SAND2020-10840C

Efforts to Integrate Autonomous and AI Capabilities into Hypersonic Systems:

Developing Artificially Intelligent Aerospace Systems

PRESENTED BY

Dr. Alex Roesler

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

SAND2020-3187 C

Overview

- Artificial Intelligence and Autonomy for National Security
- Autonomy for Hypersonics
  - Future State Reference Missions
  - Hypersonics of the Future Roadmap
  - “Third Wave” Concepts for Hypersonics

AI and Autonomy for National Security

DARPA's Three Waves of Artificial Intelligence

Handcrafted Knowledge

DARPA's Three Waves of Artificial Intelligence

Handcrafted Knowledge

Statistical Learning

DARPA's Three Waves of Artificial Intelligence

Handcrafted Knowledge

Statistical Learning

Contextual Adaptation

Autonomous Systems

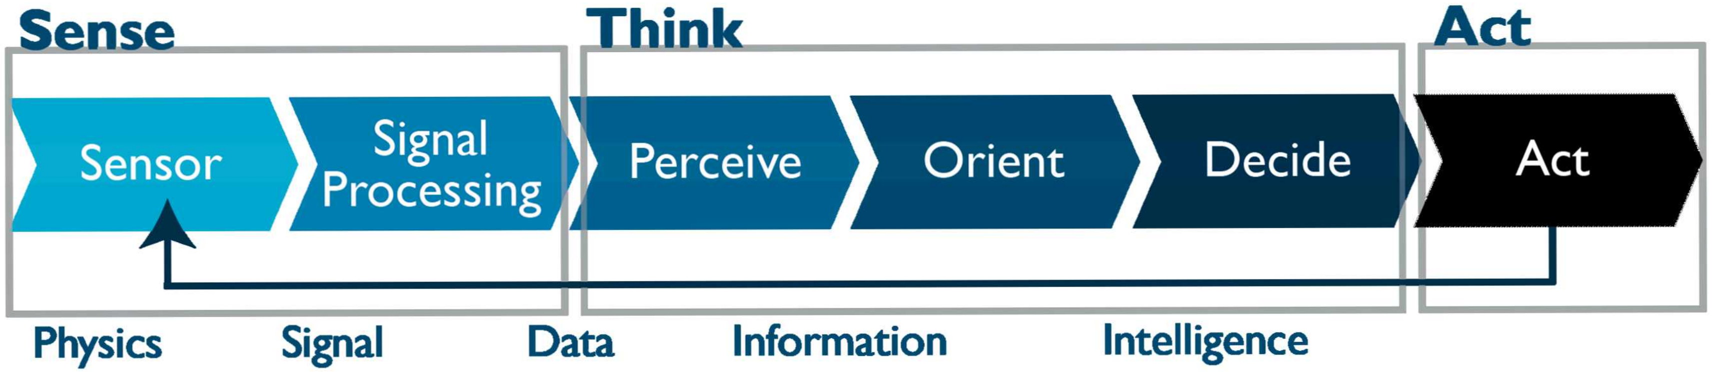

Will AI tech plug-n-play for defense?

- Andrew Ng,
Harvard Business Review

The AI community is remarkably open, with most top researchers publishing and sharing ideas and even open-source code. In this world of open source, the scarce resources are therefore:

Data.

Among leading AI teams, many can likely replicate others' software in, at most, 1–2 years. But it is exceedingly difficult to get access to someone else's data. ***Thus data, rather than software, is the defensible barrier for many businesses.***

Talent.

Simply downloading and "applying" open-source software to your data won't work. ***AI needs to be customized to your business context and data.*** This is why there is currently a war for the scarce AI talent that can do this work.

Commercial

VS

- Structured environments
- Large tolerance for error
- Large labeled training datasets for accuracy
- Can deal with object classes (car, pedestrian, etc.)
- Short-range imaging modalities (e.g. RGB iPhone)
- Can typically rely on GPS and network connectivity, which allows off-board processing and simplifies C2

Defense

- Unstructured, adversarial environments
- Low tolerance for error
- Lack of training data
- Requires precise object identification
- Remote EO/IR/SAR imaging modalities
- Operation in potentially GPS-denied environment with minimal to no network connectivity

Defense applications require different performance characteristics than their commercial counterparts, while managing SWaP and bandwidth limitations.

