

Learning to Fly: Design and Construction of an Autonomous Airplane

Presented at the 1999 Florida Conference on Recent Advances in Robotics, April, 1999.
[This version edited June, 2000]

Stephen B. Stancliff
stancliff@mil.ufl.edu

Jennifer L. Laine
jenny@mil.ufl.edu

Michael C. Nechyba
nechyba@mil.ufl.edu

Machine Intelligence Laboratory
Electrical and Computer Engineering
University of Florida
Gainesville, FL 32611-6200

Abstract

Much work has been done in recent years to abstract computational models of human control strategy that are capable of accurately emulating dynamic human control behaviors. Land-based autonomous vehicles, both in simulation and on real roads, have made successful use of this modeling formalism. Little work has been done, however, in attempting such skill transfer from humans to aerial robotic vehicles. Although control of an aerial vehicle is quite different from that of ground vehicles, we contend that human pilots can potentially serve as excellent guides in the development of intelligent autonomous aerial vehicles. As a first step in modeling human control strategies in aerial vehicles, we are developing a robotic airplane as an experimental platform for studying human-to-machine skill transfer in aerial vehicles. This paper describes the design of this airplane, the status of the project, and future planned experiments.

1. Introduction

1.1 General description of platform

The platform is constructed around a radio-control airplane kit. The kit used was chosen for its stability, its large wing area (to allow for lifting the large electronics payload), and the ready availability of information on powering it electrically. The airplane is propelled by a brushless electric motor, which is powered from nickel cadmium (NiCd) batteries. Manual control is provided through standard radio control (R/C) servo motors and a six-channel digital proportional R/C radio system. Autonomous control will be provided by a 386-class processor in conjunction with a Motorola 68HC11 processor. The HC11 will interface with those sensors that have digital and analog output, and with the radio/servo system, while the 386 will execute the high-level code and will

interface with those sensors which have RS-232 output.

1.2 Motivation

Over the past two decades, rapid advances in computer performance have not been matched by similar advances in the development of intelligent robots. Humans are much better at performing complex dynamic skills than at describing those skills in an algorithmic, machine-codeable way. This has limited our ability to develop intelligence in robots and other machines. This inability has limited not only the capabilities of individual machines, but also the extent to which humans and robots can safely interact and work cooperatively. There exists a profound need to abstract human skills into computational models which are capable of realistic emulation of dynamic human behaviors.

A number of researchers have worked in recent years to abstract models of human skill directly from observed human input-output data¹. Autonomous control and navigation of ground vehicles is one area of robotics research which has benefited from learning through observation of humans. Pomerleau [3,4], for example, has implemented in the ALVINN system real-time road-following using data collected from a human driver. A static feedforward neural network with a single hidden layer learned to map from coarsely digitized images of the road ahead to a desired steering angle. The ALVINN system has been demonstrated successfully at speeds up to 70 mph. Pentland and Liu [5] have applied hidden Markov models to inferring a driver's high-level intentions, such as turning and stopping. Finally, [1,2,6] address the autonomous control of a dynamic car--including steering and acceleration--through

¹ Nechyba [1,2] provides an overview of the literature.

observation and modeling of human driving using a driving simulator.

Surprisingly, little work has been done in using observation of human pilots to create intelligent autonomous aerial vehicles. Such vehicles could have application in a number of areas. Many of the activities that currently involve remotely piloted vehicles (RPVs) would benefit in some way from automation. In many applications, such as surveying, reconnaissance, and target acquisition, it may be possible to automate the entire mission. For other applications, where more sophisticated control is required, it may be that a human pilot is still required, but that the less complex parts of the mission can be automated. Adding intelligence to RPVs could reduce the amount of skill required of the human pilots, and could also allow one pilot to control multiple vehicles.

Although the control challenges in flying are quite different from those in driving, the basic paradigm of learning from humans is equally applicable. We therefore propose to extend to aerial vehicles some of the methods previously used for learning in ground vehicles. As a first step in this process, we are developing a robotic airplane as an experimental platform.

