Aerodynamic Co-Contraction in Redundant Aerial Actuation

A fiber-based theory of active aerodynamic promptness, passive aerodynamic damping, and their interplay with power consumption

Antonio Franchi

schol@r-franchi.eu
https://r-franchi.eu/

18 August 2026

Website, publications, and contact

One theory in two complementary publications

[F26a] Active promptness, fiber geometry, and the energy trade-off

A. Franchi, "Muscle Coactivation in the Sky: Geometry and Pareto Optimality of Energy vs. Aerodynamic Promptness and Multirotors as Variable Stiffness Actuators," in *Proc. 2026 IEEE International Conference on Unmanned Aircraft Systems (ICUAS)*, Corfu, Greece, June 2026, pp. 235–242. *(Best Conference Paper Award)*

DOI: 10.1109/ICUAS69441.2026.11598607

Free preprint: arXiv:2602.14222

Free preprint

[F26b] Variable aerodynamic damping and passive interaction shaping

A. Franchi, "Variable Aerodynamic Damping Actuation via Co-Contraction: A Structural Analogy with Variable Stiffness Actuation," *IEEE Robotics and Automation Letters*, Early Access, 2026.

DOI: 10.1109/LRA.2026.3723738

Free preprint: arXiv:2605.07292

Free preprint

What does actuation redundancy really buy us?

Beyond robustness and energy efficient allocation,
it provides task-equivalent internal states
with different local behaviors.

Redundancy is independent of full actuation

	non-redundant $\dim \mathscr{V} = \dim \mathscr{W}$	**redundant** $\dim \mathscr{V} > \dim \mathscr{W}$
underactuated $\dim \mathscr{W} < 6$	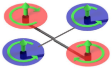 Quadrotor $\dim \mathscr{V} = \dim \mathscr{W} = 4$	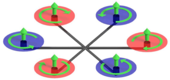 Hexarotor $6 = \dim \mathscr{V} > \dim \mathscr{W} = 4$
fully actuated $\dim \mathscr{W} = 6$	Credit: Michieletto, Ryll & Franchi 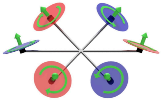 Quadrotor $\dim \mathscr{V} = \dim \mathscr{W} = 6$	Credit: Brescianini & D'Andrea 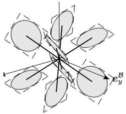 Omnicopter $8 = \dim \mathscr{V} > \dim \mathscr{W} = 6$

Redundancy means $\dim \mathscr{V} > \dim \mathscr{W}$;
it is independent of whether $\dim \mathscr{W} < 6$ or $\dim \mathscr{W} = 6$.

[F26a].

From redundant actuation to task fibers

- internal actuation state: $v \in \mathscr{V}$, $\dim \mathscr{V} = n$;
- task output: $f(v) \in \mathscr{W}$, $\dim \mathscr{W} = m < n$;
- all configurations producing the same task form a level set.

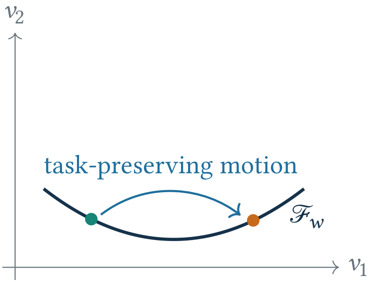

$$\mathscr{F}_w = f^{-1}(w) = \{v \in \mathscr{V} \mid f(v) = w\}$$

[F26a] Secs. II–III.

One task, an entire fiber of actuator states

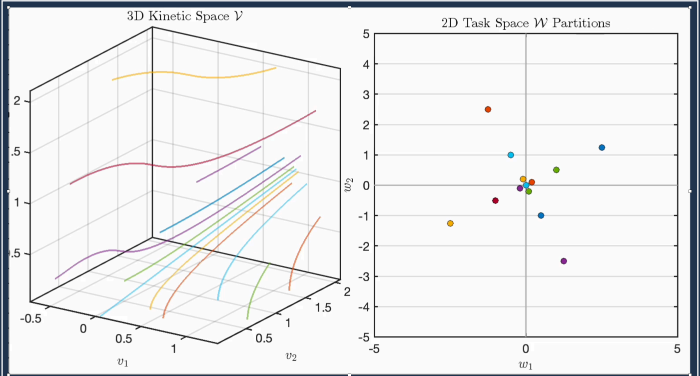

$\dim \mathscr{V} = 3$

a fiber (curve) for each task

$$\mathscr{V} \xrightarrow[f]{} \mathscr{W}$$

$\dim \mathscr{W} = 2$

a point w for each task

[F26a] Secs. II–III.

