


Action Rule Mining

Prague University of Economics and Business

What Is an Action Rule?



Undesired Part ($r_{undesired}$) - The initial condition under which the undesired target value occurs.

Desired Part ($r_{desired}$) - The condition after applying the change, under which the desired target value occurs.

Key Concepts in Frequent Itemset Mining

- **Item** - A single attribute-value pair, e.g., Age=Young, Class=2nd.
- **Itemset** - A set of items (attribute-value pairs), e.g., {Sex=Female, Age=Young, Class=2nd}.
- **Antecedent** - The if part of a rule; the conditions that need to be met.
E.g., Sex=Female $\wedge$ Age=Young $\wedge$ Class=2nd.
- **Consequent** - The then part of a rule; the predicted outcome, e.g., Survived=Yes.
- **Support** - Number of records in the dataset that contain a given itemset. Example: Support of {Sex=Female, Age=Young} is the number of passengers matching both.
- **Confidence** - Probability that the rule's consequent is true when the antecedent holds.
For rule $A \Rightarrow B$: confidence = $\text{support}(A \wedge B) / \text{support}(A)$

Action rule mining builds on definitions from frequent itemset mining


- **Support of action rule** is $\min(\text{support}(r_{\text{undesired}}), \text{support}(r_{\text{desired}}))$
- **Confidence of action rule** is $\text{confidence}(r_{\text{undesired}}) \times \text{confidence}(r_{\text{desired}})$
- **Uplift of action rule** is a new metric (more details later)

Current State – Algorithms


- Rule-based approach
 - Ras, Z. W., & Wieczorkowska, A. (2000). Action-rules: How to increase profit of a company.
 - Raś, Z. W., Wyrzykowska, E., & Wasyluk, H. (2007). Aras: Action rules discovery based on agglomerative strategy.
 - Ras, Z. W., & Tsay, L.-S. (2003). Discovering extended action-rules (system dear).
- Object-based approach
 - He, Z., Xu, X., Deng, S., & Ma, R. (2005). Mining action rules from scratch.
 - Hashemi, S., & Shamsinejad, P. (2021). Ga2rm: A ga-based action rule mining method.

Current State – Software


- LISp-Miner
 - Based on the General Unary Hypotheses Automaton (GUHA).
[Rauch Jan, Š. M. (2014). Lisp-miner, GUHA a data mining. Oeconomica]
- SCARI
 - Separate and Conquer Algorithm for Action Rules and Recommendations Induction.
[Sikora, M., Matyszok, P., & Wróbel, Ł. (2022). Scari: Separate and conquer algorithm for action rules and recommendations induction. Information Sciences, 607, 849–868.]




1) Nodes derived from stable attributes must have sufficient support for both desired and undesired classes in the consequent.




2) Groups of nodes representing flexible attributes must have sufficient support for both desired and undesired classes.




3) The required minimum number of stable or flexible attributes must be met.



4) Each path in the trie can contain at most one item derived from the same attribute.



5) The itemset is on the stop list.



Contribution: New measure for action rules

Existing measures for action rules are hard to interpret.

$$uplift = P(\text{decision} \mid \text{treatment}) - P(\text{decision} \mid \text{no treatment}) \quad [\text{Radcliffe (2007)}]$$

$$uplift(r_a) = \frac{\overbrace{(conf_{r_{desired}} - (1 - conf_{r_{undesired}}))}^{\substack{\text{Confidence of desired class} \\ \text{for desired rule}}} * \overbrace{\frac{supp_{r_{undesired}}}{conf_{r_{undesired}}}}^{\substack{\text{Confidence of desired class} \\ \text{for undesired rule}}} * \overbrace{\frac{supp_{r_{undesired}}}{conf_{r_{undesired}}}}^{\substack{\text{Number of transactions} \\ \text{that can be changed}}}}{\underbrace{|t \in D : t|}_{\substack{\text{Number of} \\ \text{transactions}}}}$$

Uplift quantifies the increase of desired target states in a dataset.

Contribution: New pruning technique for Action Rules

Dominant Action Rules

A dominant action rule is defined as a rule that cannot be extended by additional conditions without reducing its uplift. In cases where an extension of the rule leads to a higher uplift, a new dominant action rule is formed. Inspired by the concept of closed association rules (Pasquier et al., 1999).