1.3 Project outline

We have divided this project into three stages according to the level of autonomous behavior to be achieved: (1) The airplane will be able to fly straight and level, maintain a given heading, and make a turn to a new heading. (2) The airplane will be able to land and take off. (3) The airplane will be able to navigate to given map coordinates.

At the outset of the project, we set nine months and \$1000 as the time and budgetary constraints for stage one. In this paper, we describe what design choices these constraints forced, the resulting hardware and software configuration of the airplane, and the future goals for the project.

2. Platform Design

2.1 Initial configuration

The first design decision to be made was whether to use a fixed-wing or rotary-wing aircraft. Each type of platform has advantages and disadvantages for the missions envisioned. In general, a fixed-wing aircraft will have a greater payload-range, be more stable, and have fewer and simpler control mecha-

nisms. Consequently, a fixed-wing aircraft requires a much less expensive radio system and less skill from the human operator. A rotary-wing aircraft, on the other hand, has the ability to hover (useful for photography and surveillance) and can operate from smaller landing areas. For the purposes of this research, the factors of cost, simplicity, stability, and payload are the most important, and they led us to choose a fixed wing platform.

The next design decision to be made was the type of powerplant to be used—an electric motor or an internal combustion engine. A few years ago this would not have required much thought, for it is only recently that motors have attained performance comparable with engines. The advantages of the electric motor include low vibration, lower long-term cost, absence of caustic fluids (fuel and exhaust), and low noise. The internal combustion engine offers the advantages of a well-known technology, greater payload-range, and lower initial cost.

The decision in this case was not clear-cut, but the desire to maintain simplicity and reliability led us to choose an electric powerplant, in order to avoid the design effort and weight required to shield the electronics from vibration and fluids. In addition, with the electric motor, the performance of the aircraft will improve in the future as battery technology improves.

Once the decision was made to use an electric motor, a choice had to be made between brushed and brushless motors. The brushless motors available for this application are much more efficient than the brushed motors available, they produce less electromagnetic interference (EMI), and they do not require the periodic rebuilding that brushed motors do, resulting in a lower long-term cost. Brushed motors, on the other hand, are both cheaper themselves, and also allow the use of simpler motor-speed controllers, resulting in a much lower startup cost. Again, avoiding interference with the computers was deemed the most important factor, and we chose to use a brushless motor.

A final initial configuration decision was the selection of battery technology. NiCd, nickel-metal-hydride (NiMH), and lithium-ion (Li^+) technologies were considered. While the newer technologies provide a much greater energy density than NiCd, they are not capable of providing the high currents

(10-15 C²) required by the motor. Therefore NiCd was the only option for powering the motor. The more advanced batteries would significantly reduce the weight of the battery packs required for the electronics, and we are seeking donations of NiMH or Li⁺ batteries for this purpose.

At this point, we had a good estimate of the mission requirements, including the weights of the payload components. We chose to design for 3.5 lb of payload, which gave us plenty of room for error in our weight and power estimates, and also for future expansion of the sensor and computing capabilities. We also concluded that we required a time aloft of at least five minutes, to allow for sufficient data collection. Finally, we needed a stable, forgiving airframe.

Given these parameters, we sought the advice of experienced electric-model pilots about possible airframes to use. On the basis of their advice, we selected the Sig *Kadet Senior*. This airframe is well suited to the task with its large wing area (1150 in²), roomy fuselage, and highly stable flight characteristics. The one significant drawback to this model is that, in comparison with other kits, it is somewhat challenging and time-consuming to build, with many parts needing to be cut and shaped, rather than being ready to assemble.

2.2 Drivetrain design

Once the airframe was chosen, the drivetrain could be designed. The primary goals of this design process were: (1) to provide enough power to maintain flight; (2) to provide enough energy to fly for a given period of time, and (3) to minimize energy losses.

Several more-or-less scientific methods exist in the electric R/C community for sizing motors and batteries. Preliminary calculations indicated that the various methods produce similar results. The method we used was a combination of the method outlined by Orme [7] and the instructions included with the software MotoCalc [8], which is a design tool that simplifies the repetitive calculations involved.

