Symmetries and invariants in behavioral optimal control

Debra Lewis
Mathematics Department
UC Santa Cruz

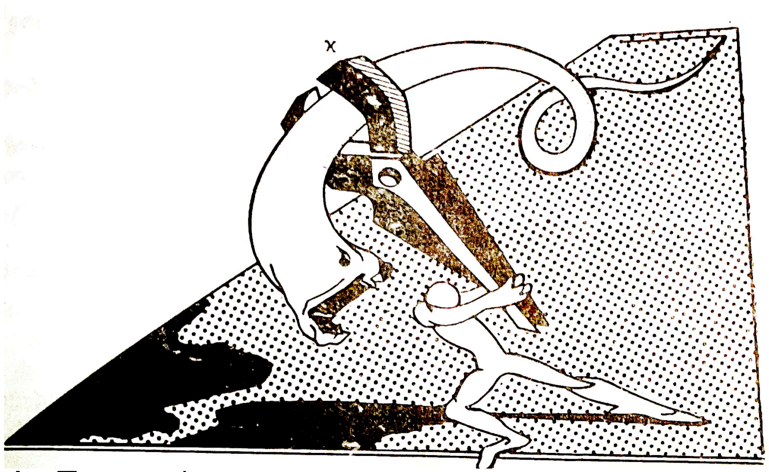

A. Fomenko

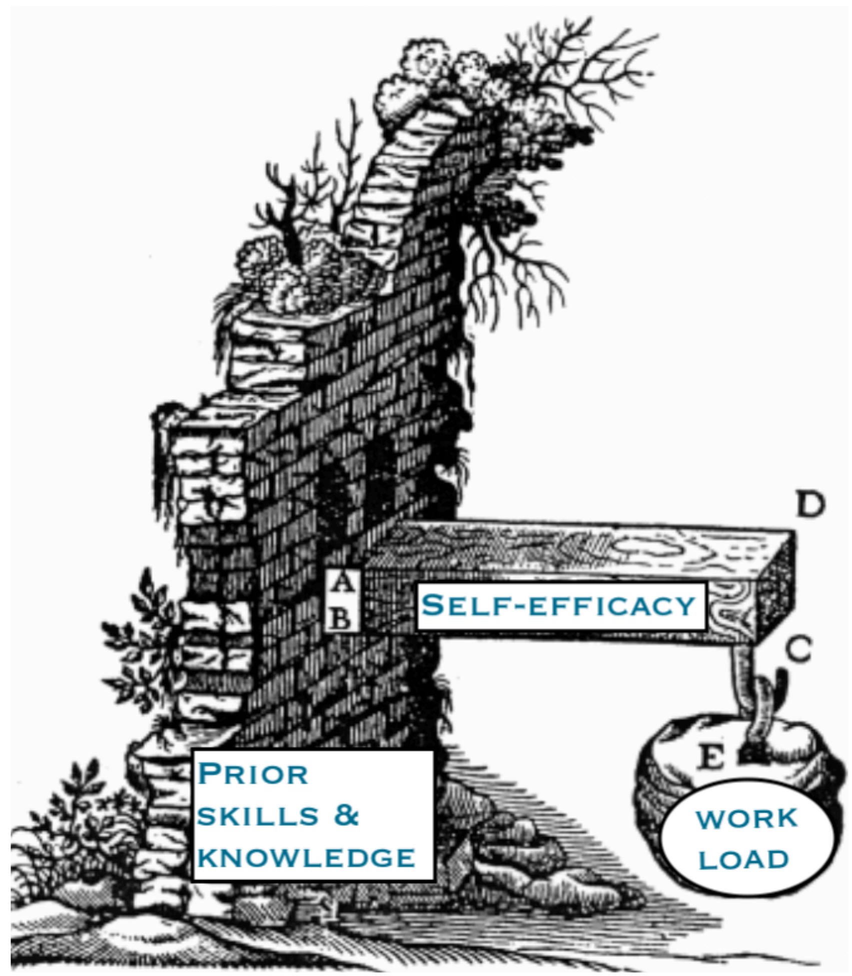

Apologies to Galileo

Why *should* students study?

Value of computation, introduced in machine learning, has been successfully used in models of human cognitive investment.