Dominant $\left\{ \text{rigor} = (\text{medium} \rightarrow \text{high}) \Rightarrow \text{evaluation} = (\text{bad} \rightarrow \text{good}) \right.$

with uplift 16.96%.

Not Dominant $\left\{ \text{rigor} = (\text{medium} \rightarrow \text{high}) \wedge \text{grammar} = (1 \rightarrow 0) \Rightarrow \text{evaluation} = (\text{bad} \rightarrow \text{good}) \right.$

with uplift 10.82%.

Contribution: High-Utility Action Rules Mining

$$u(r_{1 \rightarrow 2}) = u(r_2) - u(r_1) = (ul(\phi_2) - ul(\phi_1)) \times (\text{conf}(r_2) - (1 - \text{conf}(r_1))) + \sum_{n=1}^k (ul(\alpha_2^{(n)}) - ul(\alpha_1^{(n)}))$$

where $ul(\phi_2)$ represents the utility of the desired class, $ul(\phi_1)$ represents the utility of the undesired class, and $ul(\alpha_1^{(n)})$ and $ul(\alpha_2^{(n)})$ represent the utilities of items derived from flexible attributes before and after the recommended action. The expression $ul(\phi_2) - ul(\phi_1)$ computes the difference in utility between the desired and undesired classes. This value is multiplied by $\text{conf}(r_2) - (1 - \text{conf}(r_1))$, which represents the change in confidence of the desired class as predicted by r_2 compared to its confidence as predicted by r_1 . The final part of the formula, $\sum_{n=1}^k (ul(\alpha_2^{(n)}) - ul(\alpha_1^{(n)}))$, captures the total change in utility resulting from the recommended actions.

Item	Utility
Salary: Low	-300
Salary: Medium	-500
Salary: High	-1000
Attrition: False	700
Attrition: True	0

Contribution: Method for applying action rule mining to unstructured text data I

Feature Discovery



Feature Generation



Action Rules Mining


Feature discovery refers to the process of identifying what features should be extracted from the data — essentially discovering the names or types of features that may be relevant for a given task. In this context, an LLM analyzes the dataset and proposes which features to generate.

Feature generation is the process of extracting the values of those features for individual data instances. This also uses LLMs, which are provided with prompts to generate the specific feature values.

Contribution: Method for applying action rule mining to unstructured text data II

Rule	Area	Novelty	Rigor	Grammar	Replicability	Accessibility	Evaluation	Uplift
r_1			med $\rightarrow$ high		0 $\rightarrow$ 1		bad $\rightarrow$ good	16.96%
r_2			med $\rightarrow$ high				bad $\rightarrow$ good	16.95%
r_3				1 $\rightarrow$ 0	0 $\rightarrow$ 1		bad $\rightarrow$ good	11.26%
r_4		med $\rightarrow$ high			0 $\rightarrow$ 1		bad $\rightarrow$ good	6.73%
r_5		med $\rightarrow$ high		1 $\rightarrow$ 0			bad $\rightarrow$ good	4.36%
r_6		med $\rightarrow$ high		1			bad $\rightarrow$ good	3.96%
r_7		med $\rightarrow$ high	med $\rightarrow$ high				avg. $\rightarrow$ best	4.57%
r_8		med $\rightarrow$ high		1 $\rightarrow$ 0			avg. $\rightarrow$ best	3.58%
r_9		low $\rightarrow$ high	med $\rightarrow$ high				avg. $\rightarrow$ best	2.24%
r_{10}		low $\rightarrow$ high			0 $\rightarrow$ 1		avg. $\rightarrow$ best	1.95%
r_{11}		low $\rightarrow$ high		1 $\rightarrow$ 0			avg. $\rightarrow$ best	1.24%
r_{12}	compu.	med $\rightarrow$ high	med $\rightarrow$ high				bad $\rightarrow$ good	1.98%
r_{13}	elect.		med $\rightarrow$ high			med	bad $\rightarrow$ good	1.15%
r_{14}	mechan.	med $\rightarrow$ high	med $\rightarrow$ high	1 $\rightarrow$ 0	0 $\rightarrow$ 1		bad $\rightarrow$ good	1.51%
r_{15}	chemic.	med $\rightarrow$ high	med $\rightarrow$ high				bad $\rightarrow$ good	1.88%
r_{16}	mater.	med $\rightarrow$ high			0 $\rightarrow$ 1		bad $\rightarrow$ good	1.54%

