Three Simulation Models of Naval Air Defense

LTJG Baris Ozkan, Neil C. Rowe,
LT Sharif H. Calfee, and John E. Hiles
U.S. Naval Postgraduate School
Contact: ncrowe@nps.edu

Why isn't one simulation enough?

- Different simulations can focus on different features of a problem.
- Combining all features of a problem into one simulation may be too confusing to understand.
- Simulations can (1) simulate ordinary people, (2) simulate experts, (3) simulate good learners.
- We built three agent-based simulations of naval air defense of those kinds.
- Primary architects were Calfee for (1), Ozkan for (2), and Rowe for (3).

Naval air defense

- Goal is to protect a naval ship from air attack.
- Inputs are locations of "contacts" (platforms) obtained from radar and their observed properties.
- Outputs are orders for defensive measures.
- For U.S. Navy, air defense is done in the Combat Information Center (CIC) by 11 or more people.
- Thorough training is important, so all simulations logged results and replay of scenarios for use in a training tool.
- (Liebhaber & Smith, 2000) gives an excellent list of clues for evaluating contacts as reported by human experts from the U.S. Navy. So we used that.

Track factors cited in (Liebhaber & Smith, 2000)

- Low-altitude level flight
- Significant distance from civilian airlane
- Hostile or unknown airport of origin
- High speed
- Sharp turn
- Aircraft over water
- Not heading to civilian airport
- Military-type electronic emissions
- Nonzero or nonexistent IFF response
- Weapons system apparent
- Missile launch
- Coordination with other aircraft
- Air support
- Intelligence reports suggesting hostilities

