On the Efficacy of State Space Reconstruction Methods in Determining Causality*

Bree Cummins[†], Tomáš Gedeon[†], and Kelly Spendlove[‡]

Abstract. We present a theoretical framework for inferring dynamical interactions between weakly or moderately coupled variables in systems where deterministic dynamics plays a dominating role. The variables in such a system can be arranged into an interaction graph, which is a set of nodes connected by directed edges wherever one variable directly drives another. In a system of ordinary differential equations, a variable x directly drives y if it appears nontrivially on the right-hand side of the equation for the derivative of y . Ideally, given time series measurements of the variables in a system, we would like to recover the interaction graph. We introduce a comprehensive theory showing that the transitive closure of the interaction graph is the best outcome that can be obtained from state space reconstructions in a purely deterministic system. Our work depends on extensions of Takens' theorem and the results of Sauer et al. [*J. Stat. Phys.*, 65 (1991), pp. 579–616] that characterize the properties of time-delay reconstructions of invariant manifolds and attractors. Along with the theory, we discuss practical implementations of our results. One method for empirical recovery of the interaction graph is presented by Sugihara et al. [*Science*, 338 (2012), pp. 496–500], called convergent cross-mapping. We show that the continuity detection algorithm of Pecora et al. [*Phys. Rev. E*, 52 (1995), pp. 3420–3439] is a viable alternative to convergent cross-mapping that is more consistent with the underlying theory. We examine two examples of dynamical systems for which we can recover the transitive closure of the interaction graph using the continuity detection technique. The strongly connected components of the recovered graph represent distinct dynamical subsystems coupled through one-way driving relationships that may correspond to causal relationships in the underlying physical scenario.

Key words. state space reconstruction, delay embedding, Takens' theorem, chaotic attractor, partial order, lattice, filtration, causality, convergent cross-mapping

AMS subject classifications. 34C45, 34D45, 37C20

DOI. 10.1137/130946344

1. Introduction. Over the past few decades, the development of high throughput technologies in the sciences has enabled us to collect massive amounts of data inexpensively and efficiently. We have terabytes of data on different complex processes from global warming, changing ocean currents, microbiome metagenomic data, to single-cell gene expression. The analysis of these data remains one of the great challenges. In particular, a topic of much interest is the identification of causal relationships between measured quantities to deduce their underlying network structure.

*Received by the editors November 25, 2013; accepted for publication (in revised form) by T. Sauer December 5, 2014; published electronically March 4, 2015. This research was supported by DARPA grant D12AP200025.

<http://www.siam.org/journals/siads/14-1/94634.html>

[†]Mathematical Sciences, Montana State University, Bozeman, MT 59717 (cummins@math.montana.edu, gedeon@math.montana.edu). The work of the second author was partially supported by NSF grant DMS-1226213 and NIH R01 grant 1R01AG040020-01.

[‡]Department of Mathematics, Rutgers University, Hill Center - Busch Campus, Piscataway, NJ 08854 (kelly.spendlove@rutgers.edu).

It has been long recognized that correlation between two variables does not imply causation. Correlated variables can both be driven by a third hidden variable. At the same time, in nonlinear dynamical systems, two variables that drive each other's evolution in time can be correlated at some times, anticorrelated at others, and have no correlation at yet other times [19].

The framework of Granger causality (GC) [7] uses predictability to identify causation between two time series. A variable x is said to “Granger cause” y if the predictability of y declines when x is removed from the set of all possible causative variables [7]. The key assumption for applicability of GC is separability; that is, the information about x can be removed from the system by eliminating that variable from the model. Separability is characteristic of purely stochastic and linear systems, but not present in complex interconnected systems where many variables are mutually coupled. In such systems information about one variable x is contained in the variable y and cannot be formally removed without severely altering the dynamics of other variables.

Recently, Sugihara et al. [19] proposed a method to detect causal relationships based on the state space reconstruction of time series data. This approach, dubbed the convergent cross-mapping technique (CCM), is complementary to the GC approach. While GC is designed for stochastic processes, CCM applies to weakly or moderately coupled variables in systems where deterministic dynamics plays a dominating role. At the root of CCM is Takens' embedding theorem, which states that for an invariant manifold M , generic observation function $x(t)$, and delay T (explained in greater detail in section 2.1) a delay reconstruction of M is diffeomorphic to M itself. If one assumes that the coordinate times series $x(t)$ and $y(t)$ belong to the generic set, then the reconstructions M_x using $x(t)$ and M_y using $y(t)$ are diffeomorphic via their diffeomorphism to M . Therefore one can estimate elements of M_y using information from M_x , and vice versa. Sugihara et al. [19] assert that if one can predict states in M_y based on information in M_x , then x is *directionally causally related* to y in the sense that x drives y . The assessment of predictability is based on both the quality of the prediction and the detection of improvement of the quality as a function of the length of the time series.

In this contribution we will develop a precise mathematical theory that builds on Takens' theorem and puts CCM on a solid mathematical framework. This allows us to illuminate the strengths and limitations of CCM in both theory and practice.

Our starting point is a dynamical system with n variables, described by a deterministic map or a set of differential equations. We will develop parallel theories for both a system evolving on an invariant manifold, where we use the embedding framework of Takens [20], or a system on a global attractor, where we use the results of Sauer, Yorke, and Casdagli [16]. The structure of the dynamical system determines an *interaction graph* between variables; if a variable x occurs on the right-hand side of the evolution of variable y , then x is a *dynamical driver* of y . We assume that we have access to time series of all n variables, perhaps corrupted by a measurement noise, and our goal is to recover the structure of the dynamical drivers in the system, i.e., the interaction graph.

In the first part of the paper, we make basic definitions and mathematical structures that are similar for both the manifold and attractor approaches. Then we focus on extending Taken's invariant manifold approach, and finally we modify our results from the manifold viewpoint to the attractor viewpoint, outlining our work in a very similar fashion.

If two variables x and y are in the same connected component of the interaction graph, then provided that the projections of the dynamical system onto x and y belong to a generic set of observation functions, their reconstructions will be diffeomorphic, a result that follows immediately from Takens' Theorems 3.1 and 3.3. This motivates the definition of a coarser *component graph* whose vertices are connected components of the interaction graph. We show that the set of all potentially different reconstructions is parameterized by a *lattice of upper sets* of the component graph, which is introduced in section 2.3 and referenced frequently in the remainder of the paper. The parameterization occurs because an upper set of a set of variables includes all potential dynamical drivers of these variables.

We introduce the concept of fully resolved manifold structure where there is a different (i.e., nonhomeomorphic) manifold for each upper set in Definition 3.8. Only fully resolved manifold structures have sufficient richness of dynamics to allow the reconstruction of the full component graph. This is related to the observation of Sugihara et al. [19] that if a driver variable z drives two other variables x and y so strongly that they become synchronized, CCM will not be able to distinguish this from mutual dynamical driving between x and y . In our view this is a sign that the manifolds corresponding to the upper sets $\{Z, X\}$ and $\{Z, Y\}$ (explained in more detail in section 2.3) are homeomorphic, and thus the manifold structure is not fully resolved.