Contribution: Use of GPU data structures for action rule mining



Contribution: Python packages

Criteria / Packages	actionrules-lukassykora	action-rules
Algorithms	Baseline, ARAS, DEAR	Action-Apriori (author's algorithm)
Load classification rules	YES	NO
High-Utility Mining	YES	YES
Command-line interface	NO	YES
GPU Support	NO	YES
Multiple values per attribute	NO	YES

Contribution: User-friendly explanation in natural language


```
31 r = ar.get_rules().get_ar_notation()
32 for action_rule in r:
33     print(action_rule)
```

$r = [(gender: Female) \wedge (SeniorCitizen: 1) \wedge (InternetService*: Fiber\ optic) \wedge (OnlineSecurity: No \rightarrow Yes) \wedge (DeviceProtection: No \rightarrow Yes) \wedge (StreamingTV: No \rightarrow Yes)] \Rightarrow [Churn: Yes \rightarrow No]$ with support of undesired part: 0.073, confidence of undesired part: 0.73, support of desired part: 0.029, confidence of desired part: 0.8286, and uplift: 0.0079.


```
34 r2 = ar.get_rules().get_pretty_ar_notation()
35 for action_rule in r2:
36     print(action_rule)
```

If attribute **gender** is **Female**, attribute **SeniorCitizen** is 1, attribute (flexible is used as stable) **InternetService** is **Fiber optic**, attribute **OnlineSecurity** value **No** is changed to **Yes**, attribute **DeviceProtection** value **No** is changed to **Yes**, and attribute **StreamingTV** value **No** is changed to **Yes**, then **Churn** value **Yes** is changed to **No** with uplift: 0.007930873613111296, support of undesired part: 73, confidence of undesired part: 0.73, support of desired part: 29, and confidence of desired part: 0.8285714285714286.

Improved scalability through the use of sparse matrices and GPU




Evaluation: Benchmarking of data structures (GPU vs CPU processing)



Evaluation: Scalability

The complexity of association rule learning depends on the support threshold.


In this experiment, we observed the impact of varying the minimum support threshold on the runtime of two of the implemented algorithms (Titanic dataset).



Validation of results against LISp-Miner and performance comparison

Setting		LISp-Miner		action-rules		actionrules-lukassykora	
Min. Sup	Min. Conf	Time (s)	Number of Rules	Time (s)	Number of Rules	Time (s)	Number of Rules
70	50%	28.571	178	0.996	178	65.000	178
70	60%	26.143	32	1.070	32	13.000	32
140	50%	11.571	70	0.480	70	29.900	70
140	60%	10.000	8	0.491	8	5.520	8

Execution Time Comparison




Bootstrap

Bootstrap

- Suitable for non-parametric, rule-based systems
- Simulates model uncertainty by **resampling data**
- Produces an empirical distribution of outputs
- Provides **confidence intervals** for uplift (or other metric)

Action rule and bootstrap distribution:

$[(\text{gender: Female}) \wedge (\text{SeniorCitizen: 1}) \wedge (\text{InternetService*}: \text{Fiber optic}) \wedge (\text{OnlineSecurity: No} \rightarrow \text{Yes}) \wedge (\text{DeviceProtection: No} \rightarrow \text{Yes}) \wedge (\text{StreamingTV: No} \rightarrow \text{Yes})] \Rightarrow [\text{Churn: Yes} \rightarrow \text{No}],$
uplift: 0.007930873613111296



Links

- <https://github.com/lukassykora/action-rules>
- <https://github.com/lukassykora/actionrules>
- <https://github.com/vojtech-balek/llm-features/blob/main/notebooks/actions-M17Plus.ipynb>

Thank you for your
attention

Prague University of Economics and Business