Autonomy for Hypersonics

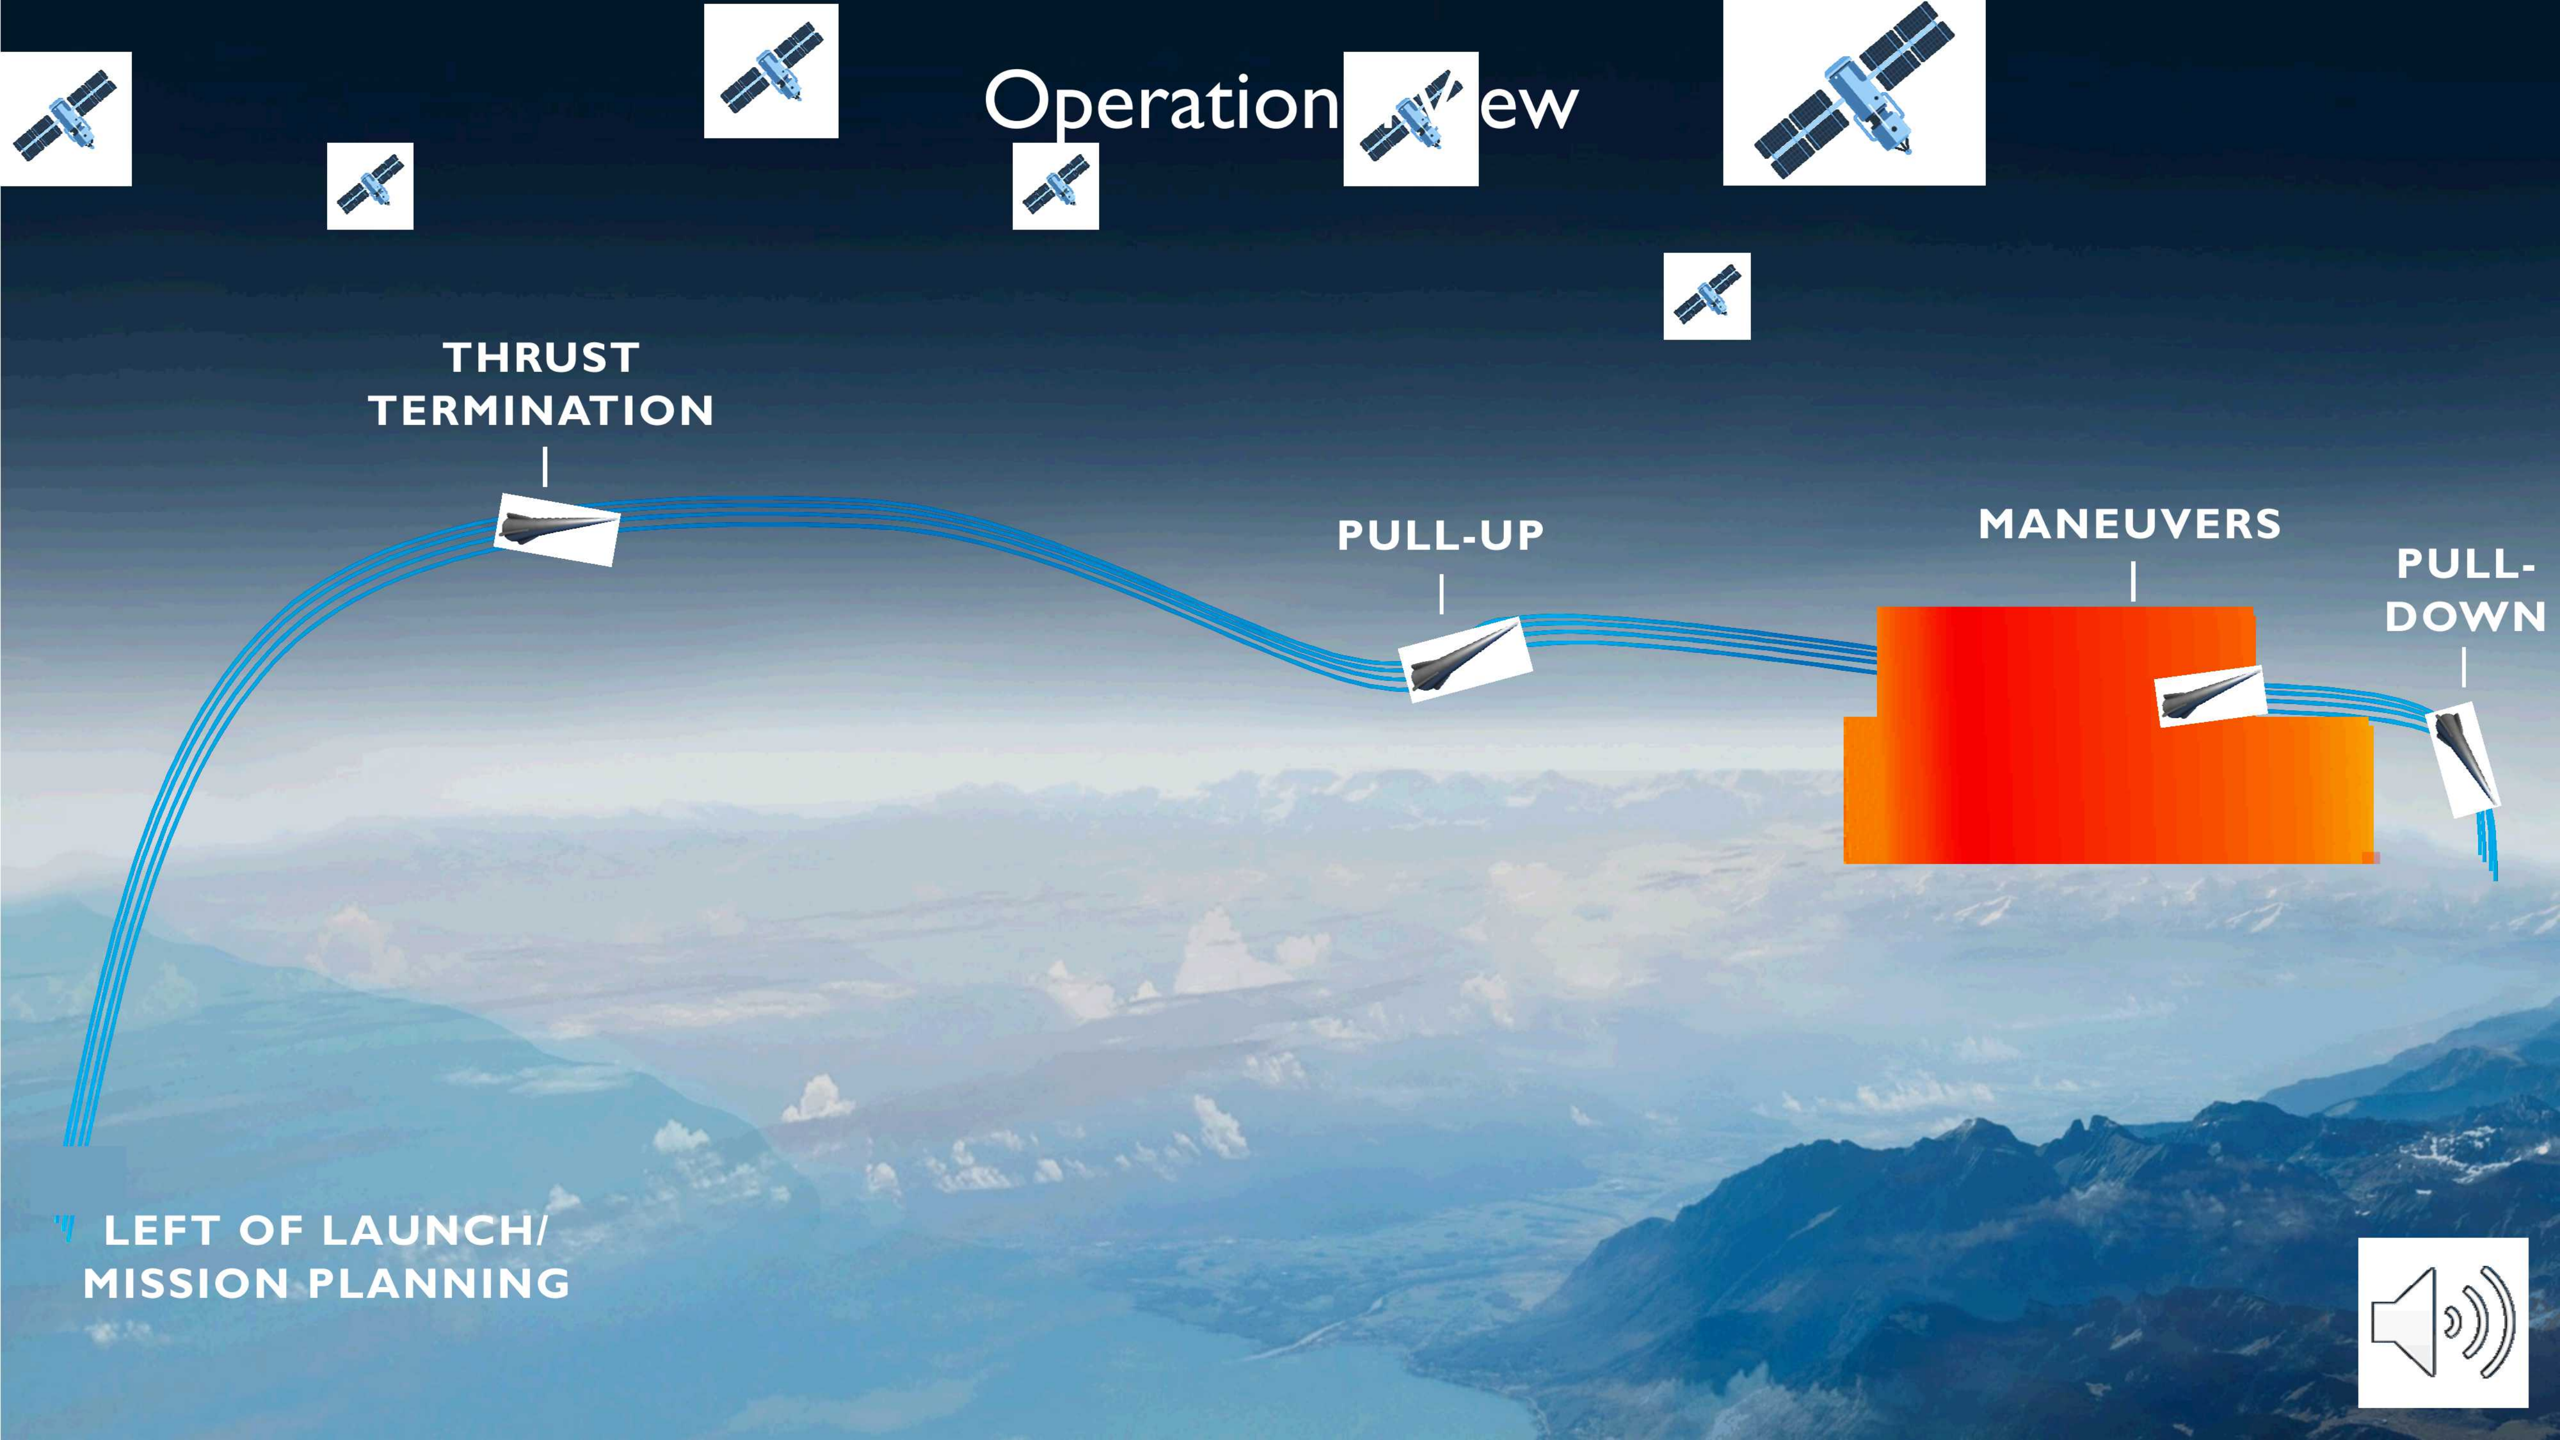
Operation[illegible]ew
THRUST
TERMINATION
PULL-UP
MANEUVERS
PULL-
DOWN
LEFT OF LAUNCH/
MISSION PLANNING