The ADC (Calfee) simulation interface

Information flow between simulated people

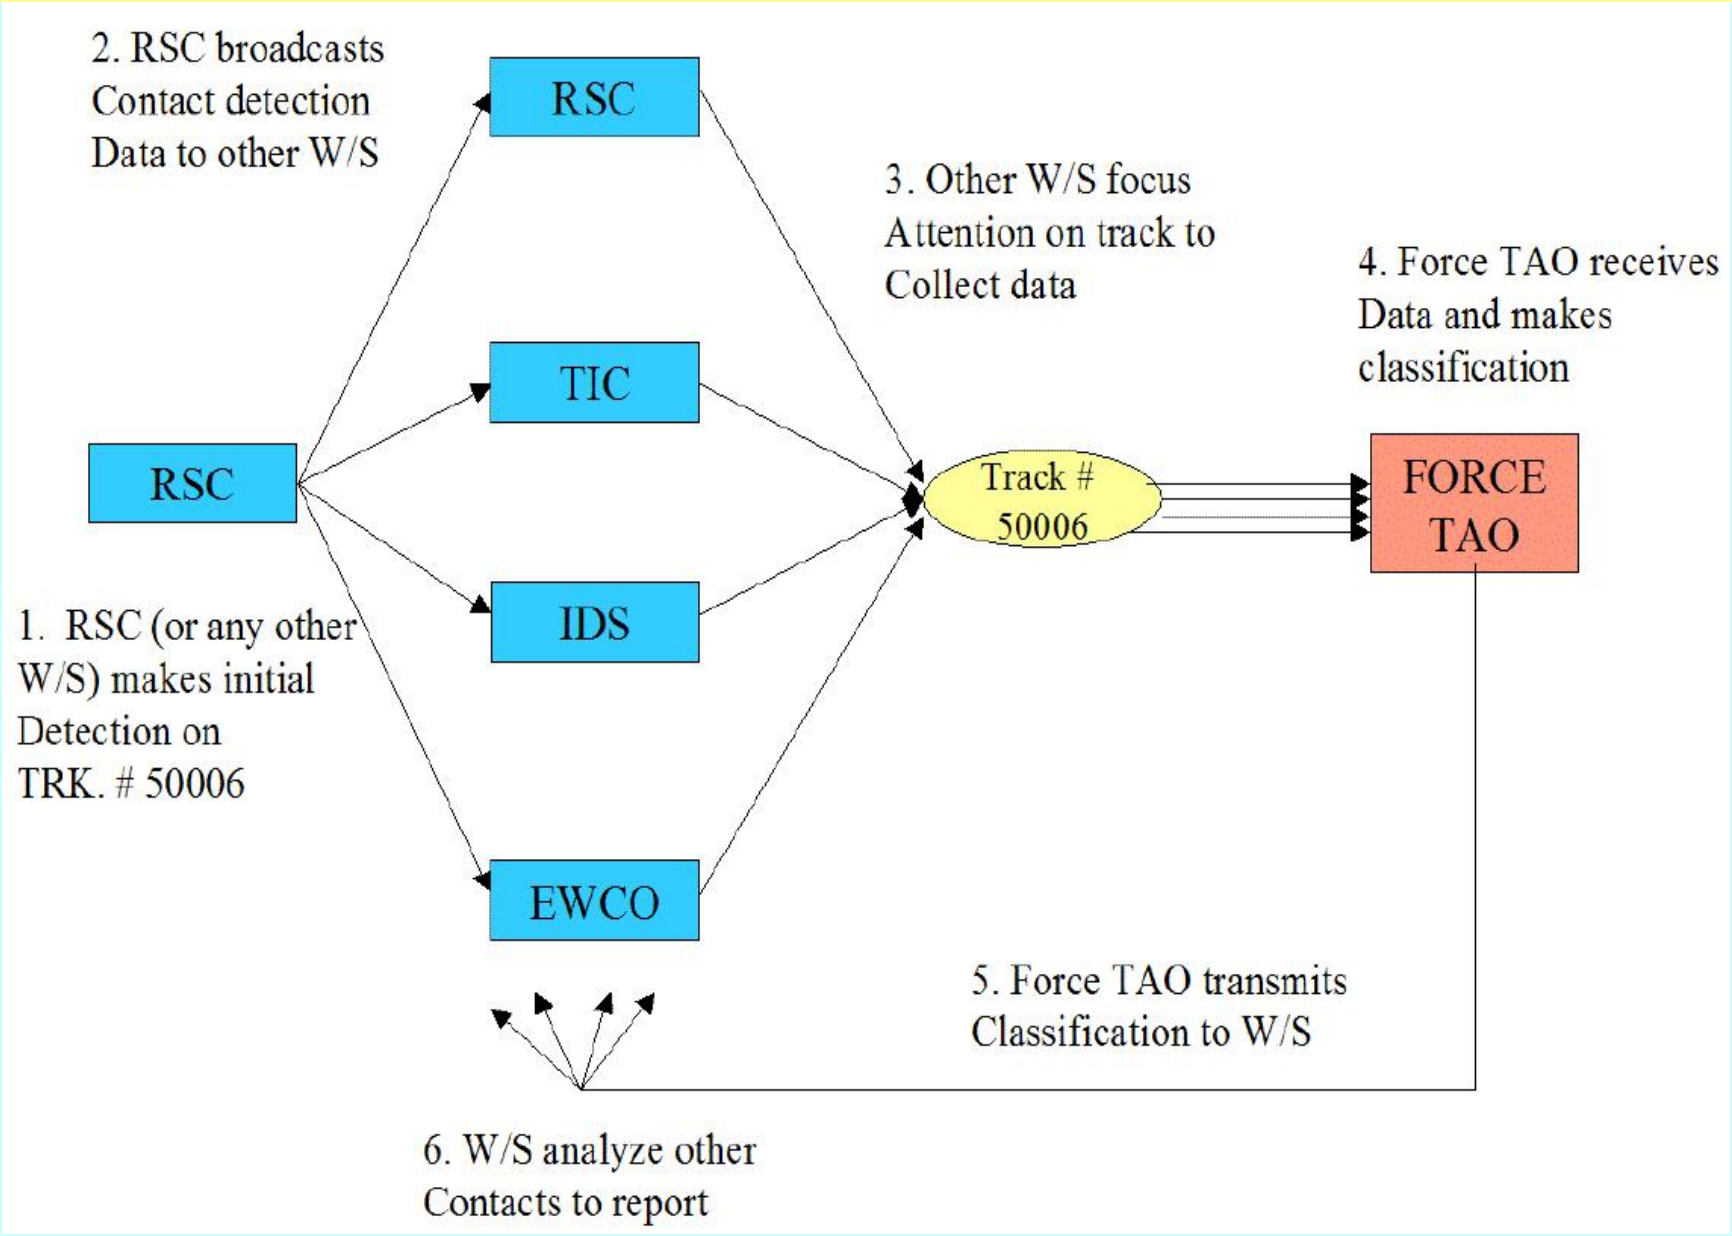

More about the Calfee simulation

- Modeled human imperfections too: Workload limits, lack of training, fatigue, equipment failure, situation criticality, weather.
- A neural network aggregated clues to classify contacts as Hostile, Suspect, Unknown, Neutral, and Friendly (at different thresholds of a single metric).
- Details were obtained from interviews with air-defense personnel.
- Subsequent tests at SPAWAR confirmed the realism of the simulation.
- The simulation is excellent for answering “what if” questions about the effect of factors like fatigue and training.

The ADL (Ozkan) simulation

- Focused on modeling the reasoning about air contacts done by the Anti-Air Warfare Commander (AAWC).
- Used “conceptual blending” to model these inferences, a form of reasoning by analogy.
- Used agents to represent the pieces of reasoning, not people.
- Modeling was done by interviews and documents.
- Experts confirmed the accuracy of the simulation.

