Challenges in Mixed-Initiative Teaming

Chris Miller, **Mark Burstein**
SIFT, LLC

A Tale of Two Associates

- **Pilot's Associate (1985-1991)**
  - Single Pilot
  - Direct pilot interaction with associate meant added workload
  - Design philosophy minimized direct pilot interaction with associate
  - Moderate user acceptance

The Pilot is ALWAYS in charge.

The Effort required of the pilot to control the associate must be less than the effort saved by the associate

Associate adapts to the pilot

- **Rotorcraft Pilot's Associate (1994-99)**
  - Two Pilots
  - 1/3 of human activity is crew coordination
  - Design philosophy included direct pilot interaction with associate
  - Improved User Acceptance

Pilot adapts/instructs the associate

Lesson Learned: The Big Tradeoff

Every-thing is done the way I like it

1. Pilot in charge of tasks
2. All needed tasks accomplished
3. Pilot in charge of information presented
4. All needed information provided
5. Stable task allocation
6. Only needed information provided
7. Tasks allocated as expected
8. Information presented as expected
9. Stable information configuration
10. Tasks allocated comprehensibly
11. Only needed tasks active

Every-thing gets done (well)

Operators want to remain in 'charge', even when they can't be fully in control

A Playbook® Approach to Delegation

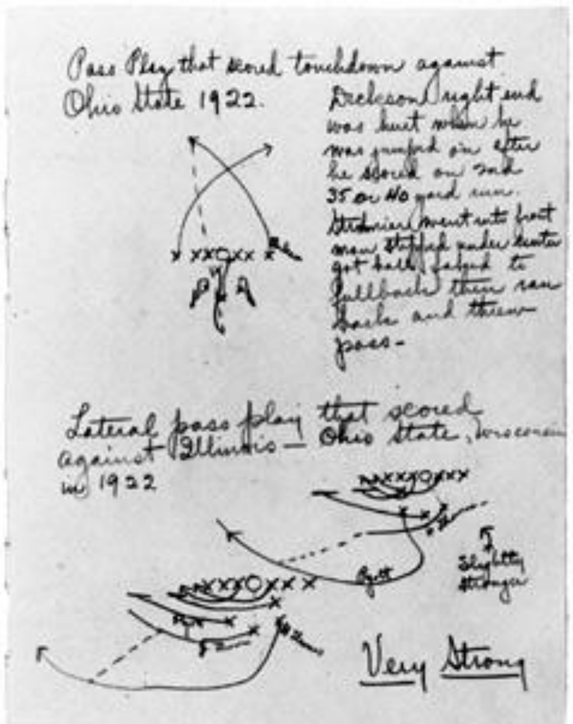

A page from Alonzo Stagg's 1927 Playbook

A means of delegation requiring a shared knowledge of domain Goals, Tasks and Actions

- Supervisor calls plays; Agents have autonomy within the play's scope
- Plays reference a defined range of plan/behavior alternatives
  - Supervisor can further constrain/stipulate
  - "Service Requests" are like plays, but without full command authority

"Plays" used to coordinate expectations, information needs, etc. between human/automation agents

- Adaptable automation and information management
- Trust impacts based on practice, detailed knowledge of range of autonomy allowed by play.

Mixed-Initiative Planning and Scheduling

- Planning/Scheduling tasks shared between human and automated 'collaborators'
  - Humans making the larger decisions about problem approach, constraints to be relaxed
  - Machine must able to evaluate and communicate about
    - Goal/Plan/Risk tradeoffs
    - Resource limitations affecting goal achievement
    - Anticipated Execution Problems
- Time limitations impact all of these
  - Communication with human is a slow channel
  - Available time to diagnosis or reach decision may be limited during execution
    (especially when faced with adversary)

Mark Burstein and Drew McDermott, "Issues in the Development of Human-Computer Mixed-Initiative Planning," Cognitive Technology, B. Gorayska and J.L. Mey (eds.), Elsevier, Jan, 1996, pp. 285-303.

