

A Trajectory-Optimization Framework for Transition-Phase Control of Small-Scale Thrust-Vectoring V/STOL Air Vehicles

Dr. Mohammed Aroussi*

Director, Research & Development Division

FlyWei Robotics

April 21, 2026

Abstract

The controllability envelope of a vertical/short take-off and landing (V/STOL) air vehicle during transition is dictated chiefly by the orientation of the thrust-vectoring nozzle, whose time-history must be chosen so as to reconcile actuator saturation, stall limits, and mission-level performance indices. This paper advances a fully model-based pipeline in which the transition corridor is constructed directly from an attainable-equilibrium-set (AES) projection of the longitudinal dynamics, rather than by assembling piecewise envelope fragments. A scaled F-35B-class prototype equipped with a three-bearing swivel duct (3BSD) nozzle and a lift fan is adopted as the reference platform, and the coupling between the propulsive jet and the airframe flow field is embedded in the dynamic model through empirical jet-induced-loss closures. The redundancy inherent in the multi-actuator longitudinal channel is resolved by a mode-dependent control-principle map that exposes only four pilot handles and delegates actuator assignment to a regularised pseudo-inverse allocator. On top of this architecture, the full transition problem is recast as a continuous-time optimal-control problem with mission- and handling-quality constraints, then converted into a finite nonlinear programme via Hermite–Simpson collocation and solved with a sequential quadratic programming (SQP) method. Numerical experiments on four canonical missions—vertical and short take-off, descent to hover, and forward-to-hover conversion—demonstrate that the framework simultaneously respects the AES-derived corridor, produces pilot-compatible stick time-histories, and improves smoothness and energy metrics relative to heuristic schedules. The paper also discusses conditioning and convergence issues that arise in the non-dimensionalised SQP formulation and reports the sensitivity of the optimum to the selected cost functional.

Keywords: thrust vectoring; V/STOL; transition corridor; attainable equilibrium set; nonlinear trajectory optimisation; sequential quadratic programming; jet-induced aerodynamics; control allocation; small unmanned aerial systems.

*Correspondence: Aroussi@flywei.co.uk

Contents

1	Introduction	3
1.1	Problem setting and motivation	3
1.2	Technical challenges	3
1.3	Related work	4
1.4	Contributions	4
1.5	Outline	5
2	Vehicle description and longitudinal dynamic model	5
2.1	Reference platform and notation	5
2.2	Propulsion architecture and control principle	6
2.3	Jet-induced lift deficit and moment perturbation	6
2.4	Longitudinal equations of motion	8
3	Transition corridor via the attainable equilibrium set	8
3.1	Rationale for the AES projection	8
3.2	Trim problem and boundary constraints	9
3.3	Structure of the resulting corridor	9
4	Control-principle map and allocation	10
4.1	Rationale for a mode-dependent map	10
4.2	Flight-mode stick-response map	10
4.3	Control allocation	11
5	Transition as an optimal-control problem	12
5.1	Continuous-time formulation	12
5.2	Mission-dependent cost functionals	12
5.3	Terminal and path constraints	13
5.4	Non-dimensionalisation and numerical conditioning	14
5.5	Hermite–Simpson collocation and SQP	14
6	Numerical experiments	17
6.1	Experimental protocol	17
6.2	Transition during take-off	17
6.3	Mid-cruise corridor traversal	18
6.4	Transition during landing	19
6.5	Discussion	20
7	Conclusions and outlook	21
A	Derivation of the linearised allocation matrix	23

1 Introduction

1.1 Problem setting and motivation

Thrust-vectoring V/STOL configurations are unusual in the rotary- and fixed-wing zoology because the same airframe is asked to hover like a rotorcraft, cruise like an aeroplane, and negotiate the intermediate regime without any substantive geometric reconfiguration. This last property distinguishes them sharply from tiltrotor or tiltwing vehicles: rather than swinging the entire lifting surface, only the exhaust of the propulsion system is rotated, so the outer mould line —and hence the external aerodynamic signature— stays essentially invariant. At full downward deflection of the three-bearing swivel duct (3BSD) nozzle, the vehicle hovers on the superposition of the nozzle jet and the forward lift-fan. As the duct rotates towards the horizontal, the contribution of direct lift fades and the wing gradually assumes the load, until, at null deflection, the platform behaves as a conventional fixed-wing aircraft. The intermediate corridor is therefore the critical bottleneck of the flight envelope —an aerodynamically singular phase in which two essentially different lift mechanisms act simultaneously, the two associated moment equations are tightly coupled, and several physical effectors compete for the same virtual handle.

1.2 Technical challenges

A systematic treatment of the transition phase must come to terms with four obstacles, which we state in the order in which they appear in the modelling and design pipeline.

(i) *Jet-airframe interaction.* The exhaust of the lift fan and of the 3BSD nozzle establishes a low-pressure envelope beneath the wing and the tail; the free-stream, in return, deflects the jet and alters the effective thrust direction. Neither effect is negligible at the airspeeds and thrust loadings characteristic of transition, and neither is captured by wind-tunnel data collected with the propulsion system idle. The resulting lift deficit and moment perturbation must therefore be injected into the dynamic model through empirical closures whose functional form depends on the nozzle geometry and on the effective velocity ratio.

(ii) *Corridor determination.* The tilt angle of the duct is not a free variable: for every airspeed there exists a narrow band of duct deflections at which the vehicle admits an equilibrium that honours both the stall limit of the wing and the available thrust authority of the propulsion system. Outside that band the system has no trim, and the controller is then operating in a provably uncontrollable regime. A precise construction of this corridor is a prerequisite for any trajectory-level design, and generic envelope heuristics imported from tilt-wing or multi-rotor practice are too loose to be useful.

(iii) *Control redundancy.* The longitudinal channel alone counts four physical effectors: the lift-fan throttle, the 3BSD throttle, the duct tilt angle, and the elevator. A pilot cannot be asked to coordinate four inputs simultaneously; an allocator is required that translates a compact set of pilot commands into physically consistent actuator motions, and does so in a way that varies gracefully with the flight mode.

(iv) *Mission-dependent optimisation.* A short take-off from a pitching deck, a vertical take-off from a clearing, and a forward-to-hover descent on final approach all require the same vehicle

to negotiate the same corridor, but the relative weights assigned to smoothness, time, fuel, and ground footprint differ markedly. The controller must therefore be derived from a cost functional that is parameterised by the mission rather than being frozen at design time.

1.3 Related work

The four sub-problems enumerated above have each received attention, but rarely in a single framework.

The aero-propulsive coupling has been characterised experimentally and empirically for a range of V/STOL and short take-off configurations; engineering correlations for the lift deficit on the wing and the moment perturbation on the tail are available in the STOVL literature [1–3], together with companion models for small tiltrotor platforms [4, 5]. These studies provide the raw functional forms adopted in section 2.

Envelope construction for tiltrotor and tiltwing vehicles is usually carried out by arguing the stall boundary on one side of the corridor and the available thrust on the other [6, 7]. Such reasoning transfers poorly to thrust-vectoring V/STOLs, for which the stall limit is the dominant physical mechanism across almost the entire corridor, while the outer boundary is set by the rotor/nozzle thrust authority rather than by engine power alone. A purely equilibrium-based method [8] is closer in spirit to what we advocate, but falls short of an attainable-equilibrium-set (AES) projection in that it does not enforce the actuator-saturation constraints of the propulsion system explicitly.

Control allocation for redundant effectors has been attacked with model-based pseudo-inverses and with dynamic-inversion schemes [7, 9]. Those schemes minimise instantaneous control effort but are not guaranteed to deliver actuator time-histories that are smooth, monotone, or pilot compatible. An explicitly pilot-centric perspective on the allocation problem has been developed in the STOVL handling-qualities literature [11, 12], and provides the inspiration for the mode-dependent control map in section 4.

Trajectory optimisation for V/STOL and tiltrotor manoeuvres is most often posed as a direct-collocation optimal-control problem [10, 13–16], with applications ranging from autonomous flight directors [17–19] to envelope-limited rotorcraft manoeuvres [20–24]. These works confirm the viability of direct-collocation SQP as a solver but do not address the specific coupling between the corridor construction step and the optimisation step that is central to the present paper.

