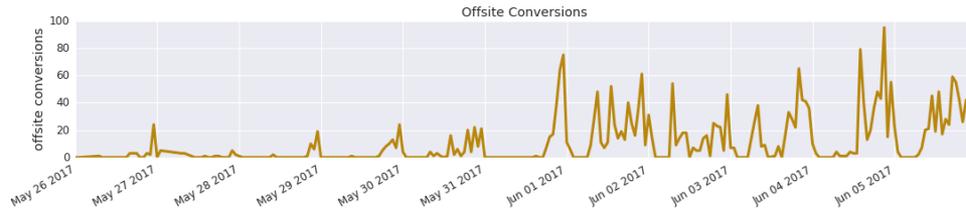
- the autopilot automatically allocates budgets between ad groups
 - the system continuously balances between exploration and exploitation
 - this means that all groups get a chance, but best performing ones receive more budget



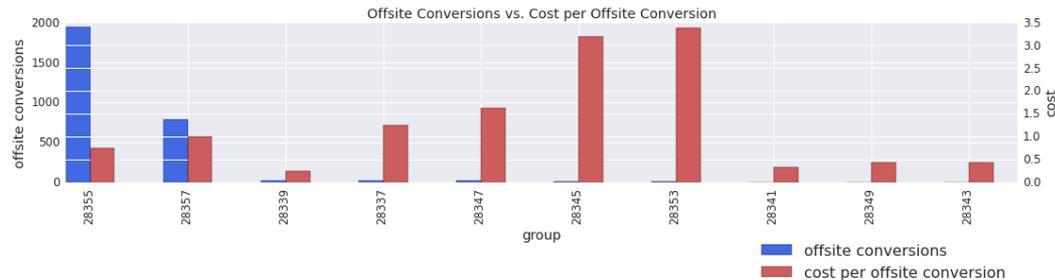
- the “hemp oil” ad groups performed best and eventually received the entire budget

CASE STUDY: RESULTS

- smarter budget allocation resulted in better campaign performance



- the “hemp oil” ad groups vastly outperformed other ad groups



CASE STUDY: SUMMARY

To summarize, here are a few takeaways from the case study:

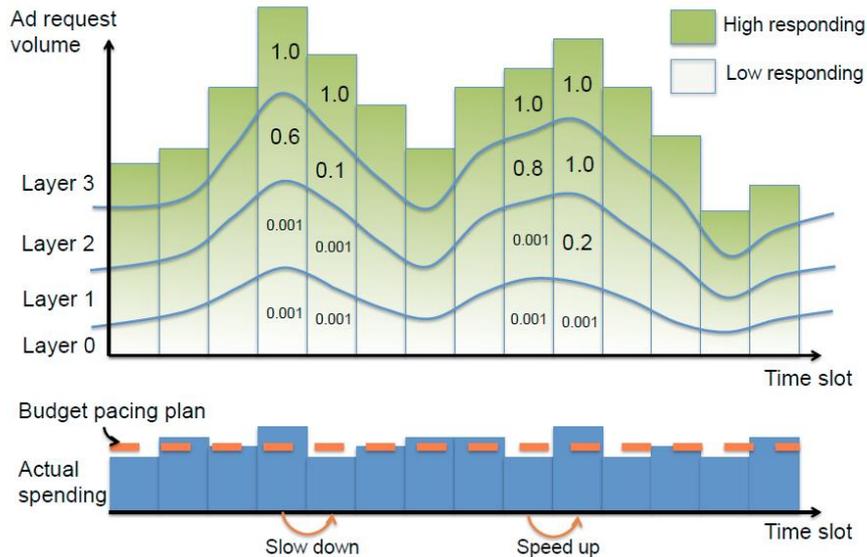
- the AdPlatform's Autopilot **automatically allocates budgets** between ad groups
- the system **continuously balances** between exploration and exploitation
- all groups get a chance, but **best performing** ones receive **more budget**

Autopilot

- can save a huge amount of manpower
- achieve better campaign performances
- avoid the risk of "junior" mistakes