Same task: can active authority still change?

Fixed current task	Variable task differential
$$f(v) = w$$ All operating points $v \in \mathscr{F}_w$ produce the same current output.	$$\delta w = J_f(v)\,\delta v$$ Different operating points $v \in \mathscr{F}_w$ can have different instantaneous task-rate authority.

The task value is fixed; its active differential need not be.

[F26a] Sec. II.

How should active “readiness” be measured?

Push the admissible actuator-rate ball through the task differential and measure the resulting task-rate volume.

How actuator-rate limits propagate to the task

$$
\begin{tikzcd}[column sep=huge, row sep=huge]
\begin{array}{c} T_v\mathscr{V} \\ \text{actuator rates} \end{array} \arrow[r, "d f_v \ (J_f)"] & \begin{array}{c} T_{f(v)}\mathscr{W} \\ \text{task rates} \end{array} \\
\begin{array}{c} T_v^*\mathscr{V} \\ \text{actuator rate covectors} \end{array} \arrow[u, "g_v^{-1} \ (G_v^{-1})"] & \begin{array}{c} T_{f(v)}^*\mathscr{W} \\ \text{task rate covectors} \end{array} \arrow[l, "d f_v^* \ (J_f^\top)"] \arrow[u, swap, "\begin{array}{l} d_v = d f_v \circ g_v^{-1} \circ d f_v^* \\ D(v) = J_f G_v^{-1} J_f^\top \end{array}"]
\end{tikzcd}
$$

Push-forward construction

The chosen metric g_v encodes actuator-rate capability;
d_v^{-1} is the induced metric on task-rate space.

[F26a] Sec. II.

Fiber density as active promptness

Let $G_v \succ 0$ encode normalized actuator-rate capability. Then

$$\delta v^\top G_v \delta v \leq 1$$

is mapped by $J_f(v)$ to a task-rate ellipsoid with co-metric

$$D(v) = J_f(v) G_v^{-1} J_f(v)^\top.$$

At regular internal states, $J_f(v)$ has full row rank and $D(v) \succ 0$.

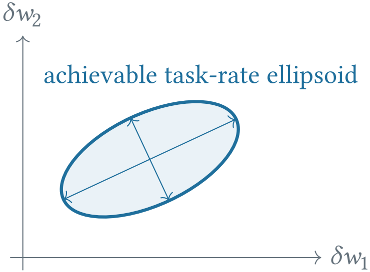

$$\rho(v) = \sqrt{\det\left(J_f(v) G_v^{-1} J_f(v)^\top\right)}$$

ρ = fiber-density volume factor = a dynamic-manipulability measure = active promptness

[F26a] Eqs. (2)–(4).

What does high fiber density mean?

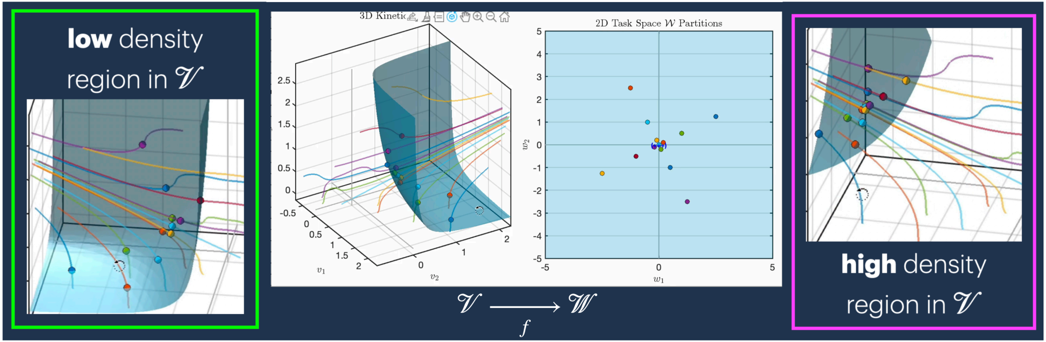

same actuator-rate budget $\Longrightarrow$
larger local task-rate change

[F26a].

