

Data and Analysis for Satisficing Models of Bayesian Theory of Mind for Explaining Behavior of Differently Uncertain Agents

Jan Pöppel

This data publication contains the code and data for the paper *Satisficing Models of Bayesian Theory of Mind for Explaining Behavior of Differently Uncertain Agents* [1] as well as more extensive statistical evaluation. Published at: <http://doi.org/10.4119/unibi/2918529>

1 Data description

This section describes the raw data contained in this publication. The experiment that produced this data is explained in [1].

1.1 Condition data

The folder *Conditions* contains compact representations of the conditions which were used by the webserver to setup each condition. This is included, as the scripts will need the information about the potential goal conditions of condition 2 in order to properly setup the models.

The filename and identifier for the condition is made up of three parts: First the map number, then the condition and lastly a variant, as there were two variants of each map and condition, which basically swapped the starting position and the goal position around.

The each contain the following topics:

- ConditionID: The identifier for the condition. This is the same as the filename, except that filenames further have the prefix "cond".
- EnvString: An ASCII representation of the maze, where "#" represent walls, and "g" represent passable grounds. Other symbols are possible, but where not used for the reported experiment.
- AlwaysVisibles: A list of positions which were always visible regardless of view radius. This was used in order to make potential goal locations always visible in conditions 2 and 3.
- ViewRadius: The number of blocks the agent was able to see around itself. A viewRadius of 0 means that the radius was not restricted and the entire maze was visible.
- TargetRadius: The distance measured in blocks an agent was able to see the color of a goal, which was only important for the second condition.
- Targets: A dictionary describing the potential targets or goal locations. The key specifies the location of the target and the value contains the targets color as well as the symbol that was to be displayed.
- Goal: A dictionary describing the actual goal for the condition. Only the type "Reach" was used in this experiment. The task descriptions varied between conditions slightly.
- StartPosition: The initial position of the agent.

1.2 Participant data

For each participant of the study, we collected all their actions within the maze. The folder *Participant_data* thus contains a folder for each participant. Within each folder, we recorded their data in individual files for each trail the participant finished. The files are named according to the condition of the trail and the first 24 lines are identical to the condition data description. Following an empty line, the behavior is recorded starting from line 26. Each line first contains a timestemp of when the action was performed. Afterwards, we record the way the action was performed (either via the keyboard –"Key"– or via the buttons displayed around the maze –"Button"). The last line should

always contain "Condition Finished" as this was recorded when the participant interacted with the goal tile.

2 Code

The code package contains the python code used to implement the different models mentioned in [1] as well as the code to evaluate the results.

The models used are implicitly implemented in `models/optimalModel.py`. For each of the four models considered, we implemented a `rate_episode` and a `rate_episode_global` method which computed the score for the last action or the entire episode, respectively.

The models all consider rational agents with a softmax policy. The value for an action was determined by the resulting distance from the considered goal which was computed via the A*-algorithm on the entire maze, or for the *World Uncertainty* condition under the free-space assumption. Both implementations can be found in `blockworld.py`

The reported experiments and their results are configured and run in the file `experiments.py`.

To run the code unzip the `src.zip` and the `data.zip` archives. Depending on where you unzip these folders you might need to adapt the path for the conditions in `playback.py` as well as the path at the bottom of `experiments.py`. Afterwards you can just use `python experiments.py` to run the script which should reproduce the experiment results.

3 Supplementary significance analysis

Due to the request of a reviewer of [1], we performed a statistical analysis using the unparameterized Wilcoxon signed-rank test for repeated measurements by relying on the corresponding `scipy.stats` function. `experiments.py` contains code to structure the results in a suitable way and perform the tests. Alternatively, one can export the results as csv files and perform other computations on them.

Note however, that all models produce deterministic outputs, thus using statistical tests on them is actually not really appropriate. A more direct comparison can be seen in Tables 1 to 4 where we report the percentage of runs one model outperformed (i.e. produced a lower or equal log likelihood) the other in the different conditions as well as how much worse it is in the other cases.

In Tables 6 to 9 you can see the statistics and p-values for all pairwise tests overall, the *No Uncertainty*, the *Goal Uncertainty* and the *World Uncertainty* conditions, respectively:

Table 1: Table reporting the percentage of runs the model on the left outperformed the model on the top in the overall condition. In brackets are the average difference (absolute mean and relative mean) in score when the model performs worse.

