An Ontology for Active and Passive Aerial Drone Threat Automatic Plan Recognition

Ronald P. Loui, Ph.D.
Department of Computer Science
University of Illinois
Springfield, IL
rloui2@uis.edu

Josh Smith
Department of Computer Science
University of Illinois
Springfield, IL

***Abstract*—This paper initiates a discussion on the design of terms, features, and descriptors that would support machine learning for automated plan recognition of drone and drone swarms engaged in threatening activity. A few prototype aerial missions for drones are discussed and semantic markers, such as distance and line of sight to potential targets, mirrored motion, path and position optimality, coordination, and formation, are noted. This semantic description of motion in terms of objectives and capabilities contrasts with a naïve description of motion in a 3d coordinate system without reference to targets; terminology is the first step in automated anomaly detection analytics. The paper further discusses active plan recognition, which selects interventions in order to force the drone or swarm to reveal its intentions. Analogies to, and distinctions from, two-dimensional active plan discernment, e.g., stalking, tailing, pursuing, and intercepting, are given.**

Keywords—drones, swarms, drone threats, plan recognition, description, databases, analytics, machine learning, artificial intelligence, interactive, stalking, mission, reconnaissance, pursuit, attack, motion semantics, interpretation, intelligence, threat assessment.

I. Introduction

A. Defending Against New Terrorist Drone Threats

This paper takes a step toward an AI and machine learning contribution to the problem of determining potentially malicious intentions of drones by analysis of in-flight behavior. This contribution is the discussion of ontological concerns, which is a prerequisite for applying well known algorithms, such as deep learning. The main message is that modeling of behavior should be done in terms relative to the drone's target, with explicit consideration of the threatening side's offensive mission requirements and objectives.

The defensive problem of asymmetric threats from unmanned aerial vehicles, drones, is a novel, tactical disruption. Offense and defense have been thrown out of balance by the availability, low cost, easy deployment, potential reprogrammability, multi-modal flight control, speed, maneuverability, inevitably increasing payload size, and small signature of commercial drones. So far, high tech forces have been the laudable beneficiaries of UAV offense, but there is a considerable new engineering problem to be solved to protect targets when one is on the defensive.

Our particular interest is the prevention of terrorist attacks on high value public targets in domestic settings (DHS concerns), but the approach taken in this paper should apply equally to force protection while operating in overseas hostile environments (DoD concerns). Hence, whether the target is a traveling President, a visiting Pope, a Navy vessel in a foreign port, or a temporary command center deployed in theater, small drones may soon be large enough, piloted well enough, and coordinated perniciously enough to cause concern. Aerial drone-threat prevention may be part of the airspace control responsibility, not merely the ground force protection obligation.

We assume that kinetic and electronic countermeasures under development will grow in effectiveness: jamming, blinding, GPS fencing and spoofing, mandatory "unhackable" manufacturer fail-safes, microwaving, SkyWall nets, shotgun scatter, air superiority drones for capture, collision, and other defeat, Naval LaWS, Boeing's HEL MD 10kw laser cannon, etc. Sensing is developing more quickly: DroneShield's audio signatures, DroneDetector's / DroneTracker's multi-sensors, DeDrone's video analytics, DroneZon's LIDAR, CellAntenna's radio frequency interception, CGG's seismic sensing; many of these appear to have had successful real-world deployments. No doubt, arrays of aerial drone-sensors and other resource-intensive strategies will be available for some temporary target protection (the modern interpolation of air cover patrol for fleets and air balloon obstacles for beach landings).

One potential contribution of AI in this arena is the discernment of intent. Determining intent is an analytic, real-time classification problem. Plan recognition has long been a fruitful subject of AI research: it involves modeling of intentions and goals, understanding the epistemology of the situated agent (what it knows and what it can see, hear, or otherwise sense). It requires probabilistic reasoning algorithms. It can benefit from machine learning. But it cannot even begin without discussion of the domain ontology

978-1-5090-3441-3/16/$31.00 ©2016 IEEE

of the problem: what level of abstraction at which to collect data, describe agents, and understand connections through inference.

One of the main lessons of the fundamental work of Agre [1, 2], which we follow here, was that multi-agent adversarial real-time pursuit and targeting activities require a description of the world in relative terms: what is the closest threat, which is the easiest to pursue, which targets are in range, what are the priority events or salient activities, etc. Using real-valued coordinate systems is too low a level of abstraction to be useful when reasoning about action or intention. It is tempting to describe aerial motion in terms of paths through grids, waypoints, collecting velocities and altitudes, but the next semantic level yields actionable inference. Our semantic abstractions are motivated by such questions as:

- What is gained by being in a place at a time?
- What capability is added or maintained with a particular motion?
- What motion supports the goal of conducting surveillance?
- What motion supports a coordinated attack?
- Why did the drone respond in the way that it did, when provoked?