We also show in Theorem 3.11(b) that for any two different manifolds M_U, M_V that are parameterized by two upper sets with $U \subset V$ there is a projection $M_V \rightarrow M_U$. This corresponds to the concept of “unidirectional causality” in the CCM method.

We show that the lattice of upper sets induces filtrations of the global invariant manifold in section 2.4. By extending the embedding theorem of Takens [20], we define a generic set of observation functions with the same lattice structure (section 3.2). If the single coordinate observation functions are in this generic set, one can recover the component graph of a fully resolved system almost completely. The recovery is based on creating the reconstructions M_x for each time series x , and then detecting which reconstructions are homeomorphic to each other (this detects mutual driving), which ones are projections of each other (detects one-directional driving), and which are unrelated (Algorithm 1). We show that this approach cannot recover self-loops, and, further, one cannot distinguish between mutual versus one-directional dynamical driving within connected components of the interaction graph in Theorem 3.13. We follow similar steps for the attractor framework in section 4, building on the theorems of Sauer, Yorke, and Casdagli [16].

Finally, we suggest an alternative approach to the CCM algorithm in section 5. While CCM concentrates on the predictability of a time series $y(t)$ based on the information from the reconstructed manifold M_x , we implement an algorithm by Pecora, Carroll, and Heagy [15], which tests the hypothesis that two reconstructions are continuous images of each other. In agreement with the CCM approach, we detect continuity based on both the quality of the fit and the convergence trends with the length of the time series. While our approach is more consistent with the underlying theory, it is closely related to CCM.

We do not claim to recover causal relationships using this method, because the assignment of causality involves the interpretation of the variables in a dynamical system. For example, suppose there are two pendulums, p_1 and p_2 , that are perfectly synchronized, perhaps by a hidden and unmeasured mechanism. The second one, p_2 , has an additional chaotic pendulum

hanging from its bob. The independent pendulum p_1 obviously is not the cause of the chaotic pendulum's motion, but because it is synchronized to the proximal cause of motion, p_2 , it makes no difference whatsoever which of p_1 or p_2 are measured. In this case, the measured motion of p_1 is a proxy variable for p_2 ; it is a (perfect) correlate of p_2 . In the terminology of this paper, both p_1 and p_2 are *dynamical drivers* of the chaotic pendulum. Our methods recover the direct and indirect dynamical drivers of a variable within a given set of measurements. These relationships form a set of hypothetical causes that must be assessed in light of their physical interpretations before causality is inferred.

The paper is organized in the following way. In section 2 we formalize the interaction and component graphs, lattices, and filtrations mentioned above. In section 3 we extend Takens' theorem regarding embedding invariant manifolds to a filtration structure. We then summarize the algorithm for recovering the interaction and component graphs of a set of time series. In section 4 we use the theorems of Sauer, Yorke, and Casdagli [16] to show that the same results apply to the attractors of the system. In section 5, we describe the application of the method of Pecora, Carroll, and Heagy [15] to the recovery of the component graph of a system and present several examples. We conclude with a brief summary and discussion.

2. Decomposition of invariant manifolds and attractors.

2.1. Preliminaries. A dynamical system is a collection of maps $\phi^t : \mathbb{R}^n \rightarrow \mathbb{R}^n$, parameterized by either $t \in \mathbb{Z}$, $t \in \mathbb{Z}^+$, or $t \in \mathbb{R}$. In the first two cases the dynamical system may be generated by iterates of either a diffeomorphism or a smooth map ϕ . In the last case it may be generated by solutions of a system of differential equations.

Restricting ϕ to some compact, invariant, smooth manifold M of dimension m , the celebrated theorem of Takens shows that under very general conditions on a dynamical system (M, ϕ) , for a generic smooth observation function $\varphi : M \rightarrow \mathbb{R}$ and a generic choice of $T > 0$, the following *delay reconstruction map* is an embedding of M into $\mathbb{R}^{2m+1}$:

$$\Phi_{(\phi^T, \varphi, 2m)} : M \rightarrow \mathbb{R}^{2m+1} := (\varphi(x), \varphi(\phi^{-T}(x)), \varphi(\phi^{-2T}(x)), \dots, \varphi(\phi^{-2mT}(x))).$$

Note that the components correspond to time-lagged observations of the dynamics on M moderated by scalar observation function φ . For this paper, any C^k scalar function $\varphi : M \rightarrow \mathbb{R}$ will be referred to as an *observation function*. As will be pointed out later, for Takens' framework $k = 2$ is needed, but in the Sauer, Yorke, Casdagli [16] approach $k = 1$ is sufficient.

The following discussion will be carried out for a general dynamical system on a compact m -dimensional invariant smooth manifold M that is a submanifold of Euclidean space $\mathbb{R}^n$. While the observation function in Takens' theorem can be arbitrary, it is often the case that the observation function is a projection onto a single variable, or, in other words, a coordinate in $\mathbb{R}^n$. Takens' theorem states one may be able to reconstruct the invariant manifold from observations of a single variable of the system (see the Lotka–Volterra example below). If there exists such an observation function obtained from x , we say variable x *reconstructs* M .

Of course there are exceptions, and some variables cannot reconstruct the original manifold. One of the reasons for this may be symmetry in the manifold which causes a particular variable w to lie in a fixed subspace of a symmetry transformation. Then the variables related by symmetry, but not fixed by the symmetry, cannot be distinguished based on the observation of w .

Another reason a variable may not reconstruct M , which is the main focus of this paper, is an *interaction* structure of the set of variables. If the dynamics of the observed variable w is unaffected by another variable x , then the reconstruction from w cannot reconstruct dynamics of variable x . Sugihara et al. [19] have recently introduced a method based on Takens' theorem which attempts to infer the interaction relationship between two variables.

Often one may be more interested in embedding an attractor of a system, rather than some invariant manifold it lies on. In many cases, the attractor is not a compact submanifold of the phase space (indeed, they can be notoriously complex objects), and thus one cannot invoke Takens' theorem. In Sauer, Yorke, and Casdagli [16], results were established showing that if m is larger than twice the box-counting dimension of attractor A , then for U open with $A \subset U$, "almost every" map from $U \rightarrow \mathbb{R}^m$ is one-to-one on A and an immersion on compact subsets of smooth manifolds contained within A . Both Takens' theorem and the work of Sauer, Yorke, and Casdagli [16] will play a key role in this work.

2.2. Interaction graph. We will now introduce the notion of an *interaction graph*, which is a directed graph reflecting the dynamical interaction between variables in the dynamical system.