But can we take the determinant of a co-metric?

The push-forward construction gives

$$d_v = df_v \circ g_v^{-1} \circ df_v^* \; : \; T^*_{f(v)}\mathscr{W} \longrightarrow T_{f(v)}\mathscr{W}, \qquad D(v) = J_f(v) G_v^{-1} J_f(v)^\top.$$

The objection

d_v maps covectors to vectors. It is not an endomorphism:

$$T^*_{f(v)}\mathscr{W} \neq T_{f(v)}\mathscr{W}.$$

Therefore, $\det D(v)$ and the eigenvalues of $D(v)$ are not intrinsic quantities by themselves.

What identifies task vectors with task-rate covectors?

[F26a] Sec. II-A.

The task metric closes the map

$$
\begin{tikzcd}[column sep=huge, row sep=large]
\begin{array}{c} T_v\mathscr{V} \\ \text{actuator rates} \end{array} \arrow[r, "d f_v"] & \begin{array}{c} T_{f(v)}\mathscr{W} \\ \text{task rates} \end{array} \arrow[r, "a_v = d_v \circ h_{f(v)}", "\begin{array}{c}\text{task-space}\\ \text{endomorphism}\end{array}"'] \arrow[d, shift right, "h_{f(v)}"'] & \begin{array}{c} T_{f(v)}\mathscr{W} \\ \text{task rates} \end{array} \\
\begin{array}{c} T^*_v\mathscr{V} \\ \text{actuator-rate covectors} \end{array} \arrow[u, "g_v^{-1}"] & \begin{array}{c} T^*_{f(v)}\mathscr{W} \\ \text{task-rate covectors} \end{array} \arrow[l, "d f_v^*"] \arrow[u, shift right, "d_v"'] &
\end{tikzcd}
$$

$$a_v = d_v \circ h_{f(v)} \; : \; T_{f(v)}\mathscr{W} \longrightarrow T_{f(v)}\mathscr{W}$$

Now determinants and eigenvalues have an intrinsic meaning.

[F26a] Sec. II-A.

But do the task chart or metric change the allocation?

Change of task chart

For
$$\widetilde{D}(v) = T_W(w)D(v)T_W(w)^\top,$$
the unweighted objective becomes

$$\sqrt{\det \widetilde{D}(v)} = |\det T_W(w)| \sqrt{\det D(v)}.$$

On $\mathcal{F}_w$, the factor is positive and constant.

Choice of task metric

The intrinsic objective at $f(v)$ satisfies

$$\begin{aligned}\rho_H(v) &= \sqrt{\det\big(D(v)H(f(v))\big)} \\ &= \sqrt{\det H(w)}\sqrt{\det D(v)}.\end{aligned}$$

Again, on $\mathcal{F}_w$, the additional factor is positive and constant.

We can safely use the coordinate determinant proxy for fiberwise optimization

$$\underset{v\in\mathcal{F}_w}{\arg\operatorname{ext}}\, \rho_H(v) = \underset{v\in\mathcal{F}_w}{\arg\operatorname{ext}} \sqrt{\det \widetilde{D}(v)} = \underset{v\in\mathcal{F}_w}{\arg\operatorname{ext}} \sqrt{\det D(v)}.$$

[F26a] Sec. II-A.

What does this geometry become for a multirotor?

The nonlinear signed-square rotor map makes promptness depend on the current rotor-speed configuration.

The nonlinear multirotor task map

- v_i: signed angular speed of rotor i;
- $a_i \in \mathbb{R}^m$: unit signed-square wrench contribution;
- $A = [a_1 \ \cdots \ a_n]$.

Key physical point

The static wrench is quadratic in speed, while the instantaneous wrench-rate gain is linear in the current speed magnitude.

$$f(v) = \sum_{i=1}^{n} a_i v_i |v_i|, \qquad J_f(v) = 2A\,\mathrm{diag}(|v|)$$

A linear allocation model Ax with $x_i := v_i|v_i|$ would erase the operating-point dependence that we want to exploit.

