Episodic Learning with Control Lyapunov Functions for Uncertain Robotic Systems

Andrew Taylor Victor Dorobantu Hoang Le
Yisong Yue Aaron D. Ames

Computing and Mathematical Sciences
California Institute of Technology

November 7^{th}, 2019

Caltech

Control in the real world is hard

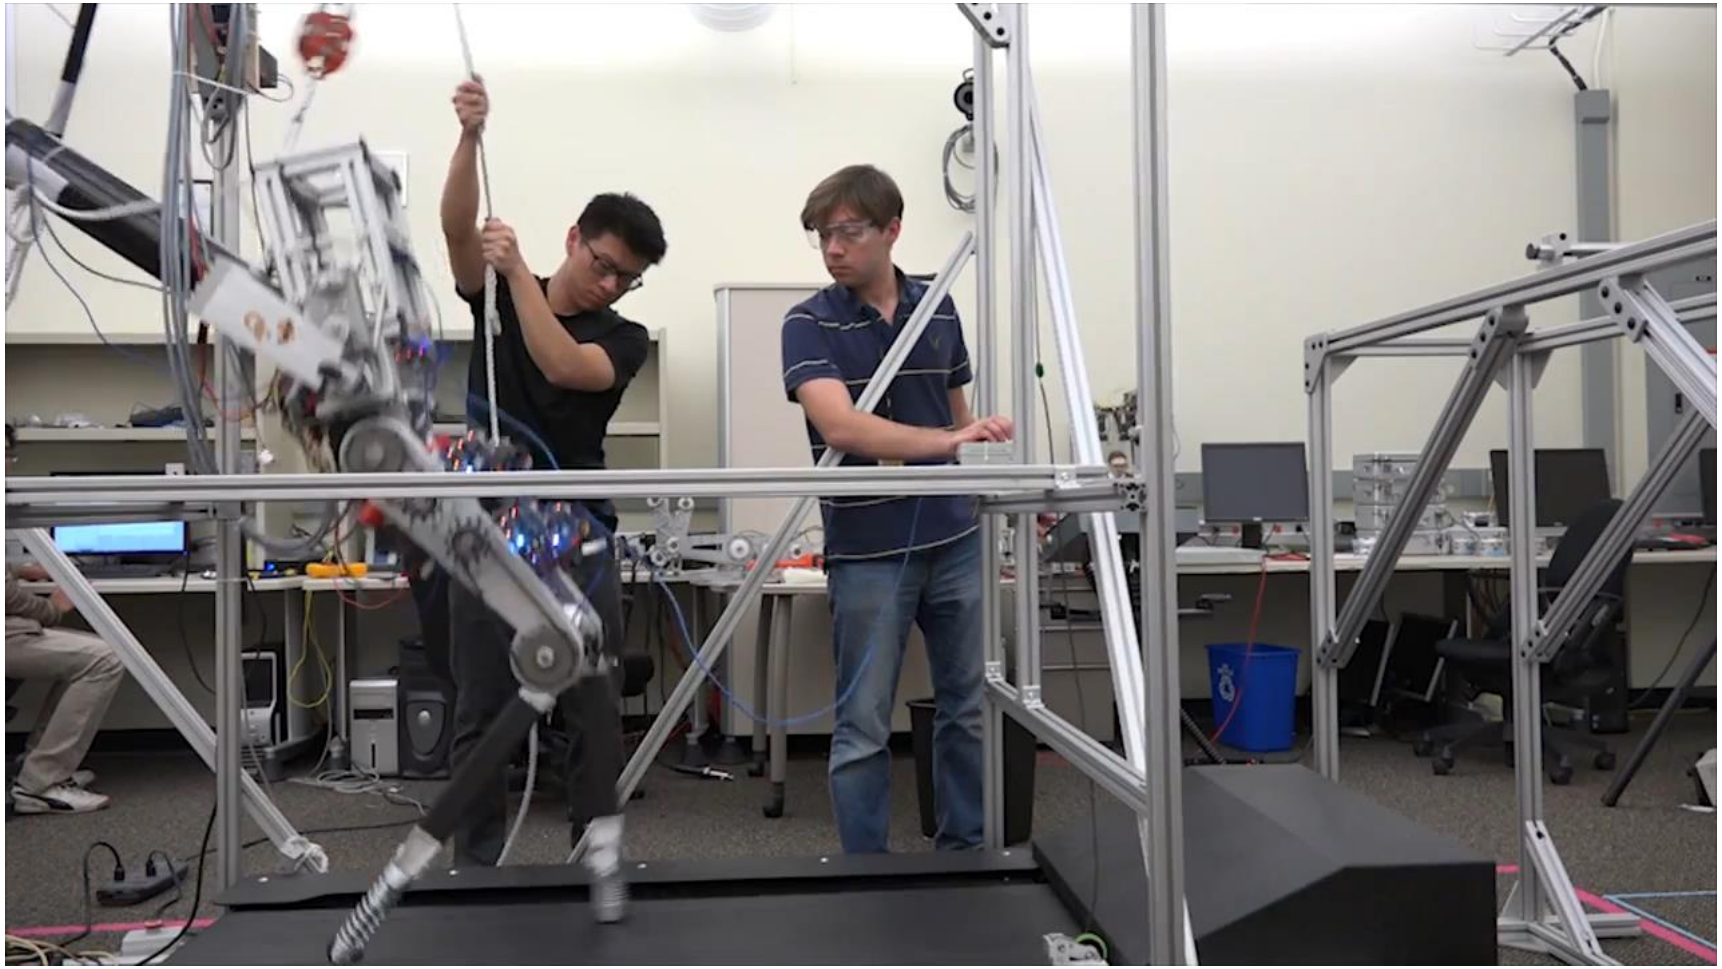

Caltech
Andrew J. Taylor

Control in the real world is hard

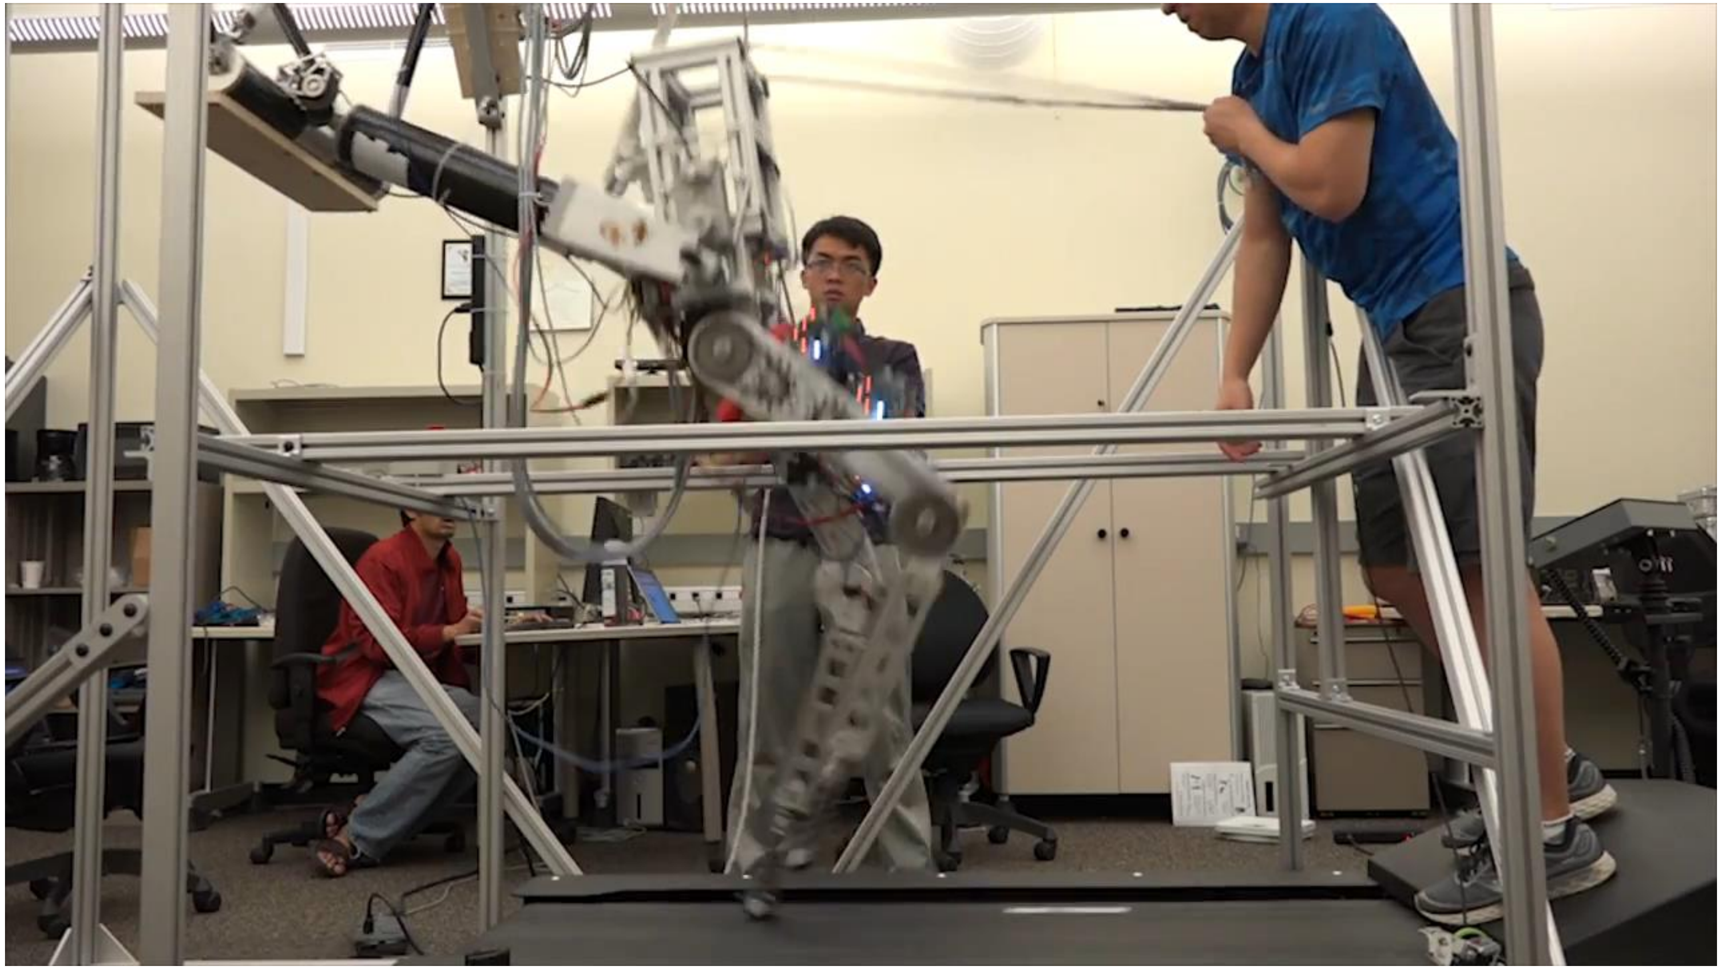

Caltech

Andrew J. Taylor

But: Pretty when it works…

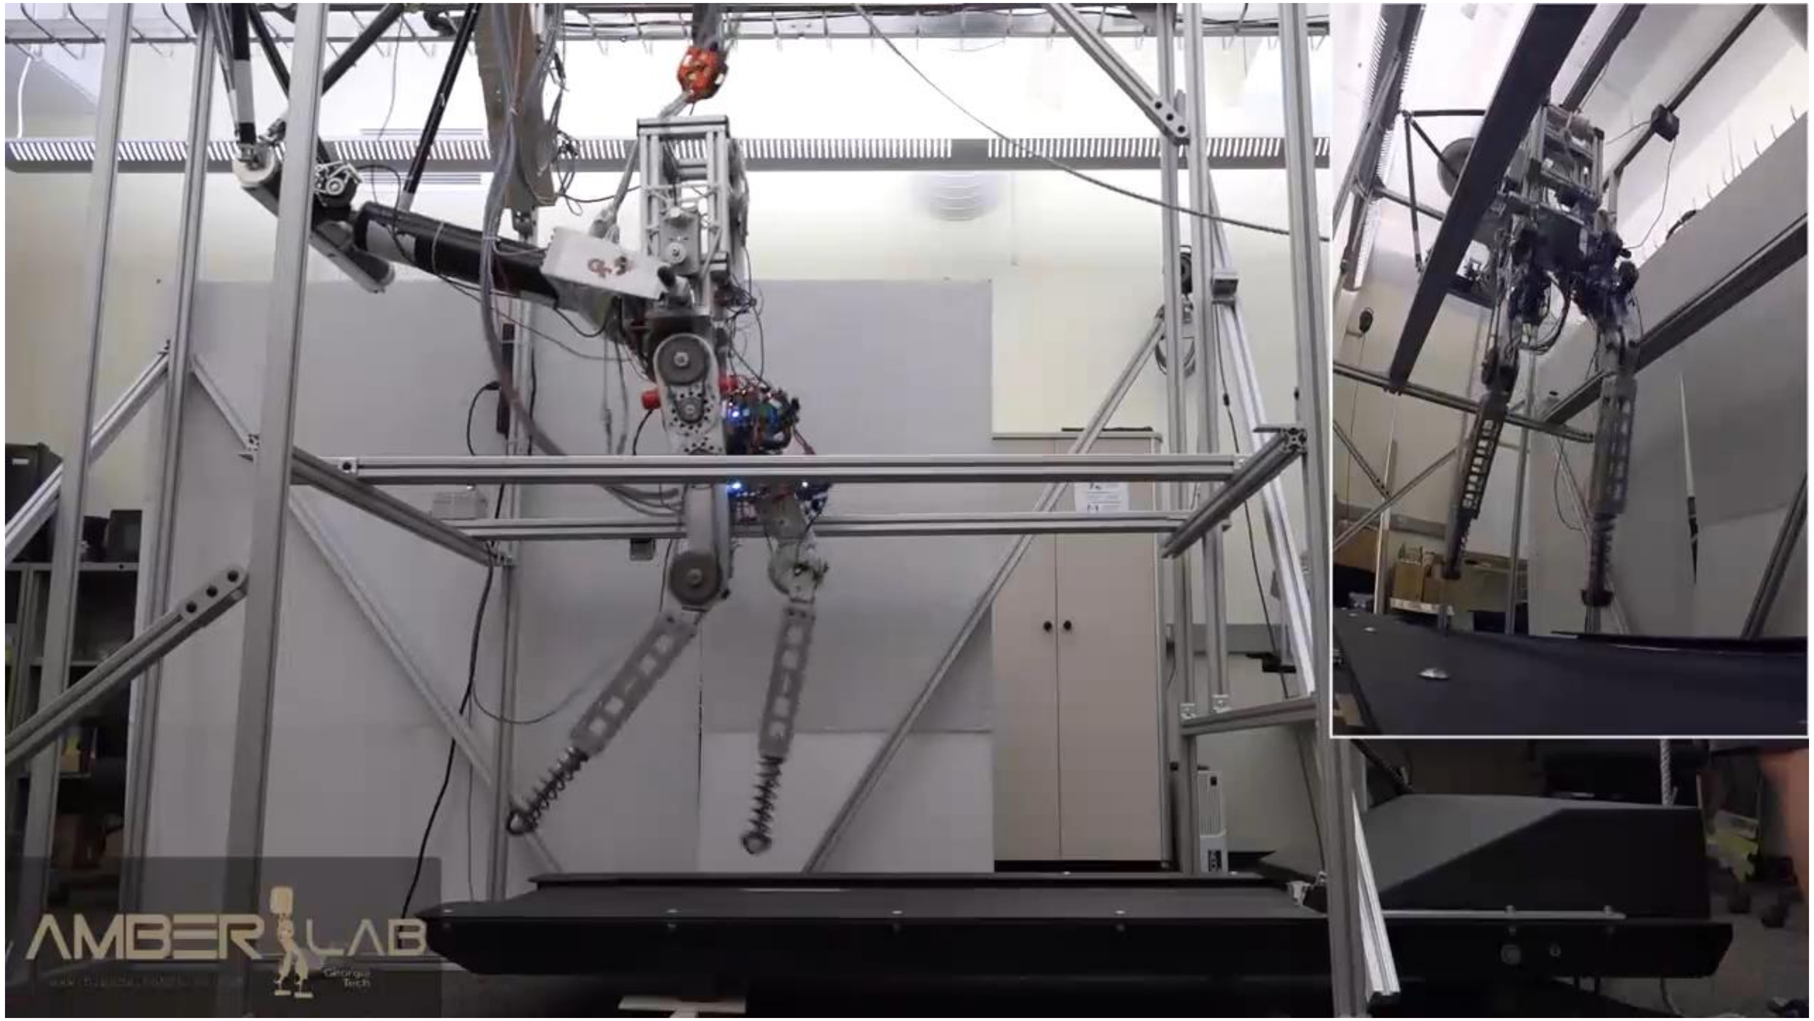

W. Ma, et al., Bipedal robotic running with durus-2d: Bridging the gap between theory and experiment

Claim: Need to Bridge the Gap

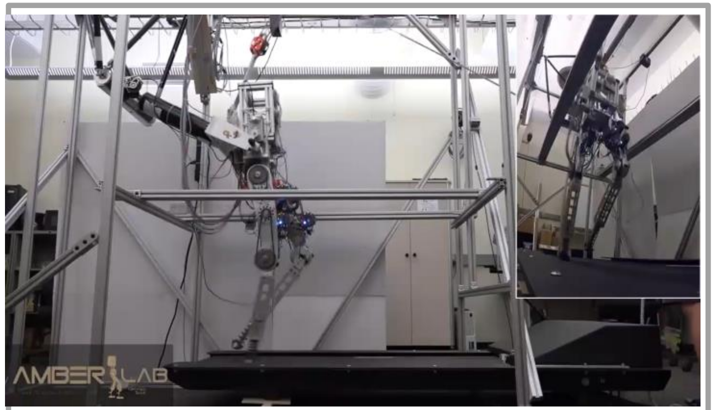

$$\mathbf{k}(\mathbf{q}, \dot{\mathbf{q}}, t) = \underset{\mathbf{u} \in \mathbb{R}^m}{\operatorname{argmin}} \ \|\mathbf{u}\|_2$$
$$\text{s.t.} \quad \dot{V}(\mathbf{q}, \dot{\mathbf{q}}, t, \mathbf{u}) \leq -\alpha V(\mathbf{q}, \dot{\mathbf{q}}, t)$$