> *we are constitutively reluctant to mobilize all available cognitive resources. That is, mental effort is inherently aversive or costly...*
>
> *individuals can increase their control allocation when higher incentives are on offer (i.e., they are not constrained by ability) but hold back from doing so*
>
> A. Shenhav et al. (2017)

Many students lack confidence in the future benefits of their academic progress relative to the costs of studying.

Math anxiety and stereotype threat make some cognitive tasks torturous for many students, leading to academic disengagement.

Modeling aggregated course sequence outcomes

Attrition from required STEM sequences is costly for both students and universities.

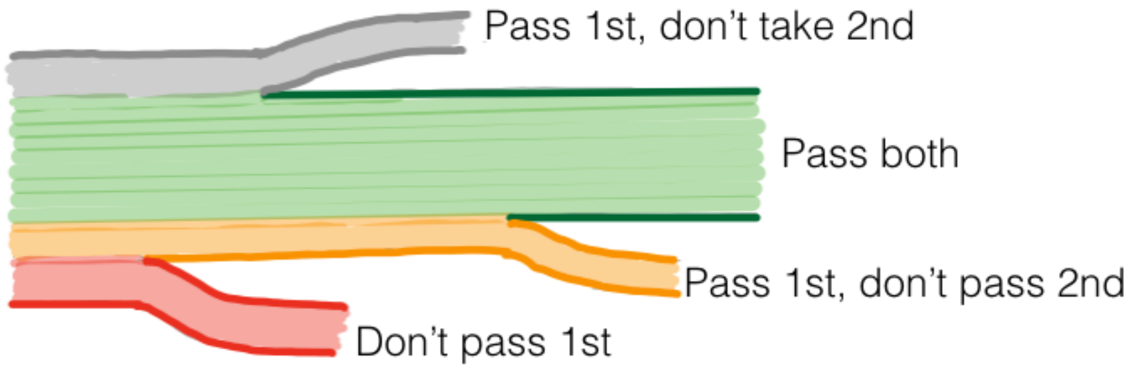

If 80% of students who take each course pass it, and 80% of those who pass the first course take the second one, 51% of the students who start the sequence complete it.

Replacing 80 with 90 yields a 73% completion rate.

Q: Does math placement predict persistence rates for two-course calculus sequences?

Attrition from Calculus for the Physical Sciences, Math & Engineering

Socio-economic   Solid: Low income, 1st Gen, or vet. Dashed: "privileged"