Left of Launch/Mission Planning

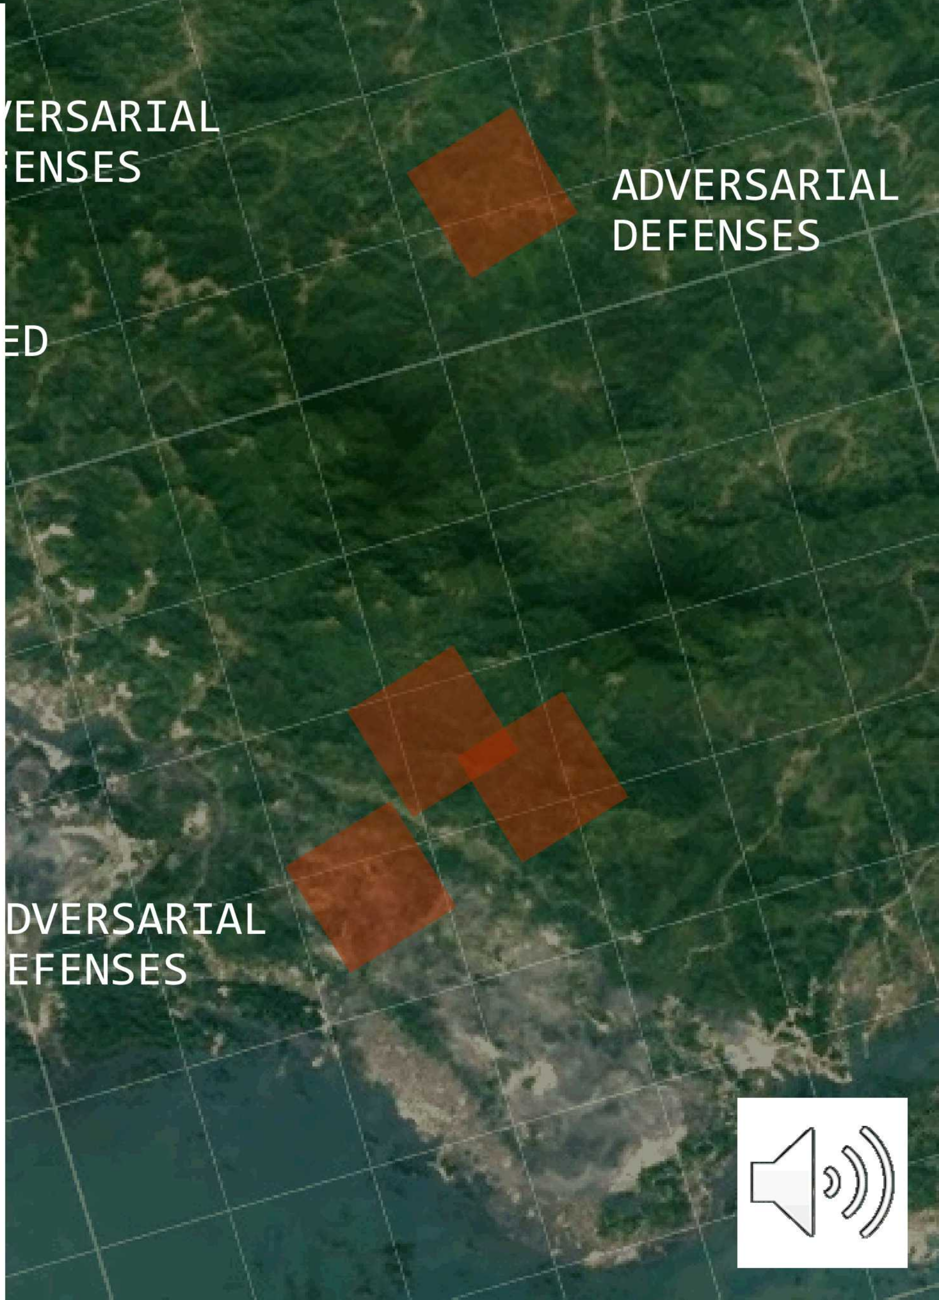

Target Approach Phase

Sandia's Hypersonics of the Future Roadmap

PRE-PROGRAMMED

POSITIONALLY AWARE

POSITION ADAPTING

TARGET HUNTING

SITUATIONALLY AWARE

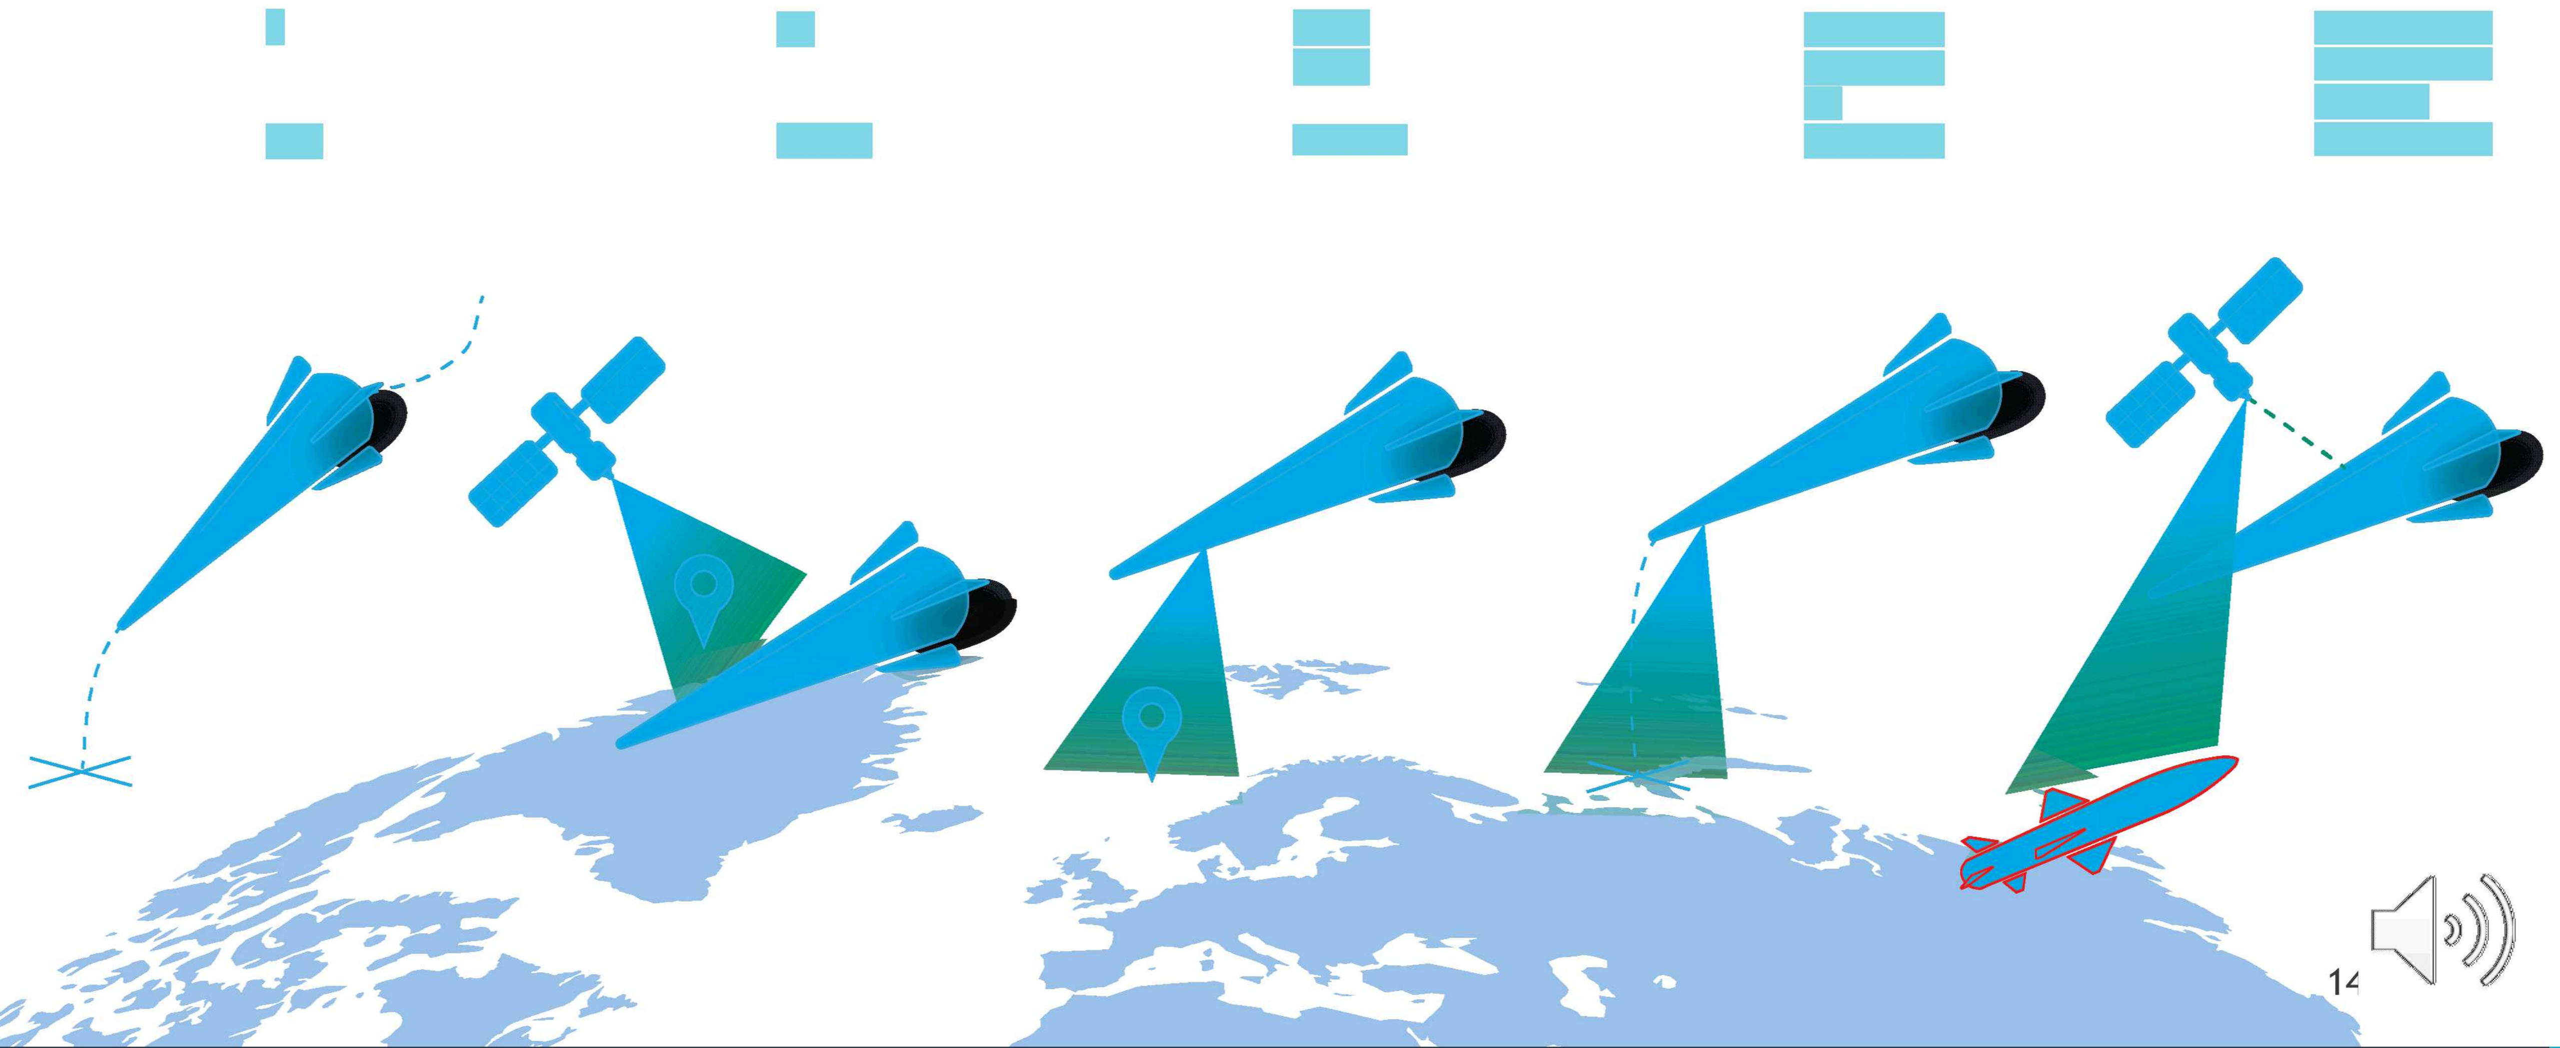

COORDINATE SEEKING CAPABILITY THAT IS ROBUST TO THE GPS CONTESTED ENVIRONMENT

POSITIONALLY AWARE

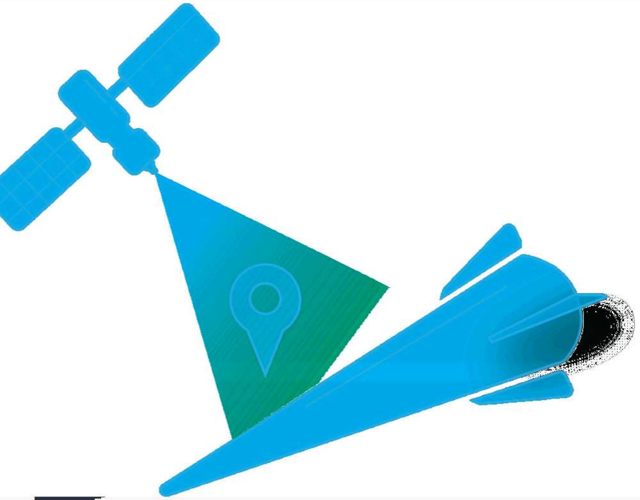

- Senses vehicle position throughout flight
- Delivers warhead to coordinates that are specified prelaunch
- Requires GPS for a substantial portion of flight
- GPS robust against spoofing and modest jamming environments
- Leverages simple sensors to enhance accuracy

RESEARCH CHALLENGES

- Rapid trajectory generation
- Adaptive control algorithms

POSITION ADAPTING

COORDINATE SEEKING CAPABILITY THAT IS ROBUST IN THE NON-GPS ENVIRONMENT

- Senses vehicle position throughout flight
- Initial target coordinates are specified prelaunch
- Leverages GPS when available
- Employs alternate navigation scheme(s) to determine vehicle position
