

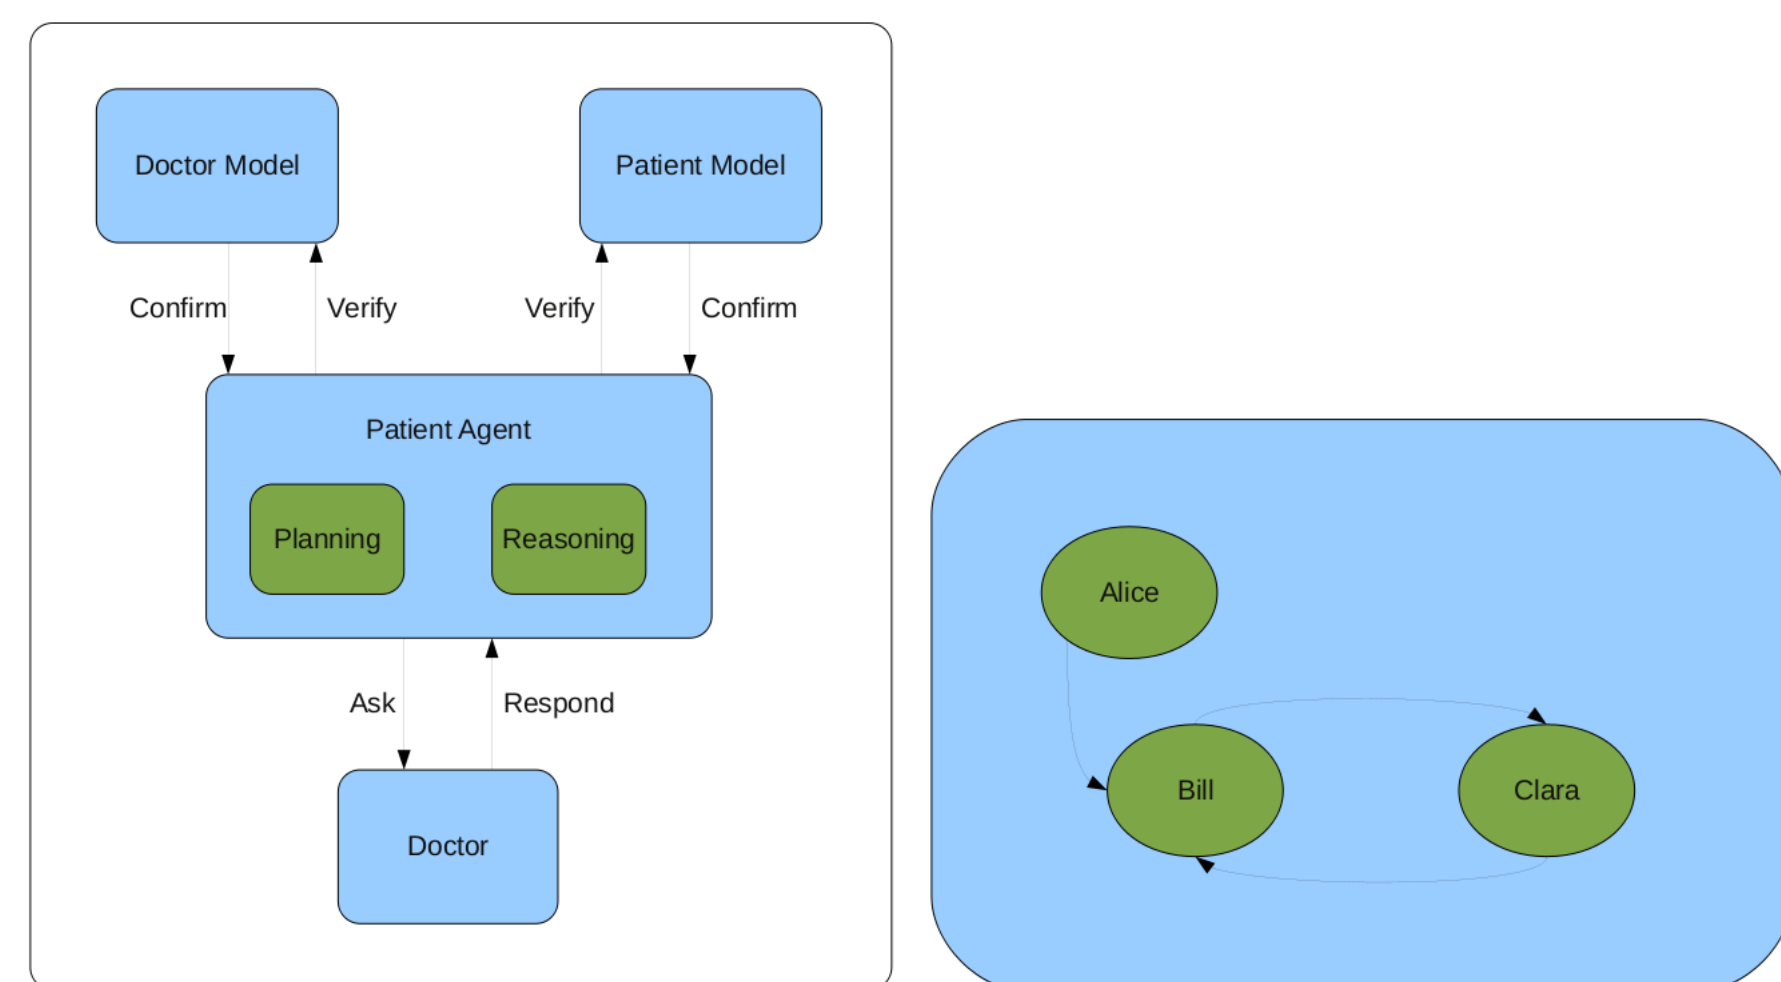
1. Introduction

- ▶ Novel approach to scheduling doctors, with a focus on mass casualty incidents.
- ▶ Based on multiagent resource allocation (MARA) and Transfer-of-Control strategies.
- ▶ Incorporation of user models and learning.

2. Patient scheduling as resource allocation

- ▶ Patient scheduling is a resource allocation problem, timeslots are the “resources”.
- ▶ Patients value times based on personal preferences and their condition.