As the first step of the design process, we estimated the number of sub-C size NiCd cells required. Orme [7] has a rule-of-thumb that a large, slow, trainer-type aircraft needs one cell per 50 in² of wing area.

This rule-of-thumb results in a wing loading of about 20 oz/ft². Applying this rule to our airframe gave us a requirement of about 24 cells (28.8 V).

Using this approximate cell count, we were able to select several motors for evaluation, based on their voltage ratings. Preference was given to motors which had been flown by others in the *Kadet Senior*.

Design from this point was performed with the MotoCalc program. MotoCalc includes a database of common motors, motor controllers, batteries, and airframes. It uses this information to calculate power usage, efficiencies, and flight envelope. All of the motors and controllers we considered were found in the MotoCalc database. We used 24 sub-C NiCd batteries for the preliminary design, since specific batteries had not yet been chosen. The current design uses five 6 V camcorder batteries, which is equivalent in voltage to 25 sub-C cells. Our airframe was not in the database, but we were able to estimate the unknown aerodynamic data from similar airframes in the database. These estimates were later refined by measurement of related aircraft dimensions. We incorporated the weight of the payload by including it in the empty weight of the airframe.




Figure 1 - Motor Efficiency

The first step in design with MotoCalc was to determine the most efficient current for a given motor and number of cells. The program does this by calculating the losses in the motor and wiring at various currents. For the motor we ended up choosing (Aveox 1406/4YSE), the peak motor efficiency with 25 cells occurs at about 18 A (Figure 1) and is about 89%. The efficiency is greater than 88% for currents from 14 A to 25 A (a characteristic typical of brushless motors), giving us plenty of

² C is the one-hour current capacity of the cell.

leeway in the selection of propeller, cells, and gearbox.

At the desired current of 18 A, the motor will turn at about 33,000 rpm. The propeller should usually spin at 6,000-12,000 rpm at full throttle, so a gearbox was needed. A gearbox with a ratio of 3.69:1 was selected. The calculations were repeated with this gearbox ratio, and for a range of propeller sizes (Figure 2). The two propeller sizes which came closest to the desired current, while being able to maintain flight, were 11x7³ and 12x6.




Figure 2 - Propeller Selection

Performance predictions for the two selected propeller sizes are given in Table 1. By comparing initial rates of climb and flight durations, it is apparent that the 11x7 propeller will provide greater power for a shorter period of time.

Table 1 - Propeller performance

Propeller size	11 x 7	12 x 6
Motor current (A)	18.6	20.6
Motor efficiency (%)	89.0	88.9
Prop speed (rpm)	8,853	8,489
Stall speed (mph)	21	21
Max. level speed (mph)	43	41
Rate of climb (fpm)	359	342
Duration (min)	14:26	16:41

The flight durations given in Table 1 are upper limits, and cannot be achieved in reality. They are based on level flight at a partial-throttle setting (in this case about 89%), and therefore do not account for the energy required for takeoff and climb to altitude. The worst-case duration, based on constant 100% throttle, was also calculated, and is about four

minutes for each propeller. It is therefore reasonable to assume that we will be able to achieve flight durations of 5-10 minutes.

As a confirmation of the MotoCalc predictions, hand calculations were performed for a few of the propeller-motor-gearbox combinations, and the results were in close agreement with the MotoCalc results. In addition, the design has been revisited as we have obtained more accurate component data, and the design predictions have remained very stable as the design has solidified, giving us confidence that the platform will provide the performance that we require.

3. Intelligent Systems Design

3.1 Sensors

One of the benefits of using a controller which learns from a human is a reduction in the engineering effort required to characterize the system being controlled. The system designer does not require a highly detailed understanding of the dynamics of the system being controlled or of the sensors and actuators. It is not necessary to have explicit relationships between inputs and outputs, only a general idea of what inputs contribute to the outputs. As long as the inputs are sufficient, the neural network can determine which are necessary. From this standpoint, it is desirable to have as many sensor inputs as possible. On the other hand, with an aerial platform it is necessary to limit the number and type of sensors in order to maintain low weight.