BUDGET ALLOCATION & BUDGET PACING



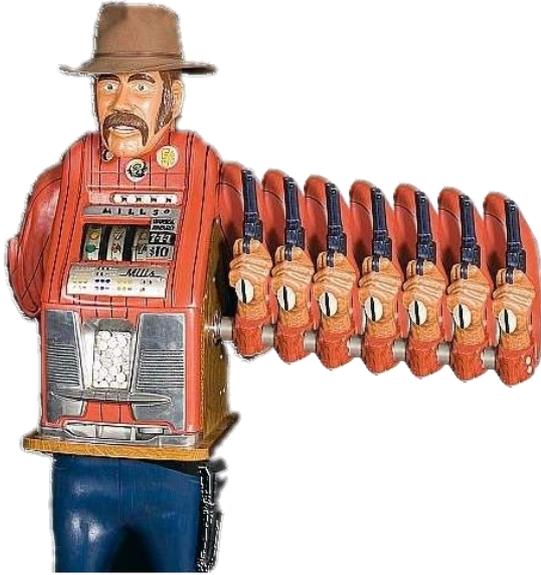
- increase budget spending on better-performing ad sets
- decrease budget spending on worse-performing ad sets
- control budget pacing

However...

- the advertising landscape is constantly changing
- need to balance between exploration and exploitation(!)

BUDGET ALLOCATION

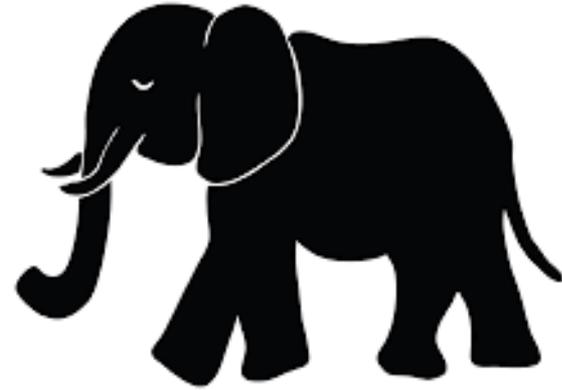
quality assurance



pick so as to get the most profit
as you can over time

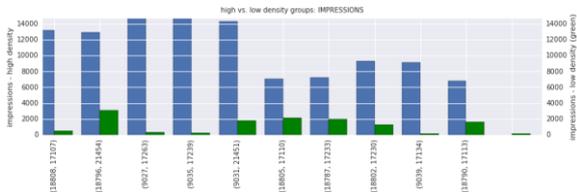
DELIVERY PACING

quantity assurance

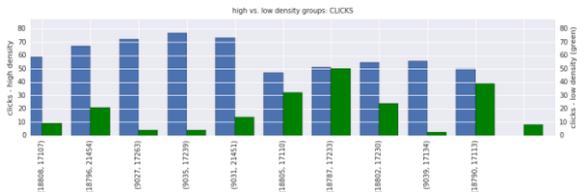


reach delivery and performance goals
execute the budget pacing plan

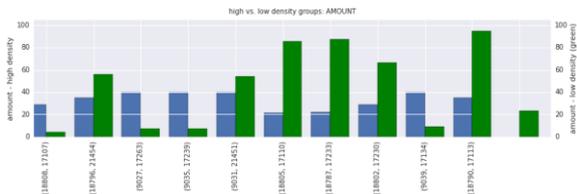
MORE BUDGET \neq BETTER PERFORMANCE



✓ blue ad sets received far **more impressions** than green ad sets



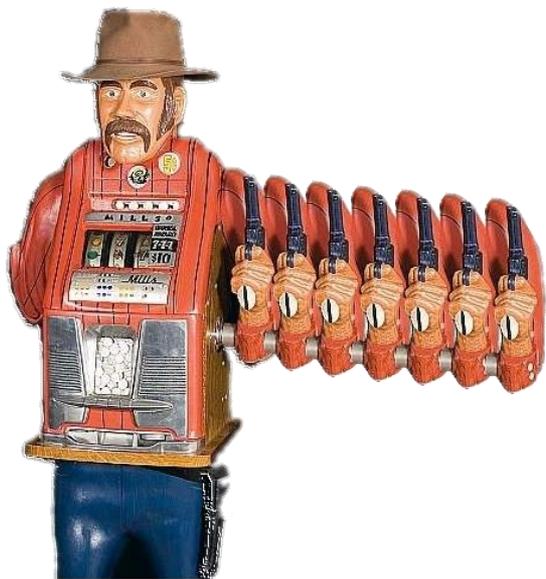
✓ blue ad sets received far **more clicks** than green ad sets



✓ blue ad sets spent far **less amount** than green ad sets

WHAT IS MULTI-ARMED BANDIT

How to pick between different arms so that you walk out with most \$\$\$ out of Casino at the end of the Night?



