Constructive Nonlinear Control of Underactuated Systems via Zero Dynamics Policies

William Compton*, Ivan Dario Jimenez Rodriguez*, Noel Csomay-Shanklin*,
Yisong Yue, Aaron D. Ames

***Abstract*— Stabilizing underactuated systems is an inherently challenging control task due to fundamental limitations on how the control input affects the unactuated dynamics. Decomposing the system into actuated (output) and unactuated (zero) coordinates provides useful insight as to how input enters the system dynamics. In this work, we leverage the structure of this decomposition to formalize the idea of *Zero Dynamics Policies (ZDPs)*—a mapping from the unactuated coordinates to desired actuated coordinates. Specifically, we show that a ZDP exists in a neighborhood of the origin, and prove that combining output stabilization with a ZDP results in stability of the full system state. We detail a constructive method of obtaining ZDPs in a neighborhood of the origin, and propose a learning-based approach which leverages optimal control to obtain ZDPs with much larger regions of attraction. We demonstrate that such a paradigm can be used to stabilize the canonical underactuated system of the cartpole, and showcase an improvement over the nominal performance of LQR.**

I. Introduction

Underactuated systems – e.g. legged robots, dexterous manipulators, and systems with strict actuator limits – are inherently challenging to control due to the present passive dynamics. Specifically, one cannot directly actuate all degrees of freedom, which results in dynamics that cannot be arbitrarily shaped. In the case when these passive dynamics are stable, constructive feedback controllers can be synthesized, e.g., input-output linearization [1] and control Lyapunov functions [2]. Yet stabilizing underactuated systems without stability assumptions on the passive dynamics remains a challenging problem. To address this problem, this paper builds upon the observation that the existence of passive (not directly actuated) dynamics does not necessarily imply a loss of stabilizability—rather achieving stable behaviors requires careful coordination of the actuated degrees of freedom.

To stabilize underactuated systems, we leverage *Zero dynamics*: a powerful tool for analyzing underactuated systems which considers the residual dynamics when all outputs have been zeroed [3]. Traditional Zero Dynamics approaches craft output coordinates whose zero dynamics are then algebraically checked for stability. Underactuation compels this guess-and-check approach due to the challenges associated with finding feedback linearizable output coordinates which span the complete state space [4]. This paper takes a different approach, wherein we view the outputs as learnable design variables, enabling stable zero dynamics by construction. We prove that stabilizing to the zero dynamics surface defined by the learned outputs results in stability of the overall system.

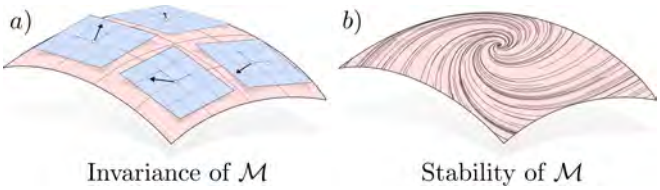

Fig. 1: The two conditions required of the zeroing manifold: a) controlled invariance, and b) stable zero dynamics.

The perspective taken in this paper has origins in the stabilization of non-minimum phase systems [5], [6], [7], and extends work proposing the design of a stabilizing zero dynamics manifold in the context of Hybrid Zero Dynamics (HZD) for bipedal walking [8]. These methods, however, are often domain-specific and challenging to synthesize, leading many practitioners to turn to optimal control techniques instead. Both approximate value function feedback [9], [10] and receding horizon Model Predictive Control (MPC) are common variations of this approach [11], [12]. Although practically useful, these methods provide limited insight as to why stabilizing underactuated systems is fundamentally difficult. To address this issue, we take inspiration from legged locomotion [13], [14], [15], where mappings between the underactuated coordinates to the actuated coordinates are central to controller synthesis. The authors exploited this connection, in conjunction with HZD, in [16] where planar biped walking was generated by enforcing barrier function certificates on the zero dynamics manifold through the use of learned output parameters. Our goal in this work is to formalize and unify the above approaches.

This paper presents a framework for the constructive stabilization of underactuated systems through *Zero Dynamics Policies (ZDPs)*. These policies take the form of a mapping from the underactuated states to the actuated states, which defines a controlled invariant and stable manifold, represented in Figure 1. Stabilizing this manifold via output tracking results in guaranteed stability of the overall system. Furthermore, we prove that for (locally) controllable nonlinear systems such a mapping exists and can be constructed analytically in a neighborhood of the origin. Finally, we leverage optimal control and machine learning as a constructive method for extending the region of validity of these ZDPs. We demonstrate that the use of ZDPs leads to a larger region of attraction than traditional control methods on the classic cartpole system.

*denotes equal contribution. Authors are with the Department of Computing and Mathematical Sciences, California Institute of Technology, Pasadena, CA 91125. This research is supported by Technology Innovation Institute (TII), NSF Graduate Research Fellowship No. DGE-1745301, NSF CPS Grant #1918655, JPL RTD 1643049.

II. Preliminaries

Consider a control affine nonlinear system

$$\dot{\mathbf{x}} = \mathbf{f}_{\mathbf{x}}(\mathbf{x}) + \mathbf{g}_{\mathbf{x}}(\mathbf{x})v \tag{1}$$

with state $\mathbf{x} \in \mathbb{R}^n$, input $v \in \mathbb{R}$, and functions $\mathbf{f}_{\mathbf{x}} : \mathbb{R}^n \to \mathbb{R}^n$ and $\mathbf{g}_{\mathbf{x}} : \mathbb{R}^n \to \mathbb{R}^n$ assumed to be continuously differentiable on $\mathbb{R}^n$. When analyzing underactuated systems, it will be useful to consider how actuation enters the system dynamics; to this end, consider an output $y : \mathbb{R}^n \to \mathbb{R}$. In order for the evolution of this output to be impacted by a controller, the input v must appear in a derivative of y in a meaningful way. Consider a time derivative of the output:

$$\dot{y}(\mathbf{x}) = \underbrace{\frac{\partial y}{\partial \mathbf{x}}\mathbf{f}_{\mathbf{x}}(\mathbf{x})}_{L_{\mathbf{f}_{\mathbf{x}}}y(\mathbf{x})} + \underbrace{\frac{\partial y}{\partial \mathbf{x}}\mathbf{g}_{\mathbf{x}}(\mathbf{x})}_{L_{\mathbf{g}_{\mathbf{x}}}y(\mathbf{x})} v$$

where $L_{\mathbf{f}_{\mathbf{x}}}y : \mathbb{R}^n \to \mathbb{R}$ and $L_{\mathbf{g}_{\mathbf{x}}}y : \mathbb{R}^n \to \mathbb{R}$ are the *Lie derivatives* of the output y with respect to the vector fields $\mathbf{f}_{\mathbf{x}}$ and $\mathbf{g}_{\mathbf{x}}$, respectively. If $L_{\mathbf{g}_{\mathbf{x}}}y(\mathbf{x}) \equiv 0$, we can attempt to continue differentiating until a higher derivative is nonzero:

$$y^{(\gamma)}(\mathbf{x}) = L^{\gamma}_{\mathbf{f}_{\mathbf{x}}}y(\mathbf{x}) + L_{\mathbf{g}_{\mathbf{x}}}L^{\gamma-1}_{\mathbf{f}_{\mathbf{x}}}(\mathbf{x})v.$$

Differentiating the output until the input appears is captured in the following notion of strict relative degree:

Definition 1. [1] An output $y : \mathbb{R}^n \to \mathbb{R}$ for the system (1) is said to have relative degree $\gamma \in \mathbb{Z}_+$ at $\mathbf{x}_0$ if:

$$L_{\mathbf{g}_{\mathbf{x}}}L^{k}_{\mathbf{f}_{\mathbf{x}}}(\mathbf{x}) \equiv 0, \qquad 0 \leq k \leq \gamma - 2$$

and $L_{\mathbf{g}_{\mathbf{x}}}L^{\gamma-1}_{\mathbf{f}_{\mathbf{x}}}(\mathbf{x}) \neq 0$.

Given an output of relative degree $\gamma \in \mathbb{Z}_+$, consider the mapping $\boldsymbol{\Phi}_{\boldsymbol{\eta}} : \mathbb{R}^n \to N \triangleq \mathbb{R}^{\gamma}$, defined as:

$$\boldsymbol{\Phi}_{\boldsymbol{\eta}}(\mathbf{x}) \triangleq \begin{bmatrix} y(\mathbf{x}) & \dot{y}(\mathbf{x}) & \cdots & y^{(\gamma-1)}(\mathbf{x}) \end{bmatrix}^{\top}. \tag{2}$$

We will subsequently take $\boldsymbol{\eta} = \boldsymbol{\Phi}_{\boldsymbol{\eta}}(\mathbf{x}) \in N$ to represent coordinates of the output space. Valid relative degree allows the constructive synthesis of controllers which exponentially stabilizes the outputs.