Bridge the Gap

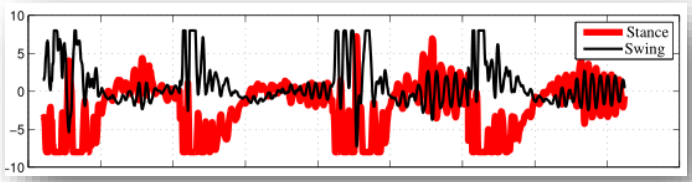

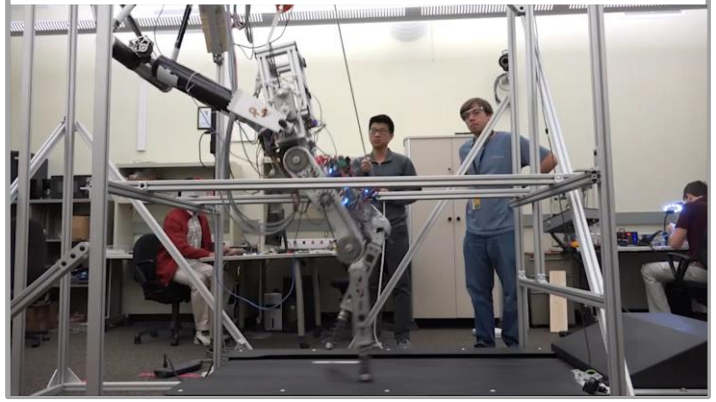

Theorems & Proofs

Experimental Realization

Caltech

Andrew J. Taylor

Robotic Dynamics

Equations of Motion

$$\widehat{\mathbf{D}}(\mathbf{q})\ddot{\mathbf{q}} + \underbrace{\widehat{\mathbf{C}}(\mathbf{q},\dot{\mathbf{q}})\dot{\mathbf{q}} + \widehat{\mathbf{G}}(\mathbf{q})}_{\widehat{\mathbf{H}}(\mathbf{q},\dot{\mathbf{q}})} = \widehat{\mathbf{B}}\mathbf{u}$$

$$\mathbf{q} \in \mathcal{Q} \subseteq \mathbb{R}^n \qquad \mathbf{u} \in \mathbb{R}^m$$

Mathematical Model

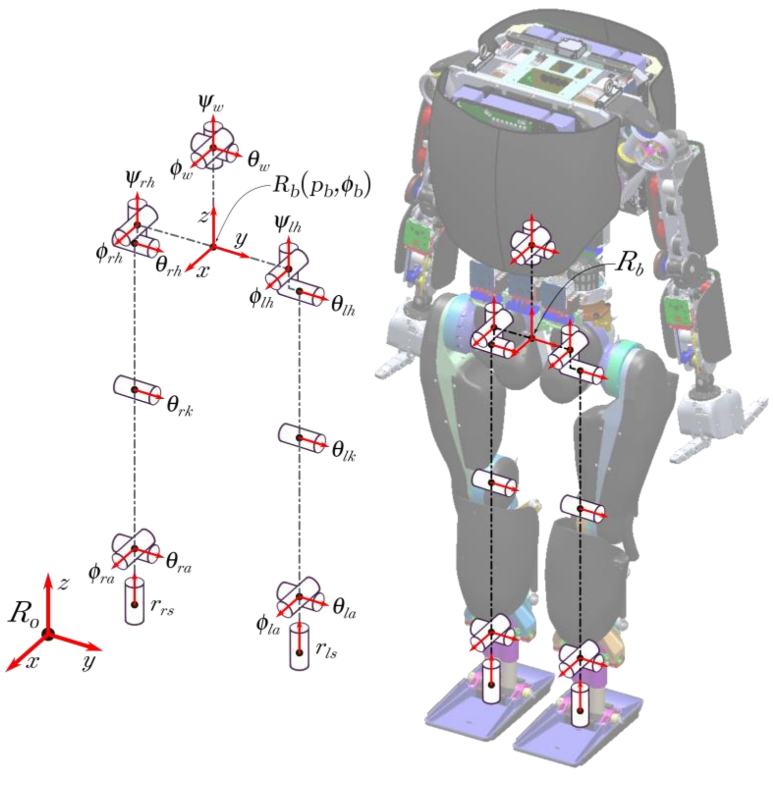

Robot Model

Robotic Dynamics

Equations of Motion

$$\underbrace{\widehat{\mathbf{D}}(\mathbf{q})\ddot{\mathbf{q}} + \widehat{\mathbf{C}}(\mathbf{q},\dot{\mathbf{q}})\dot{\mathbf{q}} + \widehat{\mathbf{G}}(\mathbf{q})}_{\widehat{\mathbf{H}}(\mathbf{q},\dot{\mathbf{q}})} = \widehat{\mathbf{B}}\mathbf{u}$$

$$\mathbf{q} \in \mathcal{Q} \subseteq \mathbb{R}^n \qquad \mathbf{u} \in \mathbb{R}^m$$

Assume Fully Actuated*

$$\widehat{\mathbf{B}} \in \mathbb{R}^{n\times n} \qquad \mathrm{rank}(\widehat{\mathbf{B}}) = n$$

Mathematical Model

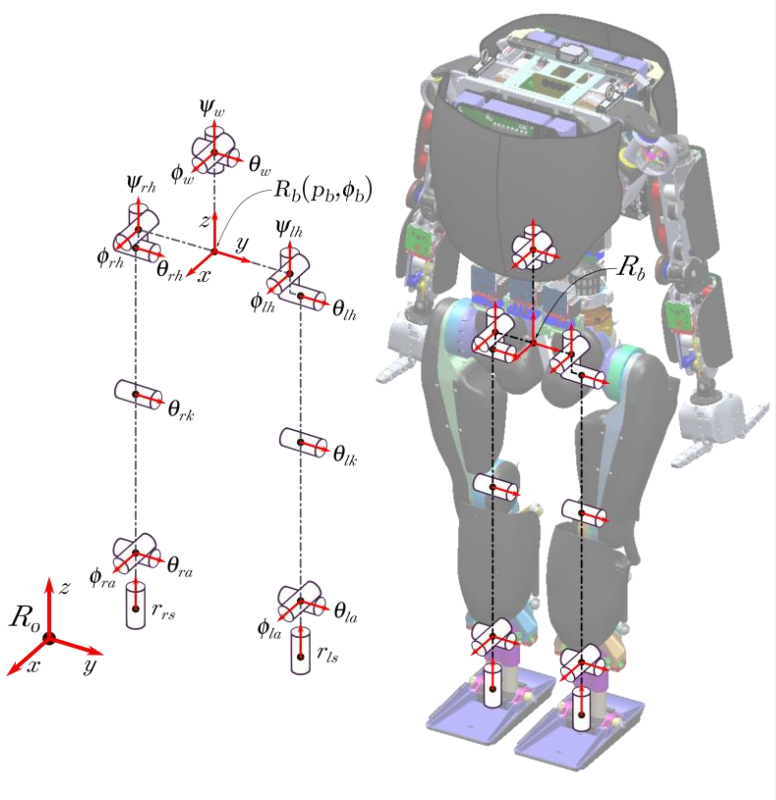

Robot Model

*Under-actuated output tracking formulation in full text.

Caltech

Andrew J. Taylor

Control Objective

Direct Collocation Trajectory Optimization

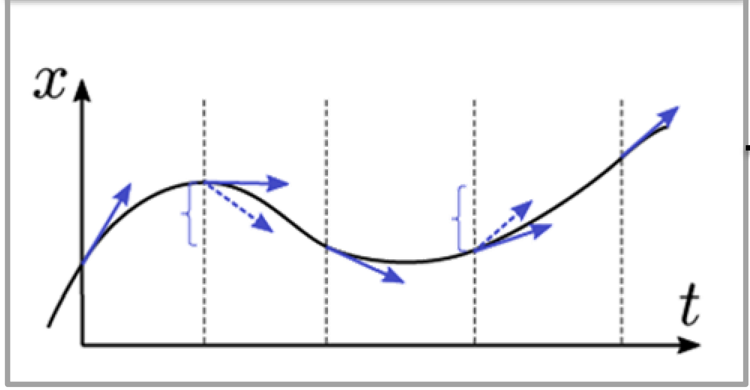

$$\begin{bmatrix}\mathbf{q}_d(t)\\ \dot{\mathbf{q}}_d(t)\end{bmatrix} \quad - \quad + \quad \begin{bmatrix}\mathbf{q}\\ \dot{\mathbf{q}}\end{bmatrix}$$

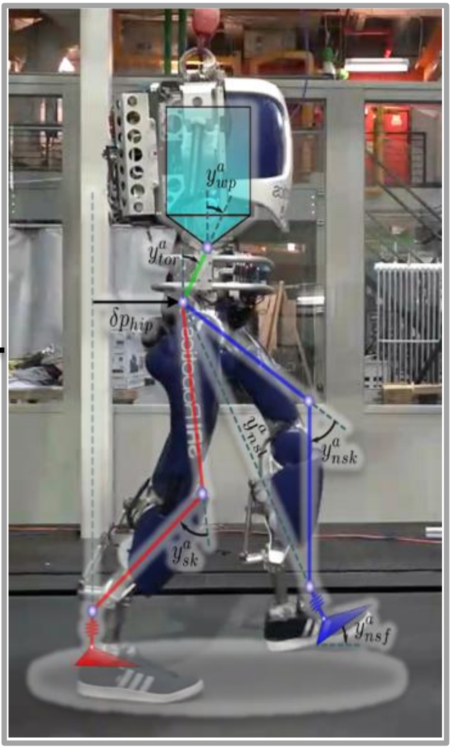

$$\boldsymbol{\eta} = \begin{bmatrix}\mathbf{q}-\mathbf{q}_d(t)\\ \dot{\mathbf{q}}-\dot{\mathbf{q}}_d(t)\end{bmatrix}$$

Error Dynamics

$$\dot{\boldsymbol{\eta}} = \frac{\mathrm{d}}{\mathrm{d}t}\begin{bmatrix}\mathbf{q}-\mathbf{q}_d(t)\\ \dot{\mathbf{q}}-\dot{\mathbf{q}}_d(t)\end{bmatrix} = \underbrace{\begin{bmatrix}\dot{\mathbf{q}}\\ -\widehat{\mathbf{D}}(\mathbf{q})\widehat{\mathbf{H}}(\mathbf{q},\dot{\mathbf{q}})\end{bmatrix}}_{\widehat{\mathbf{f}}(\mathbf{q},\dot{\mathbf{q}})} - \underbrace{\begin{bmatrix}\dot{\mathbf{q}}_d(t)\\ \ddot{\mathbf{q}}_d(t)\end{bmatrix}}_{\dot{\mathbf{r}}(t)} + \underbrace{\begin{bmatrix}\mathbf{0}\\ (\widehat{\mathbf{D}}(\mathbf{q}))^{-1}\widehat{\mathbf{B}}\end{bmatrix}}_{\widehat{\mathbf{g}}(\mathbf{q})}\mathbf{u}$$

Caltech
Andrew J. Taylor

Computed Torque

$$\dot{\boldsymbol{\eta}} = \widehat{\mathbf{f}}(\mathbf{q}, \dot{\mathbf{q}}) - \dot{\mathbf{r}}(t) + \widehat{\mathbf{g}}(\mathbf{q})\mathbf{u}$$

Computed Torque

$$\dot{\boldsymbol{\eta}} = \widehat{\mathbf{f}}(\mathbf{q}, \dot{\mathbf{q}}) - \dot{\mathbf{r}}(t) + \widehat{\mathbf{g}}(\mathbf{q})\mathbf{u}$$

Computed Torque (Feedback Linearization)

$$\mathbf{k}(\mathbf{q}, \dot{\mathbf{q}}, t) = \widehat{\mathbf{B}}^{-1}\widehat{\mathbf{D}}(\mathbf{q})\left((\widehat{\mathbf{D}}(\mathbf{q}))^{-1}\widehat{\mathbf{H}}(\mathbf{q}, \dot{\mathbf{q}}) + \ddot{\mathbf{q}}_d(t) + \boldsymbol{\nu}(\mathbf{q}, \dot{\mathbf{q}}, t)\right)$$

Computed Torque

$$\dot{\boldsymbol{\eta}} = \widehat{\mathbf{f}}(\mathbf{q}, \dot{\mathbf{q}}) - \dot{\mathbf{r}}(t) + \widehat{\mathbf{g}}(\mathbf{q})\mathbf{u}$$

Computed Torque (Feedback Linearization)

$$\mathbf{k}(\mathbf{q}, \dot{\mathbf{q}}, t) = \widehat{\mathbf{B}}^{-1}\widehat{\mathbf{D}}(\mathbf{q})\left((\widehat{\mathbf{D}}(\mathbf{q}))^{-1}\widehat{\mathbf{H}}(\mathbf{q}, \dot{\mathbf{q}}) + \ddot{\mathbf{q}}_d(t) + \boldsymbol{\nu}(\mathbf{q}, \dot{\mathbf{q}}, t)\right)$$

$$\dot{\boldsymbol{\eta}} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \boldsymbol{\eta} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} \boldsymbol{\nu}(\mathbf{q}, \dot{\mathbf{q}}, t)$$

Computed Torque

$$\dot{\boldsymbol{\eta}} = \widehat{\mathbf{f}}(\mathbf{q}, \dot{\mathbf{q}}) - \dot{\mathbf{r}}(t) + \widehat{\mathbf{g}}(\mathbf{q})\mathbf{u}$$

Computed Torque (Feedback Linearization)

$$\mathbf{k}(\mathbf{q}, \dot{\mathbf{q}}, t) = \widehat{\mathbf{B}}^{-1}\widehat{\mathbf{D}}(\mathbf{q}) \left((\widehat{\mathbf{D}}(\mathbf{q}))^{-1}\widehat{\mathbf{H}}(\mathbf{q}, \dot{\mathbf{q}}) + \ddot{\mathbf{q}}_d(t) + \boldsymbol{\nu}(\mathbf{q}, \dot{\mathbf{q}}, t) \right)$$

$$\dot{\boldsymbol{\eta}} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \boldsymbol{\eta} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} \boldsymbol{\nu}(\mathbf{q}, \dot{\mathbf{q}}, t)$$

Error PD

$$\boldsymbol{\nu}(\mathbf{q}, \dot{\mathbf{q}}, t) = -\begin{bmatrix} \mathbf{K}_p & \mathbf{K}_d \end{bmatrix} \boldsymbol{\eta}$$

Computed Torque

$$\dot{\boldsymbol{\eta}} = \widehat{\mathbf{f}}(\mathbf{q},\dot{\mathbf{q}}) - \dot{\mathbf{r}}(t) + \widehat{\mathbf{g}}(\mathbf{q})\mathbf{u}$$

Computed Torque (Feedback Linearization)

$$\mathbf{k}(\mathbf{q},\dot{\mathbf{q}},t) = \widehat{\mathbf{B}}^{-1}\widehat{\mathbf{D}}(\mathbf{q})\left((\widehat{\mathbf{D}}(\mathbf{q}))^{-1}\widehat{\mathbf{H}}(\mathbf{q},\dot{\mathbf{q}}) + \ddot{\mathbf{q}}_d(t) + \boldsymbol{\nu}(\mathbf{q},\dot{\mathbf{q}},t)\right)$$

$$\dot{\boldsymbol{\eta}} = \begin{bmatrix} \mathbf{0} & \mathbf{1} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}\boldsymbol{\eta} + \begin{bmatrix} \mathbf{0} \\ \mathbf{1} \end{bmatrix}\boldsymbol{\nu}(\mathbf{q},\dot{\mathbf{q}},t)$$

Error PD

$$\boldsymbol{\nu}(\mathbf{q},\dot{\mathbf{q}},t) = -\begin{bmatrix} \mathbf{K}_p & \mathbf{K}_d \end{bmatrix}\boldsymbol{\eta}$$

$$\dot{\boldsymbol{\eta}} = \underbrace{\begin{bmatrix} \mathbf{0} & \mathbf{1} \\ -\mathbf{K}_p & -\mathbf{K}_d \end{bmatrix}}_{\mathbf{A}_{cl}}\boldsymbol{\eta}$$

$$\mathfrak{Re}(\lambda_i(\mathbf{A}_{cl})) < 0$$
$$i = 1, \dots, n$$

Caltech

Andrew J. Taylor

Computed Torque

$$\dot{\boldsymbol{\eta}} = \widehat{\mathbf{f}}(\mathbf{q},\dot{\mathbf{q}}) - \dot{\mathbf{r}}(t) + \widehat{\mathbf{g}}(\mathbf{q})\mathbf{u}$$

Computed Torque (Feedback Linearization)

$$\mathbf{k}(\mathbf{q},\dot{\mathbf{q}},t) = \widehat{\mathbf{B}}^{-1}\widehat{\mathbf{D}}(\mathbf{q})\left((\widehat{\mathbf{D}}(\mathbf{q}))^{-1}\widehat{\mathbf{H}}(\mathbf{q},\dot{\mathbf{q}}) + \ddot{\mathbf{q}}_d(t) + \boldsymbol{\nu}(\mathbf{q},\dot{\mathbf{q}},t)\right)$$

$$\dot{\boldsymbol{\eta}} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}\boldsymbol{\eta} + \begin{bmatrix} 0 \\ 1 \end{bmatrix}\boldsymbol{\nu}(\mathbf{q},\dot{\mathbf{q}},t)$$

Error PD

$$\boldsymbol{\nu}(\mathbf{q},\dot{\mathbf{q}},t) = -\begin{bmatrix}\mathbf{K}_p & \mathbf{K}_d\end{bmatrix}\boldsymbol{\eta}$$

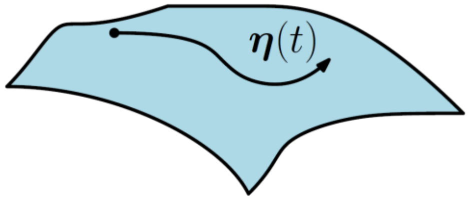

$$\dot{\boldsymbol{\eta}} = \underbrace{\begin{bmatrix} \mathbf{0} & \mathbf{1} \\ -\mathbf{K}_p & -\mathbf{K}_d \end{bmatrix}}_{\mathbf{A}_{cl}}\boldsymbol{\eta}$$

$$\mathfrak{Re}(\lambda_i(\mathbf{A}_{cl})) < 0$$
$$i = 1, \dots, n$$

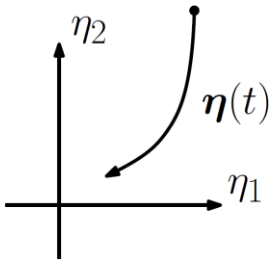

Caltech
Andrew J. Taylor

Lyapunov Functions

Continuous Time Lyapunov Equation

$$\mathbf{A}_{cl}^\top \mathbf{P} + \mathbf{P}\mathbf{A}_{cl} = -\mathbf{Q}$$
$$\mathbf{Q} \in \mathbb{S}_{++}$$

