

# Augmenting Flight Training with AI to Efficiently Train Pilots

Michael Guevarra<sup>1</sup>, Srijita Das<sup>2,3</sup>, Christabel Wayllace<sup>2,3</sup>, Carrie Demmans Epp<sup>2</sup>,  
Matthew E. Taylor<sup>2,3</sup>, Alan Tay<sup>1</sup>

<sup>1</sup> Delphi Technology Corp: guevarrm@myumanitoba.ca, alantay@delphitechcorp.com

<sup>2</sup> University of Alberta: {srijital, wayllace, demmanse, matthew.e.taylor}@ualberta.ca

<sup>3</sup> Alberta Machine Intelligence Institute (Amii)

## Abstract

We propose an AI-based pilot trainer to help students learn how to fly aircraft. First, an AI agent uses behavioral cloning to learn flying maneuvers from qualified flight instructors. Later, the system uses the agent’s decisions to detect errors made by students and provide feedback to help students correct their errors. This paper presents an instantiation of the pilot trainer. We focus on teaching straight and level flying maneuvers by automatically providing formative feedback to the human student.

There is a critical shortage of commercial pilots worldwide: according to Oliver Wyman (2022), there will be a global gap of 34,000 pilots by 2025. Part of the problem is that pilots qualified to conduct such training are in very high demand and in short supply. Currently, human instructors guide trainees using flight simulator exercises. We posit that training an AI-enabled system to provide instruction for some tasks is a viable approach to reducing instructor workload while allowing them to interact with more students. This could increase the number of students per pilot trainer, improving the throughput of training pilots and therefore increase the supply of trained pilots.

In recent human-in-the-loop research, AI agents use advice from humans in different forms to speed up learning (Bignold et al. 2021; Cui et al. 2021; Christiano et al. 2017; Da Silva\* et al. 2020). Specifically, imitation learning allows an agent to learn to mimic a human’s behavior. Further, AI has been used for airplane flying (Morales and Sammut 2004; Sandström, Luotsinen, and Oskarsson 2022) as well as inside intelligent tutoring systems for multiple tasks ranging from student skill development (Georgila et al. 2019) to improving teaching strategies (Wang 2018). This paper presents a system where an agent mimics a qualified pilot and assists students in a pilot training program. Specifically, we focus on the straight and level flight task as a proof of concept. A trained agent identifies mistakes or sub-optimal maneuvers of trainee pilots inside a flight simulator and suggests corrective actions. To the best of our knowledge, this is a first attempt to use an AI tutor to train human pilots for flight.

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

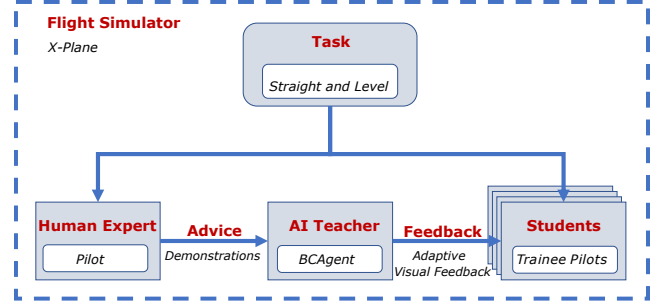


Figure 1: System Architecture

## System Architecture

Our proposed intelligent tutoring system is shown in Figure 1. It includes four general components: (1) Task, (2) Human expert, (3) AI teacher, and (4) Students. The task is sampled from a curriculum of flying maneuvers useful for learning to fly. Both the pilot and students perform tasks using the flight simulator X-Plane.<sup>1</sup> We use the fundamental flight maneuver “straight and level”<sup>2</sup> as the target task. There is a human expert (pilot) who is adept in the task and can provide advice to train the AI teacher. After the AI teacher is trained to mimic the human teacher, the AI teacher is used to guide students by providing different types of feedback based on their performance on the target task.