OBJECTIVE

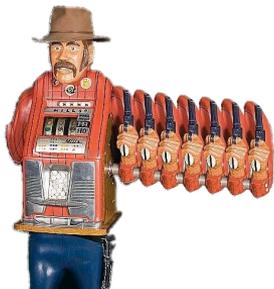
pick so as to get the most return/profit as you can over time

technical term: minimize regret

SEQUENTIAL SELECTION

- which sequence of arms to pick?
- need to sample, but do it efficiently

EXPLORATION vs. EXPLOITATION



EXPLORATION

- investment: data collection is costly
- be efficient → balance the potential value of collecting new data with exploiting what we currently know

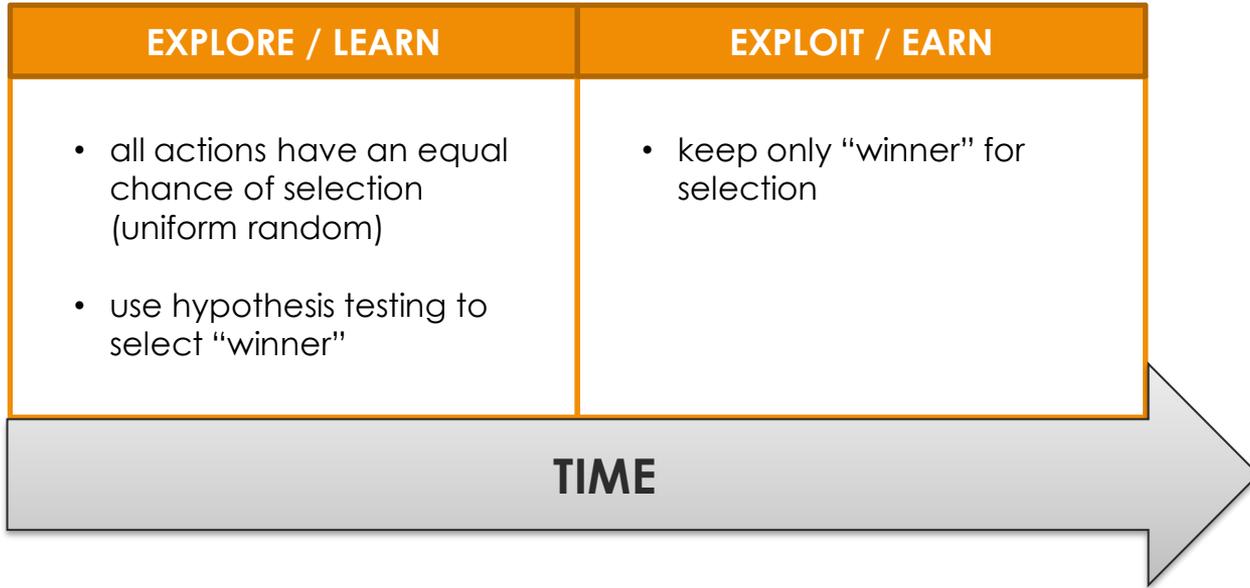
EXPLORE / LEARN

- try out different actions to learn how they perform over time
- a data collection task
- choosing actions whose benefit will come only later

EXPLOIT / EARN

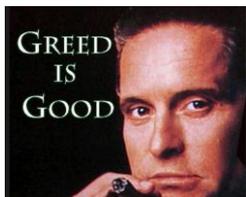
- take advantage of what you have learned to get highest payoff
- our current best guess
- choosing actions that yield immediate reward

A/B TESTING



drawback: reward distributions are usually dynamic (best combination of ads changes over time)

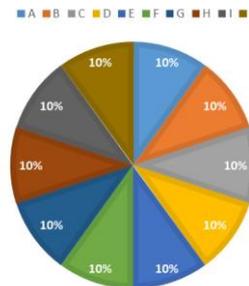
EPSILON GREEDY



greedy: make whatever choice seems best at the moment

ϵ – greedy:

- explore – randomly select action ϵ percent of the time (e.g. 20%)
- exploit – play greedy (pick the current best) $100 - \epsilon$ percent of the time (e.g. 80%)



