Active Behavior Recognition in Beyond Visual Range Air Combat

Ron Alford[1] Hayley Borck[2] Justin Karneeb[3] David W. Aha[4]

[1]ASEE/NRL Postdoctoral Fellow
ronald.alford.ctr@nrl.navy.mil

[2]Knexus Research Corporation
hayley.borck@knexusresearch.com

[3]Knexus Research Corporation
justin.karneeb@knexusresearch.com

[4]U.S. Naval Research Laboratory
david.aha@nrl.navy.mil

May 31st, 2015

Beyond Visual Range Air Combat

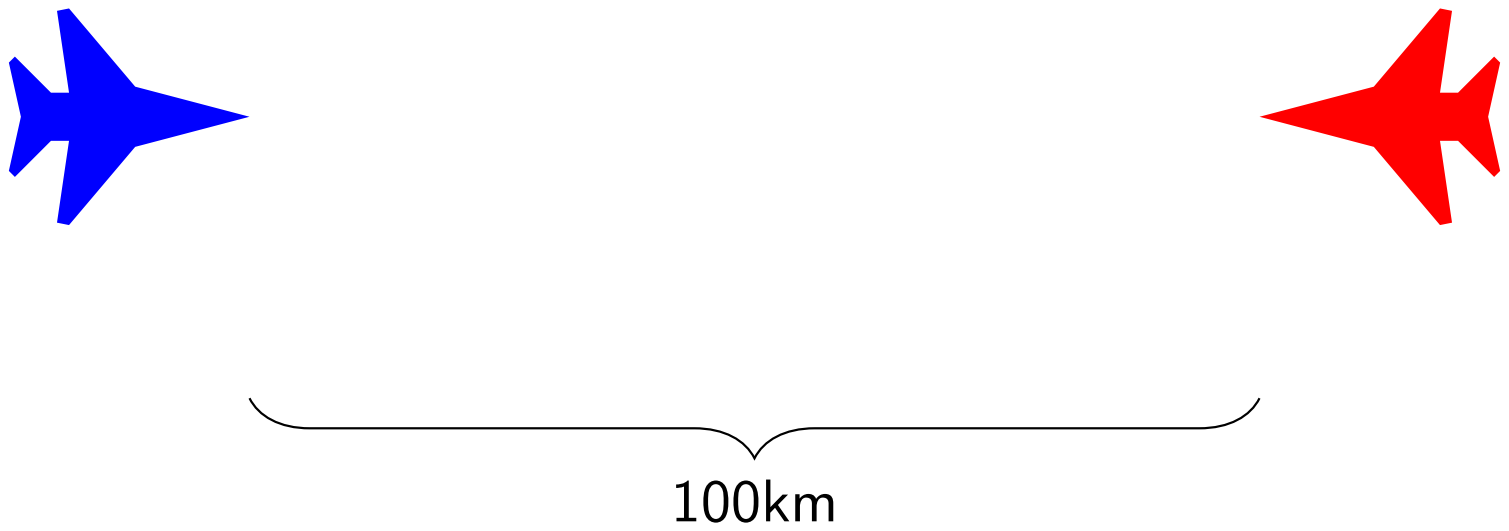

- $300 million planes

Beyond Visual Range Air Combat

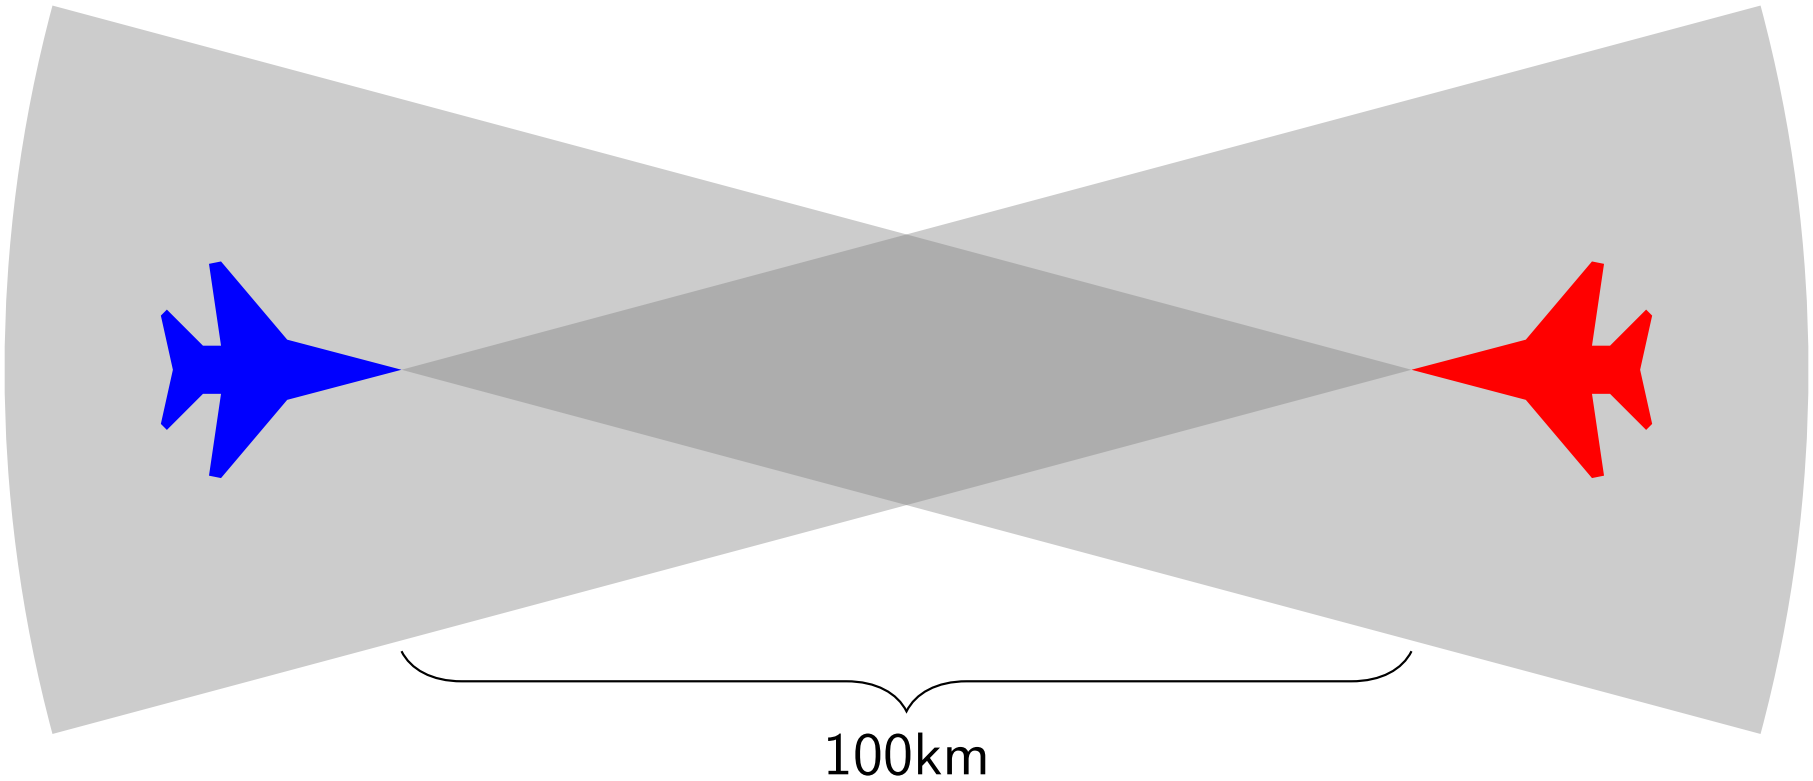

- $300 million planes
- $10 million sensor package

Beyond Visual Range Air Combat

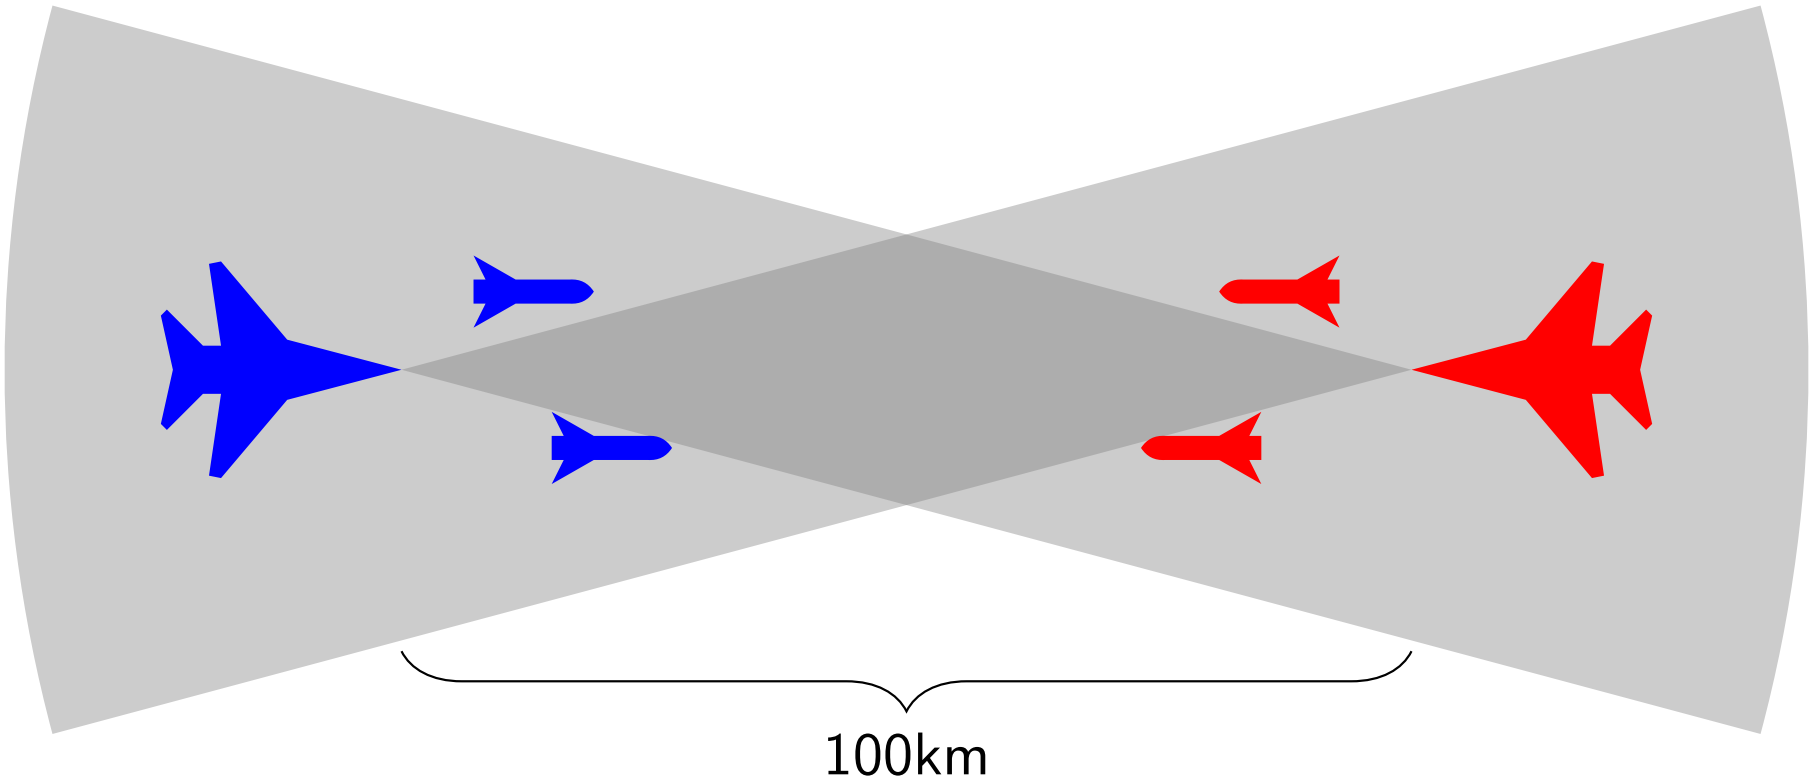

- $300 million planes
- $10 million sensor package
- $1 million $ missiles

Adding UAV wingmen to the mix

The Promise:

- More platforms per pilot
  - Better strategies
- Reduced pilot risk
- Retain (most) human judgment

The Caveat:

- Pilot is already cognitively burdened
- UAV needs to respond (or act) intelligently

Source: Dassault

Some obstacles to intelligent behavior

- Partial-observability
- Continuous action space
- Multi-agent (non-zero-sum)

Some obstacles to intelligent behavior

- ~~Partial-observability~~
  Full-observability
- ~~Continuous action space~~
  Discrete action space
- ~~Multi-agent (non-zero-sum)~~
  Single-agent with fixed (unknown) opponent

Behavior Recognition (Assumptions)

Assumptions:

- Finite set of predictive agent models
  - Used in training recognizer
  - Used to predict future states
- Agents use fixed polices
  - React to history of observations
  - Not rational nor optimal

Behavior Recognition (generic):

- Inputs:
  - Agent models
  - History of observations
- Output: A probability distribution over the models.

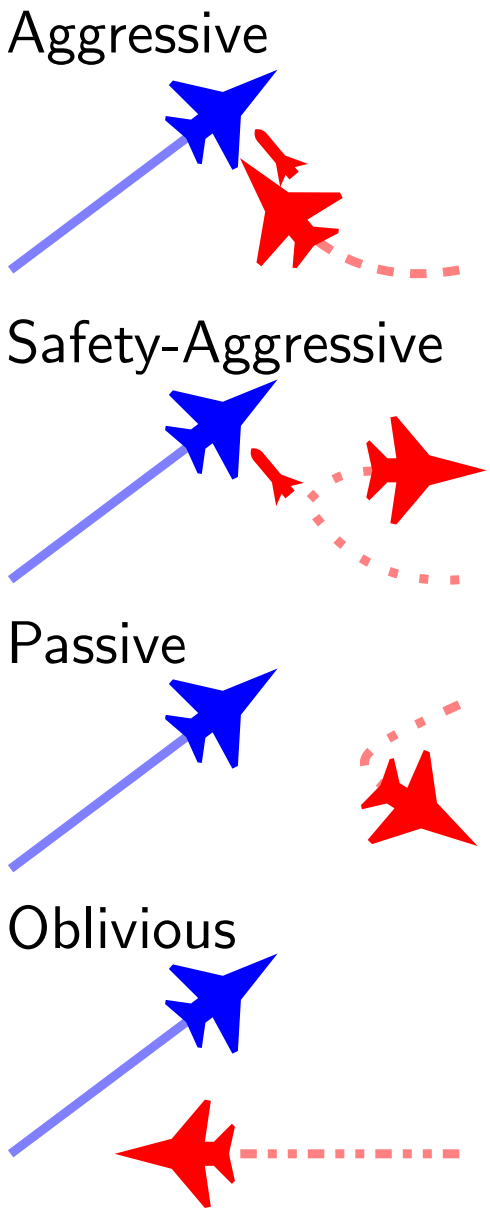

Acting depends on behavior recognition

Almost all actions in air combat are dependent (or relative) to other agents.

Safety-Aggressive vs. Aggressive

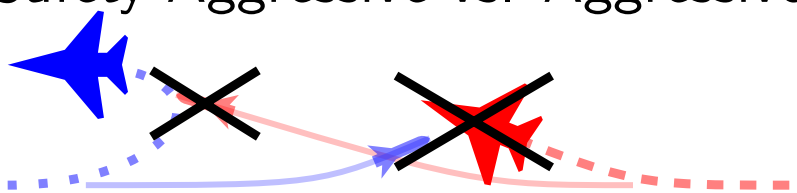

Safety-Aggressive vs. Oblivious

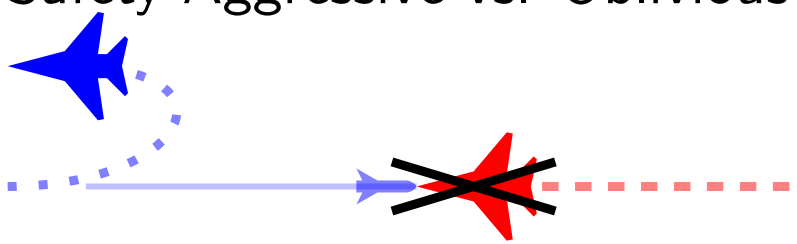

Safety-Aggressive vs. Safety-Aggressive

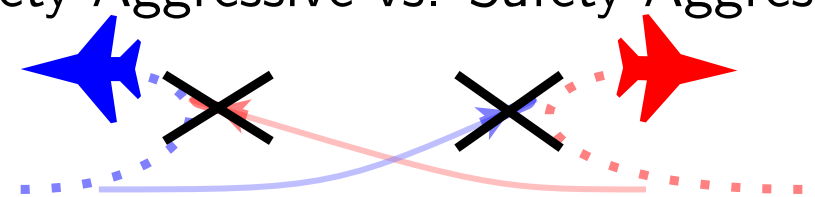

Behavior recognition depends on acting

Our actions determine what we observe.