Definition 2.1. *Given a dynamical system ϕ , and two variables x and y of ϕ , the variable x directly drives y if the dynamics of y directly depend on x . In the case of a diffeomorphism, or a set of differential equations, this means that the right-hand side of the evolution equation for y has nontrivial dependence on x , i.e., $y_{t+1} = g(x, \cdot)$ or $\dot{y} = g(x, \cdot)$. We exclude trivial dependence where changes in x do not cause any changes in $g(x, \cdot)$.*

Definition 2.2. *Given a dynamical system ϕ , the interaction graph is a directed graph $\mathcal{IG}_\phi = (V_\phi, E_\phi)$ with the set of vertices V_ϕ corresponding to variables of ϕ and a directed edge from vertex x to y if x directly drives y . To simplify the notation, we often drop the subscript ϕ when the dynamical system under consideration is clear from the context.*

As an example, consider the Lotka–Volterra model, part of a broad class of predator-prey systems, which is a pair of first order, nonlinear, differential equations used to describe the dynamics of a biological systems in which two species interact, with one being predatory and the other its prey. The equations are given by

$$\begin{aligned} (1) \quad & \frac{du}{dt} = au - buv, \\ (2) \quad & \frac{dv}{dt} = -cv + dbuv. \end{aligned}$$

Here u is the prey, v is the predator, and a, b, c, d are parameters.

For some parameter sets, the above system exhibits a periodic orbit, which is depicted in Figure 1, along with both corresponding time series.

Notice that both variables in the system are involved in each other's production: u directly drives v , and vice versa. The resulting interaction graph is shown in Figure 2.

Further, since the system does not show any obvious symmetries, observations of either u or v should yield generic observation functions. Thus Takens' theorem allows one to reconstruct the invariant 1-manifold (periodic orbit) from delay reconstruction of either time series. Figure 3 shows the two-dimensional reconstructions from each variable.

Figure 1. Lotka–Volterra model: (a) time series; (b) periodic orbit in phase space.

Figure 2. Interaction graph for the Lotka–Volterra system.

Figure 3. Delay reconstructions from (a) u variable (prey); (b) v variable (predator).

2.3. Component graph. We want to show that the interaction structure between variables in the dynamical system naturally leads to a hierarchy of invariant manifolds or attractors that reflects the driving relationships between these variables. To develop the notion of such a hierarchy, we first coarsen the idea of interaction graphs. This requires some notions from graph theory.

Definition 2.3. For directed graph $G = (V, E)$ and $u, v \in V$, u is reachable from v if there exists a directed path of edges in E from v to u . In particular, every vertex is reachable from itself via the trivial path.

Notice that v may still be reachable from u even if u does not directly drive v . In this case the dynamics of variable u *indirectly* influence the production of v . We take the sets of mutually reachable vertices to be the basic unit of interaction using the notion of a strongly connected component.

Definition 2.4. A subset of vertices $W \subset V$ is a strongly connected component if it is a maximal subset such that for any $u, v \in W$, u is reachable from v and v is reachable from u .

Note that belonging to the same strongly connected component is an equivalence relation. Thus the set of strongly connected components partitions the graph. We denote the set of strongly connected components by $\mathcal{S}$. Although two variables in the same strongly connected component may not directly drive one another, they are still governed in part by each other's dynamics due to the strongly connected component structure. This motivates the following definition.

Definition 2.5. Given a dynamical system ϕ and its interaction graph $\mathcal{IG}_\phi$, the component graph of ϕ is obtained by taking each strongly connected component $q \in \mathcal{S}$ to be a vertex and declaring an edge from $q \rightarrow p$ if, for some (and therefore any) $v \in q$ and $u \in p$, u is reachable from v . We denote this graph by $\mathbf{CG}_\phi = (\mathcal{S}_\phi, U_\phi)$.

Notice that any component graph must be an acyclic directed graph and can therefore be thought of as a partially ordered set (poset) $(\mathcal{S}, \leq)$. We will refer to the component graph as both a graph and poset, the partial order on $\mathcal{S}$ being induced from the directed graph, i.e., $q \geq p$ if there exists an edge $q \rightarrow p$ in the transitive closure of $\mathbf{CG}_\phi$. This order $\mathcal{S}$ captures a coarse interaction structure since if $q \geq p$, then there is a path along directed edges in $\mathcal{IG}_\phi$ from at least one variable in q to at least one variable in p . If we are interested in all indirect drivers to a given variable, we are lead to the notion of an *upper set* in the partial order $\mathbf{CG}$: an upper set U is a subset $U \subset \mathcal{S}$ with the property that if $p \in U$ and $q \geq p$, then $q \in U$.

As we will show, the significance of an upper set U of $\mathbf{CG}$ is that for a dynamical system ϕ with an interaction graph $\mathcal{IG}$ there is an associated well-defined dynamical system ϕ_U obtained by restricting dynamics to the variables in the upper set. To describe the relationship between the dynamical systems ϕ_U corresponding to upper sets $U \subset \mathbf{CG}$ we introduce the notion of lattice of upper sets. Note that the upper sets of $\mathbf{CG}$ form a distributive lattice in which the partial ordering is given by set inclusion: the join corresponds to union and the meet to intersection; the highest element of this lattice represents all vertices of the component graph $\mathbf{CG}$ and hence all variables of the system. We denote this lattice UL_ϕ . An example can be seen in Figure 4.

The component graph and associated lattice of upper sets encode the driving relationships between the variables of the system. These structures allow us to decompose the invariant manifold or attractor of the system in a natural fashion. We begin with a definition of a structure for decomposing the manifold, for which UL_ϕ will serve as an indexing set. We adapt slightly the notion of a filtration, which usually requires that the indexing set is well ordered; in our case UL_ϕ is only partially ordered.

2.4. Filtration of invariant manifolds and attractors.

Definition 2.6. A filtration $\mathbb{F}$ is a collection of subsets X_a of a set X indexed by a poset I , together with a collection of inclusion maps $\iota_{ab} : X_a \hookrightarrow X_b$, for every pair of indices satisfying $a \leq b$ in the partial order. This collection is consistent; that is, for all ordered triples $a \leq b \leq c$

Figure 4. Example $\mathcal{IG}_\phi$, $\mathcal{CG}_\phi$, and associated lattice of upper sets UL_ϕ . We will use this example throughout this paper.

we have

$$(3) \quad \iota_{bc} \circ \iota_{ab} = \iota_{ac}.$$

A descending filtration is a collection of subsets indexed by a poset, together with a collection of projection maps $\pi_{ba} : X_b \rightarrow X_a$, for every pair of indices satisfying $a \leq b$ in the partial order. Furthermore, this collection is consistent; that is, for all ordered triples $a \leq b \leq c$ we have

$$(4) \quad \pi_{ba} \circ \pi_{cb} = \pi_{ca}.$$

Given the lattice UL_ϕ as an index set we define a descending filtration of the phase space $\mathbb{R}^n$. To each element $U \in UL_\phi$ we define $|U|$ to be the set of variables corresponding to vertices in $\mathcal{IG}$ that belong to U . Let $\mathbb{R}_U \subset \mathbb{R}^n$ be a subspace spanned by the variables in $|U|$. If n_U is the number of variables in U , then $\mathbb{R}_U$ has dimension n_U (in other words, $\mathbb{R}_U$ is isomorphic to $\mathbb{R}^{n_U}$). Notice that the partial order on UL_ϕ implies that

$$\mathbb{R}_U \subset \mathbb{R}_V \quad \text{if and only if} \quad U < V \quad \text{if and only if} \quad |U| \subset |V|.$$

