Compromise Strategies for Action Selection

Frederick L. Crabbe

Computer Science Department
United States Naval Academy

University of Pittsburgh Intelligent Systems Program

The Problem

- Autonomous agents (animal, robot, software) pick actions to take
- Many approaches
  - Solve the problem optimally
  - Treat problem heuristically
  - Pros and cons of each
- Multiple conflicting goals introduce tough problems for optimal approach
  - Difficult to express
  - Difficult to compute
- Multiple conflicting goals introduce tough problems for heuristic approach
  - Which goal does the agent pursue?
  - How can they (should they) be combined?
- This talk is on one such combination technique: compromise

The Message

- Compromise behavior (seletcting actions that compromise between goals) is an influential concept in many areas of agents research
- Experiments here show it less beneficial than predicted
  - Infinite variations possible...
  - Experiments are based on scenarios compromise advocates say should work
  - Something's wrong with the currently accepted hypothesis
- We propose an alternate hypothesis
  - The level the decision is being made at is key to whether it is helpful
  - Compromise is more useful at higher level of decision making.

Outline

1. Introduction

Outline

1. Introduction
2. History
   - Computational
   - Biological

Outline

1. Introduction
2. History
   - Computational
   - Biological
3. Prescriptive Action Selection
   - Formulation
   - Experiments

Outline

1. Introduction
2. History
   - Computational
   - Biological
3. Prescriptive Action Selection
   - Formulation
   - Experiments
4. Proscriptive Action Selection
   - Formulation
   - Experiments

Outline

1. Introduction
2. History
   - Computational
   - Biological
3. Prescriptive Action Selection
   - Formulation
   - Experiments
4. Proscriptive Action Selection
   - Formulation
   - Experiments
5. What does it mean?
   - A new hypothesis

Outline

1. Introduction
2. History
   - Computational
   - Biological
3. Prescriptive Action Selection
   - Formulation
   - Experiments
4. Proscriptive Action Selection
   - Formulation
   - Experiments
5. What does it mean?
   - A new hypothesis
6. Future Work and Conclusion
   - Future Work
   - Conclusion

Traditional Planning

- Action selection problem = single goal in search space
  - Can have multiple parts: *Have(robot,medicine003)* $\wedge$ *In(robot,room342)*
  - Cannot be conflicting.
- Find shortest path, next action is 1st step on path
- Calculating this: hard
- Relax optimality? hard
- Still can’t handle the multiple conflicting goals

Outline

1 Introduction
2 History
- Computational
- Biological

3 Prescriptive Action Selection
- Formulation
- Experiments

4 Proscriptive Action Selection
- Formulation
- Experiments

5 What does it mean?
- A new hypothesis

6 Future Work and Conclusion
- Future Work
- Conclusion

Comparative Psychology

- Branch of animal psychology
- Derived from traditions of behaviorism
- Experiments outside of natural environment: maze, skinner box
- All animal drives not being tested are met by experimenters
- Designed to isolate matters in question
- Focus on reasoning and learning

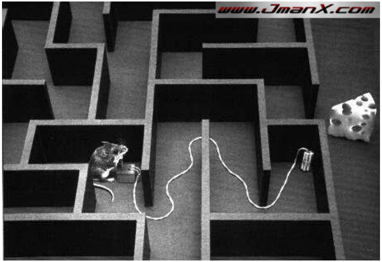

Comparative Psychology

- Branch of animal psychology
- Derived from traditions of behaviorism
- Experiments outside of natural environment: maze, skinner box
- All animal drives not being tested are met by experimenters
- Designed to isolate matters in question
- Focus on reasoning and learning

Ethology

- A different perspective
  - Observe animals in natural surroundings
  - Performing natural tasks
  - Often with multiple conflicting goals
- Fixed Action Patterns
  - Animals often react to external stimuli with hard-coded behaviors
  - What happens when multiple FAPs are active?
  - A focus of ethology research

Ethology

- A different perspective
  - Observe animals in natural surroundings
  - Performing natural tasks
  - Often with multiple conflicting goals
- Fixed Action Patterns
  - Animals often react to external stimuli with hard-coded behaviors
  - What happens when multiple FAPs are active?
  - A focus of ethology research