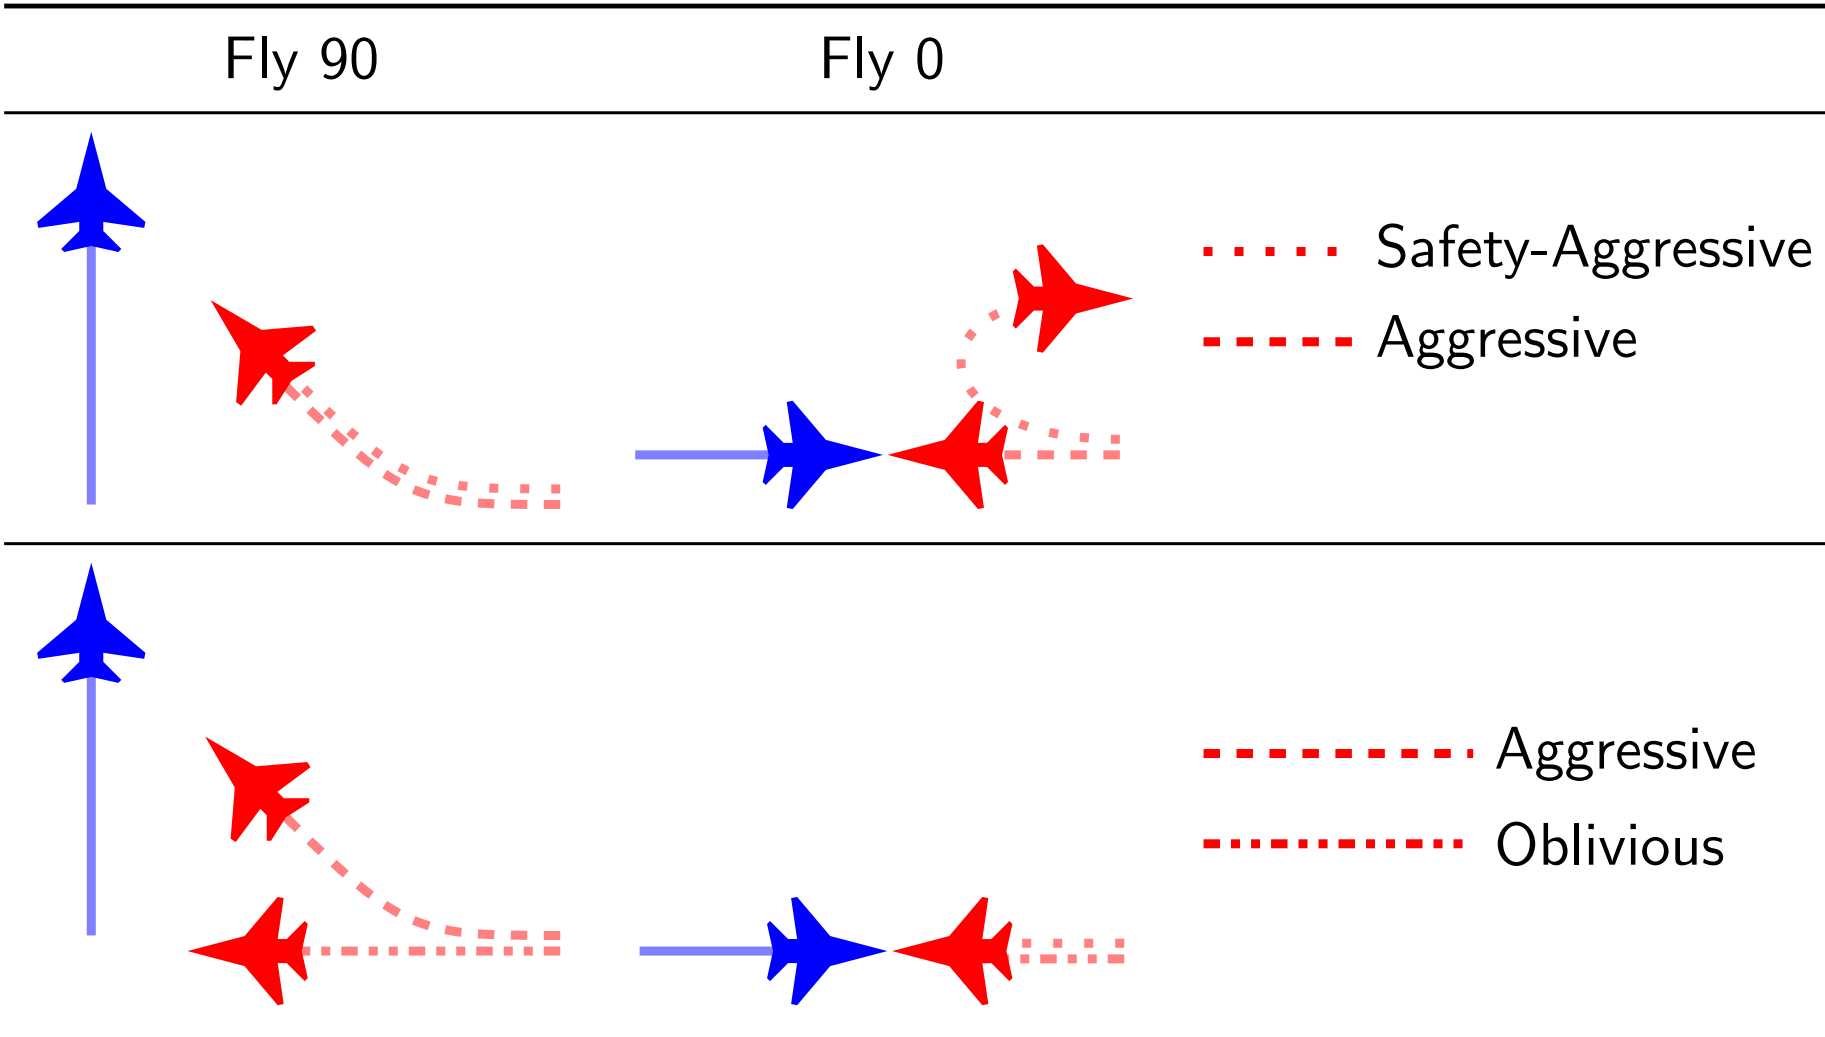

Case-based Behavior Recognition

The rough algorithm:

- During training:
  - Run a number of randomized trials
  - Project states to a feature space
  - Record short histories of features and their associated models as cases
- During recognition:
  - Retrieve cases with similar histories
  - Treat the relative frequency of agent models as a probability distribution

