

Exploring Methods of Sequential Transfer Learning for Re-Entry Guidance and Control

Benjamin Mitchell*

University of Central Florida, Orlando, FL 32816, USA

Tarek A. Elgohary†

University of Central Florida, Orlando, FL 32816, USA

Traditional transfer learning is a method of machine learning in which the neural network can learn from previously trained neural networks. Thus, the neural network can be fully trained with much less data and time than traditional machine learning requires. Transfer learning is a method of machine learning in which the neural network can learn from previously trained neural networks. Sequential Transfer Learning can be described as applying transfer learning iteratively to missions with parameters that vary increasingly in a certain direction within the design space. Thus, missions that may be well outside of the original mission's design space for the traditional transfer learning will still benefit from reduction in required data as well as training time. The neural network selected for this application is the long short-term memory neural network (LSTM). First, training data sets are generated for three missions, each with their own vehicle mass and geometric parameters. The largest of these three data sets is called the "original" mission. Both remaining data sets have the same amount of data, that is significantly less data than the original, one of them, the "link" mission, has mass and geometric parameters that are similar enough to apply transfer learning with the original mission. The other data set, the "target" mission, is much too different from the original mission to apply transfer learning. First the LSTM is trained on the "original" mission, then using transfer learning, the model is trained on the link mission. Ending the sequential process, the model is then uses transfer learning and trained on the "target" mission's data. The results show that the method developed in this paper has the capability and robustness to control various re-entry vehicles given different re-entry positions to within the desired landing accuracy requirements and meets the quick turnaround time required for high launch cadence missions.

I. Introduction

Recent advancements in rocketry, such as SpaceX's Starship, have made it possible to transport metric tons of supplies across the globe faster than any ship or airplane. This capability opens up new possibilities for responding to crises like natural disasters, where humanitarian aid is urgently needed. As the frequency of these events increases, the demand for rapid and effective response systems grows. However, the success of this approach hinges on the ability to quickly and accurately plan trajectories. Upon notification of a disaster, every second that the launch vehicle spends en route is another second that people remain in danger—making time a critical factor. Rapid mission design, therefore, becomes essential to ensure timely delivery of aid.

Traditional guidance and control algorithms are often tailored to a specific mission, requiring considerable time to adapt to new scenarios. The advent of Transfer Learning methods has enabled the usage of previously generated or recorded data from past missions to inform new models that are trained to predict trajectories for new missions. This adaptability allows the model to account for variability in mission parameters, such as initial conditions or target locations. However, the amount of additional training data and time required depends on the magnitude of variation between the new mission and the existing data, which is critical given the emphasis on rapid mission design.

A new sequential transfer learning is introduced to potentially improve in the transfer learning method and allow for greater variation between the missions with large and easily accessible datasets and missions that have little to no previous exploration or simulation. This is will be achieved by providing or accessing multiple "step" or "link" mission between the original mission and the objective mission. Within the design space, the step missions will exist somewhere

*Undergraduate Research Assistant, Mechanical and Aerospace Engineering, University of Central Florida, 12760 Pegasus Drive, Orlando, FL

†Associate Professor, Mechanical and Aerospace Engineering, University of Central Florida, 12760 Pegasus Drive, Orlando, FL

between the original and target mission. For example, the target mission may have a vehicle with a frontal surface area of 10 and the target mission will have 12. The step mission's vehicle may have an area of 11, in order to provide the neural network a path of training data that leads towards the target mission.

