El Farol Bar problem using learning automata

By Harald Unander & Torbjørn Vaaje

The El Farol Bar Problem

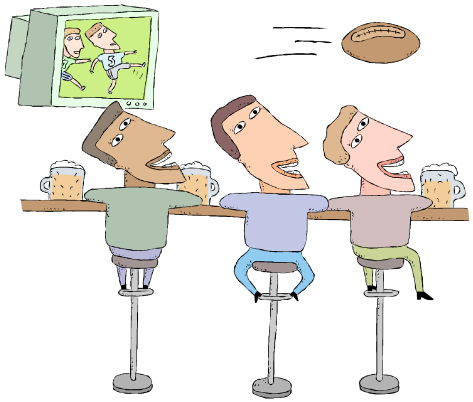

More than 60 person present

Less than 60 persons present

- Nobody knows in advanced what the others will do
- Only historical data available

Project Goals

- To see if learning automata technique is suitable for solving the El Farol Bar Problem
- Compare learning automata with the traditional solution to the El Farol Bar Problem (as described by Jeannequin/Arthur)

The Jeannequin/Arthur Agent

- Have a number of predictors to choose between
- Has an algorithm to choose the best method at a given time
- The agents do not have the same set of predictors

The Jeannequin/Arthur Agent Implementation

- Need quite a few predictors
- Need some good predictors
- Not a problem if there are bad predictors
- A predictor can be good for one data set, but bad for another data set (may not scale well)

The Jeannequin/Arthur Agent results

No	AgentCount	Rounds	Goal	Mean	SD	95%L	95%U	Zero	Full	Lower	Upper	Number of methods
1	100	1000	60	59	7	45	73	0	0	43	75	2
2	100	1000	60	59	7,2	45	73	0	0	38	82	4
3	100	1000	60	60	8,8	42	78	0	0	11	80	8
4	1000	1000	60	78	43,8	-10	166	37	0	4	140	2
5	1000	1000	60	60	6,9	46	74	22	0	42	82	4
6	1000	1000	60	85	59,3	-34	204	0	0	24	377	8
7	1000	10000	60	85	29,9	25	145	0	0	3	272	2
8	1000	10000	60	60	9,1	42	78	0	0	25	284	4
9	1000	10000	60	94	63,9	-34	222	0	0	25	322	8
10	10000	10000	60	236	320,9	-406	878	0	0	38	2577	2
11	10000	10000	60	86	117,5	-149	321	0	0	21	1477	4
12	10000	10000	60	60	7,9	44	76	0	0	31	94	8

TV Agent

- Choose a strategy and keep if for some time
- The closer the group are to the goal the longer the agent will keep its strategy
- The closer to the goal the group are it is less likely that the agent will change strategy
- No random implemented

The Tsetlin Automaton

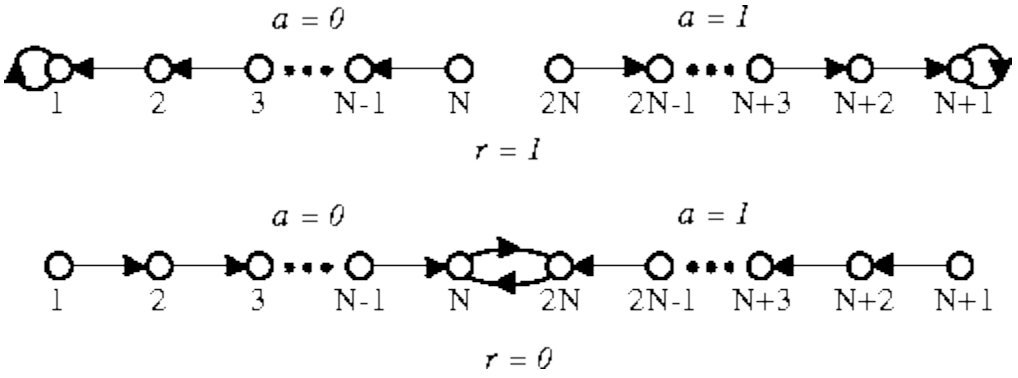

- A Tsetlin automaton with $2N$ states. The top row shows the state transitions that are made when the previous action resulted in a reward; the bottom row shows transitions after a penalization. In states in the left half of the figure, action “attend” is taken; in those on the right, action “don’t attend” is taken.