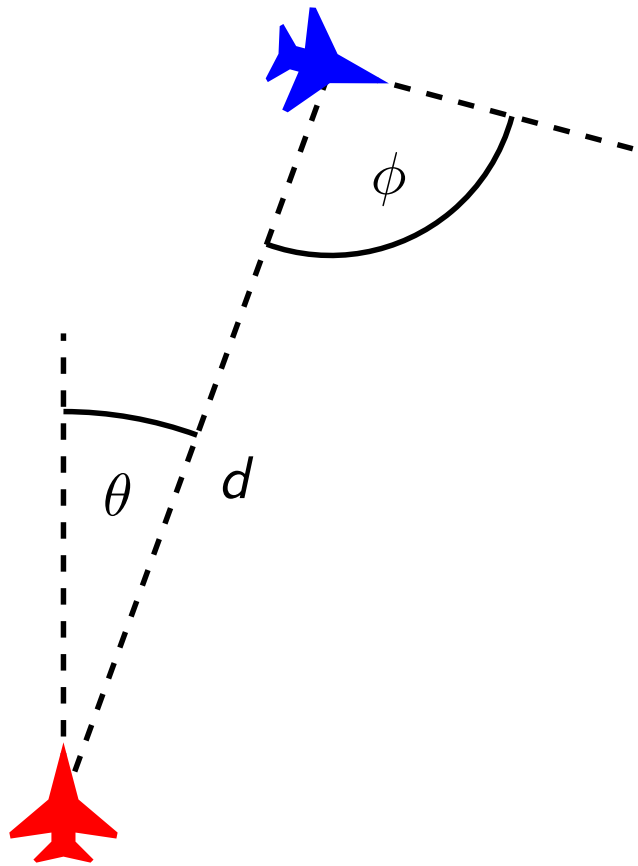

How acting influences Case-based Behavior Recognition

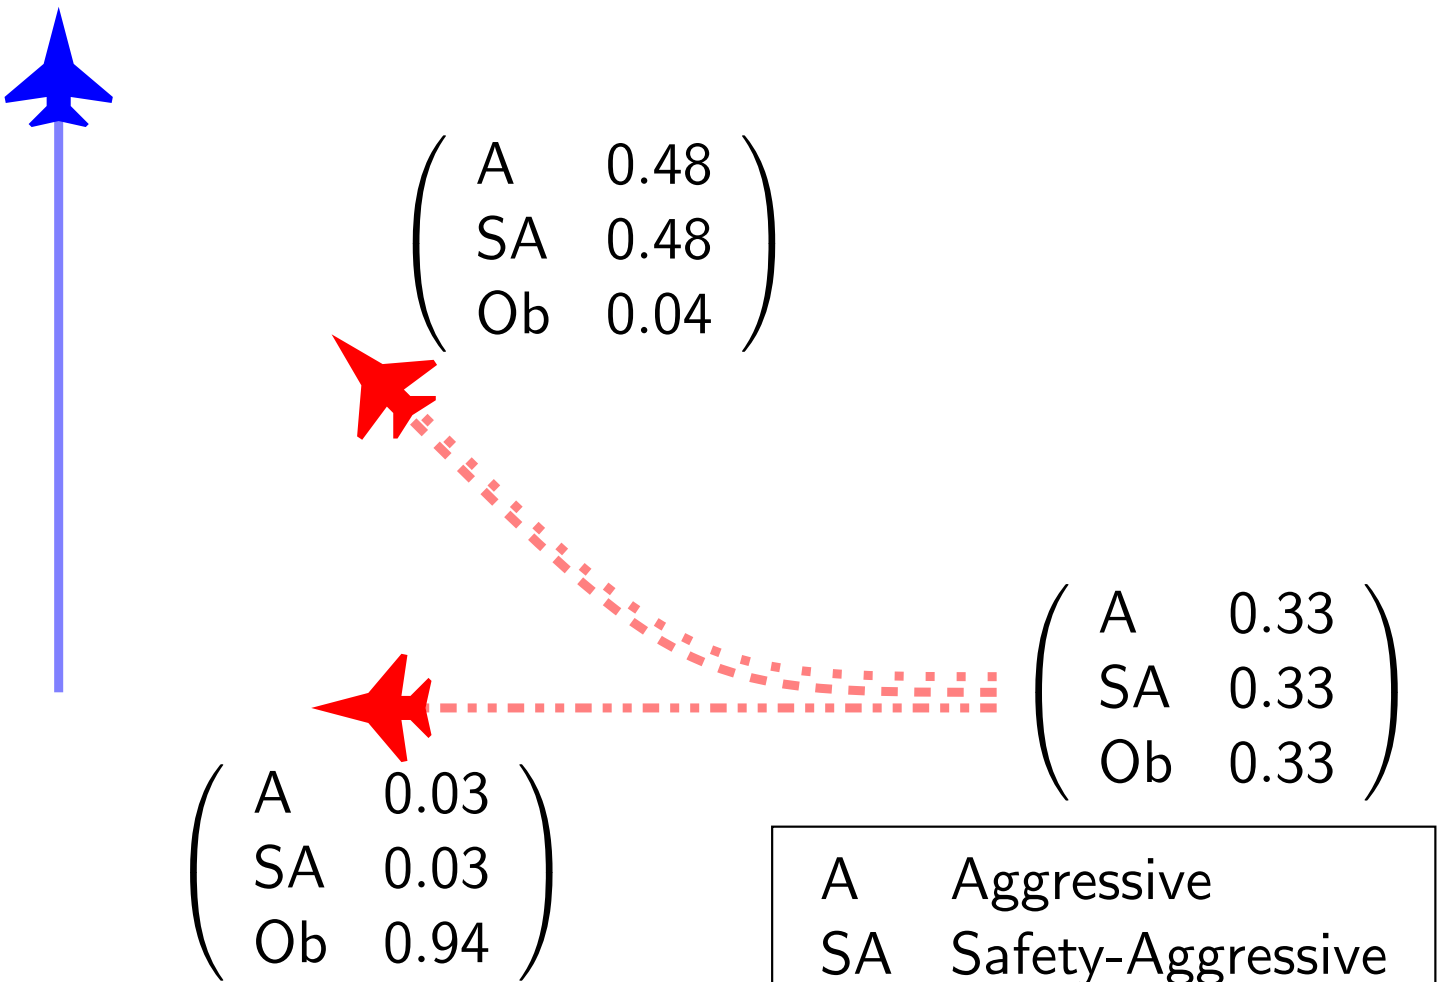

How acting influences Case-based Behavior Recognition

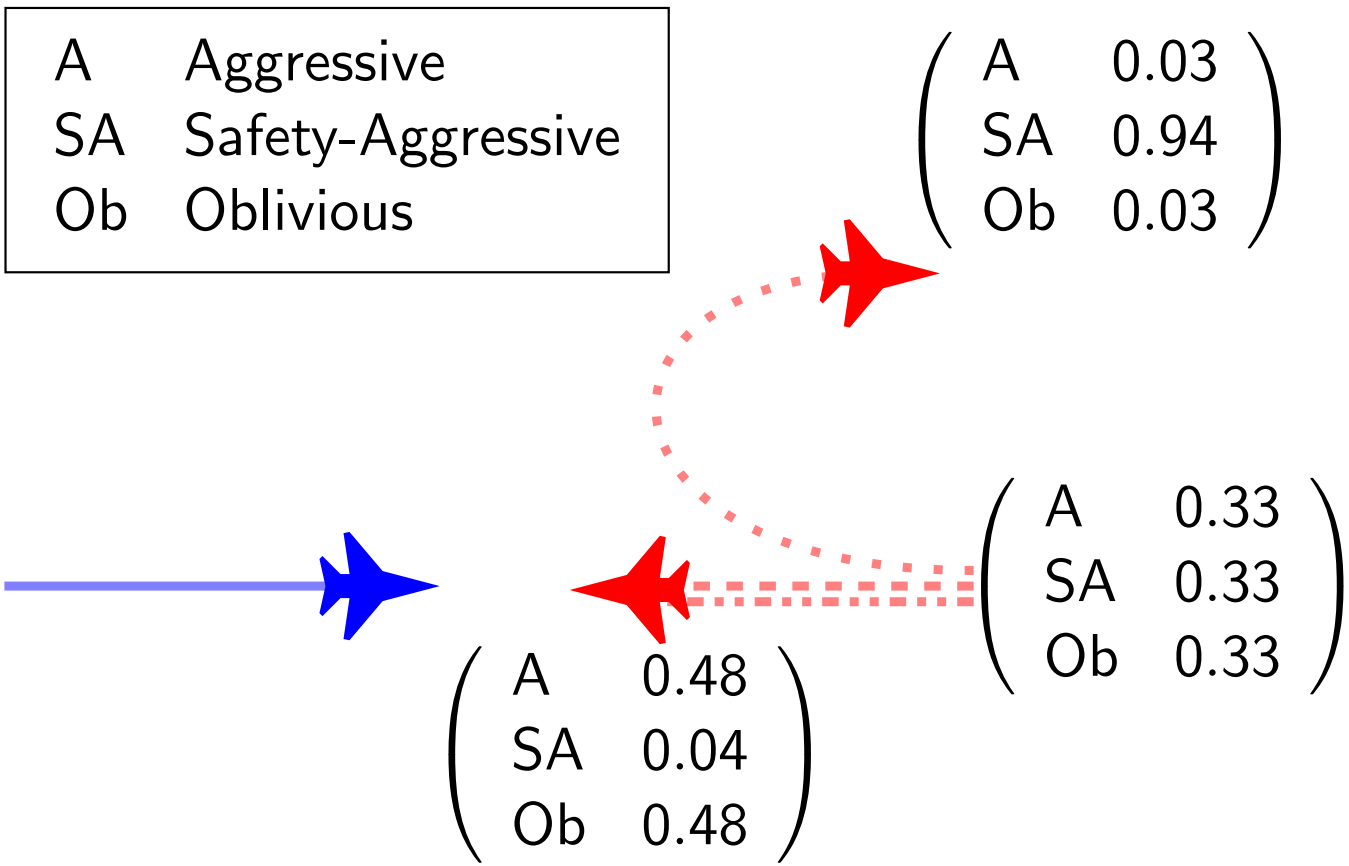

How acting influences Case-based Behavior Recognition

- Acting and Behavior Recognition:
  - Head-long flight disambiguates Safety-Aggressive
  - Perpendicular flight disambiguates Oblivious
  - Need both to make a confident prediction
- Similar to a POMDP
  - Move with uncertainty about other agents' behaviors