B. *UAV Threat Assessment: Early Warning, Deception, Mobile Targets*

The pragmatics of drone intention recognition are not currently problematic: the cost of false positives is low largely because the cost of drones is low. Most drones can simply be shot down or electronically defeated upon entering restricted airspace with little consequence, no questions asked. The main costs have to do with cost of deploying the countermeasure, potential collateral damage from the interaction, disruption of legitimate drone flight purposes (e.g., the defeated device might be part of a different security team's defensive scheme), and legal or socio-political consequences. A lone incoming drone is as unambiguous as an incoming missile or mortar shell. However, there are important current and future scenarios where drone intention recognition can be an improvement.

First, assessing with high probability what a drone is doing provides early warning of attack potential. This is especially important when actions are deliberately deceptive; some attacks do not look like attacks, until the moment of attack, so suspicion must be assigned early. Insider, treasonous, and terrorist actors are very likely to mask themselves with benign behaviors in order to preserve or acquire critical capabilities, such as proximity or preferred angle of attack.

Moreover, when drones are part of a multi-stage, sequential or parallel coordinated attack, knowledge of that fact can reduce vulnerabilities exposed during the initial response.

Importantly, legal and ethical cover can be gained by just the attempt to divine the intentions of a drone. This is part of due diligence, even when countermeasures are triggered on low probability threats.

Finally, looking forward, aerial drones may become essential parts of infrastructure, performing important functions in complex systems. It may in the future be impractical to shut down airspace to all drone traffic, at the distances and altitudes desired. This is especially relevant when protecting mobile targets, and the defensive challenge is worse when the target's motion is not known in advance.

II. AUTOMATION

A. *Plan Recognition*

The basic idea of plan recognition is to have a library of plans that can be attributed to an actor. Rather than infer in a purely probabilistic manner that observations correlate with candidate purposes, plan recognition hypothesizes each known purpose in turn, abductively testing observations for coherence with hypothetical desires. It seeks explanations, not just correlations. In the design of observational primitives for learning and classification, knowledge of potential plans also determines relevance and importance of surface measurements. Hybrid abductive (constraint-satisfaction/best-explanation) and inductive (probabilistic) methods are generally the strongest.

The earliest work in this area was done in natural language dialogue understanding [6, 7, 15], but the principles are the same for non-verbal actors [5, 8, 11, 16, 23, 24, 26].

B. *Knowledge, Rules, Machine Learning, and Ontology*

An AI classifier requires an input layer, or a training set of data described in terms that properly hint the machine learning [21]. Theoretically, one could start with an observational terminology that is unhelpful and contains no human insight, then through training, a hidden layer of concepts can eventually emerge. This would be roughly equivalent, functionally, to the human-tailored terminology. In some cases, it could even be more discerning (or appear to be so, when there are small test sets and possible over-fitting). However, in settings where human knowledge of important details can be brought to bear, in addition to emergent concepts, classification can be done more quickly and robustly with shorter development cycles if the human knowledge is "baked into" the terminology.

The human-tailored or human-hinted ontological starting point is also the terminology that might be used in a rule based expert system, or in a probabilistic reasoning system.

Although much is made of the algorithms for inference and learning in AI, the domain ontology is where work needs to be done, and where the power of information processing resides.

C. Naive 3D Ontology of Location and Motion

The first main point of this paper is that one cannot expect to automate intention recognition, whether in a rule-based or learning-based approach, using the naïve ontology of 3d coordinate location and velocity, or using GPS location. This terminology is too low level. An aerial threat's appearance in named parts of an airspace might be more useful with respect to some capabilities, but does not have the detail needed to make use of rapidly changing position, attitude, path constancy, and response to inputs. A history of xyz-t coordinates does not have the granularity of measurement, even for the simple determination of a repeating path or semi-stationary position.

What is needed is a language that takes into account the requirements of various missions that could be attributed to drones.

D. Practical Semantically Relevant Ontology

Our ontology is based on four main concerns:

- *Targets*

We start with the idea that there is a predetermined target, that the target is dynamic in location and in presentation, and that the drone or swarm is uniquely interested in the target. Hence, most of the observational language is a measurement in relation to the target.

- *Spaces*

Action takes place in space, which includes distance from target, time to target, angle, and path. Altitude is not a distinct dimension here, except as it might relate to sensing. For aerial drones, terrain is not relevant except when low altitude impacts the path to the target, or introduces intermittent visual contact.