## Modeling the Pilot: Agent Training

We trained a decision-making agent to learn from pilot demonstrated trajectories. Experts demonstrated the “straight and level task” inside the flight simulator for 12.5 minutes. They operate under visual flight rules, i.e., clear weather to fly towards the target. The target direction is changed between every trial demonstrated by the pilot to account for diversity in trajectory collection. The final dataset consists of 25 trials of 30 seconds each with a randomized

<sup>1</sup><https://developer.x-plane.com/article/airport-data-apt-dat-file-format-specification/>

<sup>2</sup>Straight and level flight is a flight in which a constant heading and altitude are maintained. It is an essential flight maneuver used to form correct habits. All other flight maneuvers derive from straight and level.

goal between  $\pm 30$  degrees from the starting heading.

After collecting demonstrations, we trained a behavior cloning (BC) agent (Pomerleau 1988) to mimic the pilot’s policy. It is the simplest imitation learning technique where supervised learning is used to mimic the actions of an expert. The BC loss is defined as

$$L(\theta) = \sum_{i=1}^N \|\pi_{\theta}(s_i) - \pi_e(s_i)\|^2 \quad (1)$$

where  $\pi_{\theta}(s_i)$  is the current policy,  $\pi_e(s_i)$  is the expert policy,  $\theta$  is the training model parameter, and  $N$  refers to the number of state-action pairs in the training set. We used stable baselines (Raffin et al. 2019) to train the agent from demonstrations. The agent predicts the pitch and roll of the aircraft yoke. We evaluate the agent by measuring the average heading error over 10 evaluation trials with randomized heading as seen in Figure 2 and terminate when the error stops improving.

### Deploying the Teacher and Guiding Students

We deploy an agent when it makes similar decisions to the expert in previously unseen trials. Mimicking an expert pilot’s behavior is not enough to teach students. A teacher should be able to detect students’ mistakes and provide feedback to correct them. Therefore, we: (1) Recorded students’ poorly performed flights. (2) Asked the pilots to prepare annotated critiques on errors made. (3) Identified two main types of errors. (4) Used simple distance metrics to decide whether the agent agreed or not with the student’s decisions.

The identified student errors are due to (1) Not keeping altitude and airspeed constant and (2) Overshooting the target. To identify them, we compared how the agent and the student controlled the pitch and roll.

Let  $p_a(t)$ ,  $r_a(t)$ ,  $p_s(t)$ , and  $r_s(t)$  represent the pitch and roll produced by the agent and the student at a given time  $t$ . Then, for the first type of error:

$$|p_a(t) - p_s(t)| \geq D_1 \quad (2)$$

where  $D_1 \geq 0$  is a user-defined threshold. That is, the pitch difference is larger than the user-defined threshold.

For the second type of error:

$$|r_a(t) - r_s(t)| \geq D_2 \quad (3)$$

where  $D_2 \geq 0$  is a user-defined threshold. That is, the roll difference is larger than the user-defined threshold.

The agent uses Eqs. 2 and 3 to determine when to provide feedback. With respect to the type of feedback, this paper presents an instance of *informative tutoring*, a type of formative feedback that presents verification feedback, error flagging, and strategic hints on how to proceed (Shute 2008). Every time the user exceeds a threshold, the system displays a black square containing two lines denoting the user and the agent’s status (Figure 3). The position  $(x, y)$  of the middle of each line corresponds to  $(r_a(t), p_a(t))$  and  $(r_s(t), p_s(t))$  respectively, and the slope represents the roll angle.

This visualization shows how far apart the correct trajectory and position are. Instinctively, the student would aim to overlap both lines. Confirming the effectiveness of this and other types of feedback is part of future work.