The ADL simulation interface

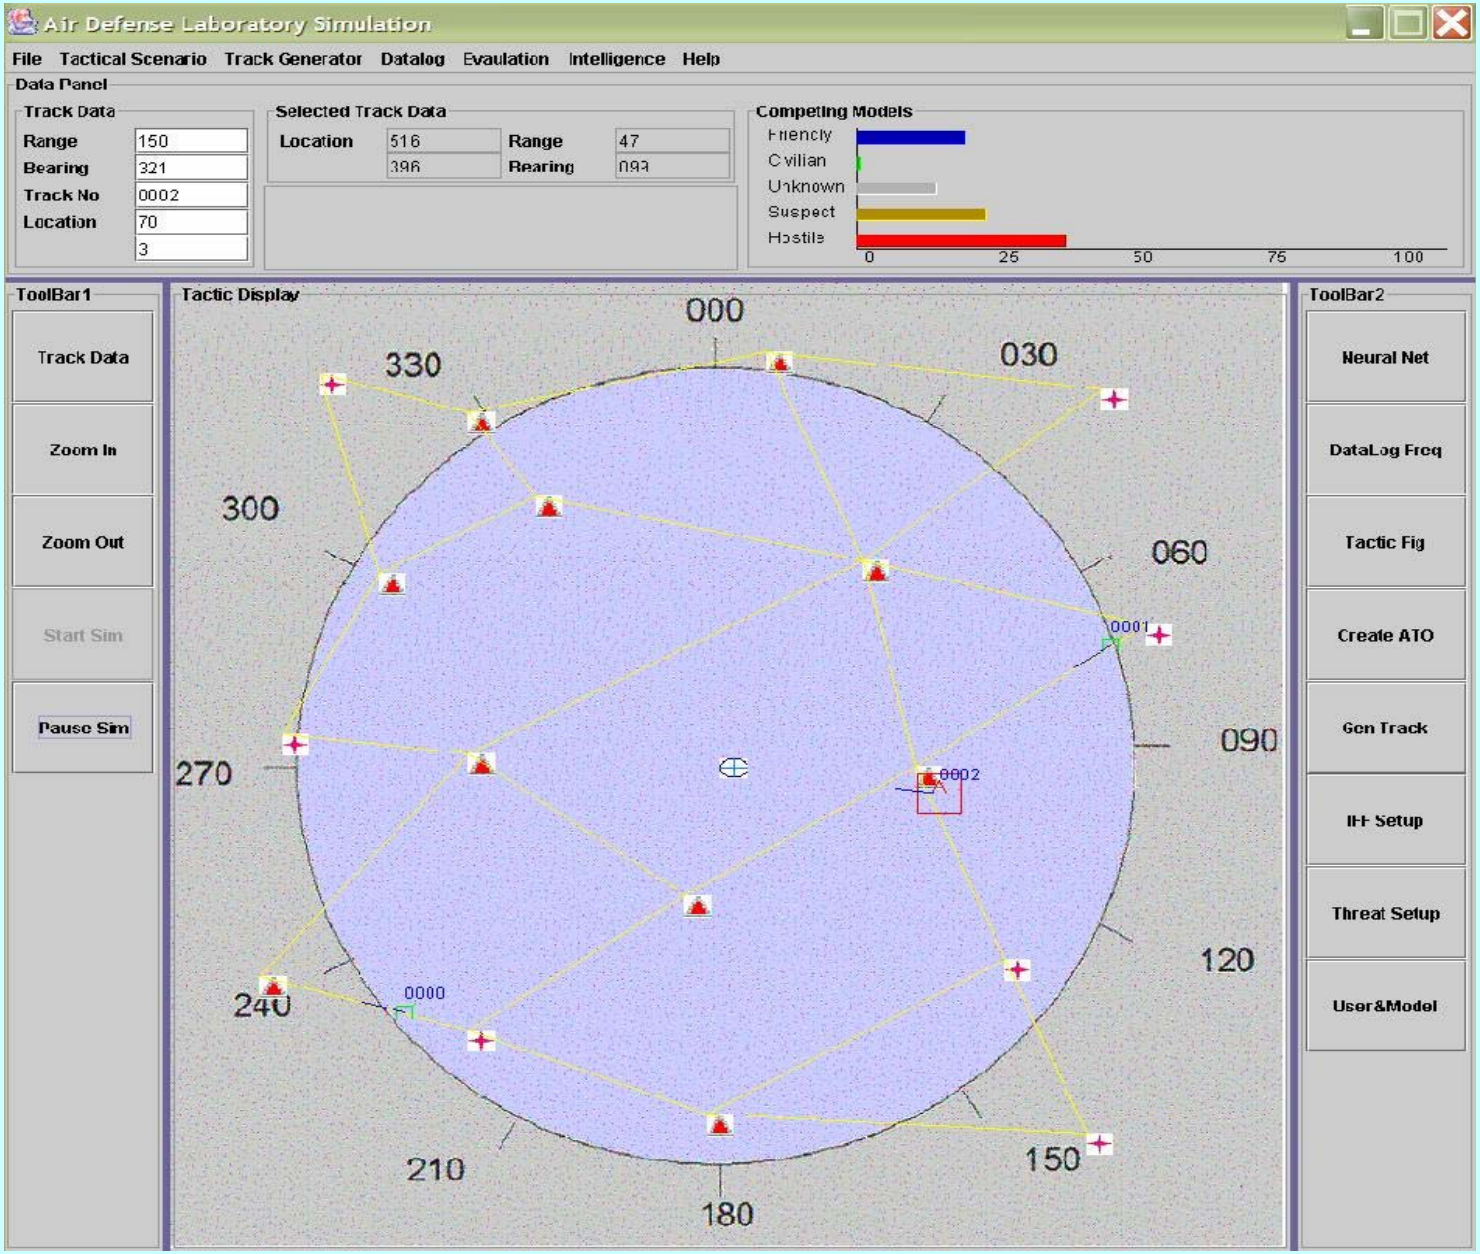

Agents in the ADL Simulation

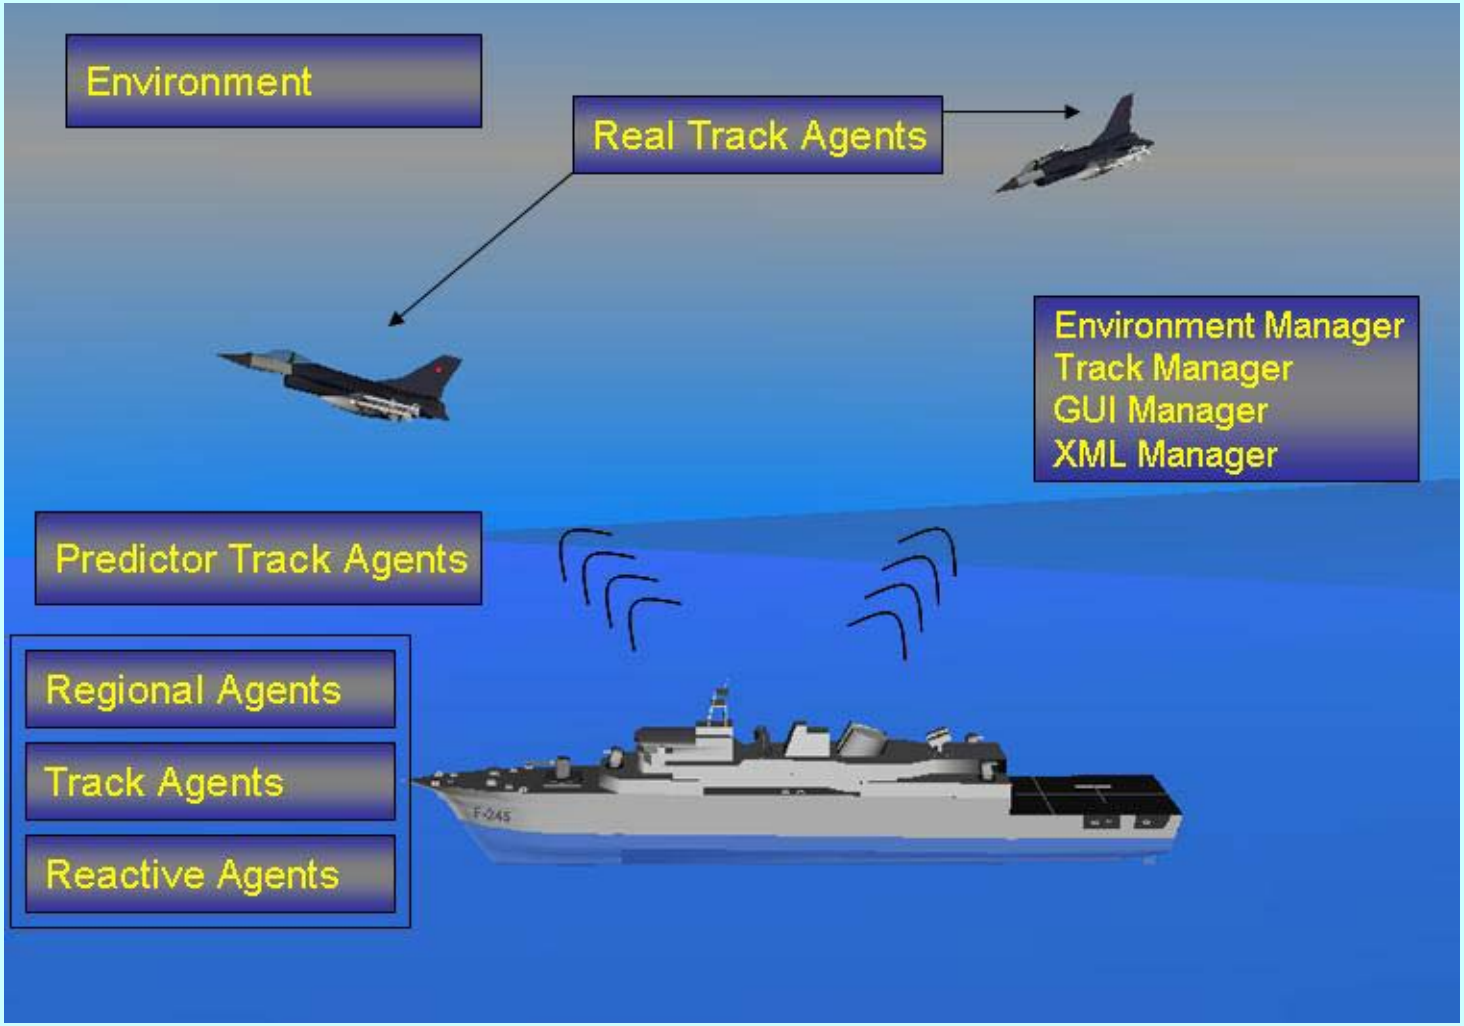

Principles of Blending

Conceptual blending is a set of operations for combining cognitive models in a network of mental spaces.