- Accepts updated target coordinates during flight

RESEARCH CHALLENGES

- Non-GPS navigation (sensors and algorithms)
- Mission planning with sensor constraints
- Real-time trajectory generation (RTTG)

ROBUST CAPABILITY TO ADDRESS RELOCATABLE AND MOBILE TARGETS

- Approximate target coordinates and target signature are specified prelaunch
- Employs GPS and/or alternate navigation to localize
- Accepts updated target information during flight
- Employs a terminal sensor(s) to identify target

RESEARCH CHALLENGES

- Left-of-launch mission planning and analysis
- Sensor systems and window materials
- Vehicle perception—Image processing / Automatic Target Recognition (ATR) algorithms
- Sensor-aided terminal guidance and control

SITUATIONALLY AWARE

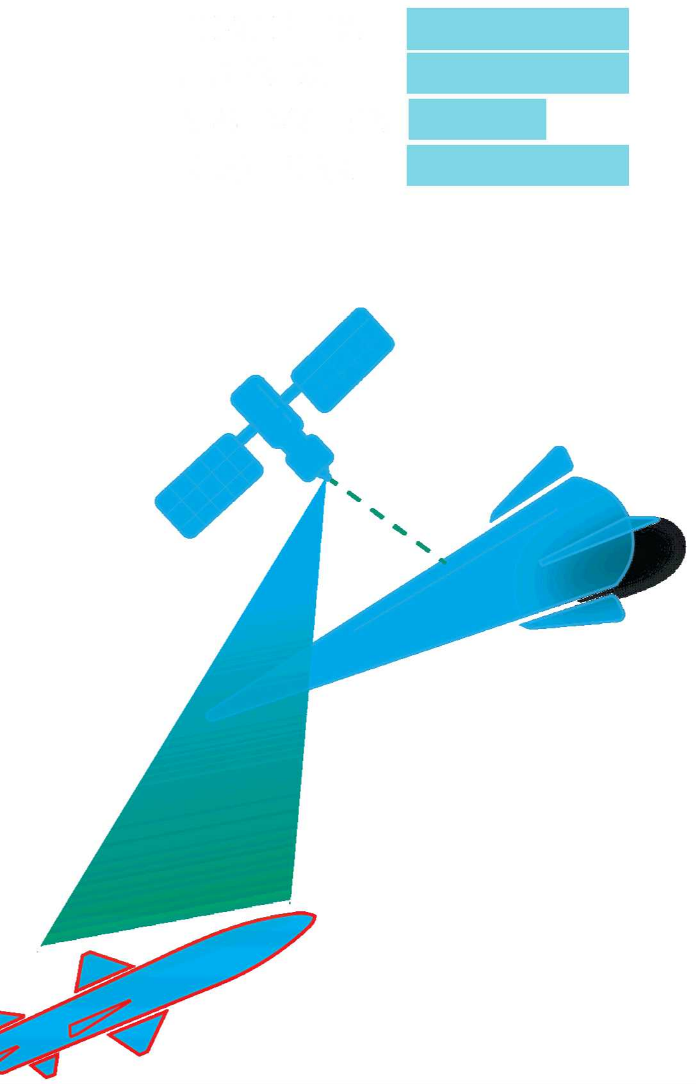

AUTONOMOUS ADAPTATION TO MAXIMIZE STRIKE EFFECTIVENESS

- Senses many elements of its environment
- Fuses data from off-board sensors
- Learns from the experiences of other strike vehicles
- Develops holistic view of mission challenges
- Adapts flight plan for optimal engagement

RESEARCH CHALLENGES

- AI-enabled mission analysis
- Autonomous mission planning (left of launch)
- Human-machine teaming
- Cooperative sensor fusion and exploitation
- Dynamic mission re-planning (right of launch)

Third Wave Concepts for Hypersonics

Lifelong learning by letting the system "dream" and constantly scrimmage

Encoding physics constraints directly into the machine learning

Human-AI symbiosis

Adversarial Reinforcement Learning: Google's AlphaStar

In December 2018, AlphaStar was used to beat a professional StarCraft II player

- StarCraft II is one of the most challenging "Real-Time Strategy" games and demonstrates a huge advancement in reinforcement learning development

Mastering this problem required breakthroughs in several AI research challenges including:

GAME THEORY:

There isn't a single best strategy in StarCraft.
As such, an AI training process needs to continually explore and expand the frontiers of strategic knowledge.

