

Predicting Aircraft Trajectories via Imitation Learning

Theocharis Kravaris, Alevizos Bastas, George A. Vouros
University of Piraeus,
Greece

Problem Background

- Complex airspace (30000 aircraft)
- ATM needs to handle:
 - Greater Complexity
 - Larger Volumes of Traffic
- Safety Guarantees
- Solution
 - Efficient Planning
 - High fidelity aircraft trajectory prediction
- New trajectory-based ATM paradigm
- Improve the predictability of aircraft trajectories



Traffic over Europe (flihtadar24.com)

Objective

- We aim at learning a policy that models the way stakeholders shape the evolution of aircraft trajectories w.r.t. contextual conditions (e.g. weather conditions), considering historical enriched trajectories as experts' demonstrations.
- In the aviation domain, a *trajectory* is defined as the description of movement of an aircraft both in the air and on the ground (4D).
- Enriched with various data, producing the *enriched trajectory*.

$$s_{r,i} = \langle p_i, t_i, v_i \rangle, i \in [0, |T| - 1]$$

(Data-driven) Problem Formulation

Given the historical trajectories and a cost function the general objective is:

$$T_{\pi} = \underset{\pi}{\operatorname{argmin}} \mathbb{E}_{\pi} [c(\langle p, t, v \rangle, a)]$$

Question: What c models and how is it engineered/learned?

Approach

- **Question:** Does inverse RL help?
- We aim to imitate trajectories demonstrated by “experts”.

In this work we explore and demonstrate the potential of GAIL*

- GAIL employs a generative trajectory model G that models π and a discriminative classifier D that distinguishes between the distribution of state action pairs generated by π and the demonstrated data.
- Learns the optimal policy from expert demonstrations directly, efficiently, scaling to large, continuous state-action spaces, and does not need to derive a cost function.

- It aims to bring the distribution of the state-action pairs of the imitator as close as possible to that of the expert.

- GAIL’s Objective:
$$\min_{\pi} \max_{D \in (0,1)^{S \times A}} \mathbb{E}_{\pi} [\log D(s, a)] + \mathbb{E}_{\pi_E} [\log(1 - D(s, a))] - \lambda_H H(\pi)$$

* Generative Adversarial Imitation Learning

[Ho Jonathan, Stefano Ermon 2016]

Implementation of GAIL

➤ Use of

- Trusted Region Policy Optimization (TRPO)
- Generalized Advantage Estimation (GAE)