  - Our actions give evidence about that behavior
  - POMDPs are hard
- Approximate as an MDP

Planning Domain

Plan over histories of observations

- Observations divided into 60 second epochs
- Actions:
  - Four discrete actions
  - Four possible outcomes (agent models)
    - Probability dependent on behavior recognizer and current history
- Use flight simulator (AFSIM) applying action to a history
- Purpose: Maximize a utility function over finite horizon

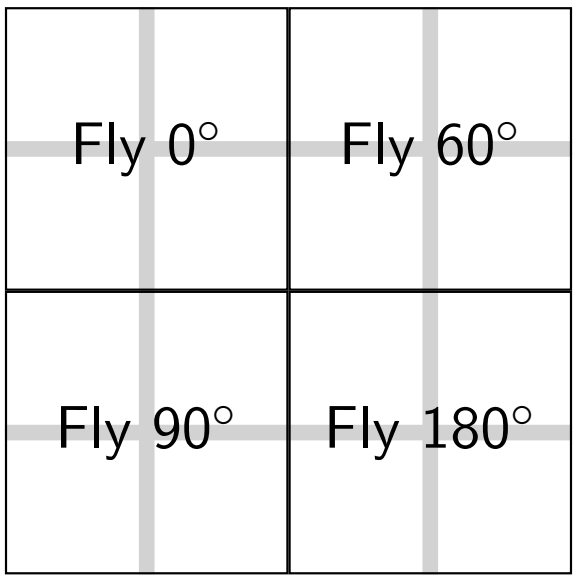

Sample-based planning (PROST)