IMPERFECT INFORMATION:

Unlike chess where players see everything, crucial information is hidden from a StarCraft player and must be actively discovered by "scouting".

LONG TERM PLANNING:

Like many real-world problems cause-and-effect is not instantaneous. Games can take up to one hour to complete, meaning actions taken early in the game may not pay off for a long time.

REAL TIME:

Unlike traditional board games where players alternate turns between subsequent moves, StarCraft players must perform actions continually as the game clock progresses.

LARGE ACTION SPACE:

Hundreds of different units/buildings must be controlled at once, in real-time, resulting in a combinatorial space of possibilitie

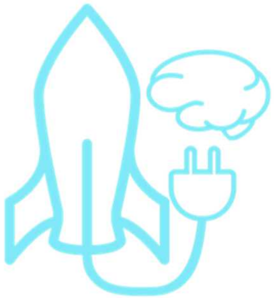

Lifelong Learning

A future where hypersonics are plugged in and constantly training for their missions

- Hypersonic systems are envisioned to penetrate and disintegrate enemy A2/AD systems
- Just like troops constantly train for their mission, so too should AI-based adaptive systems
- Future conflicts will be decided in hours—need to constantly train to be agile and adaptive
- This learning environment can be continually updated based on real-world situation awareness, e.g. regular updates based on space-based imagery analysis

Adversarial Reinforcement Learning

A form of unsupervised machine learning in which you provide the computer a goal, which it seeks to optimize.

- Let the computer generate its own training data
- Often done in a simulated environment (although not always)

- Distance to target
- Angle between distance vector & HGV velocity
- HGV altitude & speed

- Trajectory
- 6DOF control of vehicle

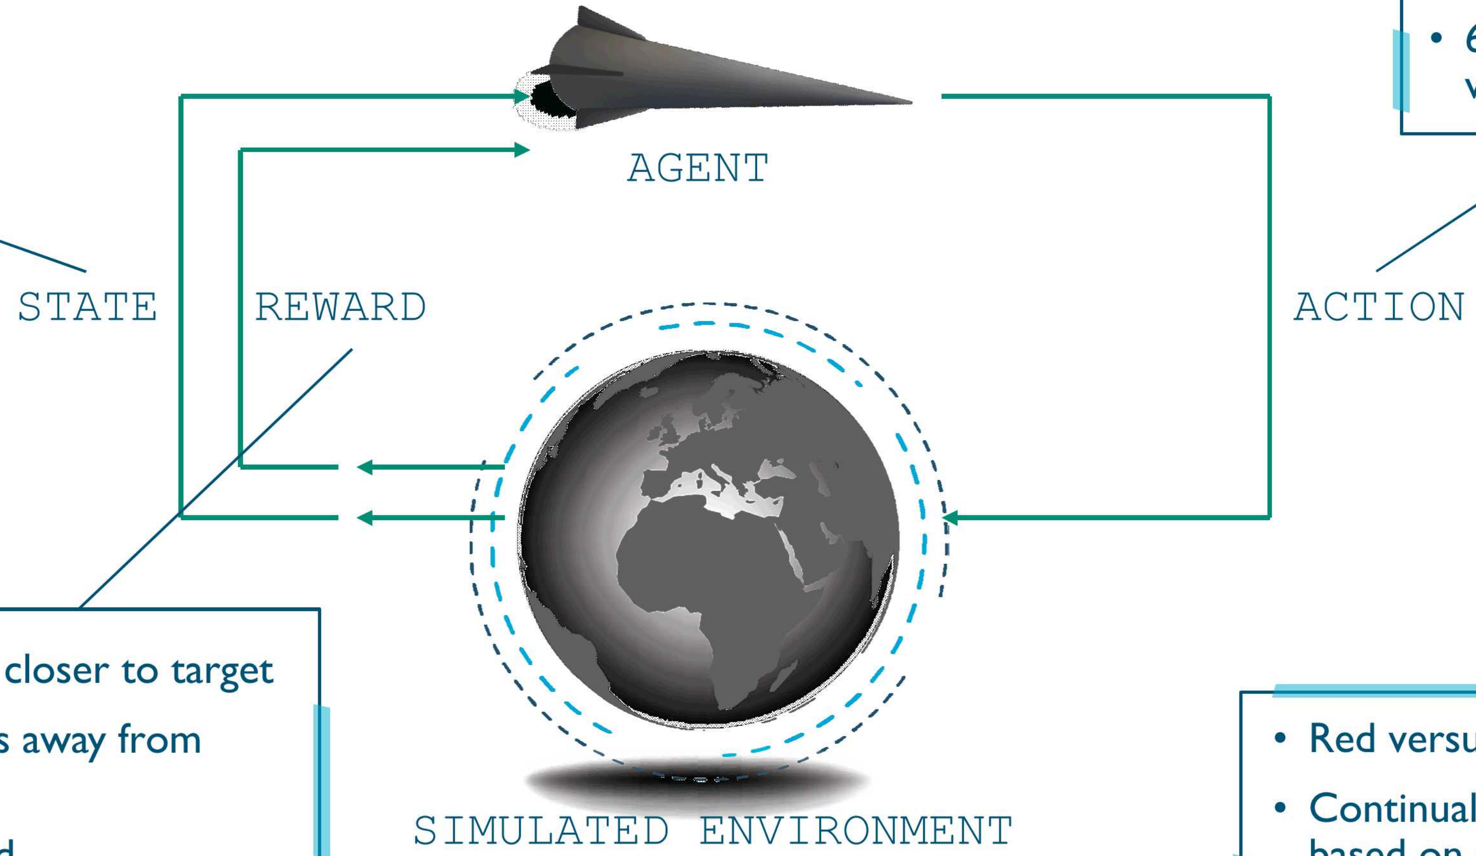

- Positive when HGV gets closer to target
- Negative when HGV gets away from target
- When target is destroyed
- Penalty for time to reach the target

- Red versus Blue
- Continual updates based on new intelligence

Encoding Physics Constraints

Knowledge representations of physics-based constraints will allow machines to explore large action spaces for complex, dynamic systems