- *Capabilities*

Capabilities are both sensory and physical. Attention is paid to the continuity of capability, and the effort the drone displays in maintaining or restoring capability. A capability-path in time/space is inherently more interesting than a location-path in-time/space.

Capabilities are derived from missions, which may include:

- delivering/firing
- intercepting
- tailing
- stalking
- finding
- fixing (i.e., fixing the position of the target)
- surveying
- information-gathering

.

- *Responses*

Events that might frustrate desires, and cause a drone to reposition (plan repair) are crucial. They might be chance events or adversarially-generated events. To reveal the drone's sensing objectives, one can actively conceal or reposition the target, block the line of sight, or even introduce a decoy or other ambiguity. To reveal the drone's attack objectives, interpose, crowd so as to reduce freedom of movement, move the target, or take cover. Other active probing is possible – this is just a start.

II. Multi-Agent Air and Ground Threats / Swarms and Forward Observers

Drones that are members of teams, swarms, or other coordinated action are of particular concern.

One problem is that coordination can occur across surprisingly long time spans: a second strike can occur hours after a first strike. Also, coordination can be mutli-modal: a human forward observer with a laser can mark a target for an aerial drone, for example, or a fully autonomous swarm can follow the path determined by a remotely controlled lead drone. Determining members of a team is difficult because roles can be parceled into small individual contributions. For example, creating distractions seems innocent, but may be precursor to subsequent lethal action. Or a suicide drone might have only the purpose of revealing a defense's strategy and deployment locations.

One way that coordinated actors can reveal their intentions is through unusual attachment to a specific time and place. Another more obvious way is through the maintenance of relative position, or even obvious display of aerial formation.

III. Active Revelation of Intentions

As noted earlier, many intentions can be revealed by actively probing responses through specific interactions. Our aim here is not to invent all ways drones can be pricked, but to note that both active and passive information gathering should be reflected in a vocabulary for describing drone behaviors. It is a separate question how one might automate a protocol for probing aerial actors to reveal their intentions. The scope of the ontology should support automatic classification regardless of whether observations and measurements are made passively or actively.

IV. An Initial Observational Terminology For Drone Threat Classification

A. Reconnaissance

Maintains-line-of-sight
Maintains-watchful-min-distance-and-max-distance
Matches-changes-in-motion
Responds-to-blocking-with-angular-deviation
Responds-to-sensory-obscuration
Responds-to-target-concealment
Exhibits-ground-survey-pattern
Follows-target-intermittently

B. Attack

Maintains-minimum-distance
Maintains-altitude
Maintains-straight-line-clear-path
Follows-target-constantly
Responds-to-blocking-with-robust-deviation
Responds-to-target-ambiguity-with-choice-hesitation
Exhibits-load-bearing-dynamics

C. Intercept

Accelerates-to-vector-to-target
Approaches-constantly
Approaches-quickly
Dives
Responds-to-blocking-with-vector-to-target-correction
Responds-to-target-ambiguity-with-hesitation
Ignores-target-concealment

D. Coordination

Maintains-formation
Maintains-angle-or-distance
Maintains-distance-to-target
Exhibits-sudden-parallel-acceleration
Exhibits-serialized-flight on-same-path
Exhibits-mutual-statistical-anomaly-of-motion
Exhibits-mutual-statistical-anomaly-of-configuration
Marks-target

E. General Anomaly (Individual)

Deviates-from-known-path
Shows-deviant-angle-to-target
Shows-deviant-distance-to-target
Shows-deviant-claiming-of-freedom-of-motion
Shows-deviant-velocity
Shows-deviant-acceleration
Shows-deviant-altitude-change

F. Different Pattern (Social)

Deviates-from-others-on-site
Shows-deviant-communications
Shows-deviant-dynamic-envelope
Deviates-from-path-boundary-history
Ignores-normal-refocusing-of-attention

V. Some Discussion

The concepts introduced in A might also be relevant to B, and so forth; these are neither exhaustive nor exclusive.

These are mid-level concepts. Some admit immediate refinement into subclasses, e.g.,

Responds-to-sensory-obscuration:

>> Responds-to-visual-obscuration
>> Responds-to-audio-obscuration
>> Responds-to-EM-obscuration

and the responses could be orthogonal:

>> Responds-by-waiting
>> Responds-by-changing-angle-to-target
>> Responds-by-closing-distance

Similarly, subclasses could be defined for

Ignores-normal-refocusing-of-attention:

>> Ignores-planned-event
>> Ignores-chance-event
>>>> Ignores-visual-spectacle
>>>> Ignores-weather-event
>> Ignores-waypoint-of-interest
>> Ignores-popular-viewpoint

although it may not matter which particular subclass is exhibited; it may matter more that it shows many kinds of subclass behaviors.