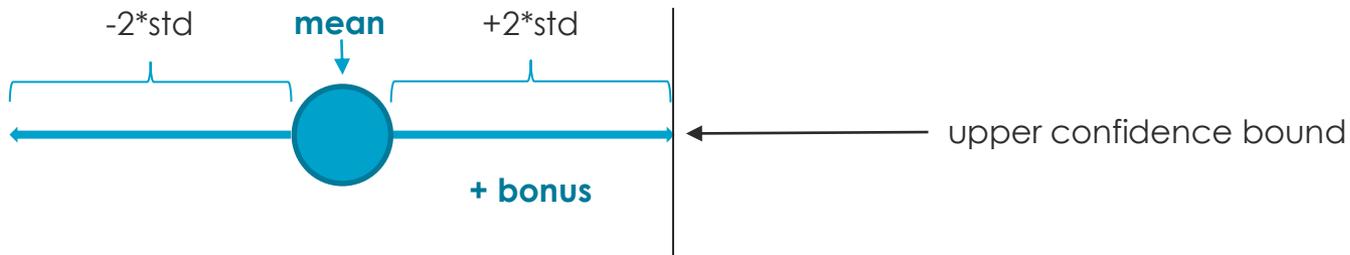
A	\$5.49
B	\$4.78
C	\$3.77
D	\$3.25
E	\$3.11
...	...



UPPER CONFIDENCE BOUND

BASIC IDEA

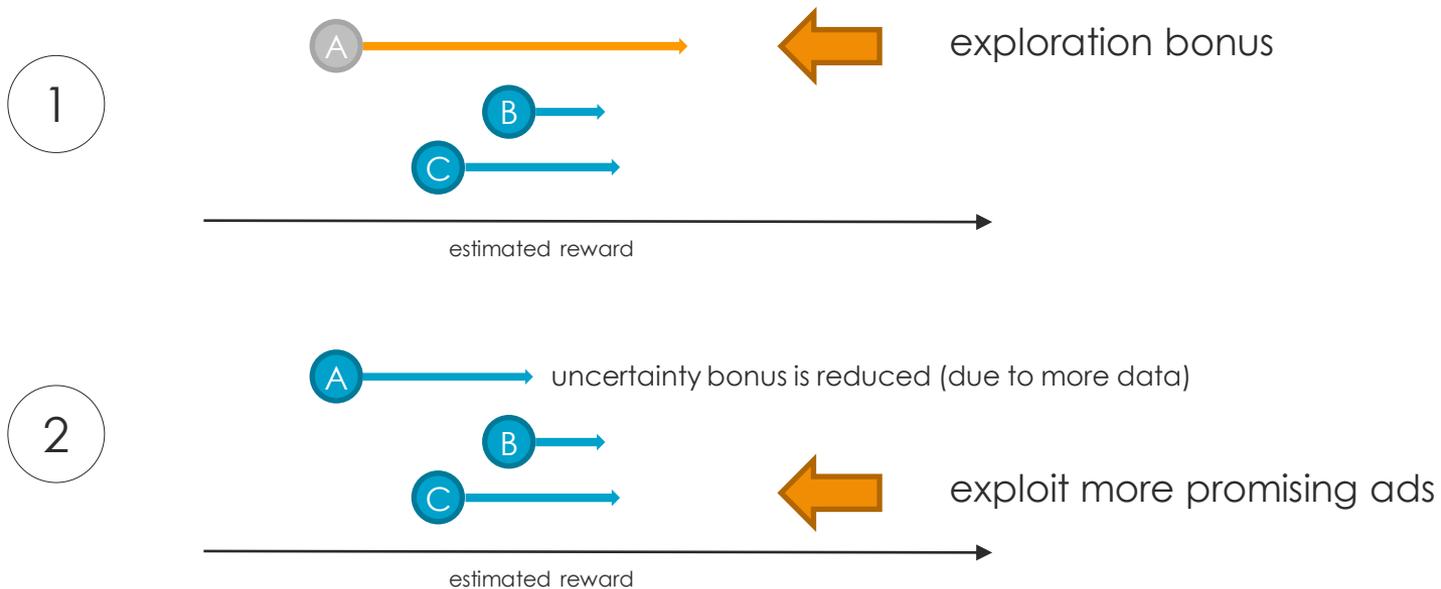
- calculate both mean and a measure of uncertainty (variance) for each action
- make greedy selections based on mean + uncertainty bonus



score each option using the upper portion of the confidence interval as a bonus

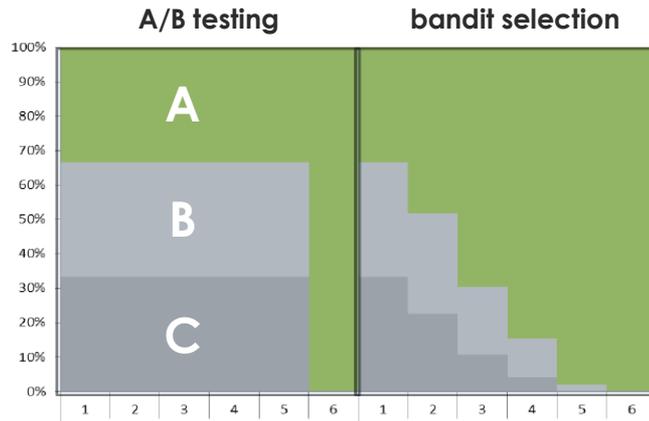
EXPLORATION BONUS

- reduce uncertainty by collecting more data
- strive towards statistical significance