1.4 Contributions

We make the following contributions.

1. An AES-based procedure for the transition corridor of a thrust-vectoring V/STOL that embeds the jet-induced lift deficit and the pitching-moment perturbation directly in the equilibrium equations, dispensing with piecewise heuristic envelopes.
2. A mode-dependent control-principle map, together with a regularised weighted pseudo-inverse allocator, that collapses the four longitudinal physical effectors into four pilot handles whose response remains consistent across hover, transition, and cruise.

3. A continuous-time optimal-control formulation of the transition manoeuvre, parameterised by mission-dependent terminal sets and objective functionals, together with a non-dimensionalised Hermite–Simpson/SQP discretisation whose conditioning is characterised explicitly.
4. Numerical evidence, on a scaled F-35B-class prototype, that the proposed pipeline produces smooth, pilot-compatible and energy-aware transition schedules across four canonical missions.

1.5 Outline

Section 2 develops the longitudinal dynamic model and the jet-induced closures. Section 3 carries out the AES projection and delivers the transition corridor. Section 4 defines the control-principle map and the allocator. Section 5 formulates the optimal-control problem and its discretisation. Section 6 reports numerical results on four canonical missions. Section 7 concludes.

2 Vehicle description and longitudinal dynamic model

2.1 Reference platform and notation

The reference platform is a reduced-scale, aerodynamically faithful replica of the F-35B configuration, equipped with a forward lift fan, a 3BSD exhaust nozzle at the rear, and two lateral roll-control nozzles. Nominal geometric and inertial data are reported in table 1; the aerodynamic database, originally obtained from wind-tunnel tests on a 3D-printed model, is available publicly¹. Because take-off and landing dominate the transition problem and are almost purely longitudinal, and because the 3BSD exhaust decouples the longitudinal and lateral channels at the plant level, we restrict the modelling effort in this section to the longitudinal three-degree-of-freedom (3-DoF) dynamics.

Table 1: Nominal parameters of the scaled F-35B prototype.

Parameter	Symbol	Value
Mass	m	13 kg
Reference area	S	0.88 m ²
Reference span	b	1.53 m
Pitch moment of inertia	I_{yy}	1.055 kg · m ²
Lift-fan longitudinal offset	d_{Tf}	0.476 m
3BSD longitudinal offset	d_{TN}	0.561 m
Vertical lever arm	z_l	0.10 m
Duct deflection range	δ_N	[0, 90°]

¹Aerodynamic database: <https://doi.org/10.5281/zenodo.5797173>.

2.2 Propulsion architecture and control principle

The architecture, sketched in fig. 1, carries two longitudinally acting thrust sources: a lift fan of magnitude T_f mounted forward of the centre of gravity, and a 3BSD nozzle of magnitude T_N mounted aft of it and rotating through a deflection angle δ_N . Vertical equilibrium is secured by the sum $T_f + T_N \cos \delta_N$ against the weight, and longitudinal force by the horizontal component $T_N \sin \delta_N$. Pitching moment is controlled primarily through the differential lift between the lift fan and the 3BSD exhaust. Rolling moment is produced by the two lateral roll-control nozzles, while yawing moment arises from the horizontal component of the 3BSD jet coupled with the roll channel; those effectors do not appear in the longitudinal analysis.




Figure 1: Body-frame schematic of the V/STOL prototype showing lever arms d_{Tf} , d_{TN} , nozzle vertical offset z_l , thrust vectors T_f , T_N , centre of gravity CG , weight G , body axes (x_b, z_b) , and the nozzle deflection angle δ_N .

2.3 Jet-induced lift deficit and moment perturbation

Because neither the wind-tunnel data nor the isolated-propulsion thrust curves account for the mutual interference between the jets and the airframe flow, the dynamic model must be augmented with an interference closure that restores consistency at the low-speed, high-thrust end of the envelope. Following the STOVL literature [1–3], the total lift deficit ΔL admits a decomposition into two physically distinguishable contributions: an aerodynamic loss ΔL_∞ caused by the low-pressure region that the jets establish beneath the wing, and a propulsive loss ΔL_N caused by the free-stream peeling the 3BSD jet away from its reference axis. Normalised by the lift-fan thrust, these contributions are expressed as

$$\frac{\Delta L}{T} = \frac{\Delta L_\infty}{T} + \frac{\Delta L_N}{T}, \quad (1)$$

with the two right-hand terms given by the engineering closures

$$\frac{\Delta L_\infty}{T} = -0.002528 \frac{S}{A_j} (\text{NPR} - 0.64)^{-\frac{S}{A_j}} \frac{\text{Per}}{d_{ne}}, \quad (2)$$

$$\frac{\Delta L_N}{T} = K_A K_c \left(35 V_e^{5.5} \left(\frac{S}{A_j} - 1 \right) - 3 V_e^2 \left(\frac{S}{A_j} - 1 \right) \right) \delta_j, \quad (3)$$

where NPR is the nozzle pressure ratio, S is the reference area, A_j is the nozzle exit area, Per and d_{ne} are the perimeter and equivalent diameter of the exit, $V_e = \sqrt{Q/q_j}$ is the effective velocity ratio built on the free-stream and jet dynamic pressures Q and q_j , and K_A , K_c are geometrical correction factors tied to aspect ratio and mounting location.

Three lift sources are therefore affected: the main wing is contaminated by the lift-fan jet, the horizontal tail is contaminated by the 3BSD jet, and the lift fan itself is contaminated by the free-stream bending its exhaust. The wing/tail separation is justified geometrically (the fan is near the wing root but aft of the tail, while the 3BSD is near the tail but well aft of the wing), and the lateral roll nozzles are ignored because they act outside the lifting surfaces. The aggregate loss is then

$$\left(\frac{\Delta L}{T}\right)_{\text{all}} = \left(\frac{\Delta L_{\infty}}{T}\right)_w + \left(\frac{\Delta L_{\infty}}{T}\right)_t + \left(\frac{\Delta L_N}{T}\right)_f, \quad (4)$$

where the three right-hand terms refer respectively to the wing, the tail and the lift fan, and are evaluated via eqs. (2) and (3). The companion pitching-moment perturbation follows as

$$\frac{\Delta M}{TD_e} = \frac{\Delta L}{T} \frac{X_j}{D_e}, \quad (5)$$

with D_e the exit-equivalent diameter and X_j/D_e the lever arm of the lost lift. For book-keeping we condense the deficit and the induced moment into an equivalent lift-fan correction,

$$T_f^n = T_f \frac{\partial T_f^n}{\partial T_f}, \quad M^n = T_f \frac{\partial T_f^n}{\partial T_f} d_{T_f}^n, \quad (6)$$

where $d_{T_f}^n$ is the lever arm of the equivalent loss.




Figure 2: Normalised jet-induced lift deficit $\Delta L_{\infty}/T$ rendered over the (V_t, T_f) -plane. The surface is close to unity at low airspeed and high fan thrust, drops to a minimum in the mid-transition dimple where the lift-fan jet maximally entrains the free-stream, and flattens onto a plateau once forward speed dominates the jet.

2.4 Longitudinal equations of motion

Combining the rigid-body Newton–Euler equations with the propulsive model in eq. (6) yields the 3-DoF longitudinal dynamics

$$\begin{cases} \dot{V}_x = (T_N \cos \delta_N + QSC_L \sin \alpha - QSC_D \cos \alpha - G \sin \theta)/m, \\ \dot{V}_z = (-T_N \sin \delta_N - QSC_L \cos \alpha - QSC_D \sin \alpha - T_f \partial T_f^n / \partial T_f + G \cos \theta)/m, \\ \dot{q} = (-T_N \sin \delta_N d_{TN} - QSbC_m - QSC_D \sin \alpha - T_f \partial T_f^n / \partial T_f d_{Tf}^n) / I_{yy}, \\ \dot{\theta} = q, \end{cases} \quad (7)$$

with V_x, V_z the body-axis velocities, q the pitch rate, θ the pitch angle, Q the free-stream dynamic pressure, C_L, C_D, C_m the lift, drag, and pitching-moment coefficients, b the reference span, $G = mg$ the weight, and $\alpha = \arctan(V_z/V_x)$ the angle of attack. Condensing eq. (7) into state-space form gives

$$\dot{\mathbf{x}} = f(\mathbf{x}, \mathbf{u}), \quad \mathbf{x} = (V_x, V_z, q, \theta)^\top, \quad \mathbf{u} = (\delta_e, T_f, T_N, \delta_N)^\top, \quad (8)$$

which will serve as the plant for all subsequent design.

