Intensity—A Metric Approach to Quantifying Attractor Robustness in ODEs[*]

Katherine J. Meyer[†] and Richard P. McGehee[‡]

Abstract. Although mathematical models do not fully match reality, robustness of dynamical objects to perturbation helps bridge from theoretical to real-world dynamical systems. Classical theories of structural stability and isolated invariant sets treat robustness of qualitative dynamics to sufficiently small errors. But they do not indicate just how large a perturbation can become before the qualitative behavior of our system changes fundamentally. Here we introduce a quantity, intensity of attraction, that measures the robustness of attractors in metric terms. Working in the setting of ordinary differential equations on $\mathbb{R}^n$, we consider robustness to vector field perturbations that are time dependent or independent. We define intensity in a control-theoretic framework, based on the magnitude of control needed to steer trajectories out of a domain of attraction. Our main result is that intensity also quantifies the robustness of an attractor to time-independent vector field perturbations; we prove this by connecting the reachable sets of control theory to isolating blocks of Conley theory. In addition to treating classical questions of robustness in a new metric framework, intensity of attraction offers a novel tool for resilience quantification in ecological applications. Unlike many measurements of resilience, intensity detects the strength of transient dynamics in a domain of attraction.

Key words. ordinary differential equations, resilience, transient dynamics, control, Conley theory, attractors

AMS subject classifications. 34D30, 37C70, 92D40, 93B03

DOI. 10.1137/20M138689X

1. Introduction. Although differential equations are imperfect models of reality, they capture the essential dynamics of many real-world systems remarkably well. This utility of dynamical systems theory stems in part from studying robust dynamical objects. For example, classical results about structurally stable systems [1, 23, 31] and isolated invariant sets [9, 8] guarantee that qualitative dynamics persist through sufficiently small errors, under suitable conditions of nondegeneracy. In this work we push beyond "sufficiently small" to explore just how wrong our models can be while still providing meaningful information about the long-term behavior of a system. More precisely, we ask in metric terms how different a vector field can become while retaining—in some sense—an original attractor.

To answer this question, we adapt the concept of intensity of attraction from the setting of maps [19] and relations [33] to the setting of vector fields on $\mathbb{R}^n$. We define the intensity of an attractor in section 5. What makes intensity so well suited to our purpose is that it

  [†]Department of Mathematics and Statistics, Carleton College, 1 North College St., Northfield, MN 55057 USA (kjmeyer@carleton.edu).
  [‡]Department of Mathematics, University of Minnesota, 206 Church St. SE, Minneapolis, MN 55455 USA (mcgehee@umn.edu).

(i) quantifies how far an attractor persists in the context of both time-independent and time-varying changes to a vector field $f : \mathbb{R}^n \to \mathbb{R}^n$;

(ii) can be computed; and

(iii) is well-defined for ODE systems of any finite dimension.

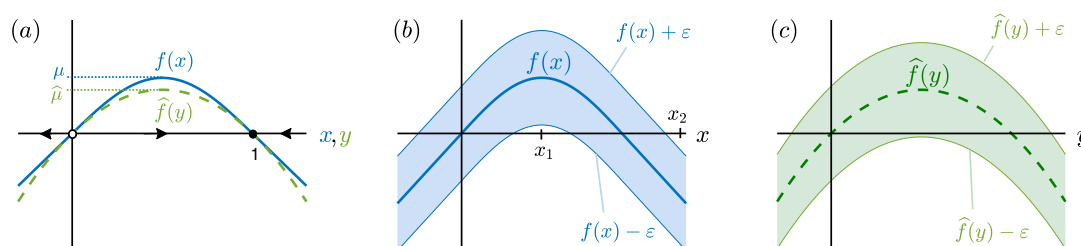Our main result, Theorem 6.11, links the time-dependent and time-independent vector field perturbations in (i): roughly speaking, if we define the intensity μ of an attractor A for a vector field f as the magnitude of time-varying control needed to escape A's domain of attraction, then the attractor continues in a related form through any time-*independent* perturbations to f that stay below μ in an appropriate metric. This connection reduces the problem of attractor persistence in an infinite-dimensional space to the well-studied problem of computing reachable sets in a control framework [2, 4, 10, 25, 26, 27, 29].

In addition to quantifying tolerable modeling errors, intensity of attraction provides a new tool for resilience quantification for applications in ecology and other fields [5, 20, 21]. In particular, intensity measures the resilience of a particular regime—modeled as a domain of attraction—to continuous-time, exogeneous disturbances. Unlike common resilience metrics like eigenvalues or basin size [20] that are based solely on invariant sets, intensity detects the strength of attracting dynamics over transient regions of state space.

We illustrate some of the preceding points about intensity of attraction in a simple example before outlining the paper's full contents.

Example 1.1. Suppose the differential equations $x' = f(x)$ and $y' = \widehat{f}(y)$ model the dynamics of biological populations x and y. The vector fields $f : \mathbb{R} \to \mathbb{R}$ and $\widehat{f} : \mathbb{R} \to \mathbb{R}$ graphed in Figure 1(a) generate topologically equivalent dynamical systems, each with a repelling equilibrium at the origin and an attracting equilibrium at 1. In the absence of error or disturbance, we expect each biological population to approach a carrying capacity at 1. But this topological picture misses a metric difference between the two models: f lies above $\widehat{f}$ over the transient interval $(0, 1)$ of state space. The intensity of each attractor $\{x = 1\}$ and $\{y = 1\}$ captures this metric difference: in this case, intensity boils down to the maximum values μ and $\widehat{\mu}$ that the functions f and $\widehat{f}$ achieve, respectively, over the interval $(0, 1)$.

The fact that the intensity μ of $\{x = 1\}$ exceeds the intensity $\widehat{\mu}$ of $\{y = 1\}$ holds implications for both population resilience and model interpretation. If we take the models at face

Figure 1. *Motivating example.* (a) *Vector fields f and $\widehat{f}$ generate topologically equivalent flows on $\mathbb{R}$, illustrated by the horizontal phase line. Intensities μ and $\widehat{\mu}$ of the mutual attractor at 1 capture metric differences between the dynamics.* (b) *If $f(x)$ is known to an error tolerance of ϵ with $0 < \epsilon < \mu$ then we have $f(x_1) \pm \epsilon > 0$ and $f(x_2) \pm \epsilon < 0$, suggesting that some trajectories $x(t)$ remain bounded and positive.* (c) *Under the same error tolerance ϵ on $\widehat{f}(y)$, we have $\widehat{f}(y) - \epsilon < 0$ for all $y \in \mathbb{R}$, so it is possible that no trajectories stay positive.*

value, then $\mu > \widehat{\mu}$ suggests that population x is more resilient to external time-dependent forcings (such as climate variability or seasonal harvests, as in [21]) than population y, based on its higher rebound rate at depressed population levels. Notably, two common resilience quantifiers based on invariant sets fail to detect this difference: f and $\widehat{f}$ share the same distance between equilibria and the same eigenvalues (slopes) at equilibria.

On the other hand, suppose models $x' = f(x)$ and $y' = \widehat{f}(y)$ originated from data on real-world populations x and y, and the values of f and $\widehat{f}$ are known to an accuracy of ε, where $\widehat{\mu} < \varepsilon < \mu$. As illustrated in Figure 1(b), we would still have topological grounds to believe that population x can remain bounded above zero in the long term, because $f(x_1) \pm \varepsilon > 0$ and $f(x_2) \pm \varepsilon < 0$. But under the same error tolerance ϵ on $\widehat{f}(y)$, we have $\widehat{f}(y) - \epsilon < 0$ for all $y \in \mathbb{R}$ (Figure 1(c)), so it is possible that no trajectories for population y remain bounded above zero. In this way, the intensities μ and $\widehat{\mu}$ measure the degree to which each attracting equilibrium ($x = 1$ or $y = 1$) persists to structure system behavior in the face of changes to the vector field—whether these are extrinsic, time-varying perturbations to population growth or intrinsic, time-independent changes to the model equations.

In the body of this paper, we generalize the notion of intensity from Example 1.1 to an analogue for higher dimensional systems that still quantifies attractor persistence and is computable. The organization is as follows.

Section 2 outlines the necessary preliminaries regarding $\mathbb{R}^n$ as a metric space and the dynamics generated by an autonomous differential equation $x' = f(x)$, $x \in \mathbb{R}^n$.

In section 3 we introduce time-dependent perturbations to the vector field f via control functions $g(t)$. In this work control functions $g(t)$ play a role analogous to discrete errors in the theory of intensity of attraction for maps [19]. In particular, studying the dynamics of $x' = f(x) + g(t)$ allows us to probe outwards from an attracting set that would otherwise be invariant under f into the transient regions of its domain of atttraction. In contrast to recent work on Morse decompositions for nonautonomous differential equations [24], we do not study attractors of $x' = f(x) + g(t)$. Instead, we use time-varying perturbations to reveal information about the attractors of $x' = f(x)$. We take controls g to be essentially bounded, allowing quantification of resilience to both continuous and discontinuous perturbations. The L^∞ norm $\|g\|_\infty$ reflects the norm in play on $\mathbb{R}^n$. In the sections that follow we quantify how large $\|g\|_\infty$ must become in order to escape the domain of attraction of an attractor under f—this is the essential ingredient for defining intensity of attraction.

Towards this end, section 4 broadens our attention from individual control functions to bounded collections of control functions. Given a set $S \subset \mathbb{R}^n$ and a bound r on control, we ask which points in state space can be reached from S using some control g with $\|g\|_\infty \leq r$. The answer is called a reachable set, a term that carries a compatible and more general meaning in the theory of control [7, 16] and differential inclusions [15, 32]. Indeed, it is possible to reframe the collection of differential equations $x' = f(x) + g(t)$, $\|g\|_\infty \leq r$ as the differential inclusion $x' \in \overline{B_r(f(x))}$, where $B_r(f(x)) = \{h \in L^\infty \mid \|h - f(x)\|_\infty \leq r\}$. Although our focus is not on the dynamics of this differential inclusion, it is interesting to note that the set reachable from an attractor for $x' = f(x)$ is invariant in the context of the differential inclusion, and in many cases fits the definition of attractor developed in [15, 32].