[F26a] Sec. V.

Two observables on every constant-wrench fiber

Energetic cost

$$P(v) = \sum_{i=1}^{n} c_i |v_i|^3$$

steady aerodynamic power proxy

Active promptness

For $G_v = I$,

$$\rho(v) = 2^m \sqrt{\det\left(A \operatorname{diag}(v^2) A^\top\right)}$$

wrench-rate volume under $\|\dot{v}\| \leq 1$

$$v \in \mathscr{F}_w \quad \longmapsto \quad \big(P(v), \rho(v)\big)$$

[F26a] Eqs. (11)–(15).

Is the energy–promptness trade-off always the same?

No. Its qualitative nature is determined
by the geometry and topology of the feasible task fiber.

Our recurring toy system: the aerodynamic dipole

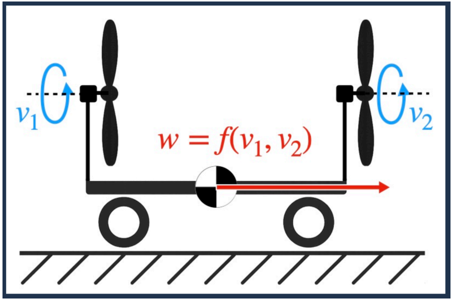

Two parallel rotors generate one scalar task:

$$w = f(v_1, v_2) = T_1(v_1) \pm T_2(v_2).$$

At zero inflow, for identical fixed-pitch rotors,

$$w = f(v_1, v_2) = k_T(v_1|v_1| \pm v_2|v_2|).$$

The $\pm$ choice depends on the sign convention of v and the rotor mounting, the geometric conclusions stay the same.

[F26a] Sec. VI; [F26b] Sec. III.

The signed-quadratic fiber family of the aerodynamic dipole

Fibers for constant w on the (v_1, v_2) space for the 2-input 1-task toy system

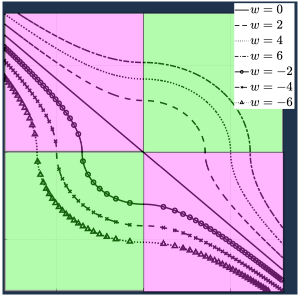

The same task map contains both
compact cooperative arcs and non-compact antagonistic branches.

[F26a] Secs. III and VI.

Cooperative versus antagonistic fibers

Cooperative contributions

$$a_1 v_1^2 + a_2 v_2^2 = w$$

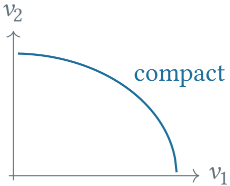

Antagonistic contributions

$$a_1 v_1^2 - a_2 v_2^2 = w$$

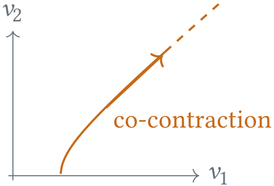

Geometric dichotomy

Compact fibers produce bounded trade-offs; antagonistic hyperbolic branches allow increasing internal loading until hardware bounds intervene.

[F26a] Secs. III and VI.

Aerodynamic dipole: exact co-contraction parameter

Restricting to the positive-speed antagonistic branch, $v_1, v_2 \geq 0$,

$$w = k_T(v_1^2 - v_2^2) = \bar{w} > 0.$$

Parameterize the positive branch by

$$v_2(s) = s, \qquad v_1(s) = \sqrt{\bar{w}/k_T + s^2}$$

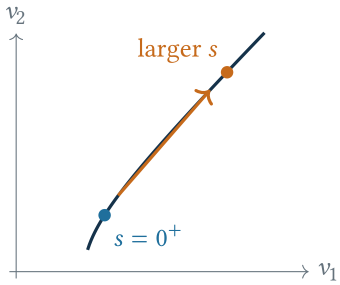

Then s parameterizes a **task-preserving family of internal states**:

$$\frac{dw}{ds} = 0, \qquad \frac{dv_1}{ds} > 0, \quad \frac{dv_2}{ds} > 0.$$

[F26a] Sec. VI; [F26b] Secs. III–IV.