Lyapunov Functions

Continuous Time Lyapunov Equation

$$\mathbf{A}_{cl}^\top \mathbf{P} + \mathbf{P}\mathbf{A}_{cl} = -\mathbf{Q}$$

$$\mathbf{Q} \in \mathbb{S}_{++}$$

Quadratic Lyapunov Function

$$V(\boldsymbol{\eta}) = \boldsymbol{\eta}^\top \mathbf{P} \boldsymbol{\eta}$$

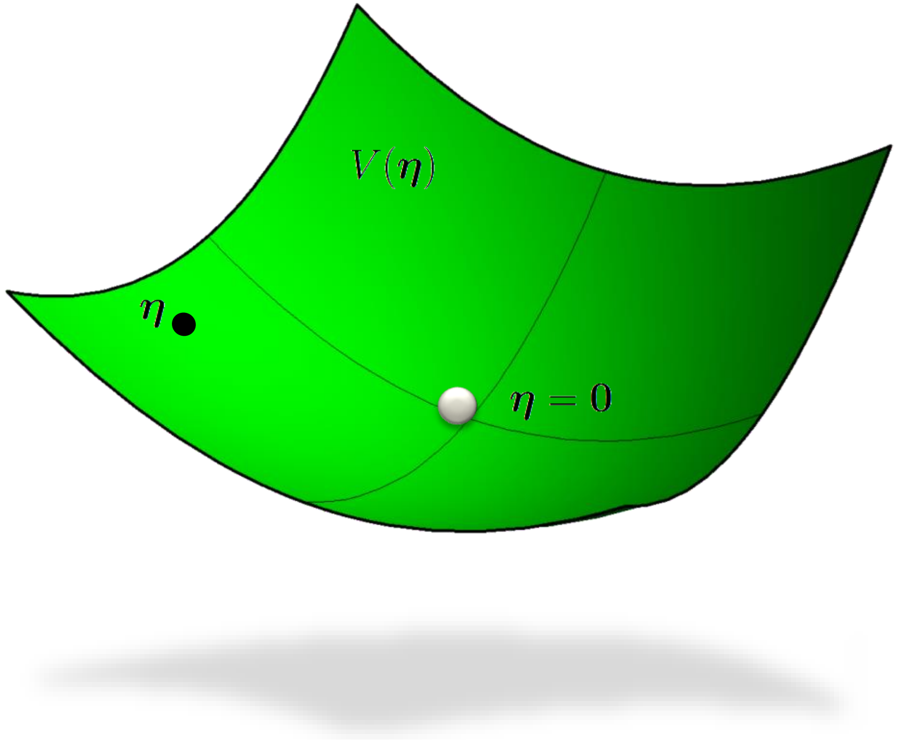

Caltech

Andrew J. Taylor

Lyapunov Functions

Continuous Time Lyapunov Equation

$$\mathbf{A}_{cl}^\top \mathbf{P} + \mathbf{P}\mathbf{A}_{cl} = -\mathbf{Q}$$
$$\mathbf{Q} \in \mathbb{S}_{++}$$

Quadratic Lyapunov Function

$$V(\boldsymbol{\eta}) = \boldsymbol{\eta}^\top \mathbf{P} \boldsymbol{\eta}$$

$$\lambda_{min}(\mathbf{P})\|\boldsymbol{\eta}\|_2^2 \leq V(\boldsymbol{\eta}) \leq \lambda_{max}(\mathbf{P})\|\boldsymbol{\eta}\|_2^2$$

$$\dot{V}(\boldsymbol{\eta}) \leq -\lambda_{min}(\mathbf{Q})\|\boldsymbol{\eta}\|_2^2$$

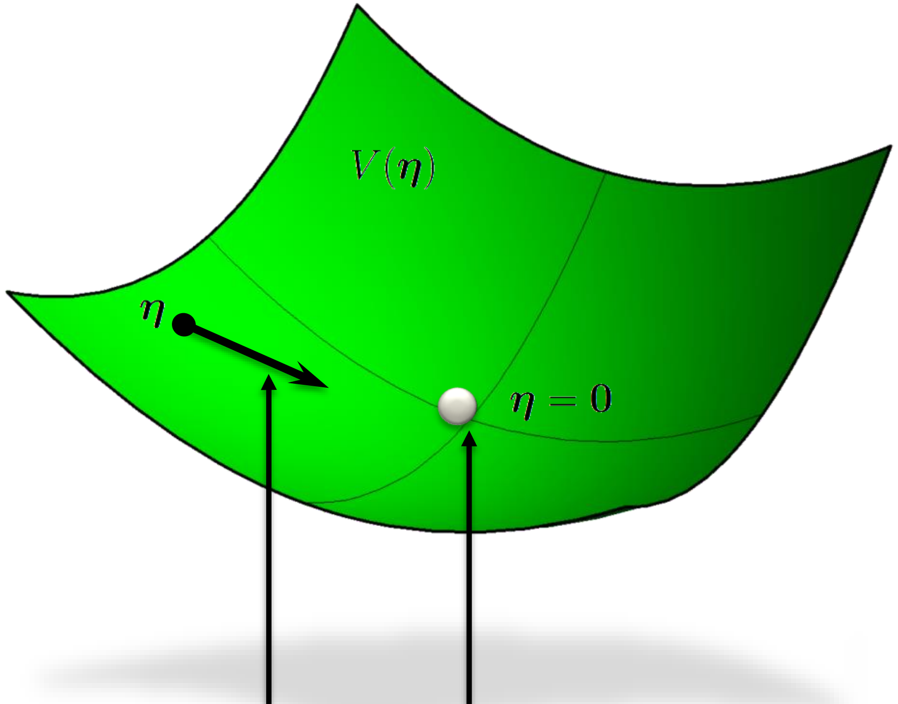

Lyapunov Functions

Continuous Time Lyapunov Equation

$$\mathbf{A}_{cl}^\top \mathbf{P} + \mathbf{P}\mathbf{A}_{cl} = -\mathbf{Q}$$

$$\mathbf{Q} \in \mathbb{S}_{++}$$

Quadratic Lyapunov Function

$$V(\boldsymbol{\eta}) = \boldsymbol{\eta}^\top \mathbf{P} \boldsymbol{\eta}$$

$$\lambda_{min}(\mathbf{P})\|\boldsymbol{\eta}\|_2^2 \leq V(\boldsymbol{\eta}) \leq \lambda_{max}(\mathbf{P})\|\boldsymbol{\eta}\|_2^2$$

$$\dot{V}(\boldsymbol{\eta}) \leq -\lambda_{min}(\mathbf{Q})\|\boldsymbol{\eta}\|_2^2$$

Certifies Exponential Stability

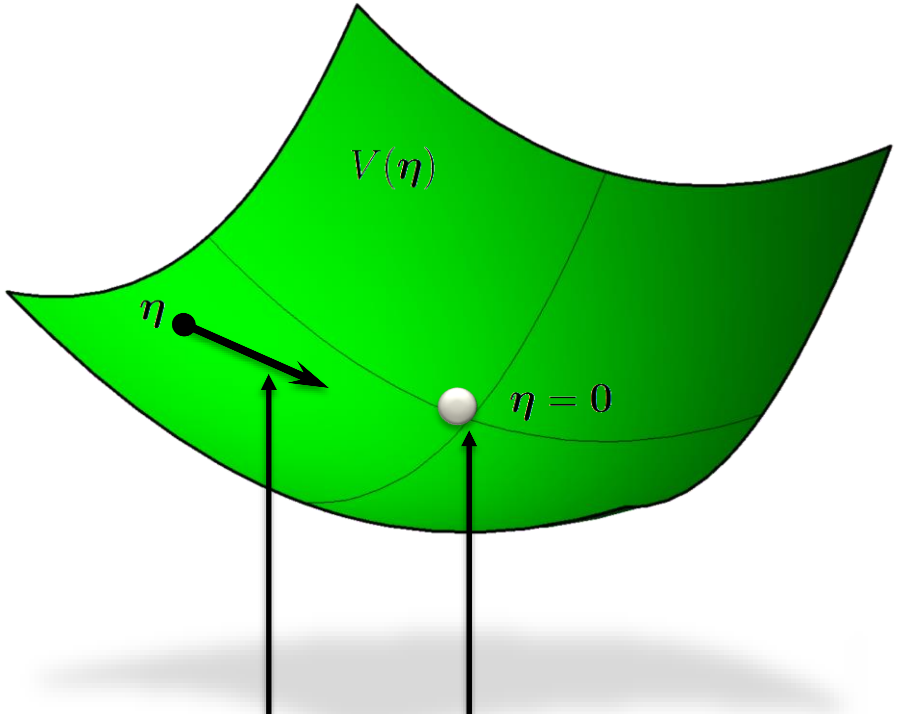

Caltech

Andrew J. Taylor

Control Lyapunov Functions (CLFs)

Computed Torque (Feedback Linearization)

$$\mathbf{k}(\mathbf{q},\dot{\mathbf{q}},t) = \widehat{\mathbf{B}}^{-1}\widehat{\mathbf{D}}(\mathbf{q})\left((\widehat{\mathbf{D}}(\mathbf{q}))^{-1}\widehat{\mathbf{H}}(\mathbf{q},\dot{\mathbf{q}}) + \ddot{\mathbf{q}}_d(t) + \boldsymbol{\nu}(\mathbf{q},\dot{\mathbf{q}},t)\right)$$

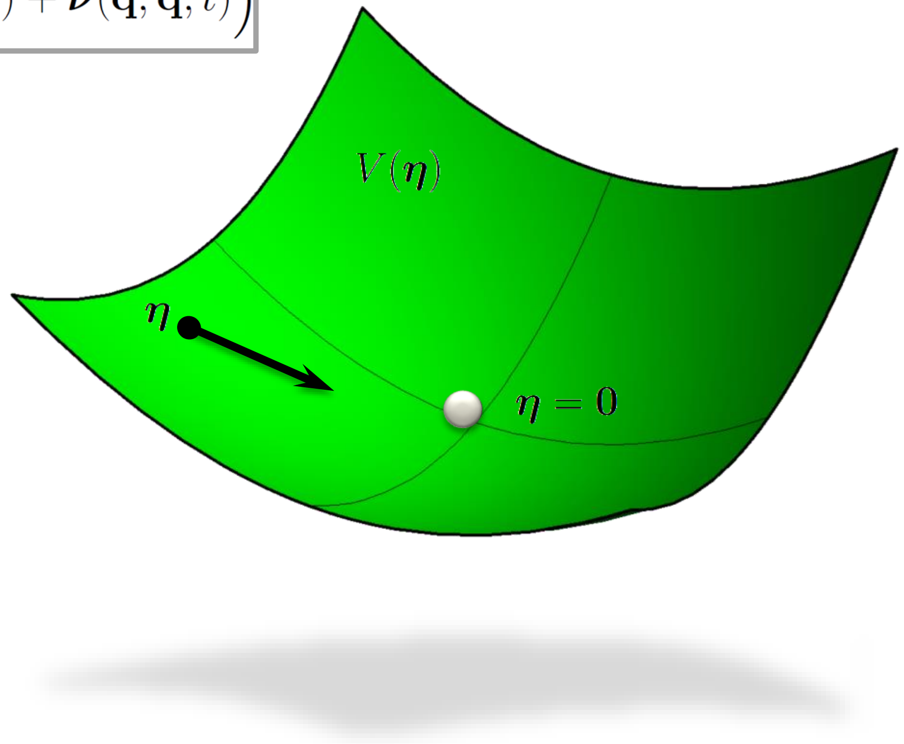

Control Lyapunov Functions (CLFs)

~~Computed Torque (Feedback Linearization)~~

~~$$\mathbf{k}(\mathbf{q},\dot{\mathbf{q}},t) = \widehat{\mathbf{B}}^{-1}\widehat{\mathbf{D}}(\mathbf{q})\left((\widehat{\mathbf{D}}(\mathbf{q}))^{-1}\widehat{\mathbf{H}}(\mathbf{q},\dot{\mathbf{q}}) + \ddot{\mathbf{q}}_d(t) + \boldsymbol{\nu}(\mathbf{q},\dot{\mathbf{q}},t)\right)$$~~

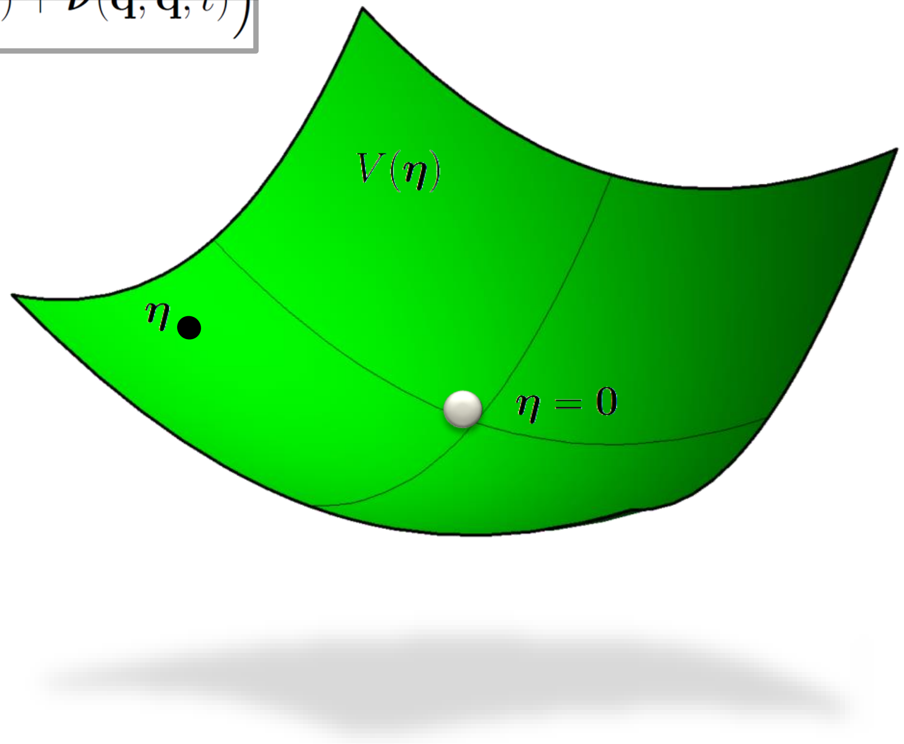

Control Lyapunov Functions (CLFs)

~~Computed Torque (Feedback Linearization)~~

~~$\mathbf{k}(\mathbf{q}, \dot{\mathbf{q}}, t) = \widehat{\mathbf{B}}^{-1}\widehat{\mathbf{D}}(\mathbf{q})\left((\widehat{\mathbf{D}}(\mathbf{q}))^{-1}\widehat{\mathbf{H}}(\mathbf{q}, \dot{\mathbf{q}}) + \ddot{\mathbf{q}}_d(t) + \boldsymbol{\nu}(\mathbf{q}, \dot{\mathbf{q}}, t)\right)$~~

Control Lyapunov Function

$$\inf_{\mathbf{u} \in \mathbb{R}^n} \dot{V}(\boldsymbol{\eta}, \mathbf{u}) \leq -\lambda_{min}(\mathbf{Q})\|\boldsymbol{\eta}\|_2^2$$

$$\dot{V}(\boldsymbol{\eta}, \mathbf{u}) = \frac{\partial V}{\partial \boldsymbol{\eta}}\left(\widehat{\mathbf{f}}(\mathbf{q}, \dot{\mathbf{q}}) - \dot{\mathbf{r}}(t) + \widehat{\mathbf{g}}(\mathbf{q})\mathbf{u}\right)$$

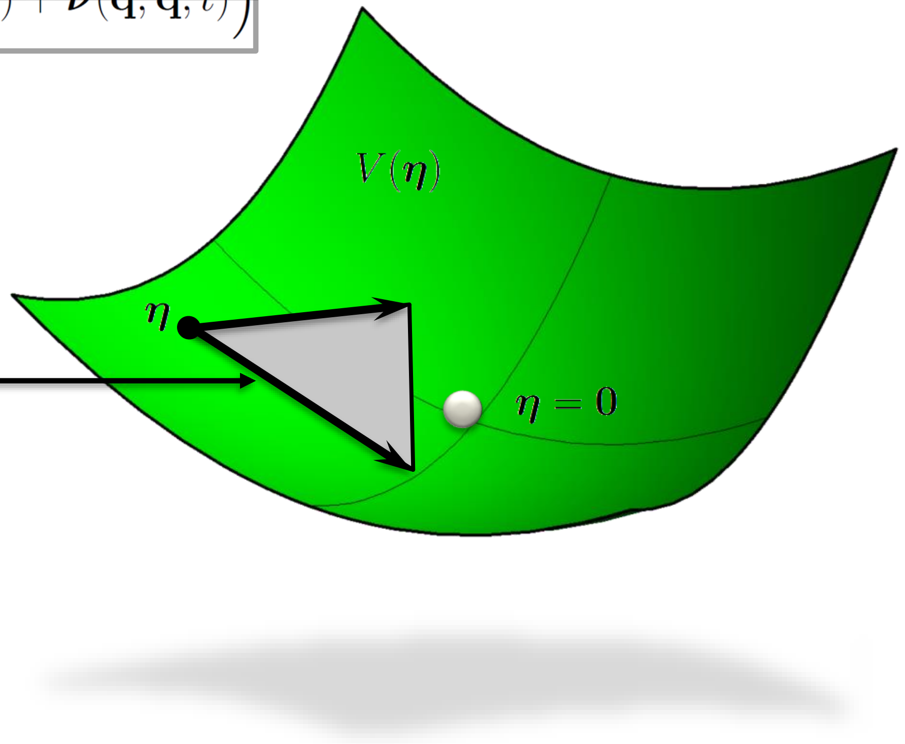

Caltech
Andrew J. Taylor

Control Lyapunov Functions (CLFs)

~~Computed Torque (Feedback Linearization)~~

$$\mathbf{k}(\mathbf{q},\dot{\mathbf{q}},t) = \widehat{\mathbf{B}}^{-1}\widehat{\mathbf{D}}(\mathbf{q})\left((\widehat{\mathbf{D}}(\mathbf{q}))^{-1}\widehat{\mathbf{H}}(\mathbf{q},\dot{\mathbf{q}}) + \ddot{\mathbf{q}}_d(t) + \boldsymbol{\nu}(\mathbf{q},\dot{\mathbf{q}},t)\right)$$

Control Lyapunov Function

$$\inf_{\mathbf{u}\in\mathbb{R}^n} \dot{V}(\boldsymbol{\eta},\mathbf{u}) \leq -\lambda_{min}(\mathbf{Q})\|\boldsymbol{\eta}\|_2^2$$

$$\dot{V}(\boldsymbol{\eta},\mathbf{u}) = \frac{\partial V}{\partial \boldsymbol{\eta}}\left(\widehat{\mathbf{f}}(\mathbf{q},\dot{\mathbf{q}}) - \dot{\mathbf{r}}(t) + \widehat{\mathbf{g}}(\mathbf{q})\mathbf{u}\right)$$

CLF Quadratic Program

$$\mathbf{k}(\mathbf{q},\dot{\mathbf{q}},t) = \underset{\mathbf{u}\in\mathbb{R}^n}{\operatorname{argmin}} \ \frac{1}{2}\|\mathbf{u}\|_2^2$$
$$\text{s.t. } \dot{V}(\boldsymbol{\eta},\mathbf{u}) \leq -\lambda_{min}(\mathbf{Q})\|\boldsymbol{\eta}\|_2^2$$

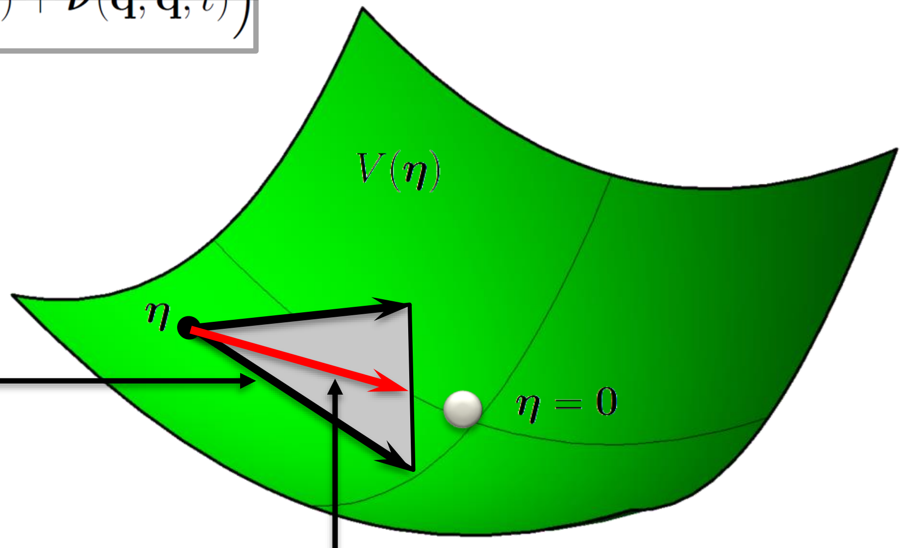

A. Ames, M. Powell, Towards the unification of locomotion and manipulation through control lyapunov functions and quadratic programs.