Possible FAP conflict strategies

- Pick one
- Intention movements
- Alternation
- Ambivalent behavior
- Common Components
- Compromise behavior
- Autonomic responses
- Displacement
- Redirection
- Regression
- Immobility
- Aggression

Traditional Planning

- Action selection problem = single goal in search space
  - Can have multiple parts: *Have(robot,medicine003)* $\wedge$ *In(robot,room342)*
  - Cannot be conflicting.
- Find shortest path, next action is 1st step on path
- Calculating this: hard
- Relax optimality? hard
- Still can't handle the multiple conflicting goals

The Behavior Based Solution

- Ethologically inspired:
  - Divide problem into FAP-like "behaviors"
  - Each dedicated to solving individual goals
  - Recombine recommendations, somehow
  - How do we recombine?
- Use the ethology list!
  - Already well studied
  - Many make intuitive sense
  - Seen in nature = good idea?
  - Need to pick and choose the good ones

Compromise Actions

- Recombination must be able to exhibit compromise behavior

Tyrell's rule 12:

[The combination mechanism must] be able to choose actions that, while not the best choice for any one sub-problem alone, are best when all sub-problems are considered simultaneously.

- Why? Council of ministers analogy
- Issues
  - Compromize can be costly (in computation AND design)
  - Actual benefit unknown

Outline

1 Introduction
2 History
- Computational
- Biological

3 Prescriptive Action Selection
- Formulation
- Experiments

4 Proscriptive Action Selection
- Formulation
- Experiments

5 What does it mean?
- A new hypothesis

6 Future Work and Conclusion
- Future Work
- Conclusion

Prescriptive Goals

- Low level, two prescriptive goal scenario
  - 2 goals to move to 2 targets
  - targets can dissappear
  - Will bet-hedging compromise be a good idea?
- Seen in nature?
  - Mating behavior Frogs (Leptodactylus ocellatus)
  - Hunting behavior of cheetahs (Acinonyx jubatus)

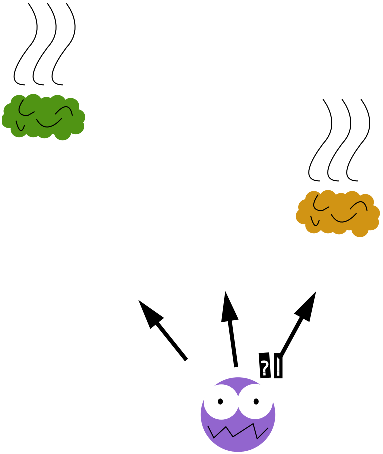

Prescriptive Goals

- Low level, two prescriptive goal scenario
  - 2 goals to move to 2 targets
  - targets can dissappear
  - Will bet-hedging compromise be a good idea?
- Seen in nature?
  - Mating behavior Frogs (Leptodactylus ocellatus)
  - Hunting behavior of cheetahs (Acinonyx jubatus)

Prescriptive Goals

- Low level, two prescriptive goal scenario
  - 2 goals to move to 2 targets
  - targets can dissappear
  - Will bet-hedging compromise be a good idea?
- Seen in nature?
  - Mating behavior Frogs (Leptodactylus ocellatus)
  - Hunting behavior of cheetahs (Acinonyx jubatus)

Modeling Approach

- Simulated environment
- Action detail level
  - Too detailed: “move left leg”
  - Too vague: “go to target”
  - Move one unit at angle θ
- Environment contains 2 stationary targets, can disappear w/ probability $1 - p$
- Measurement: utility theory

$$EU(A_i|S_j) = \sum_{S_o \in O} P(S_o|A_i, S_j) U_h(S_o)$$

$$U_h(S) = U(S) + \max_{A_i \in A} EU(A_i|S_j)$$

Formal Model

- Applying the EU equations to our situation, we get:

$$
\begin{aligned}
EU(A_i|t_a, t_b, \lambda) =& p^2 EU(A_\theta|t_a, t_b, \lambda')+\\
& p(1-p)EU(A_\theta|t_a, \lambda')+\\
& p(1-p)EU(A_\theta|t_b, \lambda'),\\
EU(A_i|t_a, \lambda') =& G_a p^{\overline{\lambda' t_a}}, \text{ and,}\\
EU(A_i|t_b, \lambda') =& G_b p^{\overline{\lambda' t_b}}
\end{aligned}
$$