UPPER CONFIDENCE BOUND

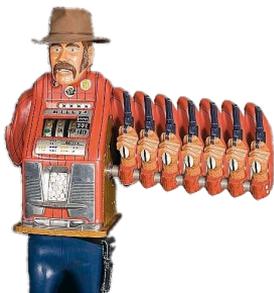
- like A/B test: uses variance measure
- unlike A/B test: no hypothesis test
- automatically balances exploration with exploitation



- more efficient learning
- automation

BUDGET-LIMITED MULTI-ARMED BANDITS

agent's actions are costly and constrained by a fixed budget



By pulling **arm i** , the agent has to pay a **pulling cost c_i** .

The agent has a **cost budget B** , which it cannot exceed during its operation time.



ONLINE ADVERTISING

Both exploration and exploitation are costly, a **combination of arms** can be pulled at once.

eCPM

eCPM or *effective CPM*: estimated earnings (or cost) for every 1000 impressions received

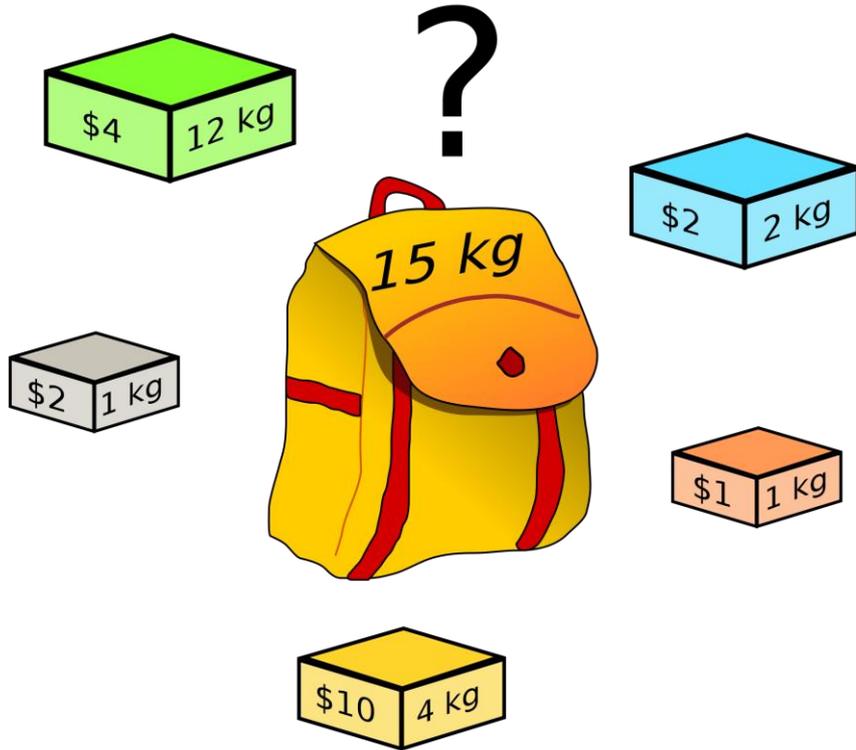
$$\frac{\text{Total Earnings}}{\text{Impressions}} \times 1000 = \text{eCPM}$$

	ad #1	ad #2
impressions	370	187
earnings	\$1.48	\$0.97
eCPM = ?	\$4.00	\$5.19

earnings can be various delivery goals:

- clicks
- impressions
- conversions
- likes
- follows
- ...