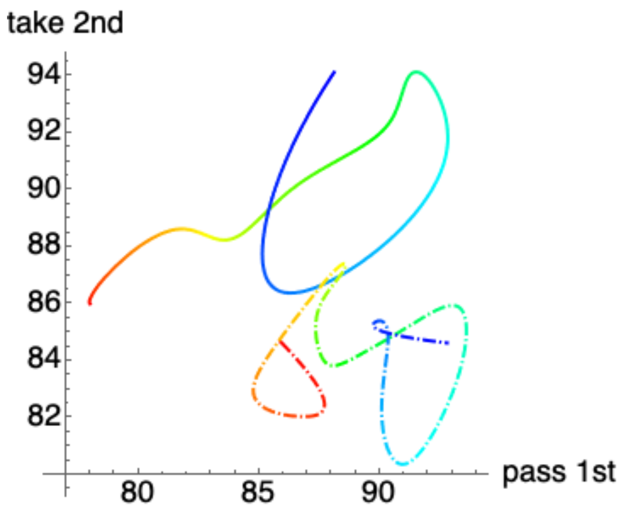

Gender   Solid: Female. Dashed: Male

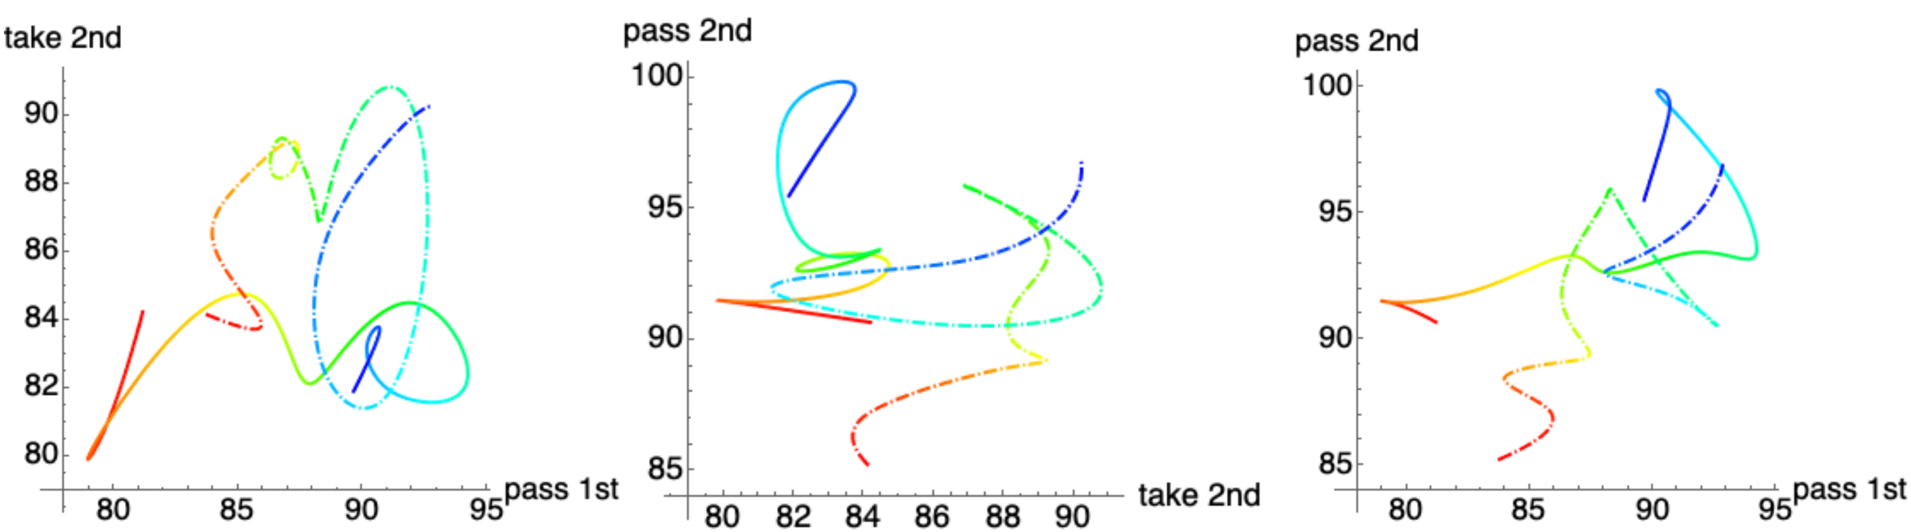

Attrition from Calculus for the Life Sciences

Socio-economic   Solid: Low income, 1st Gen, or vet. Dashed: “privileged”