For stage one, we are implementing the following sensor suite: (1) a compass, (2) a two-axis tilt sensor, (3) an altimeter, and (4) an airspeed sensor. We expect that these sensors will be sufficient to allow the system to maintain level flight on a given heading, and to make a turn to a new heading, which satisfies the stage one goals.

For heading detection, we plan to use the Honeywell HMR3000 Electronic Compass Module. This module provides an RS-232 output of heading, pitch, and roll. The heading is electronically compensated for pitch and roll angles of up to 45°.

We have constructed a tilt sensor using the Analog Devices ADXL202 dual-axis ±2g accelerometer. The tilt sensor provides a duty-cycle-modulated output signal where the duty cycle is proportional to tilt angle. Details of the construction of this tilt sensor can be found in the ADXL202 data sheet [9].

³ diameter (in) x pitch (in)

We have designed an altimeter using the Motorola ASB210 pressure sensor module. This sensor detects 0-10 in-H₂O differential pressure and provides an analog voltage output. One port will be sealed before takeoff, and the differential pressure measurement will then be proportional to altitude. We expect altitude resolution of about 7 ft using this method with the HC11. We may be able to obtain better resolution by using a separate analog to digital converter.

Obtaining or constructing a small inexpensive airspeed sensor has proven to be a great challenge. We hope to use another Motorola ASB210 differential pressure sensor with a pitot-static tube, but have not been able to find an appropriate tube. Another possibility is to use a turbine-type windspeed sensor removed from a handheld unit.

3.2 Onboard computing and control

We plan to have two onboard computers, a microcontroller to handle sensor input and control output, and a more sophisticated primary computer to implement the intelligent control .

The primary design criteria for the higher-level computer are: (1) small size and weight, (2) low power consumption, (3) a large amount of RAM (to store training data), and (4) multiple serial ports (to connect to sensors, the microcontroller, and a host terminal). In addition, a hard drive interface is desirable, since the platform will eventually employ vision, requiring significant data storage.

For the primary computer, we have chosen a 386-class processor on a PC/104 format board. The PC/104 format provides all of the features usually found on a desktop computer on a 3.6" x 3.8" board, often weighing less than 4 oz. It also provides a standard bus for easy connection of expansion boards. We have chosen to use a 386 processor rather than a 486 or Pentium since it uses less power, saving battery weight, and is also much cheaper. A 386 is sufficient for the current needs, although a more powerful computer may be necessary in the future for image processing.

The secondary computer will be a Motorola 68HC11 microcontroller on a Mekatronix MSCC11 single-chip board. This board was designed in the Machine Intelligence Laboratory (MIL), and many of our robots use it as their primary controller. The code that will run on the HC11 is readily adapted from existing code. All sensor and input data except for

the compass will be captured by this board. This board will coordinate with the main processor through a serial link. A Mekatronix MB2325 serial board will provide an RS-232 interface to the MSCC11.

In an aerial vehicle, safety of the vehicle and of nearby property and people is an important concern. It is necessary to maintain positive control at all times. Knowing that we are likely to have flaws in the computer controller at some point, we have designed a control-multiplexing system which allows a human safety pilot to take over control at any time. This system uses one channel on the radio to toggle between computer control and human control. The multiplexing system is independent of the computers and computer battery so that in the event of a computer or power failure, the human pilot may still take control of the aircraft.

3.3 Software design

In previous work [1,2,6], we have successfully modeled human control strategies (HCS) in the driving domain. We have, for example, successfully learned human driving behaviors in a driving simulator. We expect that much of that work will transfer readily to the flying domain.