Control Lyapunov Functions (CLFs)

~~Computed Torque (Feedback Linearization)~~

$$\sout{\mathbf{k}(\mathbf{q},\dot{\mathbf{q}},t) = \widehat{\mathbf{B}}^{-1}\widehat{\mathbf{D}}(\mathbf{q})\left((\widehat{\mathbf{D}}(\mathbf{q}))^{-1}\widehat{\mathbf{H}}(\mathbf{q},\dot{\mathbf{q}}) + \ddot{\mathbf{q}}_d(t) + \boldsymbol{\nu}(\mathbf{q},\dot{\mathbf{q}},t)\right)}$$

Control Lyapunov Function

$$\inf_{\mathbf{u}\in\mathbb{R}^n} \dot{V}(\boldsymbol{\eta},\mathbf{u}) \leq -\lambda_{min}(\mathbf{Q})\|\boldsymbol{\eta}\|_2^2$$

$$\dot{V}(\boldsymbol{\eta},\mathbf{u}) = \frac{\partial V}{\partial \boldsymbol{\eta}}\left(\widehat{\mathbf{f}}(\mathbf{q},\dot{\mathbf{q}}) - \dot{\mathbf{r}}(t) + \widehat{\mathbf{g}}(\mathbf{q})\mathbf{u}\right)$$

CLF Quadratic Program

$$\mathbf{k}(\mathbf{q},\dot{\mathbf{q}},t) = \underset{\mathbf{u}\in\mathbb{R}^n}{\text{argmin}} \ \frac{1}{2}\|\mathbf{u}\|_2^2$$
$$\text{s.t.} \ \dot{V}(\boldsymbol{\eta},\mathbf{u}) \leq -\lambda_{min}(\mathbf{Q})\|\boldsymbol{\eta}\|_2^2$$

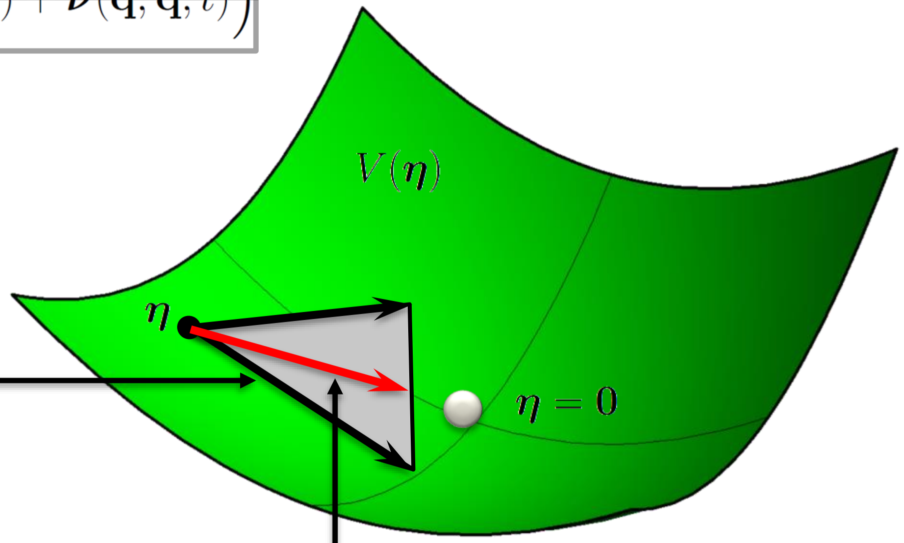

Enables Synthesis

A. Ames, M. Powell, Towards the unification of locomotion and manipulation through control lyapunov functions and quadratic programs.

Stabilizing Controllers?

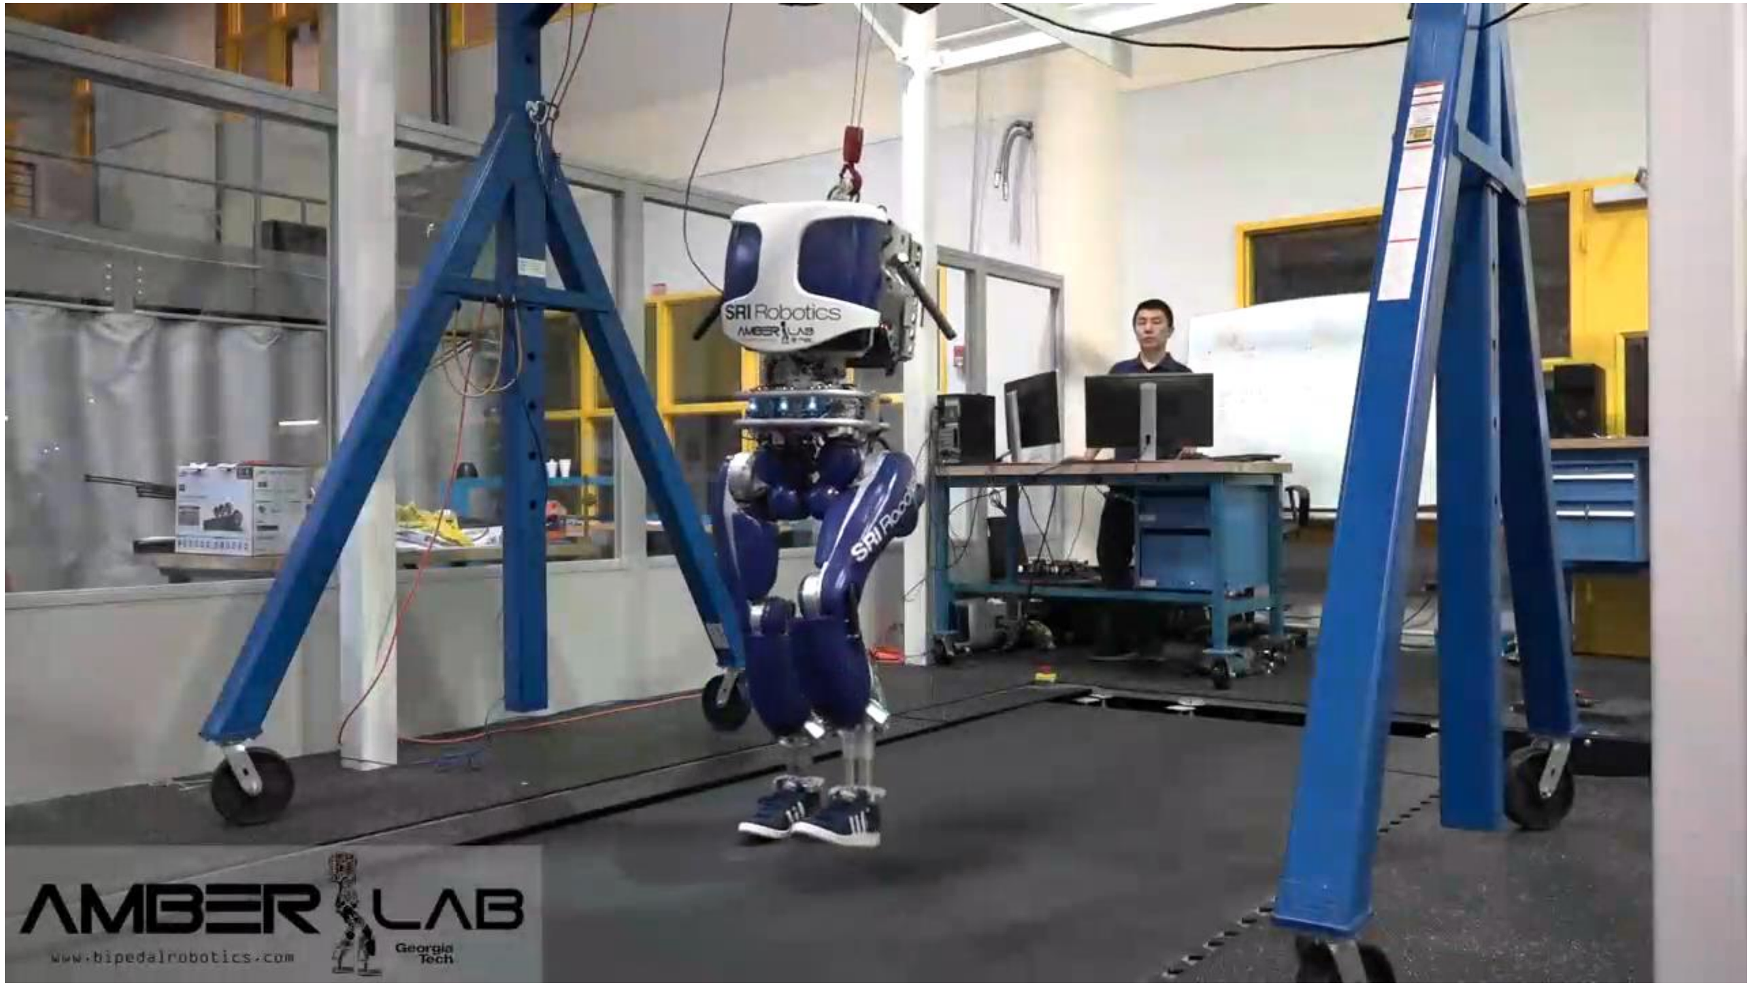

J. Reher, et al., Algorithmic foundations of realizing multi-contact locomotion on the humanoid robot DURUS

Stabilizing Controllers? (Not Quite)

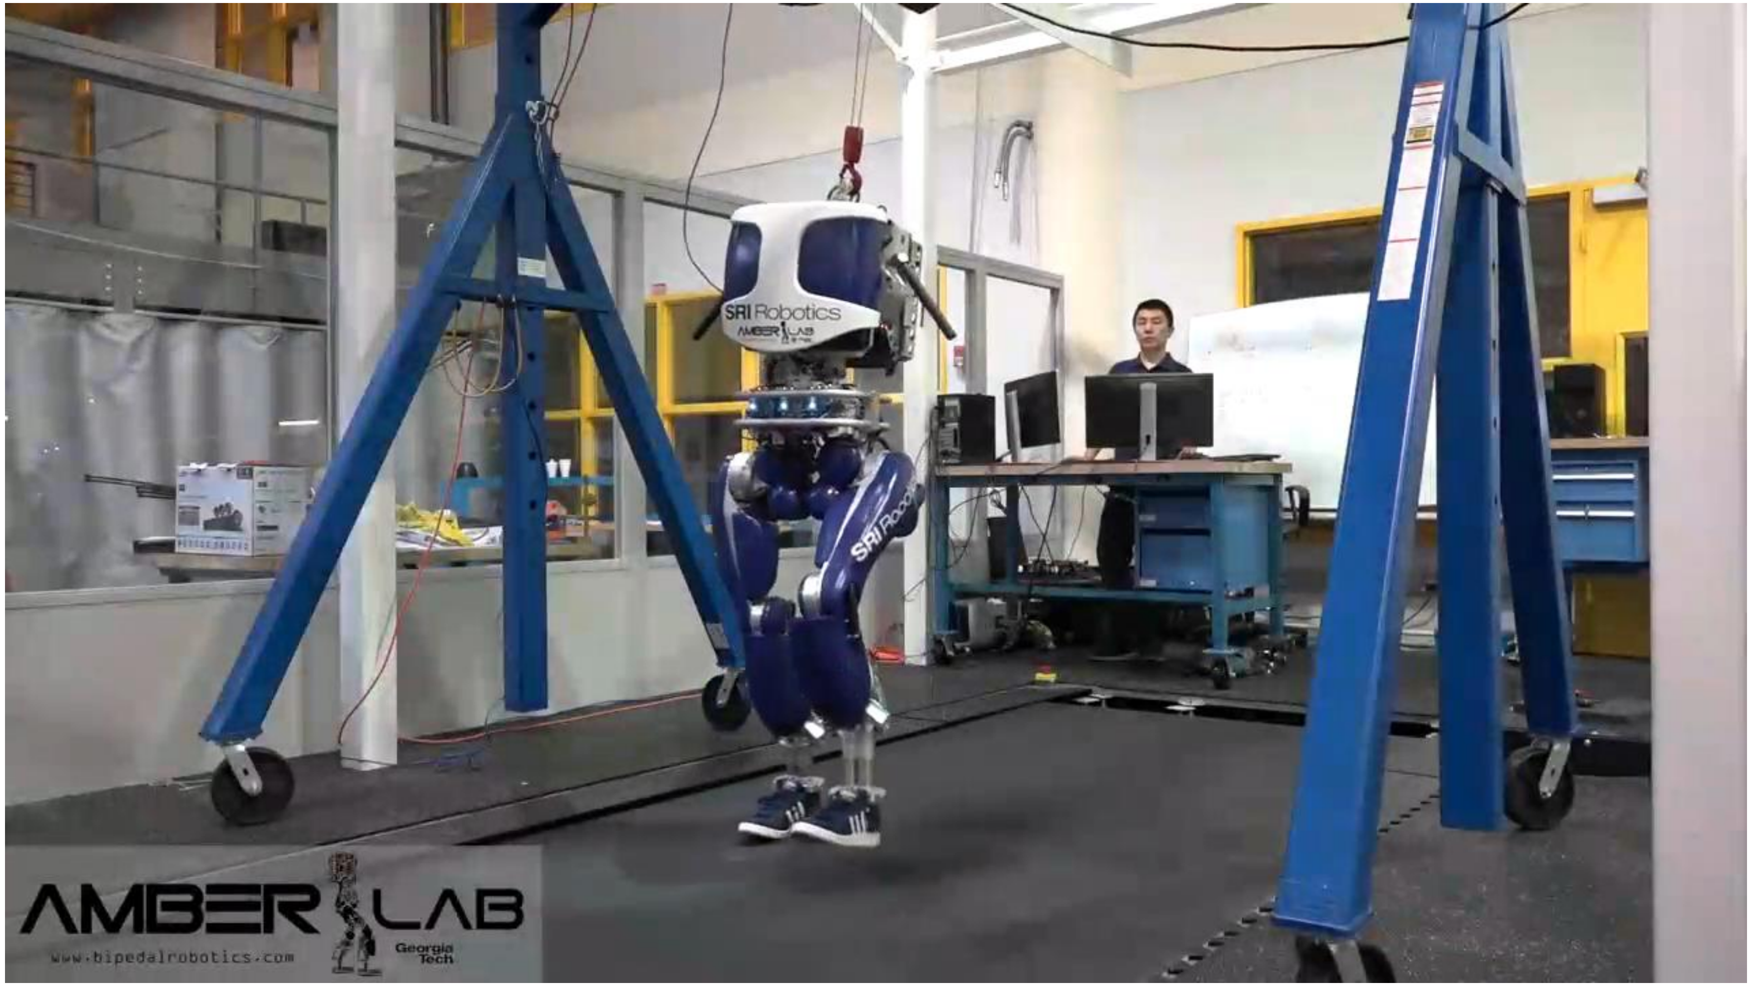

J. Reher, et al., Algorithmic foundations of realizing multi-contact locomotion on the humanoid robot DURUS

Uncertain Robotic Dynamics

Equations of Motion

$$\mathbf{D}(\mathbf{q})\ddot{\mathbf{q}} + \underbrace{\mathbf{C}(\mathbf{q},\dot{\mathbf{q}})\dot{\mathbf{q}} + \mathbf{G}(\mathbf{q})}_{\mathbf{H}(\mathbf{q},\dot{\mathbf{q}})} = \mathbf{B}\mathbf{u}$$

$$\mathbf{q} \in \mathcal{Q} \subseteq \mathbb{R}^n \qquad \mathbf{u} \in \mathbb{R}^n$$

True Dynamics

Physical Robot

Uncertain Robotic Dynamics

Equations of Motion

$$\mathbf{D}(\mathbf{q})\ddot{\mathbf{q}} + \underbrace{\mathbf{C}(\mathbf{q},\dot{\mathbf{q}})\dot{\mathbf{q}} + \mathbf{G}(\mathbf{q})}_{\mathbf{H}(\mathbf{q},\dot{\mathbf{q}})} = \mathbf{B}\mathbf{u}$$

$$\mathbf{q} \in \mathcal{Q} \subseteq \mathbb{R}^n \qquad \mathbf{u} \in \mathbb{R}^n$$

Methods

- Adaptive Control [1]
- System Identification [2]
- Machine Learning [3]
- High-gain control [4]

True Dynamics

Physical Robot

[1] M. Krstic, et al., Nonlinear Adaptive Control Design
[2] L. Ljung, System Identification
[3] J. Kober, et al., Reinforcement learning in robotics: A survey
[4] A. Ilchmann, et al., High-gain control without identification: a survey

