RECONSTRUCTING SIMULATOR CONTROL INPUTS: A MACHINE LEARNING APPROACH

Eamon Caddigan, DSoft Technology, Engineering & Analysis, Inc.

Flight simulators are valuable tools for human factors research. However, some simulation platforms fail to record all of the information relevant to the researcher. While the data produced by most simulators includes details about the position and state of the simulated aircraft, some platforms do not record pilots' control input. Missing control input data make it difficult to evaluate response times, a key behavioral measure in human factors research. Here we describe a technique that uses machine learning to reconstruct aircraft maneuvers using aircraft control surface information, which is typically available in simulator output files. This allows researchers to more accurately estimate the moment at which a pilot initiated a maneuver.

Human factors researchers often need to determine how different displays and procedures affect participants' response times (Bradley, 1971; Hartzell, et al., 1983; Rorie & Ferne, 2014; Santiago & Mueller, 2015). In aviation contexts, this is defined as the time between the onset of a stimulus (e.g., a traffic warning), and the moment that a pilot initiates a maneuver in response to the stimulus. When a simulation platform does not record the state of control inputs (which is the case for some high-fidelity training simulators), it is difficult to identify the *maneuver onset*, or moment that the pilot initiated a maneuver. The technique described in this document used the output from two Unmanned Aircraft System (UAS) simulator platforms that controlled different aircraft: The Common Open-mission Management Command and Control (ICOMC2) station controlling a Boeing Insitu RQ-21 Blackjack and a General Atomics Predator ground control station (GCS) controlling a MQ-9 Reaper (Williams, Caddigan, & Zingale, 2017).

When control input data are unavailable, maneuvers can be inferred using aircraft position information. For example, a researcher can estimate instantaneous change in altitude by calculating the ratio of the difference between adjacent altitude samples to the sample period. Sometimes thresholds can be specified *a priori*, e.g., a change in altitude of -500 ft/s could be labeled as *descent*. However, natural variance in position makes this approach unreliable; depending on how thresholds are defined and applied, time points will either alternate between maneuvers when the rate of change is close to the threshold, or the procedure will leave many time points unlabeled. Additionally, there will necessarily be a lag between the moment that the maneuver is initiated and the time at which the aircraft's position change is detectable. This document will discuss methods to best manage these challenges.

The method described below takes advantage of two facts: 1) that there is a consistent mapping between an aircraft's maneuvers and its control surfaces, and, 2) that these control surfaces change more quickly than the aircraft's position in response to pilot input. For example, in a conventional fixed wing aircraft, a pilot initiates a turn by deflecting the yoke. This will result in a nearly instantaneous change in aileron position which will induce a roll that will lead to a banking turn (Federal Aviation Administration, 2007). Response time would ideally be evaluated using the moment that the yoke was deflected. When this information is unavailable in a simulator, changes in simulated aileron position are a timelier indicator of the onset of a maneuver than changes in heading.

The method described in this paper applies thresholds to changes in simulated aircraft position to identify time points that can be associated with specific maneuvers with a high level of confidence. Control surface, engine power, and maneuver information from these time points were used to train classifiers that learned the mapping of control surfaces and engine power to maneuvers. These classifiers were then used to estimate the maneuver for each of the remaining time points, allowing for earlier identification of maneuver onsets.

PRACTICE INNOVATION

Changes in Aircraft Position

Simulator data were collected for a study on minimum information requirements for detect and avoid displays for unmanned aircraft pilots (Williams, Caddigan, & Zingale, 2017). Sixteen pilots used the Predator GCS to control a simulated MQ-9 Reaper, which has a wingspan of 20 m and a cruise speed of 169 kn. An additional sixteen pilots used ICOMC2 to control an RQ-21 Blackjack, which has a wingspan of 4.9 m and a cruise speed of 55 kn. Data were gathered in 1156 runs (613 on the Predator GCS and 543 on ICOMC2), testing different display configurations and encounter geometries, and had an average duration of 221 s (± 39 s).

One approach to recover maneuver data from a simulation is to use aircraft position information: i.e., altitude, heading, and airspeed. We estimated rates of change in heading, altitude (MSL), and true airspeed by multiplying the difference between adjacent samples by the sampling rate (5 Hz for both simulators). Heading information was unavailable for the ICOMC2 but was estimated by calculating the bearing between successive samples of latitude/longitude coordinates.