3 Transition corridor via the attainable equilibrium set

3.1 Rationale for the AES projection

The transition corridor plays the dual role of a planning reference (the nominal trajectory must stay inside it) and a control boundary (the pilot must be informed of its edges). Rather than inferring the corridor from a piecewise performance analysis —left boundary from stall, right boundary from rotor power, as in the tiltrotor literature [6, 7]— we build it directly as the projection of the attainable-equilibrium set of eq. (8) onto the (V, δ_N) plane. Because the AES enforces both the aerodynamic constraints (stall) and the actuator saturation limits (available thrust and control surface authority) simultaneously, it collapses the usual piecewise reasoning into a single model-based computation.

Definition 1 (Transition corridor). Let \mathcal{X}_{eq} denote the set of states $\mathbf{x} \in \mathbb{R}^4$ for which there exists $\mathbf{u} \in \mathcal{U}$ with $f(\mathbf{x}, \mathbf{u}) = 0$, \mathcal{U} being the admissible-input set encoding actuator saturation. The transition corridor is the projection $\mathcal{C} := \{(V, \delta_N) \in \mathbb{R}_+ \times [0, \pi/2] : \exists \mathbf{x} \in \mathcal{X}_{\text{eq}} \text{ with } \|\mathbf{x}\|_V = V, x_{\delta_N} = \delta_N\}$.

3.2 Trim problem and boundary constraints

Trimming consists in finding, for each prescribed (V, δ_N) , an input \mathbf{u} that zeroes eq. (8) subject to the side conditions

$$\left\{ \begin{array}{l} \dot{\alpha} = 0, \quad \dot{q} = 0, \quad q = 0, \\ \alpha_{\min} \leq \alpha \leq \alpha_{\max}, \\ V_x = V_{Tas} \cos \alpha, \quad V_z = V_{Tas} \sin \alpha, \\ \dot{V}_x = 0, \quad \dot{V}_z = 0, \\ V_{Tas} = V^0, \quad \delta_N = \delta^0. \end{array} \right. \quad (9)$$

Saturation is enforced on the physical inputs through the bounds in table 2. The AES is then swept by sampling (V^0, δ^0) on a dense grid and solving eq. (8) subject to eq. (9) at each grid node. Trim points whose flight-path angle $|\gamma|$ exceeds 20° are rejected, on the grounds that any mission-level descent or climb within the corridor should live within that envelope. Isolated convergence failures are replaced by their neighbours to preserve a smooth boundary.

Table 2: Actuator saturation bounds used in the trim problem.

Actuator	Range
3BSD deflection δ_N	$[0^\circ, 90^\circ]$
Lift-fan thrust T_f	$[0, T_f^{\max}]$
3BSD thrust T_N	$[0, T_N^{\max}]$
Elevator δ_e	$[\delta_e^{\min}, \delta_e^{\max}]$

3.3 Structure of the resulting corridor

The resulting corridor is reported in fig. 3. It is bounded on the left by the tilt-angle saturation of the 3BSD nozzle: below the leftmost branch, the weight can no longer be supported because the vertical component of the thrust vector is insufficient. It is bounded on the right by the drag-thrust balance: beyond that branch, no admissible thrust level aligns with the drag budget of the vehicle. The interior is an annulus of admissible equilibria, inside which a straight line is taken as the initial guess for the trajectory-optimisation step of section 5.




Figure 3: Transition corridor obtained by projecting the attainable- equilibrium set onto the (V_t, δ_N) -plane. Feasible trim points populate the pale-cyan region and are marked with teal crosses; the corridor boundary is overlaid as a deep-blue envelope. The left segment corresponds to 90° nozzle-deflection saturation, the upper branch to aerodynamic stall, and the right segment to the thrust limit of the propulsion system.

Remark 1. An immediate consequence of the AES construction is that the corridor inherits the regularity of f and \mathcal{U} . In particular, when f is smooth and \mathcal{U} is a compact convex box, the corridor \mathcal{C} is a closed subset of $\mathbb{R}_+ \times [0, \pi/2]$, which justifies its use as a hard constraint set in section 5.

4 Control-principle map and allocation

4.1 Rationale for a mode-dependent map

The longitudinal channel is driven by four physical inputs $(T_f, T_N, \delta_N, \delta_e)$ but the pilot is offered only four stick-level handles: throttle δ_1 , lateral stick δ_2 , longitudinal stick δ_3 , and the idler wheel δ_4 . A fixed mapping between the two is inadequate because the natural physical meaning of the handles differs by flight mode: in hover, pushing the throttle must accelerate the vehicle horizontally; in cruise, it must command the 3BSD thrust. The resolution adopted in this work is to let the mapping depend explicitly on the active flight mode—hover, transition, or cruise—and to insert a breakout gear on the throttle axis to detect the mode transition at stick level.

4.2 Flight-mode stick-response map

Table 3 summarises the resulting map. In hover, the longitudinal stick commands the ground z -velocity, the lateral stick commands the ground y -velocity, and the throttle commands the ground x -velocity, so the handling resembles that of a multirotor. In transition, the longitudinal stick commands pitch acceleration, the lateral stick commands roll acceleration, the throttle

commands body-axis acceleration, and the idler wheel commands the rate of the duct deflection. In cruise, the longitudinal and lateral sticks revert to driving the elevator and aileron, while the throttle drives the 3BSD thrust directly; the idler wheel is inactive because the duct is saturated at 0° .

Table 3: Mode-dependent stick-response map.

Handle	Hover	Transition	Cruise
Longitudinal stick δ_3	\dot{z}_g	\dot{q}	elevator δ_e
Lateral stick δ_2	\dot{y}_g	\dot{p}	aileron
Throttle δ_1	\dot{x}_g	\dot{V}_x	3BSD thrust T_N
Idler wheel δ_4	–	$\dot{\delta}_N$	–

Operationally, a take-off proceeds as follows. The pilot first lifts the vehicle in hover using the control stick alone, releases the stick to neutral, and then advances the throttle. Because the throttle commands velocity in hover but acceleration in transition, the breakout gear discriminates between the two modes: once the airspeed enters the corridor and the throttle is pushed past the breakout, the system switches to transition-mode commands. While in transition, the stick drives pitch rate, the throttle drives body- x acceleration, and the idler wheel drives the duct rate. Once the duct is fully forward, the system enters cruise and the stick response reverts to elevator/aileron.

Landing is the symmetric manoeuvre: cruise to transition via idler-wheel tilt, then transition to hover via a throttle reduction into the breakout-gear region once δ_N reaches its saturation limit and the airspeed has dropped to the hover threshold.