Effective, Timely Communications A Key

- Human leveraging pre-shared knowledge
  - Of plans (playbook)
  - Of graphics representations enabling rapidly interpretable and highly diagnostic situation visualizations
- Agent determining highest value information for decisionmaker
  - Requires a model of shared goals/objectives/environment
  - Reasoning about critical factors impacting decisions/tradeoffs
    - Machine should analyze tradeoffs and **GATHER ADDITIONAL INFORMATION from its decision procedures** in order to
      - EXPLAIN CHOICES IT MAKES
      - EXPLAIN TRADEOFFS IN CHOICES HUMAN MUST MAKE

Intrusion detection and the feed-forward paradigm.

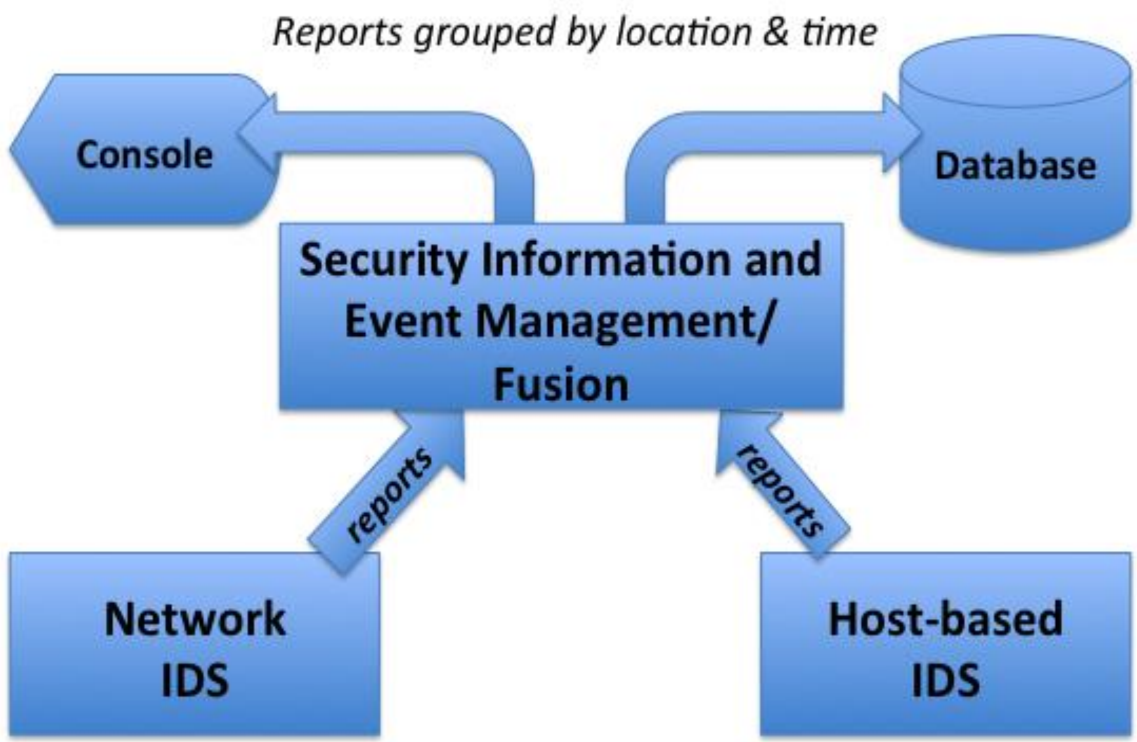

Result:
High false positive rates.
Base rate problems.

Consequence:
Network personnel disable or ignore alerts.

