

Learning for Mission Resilience: A Digital Engineering Framework for Intelligent System-of-Systems Coordination

Taylan G. Topcu, PhD* and Sajad Khodadadian, PhD

*Socio-technical Systems Engineering Lab (STSELab)
Grado Department of Industrial and Systems Engineering, Virginia Tech
email: ttopcu@vt.edu

The Fundamental Challenge of Complex Engineered Systems

- **System of Systems (SoS)** introduce significant challenges to traditional SE [Jamshidi

08', Keating and Katina 11'] due to inherent characteristics [Maier, 98' and Gorod et al. 08']

The Air Force admits the
too much. So it wants to s
Developing and procuring a brand-new nonstealth
Pentagon can defy its entire history of defense spe

Boeing to Delay Some 787 Deliveries After Finding New Glitch

- Planemaker will inspect already-built aircraft for flawed parts
- Boeing, Airbus seen facing ongoing shortages, quality issues



The Fundamental Challenge of Complex Engineered Systems

- **System of Systems (SoS)** introduce significant challenges to traditional SE [Jamshidi 08', Keating and Katina 11'] due to inherent characteristics [Maier, 98' and Gorod et al. 08']
- Despite numerous advances, **SoS remain brittle** [Sapol and Szajnfarder 21', Raz et al. 24]]
 - Extended lifecycles combined with evolving & uncertain operational environments (Oquendo '19; Lambert '17)
 - Dominant new capabilities and adversaries become available (Geurtsen et al., '23)
 - A significant chunk of costs & benefits arise **during sustainment** (Fricke '05; Walden et al., '15)

Digital Engineering (DE) and Mission Engineering Policies as a “Cure”?

- There are **two policy** initiatives calling for a paradigm shift.



Digital Engineering

DE is an integrated approach that leverages authoritative sources of system data and models (i.e., authoritative sources of truth) to manage the system lifecycle from ideation through disposal (DOD, 2018).



Mission Engineering

Mission Engineering (ME) is the deliberate planning, analyzing, organizing, and integrating of current and emerging operational and system capabilities to achieve desired mission effects.

Research Gap: Need for rigorous analytical methods to leverage this vast multi-modal space

- **Three challenges** for analytic methods [Raz et al. 24']:
 - Vast design space
 - Adaptive task allocation and reconfiguration of the SoS
 - Single point optima vs. Multiple configurations achieving a good balance




Machine Learning

Raz et al. 20', Dachowicz et al. 21', Huang et al. 26'




Graph Theory

Zhao 17', Li et.al. 18', and Covello et al. 24'



Interval Algebra & Optimization

Davendralingam et.al. 15'
Raz et al. 24'



Reinforcement Learning

Cao et.al. 25', Chen and Heydari 24' and 26'

A demonstration of this challenge for an Aerial Firefighting Mission



AC#1 Contains



AC#2 Protects
the urban areas

```
Mission Completed
Total cost of burnt area : 141615200.72 Euros
Total emissions of area burned : 1071.74 tons of CO2
Total casualties for area burned : 1826 people
Total burnt area : 649625 m^2
Time to complete mission (excluding response time): 3:17:01
```

Contained, but is this a
resilient approach??

Same tactics, same conditions,
different outcome!

We think SE4AI can help.


Research Questions

- **RQ1:** How can RL methods be synergistically integrated with high-fidelity simulation models for SoS to enable intelligent and robust mission coordination given uncertainties?
- **RQ2:** How well intelligent mission coordination agents trained on a single scenario generalize when exposed to previously unseen mission scenarios?
- **RQ3:** How do theater-level vs. isolated training compare in terms of performance and resilience?


Methodology

Research Setting and Methodology


Stage 1: Digital Model for the Firefighting Mission




Stage 2: Agent Based Firefighting Simulation



Stage 3: Training an Intelligent Mission Coordinator




Stage 4: Evaluation Across Different Scenarios




Research Setting and Methodology

Stage 1: Digital Model for the Firefighting Mission



Stage 2: Agent Based Firefighting Simulation



Stage 3: Training an Intelligent Mission Coordinator

RL Agent (PPO)



Stage 4: Evaluation Across Different Scenarios





Stage 3: Training an Intelligent Mission Coordinator

Stage 4: Evaluation Across Different Scenarios




Findings


Research Questions

- **RQ1:** How can RL methods be synergistically integrated with high-fidelity simulation models for SoS to enable intelligent and robust mission coordination given uncertainties?
- **RQ2:** How well intelligent mission coordination agents trained on a single scenario generalize when exposed to previously unseen mission scenarios?
- **RQ3:** How do theater-level vs. isolated training compare in terms of performance and resilience?

Findings: RL agents can be trained effectively for a given Scenario




Findings: RL agents can be trained effectively for a given Scenario



Research Questions

- **RQ1:** How can RL methods be synergistically integrated with high-fidelity simulation models for SoS to enable intelligent and robust mission coordination given uncertainties?
- **RQ2:** How well do intelligent mission coordination agents trained on a single scenario generalize when exposed to new mission scenarios?
- **RQ3:** How do theater-level vs. isolated training compare in terms of performance and resilience?

Findings: How does the trained agent behave when faced with new scenarios with identical settings?





Alt Scenario #1 Salamis

Takeaways:

1. **More favorable** scenario, benchmarks performing significantly better.
2. Performance of the pre-trained RL agent **is satisfying but only slightly better than Random Tactics.**

Findings: How does the trained agent behave when faced with new scenarios with identical settings?



Alt Scenario #2 Pyrenees


Takeaways:

1. More challenging scenario, benchmarks performing significantly better.
2. **Brittleness:** pre-trained RL agent is not statistically different from Random.