A similar approach to this sequential transfer learning has been used in some fields outside of engineering. In robotics, for example, one group has published a paper proposing an "iterative transfer learning" for collective motion-tuning tools. However, these methods differ, because the model used in this case is trained on data as it is generated and likely will be used on much smaller datasets. [1]

While the entire Entry, Descent, and Landing (EDL) sequence is critical for mission success, this study focuses specifically on the re-entry phase. Re-entry poses the greatest challenges due to intense heating and acceleration loads, making it a volatile segment of the trajectory. By narrowing the scope, we aim to demonstrate the accuracy of sequential transfer learning models. As with other machine learning methods, focusing the input and output data on this critical phase improves the model's ability to perform well on similar tasks, enhancing its effectiveness in stepping between missions and finally arriving at an entirely new mission. [2]

Multiple instances of EDL guidance algorithms that have emerged, the Apollo guidance algorithm being one of the first algorithm to be used in space flight application. This algorithm commands the bank angle of the reentry vehicle resulting in a decrease of effective lift. Thus, changing the range of the current trajectory, as needed, to match that of the reference trajectory. Of course, banking will also produce some lateral acceleration. The Apollo algorithm was primarily developed for low-lift vehicles like the Apollo reentry capsule however, later versions of guidance algorithms were implemented for different EDL approaches. For example, the space shuttle guidance algorithms utilized the drag acceleration. A commonly used method in the current generation of EDL guidance algorithms are predictor-corrector algorithms. These algorithms first generate a predicted EDL trajectory based on current state information and then corrects by applying necessary landing constraints to the trajectory. The Apollo guidance algorithm is generally viewed as the original algorithm which many current algorithms and models are based. It is also the algorithm used to generate trajectories in this paper. [3]

II. Methodology

A. Entry, Descent, and Landing Dynamics

First, the EDL equations of motion are presented as follows in the Earth-fixed reference frame:[4]

$$\dot{r} = {}^R V \sin \gamma \quad (1)$$

$$\dot{\theta} = \frac{{}^R V \cos \gamma \cos \psi}{r \cos \phi} \quad (2)$$

$$\dot{\phi} = \frac{{}^R V \cos \gamma \sin \psi}{r} \quad (3)$$

$${}^R \dot{V} = -\frac{D}{m} - \frac{G}{m} \sin \gamma \quad (4)$$

$$\dot{\gamma} = \left(\frac{L}{m} \cos \sigma - \frac{G}{m} \cos \gamma + \frac{{}^R V^2}{r} \cos \gamma \right) / {}^R V \quad (5)$$

$$\dot{\psi} = \left(\frac{L \sin \sigma}{m \cos \gamma} - \frac{{}^R V^2}{r} \cos \gamma \cos \psi \tan \phi \right) / {}^R V \quad (6)$$

where r is the distance of the spacecraft to the center of the planet, ${}^R V$ is the speed of the spacecraft relative to the planet, γ is the flight path angle, θ is the longitude of vehicle position, ψ is the heading angle, ϕ is the latitude of vehicle position, T is the thrust force, m is the mass of the spacecraft, ζ is the angle measured from thrust vector to velocity vector, ϵ is the angle measured from the direction of velocity to the lift-velocity plane thrust projection, D is the drag force, g is the gravitational acceleration, $\omega_{\oplus}$ is the constant planet rotational speed, σ is the bank angle, and L is the lift force.

The following assumptions are made:

- Vehicle can be approximated as a point mass
- Planet rotation is constant about north pole

- Planet-fixed and inertial coordinate systems are aligned at time $t = 0$
- Constant mass
- Drag acts in direction opposite the velocity
- Lift force is perpendicular to velocity
- Gravity force is directed along a vector from point mass to the center of the planet

B. Apollo Guidance Algorithm

The Apollo 10 guidance algorithm is the selected algorithm for study, since the flight data from this mission is publicly available. This guidance algorithm is a bank angle steering algorithm, which modulates the lift and affects the predicted range. The predicted range-to-go is computed as follows:[5]

$$R_p = R_{ref} + \frac{\partial R}{\partial D} (D - D_{ref}) - \frac{\partial R}{\partial \dot{r}} (\dot{r} - \dot{r}_{ref}) \quad (7)$$

where R_p and R_{ref} are the predicted and reference range-to-go, respectively, D and D_{ref} are the current and reference drag force, respectively, and $\dot{r}$ and $\dot{r}_{ref}$ are the current and reference altitude rate, respectively. Then, the desired vertical component of the lift to drag ratio, $(\frac{L}{D})_C$, is computed as:

$$\left(\frac{L}{D}\right)_C = \left(\frac{L}{D}\right)_{ref} + \frac{K_3(R - R_p)}{\frac{\partial R}{\partial (\frac{L}{D})}} \quad (8)$$

where $(\frac{L}{D})_{ref}$ is the reference lift to drag ratio and K_3 is the over-control gain that amplifies the range error correction.[6] Once the predicted range is close enough to the desired range, the bank angle is commanded based on a reference trajectory via the following equation:

$$\sigma = \arccos\left(\frac{(L/D)_C}{L/D}\right) * K2ROLL \quad (9)$$

where $K2ROLL$ is the bank reversal and is equal to ± 1 [5].

C. Transfer Learning

Transfer learning is a machine learning method that employs knowledge from previously trained neural networks to solve new tasks. [7] Because this method allows the neural network to draw from previously trained networks instead of starting from scratch, it addresses the long training time and large data requirements typically associated with machine learning. This method has been used in a variety of applications, from medical classification to speech recognition. [8, 9] This method has been recently used in Mars EDL applications.[3] With an application like rocket-based material transport, which requires rapid mission design with unseen landing environments, new vehicles, different centers of gravity, etc., transfer learning is a promising approach to aid in that mission.




Fig. 1 LSTM Cell [10]

In order to learn a new task via transfer learning, it is necessary that an initial task be trained using a type of neural network. The selected neural network type must match between the initial task and the transfer learned task. Due to its

ability to handle temporal data, a recurrent neural network (RNN) is selected over the traditional feed forward neural network. [11] The specific RNN selected in this study is the long short-term memory neural network (LSTM), which is shown in figure 1. This RNN cell type remedies the vanishing/exploding gradients issue attributed to standard RNN architectures. The LSTM functions with a series of gates. The first gate, the forget gate, is depicted on the left side of the figure and serves to determine the percentage of long-term memory to retain. The middle gate, the input gate, updates the long-term memory. The rightmost gate, the output gate, updates the short-term memory which is the output of the cell. The activation functions associated with the three gates are sigmoid and hyperbolic tangent functions.




Fig. 2 Transfer learning applied to rapid mission design

Transfer learning in the context of this rapid mission design problem is highlighted in figure 2. With the first mission, a Monte Carlo simulation is created, which produces simulated flight data or real flight data is obtained. With the flight data, the selected neural network is trained using python. For this problem, the altitude and vehicle relative velocity for the duration of the EDL trajectory for each mission serve as the inputs. The outputs are the computed control input (bank angle commands) associated with each input trajectory. The neural network is learning the relationship between the state of the vehicle and the desired control for each EDL mission. Once sufficient accuracy has been met for the trained neural network, new, specific mission output trajectories can be generated in a matter of seconds.

This method is sufficient for all mission descriptions similar to the original training data. For new missions that are outside of the scope of sufficient accuracy for the initially trained neural network, a small (relative to the first) set of new data will be obtained through either simulation or real flight data. Then, as before, the input and output data are structured as required by the neural network architecture (state inputs and control outputs). Unlike before, instead of beginning the training with a fresh neural network, the previously trained neural network is loaded in. The last two layers, the fully connected and regression output layers, are removed and replaced with new ones. This signals to the network to keep all upstream weights and biases the same as the previously trained network, but to prepare to solve a new task where the outputs do not necessarily match those of the previous task. With this replacement, the learning from the previous task has been transferred to the new task. Once this transfer learning step has been completed, then training may commence. Generally, this training will require much less time and data than mission 1 required. This is shown in figure 2. As before, once sufficient accuracy has been met with the transfer learned network, the user is able to input any related mission trajectory and the network will output the corresponding control trajectory.

Given this ability to solve completely new trajectories quickly and accurately based on previous missions, transfer learning is a promising method to achieve rapid mission design. While the transfer model needs less training data, it is still able to achieve a similar accuracy the initial LSTM model has with respect to its corresponding.

D. Sequential Transfer Learning

Sequential Transfer Learning has a very similar structure to traditional transfer learning. The machine learning model that is used (LSTM) also remains the same. However, this method is used to generate predictions for missions

that are significantly different from the original mission.

