


How can we automatically discover subgoals?



Advantages and disadvantages?




Figure 4: Frequency  $f_i$  of States in Trajectories




Figure 6: Distance  $D_i$




Figure 7: FD Candidacy Metric  $c_i$

# Can anyone imagine

a domain in which subgoals are related to intermediate rewards  
instead of states? action sequences instead of states?

Does it make sense to learn options after learning a working flat policy?



Mean steps to goal in 2-room gridworld