➤ Pretrained G using Behavioral Cloning

Algorithm 1 GAIL

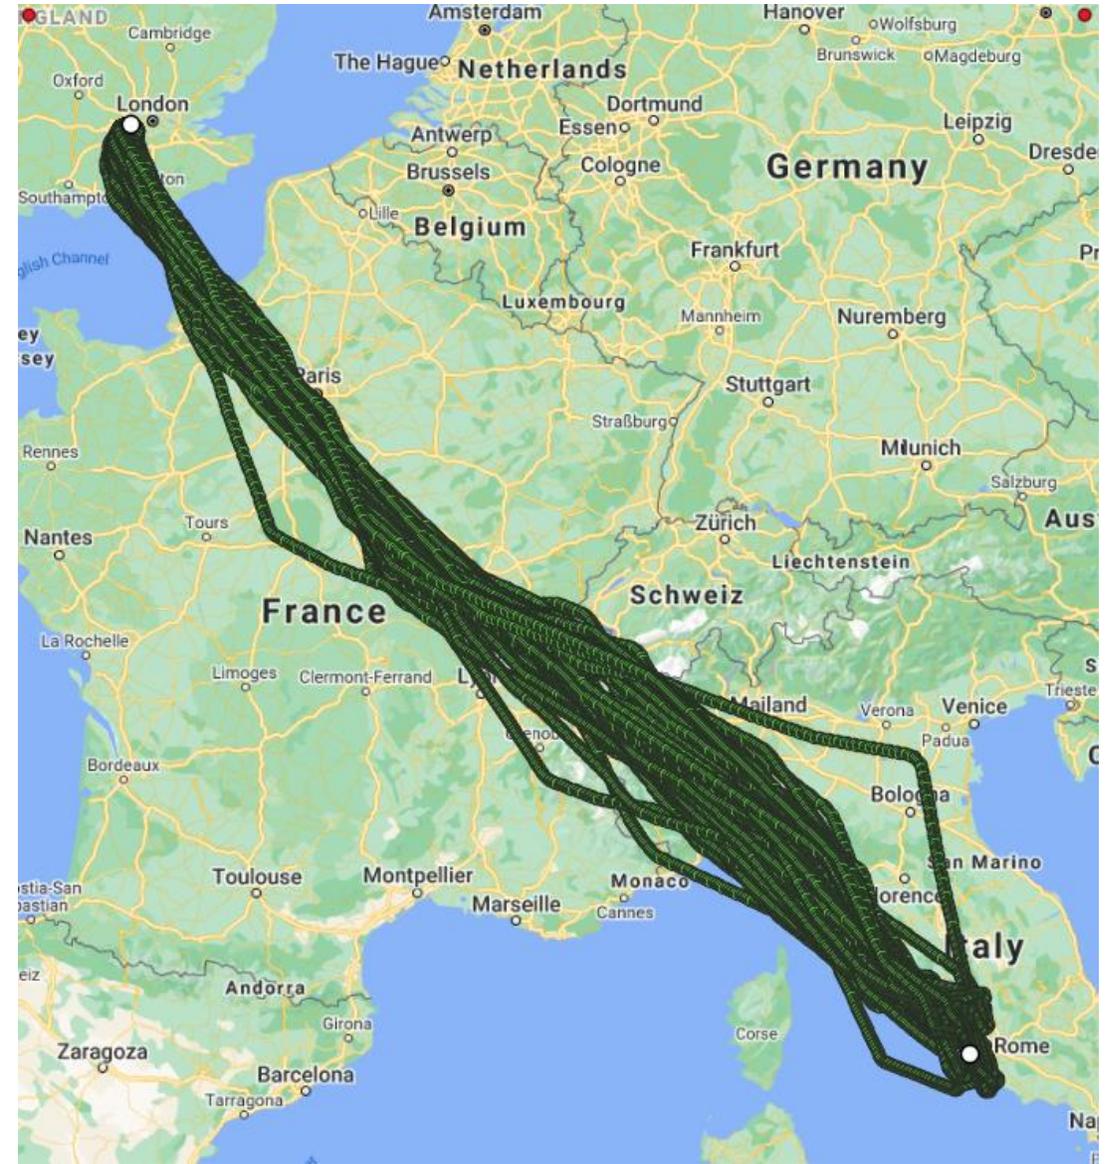
- 1: **Input:** Expert trajectories $\tau_E \sim \pi_E$, initial policy π_{θ_0} and discriminator parameters w_0
 - 2: **Output:** Policy π_{θ}
 - 3: Initialize policy using Behavioral Cloning.
 - 4: **for** $i=0,1,2,\dots$ **do**
 - 5: Sample trajectories $\tau_i \sim \pi_{\theta_i}$
 - 6: Update D parameters w with the gradient
 - 7: $\hat{\mathbb{E}}_{\tau_i} [\nabla_w \log(D_w(s, a))] + \hat{\mathbb{E}}_{\tau_E} [\nabla_w \log(1 - D_w(s, a))]$
 - 8: Estimate advantages $\hat{A}^{GAE(\gamma, \lambda)}$, according to $\pi_{\theta_{old}}$
 - 9: Take a policy step using the TRPO rule with cost function $\log(D_w(s, a))$:
 Take a KL-constrained natural gradient step
 with $\hat{\mathbb{E}}_{\tau_i} [\nabla_{\theta} \frac{\pi_{\theta}(a|s)}{\pi_{\theta_{old}}(a|s)} \hat{A}^{GAE(\gamma, \lambda)}]$ subject to
 $\mathbb{E}_{s \sim \rho_{\theta_{old}}} [D_{KL}(\pi_{\theta_{old}}(\cdot|s) \parallel \pi_{\theta}(\cdot|s))] \leq \delta$
 - 10: **end for**
-