Caltech
Andrew J. Taylor

Uncertain Robotic Dynamics

Equations of Motion

$$\mathbf{D}(\mathbf{q})\ddot{\mathbf{q}} + \underbrace{\mathbf{C}(\mathbf{q},\dot{\mathbf{q}})\dot{\mathbf{q}} + \mathbf{G}(\mathbf{q})}_{\mathbf{H}(\mathbf{q},\dot{\mathbf{q}})} = \mathbf{B}\mathbf{u}$$

$$\mathbf{q} \in \mathcal{Q} \subseteq \mathbb{R}^n \qquad \mathbf{u} \in \mathbb{R}^n$$

Methods

- Adaptive Control
- System Identification
- Machine Learning
- High-gain control

True Dynamics

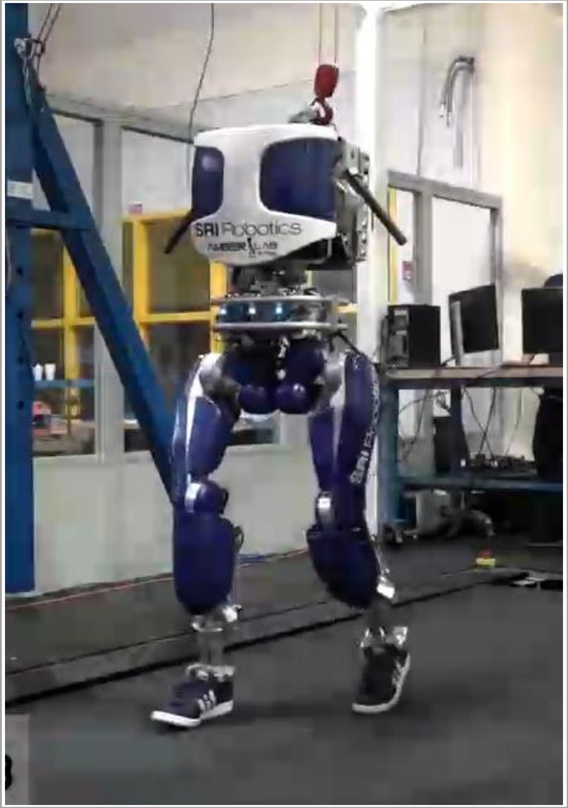

Physical Robot

Caltech
Andrew J. Taylor

Uncertain Robotic Dynamics

Equations of Motion

$$\mathbf{D}(\mathbf{q})\ddot{\mathbf{q}} + \underbrace{\mathbf{C}(\mathbf{q},\dot{\mathbf{q}})\dot{\mathbf{q}} + \mathbf{G}(\mathbf{q})}_{\mathbf{H}(\mathbf{q},\dot{\mathbf{q}})} = \mathbf{B}\mathbf{u}$$

$$\mathbf{q} \in \mathcal{Q} \subseteq \mathbb{R}^n \qquad \mathbf{u} \in \mathbb{R}^n$$

Assumptions*

- Time Invariant
- Deterministic
- Lipschitz Continuous
- $\mathrm{rank}(\mathbf{B}) = n$

True Dynamics

Physical Robot

*Under-actuated requires relative degree assumption.

Can we use the same CLF?

Can we use the same CLF?

$$\dot{\boldsymbol{\eta}} = \mathbf{f}(\mathbf{q}, \dot{\mathbf{q}}) - \dot{\mathbf{r}}(t) + \mathbf{g}(\mathbf{q})\mathbf{u}$$

Can we use the same CLF?

$$\dot{\boldsymbol{\eta}} = \mathbf{f}(\mathbf{q},\dot{\mathbf{q}}) - \dot{\mathbf{r}}(t) + \mathbf{g}(\mathbf{q})\mathbf{u}$$

Computed Torque (Feedback Linearization)

$$\mathbf{k}(\mathbf{q},\dot{\mathbf{q}},t) = \mathbf{B}^{-1}\mathbf{D}(\mathbf{q})\left((\mathbf{D}(\mathbf{q}))^{-1}\mathbf{H}(\mathbf{q},\dot{\mathbf{q}}) + \ddot{\mathbf{q}}_d(t) + \boldsymbol{\nu}(\mathbf{q},\dot{\mathbf{q}},t)\right)$$

$$\dot{\boldsymbol{\eta}} = \underbrace{\begin{bmatrix} \mathbf{0} & \mathbf{1} \\ -\mathbf{K}_p & -\mathbf{K}_d \end{bmatrix}}_{\mathbf{A}_{cl}} \boldsymbol{\eta}$$

$$\mathfrak{Re}(\lambda_i(\mathbf{A}_{cl})) < 0$$
$$i = 1, \dots, n$$

Can we use the same CLF? (We can!)

$$\dot{\boldsymbol{\eta}} = \mathbf{f}(\mathbf{q}, \dot{\mathbf{q}}) - \dot{\mathbf{r}}(t) + \mathbf{g}(\mathbf{q})\mathbf{u}$$

Computed Torque (Feedback Linearization)

$$\mathbf{k}(\mathbf{q}, \dot{\mathbf{q}}, t) = \mathbf{B}^{-1}\mathbf{D}(\mathbf{q})\left((\mathbf{D}(\mathbf{q}))^{-1}\mathbf{H}(\mathbf{q}, \dot{\mathbf{q}}) + \ddot{\mathbf{q}}_d(t) + \boldsymbol{\nu}(\mathbf{q}, \dot{\mathbf{q}}, t)\right)$$

$$\dot{\boldsymbol{\eta}} = \underbrace{\begin{bmatrix} \mathbf{0} & \mathbf{1} \\ -\mathbf{K}_p & -\mathbf{K}_d \end{bmatrix}}_{\mathbf{A}_{cl}} \boldsymbol{\eta}$$

$$\mathfrak{Re}(\lambda_i(\mathbf{A}_{cl})) < 0$$
$$i = 1, \dots, n$$

Control Lyapunov Function

$$\inf_{\mathbf{u} \in \mathbb{R}^n} \dot{V}(\boldsymbol{\eta}, \mathbf{u}) \leq -\lambda_{min}(\mathbf{Q})\|\boldsymbol{\eta}\|_2^2$$

$$\dot{V}(\boldsymbol{\eta}, \mathbf{u}) = \frac{\partial V}{\partial \boldsymbol{\eta}}(\mathbf{f}(\mathbf{q}, \dot{\mathbf{q}}) - \dot{\mathbf{r}}(t) + \mathbf{g}(\mathbf{q})\mathbf{u})$$

Can we use the same CLF? (We can!)

$$\dot{\boldsymbol{\eta}} = \mathbf{f}(\mathbf{q}, \dot{\mathbf{q}}) - \dot{\mathbf{r}}(t) + \mathbf{g}(\mathbf{q})\mathbf{u}$$

Computed Torque (Feedback Linearization)

$$\mathbf{k}(\mathbf{q}, \dot{\mathbf{q}}, t) = \mathbf{B}^{-1}\mathbf{D}(\mathbf{q})\left((\mathbf{D}(\mathbf{q}))^{-1}\mathbf{H}(\mathbf{q}, \dot{\mathbf{q}}) + \ddot{\mathbf{q}}_d(t) + \boldsymbol{\nu}(\mathbf{q}, \dot{\mathbf{q}}, t)\right)$$

$$\dot{\boldsymbol{\eta}} = \underbrace{\begin{bmatrix} \mathbf{0} & \mathbf{1} \\ -\mathbf{K}_p & -\mathbf{K}_d \end{bmatrix}}_{\mathbf{A}_{cl}} \boldsymbol{\eta}$$

$$\mathfrak{Re}(\lambda_i(\mathbf{A}_{cl})) < 0$$
$$i = 1, \dots, n$$

Control Lyapunov Function

$$\inf_{\mathbf{u} \in \mathbb{R}^n} \dot{V}(\boldsymbol{\eta}, \mathbf{u}) \leq -\lambda_{min}(\mathbf{Q})\|\boldsymbol{\eta}\|_2^2$$

$$\dot{V}(\boldsymbol{\eta}, \mathbf{u}) = \frac{\partial V}{\partial \boldsymbol{\eta}}(\mathbf{f}(\mathbf{q}, \dot{\mathbf{q}}) - \dot{\mathbf{r}}(t) + \mathbf{g}(\mathbf{q})\mathbf{u})$$

Don’t know how to choose input

CLF Derivative Uncertainty

$$\dot{\boldsymbol{\eta}} = \mathbf{f}(\mathbf{q}, \dot{\mathbf{q}}) - \dot{\mathbf{r}}(t) + \mathbf{g}(\mathbf{q})\mathbf{u}$$

CLF Derivative Uncertainty

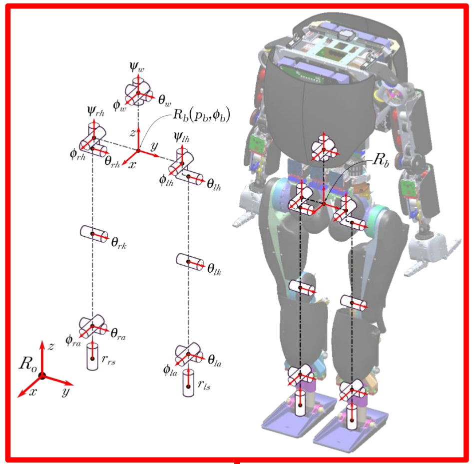

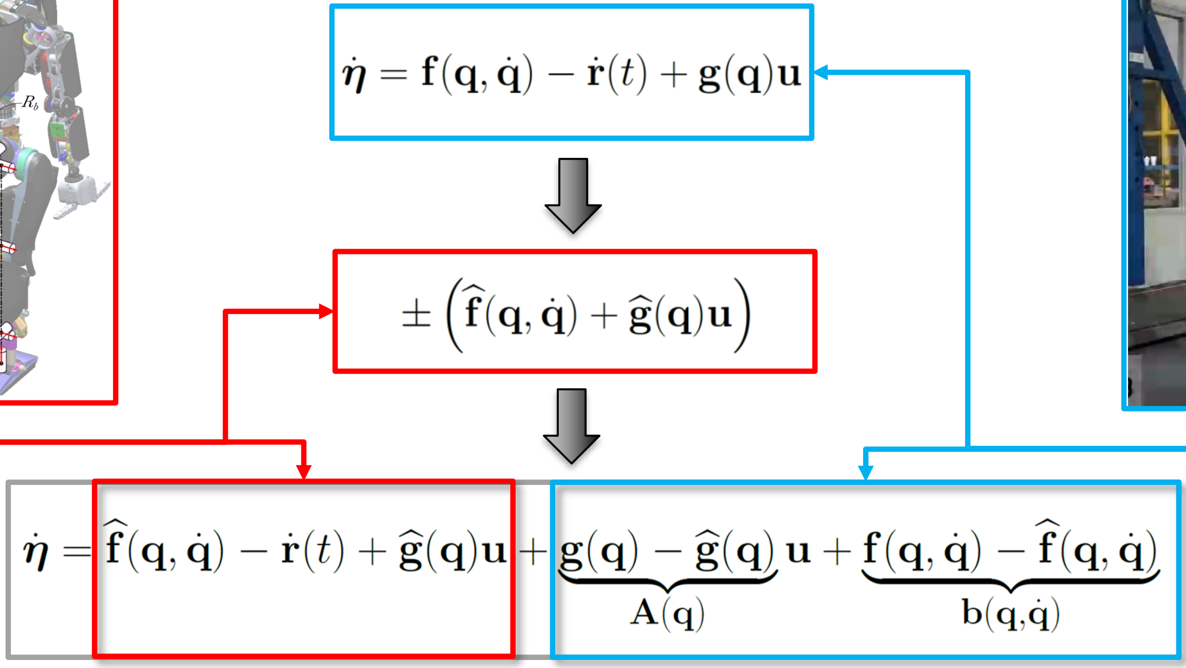

CLF Derivative Uncertainty

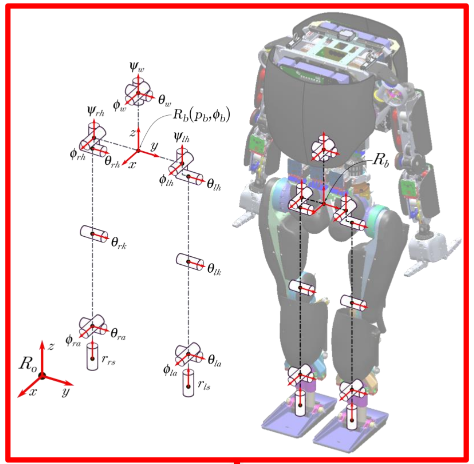

$$\dot{\boldsymbol{\eta}} = \mathbf{f}(\mathbf{q},\dot{\mathbf{q}}) - \dot{\mathbf{r}}(t) + \mathbf{g}(\mathbf{q})\mathbf{u}$$

$$\pm\left(\widehat{\mathbf{f}}(\mathbf{q},\dot{\mathbf{q}}) + \widehat{\mathbf{g}}(\mathbf{q})\mathbf{u}\right)$$

$$\dot{\boldsymbol{\eta}} = \widehat{\mathbf{f}}(\mathbf{q},\dot{\mathbf{q}}) - \dot{\mathbf{r}}(t) + \widehat{\mathbf{g}}(\mathbf{q})\mathbf{u} + \underbrace{\mathbf{g}(\mathbf{q}) - \widehat{\mathbf{g}}(\mathbf{q})}_{\mathbf{A}(\mathbf{q})}\mathbf{u} + \underbrace{\mathbf{f}(\mathbf{q},\dot{\mathbf{q}}) - \widehat{\mathbf{f}}(\mathbf{q},\dot{\mathbf{q}})}_{\mathbf{b}(\mathbf{q},\dot{\mathbf{q}})}$$

$$\dot{V}(\boldsymbol{\eta},\mathbf{u}) = \underbrace{\frac{\partial V}{\partial \boldsymbol{\eta}}\left(\widehat{\mathbf{f}}(\mathbf{q},\dot{\mathbf{q}}) - \dot{\mathbf{r}}(t) + \widehat{\mathbf{g}}(\mathbf{q})\mathbf{u}\right)}_{\widehat{\dot{V}}(\boldsymbol{\eta},\mathbf{u})} + \underbrace{\frac{\partial V}{\partial \boldsymbol{\eta}}\mathbf{A}(\mathbf{q})}_{\mathbf{a}(\boldsymbol{\eta},\mathbf{q})^\top}\mathbf{u} + \underbrace{\frac{\partial V}{\partial \boldsymbol{\eta}}\mathbf{b}(\mathbf{q},\dot{\mathbf{q}})}_{b(\boldsymbol{\eta},\mathbf{q},\dot{\mathbf{q}})}$$

Caltech
Andrew J. Taylor

CLF Derivative Uncertainty

Learn the error dynamics

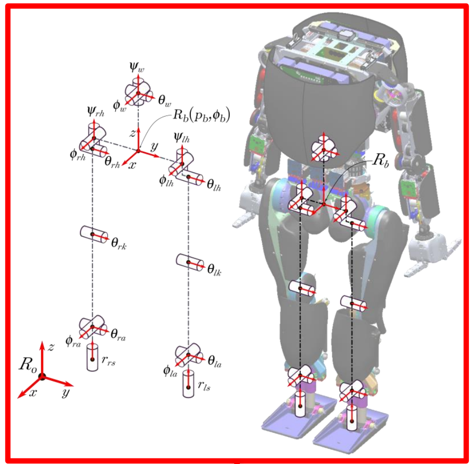

$$\dot{\boldsymbol{\eta}} = \mathbf{f}(\mathbf{q},\dot{\mathbf{q}}) - \dot{\mathbf{r}}(t) + \mathbf{g}(\mathbf{q})\mathbf{u}$$

$$\pm\left(\widehat{\mathbf{f}}(\mathbf{q},\dot{\mathbf{q}}) + \widehat{\mathbf{g}}(\mathbf{q})\mathbf{u}\right)$$

$$\dot{\boldsymbol{\eta}} = \widehat{\mathbf{f}}(\mathbf{q},\dot{\mathbf{q}}) - \dot{\mathbf{r}}(t) + \widehat{\mathbf{g}}(\mathbf{q})\mathbf{u} + \underbrace{\mathbf{g}(\mathbf{q}) - \widehat{\mathbf{g}}(\mathbf{q})}_{\mathbf{A}(\mathbf{q})}\mathbf{u} + \underbrace{\mathbf{f}(\mathbf{q},\dot{\mathbf{q}}) - \widehat{\mathbf{f}}(\mathbf{q},\dot{\mathbf{q}})}_{\mathbf{b}(\mathbf{q},\dot{\mathbf{q}})}$$

$$\dot{V}(\boldsymbol{\eta},\mathbf{u}) = \underbrace{\frac{\partial V}{\partial \boldsymbol{\eta}}\left(\widehat{\mathbf{f}}(\mathbf{q},\dot{\mathbf{q}}) - \dot{\mathbf{r}}(t) + \widehat{\mathbf{g}}(\mathbf{q})\mathbf{u}\right)}_{\widehat{\dot{V}}(\boldsymbol{\eta},\mathbf{u})} + \underbrace{\frac{\partial V}{\partial \boldsymbol{\eta}}\mathbf{A}(\mathbf{q})}_{\mathbf{a}(\boldsymbol{\eta},\mathbf{q})^\top}\mathbf{u} + \underbrace{\frac{\partial V}{\partial \boldsymbol{\eta}}\mathbf{b}(\mathbf{q},\dot{\mathbf{q}})}_{b(\boldsymbol{\eta},\mathbf{q},\dot{\mathbf{q}})}$$