- AlphaStar relies on vast amounts of training data developed from perfectly simulated games: during training, each agent experienced up to 200 years of real-time StarCraft play

- This approach maps poorly to complex, dynamic systems

- To successfully apply AlphaStar-like methods to hypersonics, knowledge representations of physics-based constraints are needed that allow the machine to quickly explore huge combinatorial action spaces for complex, real-world environments

Encoding Physics for Real-Time Evasive Maneuvers

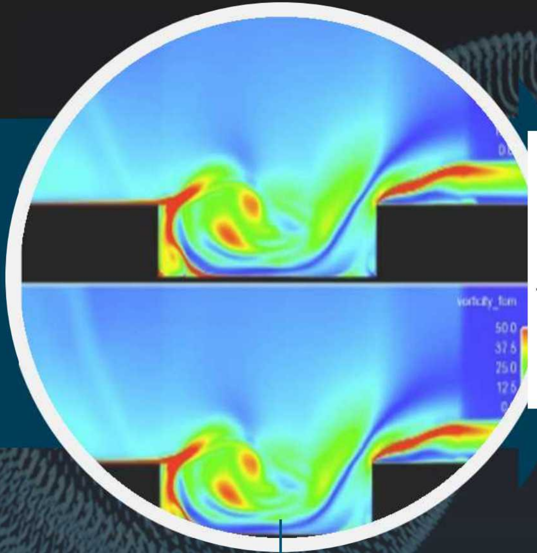

Physics-based Computational Fluid Dynamics reduced order modeling-simulation

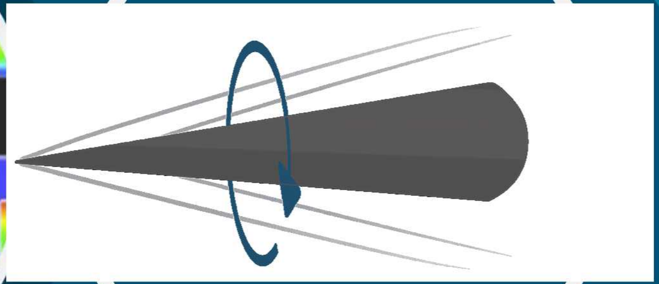

Hypersonic Wind-Tunnel Testing

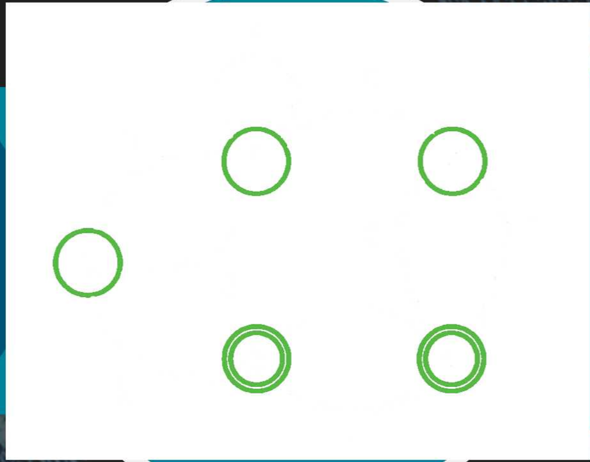

Development of Motion Primitives Library

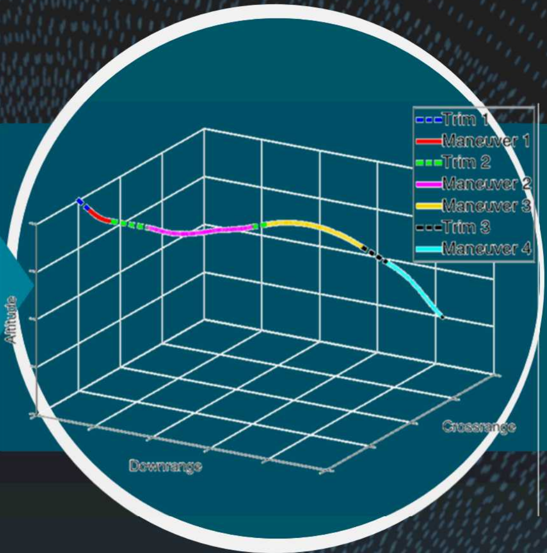

ML algorithms that rapidly generate physics-based trajectories

Human-AI Symbiosis

Humans and machines partner to better understand the complex, rapidly-evolving multi-domain battlefield environment and make the best decisions on when, why, where and how to employ autonomous hypersonic systems

- Human-AI symbiosis is a critical need, in particular for facilitating better decisions in complex, time-critical, battlefield environments

- The challenge: from a multi-domain perspective determine the best options on when, where, why and how to employ autonomous hypersonic systems