Result I: active readiness has an energetic price

Along the antagonistic fiber,

$$P(s) = c\big(v_1(s)^3 + v_2(s)^3\big), \qquad \rho(s) = 2k_T\sqrt{v_1(s)^2 + v_2(s)^2}.$$

Monotonicity

$$\frac{dP}{ds} > 0, \qquad \frac{d\rho}{ds} > 0.$$

Increasing co-contraction raises both energetic cost and active promptness.

Opposite preferred states

$$\begin{aligned} \min P(s) &: \quad s = 0, \\ \max \rho(s) &: \quad \text{largest feasible } s. \end{aligned}$$

Energy and promptness select opposite ends of the feasible branch.

Strict energy–promptness conflict

Minimum energy selects zero co-contraction;
maximum promptness selects the largest hardware-admissible co-contraction.

[F26a] Sec. VI.

Compact versus antagonistic objective profiles

Cooperative: bounded tuning

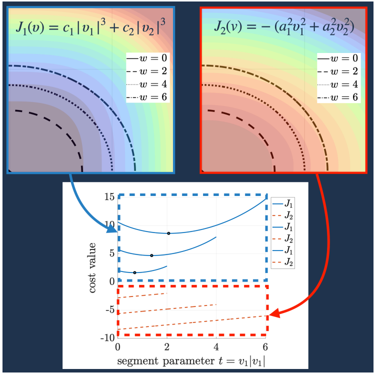

Antagonistic: strong monotone trade-off

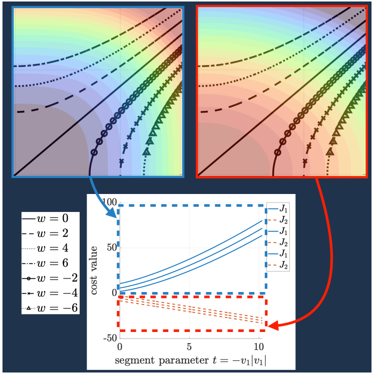

The same objectives ($J_1(v) := P(v)$ and $J_2(v) := -\rho(v)^2$)
behave differently after restriction to different fibers.

[F26a] Sec. VI.

Beyond the dipole: promptness tuning on a hexarotor

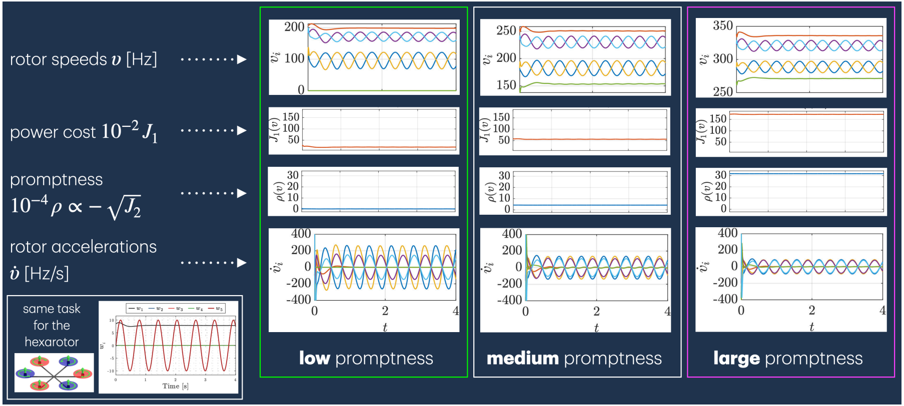

Low, medium, and high promptness allocations realize the **same wrench task**.

[F26a].

The same antagonistic map appears in a VSA

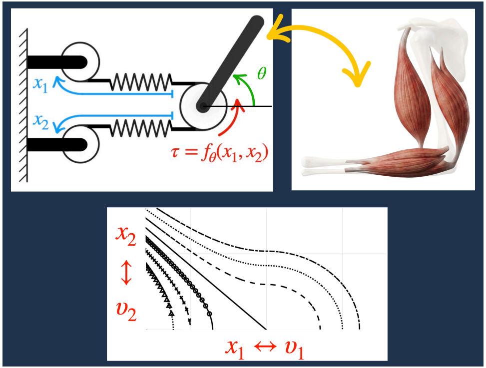

$$\tau = R[r(x_1 - R\theta) - r(x_2 + R\theta)]$$

At $\theta = 0$ with quadratic hardening,

$$f(x_1, x_2) = kR(x_1^2 - x_2^2).$$

$$(x_1, x_2) \longleftrightarrow (v_1, v_2)$$

[F26a] [F26b].