- If the previous attendance (or no attendance) is contributing to getting closer to the goal, there is a reward with a defined chance, otherwise there is a penalization (with a defined chance).

Tsetlin Results

AgentCount	Rounds	Goal	Mean	SD	95%L	95%U	Zero	Full	Lower	Higher	N	Margin	Penalize	Reward
100	1000	60	59	1,5	56	62	0	3	47	66	4	1	0.8	0.8
100	1000	60	59	1,6	56	62	1	0	55	64	8	1	0.6	0.6
100	1000	60	59	2,2	55	63	0	0	52	66	4	1	0.7	0.7
100	1000	60	59	2,5	54	64	0	0	50	68	4	1	0.7	0.7
100	1000	60	59	2,8	53	65	0	0	50	68	4	1	0.6	0.6

Table 2 Best results with agent count = 100, zero&full < 10%

AgentCount	Rounds	Goal	Mean	SD	95%L	95%U	Zero	Full	Lower	Higher	N	Margin	Penalize	Reward
1000	10000	60	60	1,7	57	63	4	0	51	74	8	1	0.7	0.7
1000	10000	60	61	2,4	56	66	0	0	53	74	12	1	0.6	0.6
1000	10000	60	61	8,1	45	77	0	0	22	147	4	1	0.9	0.9

Table 3 Best results with agent count = 1000, zero&full < 10%

AgentCount	Rounds	Goal	Mean	SD	95%L	95%U	Zero	Full	Lower	Higher	N	Margin	Penalize	Reward
10000	10000	60	63	3,6	56	70	613	0	50	79	16	1	0.6	0.6
10000	100000	60	61	1,6	58	64	248	0	54	73	20	1	0.6	0.6
10000	100000	60	63	3,6	56	70	0	0	50	82	16	1	0.6	0.6

Table 4 Best results with agent count = 10000, zero&full < 10%

Sample log

```
==============================================
Bar details:
- Agent count 100
- Rounds 1000
- Goal 60
-------------------
Agent details:
- TsetlinAgent, N=4, margin=1, penalizeLimit=0.7, rewardLimit=0.7
-------------------
Result (not including 100 rounds of stabilisation):
- Average attendance 59
- Standard deviation 2,2
- 95% interval [55-63]
- Agents never showing up 0
- Agents always present 0
- Lower participation 52
- Upper participation 66
-------------------

I0       3X   0X   1X   1X   7-   5-   1X   0X   6-   1X   3X   5-   6-   2X   1X   2X   1X   6-   2X   8-  66
I1       4X   0X   2X   2X   8-   6-   2X   0X   7-   2X   4X   4X   7-   1X   2X   3X   2X   7-   1X   8-  65
I2       5-   1X   3X   1X   7-   7-   3X   0X   8-   3X   5-   5-   6-   1X   1X   4X   1X   8-   2X   7-  55
I3       4X   1X   2X   1X   6-   6-   2X   0X   7-   2X   4X   4X   5-   1X   2X   3X   1X   8-   3X   6-  60
--
I999     7-   1X   1X   8-   1X   1X   1X   6-   4X   7-   8-   7-   1X   8-   4X   7-   2X   1X   1X   1X  61
Total   724  327  777  394  763  613  751  793  625  507  444  548  464  413  713  838  671  580  501  744
```

Tsetlin Summary

- Easy to implement
- Little tuning
- Efficient
- N must grow with the number of agents
  - Agent count 100, N=4
  - Agent count 1000, N=8 or N=12
  - Agent count 10000, N=16
- High N value gives static distribution

Comparison

- *Efficiency*
  - 10000/10000 30 seconds / 7 minutes
- *Reaching the attendance goal*
  - lower standard deviation and a smaller confidence interval.
- *Attendance distribution – see figure on next slide*
  - Tsetlin distribution slowly changing attendance state reflecting the state behaviour and randomness
  - Jeannequin reflecting the selected methods with repeating patterns
  - Tv agent rapidly reaching the static distribution
- *Implementation*
  - Tsetlin standard implemenation. Only one method needs implementation, and only the N value must be tuned with changing data sets.
  - The Jeannequin agent needs a number of methods to work well.