Sequential transfer learning utilizes the same steps of training an LSTM model on the bulk data, of the original mission, and then the model has it's last two layers removed and replaced with untrained layers. Effectively clearing all weights within these layers.




Fig. 3 Transfer learning applied to rapid mission design

The new model is first transferred to and trained on the smaller set from mission 2, this is because mission 2's parameters are only 10% greater than the original mission. The divergence from traditional transfer learning lies in the following steps. For the second time, the trained model has its last two layers, the fully connected and regression output, replaced with untrained layers. This twice-augmented sequential transfer learning model is then trained on a third and final dataset. This process is summarized in figure 3.

In some circumstances, the original and final missions differ so significantly that direct transfer learning would not allow the model to be trained on enough data that is sufficiently relevant to the final mission. However, with small sets of data for missions within the design spaces of both the original and final missions, sequential transfer learning can be utilized.

III. Results

The individual studies in this process and the missions that constitute them will be described here in detail. Most importantly, the mission parameters that deviate these mission from their origin, the Apollo 10 Mission.

The Transfer Learning study missions are as follows.

Parameter	Mission 1	Mission 2
A (Area m ²)	A	$A \times (1 + 0.2)$
m (Mass kg)	m	$m \times (1 + 0.2)$

Table 1 Mission Parameter and Transfer Learning Study

Table 1 is a thorough way to state that the second mission of this study will have all parameters identical to Apollo 10 with the exception of the vehicle mass and cross-sectional, surface area. These two parameters have a 10 percent difference from the Apollo 10 baseline.

Here, Table 2 shows a third mission and its parameters with a 20 Percent difference from the Apollo 10 baseline. Again, the motivation for this study is to see whether or not an intermediate mission can aid in the training of the transfer learning model. During both studies, Mission 1 (Apollo 10) will have 500 trajectories to train on, where Mission 2 and 3 both have 500 trajectories.

Parameter	Mission 1	Mission 2	Mission 3
A (Area m ²)	A	$A \times (1 + 0.1)$	$A \times (1 + 0.2)$
m (Mass kg)	m	$m \times (1 + 0.1)$	$A \times (1 + 0.2)$

Table 2 Mission Parameters for Sequential Transfer Learning Study

For both of these studies, the comparison of the Apollo GNC output control trajectories are compared to the LSTM predicted output control trajectories.




Fig. 4 Transfer learning Output Control Trajectory




Fig. 5 Sequential Transfer learning Output Control Trajectory

An interesting result shows that using the transfer learning method between missions 1 and 3 4, significantly

outperforms the sequential transfer learning method 5. The trend of the Apollo trajectory is captured with considerable accuracy while using transfer learning, but Sequential Transfer Learning predicts a better performing control trajectory that is, strangely dissimilar from the Apollo GNC trajectory. For a more quantified analysis, the average Mean Squared Error (MSE) between the actual and predicted control output, of all the trajectories in the Monte Carlo analysis of Mission 3, is shown in Table 3.

	Sequential Transfer Learning	Transfer Learning
Mean MSE	1.0437	0.0610

Table 3 Mean MSE between Apollo and LSTM Control Trajectories

This significant contrast in model performances between the two methods is rather unexpected, simply due to the fact that the Sequential Transfer Learning model is trained on more data overall than the Transfer Learning methods. One possible explanation for this result could be that the addition of Mission 2 training data forced the model to learn nuanced dynamics specific to that mission alone. As a result, this may reduce the potential for learning the nuances relevant to the target mission (Mission 3).

It is possible that this result shows that presenting data that are not precisely what the model is to be trained for may become a "distraction" from the tasks it is ultimately meant to predict generate predictions for.