In particular, there is a natural projection

$$\pi_{VU} : \mathbb{R}_V \rightarrow \mathbb{R}_U.$$

To describe the projection we can reorder variables in $\mathbb{R}_U$ and $\mathbb{R}_V$ in such a way that $\{x_1, \dots, x_{n_U}\}$ are variables in $\mathbb{R}_U$ and $\{x_1, \dots, x_{n_U}, x_{n_U+1}, \dots, x_{n_V}\}$ are variables in $\mathbb{R}_V$. Then

$$\pi_{VU}(x_1, \dots, x_{n_U}, x_{n_U+1}, \dots, x_{n_V}) = (x_1, \dots, x_{n_U}).$$

It is easy to see that the collections of maps $\{\pi_{VU}\}_{U,V \in UL_\phi}$ satisfy the consistency condition (4). Notice that the supremum of UL_ϕ corresponds to collection of all variables. Therefore $\mathbb{R}_{sup}$ is the entire phase space $\mathbb{R}^n$, and we use the shorthand notation $\pi_U : \mathbb{R}^n \rightarrow \mathbb{R}_U$ to denote

the projection from $\mathbb{R}^n$ to $\mathbb{R}_U$. The subspaces $\mathbb{R}_U$ indexed by elements of the lattice UL_ϕ , with natural projections discussed above, form a descending filtration that we will call the *upper set filtration*.

Proposition 2.7. *Given a dynamical system ϕ , for every pair $U, V \in UL_\phi$ with $U \leq V$ the projections π_U and π_{VU} commute with the dynamical system ϕ . Therefore there is a well defined collection of dynamical systems $\{\phi_W^t\}_{W \in UL_\phi}$, with ϕ_U^t defined on $\mathbb{R}_U$ given by*

$$(5) \quad \phi_U^t \circ \pi_U(x) = \pi_U \circ \phi^t(x) \text{ for all } x \in \mathbb{R}^n.$$

Furthermore, the collection of dynamical systems $\{\phi_W^t\}_{W \in UL_\phi}$ is consistent; i.e., it satisfies

$$\pi_{VU} \circ \phi_V^t(x) = \phi_U^t \circ \pi_{VU}(x) \text{ for all } x \in \mathbb{R}_V.$$

Proof. We need to show that ϕ_U^t is well defined. Assume that $\pi_U(x) = \pi_U(y)$ for some $x \neq y$. This means that x and y agree in all variables that correspond to vertices in $|U| \subset |V_\phi|$, where $|V_\phi|$ is the set of all vertices in the interaction graph $\mathcal{IG}_\phi$, and can differ only in the variables that correspond to vertices $|V_\phi| \setminus |U|$. The key observation is that the equations that govern variables in $|U|$ do not depend on variables in $|V_\phi| \setminus |U|$. Therefore even though $\phi^t(x)$ differs from $\phi^t(y)$, they agree on all variables in $|U|$. Thus

$$\pi_U \circ \phi^t(x) = \pi_U \circ \phi^t(y)$$

and ϕ_U^t is well defined.

To show consistency, we use the consistency of the projection maps $\{\pi_{VU}\}_{U, V \in UL_\phi}$. In particular, if $U \leq V$, we have $\pi_U = \pi_{VU} \circ \pi_V$. Using this and (5), we get

$$\pi_U \circ \phi^t(x) = \phi_U^t \circ \pi_U(x) = \phi_U^t \circ \pi_{VU} \circ \pi_V(x).$$

On the other hand

$$\pi_U \circ \phi^t(x) = \pi_{VU} \circ \pi_V \circ \phi^t(x) = \pi_{VU} \circ \phi_V^t \circ \pi_V(x).$$

Comparing the two expressions, we get the consistency statement above. $\blacksquare$

Now assume that there is a global attractor A in $\mathbb{R}^n$ under ϕ , or, alternatively, that there is an invariant manifold M for ϕ in $\mathbb{R}^n$. The next results show that these objects inherit the structure of the upper set filtration and are direct consequences of Proposition 2.7.

Corollary 2.8. *Consider ϕ restricted to some compact invariant manifold M . There is a descending filtration $\mathbb{M}$ with index set UL_ϕ , consisting of the collection $\{M_U\}_{U \in UL_\phi}$, where $M_U := \pi_U(M)$. Further, each M_U is a compact invariant set under the dynamical system ϕ_U^t , defined in Proposition 2.7.*

Proof. The continuous image of a compact set is compact, so each $M_U = \pi_U(M)$ is compact. Now let $y \in M_U$ and take any $z \in M$ such that $\pi_U(z) = y$. Then $\phi_U^t(y) = \pi_U \circ \phi^t(z) \in \pi_U(M) = M_U$. Thus M_U is a projection of M that is a compact, invariant set under ϕ_U^t . $\blacksquare$

Corollary 2.9. *Consider ϕ restricted to some compact attractor A . There is a descending filtration $\mathbb{A}$ of the attractor A with index set UL_ϕ , consisting of the collection $\{A_U\}_{U \in UL_\phi}$,*

where $A_U := \pi_U(A)$. Further, each A_U is a compact attractor of the associated dynamical system ϕ_U^t , defined in Proposition 2.7.

Proof. The argument for A_U a compact, invariant set is identical to that for M_U a compact, invariant set in Corollary 2.8. We also need to show that A_U is an attractor of ϕ_U^t . We use the definition of attractor given in Strogatz [18]: An *attractor* is a closed, invariant set that attracts an open set of initial conditions; further, it contains no proper subset with the same properties (i.e., it is minimal).

Assuming the subspace topology on $A \subset \mathbb{R}^n$, the projection π_U is an open map. Therefore, an open set $O \supset A$ of initial conditions that are attracted to A maps onto an open set $P := \pi_U(O) \supset \pi_U(A) = A_U$ under the projection. Let $y \in P$, and take any $z \in O$ such that $\pi_U(z) = y$. Then by the definition, $\phi^t(z) \rightarrow A$ as $t \rightarrow \infty$. In particular, given $\epsilon > 0$ there exists $T(\epsilon)$ such that $t > T(\epsilon)$ implies $\text{dist}(\phi^t(z), A) < \epsilon$, where $\text{dist}(\phi^t(z), A) = \inf\{\|\phi^t(z) - a\|, a \in A\}$ and $\|\cdot\|$ is the norm on $\mathbb{R}^n$. Then $\text{dist}(\phi_U^t(y), A_U) = \text{dist}(\pi_U(\phi^t(z)), \pi_U(A)) \leq \text{dist}(\phi^t(z), A) < \epsilon$ when $t > T(\epsilon)$. So $\phi_U^t(y) \rightarrow A_U$ as $t \rightarrow \infty$ and the open set $P \supset A_U$ is attracting for ϕ_U^t .