Section 5 defines intensity of attraction in terms of reachable sets: roughly, intensity reflects the largest value of r for which sets reachable from an attractor are bounded inside

its domain of attraction. We connect intensity to resilience of predator-prey dyanmics in Example 5.6. Additional examples revisit Figure 1 with greater precision (Example 5.3), illustrate how the value of intensity depends on the metric used to quantify perturbations (Example 5.4), and show how intensity relates to Colonius and Kliemann's point of first discontinuity in a map from control bounds to reachable sets [7] (Example 5.5).

Section 6 demonstrates that intensity of attraction also quantifies attractor persistence in the face of time-independent perturbations to a vector field. Theorem 6.11 gives the main result that intensity bounds from below the distance in the space of vector fields that an attractor continues immediately. The key to this result is a link between reachable sets and the isolating blocks of Conley theory, which we develop in subsections 6.1 and 6.2. Example 6.12 illustrates continuation of a predator-prey limit cycle through time-independent perturbations that don't exceed its intensity. Last, we leverage the present setting to provide a new proof that attractors are upper semicontinuous (Theorem 6.15).

Section 7 makes some concluding remarks and identifies directions for further study.

2. Preliminaries.

2.1. $\mathbb{R}^n$ as a metric space. Let $d : \mathbb{R}^n \times \mathbb{R}^n \to [0, \infty)$ be a homogeneous and translation-invariant metric on $\mathbb{R}^n$. Let $\|\cdot\|$ denote the corresponding norm with $\|x\| = d(x, 0)$ and $d(x, y) = \|x - y\|$. Though we use the Euclidean metric and norm in most examples, we intentionally develop the theory for general d and $\|\cdot\|$. In any particular application, context should inform the choice of metric and norm.

The following definitions and notations are standard in this metric setting. We say a set $X \subset \mathbb{R}^n$ is bounded if there exists an $M > 0$ such that $\|x\| \leq M$ for all $x \in X$. The distance between two sets $X, Y \subset \mathbb{R}^n$ is

$$\mathrm{dist}(X, Y) \equiv \inf\{d(x, y) \mid x \in X, y \in Y\}.$$

A single-point set $\{x\}$ is not distinguished from its element x. For $\epsilon > 0$ and $x \in \mathbb{R}^n$, we denote the ϵ-ball about x as

$$B_\epsilon(x) \equiv \{y \in \mathbb{R}^n \mid d(x, y) < \epsilon\}.$$

More generally, we denote the ϵ-neighborhood of a set $S \subset \mathbb{R}^n$ as

$$B_\epsilon(S) \equiv \{x \in \mathbb{R}^n \mid d(x, S) < \epsilon\}.$$

The metric topology is used throughout. For $S \subset \mathbb{R}^n$, $\mathrm{int}(S)$ and $\overline{S}$ denote the interior and closure of S, respectively. A set N is a neighborhood of S if $S \subset \mathrm{int}(N)$. The following lemma states that a compact set has arbitrarily close compact neighborhoods; its straightforward proof is omitted.

Lemma 2.1. *If $S \subset \mathbb{R}^n$ is compact, then for any neighborhood N of S there exists a compact set $K \subset \mathrm{int}(N)$ that is also a neighborhood of S.*

2.2. Dynamics generated by $x' = f(x)$. Consider an autonomous system of ordinary differential equations

$$(2.1) \qquad\qquad\qquad\qquad x' = f(x),$$

where $x \in \mathbb{R}^n$, x' denotes $\frac{dx}{dt}$, and f maps from an open set $U \subset \mathbb{R}^n$ to $\mathbb{R}^n$. We assume that the vector field f in the differential equation (2.1) is bounded, so that $\|f\|_{\sup} \equiv \sup_{x \in U} \|f(x)\|$ exists. We also take f to be globally Lipschitz on U, so that there exists a unique local solution $x(t)$ in U to each initial value problem

$$x' = f(x), \quad x(0) = x_0 \in U.$$

The solutions $x(t)$ generate a local flow φ on a subset D of $\mathbb{R} \times U$ given by $\varphi(t, x_0) = x(t)$. Although trajectories may in general leave U, preventing definition of a global flow $\varphi : \mathbb{R} \times U \to U$, we assume that φ is defined on any time domain of interest.

Fixing the time coordinate of φ yields a time-t map $\varphi^t : U \to U$, given by $x \mapsto \varphi(t, x)$. The collection of time-t maps for $t \in \mathbb{R}$ has the identity property $\varphi^0 = \mathrm{id}$ and the group property $\varphi^t \circ \varphi^s = \varphi^{t+s}$. Functions such as the flow and the time-t map are extended in the natural way to take sets as arguments. For example, if $T \subset \mathbb{R}$ and $S \subset \mathbb{R}^n$, then $\varphi(T, S) \equiv \bigcup_{t \in T} \bigcup_{x \in S} \varphi(t, x)$ and $\varphi^t(S) \equiv \bigcup_{x \in S} \varphi^t(x)$. We next put this notation to work in a few definitions that build to attractors, our main object of study.

Definition 2.2. *A set $S \subset U$ is* forward invariant *under the flow φ if $\varphi^t(S) \subset S$ for all $t > 0$. S is* invariant *if $\varphi^t(S) = S$ for all $t \in \mathbb{R}$.*

Definition 2.3. *The* omega limit set *of a set $S \subset U$ is $\omega(S) \equiv \bigcap_{T \geq 0} \overline{\bigcup_{t \geq T} \varphi^t(S)}$.*

Definition 2.4. *An* attractor *is a compact, nonempty, invariant set A such that $A = \omega(N)$ for some neighborhood N of A.*

The statement $A = \omega(N)$ indicates that the neighborhood N is "attracted" asymptotically to A in forward time. We will also use the following equivalent characterization of an attractor; the proof is provided in supplementary section SM1.

Lemma 2.5. *A compact, nonempty, invariant set A is an attractor if and only if there exists $\gamma > 0$ such that $d_H(\varphi^t(B_\gamma(A)), A) \to 0$ as $t \to \infty$, where d_H is Hausdorff distance.*

We call the collection of *all* points that tend to A in forward time A's domain of attraction.

Definition 2.6. *The* domain of attraction *of an attractor A is*

$$\mathcal{D}(A) = \{x \in U \mid \varnothing \neq \omega(x) \subset A\}.$$

Our primary objective is to quantify what we'll call the intensity of an attractor A, which should in some way reflect the strength of transient dynamics over $\mathcal{D}(A)$. To test that strength, we probe the domain of attraction with time-dependent control (section 3).

3. Time-dependent control functions. We introduce time-dependent perturbations to the system (2.1) via nonautonomous control functions g, yielding a perturbed system

$$(3.1) \qquad\qquad x' = f(x) + g(t).$$

The control functions g are integrable and essentially bounded with respect to a measure m; that is, given a time interval $I \subset \mathbb{R}$ we take $g \in L^\infty(I, \mathbb{R}^n)$, equipped with the norm

$$\|g\|_\infty \equiv \inf\{c \geq 0 \ : \ m(\{x \in I : \|g(x)\| > c\}) = 0\}.$$

Note that the value of $\|g\|_\infty$, which plays an important role in sections 4 to 6, depends on the norm used in $\mathbb{R}^n$.

With $g \in L^\infty(I, \mathbb{R}^n)$, the right-hand side of (3.1) satisfies the Carathéodory conditions on $U \times I$: the function $F(x,t) = f(x) + g(t)$ is measurable in t for fixed x, continuous in x for fixed t, and (without loss of generality) bounded on any compact subset of $I \times U$ by the constant and hence integrable function $\|f\|_{\sup} + \|g\|_\infty$. Furthermore, F inherits f's Lipschitz property in x. It follows (see, e.g., [6, 11]) that for any initial condition $x(t_0) = x_0$ there exists a unique local solution $x(t)$ to (3.1), in the extended sense that

$$(3.2) \qquad x(t) = x(t_0) + \int_{t_0}^t \left(f(x(s)) + g(s) \right) ds$$

$$(3.3) \qquad \text{and} \qquad x'(t) = f(x) + g(t) \text{ almost everywhere.}$$

The characterization (3.2) of solutions provides an analytic tool, while (3.3) highlights the meaning of $\|g\|_\infty$ as a maximum deviation from the vector field f. By default we assume that a unique local solution $x(t)$ to (3.1) can be continued on any desired time interval I.

Existence and uniqueness of solutions to (3.1) permit extension of standard flow notation to include the control function g in a manner similar to [7].

Definition 3.1. *For a given vector field f and time interval $[0,T]$, the controlled flow*

$$\varphi : [0,T] \times U \times L^\infty([0,T], \mathbb{R}^n) \to \mathbb{R}^n$$

is given by $\qquad (\hat{t}, x_0, g) \mapsto \varphi(\hat{t}, x_0, g; f) = x(\hat{t}),$

where $x(t)$ is the solution to $x' = f(x) + g(t)$, $x(0) = x_0$.

We omit the vector field f when it is clear from context. To emphasize dependence on initial conditions, we write $\varphi(t, \cdot, g) = \varphi_g^t(\cdot)$. In what follows, we move fluidly between describing solutions to (3.1) with the notation $x(t)$ and with the flow notation of Definition 3.1. The latter has the advantage of making g visible.

The controlled flow $\varphi(t, x, g)$ is continuous not only with respect to t and x (see [12]) but also with respect to control input g.

Lemma 3.2. *Fix a vector field $f : U \to \mathbb{R}^n$ and a time interval $[0,T]$, and let g and h be control functions in $L^\infty([0,T], \mathbb{R}^n)$. Then for any ϵ, there exists a δ such that $\|g - h\|_\infty < \delta$ implies $\|\varphi(t, x, g) - \varphi(t, x, h)\| < \epsilon$ for all $t \in [0,T]$ and all $x \in U$.*

We omit the standard proof of Lemma 3.2, which follows directly by applying Grönwall's inequality to the function $u(t) \equiv \|\varphi(t, x, g) - \varphi(t, x, h)\|$. See also Lemma 4.3.2 of [7].