Fig. 6 Error of the Sequential Transfer Learning Method




Fig. 7 Error of the Transfer Learning Method

Another surprising result shown indicates that, while the Sequential Transfer Learning method does not capture the Apollo GNC control trajectories nearly as well as the direct Transfer Learning method, the trajectories it has generated result in a significantly smaller final error 6 than it's counter 7. As exhibited in these two figures, Sequential Transfer Learning is shown to outperform Transfer Learning on all of the error components including: Down-Range, Cross-Range, and altitude. For any of the methods used, minimizing error is the primary objective, and this may show that, disregarding the fact that Sequential Transfer Learning does not capture the Apollo GNC trajectory, it may still be considered a model performing better than the Transfer Learned model.




Fig. 8 Sequential Transfer Learning Final geodetic State for all 500 Mission 3 Trajectories




Fig. 9 Transfer Learning Final geodetic State for all 500 Mission 3 Trajectories

IV. Discussion

Despite the inaccuracies in the model performance, with respect to the MSE between the Apollo and LSTM Predicted control trajectory, the Sequential transfer learning model's predicted trajectories were able to achieve a significantly lower final error state, which is the primary objective of the model. It is important to note that this Sequential Transfer Learned model has shown to achieve these low error states consistently over the 500 test trajectories. To further understanding of Sequential Transfer Learning and it's possible favorable applications, a larger study with upwards of 5-10 Training configurations is recommended. These results show a promising conclusion that suggests the possibility that Sequential Transfer Learning may offer benefits over Transfer Learning and should be further investigated. Among these advancements in the study, it will also be important to consider and vary other quantities such as the LSTM's layers and structure, as well as the training epochs, and batch size. Using control data produced using different methods, such as the Alpha-Beta method and the predictor-corrector method is also under consideration, as these GNC methods might produce trajectories that are easier patterns for the LSTM model to learn.

V. Acknowledgements

This material is based on research sponsored by Air Force Research Laboratory (AFRL) under agreement number FA9453-2 1 -2-0064. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements , either expressed or implied, of Air Force Research Laboratory (AFRL) and or the U.S. Government.

References

- [1] Abpeikar, S., Kasmarik, K., and Garratt, M., "Iterative transfer learning for automatic collective motion tuning on multiple robot platforms," *AAS Guidance, Navigation, and Control Conference*, 2023.
- [2] Chapman, D. R., "An Approximate Analytical Method for Studying Entry into Planetary Atmospheres," *NACA Technical Note*, Vol. 4276, 1958.
- [3] Jones, L., and Elgohary, T. A., "An Investigation of Transfer Learning for Guidance and Control of Mars Entry, Descent, and Landing Missions," *AAS Guidance, Navigation, and Control Conference*, 2024.
- [4] Hicks, K. D., *Introduction to astrodynamic reentry*, Books Express Publishing, 2009.
- [5] Carman, G., Ives, D., and Geller, D., "Apollo-Derived Mars Precision Lander Guidance," *AIAA*, Vol. 98-4570, 1998.
- [6] Mendeck, G., and Craig, L., "Entry Guidance for the 2011 Mars Science Laboratory Mission," *AIAA Atmospheric Flight Mechanics Conference*, 2011.
- [7] Jialin Pan, S., and Yang, Q., "A Survey on Transfer Learning," *IEEE TRANSACTIONS ON KNOWLEDGE AND DATA ENGINEERING*, Vol. 22, No. 10, 2010.
- [8] Ghosh, A., Soni, B., and Baruah, U., "Brain Hemorrhage Classification Using Leaky ReLU-Based Transfer Learning Approach," *Lecture Notes in Networks and Systems*, Vol. 698, 2023.
- [9] Ta, B., Le, N., and Do, V., "Transfer learning methods for low-resource speech accent recognition: A case study on Vietnamese language," *Engineering Applications of Artificial Intelligence*, Vol. 132, 2024.
- [10] Varikuti, M., "LSTM Networks," , 2021. URL <https://medium.com/mllearning-ai/lstm-networks-75d44ac8280f>.
- [11] Wang, H., and Elgohary, T., "A Simple and Accurate Apollo-Trained Neural Network Controller for Mars Atmospheric Entry," *International Journal of Aerospace Engineering*, Vol. 2020, No. 3793740, 2020.