THE UNBOUNDED KNAPSACK PROBLEM



knapsack capacity → budget

items → ads

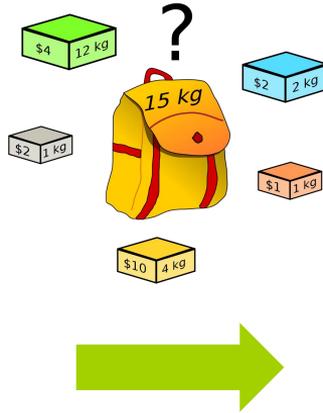
value → delivery goal (conversions, clicks, ...)

weight → amount spent

THE UNBOUNDED KNAPSACK PROBLEM

GIVEN

- k items (ads)
- each item i has a corresponding
 - value v_i
 - weight w_i
- knapsack of capacity B (budget)



OBJECTIVE

- maximize the total value of items in the knapsack
- total weight of the items should not exceed the knapsack weight capacity

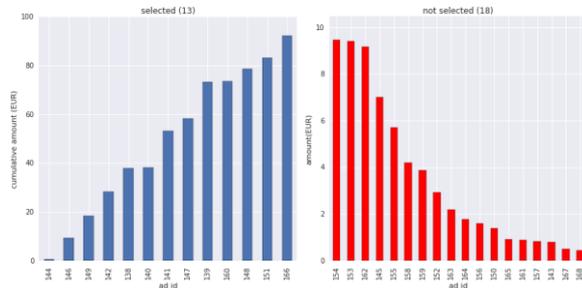
The unbounded knapsack problem is NP-hard.

An efficient approximation method for solving the knapsack problem is the **Density-Based Greedy Algorithm**.

DENSITY-BASED GREEDY ALGORITHM

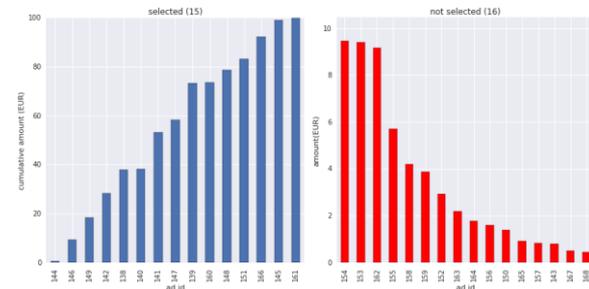
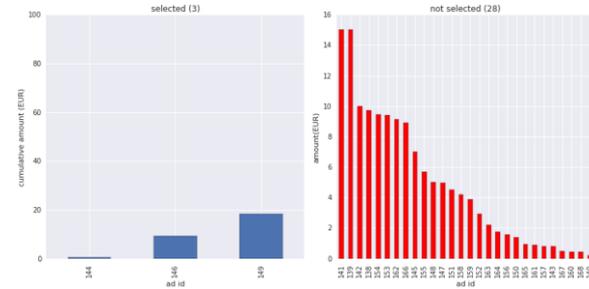
Let v_i/w_i denote the **density** of item i

- initially selected ads (ads appearing for the first time)



- the highest “density” ads are selected as is feasible without exceeding the knapsack capacity

- repeat until no feasible items left: the “densest” ad of the remaining feasible ads is selected



EXPERIMENTAL DESIGN

- Facebook data (subset)
 - 60 days (2016-02-01 to 2016-03-31)
 - 2000+ examples (key: ad, date)
- objective:
 - maximize CPC (clicks, impressions, ...)
 - do not exceed the daily budget of 100 EUR

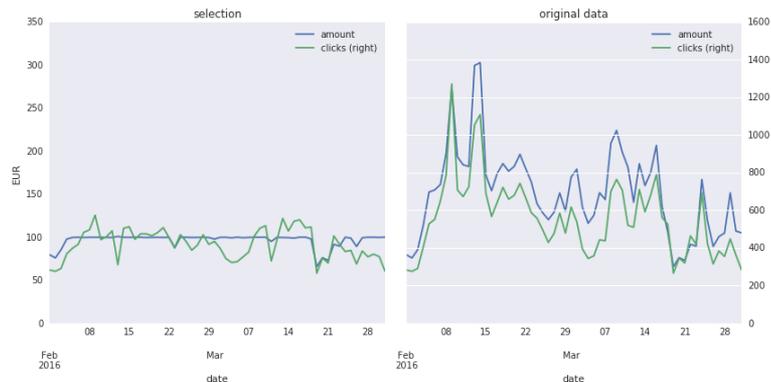
- original data, daily amount spent:

- mean: 152.07
- median: 152.04

min	65.75
25%	115.67
50%	152.04
75%	181.48
max	302.74

	date	ad_id	impressions	clicks	amount
0	2016-02-01	2320965	2757	30	4.99
1	2016-02-01	3710056	4597	53	9.71
2	2016-02-01	3710072	4903	40	9.40
3	2016-02-01	464191	28123	19	5.00
4	2016-02-01	464189	1608	18	5.00

	date	ad_id	impressions	clicks	amount
2088	2016-03-31	463455	3564	1	0.72
2089	2016-03-31	463419	1127	0	0.50
2090	2016-03-31	1861452	91	0	0.43
2091	2016-03-31	1861460	142	0	0.31
2092	2016-03-31	1861468	52	0	0.12



ORIGINAL DATA

impressions: 13,815,098
 clicks: 33563
 amount spent: 9,124.38 EUR
 CPC: 0.272 EUR
 eCPM: 2.429 (estimated clicks per 1000 impressions)

RESULTS (daily budget < 100 EUR)

impressions: 7,061,437
 clicks: 25225
 amount spent: 5,791.36 EUR
 CPC: 0.230 EUR
 eCPM: 3.572 (estimated clicks per 1000 impressions)

optimal results (based on perfect information)

UCB FORMULA

UCB formula

$$r_j = \mu + \sqrt{\frac{2 \ln n}{n_j}}$$

exploitation

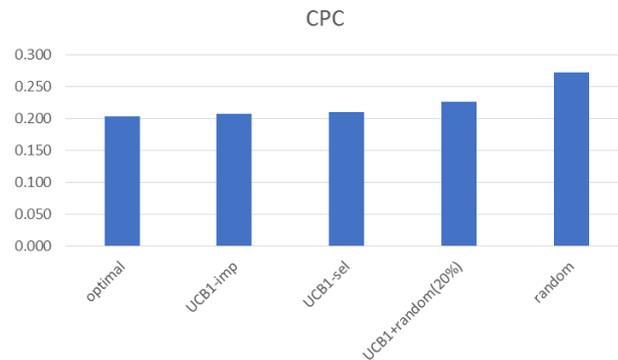
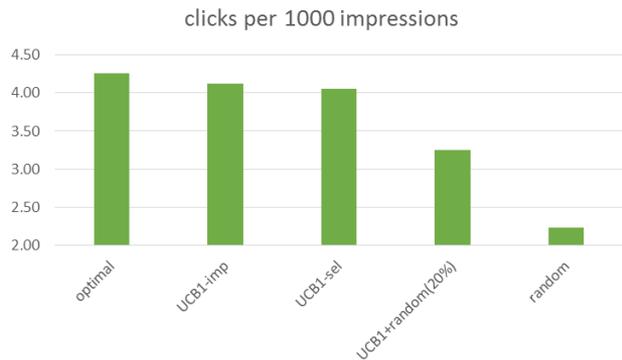
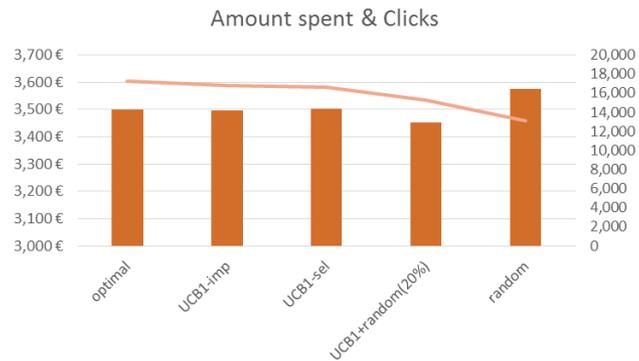
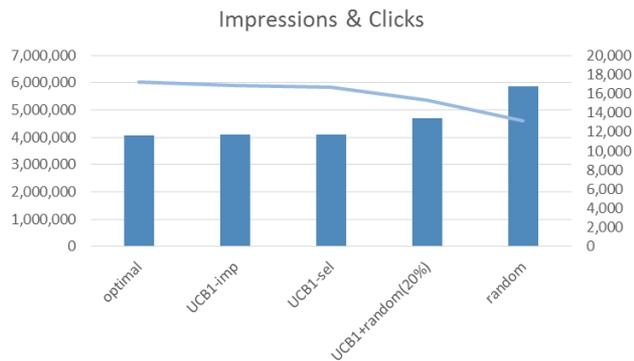
exploration