Mechanical reference: co-contraction shapes a VSA

For an antagonistic variable-stiffness actuator,

$$\tau(x) = R\big(r(x_1) - r(x_2)\big) = \bar{\tau}.$$

Active torque promptness

$$\rho(x) = R\sqrt{r'(x_1)^2 + r'(x_2)^2}$$

$\delta x \mapsto \delta\tau$

Passive stiffness

$$\sigma(x) = R^2\big(r'(x_1) + r'(x_2)\big)$$

$\delta\theta \mapsto \delta\tau$

Hardening tendons: $r'' > 0$

Every positive constant-torque co-contraction displacement increases both σ and ρ.

[F26a] Sec. VII; [F26b] Sec. II.

The "aerodynamic dipole/VSA" partial analogy

	Active task responsiveness	Passive interaction shaping
VSA / muscles	Torque promptness $\rho_x = R\sqrt{r'(x_1)^2 + r'(x_2)^2}$	Joint stiffness $\sigma_x = R^2(r'(x_1) + r'(x_2))$
Aerodynamic dipole	Aerodynamic promptness $\rho_v = \sqrt{\det(J_f G_v^{-1} J_f^\top)}$	Is there an aerodynamic passive counterpart?

[F26a] Sec. VII; [F26b] Sec. II.

What is the missing passive corner?

We need an external aerial variable that affects rotor thrust as joint displacement affects tendon force.

The missing external variable is axial inflow

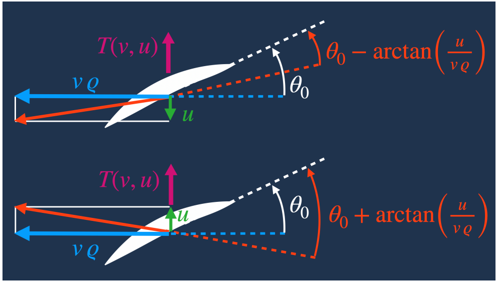

At a blade element of radius ϱ,

$$\phi(\varrho) \simeq \frac{u_{\text{in}}}{v\varrho}, \qquad \alpha(\varrho) \simeq \theta_0 \mp \frac{u_{\text{in}}}{v\varrho}.$$

With $C_L \simeq a\alpha$,

$$dT \propto \alpha(\varrho)(v\varrho)^2\, d\varrho \propto \theta_0(v\varrho)^2\, d\varrho \mp v\varrho\, u_{\text{in}}\, d\varrho.$$

zero-inflow thrust contribution ± linear inflow contribution

After integration: $T(v, u_{\text{in}}) = k_T v^2 \mp k_D v\, u_{\text{in}}.$

The key aerodynamic coupling

Thrust is locally linear in axial inflow, and
its inflow sensitivity grows linearly with rotor speed: $|\partial T/\partial u_{\text{in}}| = k_D v.$

[F26b] Sec. V.

Return to the aerodynamic dipole: the force map under axial inflow

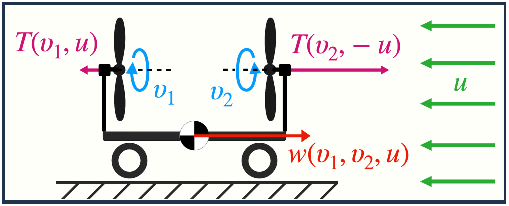

Both propellers experience the same air-relative velocity u. Because their axes are opposed, the corresponding axial inflows have opposite signs:

$$u_{\text{in},1} = u, \qquad u_{\text{in},2} = -u.$$

Neglecting mutual aerodynamic interference, the dipole force map is

$$f(v, u) = T(v_1, u) - T(v_2, -u).$$

At a fixed trim $(\bar{w}, \bar{u})$,

$$\mathscr{F}_{\bar{w}}^{\bar{u}} = \{v \mid f(v, \bar{u}) = \bar{w}\}.$$

Same trim force, different local force–inflow slope

Along $\mathscr{F}_{\bar{w}}^{\bar{u}}$, the task remains $\bar{w}$,
while $\partial f / \partial u$ can vary with the selected internal state v.