Broadly speaking, all control can be classified as either continuous or discontinuous. In the driving domain, for example, steering varies continuously with sensor inputs, while acceleration varies discontinuously with sensor inputs, because of the switching between the gas and brake pedals. We have previously found that continuous human control can be modeled well through cascade neural networks, which are powerful nonlinear function approximators which offer several advantages over more traditional neural network architectures: (1) the network architecture is not fixed prior to learning, but rather adapts as a function of learning [10]; (2) hidden units in the neural network can assume variable activation functions [11]; and (3) the weights in the neural network are trained through the fast-converging node-decoupled extended Kalman filter [11]. The flexibility of these cascade networks is ideal for HCS modeling, since few *a priori* assumptions are made about the underlying structure of the human controller.

Discontinuous human control can only be poorly approximated through a continuous learning formalism [2]. Therefore, we have developed a statistical, Markovian discontinuous learning

architecture [6], that has successfully abstracted the switching between the gas and brake pedals in human driving. The learning paradigms for both continuous and discontinuous control may be required for fully developing an intelligent autonomous airplane.

In general, we plan to proceed with the human modeling experiments as follows: (1) The plane will first be flown under human control, with the computer collecting and recording sensor and control data. (2) HCS models similar to those developed previously in the driving domain will then be trained offline using the recorded human control data. (3) The final HCS model(s) will then be downloaded to the onboard computers and the system will be tested in flight.

4. Conclusion and future work

This paper describes the first stage of development of an autonomous airplane. At the time of this writing, we have completed much of the hardware design and construction, with the airspeed indicator being the only significant remaining design issue. Final assembly of the platform awaits arrival of all of the parts, since some modification of the airframe may be necessary to provide access to all of the components while maintaining proper balance. The failsafe multiplexing system is near completion, and the development of the HC11 software has begun.

Acknowledgments

We would like to acknowledge the support of the following people and companies, without whom this project would not have been possible: (1) the members of the Machine Intelligence Laboratory; (2) Motorola, a major sponsor of the MIL; (3) Honeywell, which provided the HMR3000 compass module; (4) Eveready, which provides batteries to the MIL; and (5) Analog Devices, which provided the ADXL202 accelerometer chips.

References

- [1] M.C. Nechyba and Y. Xu, "Human Control Strategy: Abstraction, Verification and Replication," *IEEE Control Systems Magazine*, vol. 17, no. 5, pp. 48-61, 1997.
- [2] M.C. Nechyba, *Learning and Validation of Human Control Strategies*, Ph.D. thesis, Carnegie Mellon University, 1998.

[3] D.A. Pomerleau, "Neural Network Perception for Mobile Robot Guidance," Ph.D. thesis, Carnegie Mellon University, 1992.

[4] D.A. Pomerleau, "Reliability Estimation for Neural Network Based Autonomous Driving," *Robotics and Autonomous Systems*, vol. 12, no. 3-4, pp. 113-9, 1994.

[5] A. Pentland and A. Liu, "Toward Augmented Control Systems," *Proc. Intelligent Vehicles*, vol. 1, pp. 350-5, 1995.

[6] M.C. Nechyba and Y. Xu, "On Discontinuous Human Control Strategies," *Proc. IEEE Int. Conference on Robotics and Automation*, vol. 3, pp. 2237-43, 1998.

[7] M. Orme, "How to Choose the Right Motor for Your Airplane, Parts 1 and 2," Aveox Inc., Westlake Village, CA. (<http://www.aveox.com/selection.html>)

[8] "MotoCalc Manual", Capable Computing, Inc., 1998. (<http://www.motocalc.com/motocalc.htm>)

[9] "Low Cost 2 g/10 g Dual Axis iMEMS Accelerometers with Digital Output," Analog Devices Corporation, Norwood, MA. (http://www.analog.com/pdf/ADXL202_10_b.pdf)

[10] S.E. Fahlman, L.D. Baker and J.A. Boyan, "The Cascade 2 Learning Architecture," Technical Report CMU-CS-TR-96-184, Carnegie Mellon University, 1996.

[11] M.C. Nechyba and Y. Xu, "Cascade Neural Networks with Node-Decoupled Extended Kalman Filtering," *Proc. IEEE Int. Symp. on Computational Intelligence in Robotics and Automation*, vol. 1, pp. 214-9, 1997.