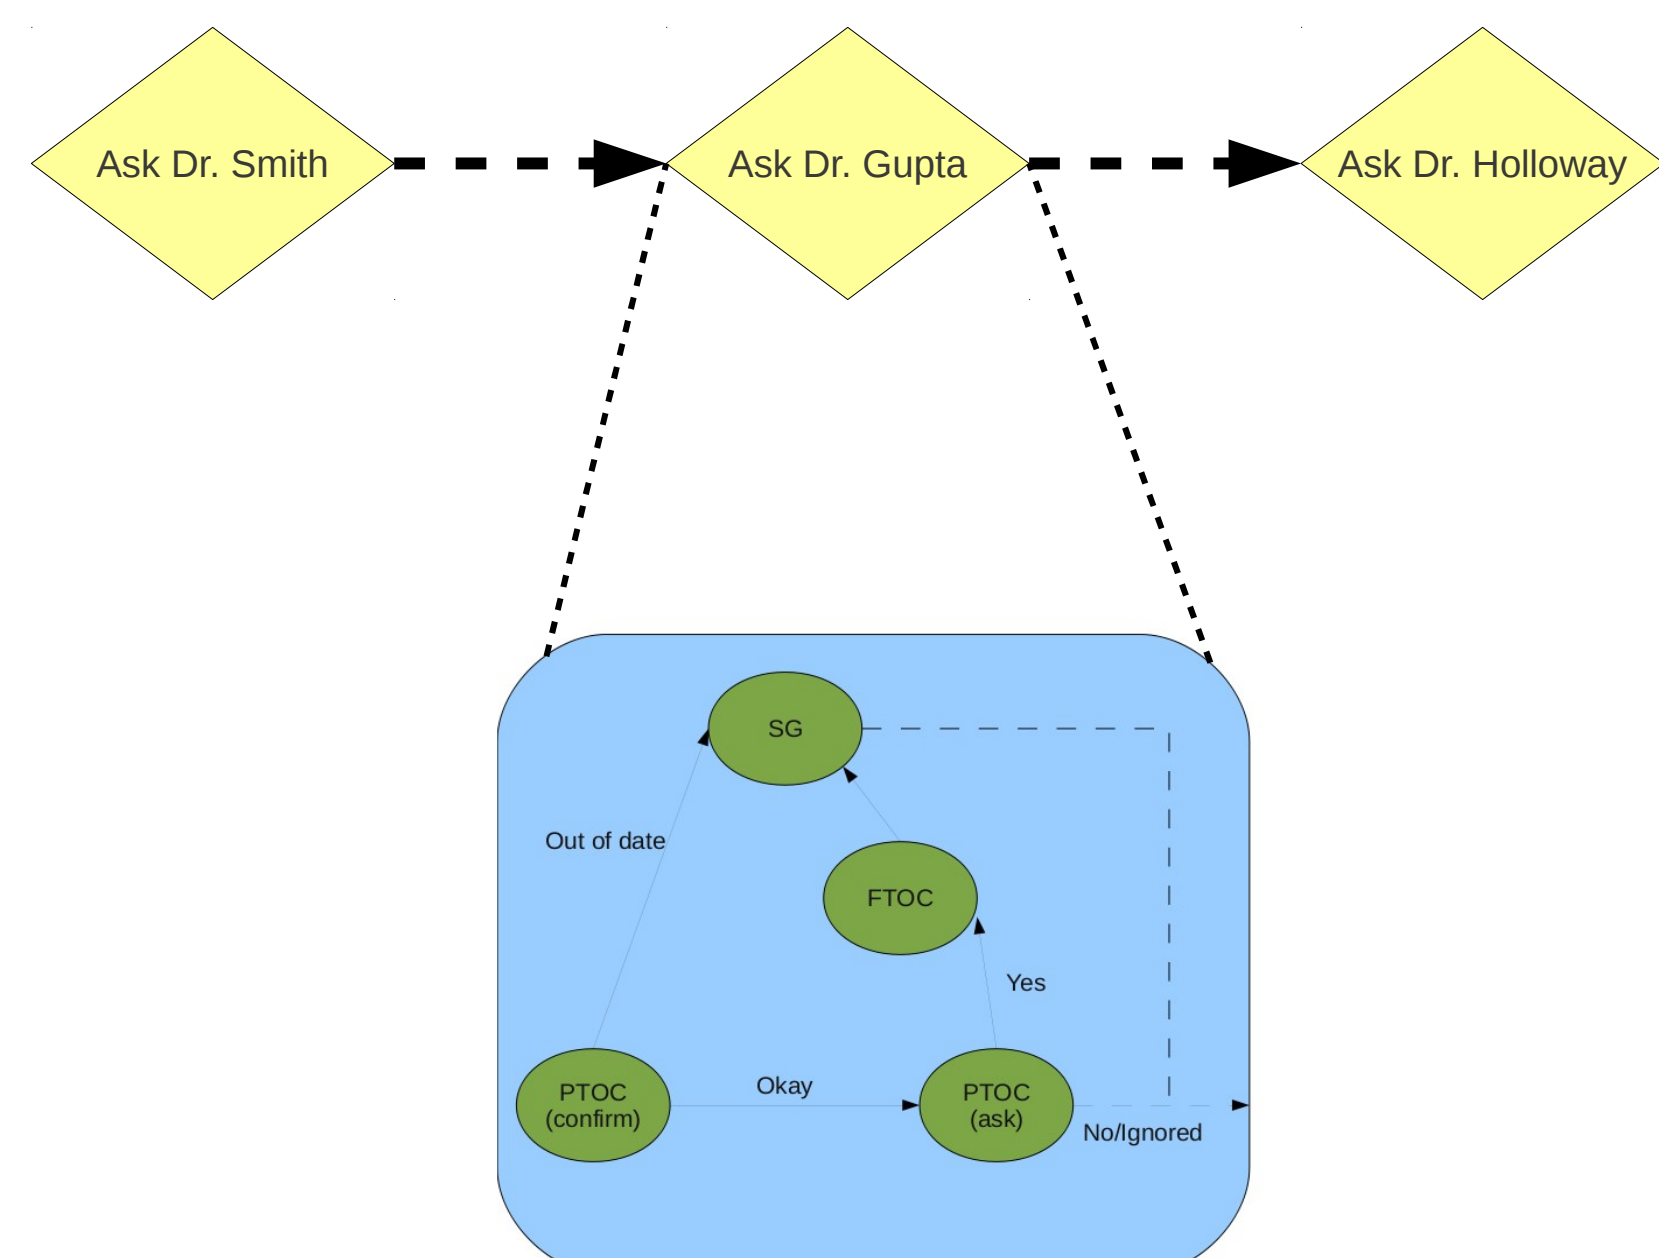
3. Multiagent resource allocation



- ▶ Each patient is assigned a software agent.
- ▶ Distributed negotiations optimize the schedule.
- ▶ Trouble estimating the cost of losing resources.

4. Planning

- ▶ Agents use pre-planned strategies in negotiation, called “Transfer-of-Control” (TOC) strategies[4][3].
- ▶ Strategies are made to maximize expected utility for a patient, and minimize the bothering doctors.
- ▶ An example TOC strategy:



- ▶ Full Transfer Of Control (FTOC): Request that the doctor take over treatment.
- ▶ Partial Transfer Of Control (PTOC): Ask the doctor a question or confirm that the plan is still valid.
- ▶ Strategy Generation (SG): Generate a new TOC strategy.

5. Estimating Costs using Transfer-Of-Control strategies

- ▶ The cost of preemption is the expected value of a contingency plan, less the benefit of the resource.
- ▶ This is a high quality estimate of the costs of changes in wait times and quality of care.
- ▶ Similar to micro-economic “Opportunity Costs”.

6. Learning

- ▶ Circumstances change with time, but agents can adapt by learning.
- ▶ Let $c(x)$ be an interaction cost, $t(x)$ be an interaction time.
- ▶ Both Costs estimates [4]:

$$BSF = \sum_{i \in I} c(i) \beta^{t(i)}$$

Where i is an interaction, β is the learning rate.

- ▶ Plan Length estimates:

$$k = \frac{1}{1 + \sum_{r \in R} c(r) \alpha^{t(r)}}$$

Where r is a regeneration, α is the learning rate.