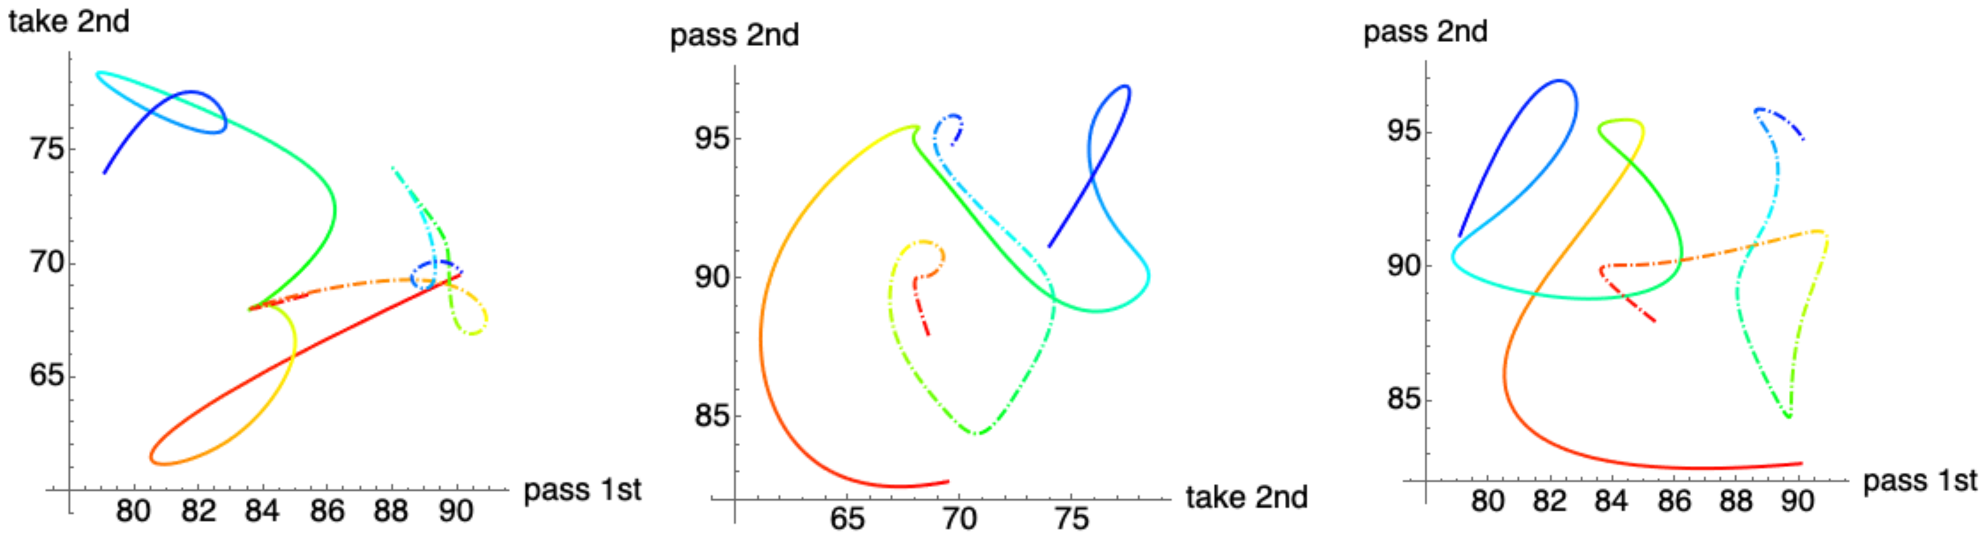

Gender   Solid: Female. Dashed: Male

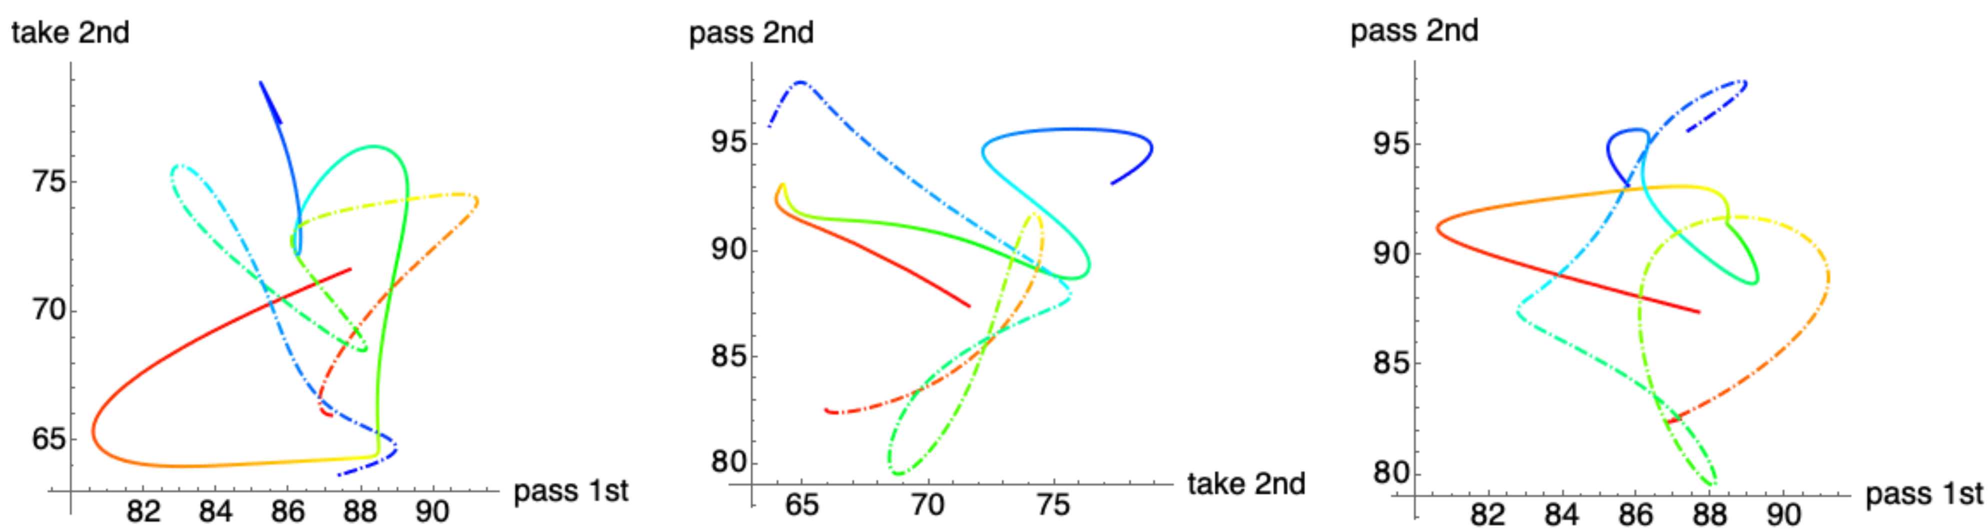

From the vaults: the Frenet frame

The *Frenet frame* of a curve $\alpha \in C^3\left([0, L], \mathbb{R}^3\right)$ is a curve $F \in C^1\left([0, L], SO(3)\right)$.

The columns of F are $\mathbf{t}$, the unit tangent; $\mathbf{n}$, the unit normal; and $\mathbf{b}$, the binormal to α. If α is parametrized by arc length, then

$$\mathbf{t} = \alpha' \qquad \mathbf{n} = \mathbb{P}_{\mathbf{t}}\alpha'' \qquad \mathbf{b} = \alpha' \times \alpha''.$$

The *Frenet equations:*

$$F' = F\widehat{X_L}, \qquad \text{for} \qquad X_L = \tau\,\mathbf{e}_1 + \kappa\,\mathbf{e}_3,$$

where

- κ and τ denote the curvature (assumed to be everywhere nonzero) and torsion of α,
- $\{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}$ is the standard Euclidean basis of $\mathbb{R}^3$, and
- $\hat{\ } : \mathbb{R}^3 \to so(3)$ is determined by $\hat{\mathbf{x}}\mathbf{y} = \mathbf{x} \times \mathbf{y} \qquad \forall\ \mathbf{y} \in \mathbb{R}^3$.

The elastica (Bernoulli, 1691; Euler, 1744)

An *elastica* is a curve $\alpha \in C^3\left([0, L], \mathbb{R}^3\right)$ minimizing

$$\int_\alpha \kappa(s)^2 ds = \int_0^L \kappa(t)^2 \left\|\alpha'\right\| dt.$$

The elastica is a highly simplified model of an inextensible elastic rod with infinitesimal uniform cross section.

Standard variational approaches yield the evolution equation

$$\alpha''' - \lambda\alpha' = c$$

for some Lagrange multiplier $\lambda : [0, L] \to \mathbb{R}$ and constant c .

Equivariance of the system with respect to Euclidean motions yields conservation laws relating κ, κ', and τ.

Curvature and torsion pairs (κ, τ) determine the Frenet frame F, which in turn determines the curve α, given initial data for (α, F).