Definition 2. The signal $\boldsymbol{\eta}(t)$ is *exponentially stable* on domain $D \subset N$ if there exists $M, \lambda > 0$ such that:

$$\boldsymbol{\eta}_0 \in D \quad \implies \quad \|\boldsymbol{\eta}(t)\| \leq Me^{-\lambda t}\|\boldsymbol{\eta}_0\|.$$

Lyapunov theory [17] states that exponential stability is one-to-one with the existence of a control Lyapunov function (CLF) $V : N \to \mathbb{R}$ satisfying:

$$\begin{aligned} k_1\|\boldsymbol{\eta}\|^2 \leq V(\boldsymbol{\eta}) \leq k_2\|\boldsymbol{\eta}\|^2 \\ \inf_{v} \dot{V}(\mathbf{x}, v) \leq -k_3 V(\boldsymbol{\eta}). \end{aligned} \tag{3}$$

for $k_i > 0$. We define $\mathcal{K} = \{k(\mathbf{x}) \mid \dot{V}(\mathbf{x}, k(\mathbf{x})) \leq -k_3 V(\boldsymbol{\eta})\}$ to be the set of all output exponentially stabilizing feedback controllers, which is nonempty under valid relative degree [18]. A common technique to stabilize outputs with valid relative degree is feedback linearization:

$$k_{\text{fbl}}(\mathbf{x}) = \left(L_{\mathbf{g}_{\mathbf{x}}}L^{\gamma-1}_{\mathbf{f}_{\mathbf{x}}}(\mathbf{x})\right)^{-1}\left(-L^{\gamma}_{\mathbf{f}_{\mathbf{x}}}(\mathbf{x}) + u\right)$$

for $k_{\text{fbl}} : \mathbb{R}^n \times \mathbb{R} \to \mathbb{R}$, with $u \in \mathbb{R}$ the auxiliary input. Under this controller, the $\boldsymbol{\eta}$ dynamics become:

$$\dot{\boldsymbol{\eta}} = \underbrace{\begin{bmatrix} \mathbf{0} & \mathbf{I} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}}_{\triangleq \mathbf{F}} \boldsymbol{\eta} + \underbrace{\begin{bmatrix} \mathbf{0} \\ 1 \end{bmatrix}}_{\triangleq \mathbf{G}} u,$$

for $\mathbf{F} \in \mathbb{R}^{\gamma \times \gamma}, \mathbf{G} \in \mathbb{R}^{\gamma}$. Once a system's available outputs are zeroed, the remaining states evolve on a manifold [3], motivating the following discussion.

Consider a differentiable function $\mathbf{h} : \mathbb{R}^n \to \mathbb{R}^p$ with $\mathbf{0}$ a regular value, i.e. $\frac{\partial \mathbf{h}}{\partial \mathbf{x}}|_{\mathbf{x}=0}$ is full rank. Then, we have that $\mathcal{M} \triangleq \{\mathbf{x} \mid \mathbf{h}(\mathbf{x}) = \mathbf{0}\}$ defines a p-dimensional embedded submanifold of $\mathbb{R}^n$ [19]. Associated with such a manifold is the notion of a tangent space. As the manifolds being considered are 0-level sets of a function $\mathbf{h}$, a vector $\mathbf{v} \in \mathbb{R}^n$ is a *tangent vector* to a manifold $\mathcal{M}$ at the point $\mathbf{x} \in \mathcal{M}$, denoted as $\mathbf{v} \in \mathsf{T}_{\mathbf{x}}\mathcal{M}$, if:

$$\frac{\partial \mathbf{h}}{\partial \mathbf{x}}^{\top} \mathbf{v} = \mathbf{0}.$$

This aligns with the classical notions of tangent vectors, as the gradient field of a function provides a basis for the annihilator of the *tangent space*, the span of all tangent vectors at a point, of the manifold defined by a level set.

A key property will be the notion of controlled invariance for such a manifold, defined as:

Definition 3. A manifold $\mathcal{M}$ is *controlled invariant* under the dynamics (1) if for all $\mathbf{x} \in \mathbb{R}^n$ there exists an input $v \in \mathbb{R}$ such that:

$$\mathbf{f}_{\mathbf{x}}(\mathbf{x}) + \mathbf{g}_{\mathbf{x}}(\mathbf{x})v \in \mathsf{T}_{\mathbf{x}}\mathcal{M}.$$

There must exist and input such that the vector field associated with the dynamics must lie in the tangent space of the manifold. Our proposed method aims to find controlled invariant manifolds with exponentially stable dynamics.

Finally, we introduce the *actuation decomposition*, which highlights the structure of actuated and unactuated states. When y is valid relative degree, each $y^{(i)}$, for $i = 0, \ldots, \gamma-1$ are linearly independent, and $\boldsymbol{\eta}$ forms a basis for γ dimensions of $\mathbb{R}^n$ [3]. We can construct a set of normal coordinates $\mathbf{z} \in Z \subset \mathbb{R}^{n_z}$, where $n_z = n - \gamma$, via $\boldsymbol{\Phi}_{\mathbf{z}} : \mathbb{R}^n \to Z$. This transform is defined such that $\boldsymbol{\Phi} : \mathbb{R}^n \to N \times Z$ given by:

$$\begin{bmatrix} \boldsymbol{\eta} \\ \mathbf{z} \end{bmatrix} = \begin{bmatrix} \boldsymbol{\Phi}_{\boldsymbol{\eta}}(\mathbf{x}) \\ \boldsymbol{\Phi}_{\mathbf{z}}(\mathbf{x}) \end{bmatrix} \triangleq \boldsymbol{\Phi}(\mathbf{x})$$

is a diffeomorphism, and $\frac{\partial \boldsymbol{\Phi}_{\mathbf{z}}}{\partial \mathbf{x}}\mathbf{g}_{\mathbf{x}} \equiv 0$ [3]. This implies that the $\mathbf{z}$ dynamics are independent of the control input:

$$\begin{bmatrix} \dot{\boldsymbol{\eta}} \\ \dot{\mathbf{z}} \end{bmatrix} = \underbrace{\begin{bmatrix} \mathbf{F}\boldsymbol{\eta} \\ \boldsymbol{\omega}(\boldsymbol{\eta}, \mathbf{z}) \end{bmatrix}}_{\mathbf{f}(\boldsymbol{\zeta})} + \underbrace{\begin{bmatrix} \mathbf{G} \\ \mathbf{0} \end{bmatrix}}_{\mathbf{g}(\boldsymbol{\zeta})} u \tag{4}$$

where $\boldsymbol{\zeta} \triangleq (\boldsymbol{\eta}, \mathbf{z}) \in \mathcal{X} \triangleq N \times Z$. We deem this transformation an *actuation decomposition*; it separates the system states into a set of directly actuated coordinates ($\boldsymbol{\eta}$), and a set of unactuated coordinates ($\mathbf{z}$). As this decomposition takes a central role in our approach to underactuated control, (4) will be the starting point for dynamics in this paper.

III. Zero Dynamics Policies

We propose a differentiable mapping $\boldsymbol{\psi} : Z \to N$, which maps from the underactuated states $\mathbf{z}$ to desired locations for the actuated states $\boldsymbol{\eta}$. This is motivated by several results in robotics, such as the Raibert Heuristic, which maps a walking robot's center of mass (unactuated) to desired foot positions (actuated). The mapping $\boldsymbol{\psi}$ induces a n_z dimensional submanifold of $\mathcal{X}$ via the zero level set of the function $\mathbf{h}(\boldsymbol{\eta}, \mathbf{z}) \triangleq \boldsymbol{\eta} - \boldsymbol{\psi}(\mathbf{z})$:

$$\mathcal{M}_{\boldsymbol{\psi}} \triangleq \{(\boldsymbol{\eta}, \mathbf{z}) \in N \times Z \mid \mathbf{h}(\boldsymbol{\eta}, \mathbf{z}) = \mathbf{0}\}, \tag{5}$$

as $\frac{\partial \mathbf{h}}{\partial \boldsymbol{\zeta}} = [\mathbf{I} \ \frac{\partial \boldsymbol{\psi}}{\partial \mathbf{z}}]$ is full row rank. With this, we can now introduce the notion of *zero dynamics*:

Definition 4. The *zero dynamics* associated with a controlled invariant manifold $\mathcal{M}_{\boldsymbol{\psi}}$ are given by:

$$\dot{\mathbf{z}} = \boldsymbol{\omega}(\boldsymbol{\psi}(\mathbf{z}), \mathbf{z}).$$

A. Required Properties of Zero Dynamics Policies

If the zero dynamics are exponentially stable, we will show that stabilizing $\mathcal{M}_{\boldsymbol{\psi}}$ stabilizes the whole system. To this end, we propose the following output:

$$y = \eta_1 - \psi_1(\mathbf{z}), \tag{6}$$

where $(\cdot)_i$ denotes the i^{th} index. The following assumption is required to ensure y can maintain relative degree γ:

Assumption 1. We have that $\frac{\partial \boldsymbol{\omega}}{\partial \eta_i} = \mathbf{0}$ for all $i = 3, \dots, \gamma$.

This assumption is trivially satisfied for $\gamma \leq 2$, a case typical for robotic systems. We now give a condition for the relative degree of (6):

Lemma 1. *The output $y(\boldsymbol{\zeta}) = \eta_1 - \psi_1(\mathbf{z})$ has valid relative degree γ if and only if $\frac{\partial \psi_1}{\partial z} \frac{\partial \boldsymbol{\omega}}{\partial \eta_2} \neq 1$.*

Proof: Taking derivatives until the input appears yields:

$$\begin{aligned} \dot{y} &= \eta_2 - \frac{\partial \psi_1}{\partial \mathbf{z}} \boldsymbol{\omega}(\boldsymbol{\eta}, \mathbf{z}) \\ y^{(i)} &= \left(1 - \frac{\partial \psi_1}{\partial \mathbf{z}} \frac{\partial \boldsymbol{\omega}}{\partial \eta_2}\right) \eta_{i+1} + W_i(\boldsymbol{\eta}_{1:i}, \mathbf{z}), \\ y^{(\gamma)} &= \left(1 - \frac{\partial \psi_1}{\partial \mathbf{z}} \frac{\partial \boldsymbol{\omega}}{\partial \eta_2}\right) u + W_\gamma(\boldsymbol{\eta}, \mathbf{z}) \end{aligned} \tag{7}$$

for $i = 2, \dots, \gamma - 1$ and where $W_i : N \times Z \to \mathbb{R}$ is introduced to hold additional terms for each derivative , and $\boldsymbol{\eta}_{i:j} = [\eta_i \ \dots \ \eta_j]^\top$. As u does not appear until $y^{(\gamma)}$, we have that $L_{\mathbf{g}} L_{\mathbf{f}}^i y \equiv 0$ for $i = 0, \dots, \gamma - 2$, and $L_{\mathbf{g}} L_{\mathbf{f}}^i y = 1 - \frac{\partial \psi_1}{\partial \mathbf{z}} \frac{\partial \boldsymbol{\omega}}{\partial \eta_2}$. Therefore, the output is relative degree γ if and only if this term is nonzero. ∎

In the case $\gamma = 1$, the output (6) has valid relative degree one. Importantly, if y is valid relative degree, each $y^{(i)}$, for $i = 0, \dots, \gamma - 1$ are linearly independent, and form a basis for γ dimensions of $\mathcal{X}$ [3]. Defining the error coordinates:

$$\mathbf{e} = \begin{bmatrix} y & \dot{y} & \cdots & y^{(\gamma-1)} \end{bmatrix}^\top \in \mathcal{E} \tag{8}$$

as y and its first $\gamma - 1$ derivatives, we can construct the associated zeroing manifold of the output $y = \eta_1 - \psi_1(\mathbf{z})$.