	No Assumption	True World and Goal Belief	True World Belief	True Goal Belief	Switching
No Assumption	-	38.87% (12.43/145.61%)	66.96% (1.24/3.86%)	44.11% (13.90/92.57%)	15.14% (12.62/120.21%)
True World and Goal Belief	61.14% (30.15/56.26%)	-	65.36% (29.42/56.13%)	52.55% (16.11/41.81%)	42.65% (23.95/56.51%)
True World Belief	33.04% (8.58/21.46%)	34.64% (17.45/175.62%)	-	40.32% (20.23/128.30%)	29.26% (21.06/21.46%)
True Goal Belief	55.90% (21.95/49.99%)	47.89% (10.44/87.94%)	59.68% (21.47/51.40%)	-	33.48% (17.75/91.67%)
Switching	84.86% (7.89/12.17%)	88.65% (0.92/2.08%)	95.78% (1.30/1.57%)	83.55% (2.34/17.66%)	-

Table 2: Table reporting the percentage of runs the model on the left outperformed the model on the top in the "No Uncertainty" condition. In brackets are the average difference (absolute mean and relative mean) in score when the model performs worse.

	No Assumption	True World and Goal Belief	True World Belief	True Goal Belief	Switching
No Assumption	-	3.91% (16.16/221.03%)	69.13% (0.53/2.09%)	32.17% (13.39/112.54%)	3.04% (16.59/222.14%)
True World and Goal Belief	96.09% (5.57/13.12%)	-	97.39% (5.65/14.62%)	80.44% (4.67/30.25%)	79.13% (3.74/15.48%)
True World Belief	30.87% (5.77/20.85%)	2.61% (19.80/268.29%)	-	29.13% (17.17/149.27%)	4.783% (20.89/279.02%)
True Goal Belief	67.83% (8.74/25.48%)	20.435% (12.51/141.72%)	70.87% (7.11/21.08%)	-	20.00% (12.53/141.16%)
Switching	96.96% (1.02/2.78%)	82.61% (0.19/0.56%)	98.26% (0.76/2.86%)	85.22% (1.62/19.86%)	-

Table 3: Table reporting the percentage of runs the model on the left outperformed the model on the top in the "Goal Uncertainty" condition. In brackets are the average difference (absolute mean and relative mean) in score when the model performs worse.

	No Assumption	True World and Goal Belief	True World Belief	True Goal Belief	Switching
No Assumption	-	81.94% (7.18/62.35%)	41.85% (1.75/5.17%)	87.23% (6.76/56.10%)	30.84% (3.44/21.75%)
True World and Goal Belief	18.06% (33.93/69.71%)	-	19.38% (32.95/68.02%)	57.71% (20.55/26.36%)	16.30% (32.04/67.10%)
True World Belief	58.15% (7.20/16.42%)	80.62% (10.60/84.43%)	-	85.02% (10.72/80.18%)	74.45% (8.74/69.92%)
True Goal Belief	12.78% (27.56/64.08%)	42.29% (9.28/28.26%)	14.98% (26.80/64.92%)	-	11.01% (26.32/68.57%)
Switching	69.16% (7.43/13.72%)	91.63% (2.72/5.84%)	89.87% (1.50/1.49%)	93.39% (2.44/6.93%)	-

Table 4: Table reporting the percentage of runs the model on the left outperformed (or performs equally well to) the model on the top in the "World Uncertainty" condition. In brackets are the average difference (absolute mean and relative mean) in score when the model performs worse.

	No Assumption	True World and Goal Belief	True World Belief	True Goal Belief	Switching
No Assumption	-	31.30% (8.59/61.72%)	89.57% (0.50/1.90%)	13.48% (15.34/82.23%)	11.74% (15.35/84.40%)
True World and Goal Belief	68.70% (23.46/26.91%)	-	78.70% (19.12/16.80%)	19.57% (16.59/52.64%)	32.17% (20.31/56.22%)
True World Belief	10.44% (11.38/24.25%)	21.30% (16.22/83.10%)	-	7.39% (24.09/119.94%)	9.13% (24.66/124.52%)
True Goal Belief	86.52% (17.71/18.53%)	80.87% (5.31/41.90%)	92.61% (17.48/17.32%)	-	69.13% (6.93/29.18%)
Switching	88.26% (10.88/10.60%)	91.74% (0.68/1.53%)	99.13% (1.60e-14/1.51e-14%)	72.17% (2.70/19.01%)	-

Table 5: Average surprise values and their standard deviations for the different models (rows) applied to the three conditions (columns); numbers in bold represent the lowest values (best).