Some representative elastica solutions

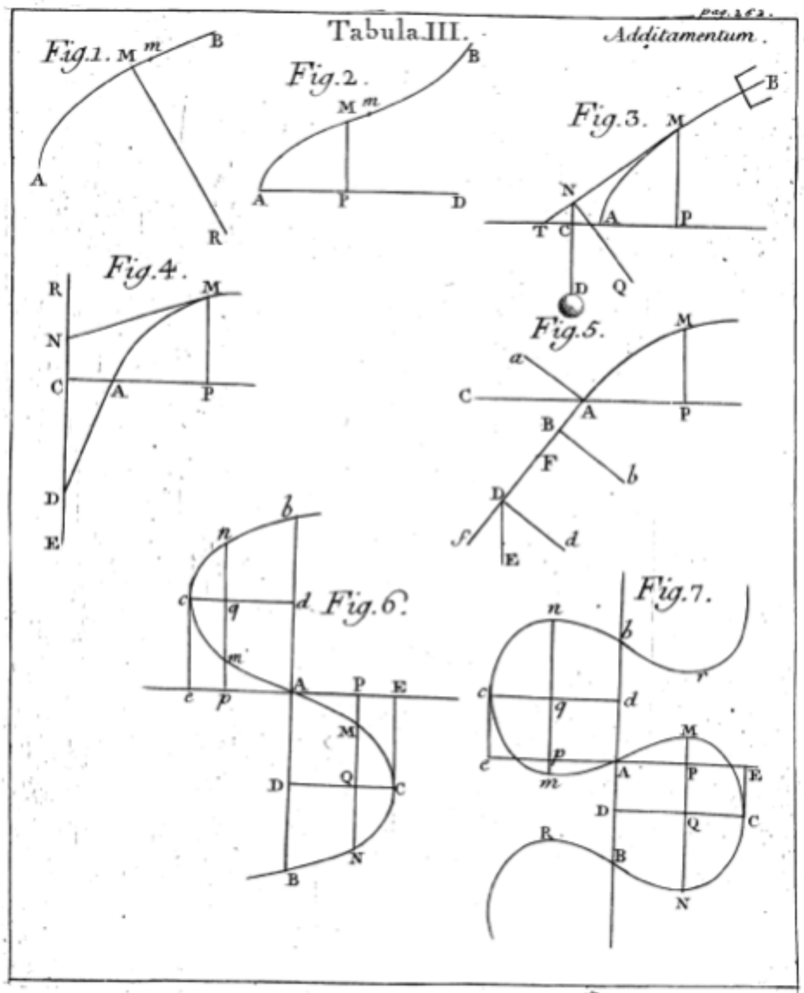

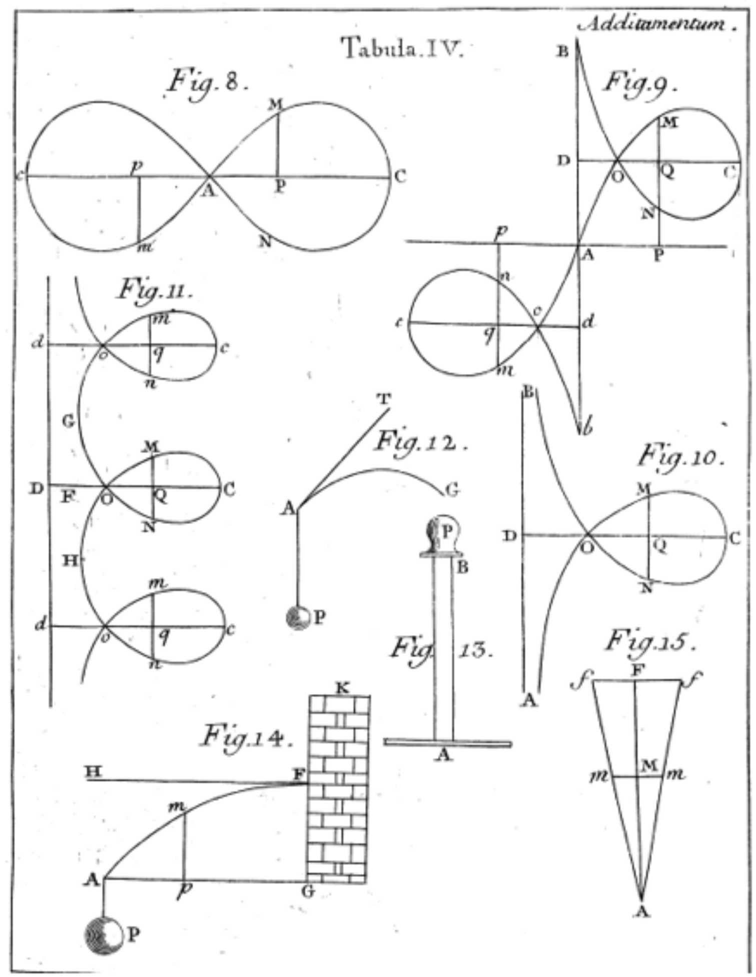

Euler, 1744, by way of “The elastica: a mathematical history”, R. Levien, EECS UC Berkeley, 2008

Generalized elasticas

Replacing κ^2 with some function of curvature, torsion, and independent parameter determining a well-defined variational problem yields *generalized elasticas*.