[F26b] Secs. III–IV and VI.

Passive observable: incremental aerodynamic damping

For the general thrust model, define the single-rotor inflow sensitivity

$$\lambda(v, u_{\text{in}}) := -\frac{\partial T}{\partial u_{\text{in}}}(v, u_{\text{in}}).$$

Then the dipole damping at trim $\bar{u}$ is

$$\sigma_a(v; \bar{u}) = -\left.\frac{\partial f}{\partial u}\right|_{\bar{u}} = \lambda(v_1, \bar{u}) + \lambda(v_2, -\bar{u})$$

Promptness: internal rate → active task-rate correction.

Damping: environmental velocity perturbation → passive task correction.

[F26b] Eqs. (15)–(16), (31)–(32).

The passive input–output map

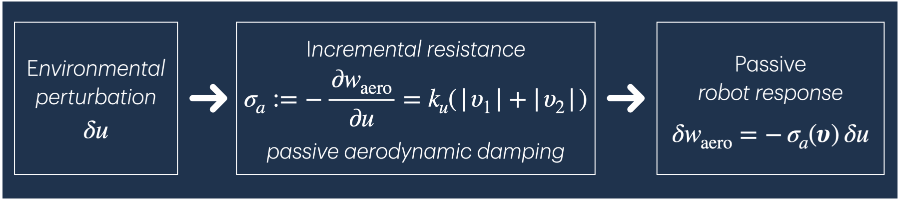

$$\delta u \longmapsto \delta w = -\sigma_a(v)\delta u$$

- input: environmental air-relative velocity perturbation;
- output: generated task response;
- allocation changes this passive map without changing the trim task.

[F26b].

Result II: co-contraction increases passive damping

Assume locally that

$$\frac{\partial T}{\partial v}(v, u_{\text{in}}) > 0, \qquad \frac{\partial \lambda}{\partial v}(v, u_{\text{in}}) > 0.$$

On a constant-task fiber,

$$0 = dw = \frac{\partial T}{\partial v}(v_1, \bar{u})\,\delta v_1 - \frac{\partial T}{\partial v}(v_2, -\bar{u})\,\delta v_2,$$

so a positive co-contraction direction has $\delta v_1, \delta v_2 > 0$. Therefore

$$d\sigma_a = \frac{\partial \lambda}{\partial v}(v_1, \bar{u})\,\delta v_1 + \frac{\partial \lambda}{\partial v}(v_2, -\bar{u})\,\delta v_2 > 0$$

Passive shaping at fixed task

Every admissible positive co-contraction displacement strictly increases incremental aerodynamic damping.

[F26b] Propositions IV.2 and VI.1.

Variable Aerodynamic Damping Actuation

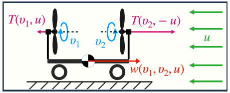

For $v_1, v_2 \geq 0$, substituting $T(v, u_{\text{in}}) = k_T v^2 - k_D v u_{\text{in}}$ into the aerodynamic-dipole task map gives

$$f(v, u) = \underbrace{k_T(v_1^2 - v_2^2)}_{\text{trim-task contribution}} - \underbrace{k_D(v_1 + v_2)u}_{\text{passive force–velocity contribution}}.$$

In this affine first-order model, define

$$\sigma_a(v) := k_D(v_1 + v_2).$$

Then

$$f(v, u) = k_T(v_1^2 - v_2^2) - \sigma_a(v)u.$$

Variable aerodynamic damping

Along a constant-task fiber, co-contraction preserves the trim task while increasing the passive aerodynamic damping $\sigma_a(v)$.

[F26b] Secs. V–VII.

The completed flying-muscle compass

	Active task responsiveness	**Passive interaction shaping**
VSA / muscles	Torque promptness $\rho_x = R\sqrt{r'(x_1)^2 + r'(x_2)^2}$	Joint stiffness $\sigma_x = R^2(r'(x_1) + r'(x_2))$
Aerodynamic dipole	Aerodynamic promptness $\rho_v = \sqrt{\det(J_f G_v^{-1} J_f^\top)}$	VADA Aerodynamic damping $\sigma_a = -\partial f / \partial u$

task-preserving co-contraction shapes both columns

Structural completion

The analogy is at the level of fiber motion and incremental input–output maps: stiffness and damping remain physically distinct.