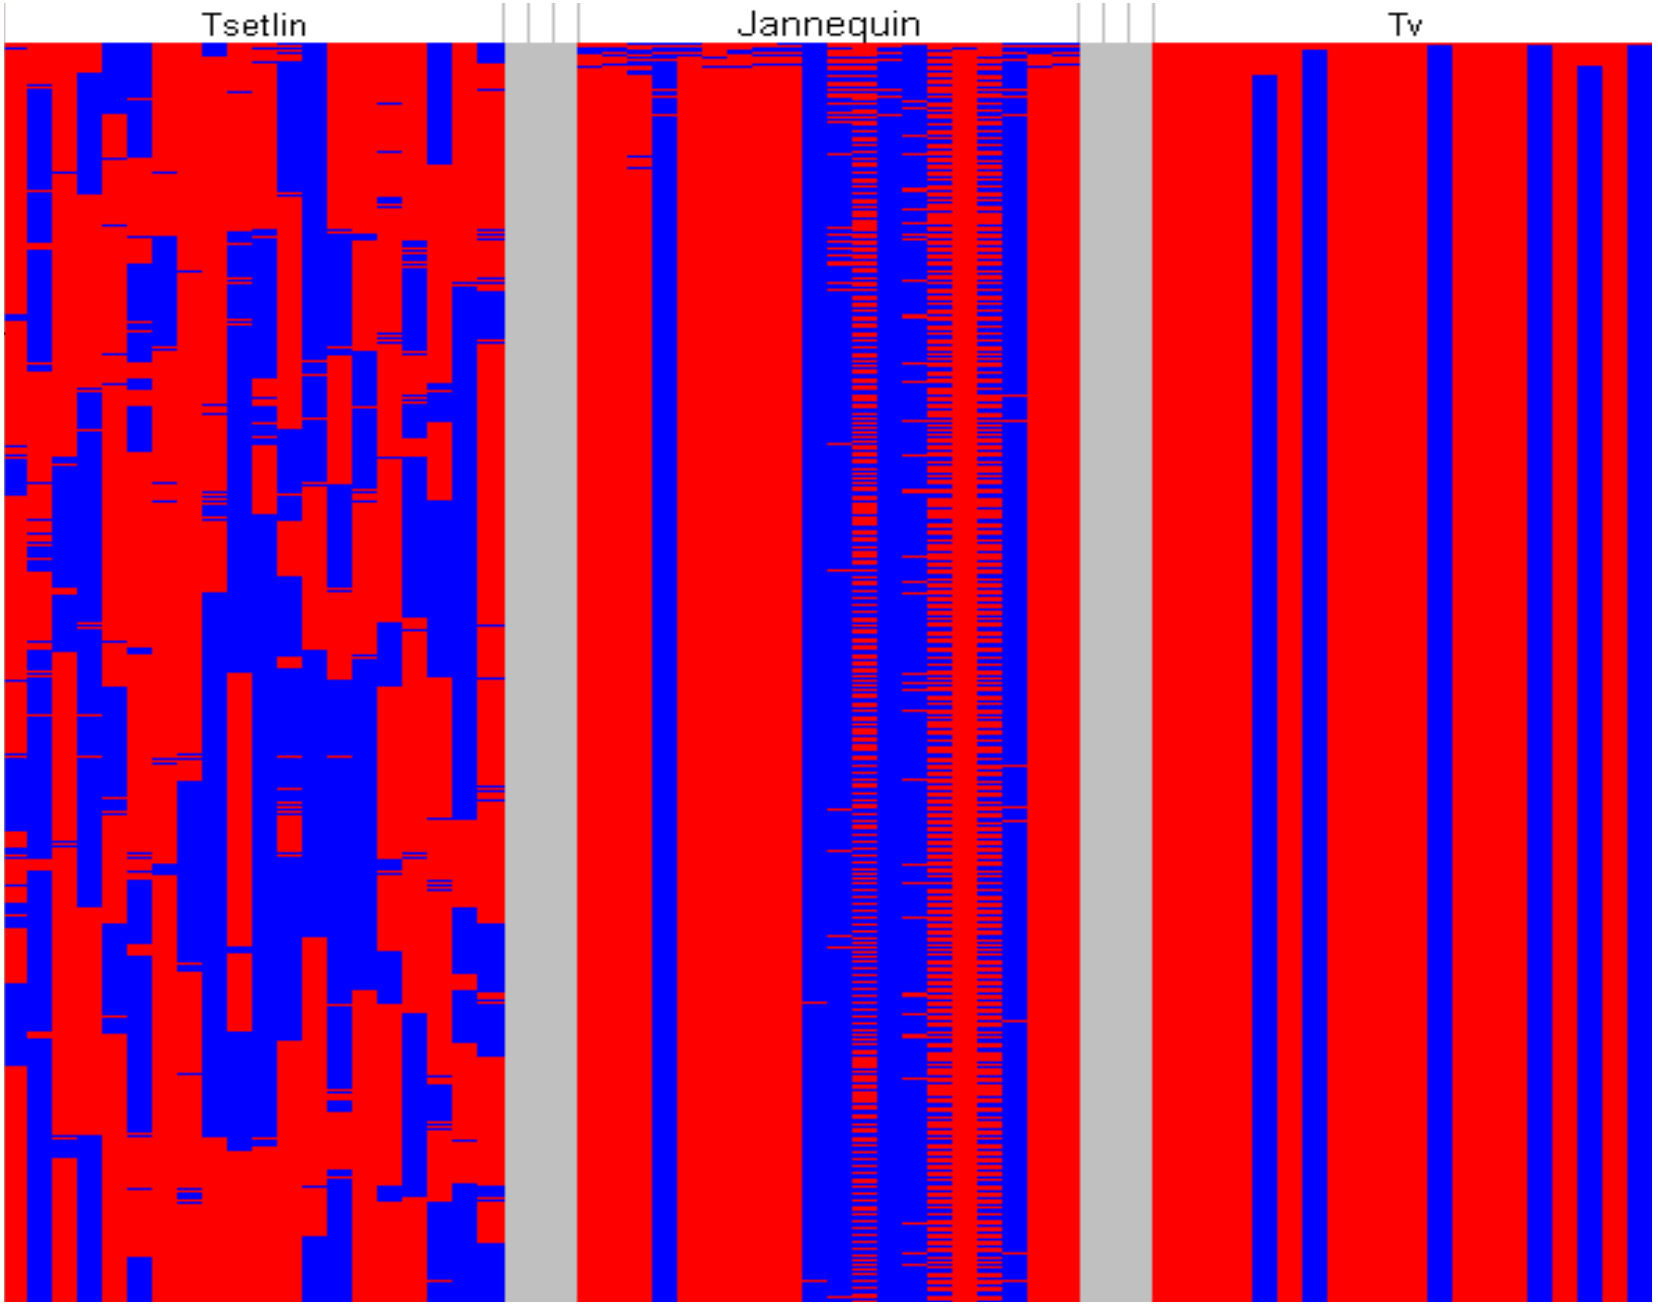
Tsetlin
Jannequin
Tv

Conclusion

- We have implemented the Tsetlin and Jeannequin agents, and even an additional agent, and they all works, reaching the desired attendance goal for the El Farol bar. Experiments show that the Tsetlin agent in all respect behaves best. It also requires less tuning of method implementations.

- The Tsetlin implementation has proved to be accurate, efficient, scaleable and not least, easy to implement and it really outperforms the competition.

- It is believed that the Jeannequin agent would gain if more effort was put into making good prediction methods. This is however also the main weakness with this agent. A lot of tuning and effort is needed to make it perform for a certain data set.

Conclusion

- There has not been time to implement mechanisms to take into account individual attendance goals of the bar customers, and we have not reached to add noise introduction, so this is left for further study.

- The implementation in Java and Eclipse has given a good framework for agent implementation, and quite some effort has been put into the logging part so that results can easily be extracted.

- By using the agent interface, new agents can easily be added, taking advantage of the provided execution and logging framework.