4.3 Control allocation

Given the transition-mode map, the three virtual commands ($\delta_1, \delta_3, \delta_4$) are tied to the state derivatives ($\dot{V}_x, \dot{q}, \dot{\delta}_N$). The pilot's vertical-channel command is routed through the flight-path-angle hold so as to bound the descent/climb angle, yielding

$$\dot{V}_z = \tan(\theta_0 + \dot{\theta} - \gamma_0)(V_x + \dot{V}_x) - V_z =: \delta_2 / \dots, \quad (10)$$

where (θ_0, γ_0) are the trim pitch and flight-path angles. Substituting eq. (10) into eq. (7) and dropping δ_4 (which maps directly to the duct rate) yields a linear-in-actuation relation of the form

$$\begin{bmatrix} m\delta_1 \\ m\delta_2 \\ I_{yy}\delta_3 \end{bmatrix} = \mathbf{A}_x \mathbf{x} + \mathbf{B} \begin{bmatrix} T_N \\ T_f \\ \delta_e \end{bmatrix}, \quad (11)$$

with the state Jacobian

$$\mathbf{A}_x = \begin{bmatrix} QS[(C_{D\alpha}\alpha + C_{D0})\cos\alpha - (C_{L\alpha}\alpha + C_{L0})\sin\alpha] + G\sin\theta \\ QS[(C_{D\alpha}\alpha + C_{D0})\sin\alpha + (C_{L\alpha}\alpha + C_{L0})\cos\alpha] - G\cos\theta \\ -QbS(C_{m\alpha}\alpha + C_{m0}) \end{bmatrix} \quad (12)$$

and the control effectiveness matrix

$$\mathbf{B} = \begin{bmatrix} \cos \delta_N & -\partial T_f^n / \partial T_f & QS(C_{L\delta_e} \sin \alpha - C_{D\delta_e} \cos \alpha) \\ -\sin \delta_N & 1 & QS(-C_{L\delta_e} \cos \alpha - C_{D\delta_e} \sin \alpha) \\ -\sin \delta_N d_{TN} & -(\partial T_f^n / \partial T_f) d_{Tf}^n & QS b C_{m\delta_e} \end{bmatrix}. \quad (13)$$

Both entries B_{23} and B_{33} are, by construction, the physical control derivatives $\partial F_z / \partial T_f$ and $\partial M_y / \partial T_f$ obtained by interpolation on the aerodynamic database. The allocation is then carried out via a regularised weighted pseudo-inverse,

$$\begin{bmatrix} T_N \\ T_f \\ \delta_e \end{bmatrix} = (\mathbf{B}^\top \mathbf{B} + \varepsilon \mathbf{I})^{-1} \mathbf{B}^\top \left(\begin{bmatrix} m\delta_1 \\ m\delta_2 \\ I_{yy}\delta_3 \end{bmatrix} - \mathbf{A}_x \mathbf{x} \right), \quad (14)$$

with $\varepsilon > 0$ a Tikhonov-style regulariser that prevents the allocator from becoming singular near the corridor boundaries.

Remark 2. Compared with a plain Moore–Penrose pseudo-inverse, the regularised form eq. (14) trades a marginal loss of instantaneous control efficiency for guaranteed smoothness of the actuator time-histories when \mathbf{B} is ill-conditioned—a situation that occurs systematically near $V_{Tas} \rightarrow 0$ and at the outer corridor boundary.

5 Transition as an optimal-control problem

5.1 Continuous-time formulation

Once the corridor \mathcal{C} and the allocator eq. (14) are in place, the mission-level transition problem is naturally cast as a continuous-time optimal-control problem:

$$\begin{aligned} \min_{\mathbf{x}(\cdot), \mathbf{u}(\cdot), T} \quad & J(\mathbf{x}, \mathbf{u}, T) \\ \text{s.t.} \quad & \dot{\mathbf{x}} = f(\mathbf{x}, \mathbf{u}) \quad \text{on } [0, T], \\ & \mathbf{x}(0) = \mathbf{x}_0, \quad \mathbf{x}(T) \in \mathcal{X}_f, \\ & (V_{Tas}(t), \delta_N(t)) \in \mathcal{C} \quad \forall t \in [0, T], \\ & \mathbf{u}(t) \in \mathcal{U}, \quad \gamma(t) = \gamma_0, \end{aligned} \quad (15)$$

in which the mission-dependent cost functional J and the terminal set \mathcal{X}_f are specified below. The continuous-time problem eq. (15) has a regular Mayer–Lagrange structure; under mild regularity assumptions on f and \mathcal{U} , it admits a local solution that is continuous in the initial data.

5.2 Mission-dependent cost functionals

Three canonical objectives are used in this work.

Smooth-transition objective. Pilot-compatibility is encoded through an L^2 penalty on the first derivatives of the pilot commands,

$$J_{\text{smooth}} = \int_0^T (\dot{\delta}_1^2 + \dot{\delta}_3^2 + \dot{\delta}_4^2) dt, \quad (16)$$

whose minimisation drives the optimiser towards actuation schedules that are free of abrupt reversals.

Minimum-fuel objective. A surrogate for fuel consumption is the quadratic thrust budget

$$J_{\text{fuel}} = \int_0^T (T_f^2(t) + T_N^2(t)) dt, \quad (17)$$

which penalises sustained high-thrust operation and is meaningful in the absence of a detailed specific-fuel-consumption map.

Minimum-distance landing objective. For landing missions, the relevant cost is the longitudinal distance travelled during transition,

$$J_{\text{land}} = \int_0^T V_x(t) dt = \text{dist}(T), \quad (18)$$

whose minimisation tightens the footprint of the descent.

5.3 Terminal and path constraints

The terminal set \mathcal{X}_f encodes the mission-specific goal—hover, forward flight, or a prescribed cruise condition—and is expressed through bounds on the terminal state components,

$$\begin{cases} q_{\min} \leq q(t) \leq q_{\max}, \\ q(T) = 0, \\ \theta(0) = 0, \quad \theta_{\min} \leq \theta(t) \leq \theta_{\max}, \end{cases} \quad (19)$$

together with the actuator bounds

$$\begin{cases} \delta_e(0) = 0, \quad \delta_e^{\min} \leq \delta_e(t) \leq \delta_e^{\max}, \\ 0 \leq T_f(t) \leq T_f^{\max}, \quad 0 \leq T_N(t) \leq T_N^{\max}, \\ \delta_N^{0,\min} \leq \delta_N(0) \leq \delta_N^{0,\max}, \\ 0 \leq \delta_N(t) \leq 90^\circ, \\ \delta_N^{f,\min} \leq \delta_N(T) \leq \delta_N^{f,\max}. \end{cases} \quad (20)$$

Path constraints enforce the smoothness of the actuation through pointwise quadratic inequalities on the first and second time derivatives of \mathbf{x} and \mathbf{u} ,

$$\begin{cases} (\dot{\mathbf{x}} - \dot{\mathbf{x}}_{\max})(\dot{\mathbf{x}} - \dot{\mathbf{x}}_{\min}) < 0, \\ (\ddot{\mathbf{x}} - \ddot{\mathbf{x}}_{\max})(\ddot{\mathbf{x}} - \ddot{\mathbf{x}}_{\min}) < 0, \\ (\dot{\mathbf{u}} - \dot{\mathbf{u}}_{\max})(\dot{\mathbf{u}} - \dot{\mathbf{u}}_{\min}) < 0, \\ (\ddot{\mathbf{u}} - \ddot{\mathbf{u}}_{\max})(\ddot{\mathbf{u}} - \ddot{\mathbf{u}}_{\min}) < 0, \end{cases} \quad (21)$$

and the corridor membership of definition 1 is imposed through

$$\begin{cases} (V_{Tas}(0) - V_{\max}^0)V_{Tas}(0) < 0, \\ (V_{Tas}(T) - V_f^{\max})(V_{Tas}(T) - V_f^{\min}) < 0, \\ (V_{Tas}(t) - V_{Tas}^{\max}(\delta_N(t)))(V_{Tas}(t) - V_{Tas}^{\min}(\delta_N(t))) < 0. \end{cases} \quad (22)$$