Caltech

Andrew J. Taylor

CLF Derivative Uncertainty

Learn the residual error dynamics

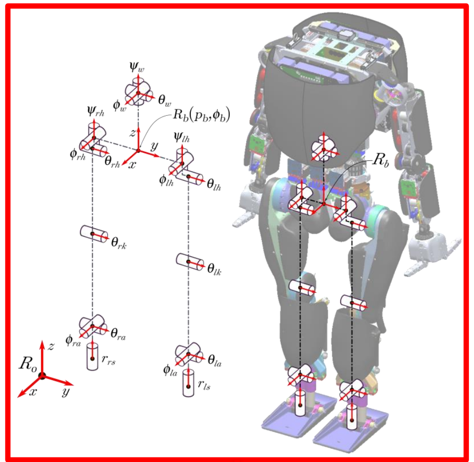

$$\dot{\boldsymbol{\eta}} = \mathbf{f}(\mathbf{q},\dot{\mathbf{q}}) - \dot{\mathbf{r}}(t) + \mathbf{g}(\mathbf{q})\mathbf{u}$$

$$\pm\left(\widehat{\mathbf{f}}(\mathbf{q},\dot{\mathbf{q}}) + \widehat{\mathbf{g}}(\mathbf{q})\mathbf{u}\right)$$

$$\dot{\boldsymbol{\eta}} = \widehat{\mathbf{f}}(\mathbf{q},\dot{\mathbf{q}}) - \dot{\mathbf{r}}(t) + \widehat{\mathbf{g}}(\mathbf{q})\mathbf{u} + \underbrace{\mathbf{g}(\mathbf{q}) - \widehat{\mathbf{g}}(\mathbf{q})}_{\mathbf{A}(\mathbf{q})}\mathbf{u} + \underbrace{\mathbf{f}(\mathbf{q},\dot{\mathbf{q}}) - \widehat{\mathbf{f}}(\mathbf{q},\dot{\mathbf{q}})}_{\mathbf{b}(\mathbf{q},\dot{\mathbf{q}})}$$

$$\dot{V}(\boldsymbol{\eta},\mathbf{u}) = \underbrace{\frac{\partial V}{\partial \boldsymbol{\eta}}\left(\widehat{\mathbf{f}}(\mathbf{q},\dot{\mathbf{q}}) - \dot{\mathbf{r}}(t) + \widehat{\mathbf{g}}(\mathbf{q})\mathbf{u}\right)}_{\widehat{\dot{V}}(\boldsymbol{\eta},\mathbf{u})} + \underbrace{\frac{\partial V}{\partial \boldsymbol{\eta}}\mathbf{A}(\mathbf{q})}_{\mathbf{a}(\boldsymbol{\eta},\mathbf{q})^\top}\mathbf{u} + \underbrace{\frac{\partial V}{\partial \boldsymbol{\eta}}\mathbf{b}(\mathbf{q},\dot{\mathbf{q}})}_{b(\boldsymbol{\eta},\mathbf{q},\dot{\mathbf{q}})}$$

Caltech

Andrew J. Taylor

CLF Derivative Uncertainty

Learn the residual CLF dynamics

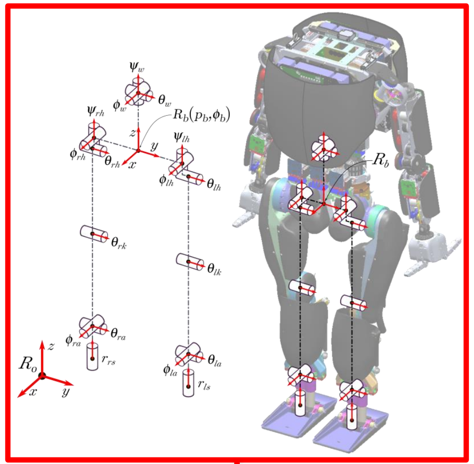

$$\dot{\boldsymbol{\eta}} = \mathbf{f}(\mathbf{q},\dot{\mathbf{q}}) - \dot{\mathbf{r}}(t) + \mathbf{g}(\mathbf{q})\mathbf{u}$$

$$\pm\left(\widehat{\mathbf{f}}(\mathbf{q},\dot{\mathbf{q}}) + \widehat{\mathbf{g}}(\mathbf{q})\mathbf{u}\right)$$

$$\dot{\boldsymbol{\eta}} = \widehat{\mathbf{f}}(\mathbf{q},\dot{\mathbf{q}}) - \dot{\mathbf{r}}(t) + \widehat{\mathbf{g}}(\mathbf{q})\mathbf{u} + \underbrace{\left(\mathbf{g}(\mathbf{q}) - \widehat{\mathbf{g}}(\mathbf{q})\right)}_{\mathbf{A}(\mathbf{q})}\mathbf{u} + \underbrace{\mathbf{f}(\mathbf{q},\dot{\mathbf{q}}) - \widehat{\mathbf{f}}(\mathbf{q},\dot{\mathbf{q}})}_{\mathbf{b}(\mathbf{q},\dot{\mathbf{q}})}$$

$$\dot{V}(\boldsymbol{\eta},\mathbf{u}) = \underbrace{\frac{\partial V}{\partial \boldsymbol{\eta}}\left(\widehat{\mathbf{f}}(\mathbf{q},\dot{\mathbf{q}}) - \dot{\mathbf{r}}(t) + \widehat{\mathbf{g}}(\mathbf{q})\mathbf{u}\right)}_{\widehat{\dot{V}}(\boldsymbol{\eta},\mathbf{u})} + \underbrace{\frac{\partial V}{\partial \boldsymbol{\eta}}\mathbf{A}(\mathbf{q})}_{\mathbf{a}(\boldsymbol{\eta},\mathbf{q})^\top}\mathbf{u} + \underbrace{\frac{\partial V}{\partial \boldsymbol{\eta}}\mathbf{b}(\mathbf{q},\dot{\mathbf{q}})}_{b(\boldsymbol{\eta},\mathbf{q},\dot{\mathbf{q}})}$$

Caltech

Andrew J. Taylor

Learning Control Lyapunov Functions

CLF Derivative Estimator

$$\dot{V}(\boldsymbol{\eta}, \mathbf{u}) \approx \widehat{\dot{V}}(\boldsymbol{\eta}, \mathbf{u}) + \widehat{\mathbf{a}}(\boldsymbol{\eta}, \mathbf{q})^\top \mathbf{u} + \widehat{b}(\boldsymbol{\eta}, \mathbf{q}, \dot{\mathbf{q}})$$

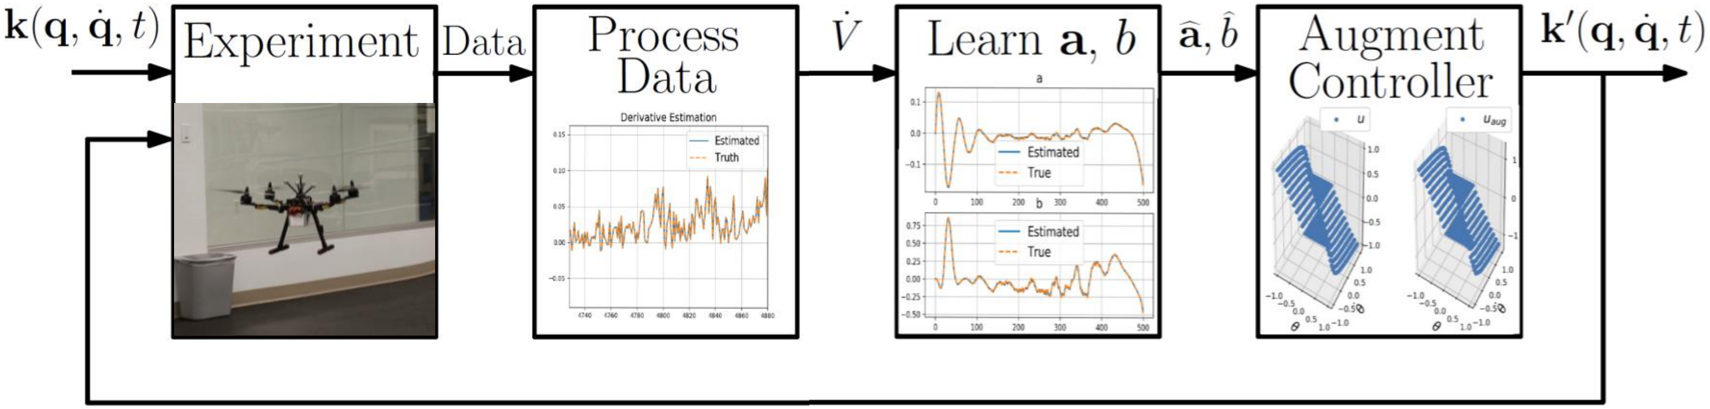

Learning Control Lyapunov Functions

Raw Data

$$D_r = \{(\mathbf{q}_i, \dot{\mathbf{q}}_i, t_i, \mathbf{u}_i)\}_{i=1}^{N}$$

Processed Data

$$D_p = \{((\mathbf{q}_i, \dot{\mathbf{q}}_i, \boldsymbol{\eta}_i, \mathbf{u}_i), \dot{V}_i)\}_{i=2}^{N-1}$$

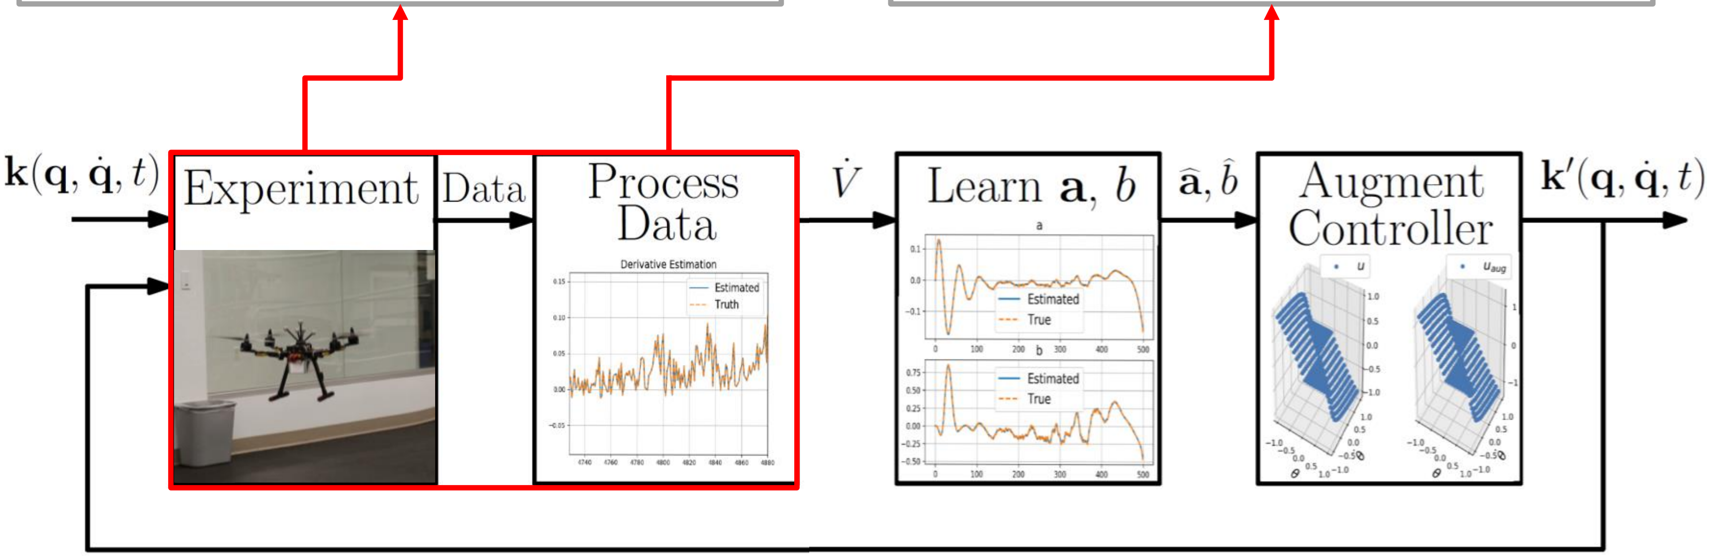

Learning Control Lyapunov Functions

Model Classes

$$\widehat{\mathbf{a}} \in \mathcal{H}_{\mathbf{a}} : \mathbb{R}^{2n} \times \mathcal{Q} \to \mathbb{R}^{n}$$
$$\widehat{b} \in \mathcal{H}_{b} : \mathbb{R}^{2n} \times \mathcal{Q} \times \mathbb{R}^{n} \to \mathbb{R}$$

Supervised Learning

$$\inf_{\substack{\widehat{\mathbf{a}} \in \mathcal{H}_{\mathbf{a}} \\ \widehat{b} \in \mathcal{H}_{b}}} \frac{1}{N-2} \sum_{i=2}^{N-1} \mathcal{L}\left(\widehat{\dot{V}}(\boldsymbol{\eta}_i, \mathbf{u}_i) + \widehat{\mathbf{a}}(\boldsymbol{\eta}_i, \mathbf{q}_i)^\top \mathbf{u}_i + \widehat{b}(\boldsymbol{\eta}_i, \mathbf{q}_i, \dot{\mathbf{q}}_i), \dot{V}_i \right)$$

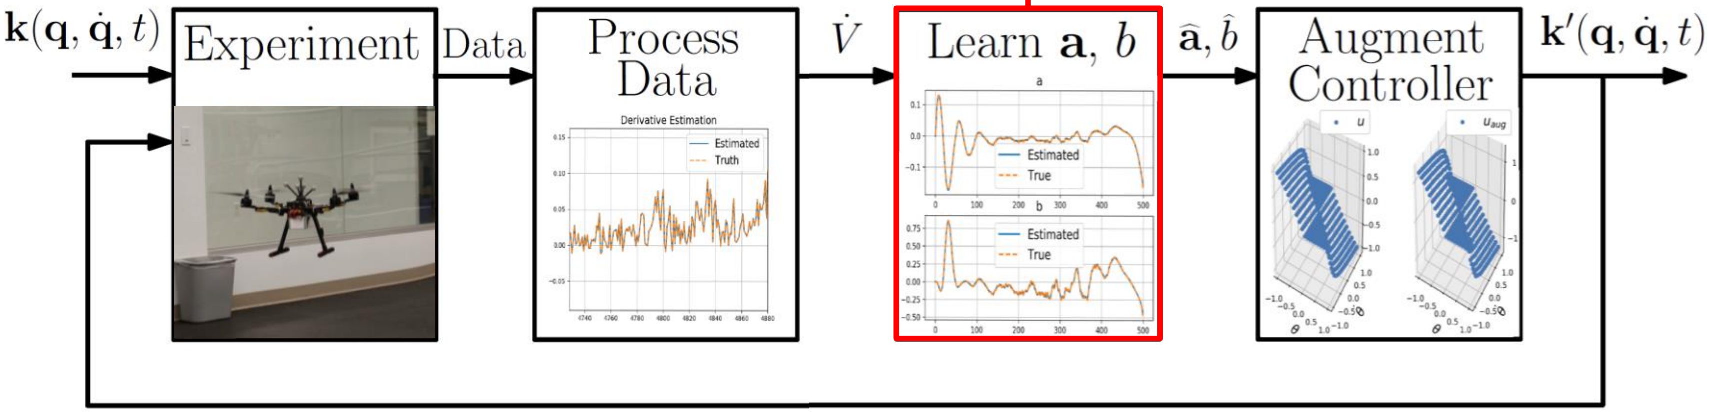

Andrew J. Taylor

Learning Control Lyapunov Functions

Learning-Augmented Controller

$$\mathbf{k}'(\mathbf{q}, \dot{\mathbf{q}}, t) = \underset{\mathbf{u}' \in \mathbb{R}^n}{\operatorname{argmin}} \; \|(\mathbf{k}(\mathbf{q}, \dot{\mathbf{q}}, t) + \mathbf{u}')\|_2^2$$

$$\text{s.t.} \quad \widehat{\dot{V}}(\boldsymbol{\eta}, \mathbf{k}(\mathbf{q}, \dot{\mathbf{q}}, t) + \mathbf{u}') + \hat{\mathbf{a}}(\boldsymbol{\eta}, \mathbf{q})^\top (\mathbf{k}(\mathbf{q}, \dot{\mathbf{q}}, t) + \mathbf{u}') + b(\boldsymbol{\eta}, \mathbf{q}, \dot{\mathbf{q}}) \leq -\lambda_{min}(\mathbf{Q}) \|\boldsymbol{\eta}\|_2^2$$

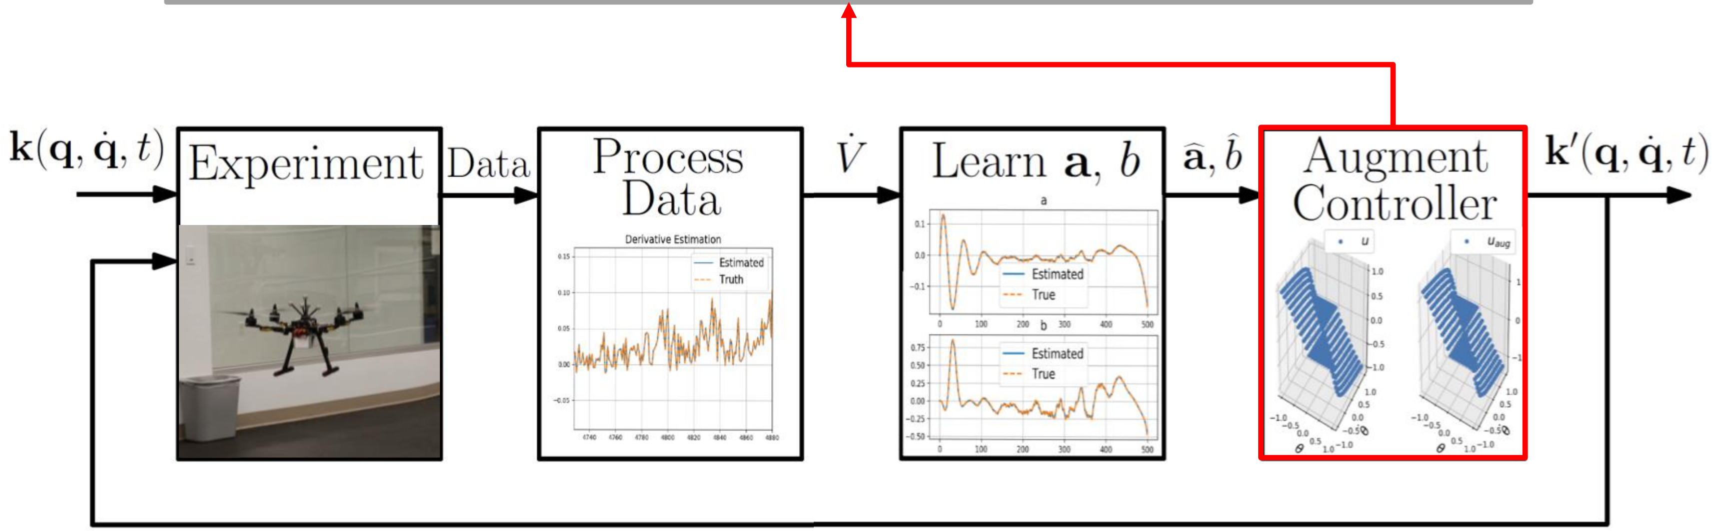

Episodic Learning

Algorithm 1 Dataset Aggregation for Control Lyapunov Functions (DaCLyF)

Require: Control Lyapunov Function V, derivative estimate $\widehat{\dot{V}}_0$, model classes $\mathcal{H}_{\mathbf{a}}$ and $\mathcal{H}_b$, loss function $\mathcal{L}$, set of initial configurations $\mathcal{Q}_0$, nominal state-feedback controller $\mathbf{k}_0$, number of experiments T, sequence of trust coefficients $0 \leq w_1 \leq \cdots \leq w_T \leq 1$