- μ is the average payoff obtained from arm j
- n_j is the number of times arm j has been pulled
- n is overall the number of pulls so far

```
def UCB1(df, row):  
    '''Return the value of UCB1 formula for the given row.'''  
  
    current_date = row['date']  
  
    # the average payoff obtained from arm j: the average "density" of this ad  
    avg_payoff = 0  
    if row['amount_cumsum'] > 0:  
        avg_payoff = row['clicks_cumsum'] / row['amount_cumsum']  
  
    # overall the number of pulls so far: cumulative sum of impressions for all active ads  
    n = df[df.date == current_date]['impressions_cumsum'].sum()  
  
    # the number of times arm j has been pulled: impressions that this ad received so far  
    nj = row['impressions_cumsum']  
  
    expl_bonus = 0  
    if nj > 0:  
        expl_bonus += math.sqrt(2 * np.log(n) / nj)  
  
    return avg_payoff + expl_bonus
```

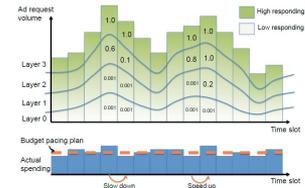
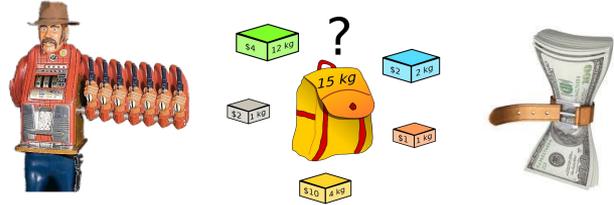
SIMULATION RESULTS: SMALLER DAILY BUDGET

Objective: optimize CPC without exceeding the daily budget of 60 EUR.

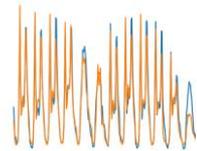


THE MACHINE LEARNING CHALLENGE

- use budget-limited multi-armed bandit approach to balance between exploration and exploitation for effective budget allocation
- use the smart pacing approach:
 - increase budget spending on better-performing ad sets
 - decrease budget spending on worse-performing ad sets

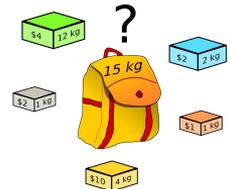


How can we benefit from predictive models?



REFERENCES: BUDGET ALLOCATION

- Niimi, Makoto, and Takayuki Ito. "Budget-Limited Multi-armed Bandit Problem with Dynamic Rewards and Proposed Algorithms." In *Advanced Applied Informatics (IIAI-AAI), 2015 IIAI 4th International Congress on*, pp. 540-545. IEEE, 2015.
- Tran-Thanh, Long, Lampros C. Stavrogiannis, Victor Naroditskiy, Valentin Robu, Nicholas R. Jennings, and Peter Key. "Efficient regret bounds for online bid optimisation in budget-limited sponsored search auctions." (2014): 809-818.
- Tran-Thanh, Long. "Budget-Limited Multi-Armed Bandits." PhD dissertation, University of Southampton, 2012.
- Tran-Thanh, Long, Archie Chapman, Alex Rogers, and Nicholas R. Jennings. "Knapsack based optimal policies for budget-limited multi-armed bandits." *arXiv preprint arXiv:1204.1909* (2012).
- Matt Gershoff. *Bandit Basics – A Different take on Online Optimization* (invited talk). eMetrics Summit, Stockholm Summit, Sweden, October 2012.
- Tran-Thanh, Long, Archie Chapman, Jose Enrique Munoz De Cote Flores Luna, Alex Rogers, and Nicholas R. Jennings. "Epsilon-first policies for budget-limited multi-armed bandits." (2010): 1211-1216.



REFERENCES: BUDGET PACING

- Xu, Jian, Kuang-chih Lee, Wentong Li, Hang Qi, and Quan Lu. "Smart pacing for effective online ad campaign optimization." In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 2217-2226. ACM, 2015.
- Geyik, Sahin Cem, Abhishek Saxena, and Ali Dasdan. "Multi-touch attribution based budget allocation in online advertising." In *Proceedings of the Eighth International Workshop on Data Mining for Online Advertising*, pp. 1-9. ACM, 2014.
- Lee, Kuang-Chih, Ali Jalali, and Ali Dasdan. "Real time bid optimization with smooth budget delivery in online advertising." In *Proceedings of the Seventh International Workshop on Data Mining for Online Advertising*, p. 1. ACM, 2013.
- Jayawardane, C. H. W., S. K. Halgamuge, and U. Kayande. "Attributing Conversion Credit in an Online Environment: An Analysis and Classification." In *Computational and Business Intelligence (ISCBI), 2015 3rd International Symposium on*, pp. 68-73. IEEE, 2015.