Suppose by contradiction that A_U is not minimal; i.e., there exists closed $B_U \subsetneq A_U$ that is invariant under ϕ_U^t and there exists open $P \supset B_U$ of initial conditions that are attracted to B_U . By continuity, the preimages $\pi_U^{-1}(B_U)$ and $\pi_U^{-1}(P)$ are closed and open, respectively, in $\mathbb{R}^n$, so that $B := \pi_U^{-1}(B_U) \cap A$ and $Q := \pi_U^{-1}(P) \cap O$ are closed and open, respectively. Let $y \in B_U$ and $z \in B$ such that $\pi_U(z) = y$. Then $\phi^t(z) \in B$ for all time t , since otherwise the projection $\pi_U(\phi^t(z)) = \phi_U^t(y)$ would not be in B_U , a contradiction. So B is a proper subset of A , because $B_U \subsetneq A_U$, which is also closed and invariant. Now let $y \in P \cap \pi_U(O)$, so that $\phi_U^t(y) \rightarrow B_U$ as $t \rightarrow \infty$. There must be $z \in Q$ such that $\pi_U(z) = y$. Then $\phi_U^t \circ \pi_U(z) \rightarrow B_U$ implies $\pi_U \circ \phi^t(z) \rightarrow B_U$ by (5), so that $\phi^t(z) \rightarrow \pi_U^{-1}(B_U)$. Since $z \in Q \subset O$, $\phi^t(z) \rightarrow A$ as $t \rightarrow \infty$, so that $\phi^t(z) \rightarrow \pi_U^{-1}(B_U) \cap A \equiv B$. This shows that Q is an open set of initial conditions that is attracted to B . This contradicts A minimal; therefore A_U must be minimal. ■

We call $\mathbb{A}$ an *attractor filtration* and $\mathbb{M}$ an *invariant set filtration* of the dynamical system ϕ . The projection $\pi_{VU} : \mathbb{R}_V \rightarrow \mathbb{R}_U$ acting on real spaces is noninjective by construction for $U < V$, but the same projection acting on invariant sets or attractors, $\pi_{VU} : M_V \rightarrow M_U$ and $\pi_{VU} : A_V \rightarrow A_U$, could be injective. This can occur, for example, if two manifolds on different levels of the filtration are homeomorphic even though the associated dynamical systems ϕ_U^t and ϕ_V^t are in different dimensions. The presence or absence of injectivity in the projections controls how much of the dynamical structure of ϕ can be recovered, and later we will discuss constraining the filtrations $\mathbb{M}$ and $\mathbb{A}$ to noninjective projections for optimal information recovery.

2.5. Join irreducibility. In the example in Figure 5, consider the attractors $A_{\{TRS\}}$, $A_{\{TR\}}$, and $A_{\{TS\}}$ corresponding to the upper sets $\{T, R, S\}$, $\{T, R\}$, and $\{T, S\}$, respectively. Since the upper set $\{T, R, S\}$ contains no variables other than those found in the sets $\{T, R\}$ and $\{T, S\}$, the attractor $A_{\{TRS\}}$ is completely determined by the attractors $A_{\{TR\}}$ and $A_{\{TS\}}$. Indeed, let

$$\pi_1 : A_{\{TRS\}} \rightarrow A_{\{TR\}} \text{ and } \pi_2 : A_{\{TRS\}} \rightarrow A_{\{TS\}}$$

Figure 5. Descending filtrations $\mathbb{A}$ and $\mathbb{M}$ associated with UL_ϕ of Figure 4.

be the projections that correspond to set inclusions $\{T, R\} \hookrightarrow \{T, R, S\}$ and $\{T, S\} \hookrightarrow \{T, R, S\}$, respectively. We want to express the observation that any trajectory $\varphi(x, t)$ in $A_{\{T, R, S\}}$ can be fully described by trajectories $\pi_1(\varphi(x, t))$ in $A_{\{T, R\}}$ and $\pi_2(\varphi(x, t))$ in $A_{\{T, S\}}$. Note that the trajectory $\varphi(x, t)$ is described by the set of coordinate scalar-valued trajectories $x_j(t)$ for all $x_j \in |T| \cup |R| \cup |S|$. If we denote the list of these trajectories by

$$\{\varphi(x, t)\} := \{x_{j_1}(t), \dots, x_{j_s}(t) \mid x_{j_i} \in |T| \cup |R| \cup |S|\},$$

then we can write

$$\{\varphi(x, t)\} = \{\pi_1(\varphi(x, t))\} \cup \{\pi_2(\varphi(x, t))\}.$$

Note that since the variables in $|T|$ are included in both projections, the sets on the right-hand side are not disjoint. This motivates the following definition.

Definition 2.10. A dynamical system $(X, \varphi(x, t))$ is a join of dynamical systems $(X_1, \varphi_1), \dots, (X_k, \varphi_k)$ if there are projections $\pi_i : X \rightarrow X_i$, $i = 1, \dots, k$, such that the flow $\varphi(x, t)$ on X admits a decomposition

$$\{\varphi(x, t)\} = \bigcup_{i=1}^k \{\pi_i(\varphi(x, t))\} =: \bigcup_{i=1}^k \{\varphi_i(x, t)\}.$$

If the ranges of the projections π_i , $i = 1, \dots, k$, are mutually disjoint, then $(X, \varphi(x, t))$ is a direct product of attractors $(X_1, \varphi_1), \dots, (X_k, \varphi_k)$ and one can write $\varphi(x, t) = (\varphi_1(x, t), \dots, \varphi_k(x, t))$.

An invariant manifold M (attractor A) is a join of manifolds $M_1, \dots, M_k$ (attractors $A_1, \dots, A_k$) if the dynamical system on M (A) is a join of the corresponding dynamical systems on $M_1, \dots, M_k$ (attractors $A_1, \dots, A_k$). This definition, combined with the notion of join-irreducible elements of a lattice, will give us a minimal representation of the dynamics.

Definition 2.11. An element of a lattice is join-irreducible if the join (least upper bound) of a and b , denoted $x = a \vee b$, implies $x = a$ or $x = b$. In other words these are elements that cannot be written as the join of two other elements of the lattice. An element is join-reducible if it is not join-irreducible.

Definition 2.12. A set of finite subsets $\{N_1, \dots, N_j\}$ of a set G will be called *irredundant* whenever all of its elements N_i are mutually incomparable (with respect to the subset ordering).

Lemma 2.13 (see [17, Theorem 1.6.2]). Let L be a distributive lattice. Then every $x \in L$ can be written uniquely as an irredundant join of join-irreducible elements.

Corollary 2.14. The filtrations $\mathbb{A}$ and $\mathbb{M}$ are fully characterized by the attractors (resp., manifolds) that correspond to join-irreducible elements in the lattice UL_ϕ . Every attractor (resp., manifold) that corresponds to join reducible element $U \in UL_\phi$ is a join of attractors (resp., manifolds) corresponding to join-irreducible elements $V_1, \dots, V_k \in UL_\phi$, for which

$$U = V_1 \vee V_2 \vee \dots \vee V_k$$

is irredundant representation of U in terms of $V_1, \dots, V_k$.