Human-Machine Teaming

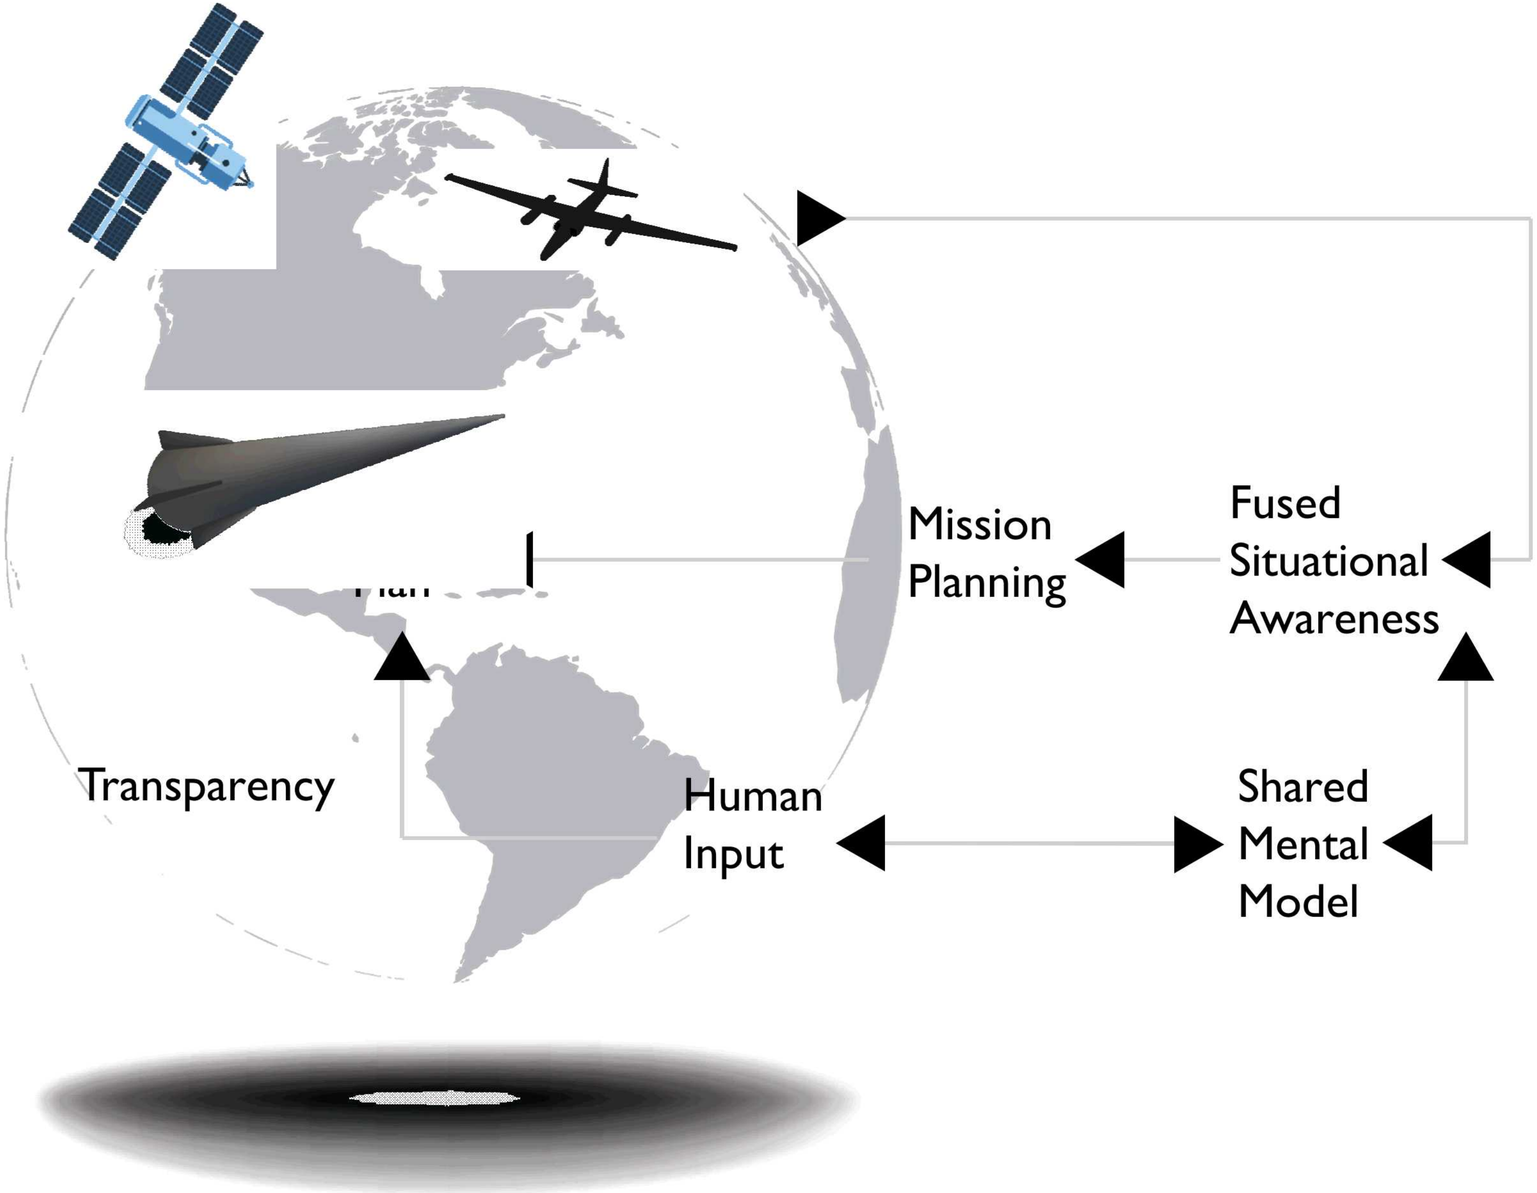

Leverage advances in AI/ML, and data science to provide real-time situational awareness to human operator.

Determine how a human and machine work collaboratively to leverage real-time intelligence to quickly build tailored mission plans

Current Research Collaborations

ILLINOIS

[redacted]

Georgia Tech

Jennifer Hasler + Craig Vineyard
Neural-Inspired Approaches and Implementations for Automatic Target Recognition

Ani Mazumdar + Katya Casper
Hypersonic Wind Tunnel Test Bed for Fault-Tolerant and Adaptive Control

Evangelos Theodorou + David Kozlowski
An Optimization and Robust Control Technique for use in Flight Control Design for Hypersonic Vehicles

Jonathan Rogers + Julie Parish
Real-Time Evasive Maneuvers in Contested, Uncertain Environments

Panos Tsiotras + Bart von Bloemen Waanders
Hyper-Differential Analysis to Mitigate Uncertainties for Control of Hypersonic Vehicles

TEXAS – The University of Texas at Austin

Ufuk Topcu + David Kozlowski
An Optimization and Robust Control Technique for use in Flight Control Design for Hypersonic Vehicles

Todd Humphreys + Julie Parish
Tightly Integrated Navigation and Guidance for Target Acquisition

Karen Willcox + Patrick Blonigan
Rapid High-Fidelity Aerothermal Responses with UQ via Reduced-Order Modeling

Renato Zanetti + Scott Jenkins
SAR Image Formation for Navigation in GPS-Denied

Maruthi Akella + Mike Grant
Autonomous 6DOF RTTG for Highly Constrained Hypersonic Missions

PURDUE UNIVERSITY

[redacted]

TEXAS A&M UNIVERSITY

John Valasek + Julie Parish
Tightly Integrated Navigation and Guidance for Target Acquisition

Johnny Hurtado + Jason Searcy
Magnetometer-Aided GPS-Denied Navigation

Johnny Hurtado + Julie Parish
Tightly Integrated Navigation and Guidance for Target Acquisition

Johnny Hurtado + Lisa Hood
Surrogate-Constrained Vehicle Modeling to Enable Rapid & Real-Time Trajectory Generation

THE UNIVERSITY OF NEW MEXICO

Meeko Oishi + John Richards
Autonomous Multi-Platform Sensor Scheduling

Don Hush + Mary Moya
Improving Model-based Training of Automatic Target Recognition for Rapid Response to Evolving Threats

NM STATE

Hyeongjun Park + Bethany Nicholson
Real-Time, Nonlinear, Optimization-Based Control Algorithms for Hypersonics

Liang Sun + Michelle Hummel
Justification and Transparency in SAR ATR using AI Rule Extraction and Fused Classification

THE UNIVERSITY OF ARIZONA

Roberto Furfaro + Bethany Nicholson
Real-Time, Nonlinear, Optimization-Based Control Algorithms for Hypersonics

KANSAS STATE UNIVERSITY

Bill Hsu + Jason Searcy
Magnetometer-Aided GPS-Denied Navigation

Visit **autonomy.sandia.gov** for additional info

Summary

- Autonomy and AI can provide a transformational capability enhancement for hypersonics
- A number of incremental advances are possible that increasingly leverage approaches across all "Three Waves of AI"

Handcrafted Knowledge | Statistical Learning | Contextual Adaptation

Questions?