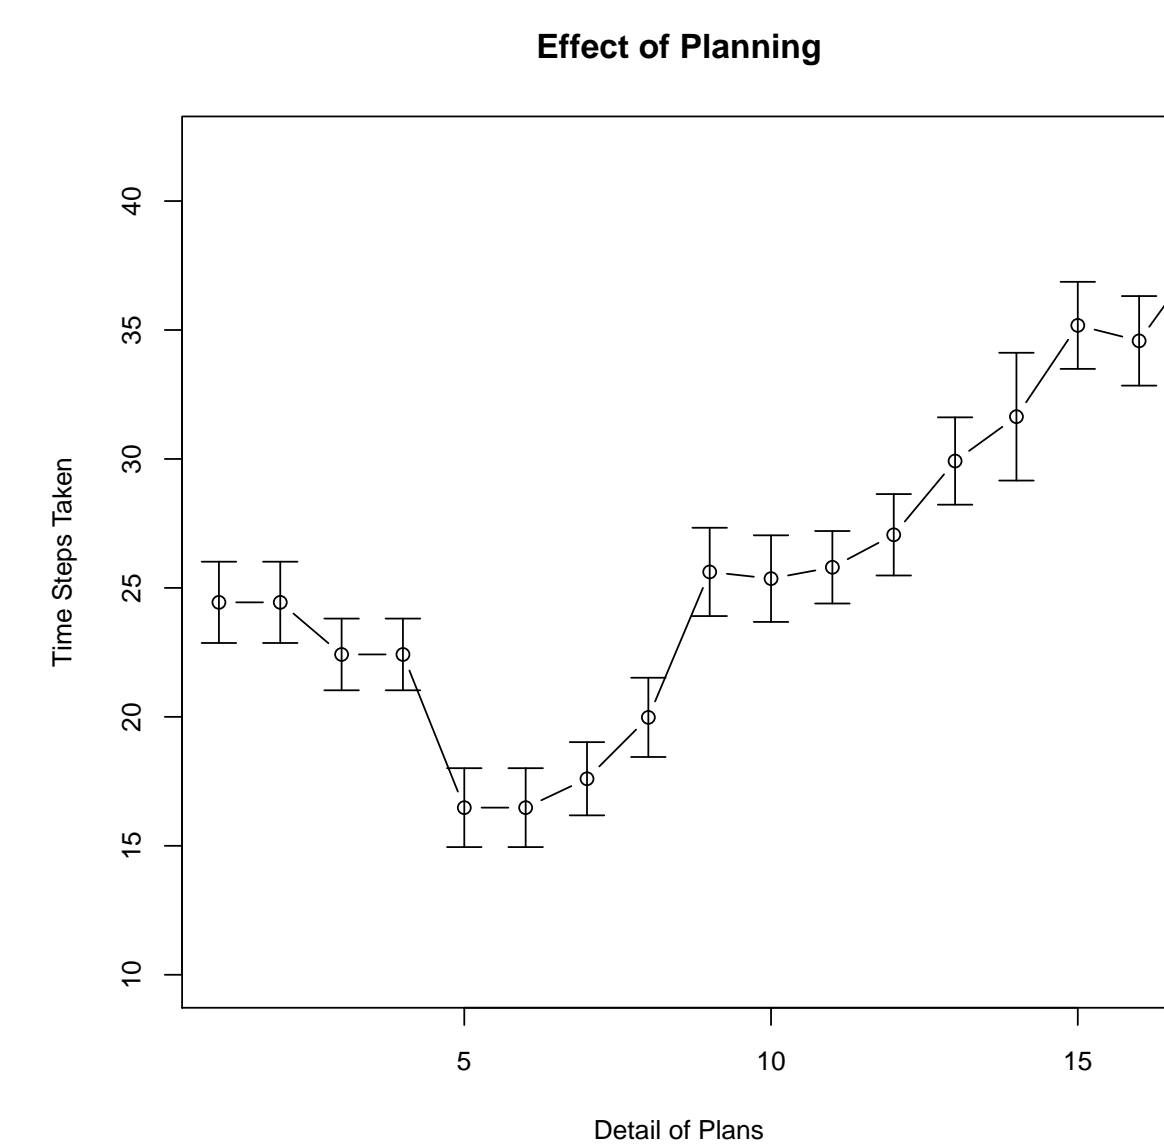
- ▶ Congestion estimates:

$$EU_{plan} = \frac{1}{1 + \sum_{p \in P} c(p) \gamma^{t(p)}}$$

Where p is a preemption, γ is the learning rate.

7. Effects of Learning

- ▶ Bother models improve resource utilization.
- ▶ Congestion and Plan length estimate prevent deadlock.
- ▶ Cost reduction from plan length optimization:



7. Benchmarking

- ▶ We simulated the system to measure its performance.
- ▶ The scenario is a mass casualty incident, where many patients arrive simultaneously at a hospital with few doctors.
- ▶ This is not intended to accurately model every detail of a real-world scenario, but to provide an initial characterization of the system.

8. Example Experiment: User Modeling

- ▶ Patients are modeled by a deterioration rate $D(c)$ and criticality (c).
- ▶ Scenario goal: Minimize the costs ($T(c)$), suffering ($S(c)$) and percentages of ‘problem patients’.
- ▶ Total cost incurred by a single patient between T_1 and T_2 is:

$$Cost(T_1, T_2) = \int_{T_1}^{T_2} S(c_t) + T(c_t).dt$$

with $D(C) = \frac{dc}{dt}$, so

$$c_t = c_{t_{init}} + \int_{t_{init}}^t D(c).dt$$

9. Example Experiment: Doctor Model

- ▶ Following [3] doctors are modeled with a type and degree of busyness.
- ▶ A bother model [4] tracks the impact of previous system interactions on doctor willingness to respond.
- ▶ Types capture the differences between, say, an intern and a heart specialist.