Proof. By Lemma 2.13 each join-reducible element U of the lattice UL_ϕ can be written as an irredundant join of join-irreducible elements. Since irredundant elements are not comparable in the lattice, in each V_j there is a variable unique to this set that does not belong to any other set. On the other hand, since U is a join of $V_1, \dots, V_k$, each variable in $|U|$ is in one of the sets $|V_1|, |V_2|, \dots, |V_k|$. Therefore the dynamical system generated by the variables in $|U|$ is a join of dynamical systems generated on the variables in $|V_1|, |V_2|, \dots, |V_k|$. The result follows. ■

Until this point, our development was identical for both Takens' [20] (manifolds) and Sauer, Yorke, and Casdagli's [16] (attractors) formulations of the delay reconstruction process. Since the notion of a "large" set is different in these two approaches (generic in Takens versus prevalent in Sauer, Yorke, and Casdagli [16]), and the assumptions guaranteeing that the set of successful observation functions is large are different as well, we will formulate our theory separately for these two cases.

3. Extension of Takens' theorem. We start with a discussion of the technical restrictions we put on the dynamical system studied in this section. These assumptions come from Takens [20]. For a diffeomorphism $\phi : M \rightarrow M$ on a smooth, compact manifold M of dimension m , we require the following properties, which are generic in $D^2(M, M)$, the space of twice differentiable diffeomorphisms on M :

- A1** The periodic points of ϕ with periods less than or equal to $2m + 1$ are finite in number.
- A2** If x is any periodic point with period $k \leq 2m + 1$, then the eigenvalues of the Jacobian $(D\phi^k)_x$ are all distinct and not equal to 1.

For a smooth flow ϕ^t on a smooth, compact manifold M of dimension m , and a time lag $T > 0$, we impose the following properties, which are generic in $C^2(M, M) \times \mathbb{R}^+$ (see Theorem A.1), where $C^2(M, M)$ is the space of twice differentiable flows on M :

- B1** The flow ϕ^t has no periodic orbits of period kT for $k \leq 2m + 1$.
- B2** If $\phi^t(x) = 0$, the eigenvalues of the Jacobian $(D\phi^T)_x$ are distinct and not equal to 1, where ϕ^T is the time T flow-induced diffeomorphism.

We offer a proof in the appendix (Theorem A.1) that pairs (ϕ^t, T) satisfying B1 and B2 form an open and dense set in $C^2(M, M) \times \mathbb{R}^+$, where the choice of T depends on the choice of ϕ^t . In fact, we prove genericity under a less restrictive version of B1.

It is not clear a priori that conditions B1 and B2 are generic, in particular dense, within a class of flows with a fixed interaction graph $\mathcal{IG}$, or a fixed component graph $\mathcal{CG}$. Restriction

to such a class provides a smaller class of perturbations than the class of all smooth flows on M . To address this concern, we show in the appendix that conditions B1 and B2 are generic in the set of flows with a fixed component graph (Theorem A.6). This result does not (as far as we know) extend to the class of flows with a fixed interaction graph $\mathcal{IG}$. Since we will later show that only strongly connected components of the interaction graph are recoverable using reconstruction methods, the lack of genericity of assumptions B1 and B2 for a fixed interaction graph does not restrict the generality of our result.

In order to apply Takens' theorem to dynamical subsystems ϕ_U^t , we further assume the following:

C1 The elements M_U of the invariant set filtration $\mathbb{M}$ are compact, invariant *smooth manifolds*.

This is not in general true, but it allows the application of Takens' theorem to the whole filtration, which we will henceforth call an *invariant manifold filtration* rather than an invariant set filtration.

3.1. Takens' theorems. We restate the results of Takens here, in slightly different language that is convenient for our purposes, and discuss the extension of his work to the current context.

Theorem 3.1 (Takens' theorem as restated by Huke [9, 10]). *Let M be a smooth, compact manifold of dimension m . Let $\phi : M \rightarrow M$ be a diffeomorphism, $\phi \in D^2(M, M)$, that fulfills the generic properties A1 and A2. Then for generic $\varphi \in C^2(M, \mathbb{R})$, the map $\Phi_{(\phi, \varphi, 2m)} : M \rightarrow \mathbb{R}^{2m+1}$ defined by*

$$\Phi_{(\phi, \varphi, 2m)}(x) = (\varphi(x), \varphi(\phi(x)), \varphi(\phi^2(x)), \dots, \varphi(\phi^{2m}(x)))$$

is an embedding.

The main difference between this statement and the original Takens' theorem is that the conditions for a generic diffeomorphism are stated explicitly in the theorem body rather than in the proof. This allows one to choose a generic diffeomorphism first and then associate to it a generic set of observation functions $\varphi \in C^2(M, \mathbb{R})$.

When considering the analogous result for flows, we must make use of flow-defined diffeomorphisms.

Definition 3.2. *For a flow ϕ^t on M , we define a diffeomorphism ϕ^T by fixing a time lag T . The composition of ϕ^T with itself, $(\phi^T)^k$, is identical to the value of the flow ϕ^t at $t = kT$.*

Takens' Theorems 2 and 4 and Corollary 5 in [20] discuss the analogous result to Theorem 3.1 for flows. We have rewritten these results in the next theorems, in which we use the wording of Huke [9, 10] as in Theorem 3.1. We have done this to make explicit the requirements on the time lag T . In Takens' original formulation, he considers the case of $T = 1$ first and then expands the result to generic $T > 0$. Our formulation chooses a time lag T together with the flow ϕ^t .

Theorem 3.3 (Takens' Theorem 2 [20]). *Let ϕ^t be a C^2 flow on a smooth, compact manifold M of dimension m . Let $(\phi^t, T) \in C^2(M, M) \times \mathbb{R}^+$ fulfill the generic conditions B1 and B2. Then for a generic observation function $\varphi \in C^2(M, \mathbb{R})$, the delay map*

$$\Phi_{(\phi^T, \varphi, 2m)}(x) = (\varphi(x), \varphi(\phi^T(x)), \varphi((\phi^T)^2(x)), \dots, \varphi((\phi^T)^{2m}(x)))$$

is an embedding $\Phi_{(\phi^T, \varphi, 2m)} : M \rightarrow \mathbb{R}^{2m+1}$.

The genericity of conditions B1 and B2 follows from the Kupka–Smale theorem [14] and the compactness of M . The proof is sketched in the appendix (see Theorem A.1), along with a proof of Theorem 3.3 assuming Theorem 3.1, which we do not prove, but refer the reader to [9, 10] and [20].

The delay map $\Phi_{(\phi^T, \varphi, 2m)}(x)$ preserves the dynamics of the flow ϕ^t , as stated in Corollary 5 of Takens [20] for a time lag in a residual set of positive numbers, restated in different language below.