4. Reachable sets. The controlled flow of Definition 3.1 gives outcomes under a single function $g \in L^\infty(I, \mathbb{R}^n)$, which highlights the effect of a known intervention in a controlled system. However, when g represents an uncertain disturbance, it is natural to consider trajectories corresponding to an entire family of such functions. We restrict our attention to families of the form

$$B_r(0) \equiv \{ g \in L^\infty(I, \mathbb{R}^n) \mid \|g\|_\infty \leq r \}.$$

By using the collection $B_r(0)$, we focus on situations in which bounds on disturbances are known, even though their exact forms are not. This approach is compatible with studies

of random dynamical systems with bounded noise (e.g., [13, 14, 35]) and with differential inclusions [4, 15, 32]. But in contrast to systems with unbounded noise such as stochastic differential equations, trajectories cannot necessarily reach any state from another.

We use $P_{r;f}(S)$ to denote the collection of states that can be reached in forward time from a point in $S \subset U$ by modifying the vector field f with r-bounded control.

Definition 4.1. $P_{r;f}(S) = \bigcup_{t \geq 0} \bigcup_{x \in S} \bigcup_{\|g\|_\infty \leq r} \varphi(t, x, g; f)$.

When no confusion could result by omitting the vector field f, we write $P_r(S)$. In the present work we will consider sets reachable from an attractor A. These can be found easily for one-dimensional systems, as the following proposition shows.

Proposition 4.2. *Consider the one-dimensional system $x' = f(x)$ with $f(0) = 0$, $f > 0$ on $(-\epsilon, 0)$, and $f < 0$ on $(0, \epsilon)$ for some $\epsilon > 0$. The reachable set from the attractor $A = \{0\}$ is $P_r(A) = (a, b)$ where*

$$a = \begin{cases} \sup\{x < 0 \mid f(x) = r\} & \text{if } f(x) = r \text{ for some } x < 0, \\ -\infty & \text{otherwise,} \end{cases}$$

$$\text{and} \quad b = \begin{cases} \inf\{x > 0 \mid f(x) = -r\} & \text{if } f(x) = -r \text{ for some } x > 0, \\ \infty & \text{otherwise.} \end{cases}$$

The natural way to think about Proposition 4.2 is by comparison to the most extreme controls possible, $g(t) \equiv \pm r$. See supplementary section SM2 for a proof.

In the case of a one-dimensional linear system, Proposition 4.2 implies that the set reachable from the origin scales in direct proportion to the control bound and in inverse proportion to the eigenvalue magnitude.

Corollary 4.3. *For the one-dimensional, linear system $x' = -\lambda x$ ($\lambda > 0$) with global attractor $A = 0$, the reachable set from A for $r > 0$ is $P_r(A) = (-\frac{r}{\lambda}, \frac{r}{\lambda})$.*

Reachable sets are, in general, harder to compute than in Proposition 4.2 and Corollary 4.3. The following example shows that the computation can be nontrivial even for a diagonal, two-dimensional linear system.
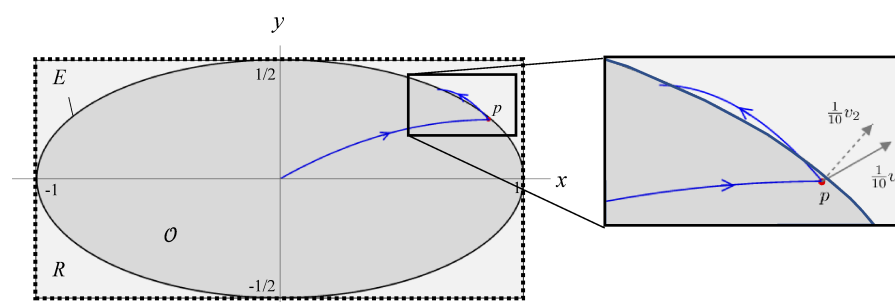
Example 4.4. The system

$$(4.1) \qquad\qquad\qquad\qquad\qquad x' = -x,$$
$$(4.2) \qquad\qquad\qquad\qquad\qquad y' = -2y$$

on $\mathbb{R}^2$ has global attractor $A = (0, 0)$. Considering each variable separately, Corollary 4.3 indicates that $P_1(A)$ is contained in the rectangular region $R = (-1, 1) \times (-\frac{1}{2}, \frac{1}{2})$. Furthermore, the family of constant control functions $\{g_{\theta,c}(t) = c \cos \theta + c \sin \theta \mid \theta \in [0, 2\pi), 0 \leq c \leq 1\}$ sends the origin asymptotically to points on ellipses $x^2 + 4y^2 = c^2$. Therefore, the reachable set $P_1(A)$ contains the open region $\mathcal{O}$ enclosed by the ellipse $E = \{(x, y) \mid x^2 + 4y^2 = 1\}$. Control functions with magnitude 1 directed opposite to the vector field also drive trajectories asympototically from $\mathcal{O}$ to E.

Because the vector field has magnitude 1 on E, one might expect that $P_1(A) = \mathcal{O}$. But the vector field is not normal to E except at the vertices of the ellipse. Therefore, adjusting

Figure 2. *Inexact bounds on a two-dimensional reachable set. The set in question consists of points reachable from the origin in the linear system $x' = -x$, $y' = -2y$ with control bound $r = 1$. This set must be nested between the open set $\mathcal{O}$ and the rectangle R. Piecewise-constant control that shifts directions from v_1 to v_2 drives a trajectory (blue) outside of $\mathcal{O}$.*

a control function's direction to push perpendicular to E can drive a solution outside $\mathcal{O}$. Figure 2 illustrates this phenomenon for piecewise constant control

$$(4.3) \qquad\qquad g(t) = \begin{cases} v_1 & \text{if } 0 \le t < 10, \\ v_2 & \text{if } 10 \le t, \end{cases}$$

where $v_1 = 0.99 \begin{bmatrix} \sqrt{3}/2 \\ 1/2 \end{bmatrix}$ and $v_2 = 0.99 \begin{bmatrix} \sqrt{3/7} \\ 2/\sqrt{7} \end{bmatrix}$ have magnitude 0.99 but different directions. We therefore have the bounds $\mathcal{O} \subsetneq P_1(A) \subset R$. Athough these bounds follow directly from analytic computation of trajectories and Corollary 4.3, a closed form for $P_1(A)$ does not.

Numerical algorithms for computing reachable sets are an active topic of research in the fields of control systems, differential inclusions, and optimization. The introductory piece [16] outlines fundamental discretization challenges nicely. Existing approaches include set-valued Euler schemes [4, 26, 27, 29], optimal control algorithms [25, 2], and level-set methods ([10] and references therein).

Though the approximation of reachable sets can be involved, analytically they obey nice nesting properties. The following lemma, which follows immediately from Definition 4.1, states that the real number ordering of control bounds induces an order by inclusion on reachable sets.

Lemma 4.5. *If $r < r'$, then $P_r(S) \subset P_{r'}(S)$.*

A stronger statement holds when the set S is an attractor A.

Proposition 4.6. *For A an attractor, $P_r(A) \searrow A$ as $r \searrow 0$, in the sense that*
(1) $r < r'$ implies $P_r(A) \subset P_{r'}(A)$;
(2) for any neighborhood V of A, there is an $r > 0$ such that $P_r(A) \subset V$; and
(3) $A = \bigcap_{r>0} P_r(A)$.

Proof. Property (1) is a direct instance of Lemma 4.5. Towards property (2), fix a neighborhood V of A and choose $\beta_1 > 0$ so that $B_{\beta_1}(A) \subset V$. By Lemma 2.5 there exists $\beta_2 > 0$ such that $d_H(\varphi^t(B_{\beta_2}(A)), A) \to 0$ as $t \to \infty$. Let $\beta = \min\{\beta_1, \beta_2\}$. It follows that

$d_H(\varphi^t(B_\beta(A)), A) \to 0$ as $t \to \infty$, so there exists $T_0 > 0$ such that $d_H(\varphi^t(B_\beta(A)), A) < \beta/2$ whenever $t \geq T_0$. We thus have $\varphi^t(B_\beta(A)) \subset B_{\beta/2}(A)$ for all $t \geq T_0$. Next choose $\gamma > 0$ such that $B_\gamma(A) \subset \varphi^{T_0}(B_\beta(A)) \subset B_{\beta/2}(A)$. Since $B_\gamma(A) \subset B_{\beta_2}(A)$, we have as before that $d_H(\varphi^t(B_\gamma(A)), A) \to 0$ as $t \to \infty$ and there exists $T_1 > 0$ such that $\varphi^t(B_\gamma(A)) \subset B_{\gamma/2}(A)$ for all $t \geq T_1$.

Applying Lemma 3.2 with $\epsilon \equiv \gamma/2$ and $T \equiv T_1$ gives a $\delta > 0$ such that $\|g\|_\infty < \delta$ implies $\|\varphi_g^t(x) - \varphi^t(x)\| < \gamma/2$ for all $t \in [0, T_1]$ and for all x in state space. We now show that choosing $0 < r < \delta$ ensures $P_r(A) \subset V$. Let g be a perturbation of size $\|g\|_\infty \leq r < \delta$. We have for $t \in [0, T_1]$ that

$$\begin{aligned}
\varphi_g^t(B_\gamma(A)) &\subset B_{\gamma/2}\left(\varphi^t(B_\gamma(A))\right) && \text{(Lemma 3.2)} \\
&\subset B_{\gamma/2}(\varphi^t(\varphi^{T_0}(B_\beta(A)))) && (B_\gamma(A) \subset \varphi^{T_0}(B_\beta(A))) \\
&= B_{\gamma/2}(\varphi^{t+T_0}(B_\beta(A))) && \text{(group property of flow)} \\
&\subset B_{\gamma/2}(B_{\beta/2}(A)) && (t + T_0 \geq T_0) \\
&\subset B_{3\beta/4}(A) && (\gamma < \beta/2) \\
&\subset B_\beta(A) && \\
&\subset V. &&
\end{aligned}$$

Furthermore, at the timepoint T_1,

$$\begin{aligned}
\varphi_g^{T_1}(B_\gamma(A)) &\subset B_{\gamma/2}\left(\varphi^{T_1}(B_\gamma(A))\right) && \text{(Lemma 3.2)} \\
&\subset B_{\gamma/2}\left(B_{\gamma/2}(A)\right) && (T_1 \geq T_1) \\
&\subset B_\gamma(A). &&
\end{aligned}$$

From the previous two chains of inclusions, we have that the forward images of the neighborhood $B_\gamma(A)$ remain inside V during the controlled flow over time $[0, T_1]$ and the forward image of $B_\gamma(A)$ returns to be contained within $B_\gamma(A)$ at time T_1. An inductive argument over the time intervals $[nT_1, (n+1)T_1]$ implies that $\varphi_g^t(B_\gamma(A)) \subset V$ for all $t \geq 0$. Since $A \subset B_\gamma(A)$, this implies $\varphi_g^t(A) \subset V$ for all $t \geq 0$ and all $\|g\|_\infty \leq r$. By Definition 4.1 we have $P_r(A) \subset V$, establishing property (2).