10. Example Experiment: Algorithm

```

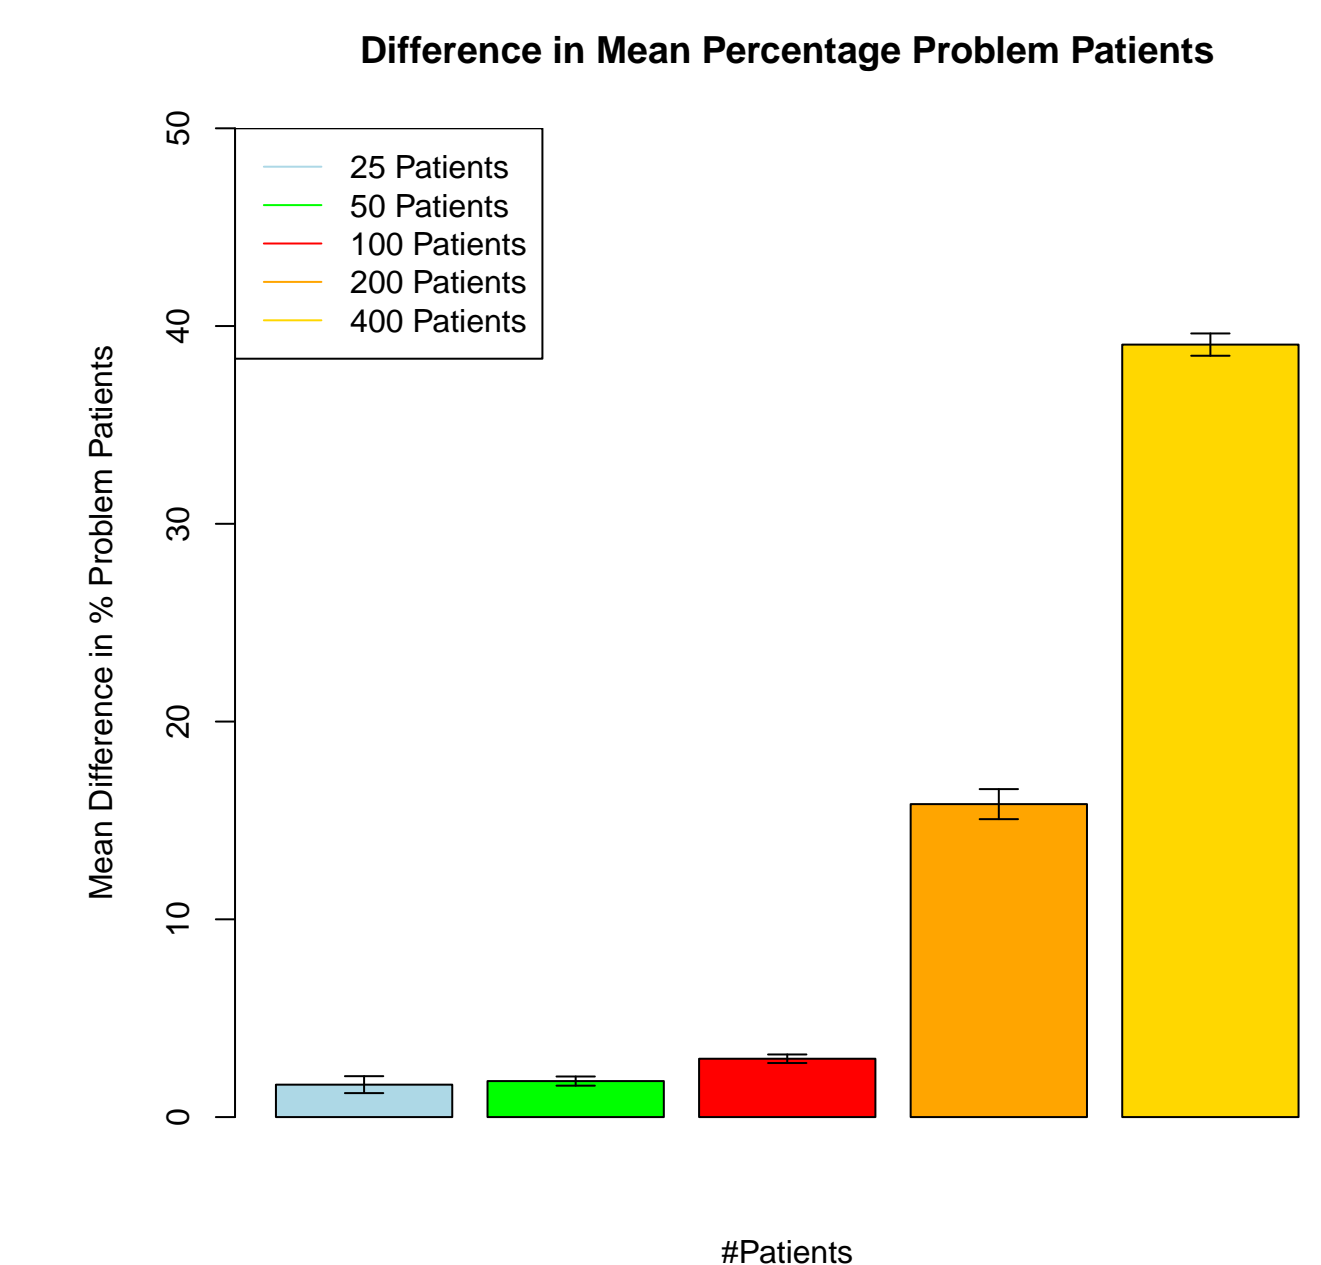
1 WHILE( there are still untreated patients )
2   FOREACH Agent A
3     //Let each agent take the next step in its TOC strategy
4     execute_plan(A)
5   ENDFOR
6
7   //Patients deteriorate based on their conditions,
8   //Doctors treat assigned patients
9   update_simulation()
10 ENDWHILE
11
12 //Subroutine for executing the next step in a plan.
13 SUB execute_plan(Agent A)
14   IF( A has no plan)
15     generate_plan(A)
16   ELSE
17     execute(A->plan->next()) //execute the next TOC world.
18   ENDIF
19 ENDSUB
    
```