The histograms of these difference values were inspected to identify peaks which we associated with periods of change and non-change. Thresholds were chosen to exclude ambiguous time points. To mitigate noise, we used lagged differences; short lag values produced noisy estimates, but long lag values removed important points. Therefore, lags and thresholds were chosen for each maneuver using visual

 DOI 10.1177/1541931218621012

inspection of the difference histograms for each aircraft position value (see Figure 1).

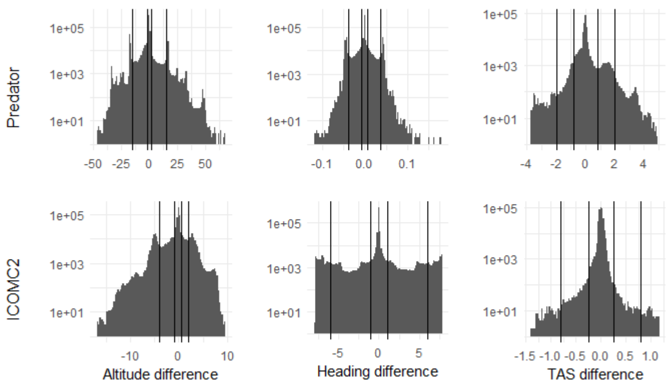

Figure 1: Distribution of lagged position changes for the Predator (top row) and ICOMC2 (bottom row). The vertical lines indicate thresholds, which categorize position samples as changing (positive or negative), unchanging, or ambiguous.

Time points were labeled with a specific maneuver using thresholds applied to changes in heading, altitude, and airspeed. The labels we created included all combinations of heading change (left turn, right turn, and straight flight), and altitude change (climb, descend, and level flight). In addition to these nine maneuvers, airspeed changes (i.e., acceleration and deceleration) were considered during straight and level flight, for a total of eleven possible maneuver labels. The maneuvers were not equally prevalent; straight and level flight was the most common label (4.5 epochs per run), while airspeed maneuvers were rare (1.2 decelerating and 1.1 accelerating epochs per run).

Classifying Maneuvers Using Control Surfaces

We used random forest classifiers to learn the mapping between simulated control surface and engine power data and maneuver labels (Liaw & Wiener, 2002). The classifiers were trained using features corresponding to aircraft control surfaces and engine power (10 features for ICOMC2, 80 features for Predator). Random forest classifiers use an ensemble of simple decision trees; each tree is built using randomly selected subsets of the available features and training data to avoid overfitting (Ho, 1995). Random forests are well suited to problems with many classes, since each class can correspond to one or more leaves. The number of trees used, B, was treated as a free parameter. We estimated the optimal value of B for each simulator using cross validation, testing the values {50, 100, 150, …, 300} to ensure that asymptotic cross validation error would be observed.

For each simulator, the runs were randomly split into a training set and a validation set, with 20% of the runs set aside for validation (109 runs for the ICOMC2 and 123 runs for the Predator). The remaining runs were randomly assigned to 8 folds (approximately 54 runs per fold for the ICOMC2, and 61 runs per fold for the Predator). The number of folds was chosen to allow for efficient parallel computation of cross validation accuracy on a standard computer workstation. For each value of B that was considered, a classifier was trained on the training data from all but one fold, and classifier performance was evaluated by calculating classifier accuracy on the remaining fold. This procedure was repeated for each fold, and accuracies were averaged to estimate classifier performance for that parameter value. Both simulators performed best with $B = 150$, with a mean cross validation error of 3% for the ICOMC2 and 4% for the Predator. Even with this level of performance, cross validation error was higher than training accuracy (1% for both simulators), which suggests that additional training data would lead to improved performance.

With the value of B selected, final classifiers were created for each simulator by training on all of the runs in the respective training sets. These classifiers provided maneuver labels for every sample that was left unlabeled after applying the thresholds to changes in aircraft position, which represented 25% of the samples. We used the aircraft position data to identify the sequence of maneuvers in the simulator data, but extended the duration of each maneuver to include all samples adjacent to the maneuver for which the classifier provided the same maneuver label. After extending the maneuver durations in this manner, the remaining unlabeled samples accounted for only 8% of the total.

FINDINGS

Classifier performance on the validation data remained good, with an error of 3% for the ICOMC2 and 4% for the Predator. This suggests that the classification accuracy observed during parameter selection did not rely on overfitting.

The primary goal of this technique is to identify the onset of aircraft maneuvers earlier than is possible by using changes in aircraft position alone. To determine this technique's effectiveness, we measured the average difference in onset times across the maneuvers in each run of a simulator. Note that our method for labeling maneuvers ensured that onset times recovered using classifiers trained on control surface and engine power data would never be later than those recovered using aircraft position alone. The results of this comparison are shown in Figure 2; the mean time difference across ICOMC2 runs was 1.53 s ($t(542) = 42.98$, $p < 0.0001$ [1.46, 1.60]), and across Predator runs was 2.38 s ($t(612) = 49.24$, $p < 0.0001$ [2.29, 2.48]).