- Can be solved using dynamic programming

Experimental Set-Up

- Select random location for 2 targets
- Select random goal values
- Select random p
- Run for 50,000 scenarios
- Calculate optimal policy
- Compare against non-compromise:
  - Closest (C)
  - Maximum Utility (MU)
  - Maximum Expected Utility (MEU)
- Compare against compromise strategies:
  - Forces (F)
  - Signal Gradient (SG)
  - Exponentially Weakening Forces (EWF)

Example

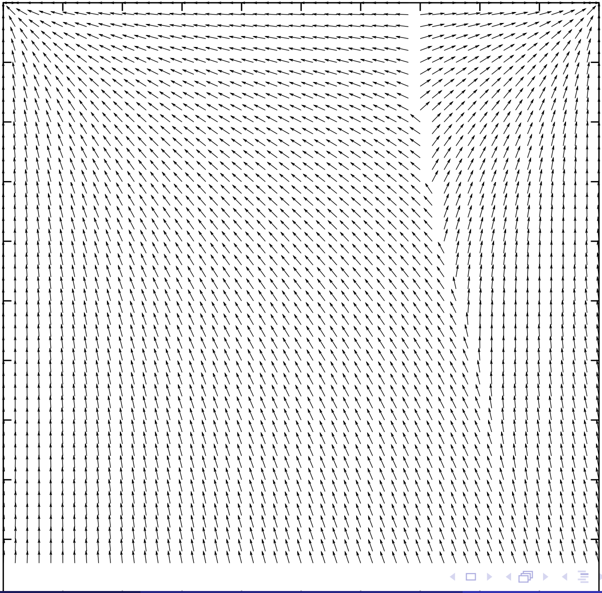

Prescriptive Results

Comparing non-compromise strategies to each other

	MU	C	MEU
% over MU	0.0	9.35	15.31
% over C	-4.13	0.0	12.62
% over MEU	-8.49	-5.96	0.0

Comparing compromise strategies to MEU

	F	SG	EWF	Optimal
avg	-4.07%	-2.79%	-2.47%	1.12%
best	4.84%	4.82 %	20.56 %	22.73%

Summary

- Standard compromise strategies worse than clever non-compromise
- Optimal only barely better than non-compromise

Extra bonus conclusion:

animals that exhibit apparent compromise in the 2 prescriptive goal case are either using some unknown strategy or are doing so for some other reason.

Outline

1. Introduction
2. History
   - Computational
   - Biological
3. Prescriptive Action Selection
   - Formulation
   - Experiments
4. Proscriptive Action Selection
   - Formulation
   - Experiments
5. What does it mean?
   - A new hypothesis
6. Future Work and Conclusion
   - Future Work
   - Conclusion

Proscriptive Goals

- Maybe the previous scenario wasn't where compromise shines
- Compromise work better with proscriptive goals?
    - Proscriptive goal is a goal to **not** do something
    - Such as, don't go near the predator
- Maybe prescriptive goal and a proscriptive one
- Move to food?
- Away from predator?
- Somewhere else?

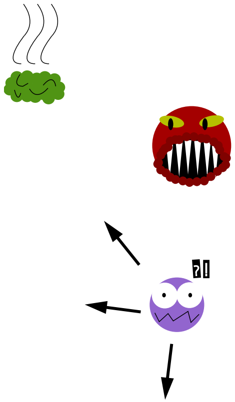

Proscriptive Goals

- Maybe the previous scenario wasn't where compromise shines
- Compromise work better with proscriptive goals?
  - Proscriptive goal is a goal to **not** do something
  - Such as, don't go near the predator
- Maybe prescriptive goal and a proscriptive one
- Move to food?
- Away from predator?
- Somewhere else?

Proscriptive Goals

- Maybe the previous scenario wasn't where compromise shines
- Compromise work better with proscriptive goals?
- Maybe prescriptive goal and a proscriptive one
- Move to food? Away from predator? Somewhere else?