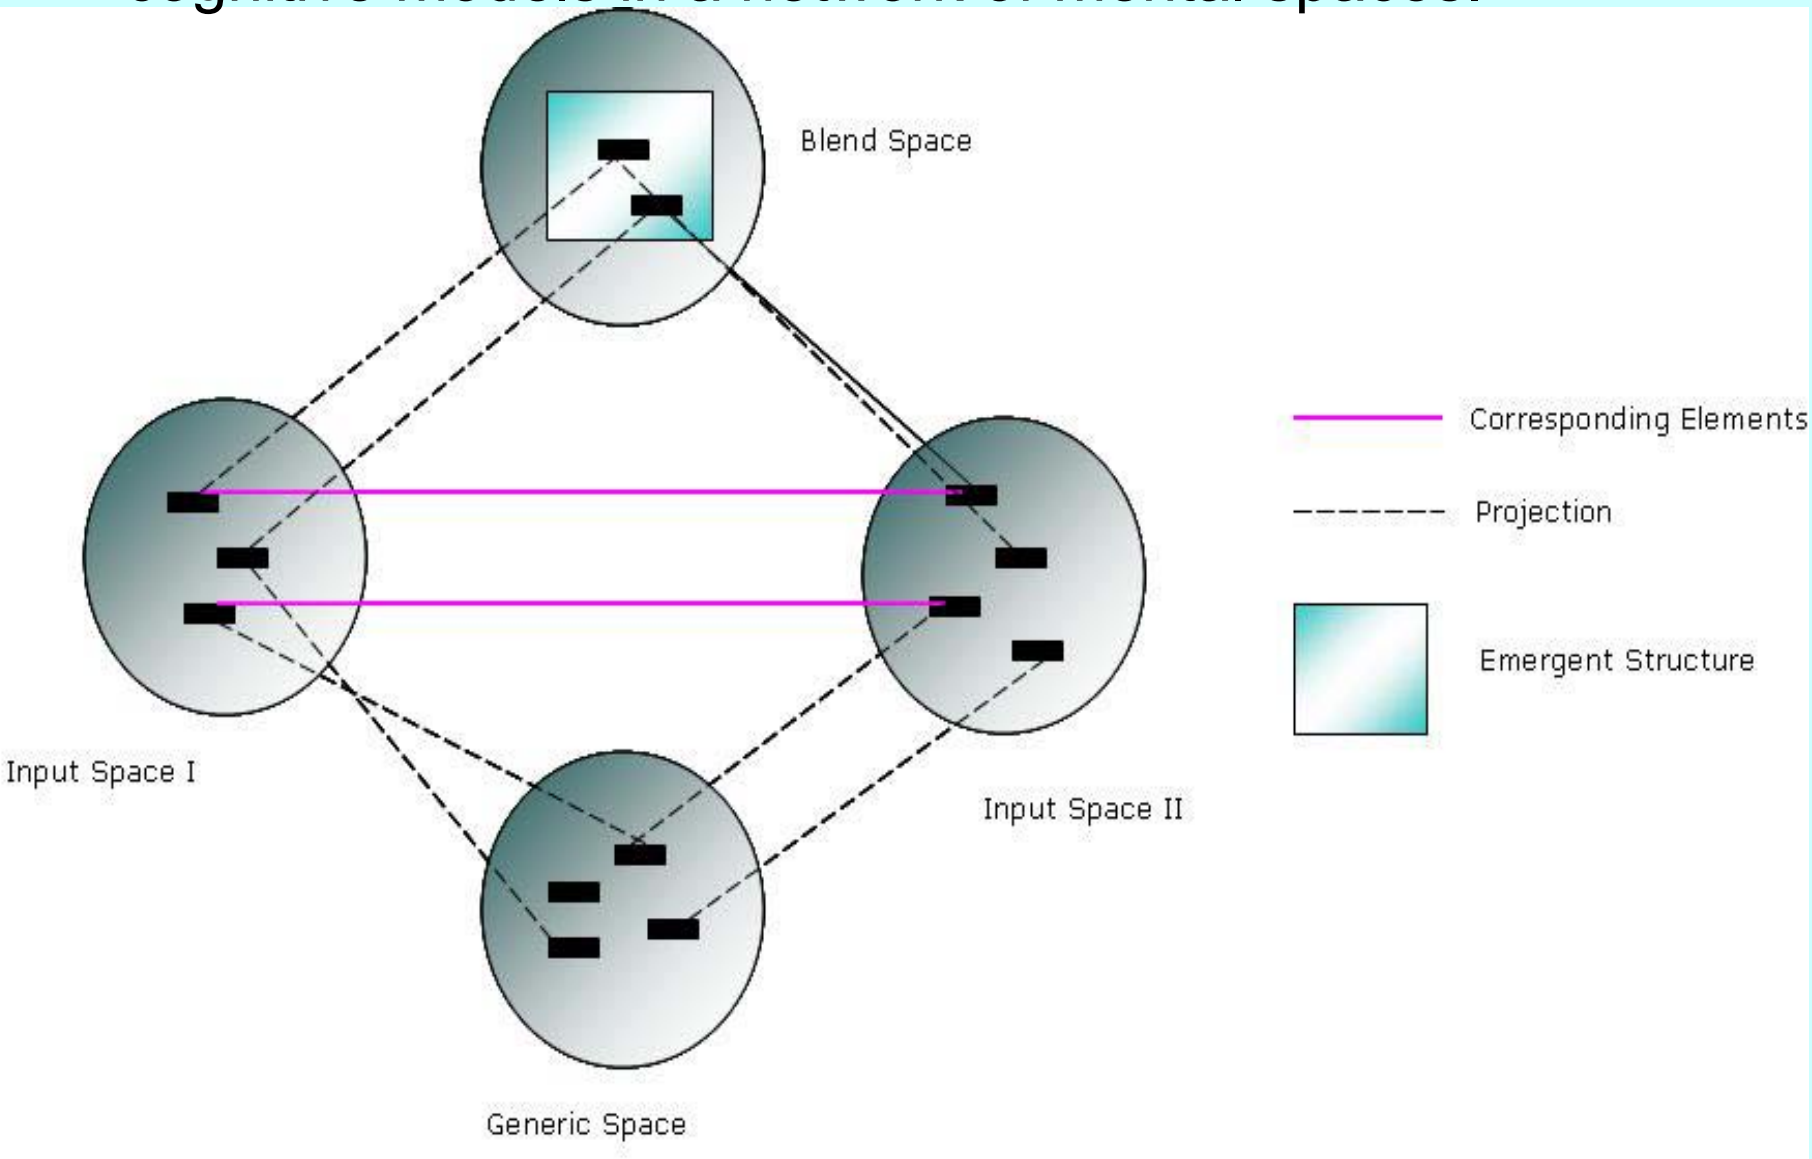

Operations of Blending

- Composition
- Completion
- Elaboration

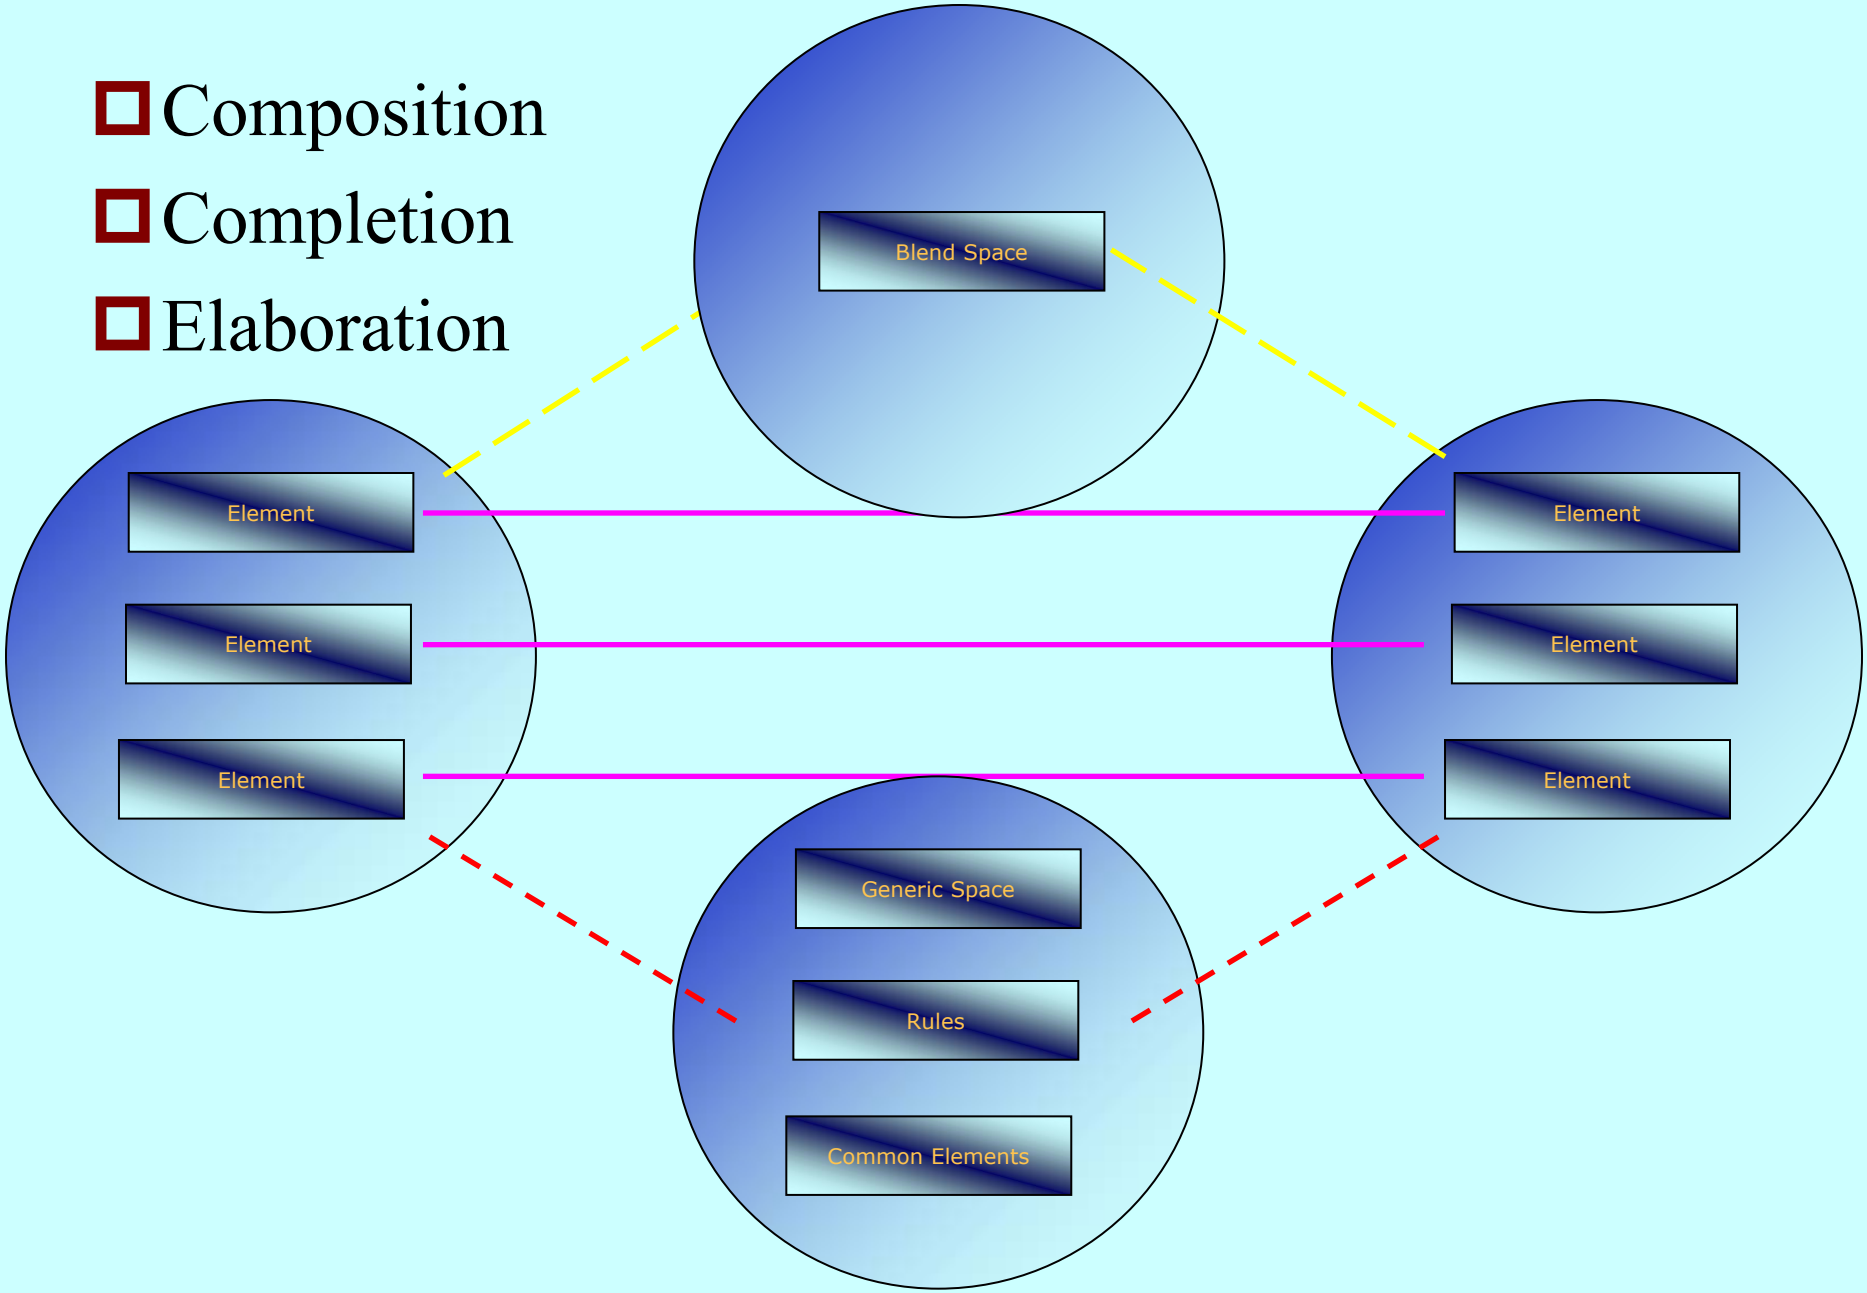

Air Defense Laboratory Simulation

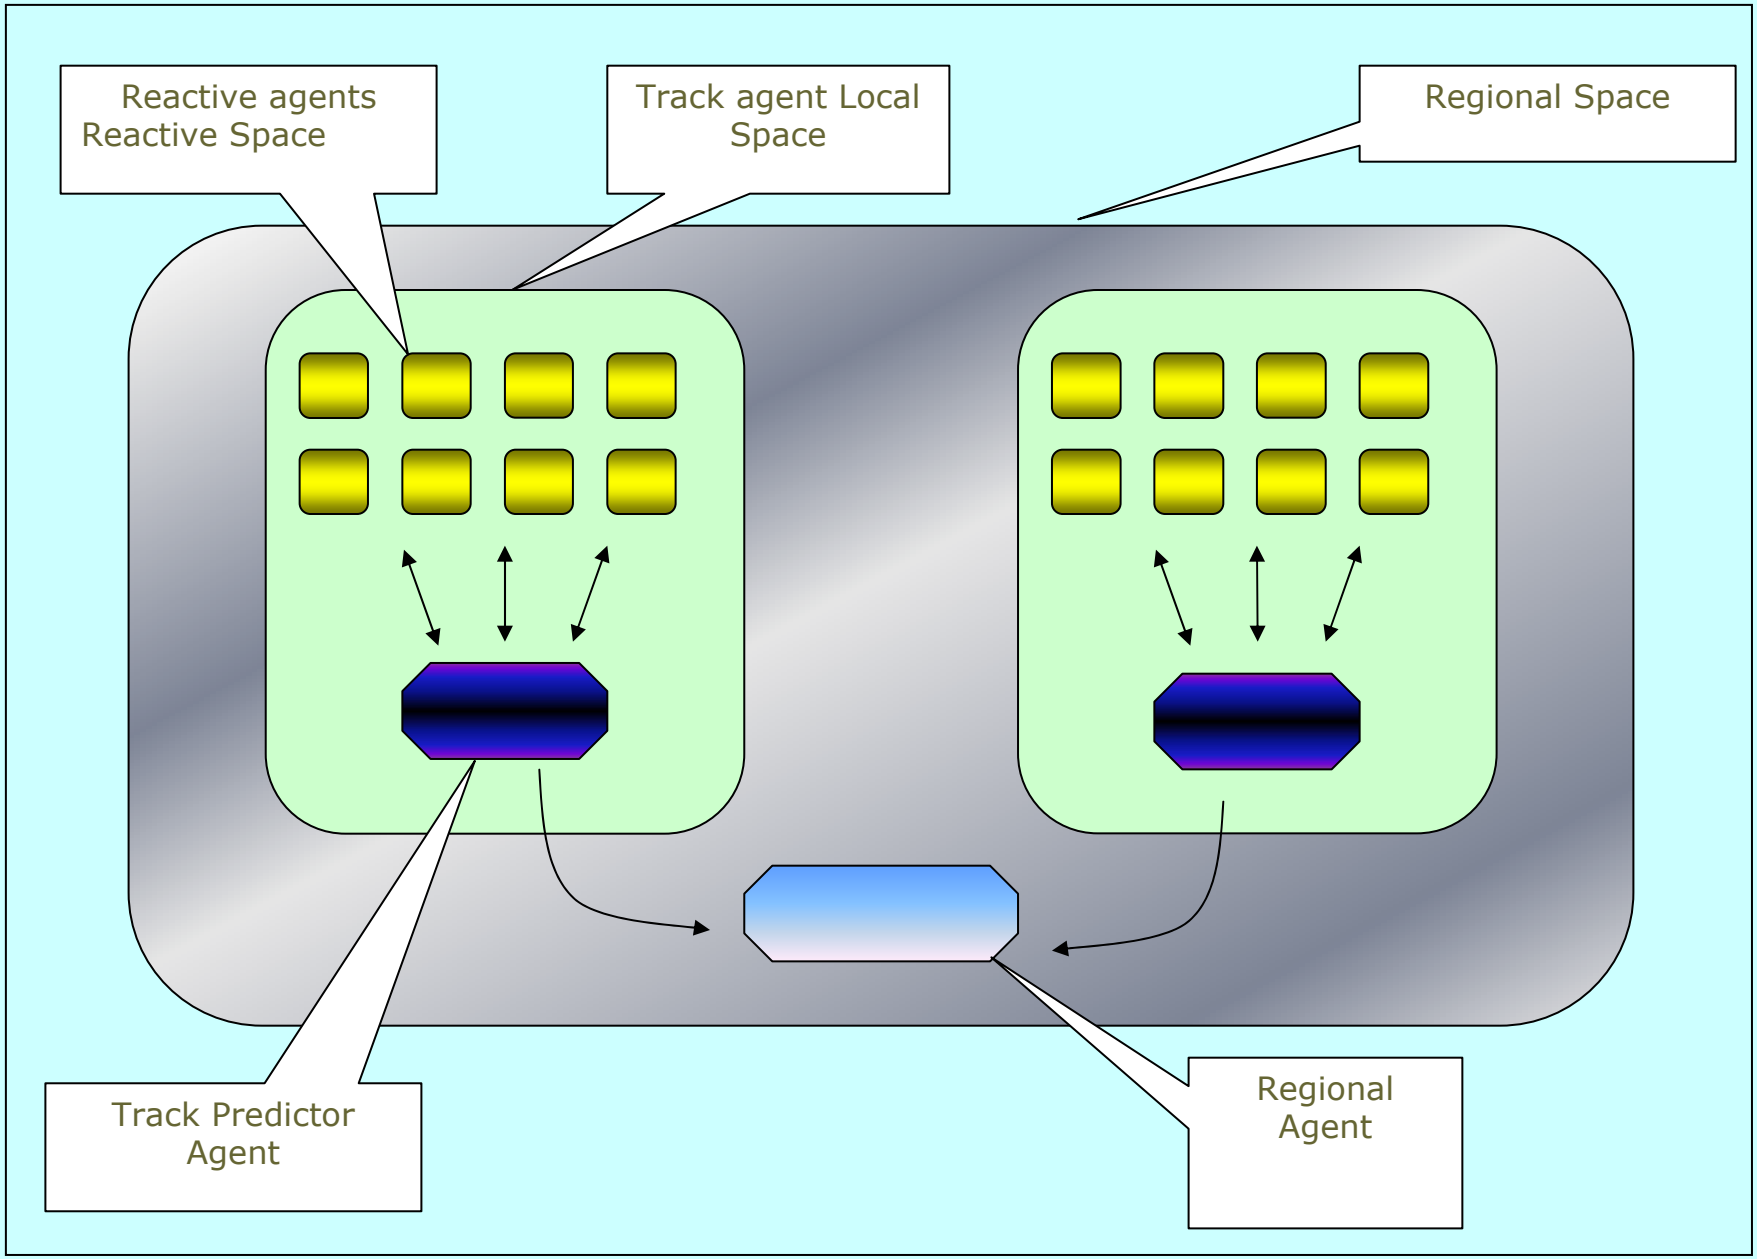

Air Defense Laboratory (ADL) Simulation

Reactive Agent Factors

Used continuously

- Location
- Heading
- Speed
- Altitude
- Range
- CPA