There are categories above this abstraction layer as well. These mid-level terms feed into the determination of upper level concerns such as **Fixation** on a subject, and mission **Accomplishment**.

Many of the terms admit degrees and durations.

VI. Comparison to 2D Motion

Insight can be gained by comparing the considerations here to tracking, intercepting, and firing in the 2D manifold or plane. The 2D situation is more familiar, providing many opportunities for comparison (e.g., in both cases, a dynamic target and a non-deterministic environment force the adversary's revelation of intention). But there are important differences.

A. Occlusion

In 2D, there is typically more terrain clutter, so lines of sight are repeatedly occluded. In 3D, clouds would be the analogues. Electromagnetic and other non-visual forms of sensing may be frustrated more often by interference from non-objects in 3D such as EM fields, heat, and precipitation. Backlighting may be a problem in 2D, which does not impact aerial drone sensing in 3D, but glare and ground-lighting may frustrate 3D visual sensing at a distance more easily than any 2D counterpart. One 3D tactic would be to use surface-level occlusions, especially concealment of certain angles, since these have low cost.

Environmental occlusion in the 2D situation provides more opportunities for analytics than in 3D environments. But the non-chance, generated tactics are more numerous in 3D.

Flying counter-drone vehicles close to the potential threat provides opportunities for other interaction, beyond occlusion (e.g., crowding and delaying). The more interruption, the more induced revelation.

B. Dynamics of Flight: Distance and Discernment

In the 2D setting, distances are more significant with respect to both sensing and attack/approach, than in 3D.

The nature of flight makes possible a quick closing of 3D distance. One of the disruptions introduced by drones is their ability to hover, loiter, and control angles with great precision. Past 3D aerial devices were inferior to 2D in this respect. Now they are markedly superior with respect to control of position. For example, angles can be improved almost at will. For the analyst interested in divining a drone's purpose, the use of this superior ability betrays the intentions of the one who uses the ability.

Visual detail at distance is likewise different in 2D and 3D because of the operating distances.

VII. Conclusion And Future Work

This work may at first appear to be more like a meditation on the BDI (belief/desire/intention) psychology and social psychology (sociopathy, deviation from norms) of drones than AI engineering. This is the nature of plan recognition (see, e.g., [26]).

Ontological design, in its initial stages, is indeed an exercise in understanding the domain. It proceeds by modeling what is potentially important. Data can't be tracked and acted upon unless it is defined and collected. There is always room for addition to a set of concepts, especially when computing and sensing are not limiting. The trick, especially for machine learning, is to find the abstraction level at which semantics and pragmatics meet. There are a lot of xyz-t data points available for drone tracking, when they can be sensed. But no one wants to be inundated with *big meaningless data*.

Thus, we have sketched our initial thoughts on how analytics of drone behaviors might be automated, focusing on the terminology that might provide useful insight into intention.

Anomaly detection usually benefits not just from modeling threats, but also from an accounting of non-threat activities. To aid machine classification, work should further be done to fill the library of non-threat intentions. However, given the nature of deceptive attack, finding markers of routine activity that are not easily spoofed by deceptive adversaries is challenging.

Threat and intention recognition are subjects well known to defense research organizations. What is changing is the likelihood that the low-altitude air spaces could soon contain more actors with new tactics, and with changing costs of disruption. Because of this, the intelligence requirements supporting command decisions may be increasing. Hence, the groundwork for drone behavioral analytics should be started.

Acknowledgements

Lucinda Caughey and Mahathi Karnamadakala contributed thoughts to this discussion.

References

[1] Agre, Philip E. *The Dynamic Structure of Everyday Life*. No. AI-TR-1085. MIT AI Laboratory, 1988.

[2] Agre, Philip. *Computational Theories of Interaction and Agency*. MIT Press, 1996.

[3] Beard, Randal W., Timothy W. McLain, Michael A. Goodrich, and Erik P. Anderson. "Coordinated target assignment and intercept for unmanned air vehicles," *IEEE Transactions on Robotics and Automation* 18:6, 2002.

[4] Boury-Brisset, Anne-Claire. "Ontological engineering for threat evaluation and weapon assignment: a goal-driven approach," *IEEE Conference on Information Fusion*, 2007.

[5] Bui, Hung Hai. "A general model for online probabilistic plan recognition," *IJCAI*, 2003.