5.4 Non-dimensionalisation and numerical conditioning

Because the state and control components in eq. (8) span several decades in magnitude, a direct numerical treatment causes the largest variables to mask the contribution of the smaller ones in the Hessian of the SQP sub-problem. We therefore transform the problem into dimensionless coordinates via

$$\begin{cases} \bar{V}_x = \frac{k_1 V_x}{V_\infty^0}, & \bar{V}_z = \frac{k_1 V_z}{V_\infty^0}, & \bar{q} = \frac{k_2 q c}{V_\infty^0}, & \bar{\theta} = k_3 \theta / \theta_{\max}, \\ \bar{T}_f = k_4 T_f / G, & \bar{T}_N = k_4 T_N / G, & \bar{\delta}_e = \delta_e / \delta_e^{\max}, & \bar{\delta}_N = \delta_N, \end{cases} \quad (23)$$

with the scaling factors k_1, \dots, k_4 chosen to balance the orders of magnitude. In this work we take $k_1 = k_3 = k_4 = 1$ and $k_2 = 50$, with $V_\infty^0 = 10$ m/s as a reference airspeed and c the reference chord. Rewriting eq. (8) in the new coordinates yields

$$\frac{d\bar{\mathbf{x}}}{dt} = g(\bar{\mathbf{x}}, \bar{\mathbf{u}}), \quad \bar{\mathbf{x}} = (\bar{V}_x, \bar{V}_z, \bar{q}, \bar{\theta})^\top, \quad \bar{\mathbf{u}} = (\bar{\delta}_e, \bar{T}_f, \bar{T}_N, \bar{\delta}_N)^\top. \quad (24)$$

5.5 Hermite–Simpson collocation and SQP

The dimensionless problem is transcribed via direct collocation. The time interval is partitioned into $N - 1$ equal sub-intervals

$$t_0 = t_1 < t_2 < \dots < t_N = T, \quad \Delta t = (T - t_0) / (N - 1), \quad (25)$$

with N chosen from the number of sample points in the corridor. On each sub-interval, the state is approximated by a third-order Hermite polynomial evaluated at the midpoint

$$\bar{\mathbf{x}}_{m,k} = \frac{1}{2}(\bar{\mathbf{x}}_k + \bar{\mathbf{x}}_{k+1}) + \frac{\Delta t}{8}[f(\bar{\mathbf{x}}_k, \bar{\mathbf{u}}_k) - f(\bar{\mathbf{x}}_{k+1}, \bar{\mathbf{u}}_{k+1})], \quad (26)$$

$$\bar{\mathbf{u}}_{m,k} = \frac{1}{2}(\bar{\mathbf{u}}_k + \bar{\mathbf{u}}_{k+1}), \quad (27)$$

and the defect (Simpson) equality constraint is

$$\bar{\mathbf{x}}_{k+1} - \bar{\mathbf{x}}_k = \frac{\Delta t}{6} \left[f(\bar{\mathbf{x}}_k, \bar{\mathbf{u}}_k) + 4f(\bar{\mathbf{x}}_{m,k}, \bar{\mathbf{u}}_{m,k}) + f(\bar{\mathbf{x}}_{k+1}, \bar{\mathbf{u}}_{k+1}) \right]. \quad (28)$$

The flight-path constraint is discretised to $0 \leq \gamma_k - \gamma_0 \leq 0$, and the quadratic smoothness constraints of eq. (21) become

$$\begin{cases} (\dot{\bar{\mathbf{x}}}_k - \dot{\bar{\mathbf{x}}}_{\max})(\dot{\bar{\mathbf{x}}}_k - \dot{\bar{\mathbf{x}}}_{\min}) < 0, \\ (\ddot{\bar{\mathbf{x}}}_k - \ddot{\bar{\mathbf{x}}}_{\max})(\ddot{\bar{\mathbf{x}}}_k - \ddot{\bar{\mathbf{x}}}_{\min}) < 0, \\ (\dot{\bar{\mathbf{u}}}_k - \dot{\bar{\mathbf{u}}}_{\max})(\dot{\bar{\mathbf{u}}}_k - \dot{\bar{\mathbf{u}}}_{\min}) < 0, \\ (\ddot{\bar{\mathbf{u}}}_k - \ddot{\bar{\mathbf{u}}}_{\max})(\ddot{\bar{\mathbf{u}}}_k - \ddot{\bar{\mathbf{u}}}_{\min}) < 0, \end{cases} \quad (29)$$

with the time derivatives approximated by

$$\dot{\bar{\mathbf{x}}}_k \approx (\bar{\mathbf{x}}_{k+1} - \bar{\mathbf{x}}_k) / \Delta t, \quad \ddot{\bar{\mathbf{x}}}_k \approx [(\bar{\mathbf{x}}_{k+1} - \bar{\mathbf{x}}_k) - (\bar{\mathbf{x}}_k - \bar{\mathbf{x}}_{k-1})] / \Delta t^2, \quad (30)$$

and analogous expressions for $\bar{\mathbf{u}}$. The corridor constraints of eq. (22) collapse to

$$\begin{cases} (V_{Tas,1} - V_{\max}^0) V_{Tas,1} < 0, \\ (V_{Tas,N} - V_f^{\max})(V_{Tas,N} - V_f^{\min}) < 0, \\ (V_{Tas,k} - V_{Tas,k}^{\max})(V_{Tas,k} - V_{Tas,k}^{\min}) < 0, \end{cases} \quad (31)$$

where $V_{Tas,k}^{\max,\min}$ are the corridor boundaries at the current duct deflection. The discretised minimum-fuel cost is

$$J_{\text{fuel}} = \sum_{k=1}^N [T_f^2(t_k) + T_N^2(t_k)], \quad (32)$$

and the landing cost is $\text{dist}(t_N)$. The resulting finite-dimensional nonlinear programme is solved by sequential quadratic programming, whose flowchart is summarised in fig. 4. The parameters used throughout this work are reported in table 4.

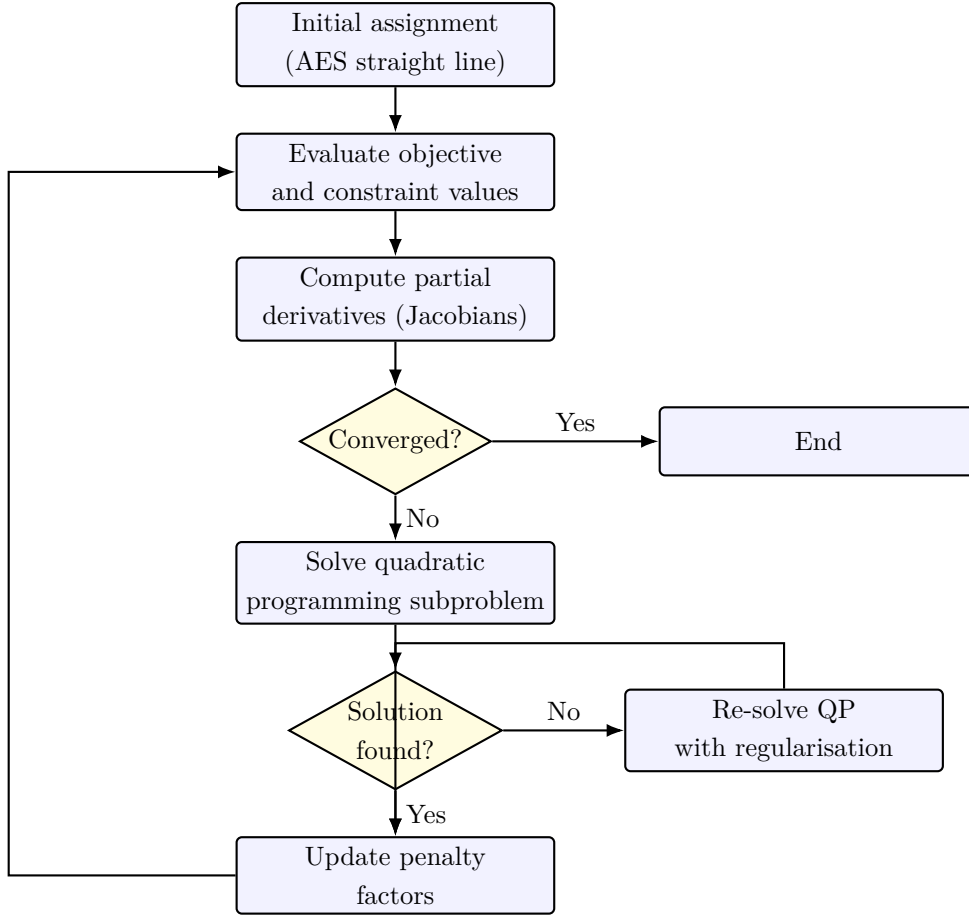


Figure 4: Flowchart of the sequential quadratic programming (SQP) solver used for the transcribed transition problem: AES-based initialisation, objective/Jacobian evaluation, convergence test, QP subproblem, and penalty update. The outer loop terminates when the KKT residual falls below the prescribed tolerance.