Theorem 3.4 (Takens' Theorem 4 and Corollary 5 [20]). *Let M and m be as above. Let $p \in M$ be an arbitrary point on the manifold. Then there exists a residual set $G(p) \subset C^2(M, M) \times \mathbb{R}^+$ such that if $(\phi^t, T) \in G(p)$, and if $\varphi \in C^2(M, \mathbb{R})$ is generic, then (1) the pair (ϕ^t, T) fulfills B1 and B2, and (2) the set of positive limit points of $\phi^t(p)$ is bijectively mapped to the limit points of the set $\{\Phi_{((\phi^T)^k, \varphi, 2m)}(p)\}_{k=0}^\infty$.*

Remark. In Theorems 3.1 and 3.3, the concept of genericity means that a property holds on an open and dense set. In Theorem 3.4, genericity for the choice of a pair (ϕ^t, T) is reduced to a residual set—a countable intersection of open and dense sets.

In the appendix, we prove that conditions B1 and B2 hold generically for (ϕ^t, T) when the class of ϕ^t is restricted to those with a fixed component graph (Theorem A.6). This means that within the class of smooth flows satisfying the constraints $\pi_{VU} \circ \phi_V^t = \phi_U^t \circ \pi_{VU}$ for any pair of upper sets in UL such that $U < V$, the pairs (ϕ^t, T) that fulfill conditions B1 and B2 form an open and dense set. Therefore we may apply Takens' theorem to each element of the manifold filtration.

Definition 3.5. *For a fixed manifold M and component graph CG , we define a subset $\mathcal{P}_{CG} \subset C^2(M, M)$ as those flows where $\pi_{VU} \circ \phi_V^t = \phi_U^t \circ \pi_{VU}$ holds for any pair of upper sets $U < V$ in the lattice UL induced by CG . Analogously, let $\mathcal{P}_{CG} \subset D^2(M, M)$ be the set of diffeomorphisms such that $\pi_{VU} \circ \phi_V^t = \phi_U^t \circ \pi_{VU}$ holds for any pair of upper sets $U < V$.*

Theorem 3.6. *For a smooth, compact manifold M , a point $p \in M$, and a component graph CG , there exists a residual subset*

$$G_{\mathcal{P}} \subset \mathcal{P}_{CG} \times \mathbb{R}^+ \subset C^2(M, M) \times \mathbb{R}^+$$

such that each pair $(\phi^t, T) \in G_{\mathcal{P}}$ satisfies the conclusions of Theorem 3.4.

Analogously for diffeomorphisms, given a smooth, compact manifold M and a component graph CG , there is an open and dense set

$$G_{\mathcal{P}} \subset \mathcal{P}_{CG} \subset D^2(M, M)$$

such that $\phi \in G_{\mathcal{P}}$ satisfies A1 and A2.

Proof. The proof of the statement for flows is in the appendix. ■

Remark. We have used the same notation $\mathcal{P}_{CG}$ and $G_{\mathcal{P}}$ for both diffeomorphisms and flows, although they belong to different spaces in the two different cases. The appropriate space should be clear from context.

It remains to show that the observation functions φ can be chosen generically in a way that respects the dynamical structure.

3.2. Filtration $\mathbb{Y}$ of observation functions. In this section we consider the filtration of spaces of observation functions with domains in $M_U \in \mathbb{M}$. Let Y_U denote the set of observation functions of M_U , $\{\varphi : \varphi \in C^2(M_U, \mathbb{R})\}$, and let $Y := \{\varphi : \varphi \in C^2(M, \mathbb{R})\}$ be the set of observation functions for the whole manifold. For $U \leq V$ in lattice UL_ϕ , define inclusion $\iota_{UV} : Y_U \hookrightarrow Y_V$ as follows. For $\varphi \in Y_U$

$$\iota_{UV}(\varphi) := \varphi \circ \pi_{VU} : M_V \rightarrow \mathbb{R}.$$

Thus $\iota_{UV}(Y_U) = \{\varphi \circ \pi_{VU} : M_V \rightarrow \mathbb{R} \mid \varphi \in Y_U\} \subset Y_V$ and $\iota_U(Y_U) = \{\varphi \circ \pi_U : M \rightarrow \mathbb{R} \mid \varphi \in Y_U\} \subset Y$. The collection of sets $Y_U = \iota_U(Y)$ forms a filtration $\mathbb{Y}$ indexed by the set UL_ϕ . This is depicted in Figure 6.

Figure 6. Example CG_ϕ , UL_ϕ , and filtration of observation functions $\mathbb{Y}$.

A subset of the observation functions in Y_U will be generic as an immediate consequence of Takens' theorem.

Corollary 3.7. Suppose that there is a generic diffeomorphism $\phi_U \in \mathcal{G}_P$ or a generic flow $(\phi_U^t, T) \in \mathcal{G}_P$ on a smooth, compact manifold M_U that is a member of the filtration $\mathbb{M}$. Then by Takens' Theorem 3.1 or 3.3, there exists a generic set $\mathcal{G}_U \subset Y_U = C^2(M_U, \mathbb{R})$ of observation functions φ such that the delay map induced by φ is an embedding of M_U .

Remark. The generic set $\mathcal{G}_U$ depends either on the given diffeomorphism ϕ (under Theorem 3.1) or on the given flow ϕ^t and the choice of time lag T (under Theorem 3.3). We omit these dependencies in the notation and note that we will assume the same time lag T for every ϕ_U^t in the dynamical system.

We now discuss the relationship between the sets of observation functions $\mathcal{G}_U$ and $\mathcal{G}_V$ for $U < V$. Consider an observation function $\varphi : M_U \rightarrow \mathbb{R}$ that belongs to $\mathcal{G}_U$. Then, as noted above, $\varphi \circ \pi_{VU} : M_V \rightarrow \mathbb{R}$ is an observation in Y_V . However, in general $\varphi \circ \pi_{VU} \notin \mathcal{G}_V$. For instance, suppose φ is a projection of M_U onto some variable in $|U|$. Then one does not expect that observations of this variable will reconstruct the full dynamics of M_V since no variable in $|V| \setminus |U|$ affects variables in U . In fact, although

$$(6) \quad \iota_{UV}(\mathcal{G}_U) = \{\varphi \circ \pi_{VU} : M_V \rightarrow \mathbb{R} \mid \varphi \in \mathcal{G}_U\} \subset Y_V,$$

we expect that $\iota_{UV}(\mathcal{G}_U)$ belongs to the exceptional set $Y_V \setminus \mathcal{G}_V$. The reason is that when $\varphi_V \in \mathcal{G}_V$, the map $\Phi_{(\phi^T, \varphi_V, 2n_V)}$ reconstructs (M_V, ϕ_V) , while when $\varphi_U \in \mathcal{G}_U$, the map $\Phi_{(\phi^T, \varphi_U, 2n_U)}$ reconstructs (M_U, ϕ_U) , which is a projection of (M_V, ϕ_V) . If the projection $\pi_{VU}(M_V) = M_U$ is proper (i.e., not injective), we indeed have

$$(7) \quad \iota_{UV}(\mathcal{G}_U) \subset Y_V \setminus \mathcal{G}_V.$$

The same argument shows that for any set U that is a proper subset of $U_{\max} = \bigcup_{U \in UL} U$, we expect that

$$\iota_U(\mathcal{G}_U) \subset Y \setminus \mathcal{G}_{\max},$$

where $\mathcal{G}_{\max}$ is open and dense in Y , and $\varphi \in \mathcal{G}_{\max}$ provides an embedding of the manifold M . We formalize the requirement for (7) to hold in the following definition.