DISCUSSION

This procedure used a machine learning method that had been trained on aircraft control surface and engine power data to reliably recover the onset of maneuvers in simulators. Maneuver onsets are identified over 1 s earlier compared to maneuver identification performed by applying thresholds to changes in aircraft position. This is a significant difference in the context of response times, where "large" differences might be on the order of hundreds of milliseconds (Donders, 1869).

A limitation to the current approach is that ground truth information for pilot intentions was unavailable; maneuvers were inferred from rates of change of simulated aircraft position. Future work should investigate whether classifier performance is improved when ground truth maneuver information is available for training. For example, researchers could generate a set of training data by asking pilots to execute a specific sequence of maneuvers at predetermined times. This approach would also provide ground truth maneuver onset times against which classification-based onset times could be compared; the current approach assumes that the true time of maneuver onset is unavailable.

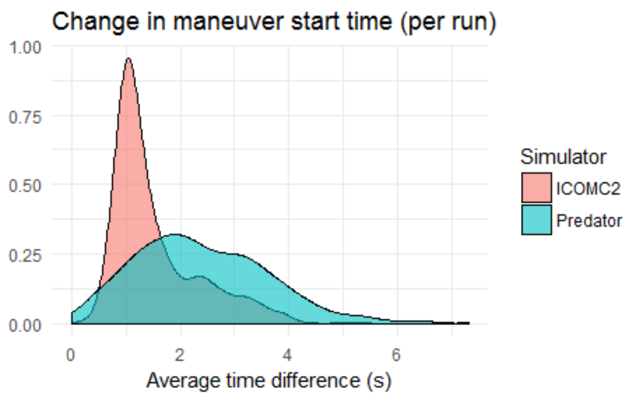

Figure 2: The distribution of mean start time differences across runs for the two simulator platforms. Maneuver start times were identified earlier using classifier results.

PRACTITIONER TAKE-AWAYS

- Human factors researchers use the onset of a maneuver to measure response time.
- High-fidelity training simulators may not record control inputs, which provide the best measure of pilots' maneuver decisions.
- Changes to an aircraft's control surfaces are better proxy for the onset of a maneuver than changes to aircraft position (e.g., heading, altitude).
- Machine learning methods can recover the mapping between control surfaces and aircraft maneuvers.
- This technique identifies the onset of maneuvers earlier than applying thresholds to changes in aircraft position alone.

REFERENCES

Bradley, J. V. (1971). Tactual coding of cylindrical knobs. In W. C. Howell & I. L. Goldstein (Eds.), *Engineering psychology: Current perspectives in research* (pp. 276-287). New York: Appleton-Century-Crofts.

Donders, F. C. (1869, trans. 1969). On the speed of mental processes (trans. W. G. Koster). *Acta Psychologica, 30*, 412-431.

Federal Aviation Administration (2007). *Pilot's Encyclopedia of Aeronautical Knowledge*. New York: Skyhorse Publishing.

Hartzell, E. J., Dunbar, S., Beveridge, R., & Cortilla, R. (1983). Helicopter pilot response latency as a function of the spatial arrangement of instruments and controls. *Proceedings of the 18th annual Conference on Manual Control* (pp. 345-364). Dayton, OH: Wright Patterson AFB.

Ho, T. K. (1995). Random decision forests. *Proceedings of the 3rd International Conference on Document Analysis and Recognition*, (pp. 278-282). Montreal, QC

Liaw, A., & Wiener, M. (2002). Classification and Regression by randomForest. *R News 2*(3), 18-22.

Rorie, R.C. & Fern, L. (2014). UAS measured response: The effect of GCS control mode interfaces on pilot ability to comply with ATC clearances. *Proceedings of the 58th Human Factors and Ergonomics Society Annual Meeting*, pp. 64-68.

Santiago, C., & Mueller, E. R. (2015). Pilot evaluation of a UAS detect-and-avoid system's effectiveness in remaining well clear. *Proceedings of the 11th USA/Europe Air Traffic Management Research and Development Seminar*, Lisbon, Portugal.

Williams, K. W., Caddigan, E., & Zingale, C. (2017). *An investigation of minimum information requirements for an unmanned aircraft system detect and avoid traffic display* (DOT/FAA/AM-17/14). Washington, DC: Office of Aerospace Medicine.