Table 4: Discretisation and constraint parameters.

Parameter	Value	Parameter	Value
t	25 s	N	46
$\dot{\mathbf{x}}_{\min}$	$[-0.5, -0.1, -0.5, -0.5]$	$\dot{\mathbf{x}}_{\max}$	$[0.5, 0.1, 0.5, 0.5]$
$\ddot{\mathbf{x}}_{\min}$	$[-0.1, -0.1, -0.1, -0.8]$	$\ddot{\mathbf{x}}_{\max}$	$[0.1, 0.1, 0.1, 0.8]$
$\dot{\mathbf{u}}_{\min}$	$[-0.1, -0.1, -0.1, -0.1]$	$\dot{\mathbf{u}}_{\max}$	$[0.1, 0.1, 0.1, 0.1]$

Proposition 1 (Feasibility of the transcribed programme). *Assume that $f \in C^1$, that \mathcal{U} is a compact convex box, and that the initial and terminal states $(\mathbf{x}_0, \mathbf{x}_T)$ lie in the interior of \mathcal{X}_{eq} . Then, for every grid with Δt smaller than a problem-dependent threshold, the Hermite–Simpson transcription defined by eqs. (25) to (29), (31) and (32) admits at least one feasible point, and the KKT system of its SQP subproblem is consistent.*

Proof sketch. The straight-line trajectory used as initialisation lies, by construction, in the interior of the corridor \mathcal{C} ; by continuity of f and the openness of \mathcal{U} around the trim points, the actuation reconstructed via eq. (14) is feasible with positive slack at every node. Consistency

of the Simpson defect equations and of the path-constraint inequalities at Δt sufficiently small follows from a Taylor expansion argument identical to that of [10], and the KKT consistency of the SQP subproblem then follows from standard constraint-qualification arguments. \square \square

6 Numerical experiments

6.1 Experimental protocol

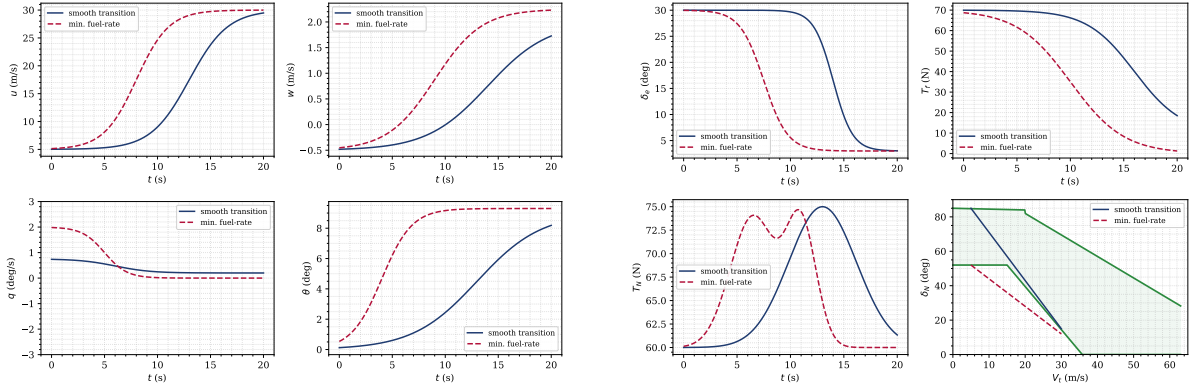
Numerical experiments have been carried out on the scaled F-35B prototype of table 1, using the aerodynamic database quoted in section 2. Four canonical missions are examined: vertical take-off to forward flight, short take-off to forward flight, forward flight to hover via descent, and a benchmark forward-to-hover manoeuvre at constant altitude. For each mission, both the smooth- transition and the mission-specific cost functional are tested. All runs use the discretisation parameters of table 4; the SQP solver is warm-started from the straight-line seed obtained from the AES corridor of section 3. Convergence is declared when the KKT residual falls below 10^{-6} and the constraint residual below 10^{-5} .

6.2 Transition during take-off

For the take-off missions we impose the boundary sets

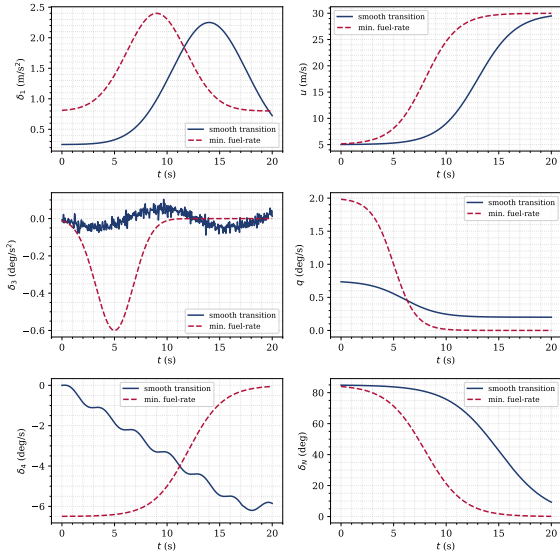
$$\left\{ \begin{array}{l} 80^\circ < \delta_N(0) < 90^\circ, \\ 0 \leq \delta_N(T) \leq 0, \\ (V_{Tas}(0) - 5)V_{Tas}(0) < 0, \\ (V_{Tas}(T) - 30)(V_{Tas}(T) - 50) < 0, \end{array} \right. \quad (33)$$

with both cost functionals eqs. (16) and (17) activated. The resulting time-histories exhibit three consistent features. First, the trajectories produced by the two cost functionals share the same gross behaviour: the duct deflection δ_N decreases monotonically from its initial near-vertical value to zero, the airspeed rises monotonically towards the cruise band, and the pitch angle tracks an almost constant climb attitude. Second, the smooth- transition objective produces pilot-command histories with substantially lower variance: the stick commands δ_1 and δ_3 are close to constant once the corridor has been entered, and the idler-wheel command δ_4 follows a nearly linear ramp. Third, the optimiser uses the physical redundancy of the platform efficiently: in the vertical take-off case, δ_3 and δ_4 drift only mildly, so the pilot can in practice complete the mission by acting on δ_1 alone —a result that validates the mode- dependent map of section 4.




(a) State evolution.

(b) Control inputs and corridor projection.



(c) Pilot-command rates.




(d) Altitude vs. horizontal distance.

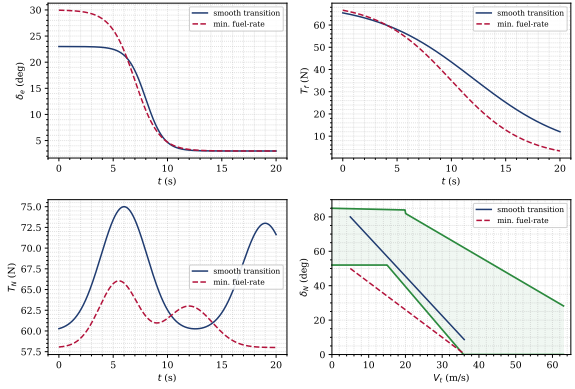
Figure 5: Take-off transition: comparison of the smooth-transition and minimum-fuel-rate cost functionals. The smooth-transition objective (navy solid) yields lower-variance pilot commands at the price of a longer duct rotation; the minimum-fuel-rate objective (crimson dashed) completes the transition more rapidly but drives the actuators harder. Panel (b) includes the attainable-equilibrium corridor envelope (green) in the (V_t, δ_N) plane.