11. Example Experiment: Strategy Generation

- ▶ Strategies are generated using a new dynamic programming approach.
- ▶ This approach requires improves upon those of prior authors ([3][4]), requiring only $O(2^n)$ steps instead of $O(n!)$.
- ▶ A linear time algorithm can also produce good approximations of the correct answer.

12. Results

- ▶ We compared our method to an implementation of the algorithm from [1].
- ▶ The systems were evaluated against 100 randomly generated sets of patients and doctors for each parameter setting.
- ▶ This graph shows the performance differences between the new algorithm and [1] in a low-critically scenario with 15 doctors.



13. Conclusions

- ▶ Learning agents adapt to changing circumstances, and provide considerable efficiency gains.
- ▶ Using contingency plans improves performance.
- ▶ More work to be done:
 - ▶ Further ablation studies
 - ▶ Comparisons with real-world data
 - ▶ Distributed implementation (e.g. on smart phones)

Bibliography

- [1] T. O. Paulussen, A. Ziller, A. Heinzl, A. Pokahr, L. Braubach, and W.o Lamersdorf, *Dynamic Patient Scheduling in Hospitals.*, Agent Technology in Business Applications, 2004
- [2] Hyunggu Jung, *Reasoning about Benefits and Costs of Interaction with Users in Real-time Decision Making Environments with Application to Healthcare Scenarios*, Master of Mathematics thesis, University of Waterloo, Waterloo, Ontario, 2010.
- [3] Robin Cohen, Hyunggu Jung, Michael Fleming, and Michael Y.K. Cheng, *A User Modeling Approach for Reasoning about Interaction Sensitive to Bother and Its Application to Hospital Decision Scenarios*, Special Issue on Personalization in e-Health, User Modeling and User-Adapted Interaction, January 2011.
- [4] Robin Cohen, Michael Y.K. Cheng and Michael W. Fleming, *Why bother about bother: Is it worth it to ask the user?*, AAAI05 Fall Symposium on Mixed-Initiative Problem-Solving Assistants, 2005