The inclusion $A \subset \bigcap_{r>0} P_r(A)$ of property (3) is immediate. The reverse inclusion follows from property (2): any point not in A can be excluded from some neighborhood of A and hence from $P_r(A)$ for some $r > 0$. Hence any point in $P_r(A)$ for all $r > 0$ must be in A. This completes the proof. ∎

5. Intensity of attraction. Intensity of attraction, introduced by McGehee for maps [19], carries over naturally to the flow setting by replacing sets reachable from an attractor under ϵ pseudo-orbits with sets reachable from an attractor under r-bounded control.

Definition 5.1. *The* intensity of attraction *of an attractor A is*

(5.1) $$\mu(A) = \sup\{r \geq 0 : P_r(A) \subset K \subset \mathcal{D}(A) \text{ for some compact } K \subset \mathbb{R}^n\}.$$

In other words, intensity of attraction reflects the control magnitude $\|g\|_\infty$ necessary to escape from every compact subset of an attractor's domain.

Proposition 5.2. *For A an attractor, $\mu(A) > 0$.*

Proof. Let A be an attractor and let N be a neighborhood of A such that $\omega(N) = A$, which implies via Definition 2.6 that $N \subset \mathcal{D}(A)$. By Lemma 2.1 there exists a compact neighborhood $K \subset \mathbb{R}^n$ such that $A \subset K \subset N \subset \mathcal{D}(A)$. Taking K as V in Proposition 4.6, we get an $r > 0$ such that $P_r(A) \subset K \subset \mathcal{D}(A)$. Definition 5.1 then implies that $\mu(A) > 0$. ∎

The following two examples, though somewhat synthetic, serve to illustrate two points about attractor intensity: it is independent of eigenvalues and basin size (Example 5.3), and its value depends on the norm $\|\cdot\|$ in play on $\mathbb{R}^n$ (Example 5.4).

Example 5.3. It is possible for two vector fields that share the same attracting equilibrium A, the same eigenvalues at A, and the same domain of attraction $\mathcal{D}(A)$ to nonetheless exhibit different intensities of attraction $\mu(A)$. Consider the pair of $C^1(\mathbb{R}, \mathbb{R})$ vector fields plotted in Figure 1:

$$
\text{(5.2)} \qquad\qquad\qquad f(x) = x - x^2,
$$

$$
\text{(5.3)} \qquad\qquad \text{and} \quad \hat{f}(x) = \begin{cases} x & \text{if } x < 0, \\ \pi^{-1}\sin(\pi x) & \text{if } 0 \le x < 1, \\ 1 - x & \text{if } x \ge 1. \end{cases}
$$

The equilibrium at $x = 1$ is an attractor for both f and $\hat{f}$, with eigenvalue -1. Its domain of attraction is $(0, \infty)$ in each case. In fact, the eigenvalue at the repelling equilibrium 0 that consitutes the boundary of $\mathcal{D}(A)$ is also 1 for both f and $\hat{f}$. Yet $\mu(A) = 1/4$ under f while $\hat{\mu}(A) = \pi^{-1} \approx 0.318$ under $\hat{f}$. This example alerts us that eigenvalues and size of a domain of attraction—two common measures of resilience in engineering and ecology [20]—can neglect differences in the strength of transient dynamics over $\mathcal{D}(A)$. Intensity of attraction detects this feature.

Example 5.4. To easily see how $\mu(A)$ depends on norm choice, consider the system

$$
\text{(5.4)} \qquad\qquad\qquad x' = \frac{\sqrt{2}}{4}(x+y)^2 - x,
$$

$$
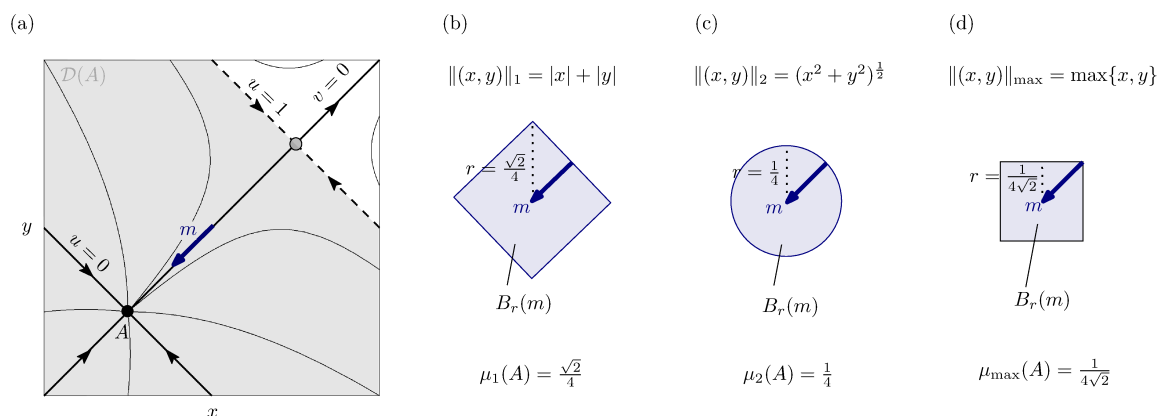\text{(5.5)} \qquad\qquad\qquad y' = \frac{\sqrt{2}}{4}(x+y)^2 - y,
$$

whose phase portrait is given in Figure 3(a). Rotating coordinates by $\pi/4$ to $u = \frac{\sqrt{2}}{2}(x+y)$ and $v = \frac{\sqrt{2}}{2}(y-x)$ gives

$$
\text{(5.6)} \qquad\qquad\qquad u' = u(u-1),
$$

$$
\text{(5.7)} \qquad\qquad\qquad v' = -v.
$$

In these coordinates it is clear that $A = (0,0)$ is an attractor with domain of attraction $\mathcal{D}(A) = \{(u,v) \mid u < 1\}$. The uncoupled form of (5.6) and (5.7) and the geometry of $\mathcal{D}(A)$ allow us to restrict attention to the invariant line $v = 0$ when calculating $\mu(A)$. On the
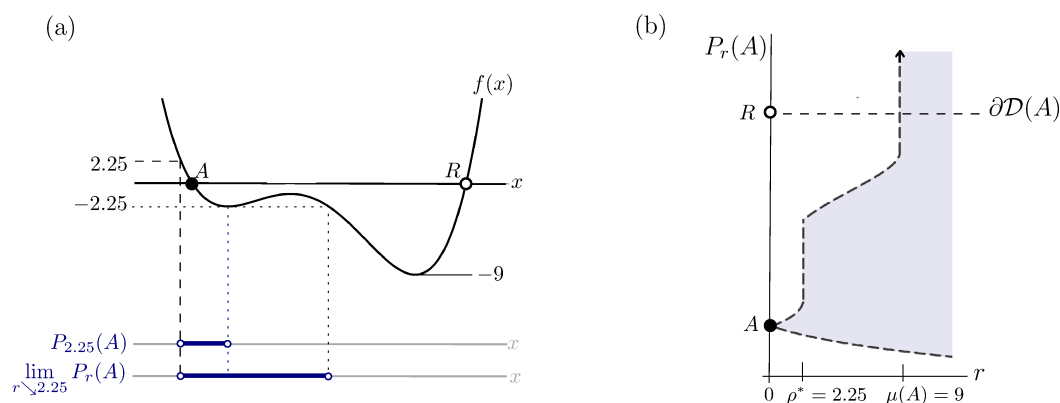
Figure 3. *Intensity depends on norm.* (a) *Phase portrait for the system* (5.4) *and* (5.5). *A is an attracting equilibrium and the dashed line marks the boundary of its domain of attraction, shaded gray. Analysis in rotated coordinates (*(5.6) *and* (5.7)*) reveals that the intensity of A is the norm of the vector m.* (b), (c), (d): *The 1-norm, 2-norm, and max norm each yield different values of intensity of attraction for A.*

interval $0 \leq u \leq 1$, the maximum strength of the vector field in the direction of the origin is $1/4$, achieved by the vector m at $u = 1/2$. A control function must push with magnitude greater than $1/4$ in the direction of $-m$ to steer a trajectory from A to $\mathcal{D}(A)^C$. The critical control bound needed for this push, $\mu(A)$, depends on the norm on (x, y) space. Under the Euclidean norm (the 2-norm), the intensity is $\mu_2(A) = 1/4$ (Figure 3(c)). The intensity of A under the 1-norm is $\mu_1(A) = \sqrt{2}/4$ (Figure 3(b)) and the intensity under the max norm is $\mu_{\max}(A) = 1/4\sqrt{2}$ (Figure 3(d)). More generally, a straightforward calculation gives that under the p-norm $\|(x, y)\|_p = (x^p + y^p)^{1/p}$, the intensity of A is $\mu_p(A) = 2^{(\frac{1}{p} - \frac{1}{2})} \cdot \frac{1}{4}$.

In specific applications, one should choose a norm compatible with the relevant notions of disturbance or uncertainty. We anticipate that the max norm will be appropriate when bounds on vector field perturbations are known componentwise rather than jointly. As in Example 5.4, one expects $\mu_p(A) \leq \mu_q(A)$ when $p > q$. Therefore, choosing a max norm may also give the most conservative estimate of attractor intensity.