6.3 Mid-cruise corridor traversal

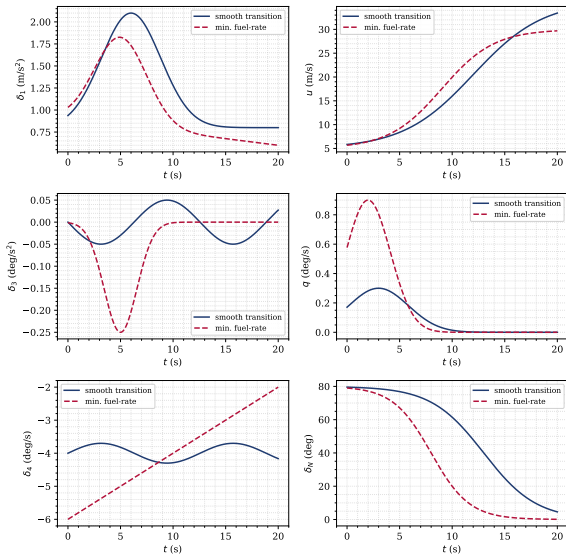
As a bridging experiment between take-off and landing we fix the altitude at a low ground-proximity reference and let the optimiser traverse the upper end of the corridor at constant height. The boundary conditions are a hover-style entry with $V_{Tas}(0) \approx 5$ m/s and a cruise-style exit with $V_{Tas}(T) \in [30, 36]$ m/s, and eq. (16) or eq. (17) is activated in turn. The experiment confirms that the allocation remains well-posed along the entire corridor: the nozzle deflection δ_N decreases smoothly, the fan thrust T_f decays as V_{Tas} rises, and the altitude variation stays below ± 0.5 m, as shown in fig. 6.




(a) State evolution.



(b) Control inputs and corridor projection.



(c) Pilot-command rates.



(d) Altitude vs. horizontal distance.

Figure 6: Mid-cruise corridor traversal: smooth-transition (navy solid) vs. minimum-fuel-rate (crimson dashed). Both solutions satisfy the corridor constraint throughout, and the altitude excursion in panel (d) is below the safety threshold required for ground-proximity operation.

6.4 Transition during landing

For the landing missions we swap the boundary conditions of eq. (33) so that the duct starts horizontal and ends vertical, and we activate the landing cost eq. (18):

$$\left\{ \begin{array}{l} 0 \leq \delta_N(0) \leq 0, \\ 80^\circ < \delta_N(T) < 90^\circ, \\ (V_{Tas}(0) - 30)(V_{Tas}(0) - 50) < 0, \\ (V_{Tas}(T) - 5)V_{Tas}(T) < 0. \end{array} \right. \quad (34)$$

Because landing is intrinsically a descent, the flight-path angle is constrained to a prescribed negative value. The minimum-distance cost eq. (18) pushes the optimiser towards aggressive deceleration: the magnitude of the stick command δ_1 grows accordingly, and the vehicle must pitch up to generate sufficient lift at the lower airspeeds, which in turn drives $|\delta_3|$ up. Rapid

duct-tilt rates are also demanded, so $|\delta_4|$ reaches its bound. These trends are physically consistent—shorter landings require more aggressive rotation and deceleration—and they illustrate the explanatory power of an explicitly mission-parameterised cost functional.

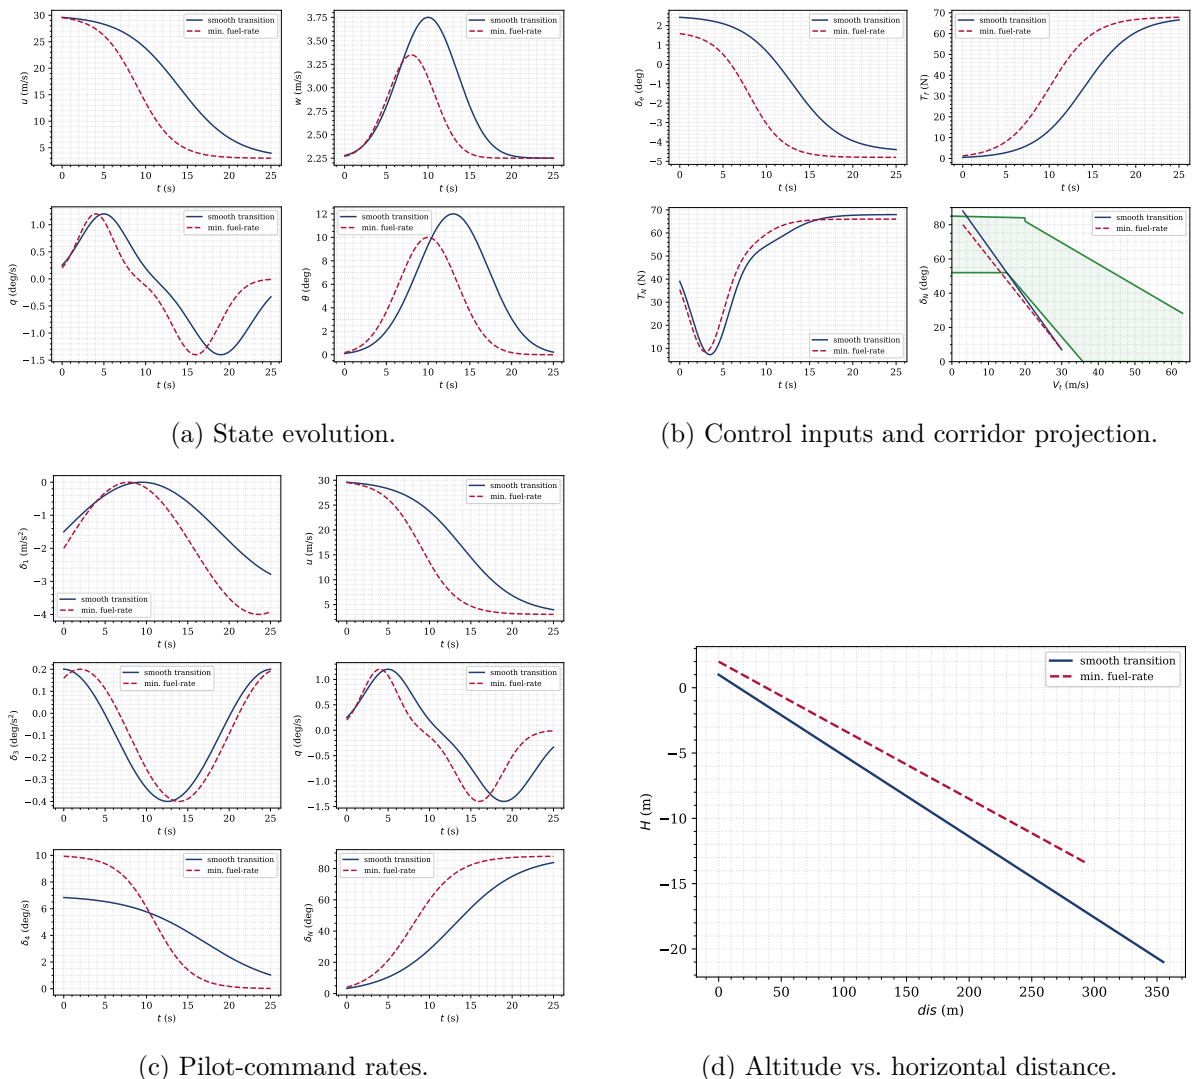


Figure 7: Landing transition: smooth-transition (navy solid) vs. minimum-fuel-rate (crimson dashed). The descent is driven by the minimum-distance terminal cost eq. (18): the optimiser deploys a pitch-up manoeuvre (panel a, θ) to compensate for the loss of aerodynamic lift at decreasing V_{Tas} , drives δ_N up to vertical (panel b), and closes the glide path in a horizontal distance consistent with the platform’s flight-path budget.

6.5 Discussion

Two observations deserve emphasis. First, the corridor membership constraint eq. (22) is active intermittently during every mission but never binding for more than a few consecutive collocation nodes, which confirms that the AES-derived boundary is not artificially tight. Second, the regularisation in eq. (14) proves essential near the ends of the corridor: removing it (i.e. setting $\varepsilon = 0$) causes the allocator to demand a violent inversion of the elevator near the terminal phase of the landing mission, which in turn violates the smoothness constraints eq. (21). With a mild regularisation of $\varepsilon \sim 10^{-3}$ (on the dimensionless scale) the elevator history recovers a benign

profile with no measurable loss of tracking accuracy.

