Thinking Too Much: Pathology in Pathfinding

Mitja Luštrek (mitja.lustrek@ijs.si)
Jožef Stefan Institute, Department of Intelligent Systems
Ljubljana, Slovenia

Vadim Bulitko (bulitko@ualberta.ca)
University of Alberta, Department of Computing Science
Edmonton, Alberta, Canada

Introduction

- Real-time pathfinding ⇒ incomplete search methods ⇒ suboptimal actions
- Deeper lookahead believed to produce better actions
- Sometimes the opposite is true: pathology

Setting

- Pathfinding in grid world on maps from computer games
- Algorithm: LRTS [Bulitko & Lee 06]

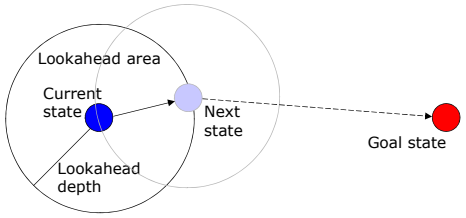

- Two types of experiments:
  - On-policy: entire path, heuristic updated
  - Off-policy: one move (from random state), heuristic not updated
- Error = probability of a suboptimal move
- Degree of pathology = number of lookahead depths where error is larger than at the previous depth

Pathology Observed

Degree of pathology	0	1	2	3	4	≥ 5
Pat. problems on-policy [%]	6.3	13.1	24.8	29.0	18.1	8.7
Pat. problems off-policy [%]	83.1	14.9	2.0	0.0	0.0	0.0

- A lot of pathology on-policy, little off-policy – why?

First Explanation

- LRTS steers the search to pathological areas
- Test: measure off-policy pathology in states visited on-policy ⇒ slightly more pathology than normally off-policy ⇒ explanation correct to a degree

Degree of pathology	0	1	2	3	4	≥ 5
Pathological problems [%]	76.8	13.8	5.7	2.3	1.0	0.4

Second Explanation

- Smaller lookahead depths benefit more from updates to the heuristic ⇒ ⇒ depths closer to each other ⇒ larger depths more likely worse than smaller
- First test: on-policy experiment, ignore updates when measuring error ⇒ ⇒ equal updates at all depths ⇒ less pathology than normally on-policy

Degree of pathology	0	1	2	3	4	≥ 5
Pathological problems [%]	29.6	13.8	5.7	2.3	1.0	0.4

- Second test: observe volume of updates to the heuristic ⇒ ⇒ larger volume at smaller depths

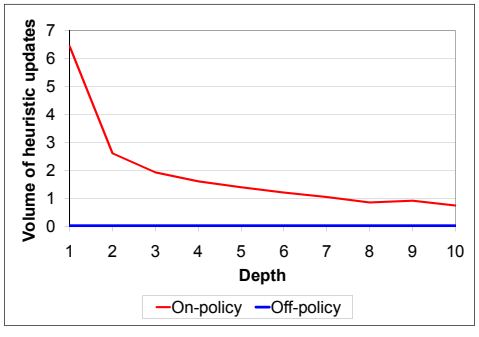

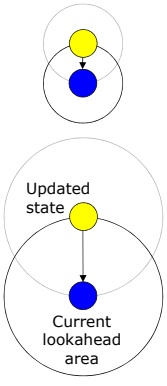

Third Explanation

- Fewer searches at larger lookahead depths ⇒ ⇒ depths closer to each other ⇒ larger depths more likely worse than smaller
- First test: on-policy experiment, search every move ⇒ less pathology than normally on-policy

Degree of pathology	0	1	2	3	4	≥ 5
Pathological problems [%]	60.2	13.8	5.7	2.3	1.0	0.4

- Second test: observe number of states generated per move when searching every move ⇒ steeper increase than normally ⇒ depths farther from each other

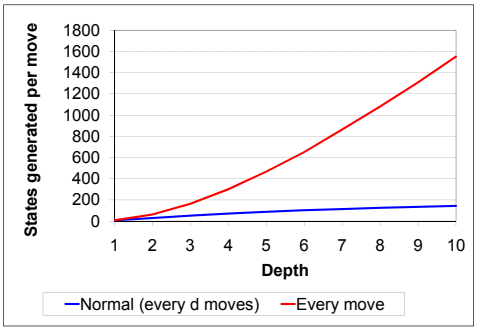

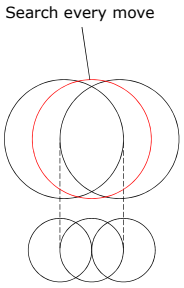

Fourth Explanation

- Optimistic (admissible) heuristic ⇒ lowest values likely have largest error ⇒ ⇒ state to move to has low value, hence large error; the more states one chooses from, the worse it gets ⇒ larger error at larger lookahead depths (more states to choose from) ⇒ pathology
- Test: use pessimistic heuristic instead
  lowest values likely have smallest error ⇒ state to move to has low value, hence small error; the more states one chooses from, the better it gets ⇒ small error at larger lookahead depths (more states to choose from) ⇒ less pathology

Degree of pathology	0	1	2	3	4	≥ 5
Pathological problems [%]	13.9	4.1	8.3	22.9	27.7	23.1

- Indeed less pathology, but more severe when it appears

Conclusions

- First explanation
  - 23.2 % instead of 16.9 % pathological states among those visited by LRTS
  - Minor factor
- Second explanation
  - 70.4 % instead of 93.7 % pathological states without updates to the heuristic (= equal benefit of updates at all lookahead depths)
  - Difficult to get around: increased intrinsic benefit of updates at large depths would reduce the number of revisits, which decreases the benefit of updates
- Third explanation
  - 34.7 % instead of 93.7 % pathological states when searching every move
  - But: *d* times more states generated per move (*d* = lookahead depth)
  - States generated per problem :
    - 4.5 times more states generated at *d* = 10 when searching every move
    - Not *d* times more, though
  - Practical solution: find some middle ground
- Fourth explanation
  - Inconclusive, further investigation needed