Used once and then only used if changed

- ESM
- IFF values
- IFF transponder status
- Radar status
- Intelligence
- Geopolitical situation
- Origin
- ATO
- Detachment
- Split
- Merge

Predictor Agent Competing Models

- Each identity is defined as a ticket.
- Tickets find a weight for each identity.
- Highest-weighted identity becomes the active model.
- Reactive agent connectors set the frames of these tickets.
- Also there are independent tickets: CPA calculator, IFF Evaluation, split detector, etc.

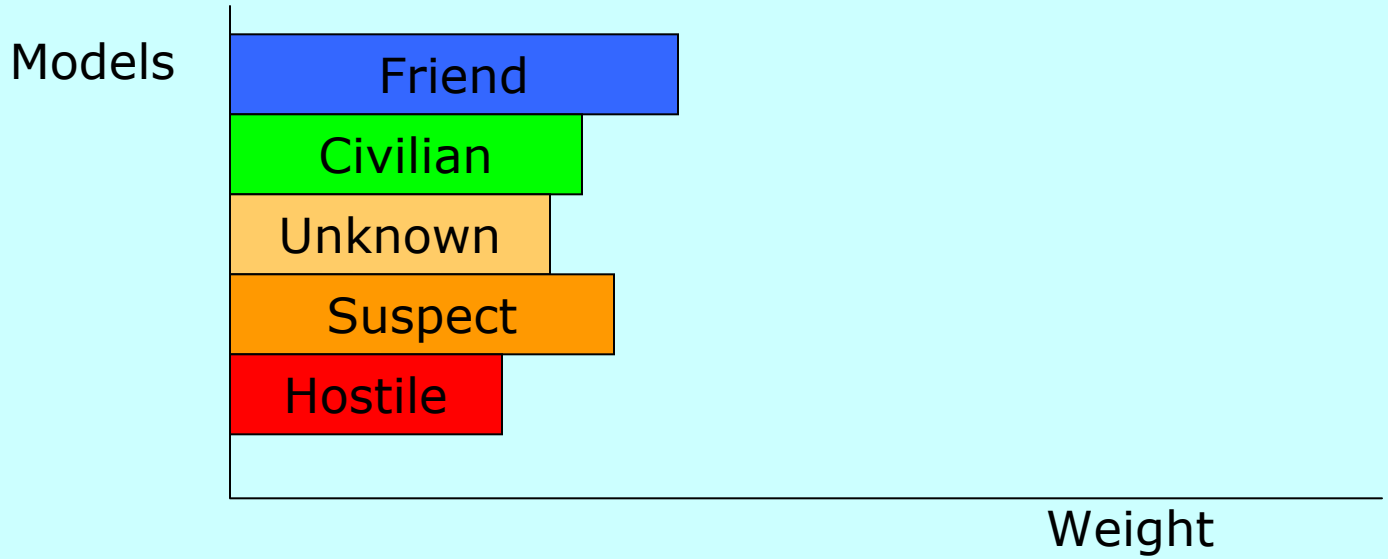

Regional Agent Activities

- Detecting merge activity
- Detecting coordinated detachment activity
- Detecting snooper coordinated attack activity
- Determining threat level