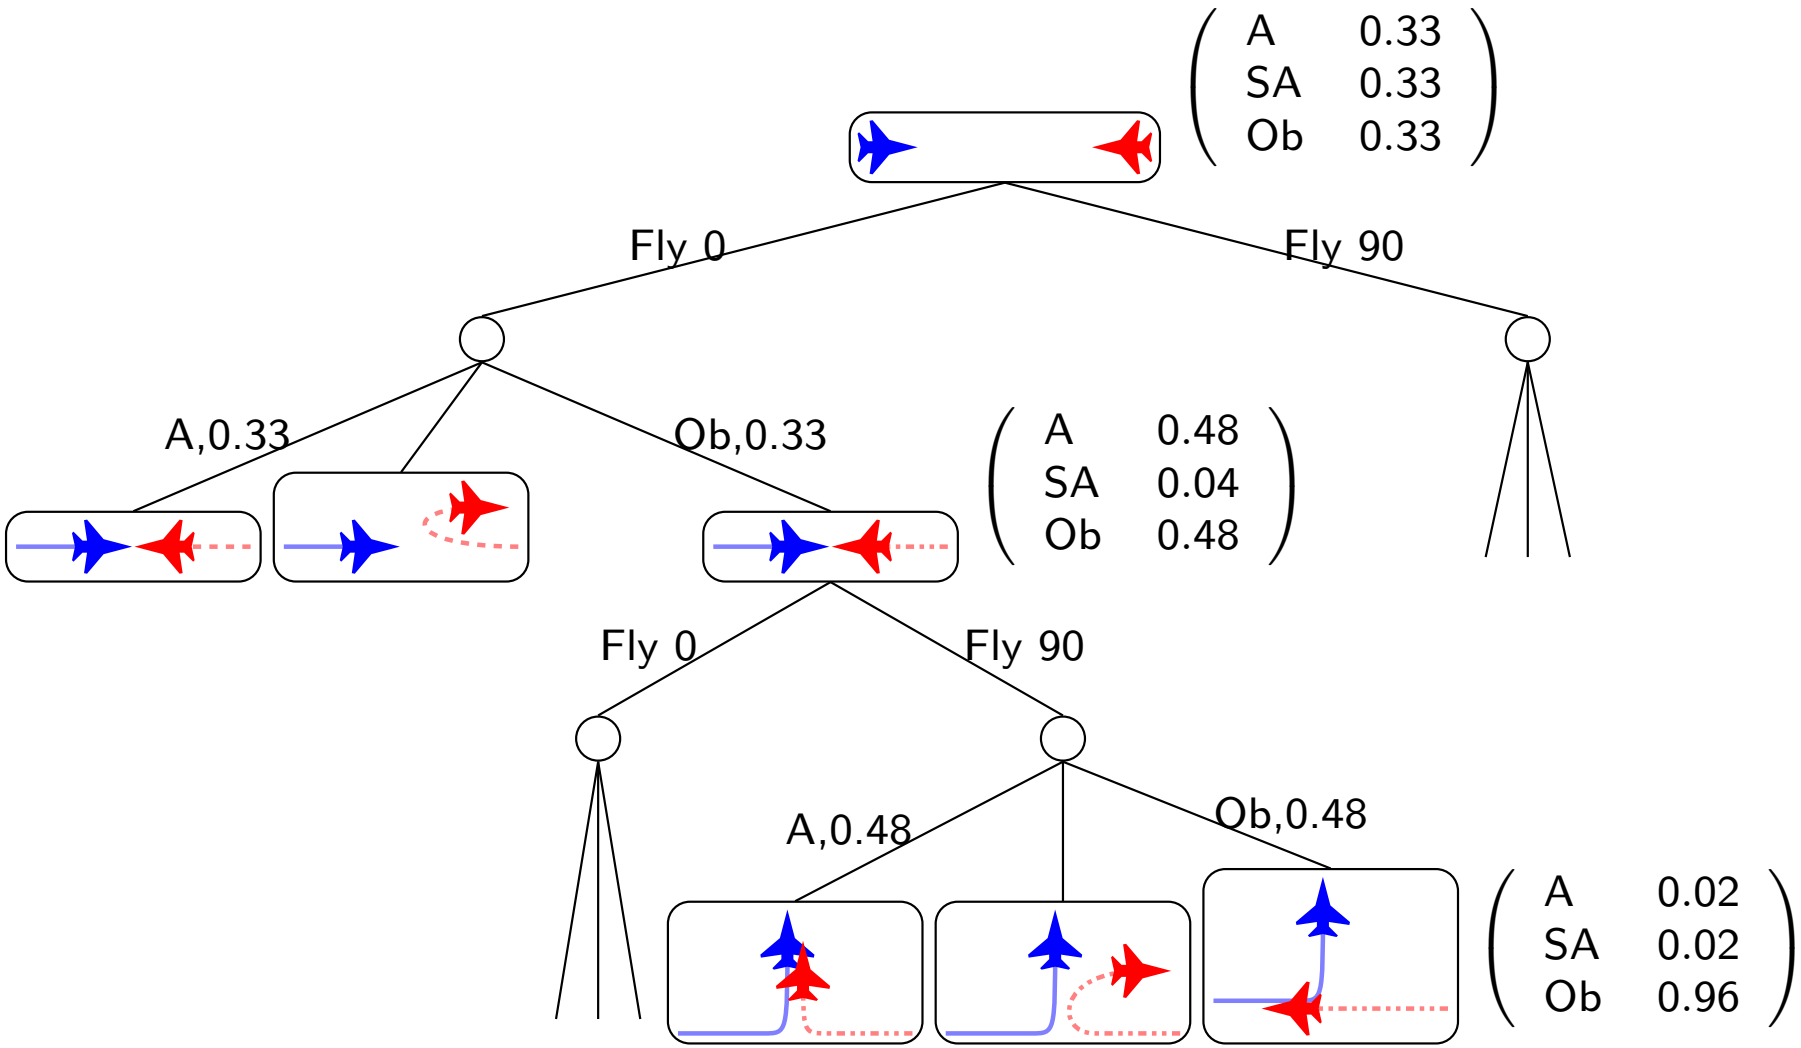

Meta-goal reasoning

We require a utility function!

- Possible mission success functions:
  - Number of “kills”
  - Air space denied
  - Reconnaissance
  - Diversion

Meta-goal reasoning

We require a utility function!

- Possible mission success functions:
  - Number of “kills”
  - Air space denied
  - Reconnaissance
  - Diversion
- Road blocks
  - Roll-outs (simulation) are slow
  - Success functions are often discontinuous

Meta-goal reasoning

We require a utility function!