Aerodynamic impedance: differential and common roles

With $w_0(v) := k_T(v_1^2 - v_2^2)$, the body dynamics become

$$m\dot{u} + k_D(v_1 + v_2)u = w_0(v) + F_{\text{ext}}$$

For $F_{\text{ext}} = 0$, equivalently,

$$u_{\text{eq}} = \frac{k_T}{k_D}(v_1 - v_2), \qquad c_{\text{app}} = k_D(v_1 + v_2).$$

differential component
sets the equilibrium

common component
sets the local damping

[F26b] Sec. VII.

The impedance correspondence at a glance

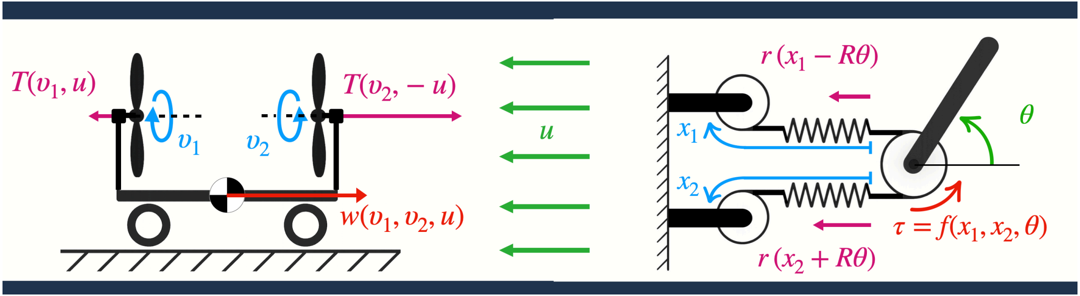

differential actuation sets equilibrium
common actuation sets impedance

[F26b] Sec. VII.

Is the effect large enough to matter?

Measured low-advance-ratio propeller data place the aerodynamic damping in the same range as ordinary low-speed body-drag damping.

Data-derived scale of the passive effect

From the standard propeller coefficients,

$$T = \rho_{\text{air}} n^2 D^4 C_T(J), \qquad J = \frac{u_{\text{in}}}{nD},$$

so an identical antagonistic pair at symmetric co-contraction has

$$\sigma_a = 2\rho_{\text{air}} n D^3 \big(-C_T'(0) \big)$$

Data set	Observed range	Median
16 propellers 62 retained runs	0.109–0.855 N s/m	0.296 N s/m

[F26b] Sec. VIII and Tables II–IV; UIUC Propeller Database.

The unified object: one fiber, three observables

$$v \in \mathscr{F}_{\bar{w}}^{\bar{u}} \quad \longmapsto \quad \Big(P(v),\ \rho(v),\ \sigma_a(v; \bar{u}) \Big)$$

Observable	Local meaning	Co-contraction trend
P	steady energetic cost	increases
ρ	active wrench-rate authority	increases
σ_a	passive force–velocity response	increases

Unified allocation problem

Select a point on the feasible trim fiber according to the required balance of endurance, active readiness, and passive interaction shaping.

[F26a] [F26b]

What remains open?

How does the aerodynamic-dipole principle extend from a scalar force module to full aerial systems, experimental models, control, and design?

What remains open?

Geometry and aerodynamic models

- full-wrench multirotor fibers;
- matrix-valued aerodynamic damping;
- coupled wakes and nonuniform inflow;
- experimentally identified thrust maps.

Control and aerial-system design

- energy–promptness–damping allocation;
- stability, passivity, and robustness;
- predictive co-contraction scheduling;
- propeller and rotor-layout co-design.

These questions define a broader research program on **redundancy as a resource for shaping aerial behavior**.

If you see a connection with your work, **please contact me**.

[F26a] [F26b]

Same task. Different behaviors.

Redundant actuation creates internal states.
Aerodynamic co-contraction uses them
to shape energy,
active promptness, and
passive aerodynamic damping.

Antonio Franchi schol@r-franchi.eu