Lemma 2. *Consider a controlled invariant manifold $\mathcal{M}_{\boldsymbol{\psi}}$ and its associated output* (6)*. If the output has relative degree, then $\mathcal{M}_{\boldsymbol{\psi}}$ is the zeroing manifold associated* (6)*.*

Proof: Valid relative degree implies the existence of a unique zeroing manifold $\mathcal{M}$, an n_z dimensional surface on which $\mathbf{e} \equiv 0$ [3]. Because $\mathcal{M}_{\boldsymbol{\psi}}$ is controlled invariant with $y \equiv 0$ (implying derivatives of y are zero), it is also an n_z dimensional zeroing surface; by uniqueness, $\mathcal{M} = \mathcal{M}_{\boldsymbol{\psi}}$. ∎

We have shown that a controlled invariant manifold $\mathcal{M}_{\boldsymbol{\psi}}$ is the zeroing manifold associated with (6) when this output is valid relative degree. If a function $\boldsymbol{\psi}$ can be found such that the zero dynamics on $\mathcal{M}_{\boldsymbol{\psi}}$ are also stable, then the system can be constructively stabilized as follows:

Theorem 1. *Consider a relative degree γ output $y = \eta_1 - \psi_1(\mathbf{z})$ with zeroing manifold $\mathcal{M}_{\boldsymbol{\psi}}$. If the zero dynamics of $\mathcal{M}_{\boldsymbol{\psi}}$ are exponentially stable, then any output stabilizing controller $\mathbf{k} \in \mathcal{K}$ renders the full state $\boldsymbol{\zeta}$ exponentially stable.*

Proof: By Lemma 2, the relative degree of y implies that holding $\mathbf{e} = 0$ renders $\mathcal{M}_{\boldsymbol{\psi}}$ invariant. Since $\mathcal{M}_{\boldsymbol{\psi}}$ has exponentially stable dynamics when rendered invariant converse Lyapunov guarantees the existence of $V_{\mathbf{z}}(\mathbf{z})$ satisfying:

$$\begin{aligned} k_{1,z}\|\mathbf{z}\|^2 \leq V_{\mathbf{z}}(\mathbf{z}) &\leq k_{2,z}\|\mathbf{z}\|^2 \\ \dot{V}_{\mathbf{z}}(\mathbf{z}) = \frac{\partial V_{\mathbf{z}}}{\partial \mathbf{z}} \boldsymbol{\omega}(\boldsymbol{\psi}(\mathbf{z}), \mathbf{z}) &\leq -k_{3,z}\|\mathbf{z}\|^2 \\ \left\| \frac{\partial V_{\mathbf{z}}}{\partial \mathbf{z}} \right\| &\leq k_{4,z}\|\mathbf{z}\| \end{aligned}$$

for $k_{i,z} > 0$. Applying any controller $k(\boldsymbol{\zeta}) \in \mathcal{K}$ (nonempty by virtue of valid relative degree) implies by converse Lyapunov the existence of a Lyapunov function on the error. Define $V_{\mathbf{e}}(\mathbf{e})$ satisfying:

$$\begin{aligned} k_{1,e}\|\mathbf{e}\|^2 \leq V_{\mathbf{e}}(\mathbf{e}) &\leq k_{2,e}\|\mathbf{e}\|^2 \\ \dot{V}_{\mathbf{e}}(\mathbf{e}) &\leq -k_{3,e}\|\mathbf{e}\|^2 \end{aligned}$$

for $k_{i,e} > 0$. We then use the implicit function theorem to establish that $\boldsymbol{\eta}$ can be written as a function of $\mathbf{e}, \mathbf{z}$:

$$\boldsymbol{\xi}(\mathbf{e}, \boldsymbol{\eta}, \mathbf{z}) \triangleq \mathbf{e} - \begin{bmatrix} y & \dot{y} & \dots & y^{(\gamma-1)} \end{bmatrix}^\top = \mathbf{0}$$

Observe from (7) that $\frac{\partial \boldsymbol{\xi}}{\partial \boldsymbol{\eta}}$ is lower triangular. Furthermore, $(\frac{\partial \boldsymbol{\xi}}{\partial \boldsymbol{\eta}})_{1,1} = 1$ and $(\frac{\partial \boldsymbol{\xi}}{\partial \boldsymbol{\eta}})_{i,i} = (1 - \frac{\partial \psi_1}{\partial \mathbf{z}} \frac{\partial \boldsymbol{\omega}}{\partial \eta_2})$ for $i = 2, \dots, \gamma$. By assumption of relative degree, the diagonal elements are nonzero and therefore $\frac{\partial \boldsymbol{\xi}}{\partial \boldsymbol{\eta}}$ is invertible. Therefore, there exists a function $\boldsymbol{\Gamma} : \mathcal{E} \times Z \to N$ such that $\boldsymbol{\xi}(\mathbf{e}, \boldsymbol{\Gamma}(\mathbf{e}, \mathbf{z}), \mathbf{z}) = 0$. We can then redefine the $\mathbf{z}$ dynamics in terms of $\mathbf{e}$ via:

$$\tilde{\boldsymbol{\omega}}(\mathbf{e}, \mathbf{z}) = \boldsymbol{\omega}\left(\boldsymbol{\Gamma}(\mathbf{e}, \mathbf{z}), \mathbf{z}\right) = \boldsymbol{\omega}(\boldsymbol{\eta}, \mathbf{z}).$$

Finally, consider the positive definite function $V(\mathbf{e}, \mathbf{z}) =$

$\sigma V_{\mathbf{e}}(\mathbf{e}) + V_{\mathbf{z}}(\mathbf{z})$, for $\sigma > 0$, whose time derivative is:

$$\begin{aligned}\dot{V}(\mathbf{e},\mathbf{z}) &= \sigma \dot{V}_{\mathbf{e}}(\mathbf{e}) + \frac{\partial V_{\mathbf{z}}}{\partial \mathbf{z}} \tilde{\boldsymbol{\omega}}(\mathbf{e},\mathbf{z}) \\ &= \dot{V}_{\mathbf{e}}(\mathbf{e}) + \frac{\partial V_{\mathbf{z}}}{\partial \mathbf{z}} \tilde{\boldsymbol{\omega}}(\mathbf{0},\mathbf{z}) + \frac{\partial V_{\mathbf{z}}}{\partial \mathbf{z}} \left(\tilde{\boldsymbol{\omega}}(\mathbf{e},\mathbf{z}) - \tilde{\boldsymbol{\omega}}(\mathbf{0},\mathbf{z})\right) \\ &\leq -\sigma k_{3,e} \|\mathbf{e}\|^2 - k_{3,z} \|\mathbf{z}\|^2 + k_{4,z} L_\omega \|\mathbf{e}\| \|\mathbf{z}\| \\ &= -\begin{bmatrix} \|\mathbf{e}\| \\ \|\mathbf{z}\| \end{bmatrix}^\top \begin{bmatrix} \sigma k_{3,e} & -\frac{k_{4,z} L_\omega}{2} \\ -\frac{k_{4,z} L_\omega}{2} & k_{3,z} \end{bmatrix} \begin{bmatrix} \|\mathbf{e}\| \\ \|\mathbf{z}\| \end{bmatrix} \end{aligned} \tag{9}$$

where L_ω is a Lipshitz constant of $\tilde{\boldsymbol{\omega}}$. Choosing $\sigma > \frac{k_{4,z}^2 L_\omega^2}{4 k_{3,e} k_{3,z}}$ renders the matrix positive definite, and the quadratic form (9) can be bounded, for $\lambda > 0$:

$$\dot{V}(\mathbf{e},\mathbf{z}) \leq -\lambda V(\mathbf{e},\mathbf{z})$$

since V can be bounded by quadratic functions, certifying that the V is a Lyapunov function by (3). Thus, the composite system is exponentially stable under any controller which exponentially stabilizes the outputs. ∎

B. Local Existence of Stabilizing Zero Dynamics Policies

Around the origin, we demonstrate via construction that $\boldsymbol{\psi}$ exists for any locally controllable nonlinear system. To achieve this, we identify a stable zeroing manifold for the linear system and leverage the relationship between a nonlinear system and its linearization to generate an output with stable zero dynamics. To this end, consider the linearization of (4) about the origin:

$$\begin{bmatrix} \dot{\eta}_1 \\ \dot{\boldsymbol{\eta}}_{2:\gamma-1} \\ \dot{\eta}_\gamma \\ \dot{\mathbf{z}} \end{bmatrix} = \underbrace{\begin{bmatrix} 0 & 1 & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} \\ 0 & 0 & \mathbf{0} & \mathbf{0} \\ \mathbf{a}_{\eta_1} & \mathbf{a}_{\eta_2} & \mathbf{0} & \mathbf{A}_{\mathbf{z}} \end{bmatrix}}_{\triangleq \mathbf{A}} \begin{bmatrix} \eta_1 \\ \eta_2 \\ \boldsymbol{\eta}_{3:\gamma} \\ \mathbf{z} \end{bmatrix} + \underbrace{\begin{bmatrix} 0 \\ \mathbf{0} \\ 1 \\ \mathbf{0} \end{bmatrix}}_{\triangleq \mathbf{B}} u \tag{10}$$

with $\mathbf{a}_{\eta_i} = \frac{\partial \boldsymbol{\omega}}{\partial \eta_i} \in \mathbb{R}^{n_z}$ and $\mathbf{A}_{\mathbf{z}} = \frac{\partial \boldsymbol{\omega}}{\partial \mathbf{z}} \in \mathbb{R}^{n_z \times n_z}$. First, we identify a manifold invariant under a stabilizing controller:

Lemma 3. *There exists a nonempty set of controllers which stabilize* (10) *and induce an n_z dimensional invariant subspace $\mathcal{S}$ such that for each $\mathbf{z}$ there exists a unique $\boldsymbol{\eta}$ such that $(\boldsymbol{\eta}, \mathbf{z}) \in \mathcal{S}$.*

Proof: By assumption the nonlinear system is locally controllable; therefore (10) is controllable [3]. As such, the system can be stabilized by pole placement. Let $u = -\mathbf{K}\boldsymbol{\zeta}$ be any controller which places the poles at unique locations on the negative real axis. Then the closed loop system, $\mathbf{A} - \mathbf{B}\mathbf{K}$, will have n unique, linearly independent eigenspaces [20]. It is therefore possible to pick n_z distinct eigenvectors such that the projection onto Z spans the Z subspace. Let $\mathbf{v}_1, \ldots, \mathbf{v}_{n_z} \in \mathbb{R}^n$ be these eigenvectors. Define

$$\mathbf{S} = \begin{bmatrix} \mathbf{v}_1 & \ldots & \mathbf{v}_{n_z} \end{bmatrix}$$

Then $\mathcal{S} = \text{span}(\mathbf{S})$ is an invariant subspace of the closed loop dynamics, since it is the span of eigenspaces. Given a point $\mathbf{z}$, the corresponding point on $\mathcal{S}$ is given by

$$\begin{bmatrix} \boldsymbol{\eta} \\ \mathbf{z} \end{bmatrix} = \mathbf{S} \left(\begin{bmatrix} 0 & \mathbf{I} \end{bmatrix} \mathbf{S} \right)^{-1} \mathbf{z} = \begin{bmatrix} \mathbf{S}_{\boldsymbol{\eta}}^\top \\ \mathbf{I} \end{bmatrix} \mathbf{z}$$

where $\mathbf{S}_{\boldsymbol{\eta}} = \begin{bmatrix} \mathbf{s}_{\eta_1} & \ldots & \mathbf{s}_{\eta_\gamma} \end{bmatrix}$, with $\mathbf{s}_{\eta_i} \in \mathbb{R}^{n_z}$. The matrix inverse is well defined since $\mathbf{S}$ is selected such that its projection onto $\mathbf{z}$ coordinates spans Z. ∎

Remark 1. While Lemma 3 is proven for a specific form of controller (pole placement), nearly all stabilizing controllers will have n_z dimensional invariant subspaces which are parameterizable by $\mathbf{z}$. Any such subspace can be chosen, and the resulting $\mathbf{S}$ can be used in the following analysis.

Lemma 3 defines an invariant manifold for the controlled linear system. In order to appeal to composite stability, there must exist a controller rendering $\mathcal{S}$ attractive. This can be achieved by constructing an associated output with valid relative degree. However, from [3], we know that there may exist zeroing manifolds that do not have valid relative degree. We now show that under the assumption of controllability, $\mathcal{S}$ is the zeroing manifold of a suitable output:

Lemma 4. *Consider a n_z dimensional subspace $\mathcal{S}$ satisfying Lemma 3. The output $y = \eta_1 - \mathbf{s}_{\eta_1}^\top \mathbf{z}$ has valid relative degree γ and $\mathcal{S}$ is the zeroing manifold for this output.*

Proof: Take $\mathbf{C} \triangleq \begin{bmatrix} 1 & \mathbf{0} & -\mathbf{s}_{\eta_1}^\top \end{bmatrix}$ such that $y = \mathbf{C}\boldsymbol{\zeta}$. Consider the closed loop matrix $\mathbf{A}_{cl} = \mathbf{A} - \mathbf{B}\mathbf{K}$. For compactness, define:

$$\begin{aligned} m_k &= -\mathbf{s}_{\eta_1}^\top \mathbf{A}_{\mathbf{z}}^k \mathbf{a}_{\eta_1} \\ \mathbf{O}_k &= -\mathbf{s}_{\eta_1}^\top \mathbf{A}_{\mathbf{z}}^k \\ q_k &= \mathbf{O}_k \left(\mathbf{a}_{\eta_1} + \mathbf{A}_{\mathbf{z}} \mathbf{a}_{\eta_2} \right) \\ p &= 1 - \mathbf{s}_{\eta_1}^\top \mathbf{a}_{\eta_2} \end{aligned}$$

Using these we construct the matrix $\mathbf{E} \in \mathbb{R}^{\gamma \times n}$:

$$\mathbf{E} = \begin{bmatrix} \mathbf{C} \\ \mathbf{C}\mathbf{A}_{cl} \\ \mathbf{C}\mathbf{A}_{cl}^2 \\ \vdots \\ \mathbf{C}\mathbf{A}_{cl}^{\gamma-1} \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & \cdots & 0 & \mathbf{O}_0 \\ m_0 & p & 0 & \cdots & 0 & \mathbf{O}_1 \\ m_1 & q_0 & p & \cdots & 0 & \mathbf{O}_2 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ m_{\gamma-2} & q_{\gamma-3} & q_{\gamma-4} & \cdots & p & \mathbf{O}_{\gamma-1} \end{bmatrix} \tag{11}$$

Observe that for $i < \gamma - 1$, $\mathbf{C}\mathbf{A}_{cl}^i \mathbf{B} = 0$, as the $E_{j,\gamma} = 0$ for $j = 1, \ldots, \gamma - 1$. Thus, the output y cannot be of relative degree less than γ. Assume for contradiction that the output is not relative degree γ. This implies:

$$L_{\mathbf{B}} L_{\mathbf{A}_{cl}\boldsymbol{\zeta}}^{\gamma-1} y = \mathbf{C}\mathbf{A}_{cl}^{\gamma-1}\mathbf{B} = p = 1 - \mathbf{s}_{\eta_1}^\top \mathbf{a}_{\eta_2} = 0. \tag{12}$$

By Lemma 3, $\mathcal{S}$ is $\mathbf{A}_{cl}$ invariant - this implies the existence of $\mathbf{J} \in \mathbb{R}^{n_z \times n_z}$ such that $\mathbf{A}_{cl}\mathbf{S} = \mathbf{S}\mathbf{J}$, i.e. the image of $\mathbf{S}$ under $\mathbf{A}_{cl}$ is contained in $\mathbf{S}$. As $\mathbf{C}\mathbf{S} = \mathbf{0}$, and $\mathbf{C}\mathbf{A}_{cl}^k\mathbf{S} = \mathbf{C}\mathbf{S}\mathbf{J}^k = 0$ we have that $\mathbf{S} \in \ker(\mathbf{E})$. Furthermore, by (12) we have $\mathbf{C}\mathbf{A}_{cl}^{\gamma-1}\mathbf{B} = 0$, and $\mathbf{B} \in \ker(\mathbf{E})$. Finally, note that $\mathbf{B} \notin \mathbf{S}$, as $\mathbf{S}$ contains an identity block in the lower n_z rows, while $\mathbf{B}$ has a corresponding $\mathbf{0}$ block. This implies $\dim(\ker(\mathbf{E})) \geq \dim(\text{span}(\mathbf{S}, \mathbf{B})) = n_z + 1$. By Rank Theorem, $\text{rank}(\mathbf{E}) + \dim(\ker(\mathbf{E})) = n$, and therefore,

$$\text{rank}(\mathbf{E}) \leq n - (n_z + 1) = \gamma - 1.$$

Therefore, $\mathbf{E}$ is rank deficient, since it has γ rows. There is a nontrivial left nullspace; there exists $\mathbf{v} = \begin{bmatrix} \alpha_0 & \ldots & \alpha_{\gamma-1} \end{bmatrix}^\top$,

such that $\mathbf{v}^\top \mathbf{E} = \mathbf{0}$. Without loss of generality[1], we take $\alpha_{\gamma-1} = 1$. As the rightmost block of $\mathbf{v}^\top \mathbf{E} = \mathbf{0}$,

$$\mathbf{O}_{\gamma-1} = -\sum_{i=0}^{\gamma-2} \alpha_i \mathbf{O}_i \tag{13}$$

Next, because $p = 0$ by (12), $\mathbf{v}^\top \mathbf{E}_{\gamma-2} = \alpha_{\gamma-1} q_0 = 0$, where $\mathbf{E}_i$ is the i'th column of $\mathbf{E}$. Therefore, $q_0 = 0$. Applying this logic on rows of $\mathbf{E}$ backward to $\mathbf{E}_2$ gives:

$$q_i = 0 \quad \forall i \in \{0, \dots, \gamma - 3\}.$$

Finally, we demonstrate via induction that $q_j = 0$, for any $j \in \mathbb{Z}_+$. Assuming that $q_j = 0$, we aim to show that $q_{j+1} = -\mathbf{s}_{\eta_1}^\top \mathbf{A}_{\mathbf{z}}^{j+1} (\mathbf{a}_{\eta_1} + \mathbf{A}_{\mathbf{z}} \mathbf{a}_{\eta_2}) = 0$. Introduce q_{j+1} by right-multiplying (13) by $\mathbf{A}_{\mathbf{z}}^k (\mathbf{a}_{\eta_1} + \mathbf{A}_{\mathbf{z}} \mathbf{a}_{\eta_2})$, for $k = j+1-(\gamma-1)$:

$$\begin{aligned} q_{j+1} &= \mathbf{O}_{\gamma-1} \mathbf{A}_{\mathbf{z}}^k (\mathbf{a}_{\eta_1} + \mathbf{A}_{\mathbf{z}} \mathbf{a}_{\eta_2}) \\ &= -\sum_{i=0}^{\gamma-2} \alpha_i \mathbf{s}_{\eta_1}^\top \mathbf{A}_{\mathbf{z}}^{i+k} \mathbf{a}_{\eta_1} - \sum_{i=0}^{\gamma-2} \alpha_i \mathbf{s}_{\eta_1}^\top \mathbf{A}_{\mathbf{z}}^{i+k+1} \mathbf{a}_{\eta_2}. \end{aligned}$$