We next compare intensity to a closely related concept examined by Colonius and Kliemann [7]: the lowest point of discontinuity in a map from control bound to the reachable set (control set) of an attractor. Example 5.5 demonstrates that this first discontinuity may differ from the intensity of the attractor.

Example 5.5. The one-dimensional system with vector field $f(x) = \frac{3}{4}x^4 - x^3 - 3x^2 - 1$ has an attractor $A = x_0$ and repeller $R = x_1$ corresponding to the roots x_0 and x_1 of the quartic (Figure 4(a)). The domain of attraction of A is $\mathcal{D}(A) = (-\infty, x_1)$. Consider the map $C : [0, \infty) \to \mathscr{P}(\mathbb{R})$ given by $r \mapsto P_r(A)$. Figure 4(b) depicts the graph of C. As the control bound r increases from zero, the first discontinuity in C occurs at $r = 2.25$, corresponding to the local minimum of f at $(-1, -2.25)$. However, reachable sets remain bounded within $\mathcal{D}(A)$ for small enough $r > 2.25$. It is not until the second discontinuity at $r = 9$, corresponding to the global minimum of f at $(2, -9)$, that reachable sets escape $\mathcal{D}(A)$. Hence the intensity $\mu(A) = 9$ exceeds the value of r at the first discontinuity ($\rho^* = 2.25$ in the notation of [7]).

Figure 4. *Intensity versus first discontinuity.* (a) *The vector field f of Example 5.5 yields an attractor A and repeller R. As the control bound r increases, the first discontinuity in the map C that sends r to $P_r(A)$ occurs at $\rho^* = 2.25$.* (b) *The graph of C illustrates that the first discontinuity at $\rho^* = 2.25$ occurs well below the intensity $\mu(A) = 9$.*

Example 5.5 shows that it is possible for reachable sets to expand discontinuously without escaping a domain of attraction. Whether the reverse is possible—reachable sets escaping a domain of attraction without expanding discontinuously—is an interesting question not pursued further here.

We close this section with an ecological application that illustrates attractor intensity for a stable limit cycle and the information about resilience it carries.
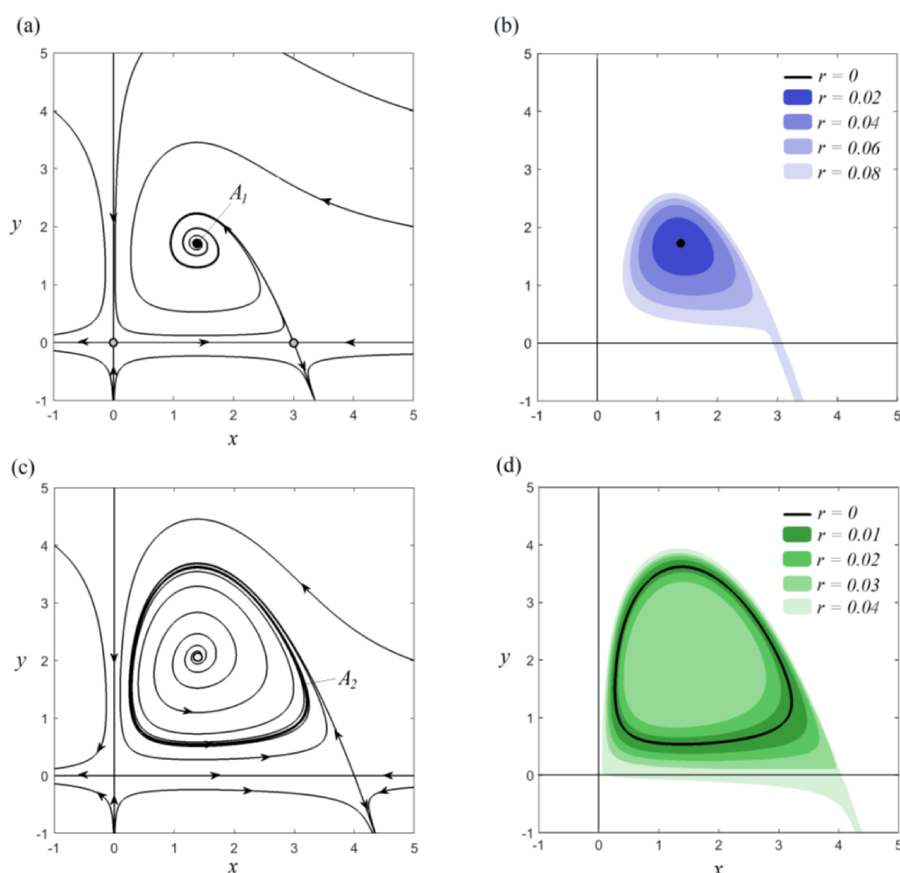
Example 5.6. Consider the following model of predator (y) and prey (x) dynamics, the details of which can be found in [18, 28]:

$$(5.8) \qquad \frac{dx}{dt} = ax\left(1 - \frac{x}{K}\right) - ky(1 - e^{-cx}),$$

$$(5.9) \qquad \frac{dy}{dt} = -by + \beta y(1 - e^{-fx}).$$

As the prey carrying capacity K grows from 3 to 4, a Hopf bifurcation in the first quadrant transforms the stable equilibrium A_1 (Figure 5(a)) into an unstable equilibrium and a stable limit cycle A_2 (Figure 5(c)). Rosenzweig used this predator-prey system to illustrate a "paradox of enrichment" in which adding nutrients to an ecosystem (thus increasing K) destabilizes the ecosystem [28]. The analysis was restricted to the dynamics near the rest point in the first quadrant, which indeed turns from stable to unstable as K increases. But in what sense is the *periodic orbit A_2* in Figure 5(c) "destabilized" relative to the spiral sink A_1 in Figure 5(a)? Both invariant sets are stable by standard mathematical definitions. In present vocabulary we might say adding nutrients lowered the resilience of the populations. Intensity of attraction provides one measure of this change in resilience.

Figure 5 plots the result of reachable set computations (described in supplement SM3 and SM4) for both the spiral sink (panel b) and stable limit cycle (panel d). Sets reachable from the spiral sink remain bounded within its domain of attraction (the first quadrant) for $r \leq 0.06$. For $r = 0.08$, the reachable set spills over the x-axis. From Figure 5(b), we estimate $0.06 < \mu(A_1) < 0.08$. Figure 5(d) illustrates sets reachable from the periodic orbit

Figure 5. *Phase portraits ((a), (c)) and reachable sets ((b), (d)) for a predator-prey equilibrium ((a), (b)) and limit cycle ((c), (d)) corresponding to (5.8) and (5.9). Reachable sets were computed using a fixed-grid set-valued Euler method at a resolution of 10^{-3} (code available in supplementary materials file M138689_02.zip [local/web 5.22KB]). Parameters, in arbitrary units, are $K = 3$ ((a), (b)), $K = 4$ ((c), (d)), and $a = 1$, $k = 0.5$, $c = 1.5$, $b = 0.5$, $\beta = 1$, $f = 0.5$ (all panels).*

A_2. These remain bounded within its domain of attraction (the first quadrant minus the unstable rest point) for $r \leq 0.02$; for $r = 0.03$, the reachable set includes the unstable rest point. Hence $0.02 < \mu(A_2) < 0.03$. If we instead consider the attractor A_3, consisting of the periodic orbit A_2 and the region it encloses, we get a slightly higher intensity $0.03 < \mu(A_3) < 0.04$. Regardless of whether we consider A_2 or A_3, the intensity estimates agree with Rosenzweig's argument that nutrient enrichment lowers stability (resilience): the attractor in the first quadrant has a lower intensity of attraction in the "enriched" scenario $K = 4$ compared to $K = 3$. As a result, we might expect that environmental variability is more likely to drive the predator and/or prey population to extinction from the cycle A_2 than from the equilibrium A_1.

6. Continuation of attractors.
We now turn from time-dependent to time-independent perturbations of a vector field. Subsection 6.1 introduces attractor blocks, an object from the

theory of isolated invariant sets that is central to the following arguments. Subsection 6.2 establishes a connection between reachable sets and attractor blocks, which we exploit in subsection 6.3 to show that intensity of attraction gives a lower bound on attractor continuation distance in the space of vector fields. Theorem 6.11, our main result, grounds the abstract question of attractor persistence in infinite-dimensional vector field space to the study of a single system's reachable sets. We conclude in subsection 6.4 with a proof of the semicontinuity of attractors based on reachable sets.

6.1. Attractor blocks. Attractors are a special type of isolated invariant set, a useful object introduced by Conley for studying robust qualitative features of flows [8, 9]. When the flow interacts nicely with the boundary of an isolating neighborhood, that neighborhood earns the name *isolating block*. Various definitions of "nicely" have been used in [8, 9, 22, 34]. Here we restrict our attention to isolating blocks associated with attractors, henceforth called attractor blocks. Roughly, the vector field should point inward to the interior of an attractor block at each point along its boundary. The following definition generalizes this notion slightly.

Definition 6.1. An attractor block B for a flow $\varphi : \mathbb{R} \times \mathbb{R}^n \to \mathbb{R}^n$ is a nonempty, compact subset of $\mathbb{R}^n$ satisfying $\varphi^t(B) \subset int(B)$ for all $t > 0$.

Any attractor block is associated with an attractor in its interior.

Lemma 6.2. For B an attractor block, $\omega(B)$ is an attractor.

Proof. $\omega(B)$ must lie interior to B because $\varphi^t(B) \subset int(B)$ for all $t > 0$. Therefore $\omega(B)$ is the omega limit set of its neighborhood B. Further, it follows directly from Definition 2.3 that $\omega(B)$ is invariant and closed; as a subset of compact B it must also be compact. By Definition 2.4, $\omega(B)$ is an attractor. ∎

Definition 6.3. If B is an attractor block and $\omega(B) = A$, we call B an attractor block associated with A.

6.2. Persistent attractor blocks from reachable sets. In this subsection we show that one can construct attractor blocks from certain reachable sets; furthermore, these blocks persist as such not just for sufficiently small perturbations to the vector field, but for perturbations smaller than r, the metric bound on control. Before stating these results we establish a definition and technical lemma.