Given outcome data with a single independent parameter, e.g. best placement score, we can (try to) identify functions for which our sequence outcome data approximate generalized elasticas.

What about initial data?
Rather than using the data for the lowest placement value, set conditions on some choice or analog of means:

- Translation—arithmetic mean of $\alpha(t)$
- Rotation—if there exists a neighborhood $\mathcal{U}$ of the origin such that $\exp|_{\mathcal{U}}$ is invertible and $F(t) \in \exp(\mathcal{U})$ for $0 \leq t \leq L$, then use the geometric mean of F, i.e. exp of the arithmetic mean of $\exp^{-1}(F(t))$. If not, ...

The generalized Frenet equations and Kirchhoff rods

A curve $(\alpha, F, \mathbf{u}) \in C^1\left([0,L], \mathbb{R}^3 \times SO(3) \times \mathbb{R}^3\right)$ satisfies the *generalized Frenet equations* if

$$F\mathbf{e}_1 = \alpha' \qquad \text{and} \qquad F' = F\hat{\mathbf{u}}.$$

A *Kirchhoff rod* is a solution of these equations minimizing the *total elastic energy*

$$\tfrac{1}{2}\int_0^L \mathbf{u}^T \mathrm{diag}[\mathbf{c}]\mathbf{u}\,ds \qquad \text{for some} \quad \mathbf{c} \in [0,\infty)^3.$$

A Kirchhoff rod models an elastic rod with centerline α and infinitesimally thick circular uniform cross-section determined by F.

- Twisting energy: $c_1 u_1^2$
- Bending energy: $c_2 u_2^2 + c_3 u_3^2$

An untwisted symmetric ($c_2 = c_3$) Kirchhoff rod is an elastica.

Control systems

The generalized Frenet equations can be interpreted as a control system with control $\mathbf{u}$.

The state space $\mathcal{S}$ of a control system is typically a manifold (possibly with boundary).

Each state $m \in \mathcal{S}$ has an associated admissible control region $\mathcal{A}_m$, which is often a manifold with boundary.

The set $\mathcal{A}$ of admissible state and control pairs is given by

$$\mathcal{A} := \{(m, u) : m \in \mathcal{S} \text{ and } u \in \mathcal{A}_m\}$$

The evolution of the state variable m is determined by a controlled vector field $X : \mathcal{A} \to T\mathcal{S}$.

A solution of the *fixed time control problem* for X, with boundary data m_0 and m_f, and duration t_f is a curve (m, u) in $\mathcal{A}$ satisfying

$$m' = X(m, u), \qquad m(0) = m_0, \quad \text{and} \quad m(t_f) = m_f.$$

Optimal control and Pontryagin's maximum principle

Given an instantaneous cost function $C \in C^0(\mathcal{A}, \mathbb{R})$, a solution (m_*, u_*) of the *optimal control problem* is a solution of the control problem that minimizes the total cost over all such solutions (m, u)

$$\int_0^{t_f} C(m_*(t), u_*(t))dt \leq \int_0^{t_f} C(m(t), u(t))dt.$$

Pontryagin's maximum principle relates optimal control to Hamiltonian dynamics, introducing an auxiliary variable in the fiber of the cotangent bundle of the state space $\mathcal{S}$.

Potentially optimal trajectories can be constructed using Hamiltonian methods—including exploitation of symmetries!

The Kirchhoff rod can be regarded as an optimal control problem, with cost function equal to the elastic energy.

We can seek coefficients **c** such that our sequence outcome data approximates a Kirchhoff rod.

Multiple measures and more general rod models

Most incoming pre-STEM UCSC students have multiple math assessments scores on record—one or more SAT Math section scores, and possibly multiple ALEKS PPL scores.

We can use this additional information about prior knowledge and math skills in our models, while still prioritizing the role of the placement score.