- **Sensors are always on**, so they must be cheap (time and space) and simple.
- **Sensors are not context-aware**:
  - Alerts triggered on known anomalies - can't distinguish software update from attacker recon.
  - Alerts triggered when target not vulnerable.
- **Sensors do not support decisions.**
  - ***Can't ask for additional information to make diagnoses/decisions.***

Active Perception: Moving Beyond SIEM

- Active perception engages senses to resolve observation uncertainty, support action choices
- Active perception uses models of context-specific information needs to
  - Do **context-dependent tuning** of sensors;
    - Incorporate contextual information into sensor interpretation;
    - Follow up initial alerts from cheap, inaccurate sensors with targeted use of expensive, accurate sensors;
  - **Deploy targeted/expensive sensors dynamically**
    - To resolve uncertainties between alternate explanatory hypotheses
    - To gather evidence to enable further interpretation of threats in situations requiring human decision-making

Active Perception Process

- We have an initial situation assessment.
- Evidence suggests revisions to situation assessment: *new hypotheses.*
- Identify what information is needed to confirm or disconfirm hypotheses.
- Allocate computational and sensing resources to optimize *Value of Information* based on decisions we must make.
- Fuse background knowledge with evidence to choose the best hypothesis.
- Choose a response based on the best interpretation(s).
  - This process is *iterative* and *recursive.*

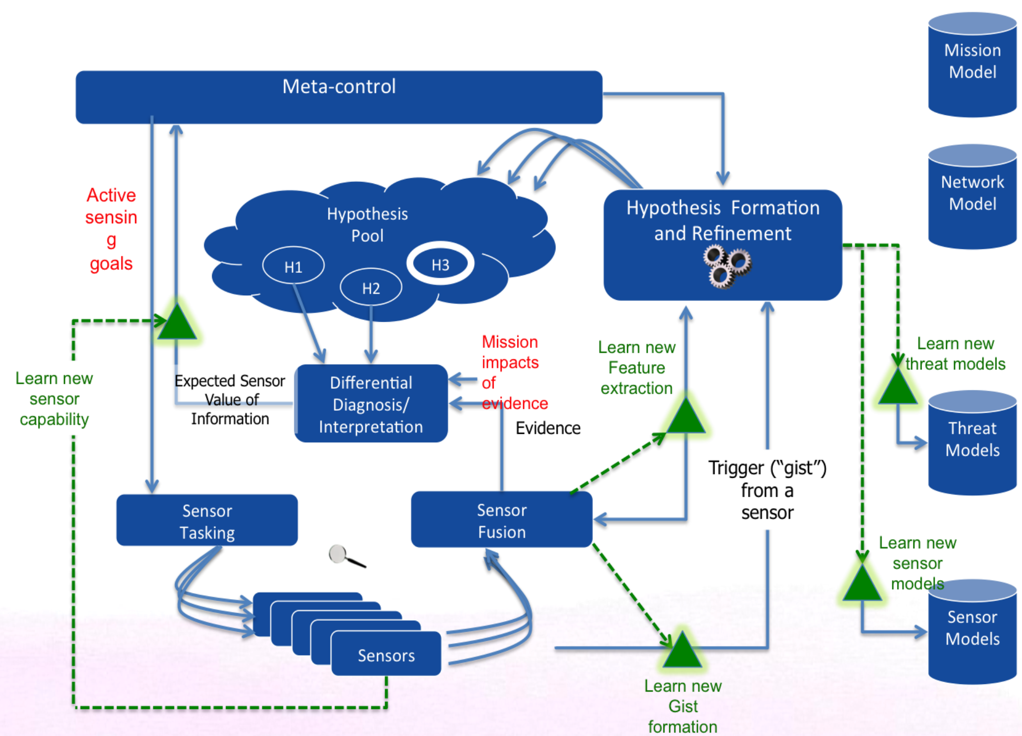

SIFT

Questions

A Priori Delegation and Function Access

- Systems CAN and should behave "fully autonomously" in some circumstances
- But this involves defining a "contract" or "sphere of authority" on which the human can't intrude
  - And to which s/he has to adapt– harder if it's unpredictable
- Choose situations wisely
- **Remember the Sheepdog!**
  - (and Wilson and Neal, 2001)