By the induction hypothesis $q_i = 0$, for $i = 0, \dots, j$ which can be rearranged:

$$-\mathbf{s}_{\eta_1}^\top \mathbf{A}_{\mathbf{z}}^i \mathbf{a}_{\eta_1} = \mathbf{s}_{\eta_1}^\top \mathbf{A}_{\mathbf{z}}^{i+1} \mathbf{a}_{\eta_2}.$$

Substituting this for terms $q_{j-(\gamma-1)}$ to q_j into the right sum:

$$q_{j+1} = -\sum_{i=0}^{\gamma-2} \alpha_i \mathbf{s}_{\eta_1}^\top \mathbf{A}_{\mathbf{z}}^{i+k} \mathbf{a}_{\eta_1} + \sum_{i=0}^{\gamma-2} \alpha_i \mathbf{s}_{\eta_1}^\top \mathbf{A}_{\mathbf{z}}^{i+k} \mathbf{a}_{\eta_1} = 0.$$

And the inductive step has been shown. The base case holds trivially for $\gamma \geq 3$ by the structure of $\mathbf{E}$, and $\mathbf{CB} = 1$ for relative degree one systems (trivially valid relative degree). It remains to show the base case holds for $\gamma = 2$.

First, examine the fact that $\mathcal{S}$ is invariant under $\mathbf{A}_{cl}$. Expanding $\mathbf{A}_{\text{cl}}\mathbf{S} = \mathbf{S}\mathbf{J}$, given that $\mathbf{K} = [k_{\eta_1} \; k_{\eta_2} \; \mathbf{k}_{\mathbf{z}}]$:

$$\begin{bmatrix} \mathbf{s}_{\eta_2}^\top \\ -k_{\eta_1} \mathbf{s}_{\eta_1}^\top - k_{\eta_2} \mathbf{s}_{\eta_2}^\top - \mathbf{k}_{\mathbf{z}} \\ \mathbf{a}_{\eta_1} \mathbf{s}_{\eta_1}^\top + \mathbf{a}_{\eta_2} \mathbf{s}_{\eta_2}^\top + \mathbf{A}_{\mathbf{z}} \end{bmatrix} = \begin{bmatrix} \mathbf{s}_{\eta_1}^\top \mathbf{J} \\ \mathbf{s}_{\eta_2}^\top \mathbf{J} \\ \mathbf{J} \end{bmatrix}$$

Right multiply the second equation by $\mathbf{a}_{\eta_2}$, and using (12):

$$\mathbf{a}_{\eta_1} + \mathbf{A}_{\mathbf{z}} \mathbf{a}_{\eta_2} = (\mathbf{I} - \mathbf{a}_{\eta_2} \mathbf{s}_{\eta_1}) \mathbf{J} \mathbf{a}_{\eta_2}$$

And note that left multiplying by $\mathbf{s}_{\eta_1}$ gives:

$$q_0 = \mathbf{s}_{\eta_1}^\top (\mathbf{a}_{\eta_1} + \mathbf{A}_{\mathbf{z}} \mathbf{a}_{\eta_2}) = \mathbf{s}_{\eta_1}^\top (\mathbf{I} - \mathbf{a}_{\eta_2} \mathbf{s}_{\eta_1}^\top) \mathbf{J} \mathbf{a}_{\eta_2} = 0$$

again leveraging $\mathbf{s}_{\eta_1}^\top \mathbf{a}_{\eta_2} = 1$. We have established the based case for induction, for relative degree $\gamma = 2$.

We now aim to demonstrate the contradiction, by showing that (12) leads to the system losing controllability. Consider the controllability matrix of the system. Controllability is preserved under feedback, so we examine $\mathscr{C}(\mathbf{A}, \mathbf{B})$:

$$\mathscr{C}(\mathbf{A}, \mathbf{B}) = \begin{bmatrix} \mathbf{0} & 1 & 0 & \dots & 0 \\ \mathbf{I} & 0 & 0 & \dots & 0 \\ \mathbf{0} & \mathbf{a}_{\eta_2} & \mathbf{Q}_0 & \dots & \mathbf{Q}_{n-(\gamma+1)} \end{bmatrix}$$
$$\mathbf{Q}_i = \mathbf{A}_{\mathbf{z}}^i (\mathbf{a}_{\eta_1} + \mathbf{A}_{\mathbf{z}} \mathbf{a}_{\eta_2})$$

We finish by showing the $\mathbf{C}$ is in the left null space of $\mathscr{C}$:

$$\mathbf{C}\mathscr{C} = \begin{bmatrix} 0 & p & q_0 & \dots & q_{n-(\gamma+1)} \end{bmatrix} = \mathbf{0}$$

And therefore $\mathscr{C}$ is not full rank. We have reached contradiction, and y must be relative degree γ. Finally, note that $\mathcal{S}$ is the zeroing surface associated with y by Lemma 2. ■

Lemma 4 constructs an output and its associated zeroing manifold such that 1) the output has valid relative degree, and 2), the zeroing manifold is exponentially stable. Now, we show that the output locally retains both of these properties under the nonlinear dynamics:

Theorem 2. *Given a nonlinear system* (4)*, the output* $y = \eta_1 - \mathbf{s}_{\eta_1}^\top \mathbf{z}$ *obtained via linearization in Lemma 4 has valid relative degree and exponentially stable zero dynamics for the nonlinear system. As such, stabilizing* $\mathbf{e} \to 0$ *results in stability of the entire system in a neighborhood of the origin.*

Proof: First, we establish that the output is relative degree γ for the nonlinear system. The output has the form $y = \mathbf{C}\boldsymbol{\zeta} = \eta_1 - \mathbf{s}_{\eta_1}^\top \mathbf{z}$. By Lemma 1, we have that $L_{\mathbf{g}} L_{\mathbf{f}}^j y(\boldsymbol{\zeta}) \equiv 0$ for $j = 0, \dots, \gamma - 2$ and:

$$L_{\mathbf{g}} L_{\mathbf{f}}^{\gamma-1} y = 1 - \mathbf{s}_{\eta_1}^\top \frac{\partial \boldsymbol{\omega}}{\partial \eta_2}.$$

For relative degree, we need $L_{\mathbf{g}} L_{\mathbf{f}}^{\gamma-1} y \neq 0$.

$$\begin{aligned} L_{\mathbf{g}} L_{\mathbf{f}}^{\gamma-1} y &= 1 - \mathbf{s}_{\eta_1}^\top \frac{\partial \boldsymbol{\omega}}{\partial \eta_2} \\ &= \underbrace{1 - \mathbf{s}_{\eta_1}^\top \mathbf{a}_{\eta_2}}_{\mathbf{C}\mathbf{A}^{\gamma-1}\mathbf{B}} + \underbrace{\mathbf{s}_{\eta_1}^\top \left(\mathbf{a}_{\eta_2} - \frac{\partial \boldsymbol{\omega}}{\partial \eta_2} \right)}_{\triangleq \Delta(\boldsymbol{\zeta})} \end{aligned}$$

where $\Delta : \mathcal{X} \to \mathbb{R}$. Lemma 4 assures that $|\mathbf{C}\mathbf{A}^{\gamma-1}\mathbf{B}| \triangleq \delta > 0$. To guarantee $L_{\mathbf{g}} L_{\mathbf{f}}^{\gamma-1} y \neq 0$, we bound $|\Delta(\boldsymbol{\zeta})| < \delta$.

$$|\Delta(\boldsymbol{\zeta})| \leq \|\mathbf{s}_{\eta_1}\| \left\| \mathbf{a}_{\eta_2} - \frac{\partial \boldsymbol{\omega}}{\partial \eta_2} \right\|$$

Note that the function $\left\| \mathbf{a}_{\eta_2} - \frac{\partial \boldsymbol{\omega}}{\partial \eta_2} \right\|$ is continuous and zero at the origin. Therefore, there exists an $\varepsilon > 0$ such that

$$\left\| \mathbf{a}_{\eta_2} - \frac{\partial \boldsymbol{\omega}}{\partial \eta_2} \right\| \leq \frac{\delta}{2\|\mathbf{s}_{\eta_1}\|} \quad \forall \; (\boldsymbol{\eta}, \mathbf{z}) \in B_\varepsilon(\mathbf{0}, \mathbf{0})$$

Inside this epsilon ball, we have $|\Delta(\boldsymbol{\zeta})| \leq \frac{1}{2}\delta$, and therefore $|L_{\mathbf{g}} L_{\mathbf{f}}^{\gamma-1} y| > \frac{1}{2}\delta$. Locally, the output has valid relative degree for the nonlinear system.