Fixing all but the third argument in the controlled flow function $\varphi(t, x, g; f)$ of Definition 3.1 yields a map from control functions to trajectory endpoints.

Definition 6.4. Let $G_{T,x_0;f} : L^\infty([0, T], \mathbb{R}^n) \to \mathbb{R}^n$ be given by $g \mapsto \varphi(T, x_0, g; f)$. The subscripts on G may be omitted when f, T, and x_0 are clear from context or are general.

Lemma 6.5. G is an open map.

Proof. It suffices to show that G maps any basis element in the metric topology on $L^\infty([0, T], \mathbb{R}^n)$—an open ball $B_r(g) \equiv \{h \in L^\infty([0, T], \mathbb{R}^n) : \|h - g\|_\infty < r\}$—to an open set in $\mathbb{R}^n$. To show that $G(B_r(g))$ is open, we will show that for any $h \in B_r(g)$ there exists an ϵ_h sufficiently small so that the open ϵ_h-ball around $G(h)$ is also in the image $G(B_r(g))$.

Fix $h \in B_r(g)$ and consider a point $G(h) + v$ in $\mathbb{R}^n$. We derive an ϵ_h so that $\|v\| < \epsilon_h$ implies there exists $k \in B_r(g)$ such that $G(k) = G(h) + v$. First construct a path from x_0 to $G(h) + v$: let $x_h(t) = \varphi(t, x_0, h)$ and let

$$\tilde{x}(t) = x_h(t) + t\frac{v}{T}.$$

Then $\tilde{x}(0) = x_h(0) = x_0$ and $\tilde{x}(T) = x_h(T) + v = G(h) + v$, so $\tilde{x}$ is a path from x_0 to $G(h) + v$. The velocity along the path $\tilde{x}$ (defined almost everywhere) is

$$\frac{d}{dt}\tilde{x}(t) = \frac{d}{dt}[x_h(t)] + \frac{v}{T} = f(x_h(t)) + h(t) + \frac{v}{T}.$$

The control $k(t)$ required to achieve the path $\tilde{x}(t)$ as a solution to $x' = f(x) + k(t)$ is the difference between the velocity vectors along the path and the underlying vector field, so

$$k(t) = \frac{d}{dt}\tilde{x}(t) - f(\tilde{x}(t)) = f(x_h(t)) + h(t) + \frac{v}{T} - f(\tilde{x}(t)).$$

We will show that for $\|v\|$ sufficiently small $\|k - g\|_\infty < r$. Let L be the Lipschitz constant for f. For almost every $t \in [0, T]$,

$$\begin{aligned}
\|k(t) - g(t)\| &= \left\| f(x_h(t)) + h(t) + \frac{v}{T} - f(\tilde{x}(t)) - g(t) \right\| \\
&\leq \|f(x_h(t)) - f(\tilde{x}(t))\| + \|h(t) - g(t)\| + \frac{\|v\|}{T} \\
&\leq L\|x_h(t) - \tilde{x}(t)\| + \|h(t) - g(t)\| + \frac{\|v\|}{T} \\
&\leq L\|v\| + \|h - g\|_\infty + \frac{\|v\|}{T} \\
&= \|v\|\left(L + \frac{1}{T}\right) + \|h - g\|_\infty.
\end{aligned}$$

Let $\epsilon_h = \frac{r - \|h - g\|_\infty}{L + 1/T}$. Then $\|v\| < \epsilon_h$ implies that $\|k - g\|_\infty < r$. Because there exists an ϵ_h-ball about any point in $G(B_r(g))$ that is also in the image $G(B_r(g))$, the image of any basis element under G is open, and the proof is complete. ∎

With Lemma 6.5 in hand, we're ready to construct attractor blocks from reachable sets and prove their persistence properties.

Proposition 6.6. *If S is a nonempty subset of $\mathbb{R}^n$ and if $P_{r;f}(S)$ is bounded, then for any globally Lipchitz and bounded vector field $\hat{f} : U \to \mathbb{R}^n$ satisfying $\|f - \hat{f}\|_{sup} < r$, the set $\overline{P_{r;f}(S)}$ (defined using f) is an attractor block for the flow $\hat{\varphi}$ generated by $\hat{f}$.*

Before proving Proposition 6.6, we note that in particular it says that boundedness of $P_{r;f}(S)$ implies that $\overline{P_{r;f}(S)}$ is an attractor block for the flow generated by f. Intuitively, one can interpret this result as follows: whenever adding bounded control g results in a bounded reachable set, the underlying vector field f must be counteracting the control at the boundary of the reachable set, "pulling inwards" as required for an attractor block. The strength of this inwards pull reflects the control bound r used to construct the block, and feeds the persistence of the block through autonomous vector field perturbations.

Proof. (Proposition 6.6) $\overline{P_{r;f}(S)}$ is compact since it is closed and bounded; it is nonempty since it contains S. To confirm that $\overline{P_{r;f}(S)}$ is an attractor block for $\hat{\varphi}$, it will suffice to show that for any $T > 0$,

$$
(6.1) \qquad \hat{\varphi}^T(\overline{P_{r;f}(S)}) \overset{\text{(I)}}{\subset} \bigcup_{\substack{x_0 \in \overline{P_{r;f}(S)} \\ \|g\|_\infty < r}} \varphi_g^T(x_0) \overset{\text{(II)}}{\subset} \text{int}(\overline{P_{r;f}(S)}),
$$

where the controlled flow in the middle term is generated by the vector field f; that is, $\varphi_g^T(x_0) = \varphi(T, x_0, g; f)$. Fix $T > 0$. For inclusion (I), we first establish that $\hat{\varphi}^T(x_0) \in \bigcup_{\|g\|_\infty < r} \varphi_g^T(x_0)$ for any point $x_0 \in \mathbb{R}^n$. Let the path $\hat{x} : [0,T] \to \mathbb{R}^n$ be given by $\hat{x}(t) = \hat{\varphi}^t(x_0)$; in other words, $\hat{x}$ solves the initial value problem $[x' = \hat{f}(x), x(0) = x_0]$. Note that $\hat{x}$ also solves the initial value problem $[x' = f(x) + g(t), x(0) = x_0]$, where $g : [0,T] \to \mathbb{R}^n$ is given by $g(t) = \hat{f}(\hat{x}(t)) - f(\hat{x}(t))$. In the notation of controlled flow for the vector field f, $\hat{x}(T) = \varphi_g^T(x_0)$. Since by hypothesis $\|g\|_\infty \leq \|f(x) - \hat{f}(x)\|_{\sup} < r$, it follows that $\hat{\varphi}^T(x_0) = \hat{x}(T) \in \bigcup_{\|g\|_\infty < r} \varphi_g^T(x_0)$. Inclusion (I) follows directly by taking the union over $x_0 \in \overline{P_{r;f}(S)}$. We argue inclusion (II) by contradiction. Suppose there exists $x^* \in \overline{P_{r;f}(S)}$ and $g^* \in L^\infty([0,T], \mathbb{R}^n)$, $\|g^*\|_\infty < r$, such that $\varphi_{g*}^T(x^*) \notin \text{int}(\overline{P_{r;f}(S)})$. Because $\|g^*\|_\infty$ is strictly bounded below r, taking $s = \frac{r - \|g^*\|_\infty}{2}$ implies $B_s(g^*) \subset B_r(0) \subset L^\infty([0,T], \mathbb{R}^n)$. The map $G_{T,x^*;f}$ is open (Lemma 6.5) so it sends $B_s(g^*)$ to an open neighborhood $\mathcal{O}$ of $\varphi_{g*}^T(x^*)$. Because $\varphi_{g*}^T(x^*) \notin \text{int}(\overline{P_{r;f}(S)})$, it must be that $\mathcal{O} \cap (\overline{P_{r;f}(S)})^c \neq \varnothing$. Therefore, there exists a control $h \in B_s(g)$ such that $\varphi_h^T(x^*) \in (\overline{P_{r;f}(S)})^c$, an open set. Continuity of $\varphi_h^T(x)$ with respect to x implies that the inverse image of $(\overline{P_{r;f}(S)})^c$ under $\varphi_h^T(\cdot)$ is an open neighborhood V of x^*. V must intersect $P_{r;f}(S)$ nontrivially because $x^* \in \overline{P_{r;f}(S)}$. Let $b \in P_{r;f}(S) \cap V$, so that $b \in P_{r;f}(S)$ but $\varphi_h^T(b) \notin P_{r;f}(S)$. The contradiction comes from concatenating controls to move from S to b to outside $P_{r;f}(S)$. In particular, $b \in P_{r;f}(S)$ implies by Definition 4.1 that there exists $\tau \geq 0$, $a \in S$, and $j \in L^\infty([0,\tau], \mathbb{R}^n)$, $\|j\|_\infty \leq r$, such that $b = \varphi_j^\tau(a)$. Let $k \in L^\infty([0,\tau + T], \mathbb{R}^n)$ be given by
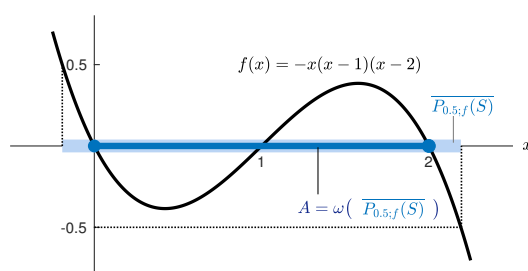
$$
k(t) = \begin{cases} j(t) & \text{if } 0 \leq t < \tau, \\ h(t - \tau) & \text{if } \tau \leq t \leq \tau + T. \end{cases}
$$

Then by construction, $\varphi_k^{\tau+T}(a) = \varphi_h^T(\varphi_j^\tau(a)) = \varphi_h^T(b) \notin P_{r;}(S)$. Yet since $\|k\|_\infty \leq r$, Definition 4.1 implies that $\varphi_k^{\tau+T}(a) \in P_{r;f}(S)$. This contradiction completes the proof. ∎

Proposition 6.6 does not guarantee any relationship between the set S used to construct the attractor block $\overline{P_{r;f}(S)}$ and the attractor associated with $\overline{P_{r;f}(S)}$. For example, if $f : \mathbb{R} \to \mathbb{R}$ is given by $f(x) = -x(x - 1)(x - 2)$ and we take a control bound $r = 0.5 > \sup_{x \in [0,2]} \|f(x)\|$, then *any* subset $S \subset [0,2]$ will yield the attractor block $\overline{P_{0.5;f}(S)}$ depicted in light blue in Figure 6, with associated attractor $A = \omega(\overline{P_{r;f}(S)}) = [0,2]$.