Research Questions

- **RQ1:** How can RL methods be synergistically integrated with high-fidelity simulation models for SoS to enable intelligent and robust mission coordination given uncertainties?
- **RQ2:** How well do intelligent mission coordination agents trained on a single scenario generalize when exposed to new mission scenarios?
- **RQ3:** How do theater-level vs. isolated training compare in terms of performance and resilience?

Theater Level Training



Theater level-training leads to **LOWEST AVERAGE REGRET**



So what?

Discussion

- We answer Raz et al. (2024)'s call for rigorous analytical methods adaptive task allocation for ME.
- RL-based Intelligent Mission Coordination Agents:
 - Outperform **in the specific scenario** they are trained in.
 - Exhibit **brittleness** when exposed to new scenarios; even when all else held constant.
 - Need to understand **which dimensions of scenario variation** (e.g., objectives, asset composition, environmental conditions) most degrade that performance?
- **THEATER LEVEL TRAINING** leads to **lowest average regret**
 - However are **computationally expensive** -> 2 independent 3-layer MLPs (512/256/128, tanh activation); ~165k weights per network
- **Mission uncertainties have strong implications:** Good algorithms can fail, bad algorithms can shine.

Conclusions and Future Directions

- We only focused on operations, need **LIFECYCLE AWARENESS**
 - Portfolio **selection** and mission **planning**
 - System design
 - Strategies and tactics
 - Interplay between the supporting infrastructure and the SoS
- Need to incorporate **disruptions** to evaluate **resilience**:
- Integrating **Explainable AI**
- Need to make **Cost – Performance – Resilience Trade-offs** for ME

Thank you! Questions are most Welcome!



STSELab Sponsors



Sajad Khodadadian, PhD
Email: sajadk@vt.edu

Contact Me!

Email: topcu@vt.edu | www.taylantopcu.com | <https://www.linkedin.com/in/taylantopcu/>

Extras

Formulation of the Deep RL Based Mission Coordinator

State Space S

Action Space A

Reward $R(s_t, a_t, s_{t+1})$

Policy $\pi(a_t|s_t)$

Max. Expected Discounted Sum of Rewards

$$V^{\pi}(s_0) = E_{\pi} \left[\sum_{t=1}^T \gamma^{t-1} * r_t | s_0 \right]$$

A Deep Reinforcement Learning Algorithm: Proximal Policy Optimization Schulman et al. (2017)





Table 1: Description of State Variables


Variable Name (Number of Variables)	Description	Unit of Measurement
Scenario identifier [†] (3)	Identifies the mission scenario in a one-hot encoded way	Dimensionless
Time since detection	Duration since the detection of the fire	Minutes
Temperature	Current environmental temperature	Celsius degrees
Humidity	Current environmental humidity	Percentage
Wind speed and direction (2)	Current wind speed and direction	m/s and degrees
Time to sunset	Duration until sunset	Minutes
Minimum distance to fire line	Shortest distance from the current fire front to the indirect fire-line segment being built now	Meters
Distance to water	Distance from the fleet to the nearest water resource	Meters
Fire-boundary distances (5)	Shortest distance from the fire boundary to urban locations, water resources, highly combustible areas, high-slope areas, and the indirect fire line	Meters
Ignition point (2)	Location of the episode's ignition	X, Y coordinates
Fire center (2)	Centroid of all burning cells in the map	X, Y coordinates
Fire extreme points (8)	Uppermost, lowermost, leftmost and rightmost coordinates of all burning cells	X, Y coordinates
Fire spread angle	Propagation angle of the fastest-spreading cell relative to the centroid	Degrees
Spread ray hit point (2)	Coordinates of the point where spread direction ray (in the direction of propagation from the fire center) hits the map boundaries	X, Y coordinates
Front severity (3)	Per-sector severity of the fire front, for each of the 3 directional sectors around the fire center	Dimensionless
Asset-protection urgency scores (15)	Per-sector urgency of defending urban (VIP), highly combustible, high-slope, water resources and indirect-line objectives threatened by that front	Dimensionless
Aircraft locations (20)	Current location of each controlled aircraft	X, Y coordinates

Assets




	Seaplane	EVTOL
Fuel	Conventional	Electric
Takeoff	From runways/water	Vertical
Cruise Speed (m/s)	100	90
Cruise Altitude (m)	3000	457
Span (m)	18	6
Takeoff distance (m)	2,760	-
Payload capacity (kg)	1,300	360
Endurance (min)	420	27
Range (km)	2,500	150
Takeoff Climb Rate (m/s)	3.81	0.51
Cruise Climb Rate (m/s)	6.96	3.56

Switch Ignition




CAS area




Current area


Expansion of the Mission Space




a) Palisades



b) Pyrenees



c) Salamis



switch-ignition-4 box (400×400 cells) ignitable area (box – water – 150 m urban keep-out)

Training specifics

- We train separate actor and critic networks of size 512, 256, and 128 units with tanh activations, sharing no weights.
- We optimize with Adam at an initial learning rate of 5×10^{-4} , decayed exponentially as the training progresses.
- We use a discount factor of 0.99, a GAE parameter of 0.95, a clipping range of 0.2, and a gradient-norm limit of 0.5.
- Training runs on a single dual-socket AMD EPYC 9454 (Genoa) node providing 96 physical cores and 755 GB of memory, on which the 128 environments run in parallel.

Action Space

- Select Point of Interest (POI)

- Water
- Urban
- Topography
- Vegetation
- Indirect

- Track POI

- Follow fire front
- Indirect

- Suppress

- Direct
- Indirect