It remains to show that the zeroing manifold, $\mathcal{M}_{\boldsymbol{\psi}}$ of this output is stable for the nonlinear system. First, given the

[1] $\mathbf{v}$ must have at least one nonzero term since the null space is nontrivial. If $\alpha_{\gamma-1} = 0$, this proof can be continued by redefining $\mathbf{E}$ to omit the last row (or as many as necessary such that the last element of $\mathbf{v}$ is nonzero). Then $\mathbf{v}$ can be scaled such that its last entry has magnitude 1.

error coordinates (8), consider $\frac{\partial \mathbf{e}}{\partial \boldsymbol{\zeta}} = \begin{bmatrix} \frac{\partial \mathbf{e}}{\partial \boldsymbol{\eta}} & \frac{\partial \mathbf{e}}{\partial \mathbf{z}} \end{bmatrix}$ evaluated at the origin. For a row of this matrix,

$$\begin{aligned}
\frac{\partial}{\partial \boldsymbol{\zeta}} \mathbf{e}_i \Big|_{\boldsymbol{\zeta}=\mathbf{0}} &= \frac{\partial}{\partial \boldsymbol{\zeta}} L_{\mathbf{f}}^i y \Big|_{\boldsymbol{\zeta}=\mathbf{0}} \\
&= \frac{\partial^2 L_{\mathbf{f}}^{(i-1)} y}{\partial \boldsymbol{\zeta}^2} \mathbf{f} \Big|_{\boldsymbol{\zeta}=\mathbf{0}} + \frac{\partial}{\partial \boldsymbol{\zeta}} L_{\mathbf{f}}^{(i-1)} y \frac{\partial \mathbf{f}}{\partial \boldsymbol{\zeta}} \Big|_{\boldsymbol{\zeta}=\mathbf{0}} \\
&= \frac{\partial}{\partial \boldsymbol{\zeta}} L_{\mathbf{f}}^{(i-1)} y \frac{\partial \mathbf{f}}{\partial \boldsymbol{\zeta}} \Big|_{\boldsymbol{\zeta}=\mathbf{0}} \qquad (14) \\
&= \mathbf{C}\mathbf{A}^i
\end{aligned}$$

where the third lines holds by $\mathbf{f}(\mathbf{0}) = \mathbf{0}$ and the last equality can be shown by induction. The base case holds as $\frac{\partial y}{\partial \boldsymbol{\zeta}} = \mathbf{C}$, and the induction step is given by (14). We have:

$$\frac{\partial \mathbf{e}}{\partial \boldsymbol{\zeta}} \Big|_{\boldsymbol{\zeta}=\mathbf{0}} = \begin{bmatrix} \mathbf{C} \\ \vdots \\ \mathbf{C}\mathbf{A}^{\gamma-1} \end{bmatrix} \qquad (15)$$

which is precisely the same as (11). Therefore, $\frac{\partial \mathbf{e}}{\partial \boldsymbol{\eta}}$ is lower triangular with nonzero elements on the diagonal, and is thus invertible. By the implicit function theorem, there exists $\boldsymbol{\psi} : Z \to N$ such that $\mathbf{e}(\boldsymbol{\psi}(\mathbf{z}), \mathbf{z}) = \mathbf{0}$. We aim to show that $\mathcal{M}_{\boldsymbol{\psi}}$ is stable. Consider that by the implicit function theorem:

$$\frac{\partial \boldsymbol{\psi}}{\partial \mathbf{z}} = -\frac{\partial \mathbf{e}}{\partial \boldsymbol{\eta}}^{-1} \frac{\partial \mathbf{e}}{\partial \mathbf{z}}. \qquad (16)$$

Furthermore, note that since $\frac{\partial \mathbf{e}}{\partial \boldsymbol{\zeta}}|_{\boldsymbol{\zeta}=\mathbf{0}} = \mathbf{E}$, (16) also holds for the linearization. Therefore, at the origin, the tangent space of $\mathcal{M}_{\boldsymbol{\psi}}$ for the nonlinear system is equal to the invariant subspace used by the linearization to design the output, i.e. $\frac{\partial \mathbf{e}}{\partial \boldsymbol{\eta}}|_{\boldsymbol{\zeta}=\mathbf{0}} = \mathbf{S}_{\boldsymbol{\eta}}$. On $\mathcal{M}_{\boldsymbol{\psi}}$ we have $\boldsymbol{\eta} = \boldsymbol{\psi}(\mathbf{z})$, or equivalently:

$$\begin{aligned}
\boldsymbol{\eta} &= \frac{\partial \boldsymbol{\psi}}{\partial \mathbf{z}} \Big|_{\boldsymbol{\zeta}=\mathbf{0}} \mathbf{z} + \left(\boldsymbol{\psi}(\mathbf{z}) - \frac{\partial \boldsymbol{\psi}}{\partial \mathbf{z}} \Big|_{\boldsymbol{\zeta}=\mathbf{0}} \mathbf{z} \right) \\
&= \mathbf{S}_{\boldsymbol{\eta}} \mathbf{z} + \underbrace{(\boldsymbol{\psi}(\mathbf{z}) - \mathbf{S}_{\boldsymbol{\eta}} \mathbf{z})}_{\boldsymbol{\Gamma}(\mathbf{z})}
\end{aligned}$$

By mean value theorem, there exists $\varepsilon > 0$ such that inside a ball $B_\varepsilon(\mathbf{0}, \mathbf{0})$ we have

$$\|\boldsymbol{\Gamma}(\mathbf{z})\| \leq N_\Gamma(\varepsilon)\|\mathbf{z}\|$$

With $N_\Gamma(\varepsilon) \to 0$ as $\varepsilon \to 0$.

Finally, given that the surface $\mathcal{S}$ is stable under the linearized dynamics from Lemma 3, converse Lyapunov guarantees the existence of a function $V_{\mathbf{z}}(\mathbf{z})$ satisfying

$$\begin{aligned}
k_1\|\mathbf{z}\|^2 \leq V_{\mathbf{z}}(\mathbf{z}) &\leq k_2\|\mathbf{z}\|^2 \\
\dot{V}_{\mathbf{z}}(\mathbf{z}) &= \frac{\partial V_{\mathbf{z}}}{\partial \mathbf{z}} \left(\mathbf{a}_{\eta_1} \mathbf{s}_{\eta_1} + \mathbf{a}_{\eta_2} \mathbf{s}_{\eta_2} + \mathbf{A}_{\mathbf{z}} \right) \mathbf{z} \\
&\leq -k_3\|\mathbf{z}\|^2 \\
\left\| \frac{\partial V_{\mathbf{z}}}{\partial \mathbf{z}} \right\| &\leq k_4\|\mathbf{z}\|
\end{aligned}$$

for $k_i > 0$. We aim to show that this function is also a Lyapunov function for the dynamics on the nonlinear zeroing manifold, where $(\boldsymbol{\eta}, \mathbf{z}) = (\boldsymbol{\psi}(\mathbf{z}), \mathbf{z})$. To this end, define $\tilde{\mathbf{A}} \in \mathbb{R}^{n_z \times n_z}$ and $\boldsymbol{\Delta}_1, \boldsymbol{\Delta}_2 : Z \to \mathbb{R}^{n_z}$:

$$\begin{aligned}
\tilde{\mathbf{A}} &= \mathbf{a}_{\eta_1} \mathbf{s}_{\eta_1} + \mathbf{a}_{\eta_2} \mathbf{s}_{\eta_2} + \mathbf{A}_{\mathbf{z}} \\
\boldsymbol{\Delta}_1(\mathbf{z}) &= \boldsymbol{\omega}\left(\mathbf{S}_{\boldsymbol{\eta}} \mathbf{z}, \mathbf{z}\right) - \tilde{\mathbf{A}} \mathbf{z} \\
\boldsymbol{\Delta}_2(\mathbf{z}) &= \boldsymbol{\omega}\left(\boldsymbol{\psi}(\mathbf{z}), \mathbf{z}\right) - \boldsymbol{\omega}\left(\mathbf{S}_{\boldsymbol{\eta}}, \mathbf{z}\right)
\end{aligned}$$

Differentiate $V_{\mathbf{z}}$ under the nonlinear dynamics:

$$\begin{aligned}
\dot{V}_{\mathbf{z}} &= \frac{\partial V_{\mathbf{z}}}{\partial \mathbf{z}} \boldsymbol{\omega}\left(\boldsymbol{\psi}(\mathbf{z}), \mathbf{z}\right) \\
&= \frac{\partial V_{\mathbf{z}}}{\partial \mathbf{z}} \tilde{\mathbf{A}} \mathbf{z} + \frac{\partial V_{\mathbf{z}}}{\partial \mathbf{z}} \left(\boldsymbol{\Delta}_1(\mathbf{z}) + \boldsymbol{\Delta}_2(\mathbf{z}) \right)
\end{aligned}$$

Noting that $\boldsymbol{\Delta}_1(\mathbf{z})$ is the error in the linearization of $\boldsymbol{\omega}$ on the linear surface $\mathcal{S}$, it can be bounded using mean value theorem as $\|\boldsymbol{\Delta}_1(\mathbf{z})\| \leq N_1(\varepsilon)\|\mathbf{z}\|$. Similarly, $\boldsymbol{\Delta}_2$ can be bounded using Lipschitz continuity of $\boldsymbol{\omega}$, and the fact that $\boldsymbol{\psi}(\mathbf{z}) - \mathbf{S}_{\boldsymbol{\eta}} \mathbf{z} = \boldsymbol{\Gamma}(\mathbf{z})$ to obtain $\|\boldsymbol{\Delta}_2(\mathbf{z})\| \leq L_{\boldsymbol{\omega}} N_\Gamma(\varepsilon)\|\mathbf{z}\| = N_2(\varepsilon)\|\mathbf{z}\|$. Letting $N(\varepsilon) = N_1(\varepsilon) + N_2(\varepsilon)$, we have

$$\begin{aligned}
\left\| \frac{\partial V_{\mathbf{z}}}{\partial \mathbf{z}} \left(\boldsymbol{\Delta}_1(\mathbf{z}) + \boldsymbol{\Delta}_2(\mathbf{z}) \right) \right\| &\leq \left\| \frac{\partial V_{\mathbf{z}}}{\partial \mathbf{z}} \right\| \left(\|\boldsymbol{\Delta}_1(\mathbf{z})\| + \|\boldsymbol{\Delta}_2(\mathbf{z})\| \right) \\
&\leq k_4 N(\varepsilon)\|\mathbf{z}\|^2
\end{aligned}$$

Choosing $\varepsilon > 0$ such that $N(\varepsilon) < \frac{k_3}{2k_4}$, we have

$$\begin{aligned}
\dot{V}_{\mathbf{z}}(\mathbf{z}) &\leq -k_3\|\mathbf{z}\|^2 + k_4 N(\varepsilon)\|\mathbf{z}\|^2 \\
&\leq -\frac{k_3}{2}\|\mathbf{z}\|^2
\end{aligned}$$

Therefore, we see that $V_{\mathbf{z}}$ is a Lyapunov function on the nonlinear zeroing manifold, with a sufficiently small ball around the origin. We have successfully demonstrated that $\boldsymbol{\psi}$ has the required properties to apply Theorem 1 and conclude local exponential stability of the composite system. ■

This proof uses linearization to design an output which has desirable properties. Locally, the zeroing manifold for this output $\mathcal{M}_{\boldsymbol{\psi}}$ is close to the linear systems zeroing manifold (captured by $\boldsymbol{\Delta}_2$), and the zero dynamics are close to the dynamics of the linear system (captured by $\boldsymbol{\Delta}_1$). Sufficiently close to the origin, relative degree and stability of $\mathcal{M}_{\boldsymbol{\psi}}$ are retained for the nonlinear system.