Tyrell:

"It is obviously preferable to combine this demand [to flee the hazard] with a preference to head toward food, if the two don't clash, rather than to head diametrically away from the hazard because the only system being considered is that of avoid hazard"

Formal Model

- Applying the EU equations to our situation, we get:

$$
\begin{aligned}
EU(O|t,d,\lambda) =& p_t p_d p_n(\lambda) EU(O|t,d,\lambda') + \\
& p_t(1-p_d) EU(O|t,\lambda') + \\
& p_d(1-p_n(\lambda)) G_d + \\
& (1-p_t) p_d p_n(\lambda) EU(O|d,\lambda') \\
EU(O|t,\lambda) =& G_t p^{\overline{\lambda,t}}, \\
EU(O|d,\lambda) =& p_n(\lambda') p_d EU(O|d,\lambda') + \\
& (1-p_n(\lambda')) G_d.
\end{aligned}
$$

- Which can be calculated using dynamic programming

Experimental Details

- Location of target: $(50, 90)$; Utility: 100
- Location of danger: $(60, 50)$; Utility: -100
- p_d in range $[0.5; 1)$, p_t in range $[0.95; 1)$
- Probability distributions to strike probability
  - Linear A: $p_n(d) = 0.04d + 0.2$ when $d \leq 20$, 1 otherwise
  - Linear B: $p_n(d) = 0.005d + .9$ when $d \leq 20$, 1 otherwise
  - Quadratic: $p_n(d) = d^2/400$ when $d \leq 20$, 1 otherwise
  - Sigmoid: $p_n(d) = 1/(1 + 1.8^{10-d})$ everywhere
- Generated 2000 scenarios
- Action Selection mechanisms
  - Optimal
  - MEU- Go directly to target
  - Active Goal- Act based on goal currently active
  - Skirt- Move toward the target, but skirt around danger zone
- Examined EU of 4 AS strategies at 200 locations per scenario

Predictions

- Compromise seen most strongly inside danger zone, with danger to one side of agent
- More compromise for linear B
- Compromise around the edges for Sigmoid and Quadratic
- More compromise with low p_d
- More compromise with low p_t