- Possible mission success functions:
  - Number of "kills"
  - Air space denied
  - Reconnaissance
  - Diversion
- Road blocks
  - Roll-outs (simulation) are slow
  - Success functions are often discontinuous
- Instead: Average confidence in most likely model
  - Confidence is generally smooth
  - Emphasize the role of planning in resolving recognition ambiguity

Experimental Setup

- Four different observer behaviors running the behavior recognizer:
  - Safety-Aggressive
  - Passive
  - Random
  - Active behavior recognition planner
- Evaluation metric:
  - Confidence in correct behavior over time.

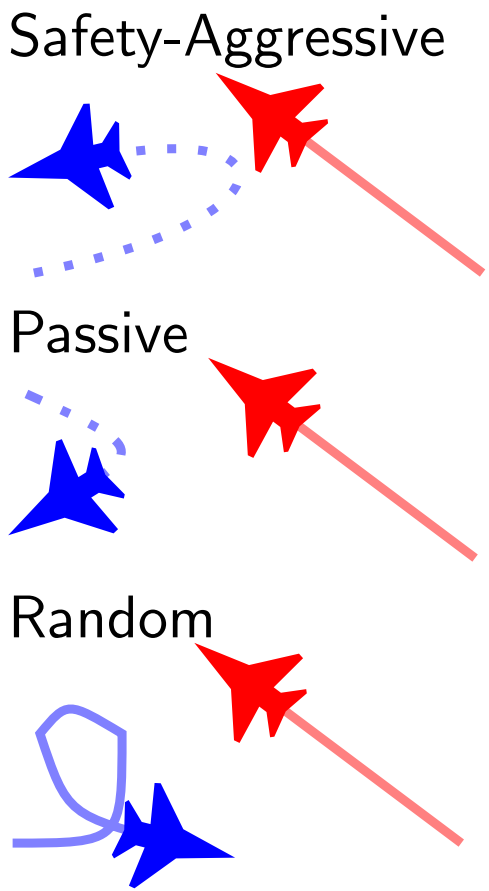

Recognition Results

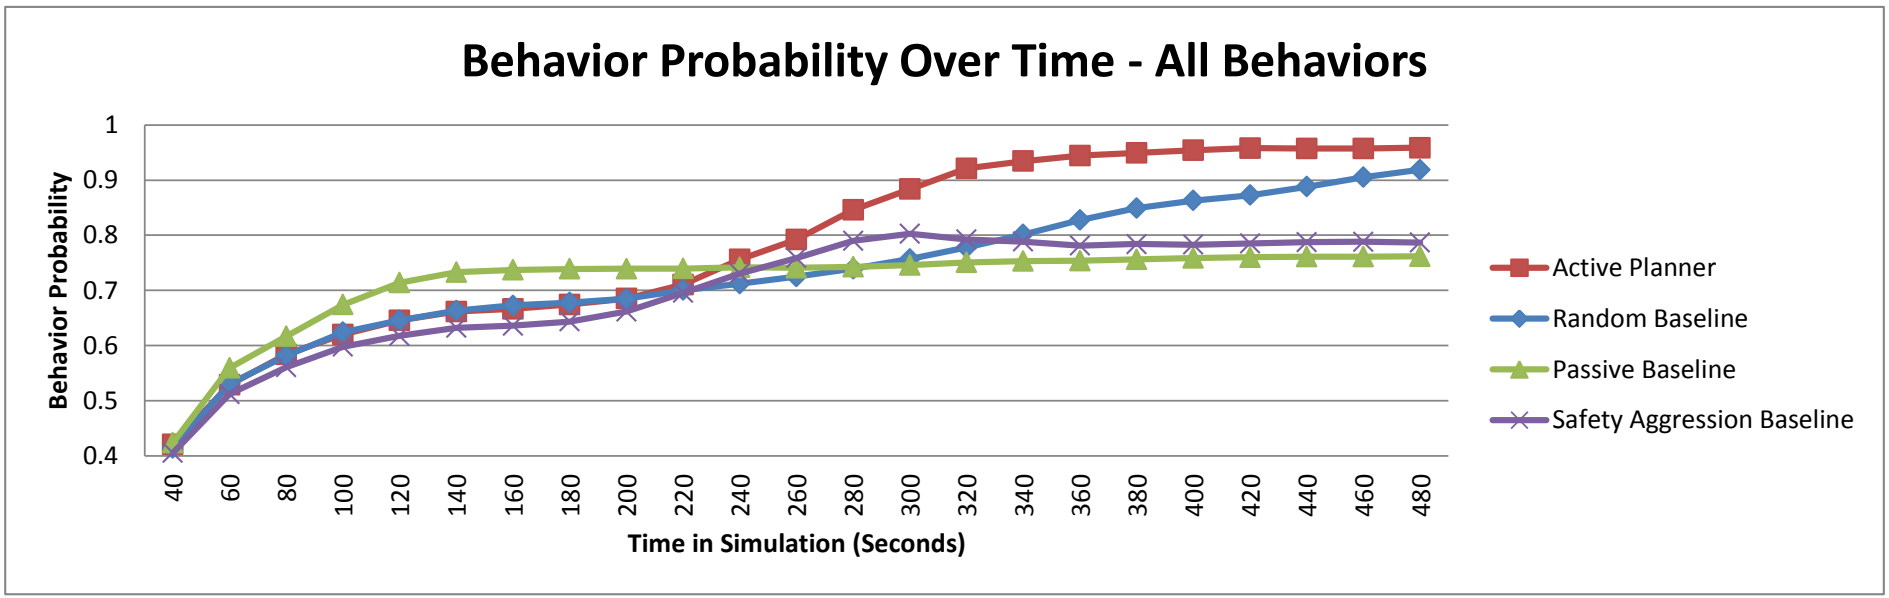

- Both Safety-Aggressive and Passive fail to disambiguate between two behaviors
- Random eventually distinguishes between all behaviors
- Planning gets good (>90%) recognition scores faster

Conclusion / Future Work

Behavior Recognition and Acting:

- Probabilistic recognition pairs well with probabilistic planning
- Need faster roll-outs to persue mission success
  - Discrete states, actions, and policies
- Game theoretic play (regret minimization)
- When do we need a separate behavior recognition component?