It is important to emphasize that Theorem 2 gives a completely *constructive method for stabilizing a broad class of underactuated systems using output stabilization*. In the results section, we demonstrate this constructive method on a canonical example of underactuation, the cartpole.

IV. Optimal Control for Learning Zero Dynamics Policies

While the linearization of a system about equilibrium can be used to construct locally stabilizing controllers, it is desirable to obtain larger regions of attraction, and leverage the nonlinear dynamics of the system. Therefore, we aim to construct $\mathcal{M}_{\boldsymbol{\psi}}$ satisfying controlled invariance and stability.

A. Optimal Control for Stabilization

As asymptotic stability is a necessary condition for optimality [21], we leverage optimal control to find $\mathcal{M}_{\boldsymbol{\psi}}$. Consider the infinite-time optimal control problem:

$$V(\boldsymbol{\zeta}_0) \triangleq \min_{\boldsymbol{\zeta},\mathbf{u}} \quad \int_0^\infty c(\boldsymbol{\zeta}(t),\mathbf{u}(t))dt \qquad (17)$$
$$\text{s.t.} \quad \dot{\boldsymbol{\zeta}} = \mathbf{f}(\boldsymbol{\zeta}) + \mathbf{g}(\boldsymbol{\zeta})\mathbf{u}$$

where $V : \mathcal{X} \to \mathbb{R}$ is the value function and $c : \mathcal{X} \times \mathcal{U} \to \mathbb{R}$ is a positive definite cost function. In order to apply Theorem 1, we will require exponential stability on the manifold. Therefore, we begin by stating conditions under which the optimal controller is exponentially stabilizing:

Theorem 3. *Let $V(\boldsymbol{\zeta})$ be the value function for the optimal control problem defined* (17)*, with quadratic cost $c(\boldsymbol{\zeta},\mathbf{u}) = \boldsymbol{\zeta}^\top\mathbf{Q}\boldsymbol{\zeta} + ru^2$, with $\mathbf{Q} \in \mathbb{R}^{n\times n}$ positive definite, $r > 0$ and compact state space $\mathcal{X}$. The nonlinear system is exponentially stable under the optimal controller.*

Proof: In a sufficiently small ball around the origin, the LQR approximation of the optimal controller, obtained by linearizing the dynamics about equilibrium, will be exponentially stabilizing for the nonlinear system [1], as it locally satisfies input bounds. This implies constants $M_{LQR},\ \lambda_{LQR},\ \delta > 0$ such that:

$$\|\boldsymbol{\zeta}_0\| \le \delta \quad \implies \quad \|\boldsymbol{\zeta}(t)\| \le M_{LQR}e^{-\lambda_{LQR}t}\|\boldsymbol{\zeta}_0\|$$

We aim to show that the trajectory emanating from an arbitrary initial condition $\boldsymbol{\zeta}_0 \in B_\delta(\mathbf{0})$ is exponentially stable. For any $M, \lambda > 0$, consider the set:

$$T = \{t \ge 0 \mid \|\boldsymbol{\zeta}(t)\| > Me^{-\lambda t}\|\boldsymbol{\zeta}_0\|\}$$

We condition on whether there is an upper bound to the elements of T:

Case 1: There exists an upper bound $\overline{T}$ such that $t < \overline{T}$ for all $t \in T$. Then consider the maximum violation ratio

$$\overline{r} = \sup_{t\in T} \frac{\|\boldsymbol{\zeta}(t)\|}{Me^{-\lambda t}\|\boldsymbol{\zeta}_0\|} \le \frac{B}{Me^{-\lambda\overline{T}}\|\boldsymbol{\zeta}_0\|}$$

Take $\overline{r} = 1$ if T is empty. Then:

$$\|\boldsymbol{\zeta}(t)\| \le \overline{r}Me^{-\lambda t}\|\boldsymbol{\zeta}_0\|$$

implying the trajectory is exponentially stable.

Case 2: There is no upper bound on the elements in T. We will establish $V(\boldsymbol{\zeta})$ is a Lyapunov function certifying exponential stability of the trajectory. Bound the decrease:

$$\begin{aligned}\dot{V}(\boldsymbol{\zeta}) &= -\left(\boldsymbol{\zeta}^\top\mathbf{Q}\boldsymbol{\zeta} + ru^2\right)\\ &\le -\underline{\lambda}(\mathbf{Q})\|\boldsymbol{\zeta}\|^2 \qquad (18)\end{aligned}$$

Next, bound V above by a quadratic function. Because LQR is suboptimal for the nonlinear system, applying it can only increase the cost relative to $V(\boldsymbol{\zeta})$:

$$\begin{aligned}V(\boldsymbol{\zeta}_0) &\le \int_0^\infty \boldsymbol{\zeta}^\top\mathbf{Q}\boldsymbol{\zeta} + r(\mathbf{K}\boldsymbol{\zeta})^2 dt\\ &\le \int_0^\infty (\overline{\lambda}(\mathbf{Q}) + r\overline{\lambda}(\mathbf{K}^\top\mathbf{K}))\|\boldsymbol{\zeta}\|^2 dt\\ &\le \int_0^\infty (\overline{\lambda}(\mathbf{Q}) + r\overline{\lambda}(\mathbf{K}^\top\mathbf{K}))M_{LQR}^2e^{-2\lambda t}\|\boldsymbol{\zeta}_0\|^2 dt\\ &= \frac{(\overline{\lambda}(\mathbf{Q}) + r\overline{\lambda}(\mathbf{K}^\top\mathbf{K}))M_{LQR}^2}{2\lambda_{LQR}}\|\boldsymbol{\zeta}_0\|^2\end{aligned}$$

with $\overline{\lambda}, \underline{\lambda}$ the maximum and minimum eigenvalues respectively. Finally, lower bound $V(\boldsymbol{\zeta})$ by a quadratic.

$$\begin{aligned}V(\boldsymbol{\zeta}_0) &= \int_0^\infty \boldsymbol{\zeta}^\top\mathbf{Q}\boldsymbol{\zeta} + ru^2 dt\\ &\ge \int_T \underline{\lambda}(\mathbf{Q})\|\boldsymbol{\zeta}\|^2 dt\\ &\ge \underline{\lambda}(\mathbf{Q})M^2\|\boldsymbol{\zeta}_0\|^2\left(\int_T e^{-2\lambda t}dt\right)\\ &= \underline{\lambda}(\mathbf{Q})M^2\|\boldsymbol{\zeta}_0\|^2\left(\int_0^\infty e^{-2\lambda t}dt - \int_{\mathbb{R}_{\ge 0}\setminus T} e^{-2\lambda t}dt\right)\\ &= \underline{\lambda}(\mathbf{Q})M^2\left(\frac{1}{2\lambda} - c\right)\|\boldsymbol{\zeta}_0\|^2\end{aligned}$$

where $\mathbb{R}_{\ge 0}\setminus T$ is the set difference between the nonnegative reals and T, and $\frac{1}{2\lambda} - c > 0$ as both integrals integrate over the same strictly positive function, but the right integral does so over a smaller domain. The bounds hold at each point on the trajectory, and V is a Lyapunov function certifying exponential stability of the trajectory.

We now extend this claim over the compact state space $\mathcal{X}$. At $V \succ 0$ and $\dot{V} \prec 0$, we have that the optimal controller is asymptotically stabilizing [21]. By compactness of $\mathcal{X}$ and (18), the time for a trajectory to enter $B_\delta(\mathbf{0})$ is bounded by:

$$T_{\max} = \frac{\sup_{\boldsymbol{\zeta}_0\in\mathcal{X}} V(\boldsymbol{\zeta}_0)}{\inf_{\boldsymbol{\zeta}_0\in\mathcal{X}\setminus B_\delta(\mathbf{0})} \dot{V}(\boldsymbol{\zeta}_0)} \le \frac{\sup_{\boldsymbol{\zeta}_0\in\mathcal{X}} V(\boldsymbol{\zeta}_0)}{\underline{\lambda}(\mathbf{Q})\delta^2}$$

Because trajectories in $B_\delta(\mathbf{0})$ converge exponentially:

$$\|\boldsymbol{\zeta}(t)\| \le Me^{-\lambda(t-T_{\max})}\|\boldsymbol{\zeta}(T_{\max})\| \quad \forall t > T_{\max}$$

By compactness of $\mathcal{X}$, trajectories are bounded by $\|\boldsymbol{\zeta}\| \le B$, and the whole trajectory can be bounded exponentially:

$$\|\boldsymbol{\zeta}(t)\| \le \frac{\max\{M,B\}e^{\lambda T_{\max}}}{\min\{1,\delta\}}e^{-\lambda t}\|\boldsymbol{\zeta}_0\|$$