- p_t **high,** p_d **high, and** $p_n(d)$ **is Linear A:**

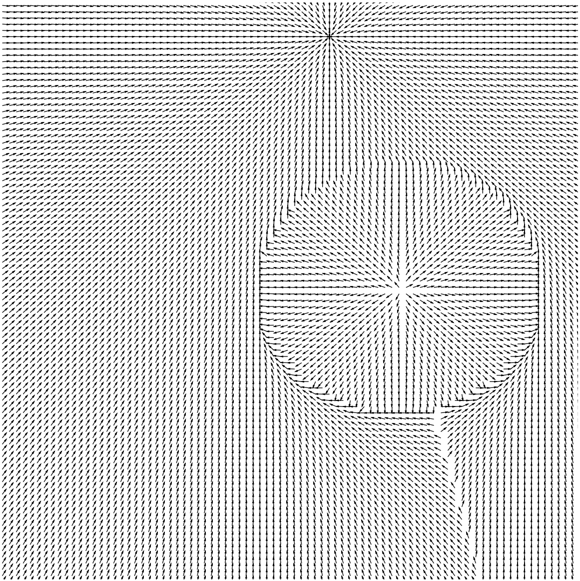

- p_t **high,** p_d **high, and** $p_n(d)$ **is Linear B:**

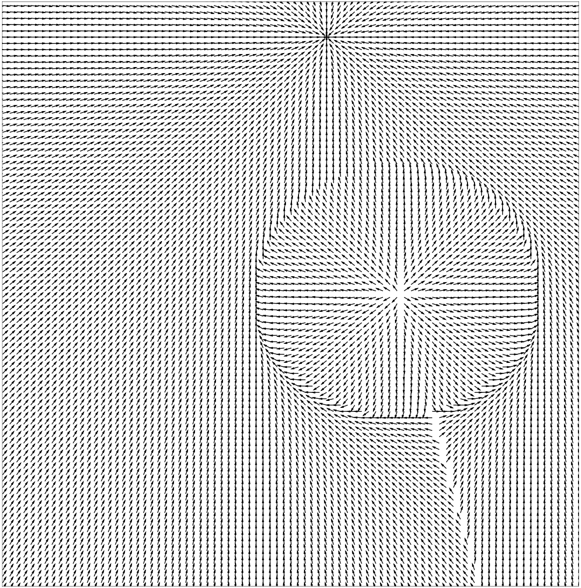

- p_t **high,** p_d **low, and** $p_n(d)$ **is Linear A:**

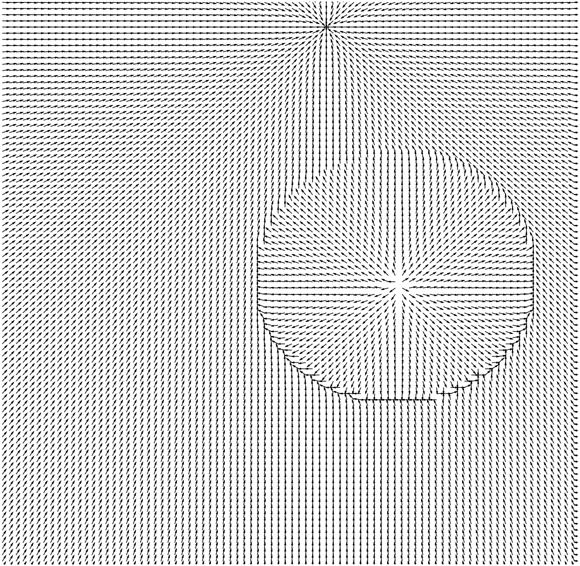

- p_t **low,** p_d **low, and** $p_n(d)$ **is Linear A:**

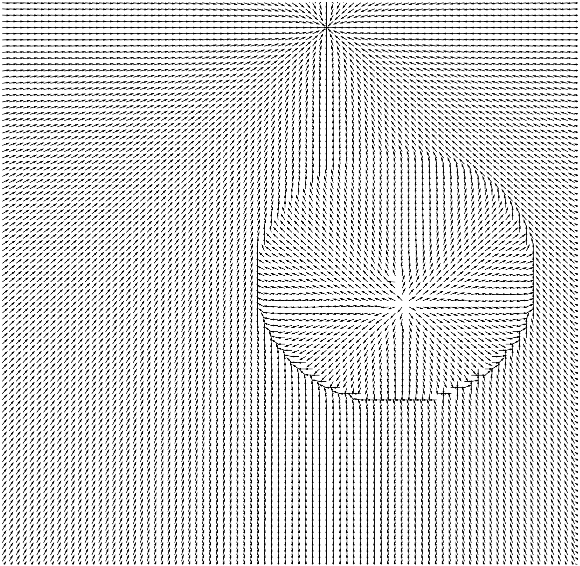

- p_t **low,** p_d **low, and** $p_n(d)$ **is Linear B:**

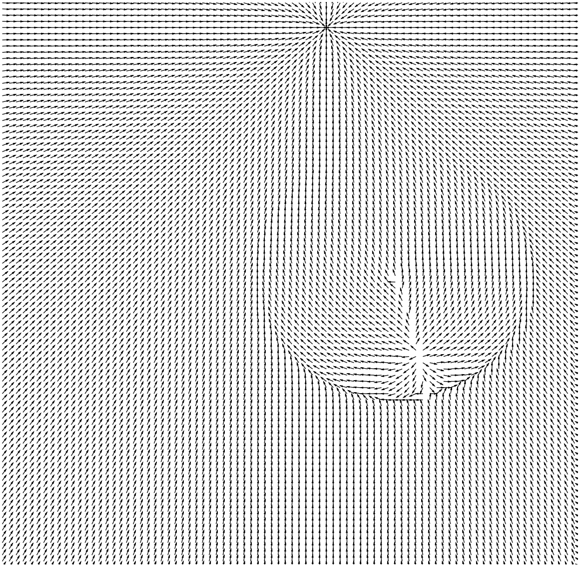

- p_t **high,** p_d **high, and** $p_n(d)$ **is Sigmoid:**

Quantitative Results

scenario	optimal over active goal	optimal over skirt	skirt over active goal
all	29.6%	0.1%	29.1%
opposite	64.9%	0.2%	63.3%
danger zone	26.2%	0.01%	26.1%