[6] Carberry, Sandra. "Techniques for plan recognition," *User Modeling and User-Adapted Interaction* 11:1-2, 2001.

[7] Charniak, Eugene, and Robert P. Goldman. "A Bayesian model of plan recognition," *Artificial Intelligence* 64:1, 1993.

[8] Cuppens, Frédéric, Fabien Autrel, Alexandre Miege, and Salem Benferhat. "Recognizing Malicious Intention in an Intrusion Detection Process," *Hybrid Intelligen Systems Conference*, 2002.

[9] Evers, Lanah, Twan Dollevoet, Ana Isabel Barros, and Herman Monsuur. "Robust UAV mission planning," *Annals of Operations Research* 222:1, 2014.

[10] Gaudiano, Paolo, Eric Bonabeau, and Ben Shargel. "Evolving behaviors for a swarm of unmanned air vehicles," *IEEE Swarm Intelligence Symposium*, 2005.

[11] Geib, Christopher W., and Robert P. Goldman. "Plan recognition in intrusion detection systems," *DARPA Information Survivability Conference & Exposition II*, 2001.

[12] Girard, Anouck R., Adam S. Howell, and J. Karl Hedrick. "Border patrol and surveillance missions using multiple unmanned air vehicles," *IEEE Conference on Decision and Control*, 2004.

[13] Jang, Myeong-Wuk, Smitha Reddy, Predrag Tosic, Liping Chen, and Gul Agha. "An actor-based simulation for studying UAV coordination," *European Simulation Symposium*, 2005.

[14] Johansson, Fredrik, and Goran Falkman. "A Bayesian network approach to threat evaluation with application to an air defense scenario," *IEEE Conference on Information Fusion*, 2008.

[15] Kautz, Henry A. "Formal Theories of Plan Recognition," *Rochester Planning Workshop: From Formal Systems to Practical Systems*, Technical Report 284 April 1989.

[16] Le Guillarme, Nicolas, Abdel-Illah Mouaddib, Xavier Lerouvreur, and Sylvain Gatepaille. "A Generative Game-Theoretic Framework for Adversarial Plan Recognition," *Journées Francophones sur la Planification, la Décision et l'Apprentissage,* 2015.

[17] Little, Eric G., and Galina L. Rogova. "An ontological analysis of threat and vulnerability," *IEEE Conference on Information Fusion,* 2006.

[18] Little, Eric, Kathryn B. Laskey, and Terry Janssen. "Ontologies and probabilities: working together for effective multi-INT fusion," *Ontologies for The Community*, 2007.

[19] McLain, Timothy W., Phillip R. Chandler, Steven Rasmussen, and Meir Pachter. "Cooperative control of UAV rendezvous," *IEEE American Control Conference,* 2001.

[20] Meersman, R. A. "Semantic ontology tools in IS design," In *Foundations of Intelligent Systems*, Springer Berlin Heidelberg, 1999.

[21] Michalski, Ryszard S., Jaime G. Carbonell, and Tom M. Mitchell, eds. *Machine Learning: An Artificial Intelligence Approach*. Springer Science & Business Media, 2013.

[22] Pascarella, Domenico, Salvatore Venticinque, and Rocco Aversa. "Agent-based design for UAV mission planning," P*2P, Parallel, Grid, Cloud and Internet Computing*, 2013.

[23] Pynadath, David V., and Michael P. Wellman. "Accounting for context in plan recognition, with application to traffic monitoring," *Uncertainty in Artificial Intelligence Conference*, 1995.

[24] Qin, Xinzhou, and Wenke Lee. "Attack plan recognition and prediction using causal networks," *IEEE Computer Security Applications Conference*, 2004.

[25] Schade, Ulrich, Joachim Biermann, and Miloslaw Frey. *Towards Automatic Threat ecognition*. No. RTO-MP-IST-055. FGAN-FKIE WACHTBERG (GERMANY), 2006.

[26] Schmidt, Charles F., N. S. Sridharan, and John L. Goodson. "The plan recognition problem: An intersection of psychology and artificial intelligence," *Artificial Intelligence* 11:1-2, 1978.

[27] Shima, Tal, and Steven J. Rasmussen, eds. *UAV Cooperative Deision and Control: Challenges and Practical Approaches*, 2009.

[28] Tulum, Kamil, Umut Durak, and S. K. Yder. "Situation aware UAV mission route planning," *IEEE Aerospace Conference*, 2009.

[29] Vandepeer, Charles. "Intelligence analysis and threat assessment: towards a more comprehensive model of threat," unpublished online paper, 2011.