Model	Overall	No Uncert.	Goal Uncert.	World Uncert.
No Assumption	43.6975 (41.98)	29.1360 (11.04)	48.8868 (38.29)	53.1374 (58.02)
True World and Goal Belief	47.8146 (59.76)	13.8305 (12.54)	75.3903 (56.71)	54.5828 (73.36)
True World Belief	49.0316 (44.56)	32.9640 (13.26)	50.8797 (39.45)	63.2752 (61.28)
True Goal Belief	45.6111 (51.67)	22.8657 (15.75)	72.0591 (47.27)	42.2535 (65.58)
Switching	34.1854 (46.29)	13.0823 (10.38)	48.7999 (40.96)	40.8646 (62.75)

Table 6: Wilcoxon signed-rank test statistics of the average performance of the different models over all conditions. p-values are reported in the brackets.

	No Assumption	True World and Goal Belief	True World Belief	True Goal Belief	Switching
No Assumption	-	111682 (0.2129)	40605 (3.09e-50)	206966 (3.97e-05)	27803 (1.53e-67)
True World and Goal Belief	111682 (0.2129)	-	90368 (9.22e-08)	114301.5 (0.6291)	6360 (4.60e-63)
True World Belief	40605 (3.09e-50)	90368 (9.22e-08)	-	93513 (2.17e-06)	1308 (1.21e-83)
True Goal Belief	206966 (3.97e-05)	114301.5 (0.6291)	93513 (2.17e-06)	-	16602 (2.03e-63)
Switching	27803 (1.53e-67)	6360 (4.60e-63)	1308 (1.21e-83)	16602 (2.03e-63)	-

Table 7: Wilcoxon signed-rank test of the average performance of the different models in the "No Uncertainty" condition. p-values are reported in the brackets.

	No Assumption	True World and Goal Belief	True World Belief	True Goal Belief	Switching
No Assumption	-	348 (1.55e-37)	3048 (4.02e-24)	6300 (4.77e-12)	60 (3.79e-39)
True World and Goal Belief	348 (1.55e-37)	-	204 (2.45e-38)	2765.5 (5.84e-25)	1915 (0.1720)
True World Belief	3048 (4.02e-24)	204 (2.45e-38)	-	4625 (1.03e-17)	20 (1.49e-38)
True Goal Belief	6300 (4.77e-12)	2765.5 (5.84e-25)	4625 (1.03e-17)	-	1368 (5.62e-30)
Switching	60 (3.79e-39)	1915 (0.1720)	20 (1.49e-38)	1368 (5.62e-30)	-

Table 8: Wilcoxon signed-rank test statistics of the average performance of the different models in the "Goal Uncertainty" condition. p-values are reported in the brackets.

	No Assumption	True World and Goal Belief	True World Belief	True Goal Belief	Switching
No Assumption	-	2179 (1.74e-27)	11604 (0.1777)	1223 (2.81e-32)	10207 (0.0058)
True World and Goal Belief	2179 (1.74e-27)	-	2841 (2.10e-24)	12546.5 (0.7768)	543.5 (2.91e-33)
True World Belief	11604 (0.1777)	2841 (2.10e-24)	-	2045 (3.91e-28)	638.5 (2.85e-07)
True Goal Belief	1223 (2.81e-32)	12546.5 (0.7768)	2045 (3.91e-28)	-	380 (1.89e-35)
Switching	10207 (0.0058)	543.5 (2.91e-33)	638.5 (2.85e-07)	380 (1.89e-35)	-

51

Table 9: Wilcoxon signed-rank test statistics of the average performance of the different models in the "World Uncertainty" condition. p-values are reported in the brackets.

	No Assumption	True World and Goal Belief	True World Belief	True Goal Belief	Switching
No Assumption	-	10853 (0.0162)	552 (2.06e-36)	3293 (4.67e-23)	2477 (1.06e-26)
True World and Goal Belief	10853 (0.0162)	-	5052 (3.72e-16)	3238.5 (4.47e-23)	173.5 (5.49e-28)
True World Belief	552 (2.06e-36)	5052 (3.72e-16)	-	1390 (5.43e-32)	9 (2.57e-36)
True Goal Belief	3293 (4.67e-23)	3238.5 (4.47e-23)	1390 (5.43e-32)	-	4963 (0.6082)
Switching	2477 (1.06e-26)	173.5 (5.49e-28)	9 (2.57e-36)	4963 (0.6082)	-

References

- [1] Jan Pöppel and Stefan Kopp. Satisficing Models of Bayesian Theory of Mind for Explaining Behavior of Differently Uncertain Agents. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, IFAAMAS*, 9 pages