Precision Identification and Persistent Sensing

Thomas M. Strat

Precision warfare requires precision targeting.
But identifying valid targets in imagery can be almost impossible.
For example, here are some images - which one is a valid target?

Not this one; this is a French Roland air defense missile system - used by both sides during the conflict in Serbia.

While Taliban forces used white Land Rovers extensively in Afghanistan, this one is being driven by an anti-Taliban warlord.

This is an Iraqi T-72 during the Gulf War, but it has been captured and is being operated by Coalition forces.

And some vehicles are targets no matter what they look like!

Identifying targets in modern asymmetric warfare is exceedingly difficult. Yet, that's exactly what we ask our warfighters to do repeatedly, often under conditions of extreme stress, and with severe time pressure.
Automated assistance is desperately needed.

Fortunately, US and Coalition forces are deploying a proliferation of sensors and platforms to carry them.
Some see the corresponding increase in sensor data as a burden - who is going to analyze all that imagery?
But it's not a burden.
The advent of persistent and ubiquitous sensing is an opportunity to automate image exploitation with reliability and accuracy that has never been possible before.
It fosters a new paradigm for sensor exploitation that can help us reduce the fog of war.

Today's image analysis algorithms do not take advantage of persistent sensing, so new research in image understanding is needed to create the foundations upon which we can build the next generation of sensor exploitation systems.
I will present some thoughts on how to do that, but we need your ideas on how to turn these promising generalities into reality.

Precision Identification is concerned with determining the validity of a target.
It's a concept that requires reasoning about form and function, about context and history, about salience and intent.
It requires an understanding of linguistic concepts, such as 'artillery' and 'reinforcement'.

At DARPA, we have invested in an array of programs to engineer solutions for particular sensor interpretation tasks: for example, AMSTE to track moving vehicles using Doppler radar, E3D to recognize military vehicles using multiple Laser radar images, and VIVID to confirm the identity of a target using airborne video observations

The field of Automatic Target Recognition has made steady advances toward its goal of identifying the make and model of target vehicles by inverting the image-formation process.
For example, the MSTAR program has achieved 80-90% recognition rates using realistic synthetic aperture radar imagery.
Yet Precision Identification requires more than determining the make and model of target vehicles.
Even if ATR worked perfectly, we would still need a way to determine the validity of a target.
Precision ID requires more than inverting the image-formation process.

So what can be done?

The answer lies in taking full advantage of Persistent Sensing--the tremendous capacity we have today to collect sensor data.
Until now, timely battlefield imagery has been a scarce commodity.
The small number of aerial and space-borne collection platforms was unable to satisfy the collection requirements of the many military units that relied on them.
Images were exploited in isolation, and even so, the exploitation community was hard-pressed to keep up with the flow.

But the advent of Predator, Global Hawk, and the vast array of tactical UAVs is changing all that.

Modern radars can surveil vast areas in SAR and MTI modes.
Global Hawk alone produces 10 times as much sensor data as the U-2.

And more video data is being collected today than by all the other modalities combined.
This proliferation of sensor data should not be seen as an additional burden on the exploitation community - instead, it is an unprecedented opportunity to change the way we go about identifying targets.

Imagine a future battlefield in which we employ near-ubiquitous, persistent sensing.
Droves of UAVs bristling with sensors.
We watch everything that moves.
And when it stops, we watch it until it moves again.
We keep track of every known piece of military hardware.
That's a lot of stuff, but it's not nearly as large as all the terrain we would have to search every time we needed to know where a ballistic missile launcher or an air-defense radar had moved to.
For example, Saddam Hussein's Iraq had roughly 4000 mobile targets of interest, such as tanks, TELs, and SAMs.
Even if we dedicated one sensor to each of the 4000 targets, that's far less than what would be needed to repeatedly search Iraq's 437,172 square kilometers of terrain.

Fully exploiting this flood of imagery is beyond the reach of current systems and staffing levels.
We will need automation to extract the relevant information from the sea of data.
But the good news is that this dense sampling of time and space actually makes automated image interpretation easier, not harder.

Instead of trying to make difficult determinations from isolated pieces of data, we can use more complete sets of data to ease the exploitation tasks.

Here are some examples:

* Tracking versus searching.
 Let's track vehicles rather than repeatedly search for them - continuous tracking is easier than searching for vehicles over broad areas, especially in the face of camouflage and deception.

* Maintaining ID.
 When we are clever enough or lucky enough to have recognized the class of a vehicle, let's maintain that ID over time - we only need to recognize each target once! We might even consider doing the recognition manually, if we can automate the tracking process, because we won't have to recognize new targets very often.

* Detecting change.
 Rather than attempt change detection over long spans of time, let's continuously monitor for change over closely spaced intervals.
 Watching a delivery to a chemical munition plant is a lot easier than trying to infer that it happened after the fact.

* Learning.
 Large collections of repeated observations open the door to algorithms that learn.
 Today we build an exploitation system in the factory, deploy it, and use it until we ship a new version from the factory.
 Tomorrow, our exploitation systems will add target types on-the-fly, they will adapt through experience, and their performance will improve over time.
 Our automated systems should never make the same mistake twice.

The typical sensor exploitation system that ingests data, analyzes it, and provides a response is handicapped from the start.
 It has no context, no notion of what it has seen before, no ability to improve its performance over time, and no sense of awareness.
 Every processing task is independent of the others -- the same inputs will always produce the same outputs, and it is doomed to repeating the same mistakes endlessly.

The new opportunities arise because of persistent, ubiquitous sensing.
 These are closely related concepts: Ubiquity can be thought of as spatial persistence-and persistence is temporal ubiquity.

Although persistent, ubiquitous sensing exists only on a small scale today, the day when it becomes practical for real-world problems is not far off.
 For example, a single Global Hawk mission can collect more than 4,000 large-format radar images.
 A single video camera collects more than 2.5 million image-frames per day, and every ground combat vehicle soon will be carrying multiple video cameras.
 Many automated capabilities for interpreting sensor data already exist, such as GMTI tracking, foliage penetrating radar, and infrared target detection, but we

don't have much that takes advantage of the unique opportunities that leverage persistence and ubiquity.
Worse, we don't even have solutions at a laboratory level to such problems as reasoning from context, monitoring change, and learning from experience.

What we need is a concerted effort, at the level of basic research, to create the algorithms and the software architectures -- the foundation upon which we can conceive and design and build exploitation systems that perform Precision Identification at human levels of competence -- or better.

Persistent, ubiquitous sensing changes the problem.
Now is the time to explore approaches to create perceptually aware intelligent systems for Precision ID.
Rather than search broad areas for camouflaged targets, we can monitor those targets for changes.
Rather than attempt to detect and recognize targets in unfamiliar terrain, we can track them so we know where they are at all times.
Rather than deploy exploitation systems that are difficult to modify, we can develop systems that improve their performance through experience.
If we do this, we can turn the ever-increasing ability to sense the battlefield into the key that solves the precision identification problem.

Together with our partners in government and industry, we are working to make automated Precision Identification a reality.
I invite you to help us lift the fog of war and replace it with a more lethal environment for our adversaries.