On the other hand, if we construct the attractor block $\overline{P_{r;f}(A)}$ from an attractor A, we recover A as the associated attractor under one further condition.

Corollary 6.7. *If A is an attractor and if $P_r(A)$ is contained in a compact subset K of $\mathcal{D}(A)$, then $\overline{P_r(A)}$ is an attractor block associated with A.*

Figure 6. *Relationship between reachable sets and attractors. For S any subset of $[0, 2]$, the reachable set under control bound $r = 0.5$ is the light blue interval, and the associated attractor is the dark blue interval $[0, 2]$.*

Proof. Proposition 6.6 gives that $\overline{P_r(A)}$ is an attractor block. It follows directly from Definition 2.3 that the omega limit set preserves inclusion, and so $\omega(A) \subset \omega(\overline{P_r(A)}) \subset \omega(K)$. By invariance of A, $\omega(A) = (A)$. Further, an open cover argument gives that for K a compact subset of $\mathcal{D}(A)$, $\omega(K) \subset A$. It follows that $\omega(\overline{P_r(A)}) = A$, so the attractor block $\overline{P_r(A)}$ is associated with the attractor A. ∎

We have already seen in Lemma 6.2 that any attractor block is associated with an attractor. We present a new proof of the converse using reachable sets.

Proposition 6.8. *For any attractor A there exists an attractor block B associated with A.*

Proof. Because $\mu(A) > 0$ (Proposition 5.2), by Definition 5.1 there exists $r > 0$ such that $P_r(A) \subset K \subset \mathcal{D}(A)$ for some compact set K. By Corollary 6.7, $\overline{P_r(A)}$ is an attractor block associated with A. ∎

6.3. Intensity and continuation distance. In this subsection we leverage the persistent attractor blocks of Proposition 6.6 and Corollary 6.7 to derive bounds on attractor continuation distance. First, we make precise the notion of continuation. Suppose that when a vector field f is perturbed to $\hat{f}$, the original attractor A shifts to another attractor $\hat{A}$, with qualitatively similar features. In what sense has A continued through the perturbation? We use the following definition, adapted from [19].

Definition 6.9. *Given an attractor A for vector field f and an attractor $\hat{A}$ for vector field $\hat{f}$, we say A continues immediately to $\hat{A}$ if there exists a set B that is an attractor block associated with A under the flow generated by f and is an attractor block associated with $\hat{A}$ under the flow generated by $\hat{f}$.*

Remark 6.10. Because A and its immediate continuation $\hat{A}$ share a common attractor block, they must also share any topological properties gleaned from the attractor block.

The next theorem gives our central result: an attractor's intensity bounds from below the distance in vector field space over which that attractor continues immediately.

Theorem 6.11. *If A is an attractor for vector field f with intensity $\mu(A)$, then for any second vector field $\hat{f}$ satisfying $\|f - \hat{f}\|_{sup} < \mu(A)$, A continues immediately to an attractor $\hat{A}$ for $\hat{f}$.*
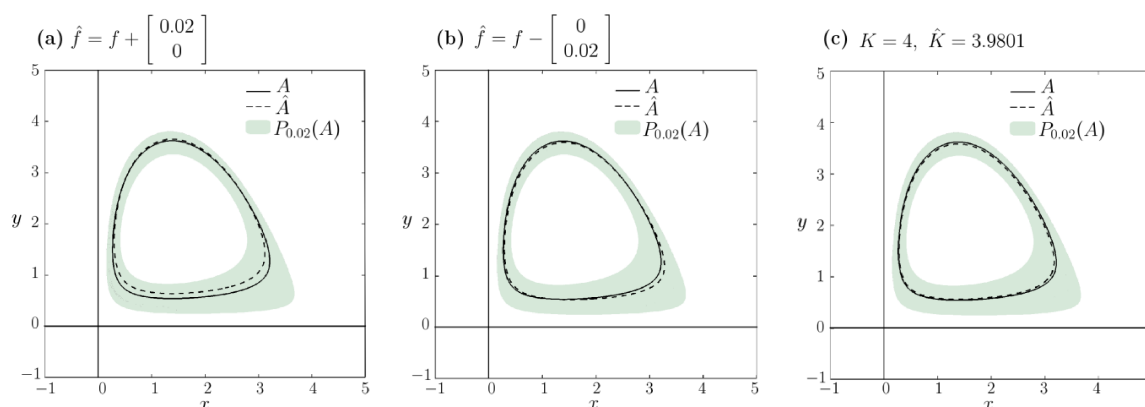
Proof. If $\|f - \hat{f}\|_{\sup} < \mu(A)$, then there exists a real number $r > 0$ such that $\|f - \hat{f}\|_{\sup} < r < \mu(A)$. It follows from Definition 5.1 and Lemma 4.5 that $P_r(A) \subset K \subset \mathcal{D}(A)$ for some compact set $K \subset \mathbb{R}^n$. This implies, by Corollary 6.7, that $\overline{P_r(A)}$ is an attractor block associated with A for vector field f. Proposition 6.6 gives that $\overline{P_r(A)}$ is also an attractor block for vector field $\hat{f}$. Defining $\hat{A}$ to be the omega limit set of $\overline{P_r(A)}$ under $\hat{f}$, we have that $\overline{P_r(A)}$ is an attractor block associated with attractors A and $\hat{A}$ under their respective flows. Hence A continues immediately to $\hat{A}$, as claimed. ∎

The following example illustrates continuation of the predator-prey limit cycle from Example 5.6.

Example 6.12. Consider the predator-prey system (5.8) and (5.9), which for suitable parameters features a stable limit cycle A in the first quadrant (see Figure 5(c)). Based on reachable set computations with the Euclidean norm, we estimated $0.02 < \mu(A) < 0.03$ in Example 5.6. By Theorem 6.11, A should continue immediately to an attractor $\hat{A}$ for any second vector field that differs from the original by up to 0.02 in Euclidean sup-norm. Figure 7 illustrates the immediate continuation $\hat{A}$ for three such vector field perturbations: (a) adding 0.02 to dx/dt, (b) subtracting 0.02 from dy/dt, and (c) changing the prey carrying capacity parameter from $K = 4$ to $K = 0.39801$ (this parameter change perturbs the vector field by less than 0.02 in a neighborhood of the attractor block $\overline{P_{0.02}(A)}$). The original cycle A is shown with a solid line, the perturbed attractor $\hat{A}$ is shown with a dashed line, and their common attractor block $\overline{P_{0.02}(A)}$ is shaded in green.

Example 6.12 highlights the strength of Theorem 6.11 in guaranteeing attractor continuation across a variety of perturbation forms, rather than perturbation of a specific parameter.

The closing example in this subsection demonstrates that attractors may continue for vector field perturbations that *exceed* intensity of attraction. This does not contradict Theorem 6.11, but indicates that intensity does not give an upper bound on attractor continuation distance.

Figure 7. *Continuation of a predator-prey limit cycle through three time-independent pertubations to the vector field f. The perturbation is given above each panel. Solid and dashed lines give the original and perturbed limit cycle, respectively. The reachable set from the original limit cycle with control bound $r = 0.02$ is shaded in green.*

Example 6.13. Let $x' = f(x) = x(1 - x)$. Then $A = 1$ is an attractor for f with intensity $\mu(A) = 0.25$. The set $B = \left[\frac{1}{2}, \frac{3}{2}\right]$ is an attractor block associated with A. Consider a second system $x' = \hat{f}(x) = cx(1 - x)$ with $c > 0$. B remains an attractor block associated with $\{1\}$ for any such system. Yet the distance $\|f - \hat{f}\|_{\sup}$ may exceed 0.25 by taking c sufficiently large. Thus the attractor $A = 1$ continues immediately to $\hat{A} = 1$ despite $\|f - \hat{f}\|_{\sup}$ exceeding $\mu(A)$.

6.4. Upper semicontinuity of attractors. Without too much extra effort, the present framework yields a new proof that attractors are upper semicontinuous—roughly, they do not expand suddenly in response to changes in a vector field. Formally, we have the following.

Definition 6.14. *An attractor A corresponding to the vector field f is* upper semicontinuous *if for any neighborhood V of A there exists a $\delta > 0$ such that for any second vector field $\hat{f}$ with $\|f - \hat{f}\|_{sup} < \delta$, A continues to an attractor $\hat{A} \subset V$.*

Theorem 6.15. *Attractors in the present setting are upper semicontinuous in the sense of Definition 6.14.*

Proof. Fix a neighborhood V of an attractor A for $x' = f(x)$. We will show that there exists a δ such that $\|f - \hat{f}\|_{\sup} < \delta$ implies that A continues immediately to an attractor $\hat{A} \subset V$ for $\hat{f}$. Using Lemma 2.1, let K be a compact set with $A \subset \text{int}(K) \subset K \subset V \cap \mathcal{D}(A)$. Part 2 of Proposition 4.6 gives an $r > 0$ such that $P_r(A) \subset \text{int}(K)$. Then $\overline{P_r(A)} \subset K \subset \mathcal{D}(A)$. Let $\delta = r$. By Proposition 6.6, $\|f - \hat{f}\|_{\sup} < r$ implies that $\overline{P_r(A)}$ is an attractor block for $\hat{f}$. Let $\hat{A}$ be the omega limit set of $\overline{P_r(A)}$ under $\hat{f}$. Then $\hat{A}$ is an attractor under $\hat{f}$ (Lemma 6.2), is an immediate continuation of A (Definition 6.9), and is contained in V, as desired. ∎

7. Conclusions. We have defined a quantity, intensity of attraction, which gives a lower bound on the magnitude of vector field perturbations through which an attractor continues to structure the long-term behavior of trajectories. Intensity is an inherently metric property of an attractor, and its value depends on the metric used on $\mathbb{R}^n$. Provided that time-dependent perturbations to a vector field are essentially bounded below the intensity of an attractor, trajectories that start at the attractor remain within its domain of attraction. And when one autonomous vector field is perturbed to a second one, an attractor for the first continues immediately to the second as long as the time-independent perturbation does not meet or exceed its intensity.