The optimal controller is exponentially stabilizing. ∎

B. Learning Zero Dynamics Policies

We propose learning $\boldsymbol{\psi}$ so $\mathcal{M}_{\boldsymbol{\psi}}$ is controlled invariant and stable. Let $u^*(\boldsymbol{\zeta})$ solve the optimal control problem (17). Then $\mathcal{M}_{\boldsymbol{\psi}}$ is invariant under $u^*(\boldsymbol{\zeta})$ if for $\boldsymbol{\zeta}_{\boldsymbol{\psi}} = (\boldsymbol{\psi}(\mathbf{z}),\mathbf{z})$:

$$\begin{bmatrix}\mathbf{I} & -\frac{\partial\boldsymbol{\psi}}{\partial\mathbf{z}}\end{bmatrix}\begin{bmatrix}\hat{\mathbf{f}}(\boldsymbol{\zeta}_{\boldsymbol{\psi}}) + \hat{\mathbf{g}}(\boldsymbol{\zeta}_{\boldsymbol{\psi}})u^*(\boldsymbol{\zeta}_{\boldsymbol{\psi}})\\ \boldsymbol{\omega}(\boldsymbol{\zeta}_{\boldsymbol{\psi}})\end{bmatrix} = \mathbf{0}$$

Given a neural network parameterization of $\boldsymbol{\psi}_{\boldsymbol{\theta}}$, we define the loss function:

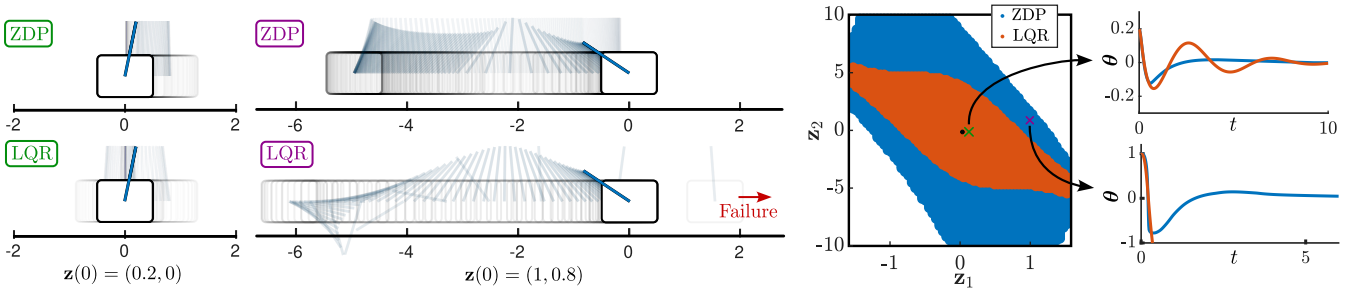

Fig. 2: Zero Dynamics Policies (ZDPs) compared to LQR for the nonlinearly damped cartpole. Left: Simulated initial conditions for the cartpole, where LQR has slow response times and instabilities. Right: Region of attraction ($x = \dot{x} = 0$) for the two methods, as well as pendulum angle over time.

$$\mathcal{L}(\boldsymbol{\theta}) = \mathop{\mathbb{E}}_{\mathbf{z}\sim D} \left\| \hat{\mathbf{f}}(\boldsymbol{\zeta}_{\boldsymbol{\theta}}) + \hat{\mathbf{g}}(\boldsymbol{\zeta}_{\boldsymbol{\theta}}) u^*(\boldsymbol{\zeta}_{\boldsymbol{\theta}}) - \frac{\partial \boldsymbol{\psi}_{\boldsymbol{\theta}}}{\partial \mathbf{z}} \boldsymbol{\omega}(\boldsymbol{\zeta}_{\boldsymbol{\theta}}) \right\| \quad (19)$$

for $\boldsymbol{\zeta}_{\boldsymbol{\theta}} = (\boldsymbol{\psi}_{\boldsymbol{\theta}}(\mathbf{z}), \mathbf{z})$. Zero loss implies invariance of $\mathcal{M}_{\boldsymbol{\psi}}$ under the optimal control which gives stability, by Theorem 3. We minimize this loss using stochastic gradient descent.

C. Application to the Cartpole

We deploy the ZDPs on a classic underactuated system: the cartpole. We add nonlinear damping to the base coordinates of the form $d(\dot{x}) = \sigma(\dot{x})\dot{x}$ where $\sigma(\dot{x}) = 0$ if $|\dot{x}| < 1e^{-3}$ and 1 otherwise. This helps explore the effect of nonlinearities on the degradation of LQR performance, and how our nonlinear method compares. We take $\mathbf{Q} = \mathbf{I}$ and $\mathbf{r} = 0.01$. The ZDP policy was trained in the JAX module using iLQR to approximate u^* and its gradient for training. A 2 layer, 256 neuron feedforward neural network with ReLU activations was pretrained with LQR and then minimized (19). We stabilized $\mathcal{M}_{\boldsymbol{\psi}}$ with a PD controller and a feedforward term. Our code can be found at [22].

Observe the performance of LQR versus ZDPs in Figure 2. Even initial conditions close to the origin and within the domain of attraction of LQR, the modified cartpole's unstable nonlinear damping significantly slowed the controller's by inducing oscillations. In comparison, ZDPs have smoother behavior and a larger region of attraction.

V. Conclusion

We proposed a method of constructing feedback controllers for underactuated systems. We split the controller design process into two steps: 1) learning a manifold that is invariant under optimal control and 2) applying output feedback to this manifold. We proved that such a manifold exists for a broad class of nonlinear systems, and demonstrated the effectiveness of the ZDP method towards stabilizing the canonical underactuated system of the cartpole. Future work includes proving the existence of such a manifold over a larger domain, and applying this method to hardware.

References

[1] S. Sastry, "Linearization by State Feedback," in *Nonlinear Systems: Analysis, Stability, and Control*, ser. Interdisciplinary Applied Mathematics, S. Sastry, Ed. New York, NY: Springer, 1999, pp. 384–448.

[2] E. D. Sontag, "Control-Lyapunov functions," in *Open Problems in Mathematical Systems and Control Theory*, V. Blondel, E. D. Sontag, M. Vidyasagar, and J. C. Willems, Eds. Springer, 1999, pp. 211–216.

[3] A. Isidori, "Elementary Theory of Nonlinear Feedback for Single-Input Single-Output Systems," in *Nonlinear Control Systems*, ser. Communications and Control Engineering, A. Isidori, Ed. London: Springer, 1995, pp. 137–217.

[4] N. S. Bedrossian, "Nonlinear Control Using Linearizing Transformations," 1991.

[5] S. Devasia, Degang Chen, and B. Paden, "Nonlinear inversion-based output tracking," *IEEE TAC*, vol. 41, no. 7, pp. 930–942, Jul. 1996.

[6] C. Tomlin, J. Lygeros, L. Benvenuti, and S. Sastry, "Output tracking for a non-minimum phase dynamic CTOL aircraft model," in *Proceedings of 1995 34th IEEE Conference on Decision and Control*, vol. 2. New Orleans, LA, USA: IEEE, 1995, pp. 1867–1872.

[7] A. Isidori and C. Byrnes, "Output regulation of nonlinear systems," *IEEE TAC*, vol. 35, no. 2, pp. 131–140, Feb. 1990.

[8] E. Westervelt, J. Grizzle, and D. Koditschek, "Hybrid zero dynamics of planar biped walkers," *IEEE Transactions on Automatic Control*, vol. 48, no. 1, pp. 42–56, Jan. 2003, conference Name: IEEE Transactions on Automatic Control.

[9] T. Miki, J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning robust perceptive locomotion for quadrupedal robots in the wild," *Science Robotics*, vol. 7, no. 62, 2022.

[10] J. Siekmann, K. Green, J. Warila, A. Fern, and J. Hurst, "Blind bipedal stair traversal via sim-to-real reinforcement learning," *arXiv preprint arXiv:2105.08328*, 2021.

[11] N. Csomay-Shanklin, V. D. Dorobantu, and A. D. Ames, "Nonlinear model predictive control of a 3d hopping robot: Leveraging lie group integrators for dynamically stable behaviors," in *2023 IEEE (ICRA)*. IEEE, 2023, pp. 12 106–12 112.

[12] R. Grandia, F. Jenelten, S. Yang, F. Farshidian, and M. Hutter, "Perceptive locomotion through nonlinear model-predictive control," *IEEE Transactions on Robotics*, 2023.

[13] S. Kuindersma, R. Deits, M. Fallon, A. Valenzuela, H. Dai, F. Permenter, T. Koolen, P. Marion, and R. Tedrake, "Optimization-based locomotion planning, estimation, and control design for the atlas humanoid robot," *Autonomous robots*, vol. 40, pp. 429–455, 2016.

[14] R. J. Full and D. E. Koditschek, "Templates and anchors: neuromechanical hypotheses of legged locomotion on land," *Journal of experimental biology*, vol. 202, no. 23, pp. 3325–3332, 1999.

[15] M. H. Raibert, "Hopping in legged systems—modeling and simulation for the two-dimensional one-legged case," *IEEE Transactions on Systems, Man, and Cybernetics*, no. 3, pp. 451–463, 1984.

[16] I. D. J. Rodriguez, N. Csomay-Shanklin, Y. Yue, and A. D. Ames, "Neural Gaits: Learning Bipedal Locomotion via Control Barrier Functions and Zero Dynamics Policies," in *Proceedings of The 4th Annual L4DC*. PMLR, May 2022, pp. 1060–1072, iSSN: 2640-3498.

[17] E. D. Sontag, "A 'universal'construction of artstein's theorem on nonlinear stabilization," *Systems & control letters*, vol. 13, no. 2, pp. 117–123, 1989.

[18] A. D. Ames, K. Galloway, K. Sreenath, and J. W. Grizzle, "Rapidly exponentially stabilizing control lyapunov functions and hybrid zero dynamics," *IEEE TAC*, vol. 59, no. 4, pp. 876–891, 2014.

[19] M. Spivak, *Calculus on manifolds: a modern approach to classical theorems of advanced calculus*, ser. Mathematics monograph series. CRC Press, Taylor & Francis Group, 2018.

[20] K. J. Astrom and R. M. Murray, *Feedback Systems: An Introduction for Scientists and Engineers*. USA: Princeton University Press, 2008.

[21] D. Liberzon, *Calculus of Variations and Optimal Control Theory: A Concise Introduction*. Princeton University Press, 2012.

[22] "Code," 2024. [Online]. Available: https://github.com/ivandariojr/LearnedZeroDynamicsPolicies