Automatic target recognition with passive bistatic radars, with applications to the detection of anomalies in the civilian air traffic

Jonathan Pisane (1,2), Sylvain Azarian (2), Marc Lesturgie (2), Jacques G. Verly (1)
1: University of Liège, Dept. of Electrical Engineering and Computer Science, INTELSIG Lab., Liège, Belgium
jpisane@ulg.ac.be
2: SUPELEC /E3S - EA4454/SONDRA, Gif-sur-Yvette, France

Abstract

We consider the automatic recognition of non-cooperative airplanes observed with a passive bistatic radar, thus using illuminators of opportunity ,and in particular navigation aids (VOR's). The testbed we deployed around the Orly airport, France, allowed us to compute the bistatic radar cross-section of real airplanes. We then performed the recognition of these airplanes using their radar cross-sections and subspace methods, achieving a percentage of correct recognition of about 83%.

Today, air traffic controllers can detect airplanes by using primary surveillance radars, and identify them by using secondary surveillance radars. However, a non-cooperative airplane – i.e. an airplane without an ADS-B transponder (or equivalent) or with one that is defective and/or not responding to radio calls – flying in, and perhaps intruding into, controlled airspace, cannot be identified by secondary surveillance radars. It is thus important to develop systems that work independently from the air traffic control system, are easy to deploy, and are inexpensive (which excludes microwave imaging radars). A potential solution is to use passive radar systems using illuminators of opportunity, such as radio and TV stations, and navigation aids. We address this problem and demonstrate the potential of the approach.

In radar applications, targets are mainly characterized by their radar cross-section. We thus perform the recognition of airplanes by using their bistatic radar cross-section (RCS), denoted by σ.

We designed, deployed, and successfully used an experimental passive radar testbed in the vicinity of the Orly Airport, south of Paris, France. The initial plan called for using two VOR (VHF Omni Range) navigation aids as transmitters (Tx's) of opportunity, and a software-defined radio (SDR) of our own design as receiver (Rx). The VOR's were those of Rambouillet (with call-sign RMB, and operating at 114.7 MHz) and Epernon (EPR, 115.6 MHz). However the signal-to-noise ratio of the signal received from EPR proved to be too small to be usable. Our effective testbed thus comprised a single (Tx, Rx) pair. Each such pair defines a bistatic (BS) radar. We also set up an ADS-B receiver to obtain the identity (ICAO call-sign) and position of each airplane in view, each time it is interrogated by a secondary surveillance radar. By tracking the airplane's position over time, we obtained its measured trajectory, which allowed us to estimate the heading of the airplane at any time, an important piece of information for recognition. We recorded the signals received by our SDR Rx almost continuously for ten days, resulting in measurement data corresponding to 1,329 airplanes of 41 different types, 1,329 trajectories, and 54,154 sampling points. The first third of this collected data constituted the learning set, which we used to build the recognition model. The two other thirds made up the test set, which we used to test our recognition model and to quantify its performance.

For each sampling point, we separated, by using a Doppler filter in the frequency domain, the direct signal coming from the VOR and the scattered signal coming from the airplane, which allowed us to deduce the BS RCS σ of the airplane. It is well-known that this BS RCS varies with the BS angle β and with the aspect angle α (say, measured from the bisector of β). The knowledge of the position of the airplane and of its trajectory, and of the positions of the Tx and Rx, allowed us to compute the values of α and β associated with each sample point. Each such point is thus characterized by a specific triplet (α, β, σ). One can thus map the physical trajectory of each airplane into a trajectory in an (α, β) plane, with a value of σ associated to each sampling point along this parameter-plane trajectory.

A significant feature of our recognition system is that we partition the (α, β) plane into regions (which are generally rectangular, but can be of arbitrary shape), and that we build a specific recognizer for each region. It is useful to imagine that there is a distinct (α, β) plane for each pre-defined class of airplanes (assuming supervised learning), each being partitioned in the same way. To build the recognizer for each specific class and region, we built feature vectors (FV's) of RCS's from all trajectories in each such region. For each class of airplanes and for each region, we computed, via a singular value decomposition (SVD), the best corresponding subspace. Our recognition model thus consists of a list of subspaces, each subspace characterizing a class of airplanes for one region. During operational use, we would also build one or more FV's from operationally obtained trajectories, and we would project them in the subspaces of all classes. The best projection metric would determine the class of the observed airplane.

In our recognition experiments, we considered the three broad classes of large, medium, and small airplanes, most from the two major airplane manufacturers. We achieved an overall correct recognition rate of 83%, which demonstrates the validity of the approach for the detection of non-cooperative airplanes intruding into controlled airspace. Since the correct recognition rate varies according to the region of the (α, β) plane, higher recognition rates can be achieved by optimizing the position(s) of the receiver(s).