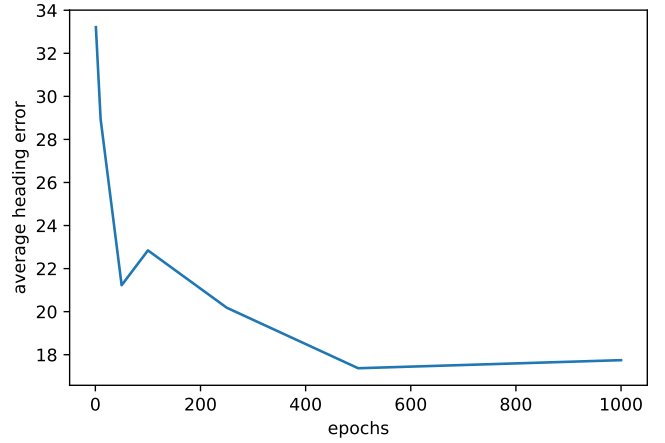


Figure 2: Average training time error of the behavior cloning teacher



Figure 3: Green and blue bars show visual tutor feedback

### Conclusion and Future Work

We presented an intelligent tutoring framework to autonomously train pilots inside a flight simulator using a simulated teacher. The teacher can provide different kinds of visual feedback to help students correct their mistakes. As future work, we will extend the simulated teacher to learn other complicated flight maneuvers like climbing and turning. We also plan to replace the BC teacher with reinforcement learning so that it can discover new policies not directly mentioned by the pilots. Lastly, we will study the impact of various kinds of feedback on student learning.

### References

- Bignold, A.; Cruz, F.; Taylor, M. E.; Brys, T.; Dazeley, R.; Vamplew, P.; and Foale, C. 2021. A conceptual framework for externally-influenced agents: an assisted reinforcement learning review. *Journal of Ambient Intelligence and Humanized Computing*.
- Christiano, P. F.; Leike, J.; Brown, T. B.; Martic, M.; Legg, S.; and Amodei, D. 2017. Deep Reinforcement Learning

from Human Preferences. In *Neural Information Processing Systems*, 4299–4307.

Cui, Y.; Koppol, P.; Admoni, H.; Niekum, S.; Simmons, R. G.; Steinfeld, A.; and Fitzgerald, T. 2021. Understanding the Relationship between Interactions and Outcomes in Human-in-the-Loop Machine Learning. In *IJCAI*, 4382–4391.

Da Silva\*, F. L.; Hernandez-Leal, P.; Kartal, B.; and Taylor, M. E. 2020. Uncertainty-Aware Action Advising for Deep Reinforcement Learning Agents. In *Proceedings of AAAI Conference on Artificial Intelligence*.

Georgila, K.; Core, M. G.; Nye, B. D.; Karumbaiah, S.; Auerbach, D.; and Ram, M. 2019. Using reinforcement learning to optimize the policies of an intelligent tutoring system for interpersonal skills training. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*.

Morales, E. F.; and Sammut, C. 2004. Learning to fly by combining reinforcement learning with behavioural cloning. In *Proceedings of the twenty-first international conference on Machine learning*, 76.

Oliver Wyman. 2022. After COVID-19, Aviation Faces a Pilot Shortage. <https://www.oliverwyman.com/our-expertise/insights/2021/mar/after-covid-19-aviation-faces-a-pilot-shortage.html>. Accessed: 2022-09-15.

Pomerleau, D. A. 1988. Alvin: An autonomous land vehicle in a neural network. *Advances in neural information processing systems*, 1.

Raffin, A.; Hill, A.; Ernestus, M.; Gleave, A.; Kanervisto, A.; and Dormann, N. 2019. Stable Baselines3. <https://github.com/DLR-RM/stable-baselines3>.

Sandström, V.; Luotsinen, L.; and Oskarsson, D. 2022. Fighter Pilot Behavior Cloning. In *2022 International Conference on Unmanned Aircraft Systems (ICUAS)*, 686–695. IEEE.

Shute, V. J. 2008. Focus on formative feedback. *Review of educational research*, 78(1): 153–189.

Wang, F. 2018. Reinforcement learning in a pomdp based intelligent tutoring system for optimizing teaching strategies. *International Journal of Information and Education Technology*.