(Data-driven) Problem Formulation

- State (continuous)
 - Spatiotemporal Information
 - Longitude (degrees)
 - Latitude (degrees)
 - Altitude (feet)
 - Timestamp
 - Weather Information, aircraft models
- Actions (continuous) in Δt
 - Δ Longitude (degrees)
 - Δ latitude (degrees)
 - Δ altitude (feet)

Data Sources

- Radar tracks for flights between 3 Origin-Destination pairs considering **short and long trajectories**
 - Barcelona to Madrid (BCN-MAD) during July 2019 (308 trajectories)
 - London Heathrow to Rome Fiumicino (LHR-FCO) during July 2019 (219 trajectories)
 - Helsinki to Lisbon (HEL-LIS) during July 2019 (44 trajectories)
- Weather data obtained from National Oceanic and Atmospheric Administration (NOAA)
- Aircraft models



Experimental Results (1/2)

- We measure the **prediction accuracy** achieved at the **pre-tactical phase** and at **the tactical phase** introducing a parameter M in $\{0, 0.2, 0.5, 0.7\}$ showing the **ratio of trajectory length where the starting prediction state (from the origin) is**

- We report on accuracy using the following measures:

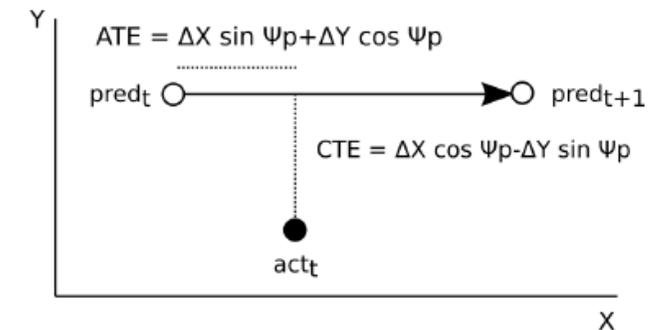
- **Root Mean Square Error (RMSE)** in meters

$$RMSE(var) = \sqrt{\frac{1}{N} \sum_{i=1}^N (var_{pred} - var_{actual})^2}$$

$$RMSE_{3D} = \sqrt{\frac{1}{N} \sum_{i=1}^N \frac{\sqrt{\sum_{d=1}^{Dim} (var_{d_{pred}} - var_{d_{actual}})^2}}{Dim}}$$

- **Along-Track Error (ATE), Cross-Track Error (CTE), and Vertical deviation (V)**

- **Estimated Time of Arrival ($|T| * \Delta t$)**



Experimental Results (2/2)

BCN-MAD									
	M	Long	Lat	Alt	3D	ATE	CTE	VE	ETA
GAIL	0	11994.99	6214.72	282.63	8005.49	13923.3	6392.61	574.13	263.06
	0.2	10021.46	5549.65	317.87	6774.39	11541.14	6159.43	516.15	259.81
	0.5	8680.78	5103.39	391.5	5977.16	11360.9	6330.23	510.78	239.35
	0.7	6327.37	4078.96	273.41	4519.16	9751.45	5284.18	375.73	158.9
LHR-FCO									
	M	Long	Lat	Alt	3D	ATE	CTE	VE	ETA
GAIL	0	23371.12	20888.65	372.33	18351.99	18689.42	17874.3	636.2	457.1
	0.2	24427.65	20568.25	359.35	18733.01	19273.79	19362.4	621.78	615.12
	0.5	20274.53	18497.25	370.98	16209.15	15758.31	16399.46	629.06	791.83
	0.7	14313.57	14444.13	539.25	12126.93	13205.18	11294.95	659.35	910.28
HEL-LIS									
	M	Long	Lat	Alt	3D	ATE	CTE	VE	ETA
GAIL	0	88448.14	95173.02	1096.41	75950.75	77341.27	59731.61	1074.71	801.44
	0.2	91184.7	100957.56	1062.17	79921.51	81309.09	52941.16	1052.04	978.19
	0.5	90334.3	92006.24	1090.32	76575.08	81468.87	49669.47	1252.01	1080.75
	0.7	77587.38	76998.23	1966.88	64771.15	81990.64	46206.48	1691.44	1113.12