ADL Simulation and Blending

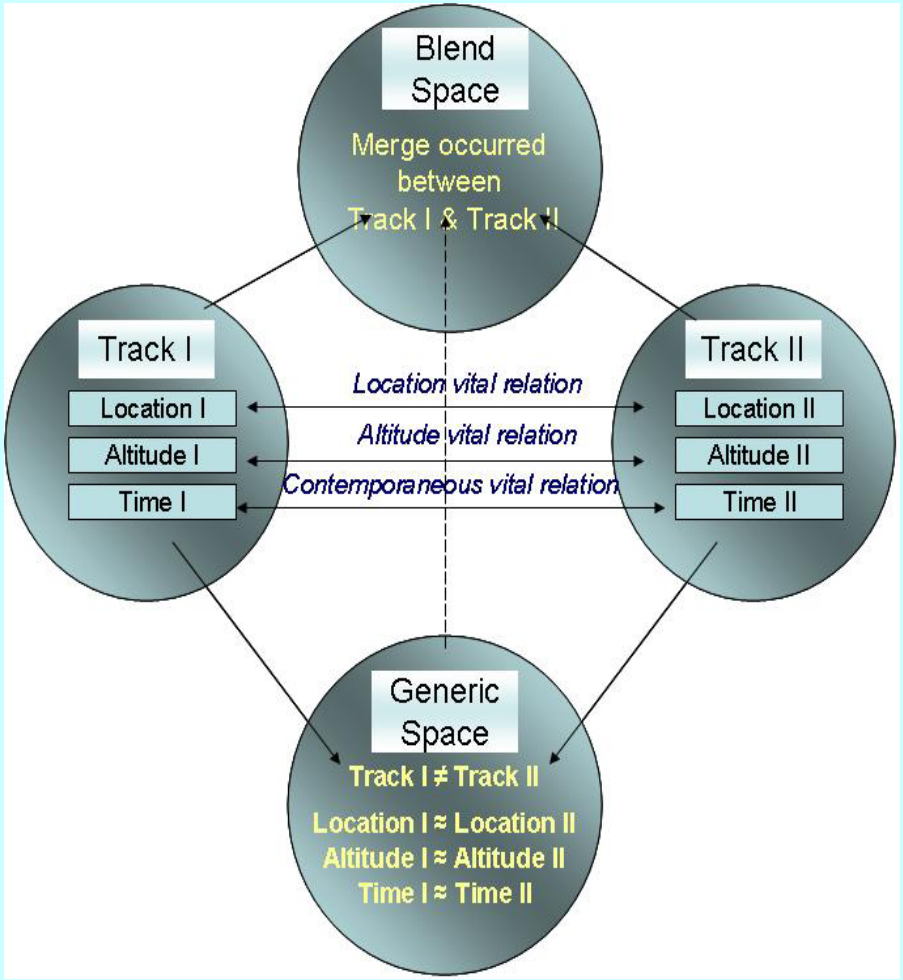

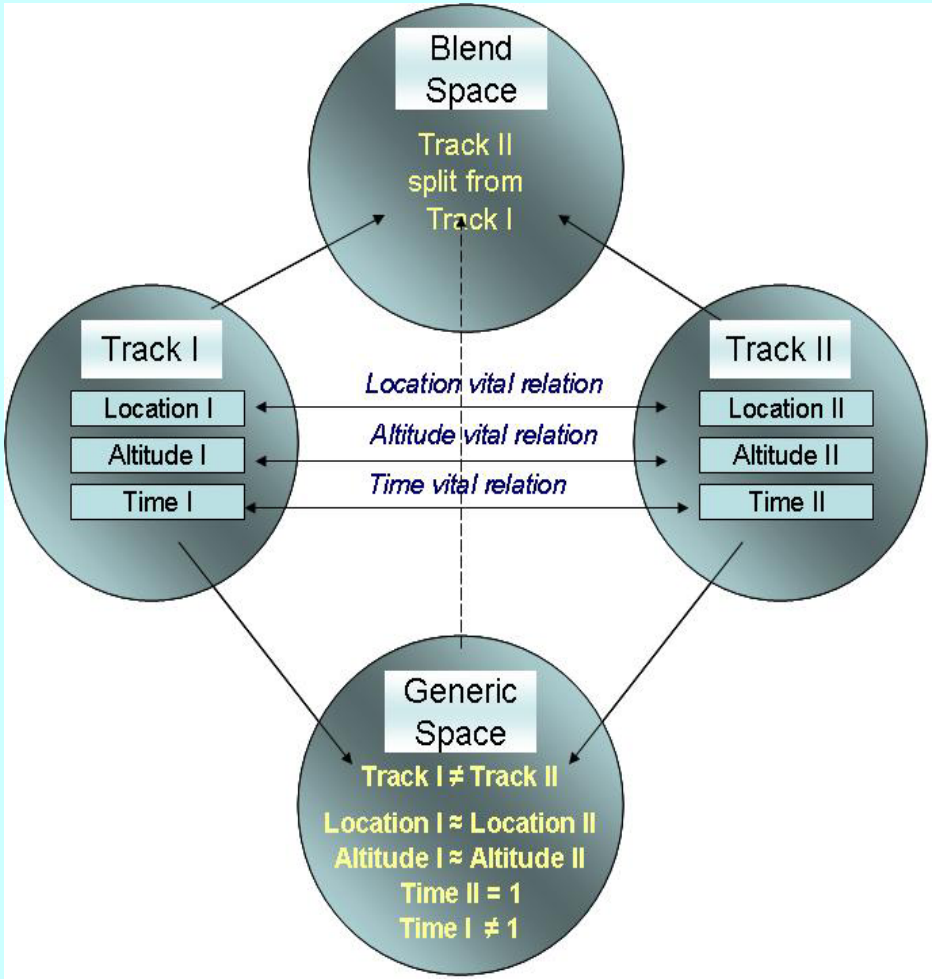

Is air defense too simple for conceptual blending?

- This is pretty far from the original use of conceptual blending to explain linguistic metaphors.
- In particular, you need to blend rather different mental spaces to get some power -- here we're blending the same kind, track data.
- Thus it's more appropriate to view this as inheritance rather than reasoning by analogy (which is what blending is).
- Inheritance: For every X and Y and for some Z, p(X,Y) <- r(X,Z), p(Z,Y).

The Bayesian (Rowe) simulation

- Neither of the previous simulation learn much from experience.
- A simulation could keep statistics from exercises to learn what clues signal hostile behavior. Use final assessment of a track and find what clues appeared in the course of the track.
- Such a simulation could be quite simple since it wouldn't need a lot of initial knowledge.
- Bayesian reasoning is the simplest way to implement such a learning system.

Naïve-Bayes odds calculation

$$o(H \mid (E1 \,\&\, E2 \,\&\, \ldots \,\&\, En)) =$$
$$[o(H \mid E1) / o(H)] * [o(H \mid E2) / o(H)] * \ldots * [o(H \mid En) / o(H)] * o(H)$$

This gives a formula for the odds that an aircraft is hostile (H) given evidence E1, E2, etc.

Here o represents odds or p/(1-p) and "|" means "given that".

Improvement if there are many time steps: Take to 1/M power each of the bracket ratios, where M is the "time window". This means updates don't change values as fast.