$D = \emptyset$ ▷ Initialize data set
for $k = 1, \dots, T$ **do**
  $(\mathbf{q}_0, \mathbf{0}) \leftarrow \text{sample}(\mathcal{Q}_0 \times \{\mathbf{0}\})$ ▷ Get initial condition
  $D_k \leftarrow \text{experiment}((\mathbf{q}_0, \mathbf{0}), \mathbf{k}_{k-1})$ ▷ Run experiment
  $D \leftarrow D \cup D_k$ ▷ Aggregate data set
  $\widehat{\mathbf{a}}, \widehat{b} \leftarrow \text{ERM}(\mathcal{H}_{\mathbf{a}}, \mathcal{H}_b, \mathcal{L}, D, \widehat{\dot{V}}_0)$ ▷ Fit estimators
  $\widehat{\dot{V}}_k \leftarrow \widehat{\dot{V}}_0 + \widehat{\mathbf{a}}^\top \mathbf{u} + \widehat{b}$ ▷ Update derivative estimator
  $\mathbf{k}_k \leftarrow \mathbf{k}_0 + w_k \cdot \text{augment}(\mathbf{k}_0, \widehat{\dot{V}}_k)$ ▷ Update controller
end for
return $\widehat{\dot{V}}_T, \mathbf{u}_T$

S. Ross, et al., A reduction of imitation learning and structured prediction to no-regret online learning.

Segway System

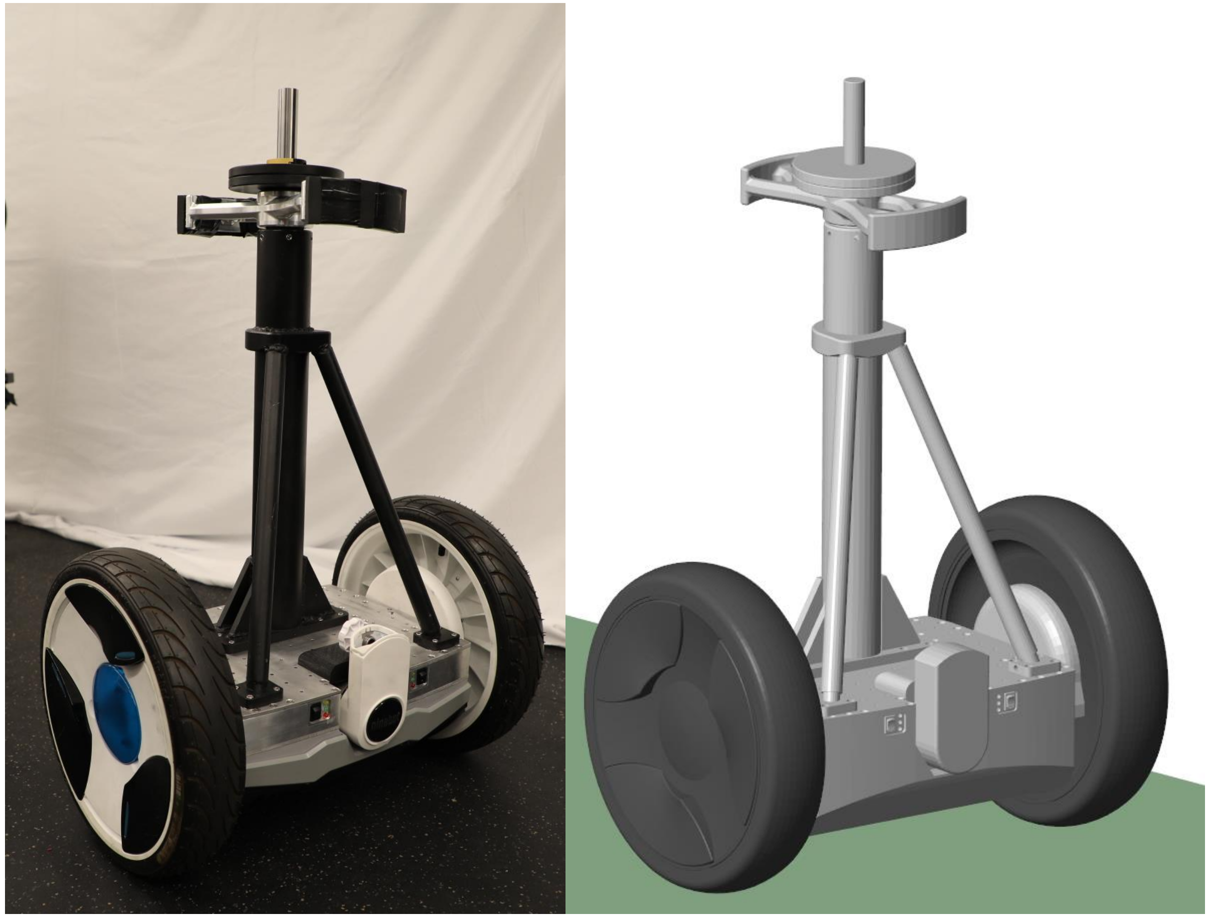

Segway Simulation

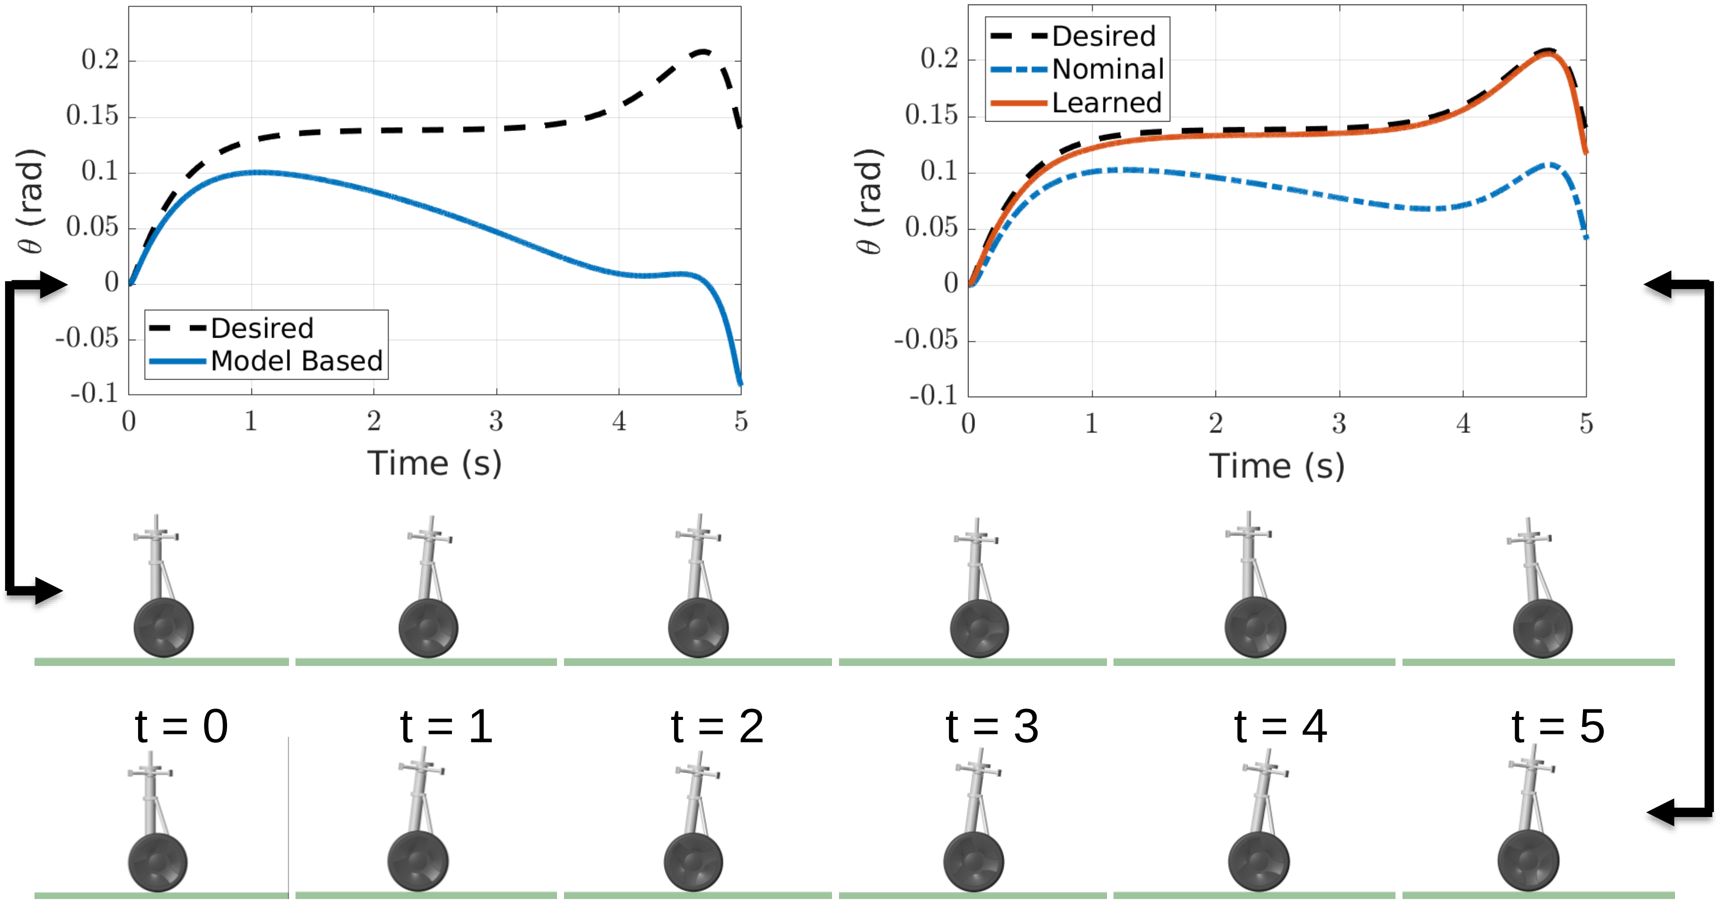

Caltech
Andrew J. Taylor

Segway Simulation

Dataset Aggregation
for
Control Lyapunov Functions
Segway Simulation

(additional PD stabilization
upright, x1.5 Speed)

Future Work

- Compare learning at different levels of dynamics
  - Evaluate low-dimensional learning against lifted methods (RKHS, Koopman Operators)
  - Explore theoretical/empirical implications of low-dimensional learning on sample-complexity
- Implement episodic learning framework on Segway hardware
  - Understand sensitivity of algorithm to noise / filtering
  - Certify validity of assumptions on dynamic uncertainty
- Study convergence of models in episodic framework
  - Understand need for structured exploration in data acquisition
  - Develop trust coefficients for estimators across episodes

Thank You!

Episodic Learning with Control Lyapunov Functions for Uncertain Robotic Systems

Andrew Taylor   Victor Dorobantu   Hoang Le
Yisong Yue   Aaron D. Ames

Projection -to-State-Stability (PSS)

- Appearing at CDC 2019:

$$\dot{V}(\boldsymbol{\eta},\mathbf{u}) = \underbrace{\frac{\partial V}{\partial \boldsymbol{\eta}}\left(\widehat{\mathbf{f}}(\mathbf{q},\dot{\mathbf{q}}) - \dot{\mathbf{r}}(t) + \widehat{\mathbf{g}}(\mathbf{q})\mathbf{u}\right)}_{\widehat{\dot{V}}(\boldsymbol{\eta},\mathbf{u})} + \underbrace{\frac{\partial V}{\partial \boldsymbol{\eta}}\mathbf{A}(\mathbf{q})}_{\mathbf{a}(\boldsymbol{\eta},\mathbf{q})^\top}\mathbf{u} + \underbrace{\frac{\partial V}{\partial \boldsymbol{\eta}}\mathbf{b}(\mathbf{q},\dot{\mathbf{q}})}_{b(\boldsymbol{\eta},\mathbf{q},\dot{\mathbf{q}})}$$

Supervised Learning

$$\begin{aligned}\dot{V}(\boldsymbol{\eta},\mathbf{u}) = {} & \widehat{\dot{V}}(\boldsymbol{\eta},\mathbf{u}) + \widehat{\mathbf{a}}(\boldsymbol{\eta},\mathbf{q})^\top\mathbf{u} + \widehat{b}(\boldsymbol{\eta},\mathbf{q},\dot{\mathbf{q}}) \\ & + \underbrace{\left(\mathbf{a}(\boldsymbol{\eta},\mathbf{q}) - \widehat{\mathbf{a}}(\boldsymbol{\eta},\mathbf{q})\right)^\top\mathbf{u} + b(\boldsymbol{\eta},\mathbf{q},\dot{\mathbf{q}}) - \widehat{b}(\boldsymbol{\eta},\mathbf{q},\dot{\mathbf{q}})}_{\delta(\boldsymbol{\eta},\mathbf{q},\dot{\mathbf{q}},\mathbf{u})}\end{aligned}$$

PSS

$$\|\boldsymbol{\eta}(t)\| \leq \beta(\|\boldsymbol{\eta}(0)\|, t) + \gamma\left(\sup_{\tau \geq 0}\|\delta(\tau)\|\right)$$

Taylor, Dorobantu, Le, Yue, Ames, A Control Lyapunov Perspective on Episodic Learning via Projection to State Stability, CDC 2019