Concluding Remarks

The **lack of benchmarks and study of different cases** hinders the systematic comparison of trajectory prediction methods.

Some **observations regarding state-of-the-art approaches** are as follows:

- **Aircraft Trajectory Prediction Made Easy with Predictive Analytics** [S. Ayhan and H. Samet. 2016]
 - computes the most likely sequence of states derived by a Hidden Markov Model
 - authors conclude that the mean value for the cross-track error is 12.601km
- **Predicting Aircraft Trajectories: A Deep Generative Convolutional Recurrent Neural Networks Approach** [Y. Liu and M. Hansen. 2018]
 - Tree-based matching algorithm combined with convolutional recurrent neural network
 - 3D RMSE increased by 1.21
- **AppLearn** [C. Spatharis, K. Blekas and G. A. Vouros 2020]
 - Apprenticeship learning imitation learning approach assumes linearity on the cost function and can be **considered as a baseline approach (applied to the same OD pairs)**
 - The 3D RMSE reported is 2 times greater than our approach for the BCN-MAD pair, and approximately 7.2 and 21.3 times, for the LHR-FCO and HEL-LIS pairs

Thank you

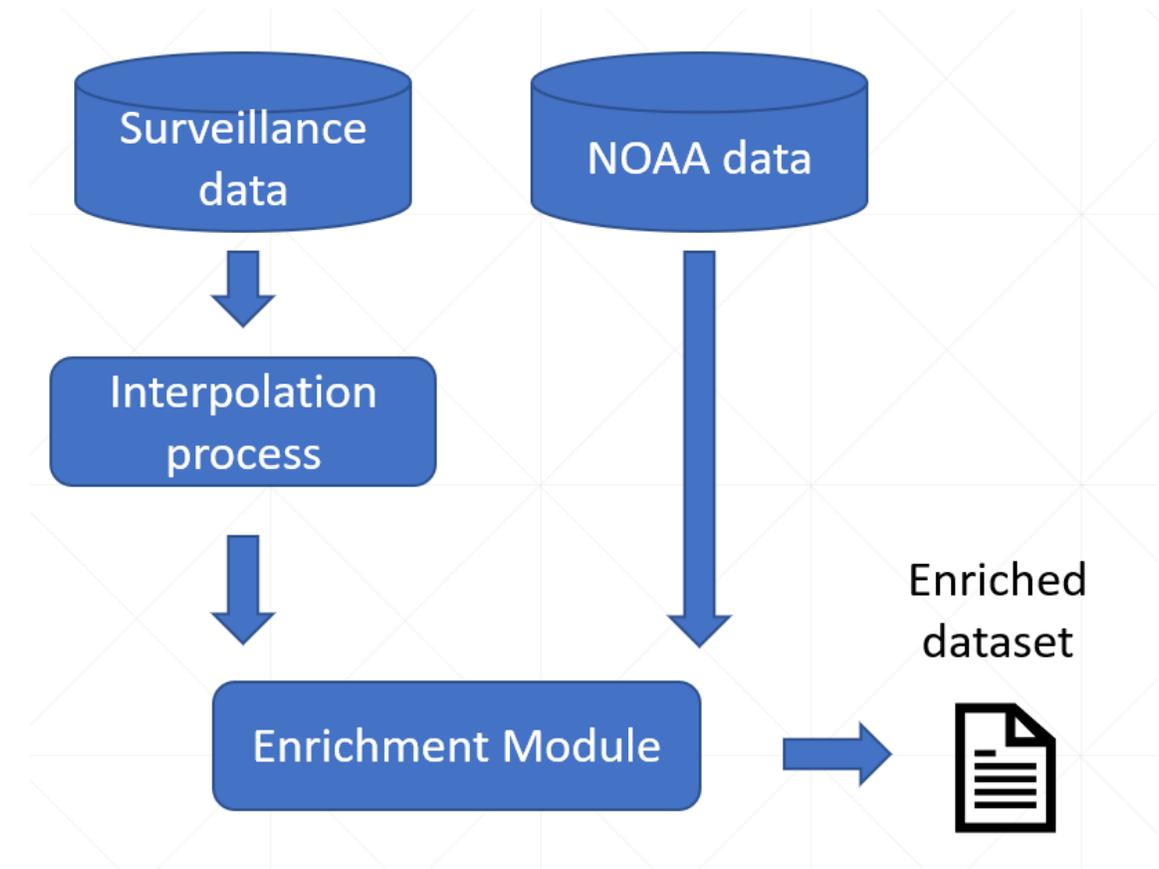
This work has been funded and completed in the context of the ENGAGE KTN Catalyst Fund project entitled "Data-Driven Trajectory Imitation with Reinforcement Learning", and it is partially funded by University of Piraeus Research Center.