The Bayesian simulation interface

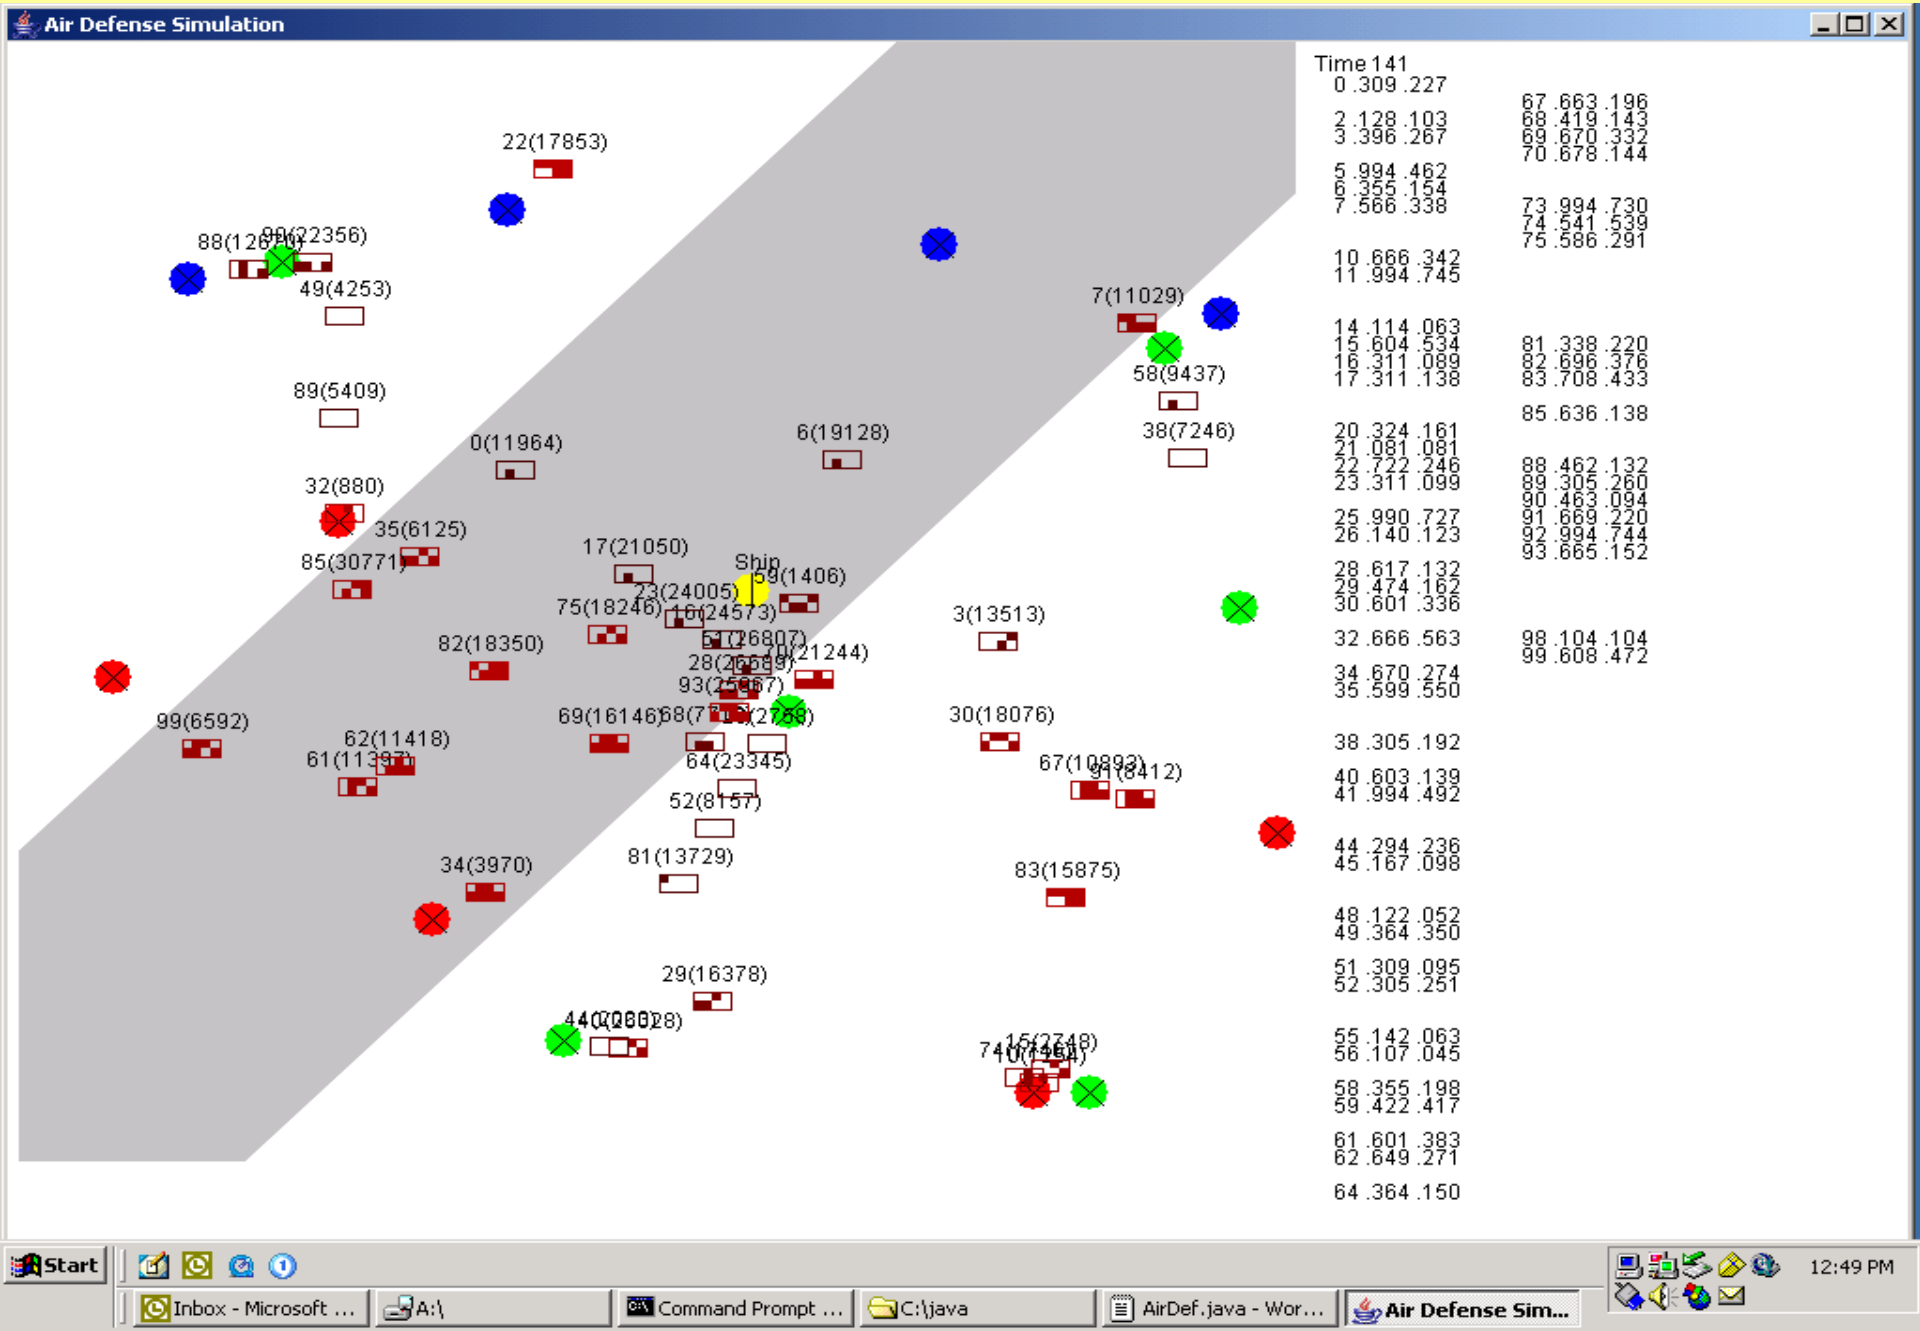

Testing of the Bayesian simulation

- We generated a variety of test scenarios involving civilian, friendly, and hostile aircraft. We can easily do > 100 aircraft at one time.
- Tracks could be scheduled flights (some originating outside radar range), military reconnaissance, "snoopers", outright attacks, and hijacked civilian aircraft.
- Results showed the system could improve with experience as it learned clues.
- Results showed it could learn how to respond to a new threat it had not seen before, the hijacked civilian aircraft, when first trained on scenarios without it.

Conclusions

- Occam's Razor applies: Bayesian simulation seems to do almost everything the Ozkan simulation does, in 60 times less code. Thus we should prefer the former to automate air defense (but not to study it).
- Air defense may be too far from linguistics, the original domain for conceptual blending.
- The Calfee simulation addresses a different problem, of modeling personnel. But Bayesian simulation suggests automating much of what those 11 people do.
- Bayesian simulation requires good training of program, which may be hard to set up.
- Bayesian simulation can be fooled by deliberately deceptive enemies, but so can people.

Automation reduces the need for training

- If we can significantly automate parts of air defense, it simplifies the tasks of the remaining personnel.
- That means less training is required.
- Thus the goal of Mike Zyda's USC-ISI group is self-contradictory: If we can develop wonderful virtual environments for training, we can usually automate the tasks taught and have no need for training.
- Exceptions would be skills requiring human judgment -- human vanity exaggerates their extent.
- But human judgment shows many suboptimal biases.
- The USS Stark and USS Vincennes incidents illustrate that people can have poor decision-making in air defense -- a computer might manage it better.