Summary

- As predicted
  - More compromise for linear B
  - Compromise at edges for Sigmoid and Exponential
  - Optimal compromise out performs Active Goal.
- Not predicted
  - Compromise not seen inside danger zone at all in many cases.
  - Compromise behind danger zone with high p_d.
  - p_t has less effect than p_d.
  - Optimal compromise does not out perform skirt.

Outline

1. Introduction
2. History
   - Computational
   - Biological
3. Prescriptive Action Selection
   - Formulation
   - Experiments
4. Proscriptive Action Selection
   - Formulation
   - Experiments
5. What does it mean?
   - A new hypothesis
6. Future Work and Conclusion
   - Future Work
   - Conclusion

Why do we care?

- Why not use optimal for our agents
  -
  -
- Why not use a faster compromise strategy?
  -
  -

Why do we care?

- Why not use optimal for our agents
  - Too slow
  -
- Why not use a faster compromise strategy?
  -
  -

Why do we care?

- Why not use optimal for our agents
  - Too slow
  - Small benefit
- Why not use a faster compromise strategy?
  -
  -

Why do we care?

- Why not use optimal for our agents
  - Too slow
  - Small benefit
- Why not use a faster compromise strategy?
  - Complicates system
  -

Why do we care?

- Why not use optimal for our agents
  - Too slow
  - Small benefit
- Why not use a faster compromise strategy?
  - Complicates system
  - Small benefit

Blending vs. Voting

- Compromise in experiments here resembles “blending” of actions
  - Matches the descriptions in ethology literature
  - Result is similar to the recommended actions of the two subgoals
- Compromise often justified as voting scheme:
  - Subgoal votes for top *n* actions from finite set
  - Action with most votes selected
  - Resulting actions different from best for each subgoal
- Confusion from equivocation on definition of compromise action
  - high vs. low level action
  - high level actions: small, discrete set; amenable to voting
  - low level actions: continuous, infinite set; result in blending

Compromise Behavior Hypothesis

1. Low level compromise action less useful than high level compromise
2. Higher the decision level, the more useful is compromise
    1. At low levels, compromise actions similar to non-compromise actions
    2. At high levels, compromise actions can be very different from non-compromise actions
    3. In complex environments, optimal or even very good non-optimal low-level actions are prohibitively difficult to calculate

You want to compromise in the selection of the high-level

Outline

1. Introduction
2. History
   - Computational
   - Biological
3. Prescriptive Action Selection
   - Formulation
   - Experiments
4. Proscriptive Action Selection
   - Formulation
   - Experiments
5. What does it mean?
   - A new hypothesis
6. Future Work and Conclusion
   - Future Work
   - Conclusion

Future Work

- Test against non-optimal compromise behaviors
- Test the compromise behavior hypothesis

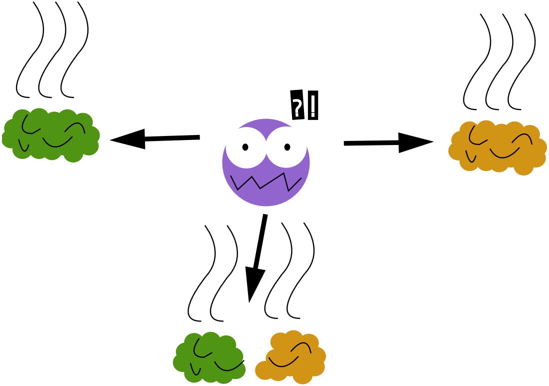

Conclusion

- In prescriptive goal scenarios
  - Optimal compromise marginally useful
  - Sub-optimal compromise harmful
- In proscriptive goal scenarios
  - Optimal compromise behavior is different from what we expected
  - Less beneficial than expected, and only in some situations
- Equivocation on definition of compromise action
  - high level vs low level actions
- Compromise Behavior hypothesis may explain what is going on