Covariate Shift

Covariate Shift

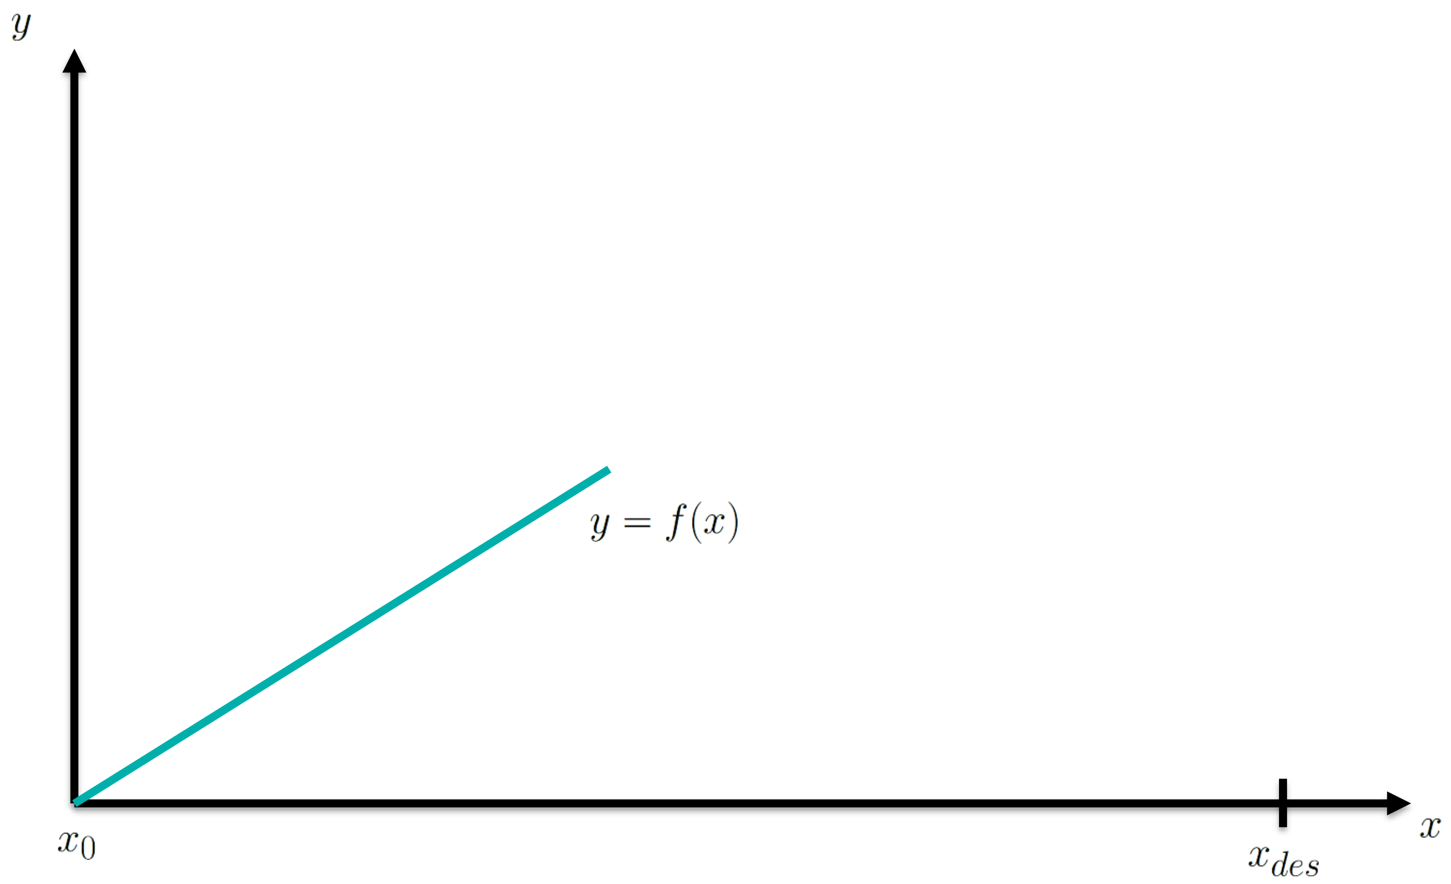

Covariate Shift

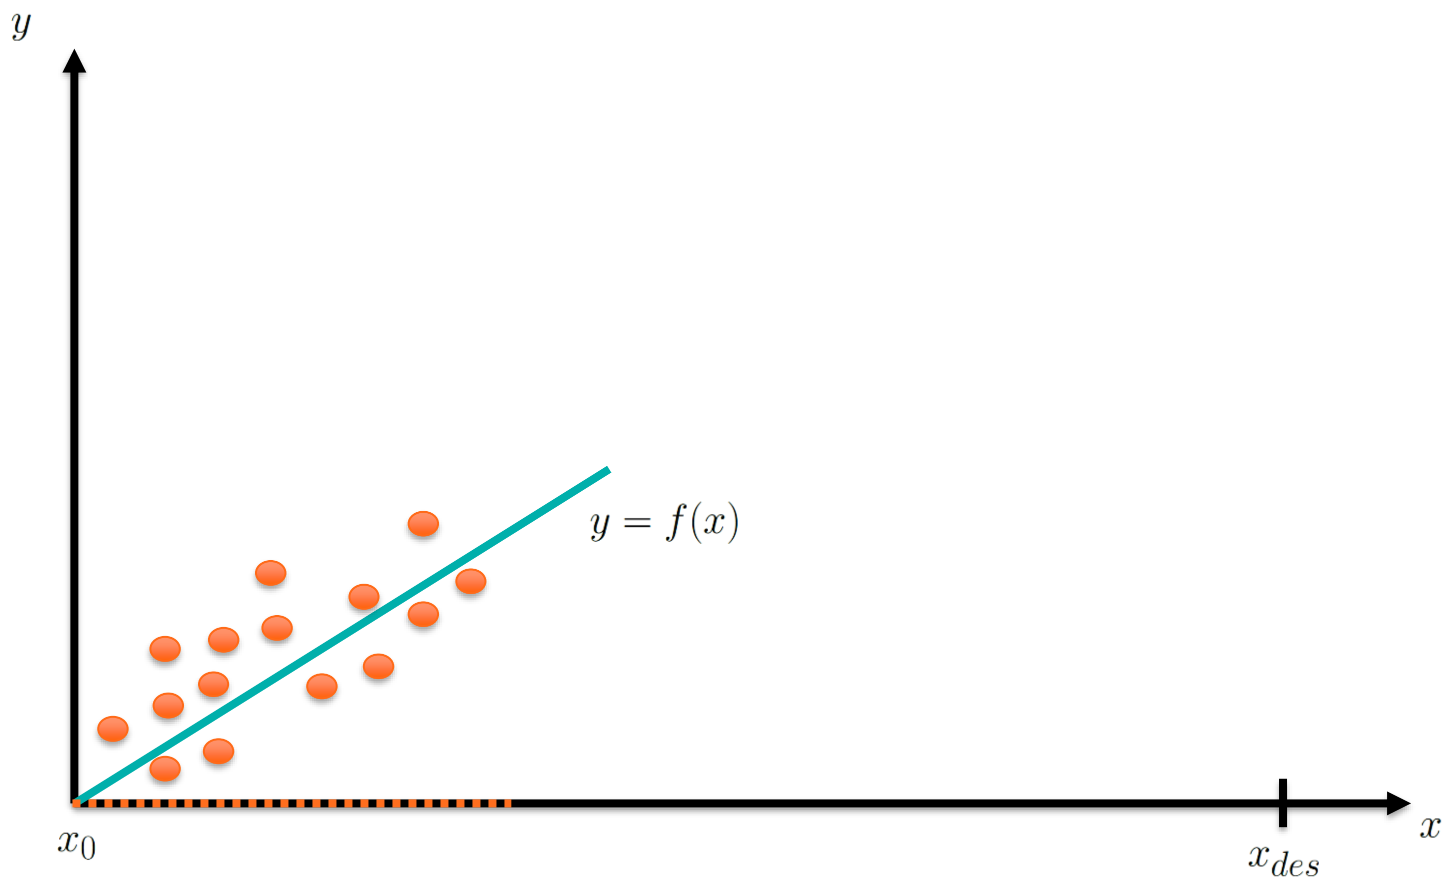

Covariate Shift

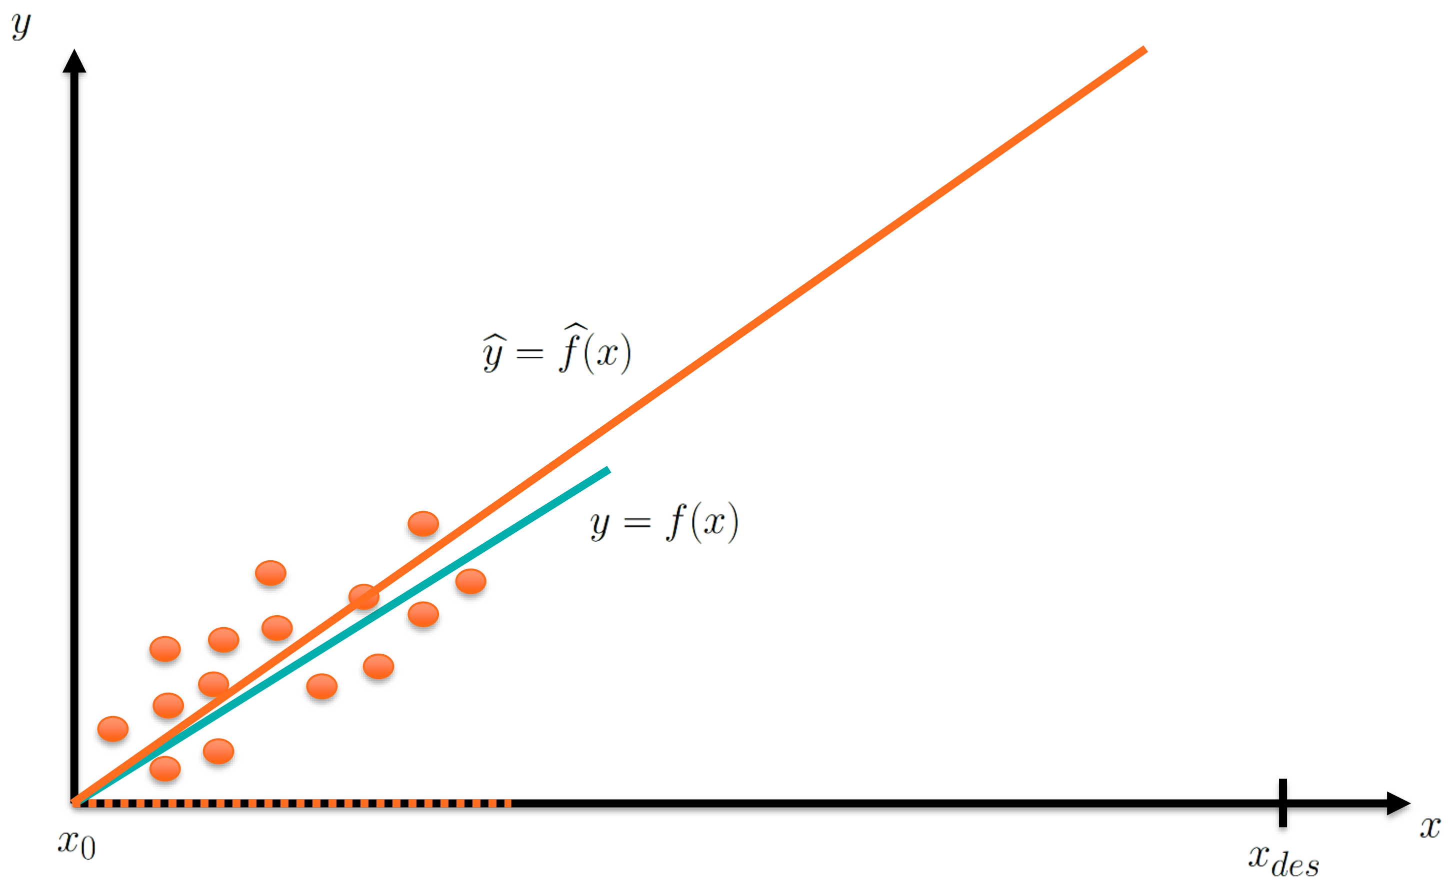

Caltech

Andrew J. Taylor

Covariate Shift

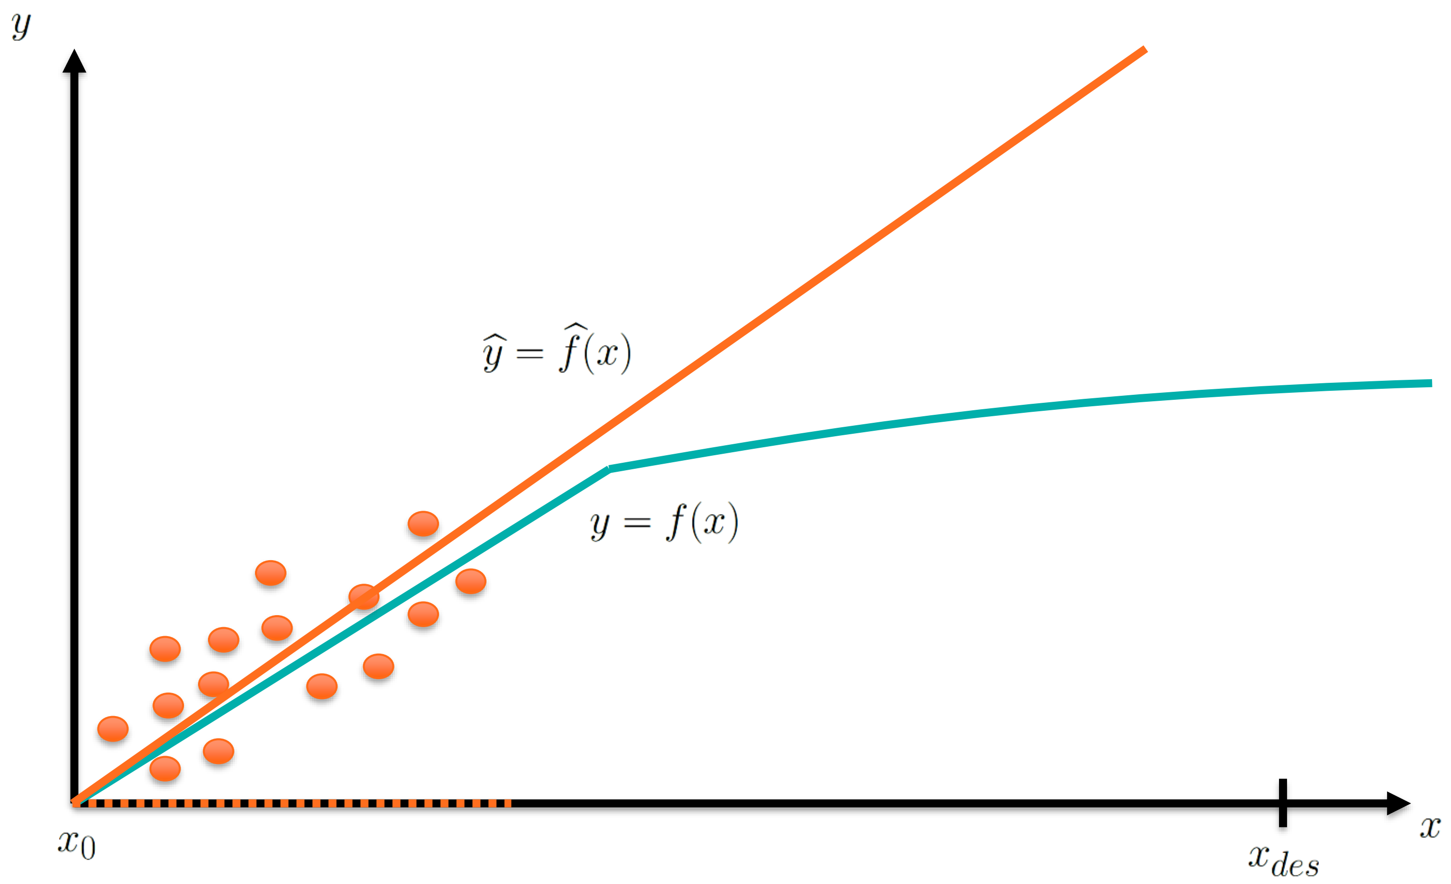

Covariate Shift

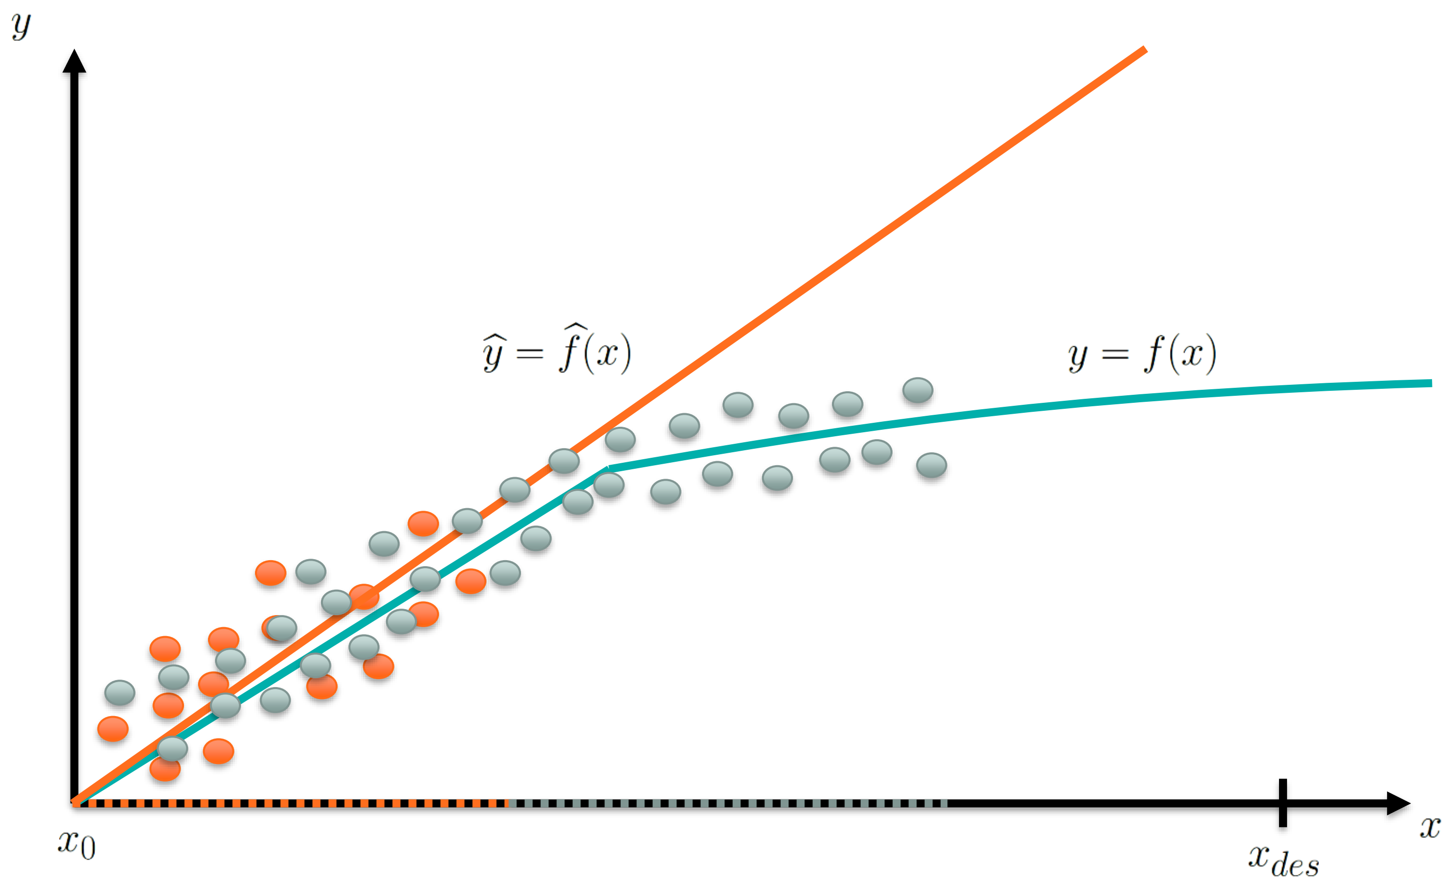

Caltech

Andrew J. Taylor

Covariate Shift

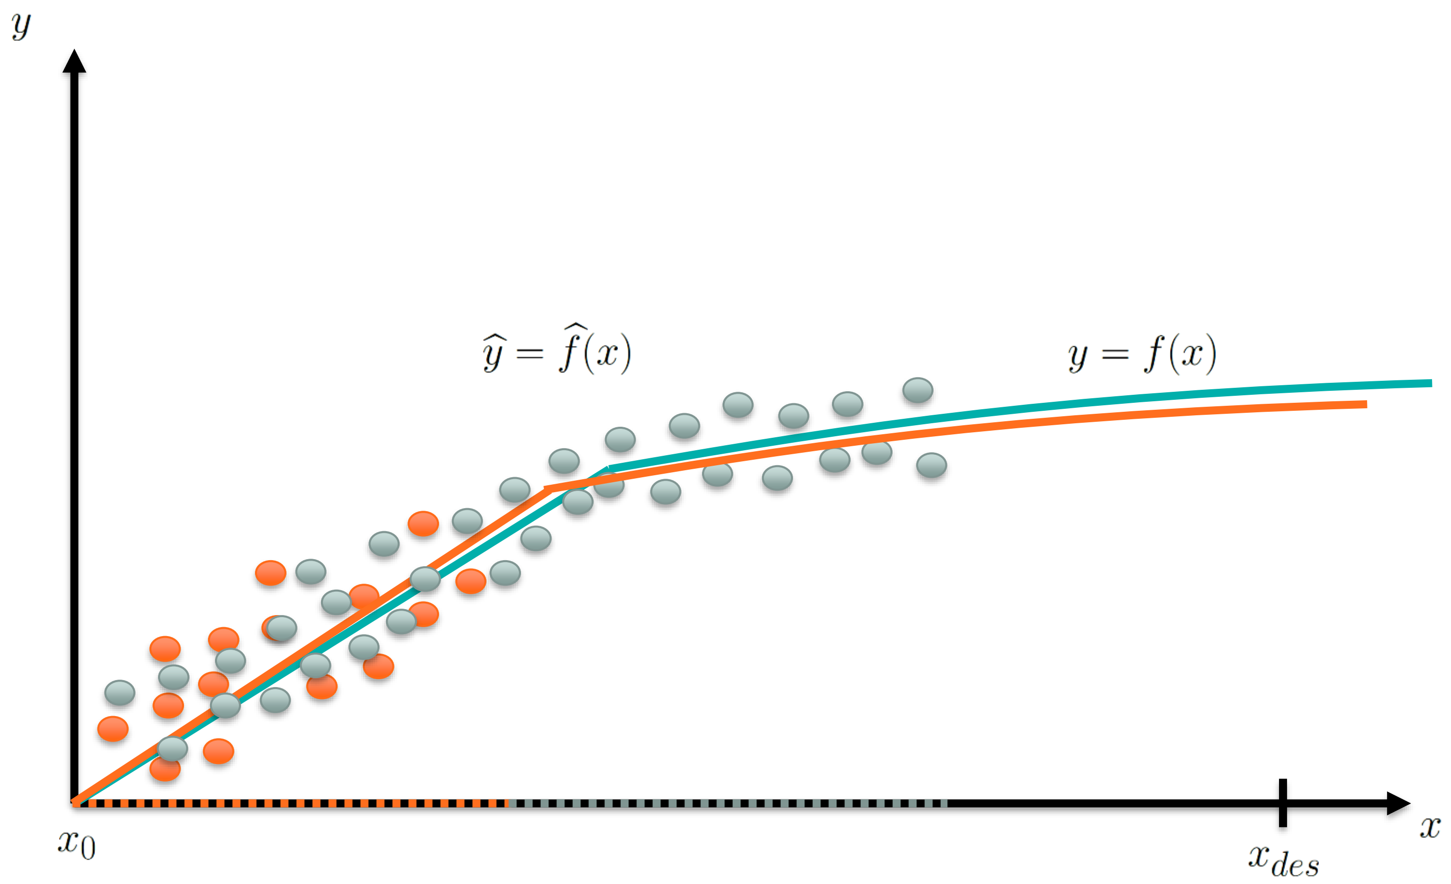

Caltech

Andrew J. Taylor