Ansatz:
There are strong, but not overwhelming, correlations between placement score (best ALEKS PPL score) and other pre-arrival math assessment scores.

Small changes in placement score result in small changes in the other measures aggregated by placement score.

Strategy:
Use rod models involving an assignment of a basis of $\mathbb{R}^3$ at each point along the centerline α.

Process:

1. Determine the curve α as before.
2. Compute mean initial and best SAT Math section scores, and initial ALEKS PPL scores, aggregated by demographic group and placement score.
3. Fit matrices $F(t)$ to the 'infinitesimal' changes in these input means and the outcome means as t increases.
4. Use rod models that require only that $F(t)$ be nonsingular (so far, so good) and $\alpha'(t)^T F(t)\mathbf{e}_1 > 0$.
5. Fit elastic energies to the (α, F) outcome curves, such that these curves are approximate solutions of the variational or optimal control problem determined by the corresponding elastic energy.

Equally weighted outcomes $\rightsquigarrow$ *left* invariance under isometries. Assumption of equal significance of changes in the multiple measures $\rightsquigarrow$ *right* invariance under isometries.

Modeling individual academic engagement

When modeling individual study effort investment over time using optimal control theory, we consider instantaneous cost functions combining three terms:

- a solely state-dependent term that models the costs (both internal and external) of the current state,
- a solely control-dependent term that models 'life balance' costs—the time and cognitive effort invested in studying isn't available for other purposes (e.g. job, child care, recreation)
- a term modeling the psychological costs of studying; e.g., anxiety and stress resulting from low self-efficacy.

For simplicity, we focus on control-affine vector fields $X(m, u)$, with drift field modeling decay of the state variables when little or no effort is invested, and linear life balance costs.

We assume that the controls lie in unit balls with respect to state-dependent inner products, and consider psychological costs ψ of the form

$$\psi(\mathbf{m},\mathbf{u}) = \psi_{\mathrm{pe}}(\mathbf{m}) - \mu(\mathbf{m})\xi(\|\mathbf{u}\|_{\mathbf{m}}) \,,$$

where $\psi_{\mathrm{pe}} \in C^1(\mathcal{S},\mathbb{R})$ denotes the psychological cost of maximal study effort $\|\mathbf{u}\|_{\mathbf{m}} = 1$, $\mu \in C^1(\mathcal{S},\mathbb{R})$, and $\xi \in C^0([0,1],[0,1])$

1. is continuously differentiable on $[0,1)$,
2. satisfies $\xi(0) = 1$, $\xi'(0) = 0$, and $\xi(1) = 0$,
3. either $\xi'(1)$ exists or $\lim_{s\to 1} |\xi'(s)| = \infty$.

The special case

$$\xi_\circ(s) := \sqrt{1-s^2}$$

yields particularly simple expressions, while enforcing the constraint $\|\mathbf{u}\| \leq 1$.

Symmetry for the sake of tractability

For learning models with a scalar state (subject mastery) and sufficiently simple costs, standard undergraduate level methods for reducing conservative systems to quadrature can be used.

Bifurcation analyses can then be applied to decompose the associated parameter (cost coefficients) and initial data space into regions with qualitatively similar behavior.

Cost terms invariant with respect to a group action on the state space yield optimal control systems to which all the standard techniques for Hamiltonian systems with symmetry—Noether's Theorem, (singular) reduction, etc.—can be applied.

The aggregate sequence outcome curves suggest that subject mastery doesn't do all the heavy lifting in academic success $\rightsquigarrow$ learning costs invariant with respect to actions combining mastery, self-efficacy, etc. seem plausible.

References

J. Alexander and S. Antman (1982) The ambiguous twist of Love

T. Griffiths, F. Lieder, and N. Goodman (2015), Rational use of cognitive resources: Levels of analysis between the computational and the algorithmic

D. Lewis (2015), A soothing invisible hand: moderation potentials in optimal control

D. Lewis (2022), Modeling student engagement using optimal control theory

L. Pontryagin et al. (1962) *The Mathematical Theory of Optimal Processes*

A. Shenhav et al. (2017) Toward a rational and mechanistic account of mental effort

D. Singer (2007), Lectures on Elastic Curves and Rods