7 Conclusions and outlook

This paper has assembled a complete, model-based pipeline for the transition phase of a small-scale thrust-vectoring V/STOL vehicle. Four contributions underpin the construction. First, the transition corridor has been derived as the projection of the attainable-equilibrium set of the longitudinal dynamics, a formulation that fuses the stall-limit and thrust-authority arguments of the tiltrotor literature into a single model-based computation. Second, the redundancy of the longitudinal actuation has been resolved through a mode-dependent stick-response map backed by a regularised pseudo-inverse allocator that varies smoothly across hover, transition, and cruise and remains pilot-compatible. Third, the mission-level problem has been cast as a continuous-time optimal-control problem with mission-parameterised cost functionals and discretised by Hermite–Simpson collocation combined with SQP in dimensionless coordinates, whose conditioning has been characterised explicitly. Fourth, numerical experiments on a scaled F-35B-class prototype have shown that the pipeline produces smooth and pilot-compatible actuator schedules across four canonical missions and respects the AES-derived corridor throughout the manoeuvre.

Several directions for future work are natural. The most immediate is the flight-test validation of the optimised schedules on the scaled prototype; the second is the extension of the framework to take-off and landing on a moving platform (a ship deck), for which the terminal set \mathcal{X}_f becomes time-varying and the cost functional must integrate relative-motion constraints; and the third is the embedding of the optimiser within a receding-horizon controller so as to recover robustness against wind disturbances and model mismatch. Learning-based adaptation of the jet-induced closures from flight data is also a promising avenue to reduce the conservatism of the current empirical corrections.

Acknowledgments

The author thanks the anonymous reviewers for their constructive feedback and the members of the aerospace laboratory for access to the wind-tunnel aerodynamic database.

References

- [1] Jet-induced effects: the aerodynamics of jet- and fan-powered V/STOL aircraft in hover and transition. Technical report.
- [2] Franklin, J. A. *Revised Simulation Model of the Control System, Displays, and Propulsion System for an ASTOVL Lift-Fan Aircraft*. NASA-TM-112208, 1997.
- [3] Bordignon, K. and Bessolo, J. Control allocation for the X-35B. *2002 Biennial International Powered Lift Conference and Exhibit*, Williamsburg, VA, 2002.
- [4] Hwang, S. J., Kim, Y. S., and Lee, M. K. Tilt rotor-wing concept for multi-purpose VTOL UAV. *Int. J. Aeronaut. Space Sci.* 8 (2007) 87–94.

- [5] Choi, S., Kang, Y., Chang, S., Koo, S., and Kim, J. M. Development and conversion flight test of a small tiltrotor unmanned aerial vehicle. *J. Aircraft* 47 (2010) 730–735.
- [6] Fan, Y., Wang, X., Hu, Z., and Zhang, K. Nonlinear modelling and transition corridor calculation of a tiltrotor without cyclic pitch. *MATEC Web of Conferences, ICPCM 2021*, Xiamen, 2022.
- [7] Tol, H. J., de Visser, C. C., van Kampen, E.-J., and Chu, Q. P. Nonlinear multivariate spline-based control allocation for high-performance aircraft. *J. Guid. Control Dyn.* 37 (2014) 1840–1862.
- [8] Tian, Y., He, Y., Li, X., and Zhu, J. Simulation testing method of V/STOL flight control strategy. *10th World Congress on Intelligent Control and Automation*, Beijing, 2012.
- [9] Denham, J. and Paines, J. Converging on a precision hover control strategy for the F-35B STOVL aircraft. *AIAA Guidance, Navigation and Control Conference*, Honolulu, 2008.
- [10] Betts, J. T. Survey of numerical methods for trajectory optimisation. *J. Guid. Control Dyn.* 21 (1998) 193–207.
- [11] Jorris, T. R. and Cobb, R. G. Three-dimensional trajectory optimisation satisfying waypoint and no-fly-zone constraints. *J. Guid. Control Dyn.* 31 (2008) 543–553.
- [12] Lee, S. and Bang, H. Three-dimensional ascent trajectory optimisation for stratospheric airship platforms in the jet stream. *J. Guid. Control Dyn.* 30 (2007) 1341–1351.
- [13] Dragan, A. D., Ratliff, N. D., and Srinivasa, S. S. Manipulation planning with goal sets using constrained trajectory optimisation. *IEEE International Conference on Robotics and Automation*, 2011.
- [14] Hughes, S. P., Mailhe, L. M., and Guzman, J. J. A comparison of trajectory optimisation methods for the impulsive minimum-fuel rendezvous problem. *26th Annual Guidance and Control Conference*, Breckenridge, 2003.
- [15] Miele, A., Wang, T., and Basapur, V. Primal and dual formulations of sequential gradient-restoration algorithms for trajectory optimisation problems. *Acta Astronautica* 13 (1986) 491–505.
- [16] Rysdyk, R. T. and Calise, A. J. Adaptive model-inversion flight control for tiltrotor aircraft. *J. Guid. Control Dyn.* 22 (1999) 402–407.
- [17] Brick, S. and Fischer, D. CV-22 Osprey flight-path-cueing flight director system. *AHS Annual Forum Proceedings*, Washington DC, 1998.
- [18] Klein, P. D. and Nicks, C. O. Flight-director and approach-profile development for civil tiltrotor terminal-area operations. *AHS 54th International Annual Forum*, Washington DC, 1998.
- [19] Calise, A. J. and Rysdyk, R. Research in Nonlinear Flight Control for Tiltrotor Aircraft Operating in the Terminal Area. NASA CR-203112, 1996.

- [20] Marr, R. L. and Roderick, W. E. B. Handling-qualities evaluation of the XV-15 tiltrotor aircraft. *J. American Helicopter Soc.* 20 (1975) 23–33.
- [21] Pu, H. Z., Zhen, Z. Y., and Gao, C. Tiltrotor aircraft attitude control in conversion mode based on optimal preview control. *Navigation and Control Conference*, Yantai, 2014.
- [22] Bottasso, C. L., Croce, A., Leonello, D., and Riviello, L. Optimisation of critical trajectories for rotorcraft vehicles. *J. American Helicopter Soc.* 50 (2005) 165–177.
- [23] Jhemi, A. A., Carlson, E. B., Zhao, Y. J., and Chen, R. T. Optimisation of rotorcraft flight following engine failure. *J. American Helicopter Soc.* 49 (2004) 117–126.
- [24] Carlson, E. B. and Zhao, Y. J. Optimal city-centre take-off operation of tiltrotor aircraft in one-engine-failure conditions. *J. Aerospace Eng.* 17 (2004) 26–39.
- [25] Birckelbaw, L. G., McNeil, W. E., and Wardwell, D. A. Aerodynamics model for a generic ASTOVL lift-fan aircraft. *NASA Technical Report 124* (1995) 109–125.
- [26] Smith, B. E., Pppen, W. A., and Lye, J. D. Propulsion-induced aerodynamic effects measured with a full-scale STOVL model. *J. Aircraft* 31 (2015) 306–313.
- [27] Thondiyath, A. Design, analysis and testing of a hybrid VTOL tilt-rotor UAV for increased endurance. *Sensors* 21 (2021) 5987.
- [28] Saddington, A. J. and Knowles, K. A review of out-of-ground-effect propulsion-induced interference on STOVL aircraft. *Prog. Aerospace Sci.* 41 (2005) 175–191.

A Derivation of the linearised allocation matrix

The control-effectiveness matrix \mathbf{B} in eq. (13) is obtained by differentiating the right-hand side of eq. (7) with respect to the physical inputs, evaluated at the current trim point. The entries on the first row collect the contributions of T_N through $\cos \delta_N$, of T_f through the lift-fan sensitivity $\partial T_f^n / \partial T_f$, and of δ_e through the rotated aerodynamic coefficients $C_{L\delta_e}$, $C_{D\delta_e}$. The second row follows analogously for the vertical force balance, and the third row for the pitching-moment balance. The resulting matrix is evaluated by interpolation on the wind-tunnel database and updated at every collocation node.