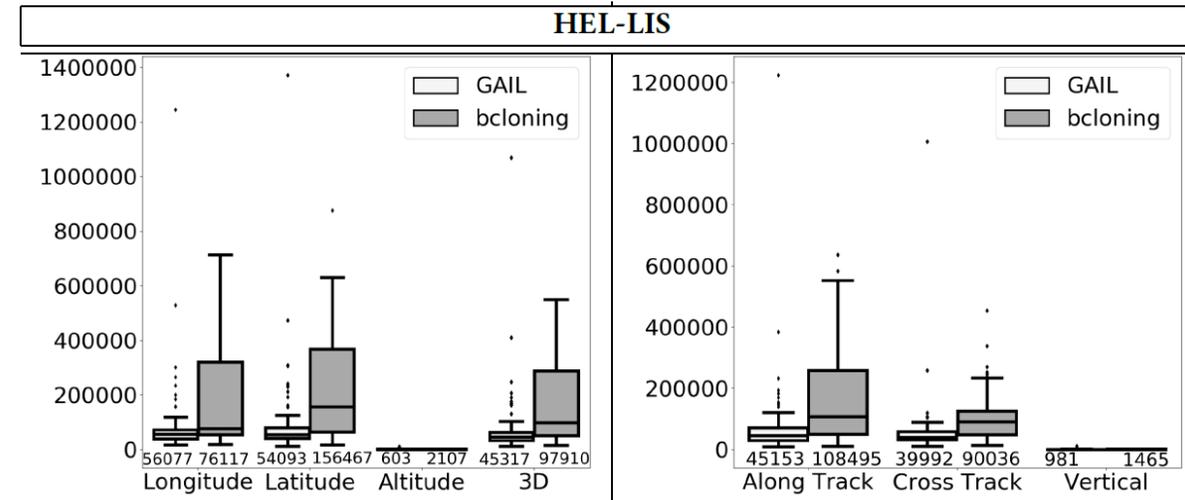
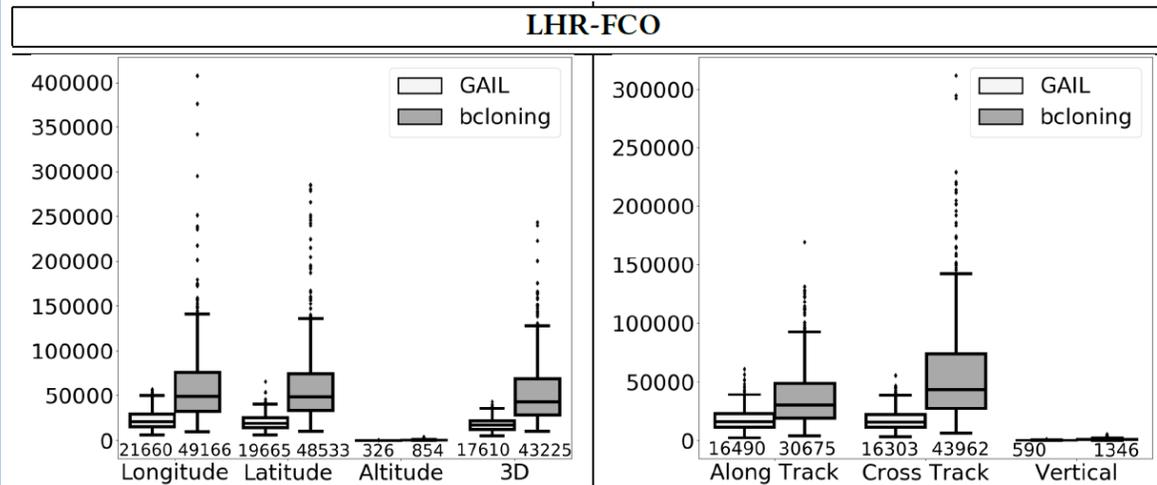
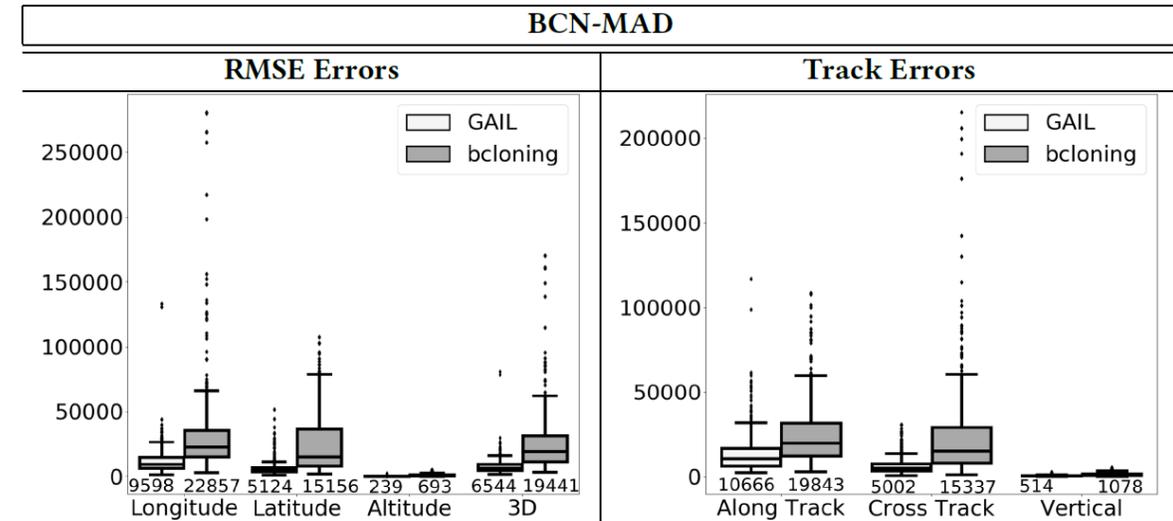
Data Sources: Processing

Per point:

- Spatiotemporal Information
 - Longitude (degrees)
 - Latitude (degrees)
 - Altitude (feet)
 - Timestamp
- Weather Information
 - Pressure surface (Pa)
 - Relative humidity (%)
 - Temperature (K)
 - Wind speed gust (m/s)
 - U component of wind (m/s)
 - V component of wind (m/s)
- Aircraft Model



Experimental Results (3/3)



Goal of the Paper

- Improve the predictability of aircraft trajectories
- Incorporate preferences, interests, constraints of different stakeholders (AUs, ATCs, etc.)
- Reduce the mismatch between planned and flown trajectories
- Reduce uncertainty
- Reduce flight inefficiencies
- Better planning of operations for Airspace Users