An attractor and its immediate continuation share the topological properties encoded in their common attractor block—and these properties are generally more coarse than the details treated in earlier works on structural stability [1, 23, 31]. By accepting this coarser lens on system structure we get attractor persistence not just for sufficiently small vector field perturbations, as in [9, 8], but specifically for perturbations smaller than the attractor's intensity.

For a complementary approach to translating between physical or biological systems and mathematical dynamics that bypasses an explicit deterministic model, see [3]. This recent work constructs a multivalued map from data and studies its isolated invariant sets and their robustness. Like the present work, [3] achieves lower bounds on persistence distance. Unlike the present work, the bounds in [3] are derived from the binning resolution of discretized state

space. A fixed discretization would yield uniform persistence distances across isolated invariant sets, whereas alternative attractors in our setting could have different intensities. Developing a concept of variable intensity for data-derived multivalued maps could help prioritize structures for further experimental study, based on their robustness.

We envision several additional extensions of the present work. Efficient numerical algorithms will be an important bridge between intensity theory for differential equations and the real-world modeling problems that motivated its development. To compute the intensity of an attractor we need information about both the domain of attraction and a nested collection of sets reachable from the attractor. Delineating a domain of attraction in a nonlinear system can be a nontrivial problem in its own right, and a variety of computational approaches have been developed (see [17] and references therein). To compute reachable sets quickly, one could pursue numerical strategies described in [30]. To achieve rigorous bounds on reachable sets, one could use methods based on interval arithmetic ([10]).

Applying intensity theory to a real-world system will require selecting a metric that best reflects the perturbations of interest. Here we have considered homogenous, translation-invariant metrics, carrying the same information as a norm. However, the relevant scale of perturbations might vary across state space, requiring a metric in its full generality. Furthermore, in certain settings one may want to restrict perturbations to a single coordinate while leaving others unperturbed. This could be accomplished with an extended metric. Whether intensity theory goes through for general or extended metrics remains to be determined.

Another theoretical question with practical implications concerns the relationship between intensity of attraction for discrete [19] and continuous-time dynamics. We conjecture that $\mu(A) = \lim_{t \to 0} \mu_t(A)/t$, where μ gives the intensity of an attractor under a flow φ (Definition 5.1) and μ_t gives the intensity of the attractor under the map φ^t (defined in [19, section 5]). Proving this connection would provide a theoretical justification for discretizing time in numerical computations of intensity μ.

In this work we chose an intentionally broad class of perturbation functions $g \in L^\infty$ to support versatile resilience quantification. It remains open whether restricting to a smaller class of controls (e.g., continuous) generates an equivalent theory, perhaps with simpler arguments.

Last, we aim to extend the theory of intensity to measure persistence properties of repellers and other isolated invariant sets. If we can achieve bounds on the persistence of local dynamic structures, it might be possible to describe how the global structure changes, in stages, under perturbations of increasing magnitude. We conjecture that this series of changes will reveal a Morse decomposition of the dynamical system.

REFERENCES

[1] A. ANDRONOV AND L. PONTRYAGIN, *Rough systems*, in Dokl. Akad. Nauk SSSR, 14 (1937), pp. 247–250.

[2] R. BAIER, M. GERDTS, AND I. XAUSA, *Approximation of reachable sets using optimal control algorithms*, Numer. Algebra Control Optim., 3 (2013), pp. 519–548, dx.doi.org/10.3934/naco.2013.3.519.

[3] B. BATKO, K. MISCHAIKOW, M. MROZEK, AND M. PRZYBYLSKI, *Conley index approach to sampled dynamics*, SIAM J. Appl. Dyn. Syst., 19 (2020), pp. 665–704.

[4] W.-J. BEYN AND J. RIEGER, *Numerical fixed grid methods for differential inclusions*, Computing, 81 (2007), pp. 91–106.

[5] S. CARPENTER, B. WALKER, J. M. ANDERIES, AND N. ABEL, *From metaphor to measurement: Resilience of what to what?*, Ecosystems, 4 (2001), pp. 765–781.

[6] E. A. CODDINGTON AND N. LEVINSON, *Theory of Ordinary Differential Equations*, McGraw-Hill, New York, 1955.

[7] F. COLONIUS AND W. KLIEMANN, *The Dynamics of Control*, Birkhäuser, Boston, 2000.

[8] C. CONLEY, *Isolated invariant sets and the Morse index*, CBMS Reg. Conf. Ser. Math. 38, American Mathematical Society, Providence, RI, 1978.

[9] C. CONLEY AND R. EASTON, *Isolated invariant sets and isolating blocks*, Trans. Amer. Math. Soc., 158 (1971), pp. 35–61.

[10] E. DE WEERDT, E. VAN OORT, E.-J. VAN KAMPEN, P. CHU, AND J. MULDER, *Level set method based on interval analysis*, in AIAA Guidance, Navigation, and Control Conference, 2011, AIAA, Reston, VA, 2011, AIAA 2011–6258, https://doi.org/10.2514/6.2011-6258.

[11] J. K. HALE, *Ordinary Differential Equations*, Krieger, Malabar, Florida, 1980.

[12] M. HIRSCH, S. SMALE, AND R. L. DEVANEY, *Differential Equations, Dynamical Systems, and an Introduction to Chaos*, 2nd ed., Academic Press, San Diego, 2004.

[13] A. J. HOMBURG AND T. R. YOUNG, *Bifurcations of random differential equations with bounded noise on surfaces*, Topol. Methods Nonlinear Anal., 35 (2010), pp. 77–97.

[14] J. LAMB, M. RASMUSSEN, AND C. RODRIGUES, *Topological bifurcations of minimal invariant sets for set-valued dynamical systems*, Proc. Amer. Math. Soc., 143 (2015), pp. 3927–3937.

[15] D. LI, *Morse decompositions for general dynamical systems and differential inclusions with applications to control systems*, SIAM J. Control Optim., 46 (2007), pp. 35–60.

[16] O. MALER, *Computing Reachable Sets: An Introduction*, CNRS-VERIMAG, http://www-verimag.imag.fr/~maler/Papers/reach-intro.pdf.

[17] L. G. MATALLANA, A. M. BLANCO, AND J. A. BANDONI, *Estimation of domains of attraction: A global optimization approach*, Math. Comput. Model., 52 (2010), pp. 574–585.

[18] R. M. MAY, *Limit cycles in predator-prey communities*, Science, 177 (1972), pp. 900–902.

[19] R. McGEHEE, *Some metric properties of attractors with applications to computer simulations of dynamical systems*, Unpublished report, (1988), http://www-users.math.umn.edu/~mcgehee/publications/McGehee1988p/index.html.

[20] K. MEYER, *A mathematical review of resilience in ecology*, Nat. Resour. Model., 29 (2016), pp. 339–352.

[21] K. MEYER, A. HOYER-LEITZEL, S. IAMS, I. KLASKY, V. LEE, S. LIGTENBERG, E. BUSSMANN, AND M. L. ZEEMAN, *Quantifying resilience to recurrent ecosystem disturbances using flow-kick dynamics*, Nat. Sustain., 1 (2018), pp. 671–678.

[22] K. MISCHAIKOW AND M. MROZEK, *Conley index*, in Handbook of Dynamical Systems, B. Fiedler, ed., vol. 2, Elsevier, Amsterdam, 1999, pp. 393–460.

[23] M. M. PEIXOTO, *Structural stability on two-dimensional manifolds*, Topology, 1 (1962), pp. 101–120.

[24] M. RASMUSSEN, *Morse decompositions of nonautonomous dynamical systems*, Trans. Amer. Math. Soc., 359 (2007), pp. 5091–5115.

[25] M. RASMUSSEN, J. RIEGER, AND K. N. WEBSTER, *Approximation of reachable sets using optimal control and support vector machines*, J. Comput. Appl. Math., 311 (2017), pp. 68–83, https://doi.org/10.1016/j.cam.2016.06.015.

[26] J. RIEGER, *Robust boundary tracking for reachable sets of nonlinear differential equations*, Found. Comput. Math., 15 (2015), pp. 1129–1150.

[27] J. Rieger, *The Euler scheme for state constrained ordinary differential inclusions*, Discrete Contin. Dyn. Syst. Ser. B, 21 (2016), pp. 2729–2744.

[28] M. Rosenzweig, *Paradox of enrichment: Destabilization of exploitation ecosystems in ecological time*, Science, 171 (1971), pp. 385–387.

[29] M. Sandberg, *Convergence of the forward Euler method for nonconvex differential inclusions*, SIAM J. Numer. Anal., 47 (2009), pp. 308–320.

[30] J. A. Sethian, *Level Set Methods and Fast Marching Methods: Evolving Interfaces in Computational Geometry, Fluid Mechanics, Computer Vision, and Materials Science*, Vol. 3, Cambridge University Press, Cambridge, 1999.

[31] S. Smale, *Differentiable dynamical systems*, Bull. Amer. Math. Soc., 73 (1967), pp. 747–817.

[32] C. Thieme, *Conley index theory and the attractor-repeller decomposition for differential inclusions*, arXiv preprint arXiv:2009.00696 (2020).

[33] T. Wiandt, *Intensity of attractors for closed relations on compact Hausdorff spaces*, Int. J. Difference Equ., 11 (2016), pp. 215–223.

[34] F. W. Wilson and J. Yorke, *Lyapunov functions and isolating blocks*, J. Differential Equations, 13 (1973), pp. 106–123.

[35] H. Zmarrou and A. J. Homburg, *Bifurcations of stationary measures of random diffeomorphisms*, Ergod. Theory Dynam. Systems, 27 (2007), pp. 1651–1692.