Definition 3.8. *The dynamical system (M, ϕ) has a fully resolved manifold filtration $\mathbb{M}$ if the following two conditions are met:*

1. *Given comparable U, V in the lattice UL_ϕ with $U < V$, the corresponding projection $\pi_{VU} : M_V \rightarrow M_U$ is noninjective.*
2. *Given noncomparable U, V in the lattice UL_ϕ , there exist no continuous surjections $f : M_U \rightarrow M_V$ or $g : M_V \rightarrow M_U$.*

The first condition ensures that comparable, nonequal upper sets are associated with distinguishable invariant manifolds. In particular, the first condition is enough to ensure (7). As we will see later, the second condition is required to recover the full set of strongly connected components in $\mathbf{CG}$.

We consider the union of the generic sets across the manifold filtration

$$(8) \quad \mathcal{G} := \bigcup_{U \in UL} \iota_U(\mathcal{G}_U).$$

The set $\mathcal{G}$ is generic in Y , because it contains the open and dense set $\mathcal{G}_{\max}$ along with some of its limit points. $\varphi \in \mathcal{G}$ may reconstruct any of the manifolds in the filtration.

Definition 3.9. *We say that φ is a proper observation function of V with respect to $\mathcal{G}$ if*

1. $\varphi \in Y_V$;
2. $\iota_V(\varphi) \in \mathcal{G}$;
3. *and there is no $U < V$ with $\gamma \in Y_U$ such that $\iota_{UV}(\gamma) = \varphi$.*

Corollary 3.10. *Let M be a smooth, compact manifold of dimension m , and let there be either a generic flow $(\phi^t, T) \in \mathbf{G}_P$ or a generic diffeomorphism $\phi \in \mathbf{G}_P$ as in Definition 3.5. Let the associated manifold filtration $\mathbb{M}$ be fully resolved and fulfill property C1. Let the generic set $\mathcal{G} \subset \bigcup Y_U$ be as in (8). Then if φ is a proper observation function of V with respect to $\mathcal{G}$, the delay reconstruction map*

$$\Phi_{(\phi^T, \varphi, k)}(x) := (\varphi(x), \varphi(\phi^T(x)), \dots, \varphi((\phi^T)^k(x)))$$

for flows or

$$\Phi_{(\phi, \varphi, k)}(x) := (\varphi(x), \varphi(\phi(x)), \dots, \varphi(\phi^k(x)))$$

for diffeomorphisms is an embedding of M_V into $\mathbb{R}^{k+1}$ when $k \geq 2n_V$.

Proof. We need to show that $\varphi \in \mathcal{G}_V$; then the statement follows from Theorems 3.1 and 3.3. From the definition of $\mathcal{G}$ in (8) and the inclusion relation in (6), the assumptions $\varphi \in Y_V$ and $\iota_V(\varphi) \in \mathcal{G}$ imply that $\varphi \in \mathcal{G}_V$ or $\varphi \in \iota_{UV}(\mathcal{G}_U)$ for some $U < V$. These states cannot be simultaneously true because $\mathbb{M}$ is fully resolved, and so the inclusion image of the generic set in U , $\iota_{UV}(\mathcal{G}_U)$, is in the excluded set $Y_V \setminus \mathcal{G}_V$ for $U < V$. By assumption, there is no $\gamma \in \mathcal{G}_U$ such that $\iota_{UV}(\gamma) = \varphi$, thus $\varphi \in \mathcal{G}_V$. ■

Corollary 3.10 partitions the set of observation functions Y . Consider an observation function $\varphi \in Y$ such that there exists $\gamma \in Y_U$ with $\iota_U(\gamma) = \varphi$. If φ is in the generic set $\mathcal{G}$ and the conditions of Corollary 3.10 are met, then φ provides a reconstruction of M_U . In particular, if $U = U_{\max}$, then $\varphi \in \mathcal{G}_{\max}$ reconstructs a diffeomorphism to M ; if $U < U_{\max}$, then $\varphi \in \mathcal{G} \setminus \mathcal{G}_{\max}$ (or $\gamma \in \mathcal{G}_U$) reconstructs a diffeomorphism to M_U ; and if $\varphi \in Y \setminus \mathcal{G}$ (or $\gamma \notin \mathcal{G}_U$), then it does not reconstruct M_U , although its reconstruction may be diffeomorphic to some other M_V .

3.3. Homeomorphisms vs. projections. Theorems 3.1 and 3.3 provide a basis for identifying dynamical drivers in deterministic systems when applied to the filtration structure $\mathbb{M}$ of the system. In this section we present our major result, a theorem that suggests an algorithm for recovering the transitive closures of the interaction and component graphs of the dynamical system. At this point, we will relax the idea of diffeomorphic manifolds to the notion of topological equivalence (see, for example, Arnol'd [1]). Two topologically equivalent systems have homeomorphic phase spaces, rather than diffeomorphic phase spaces. The presence or absence of homeomorphisms between manifolds in $\mathbb{M}$ is the property that allows the recovery of dynamical relationships.

Theorem 3.11. *Let M be a smooth, compact manifold of dimension m , and let there be either a generic flow $(\phi^t, T) \in \mathcal{G}_P$ or a generic diffeomorphism $\phi \in \mathcal{G}_P$. Let the associated manifold filtration $\mathbb{M}$ be fully resolved and fulfill property C1. Let $\mathcal{G}$ be the generic set of observation functions defined by (8). Consider a proper observation function φ_1 of U , and a proper observation function φ_2 of V . Denote the reconstructions $M_i := \Phi_{(\phi^T, \varphi_i, k_i)}(M)$ for flows or $M_i := \Phi_{(\phi, \varphi_i, k_i)}(M)$ for diffeomorphisms, $i = 1, 2$, with $k_1 \geq 2n_U$ and $k_2 \geq 2n_V$. Then the following hold:*

- (a) $U = V$ (i.e., both observation functions are from the same class $\mathcal{G}_U$) if and only if there is a homeomorphism

$$\Psi_{1,2} : M_1 \rightarrow M_2.$$

- (b) $U < V$ if and only if there is a continuous surjective, noninjective map

$$\Pi_{2,1} : M_2 \rightarrow M_1.$$

Proof. In the following, we will use the notation for a flow, with the understanding that the proof holds for either flows or diffeomorphisms.

- (a) *Forward direction:* As $U = V$ we will only use U . By Corollary 3.10, $\Phi_{(\phi^T, \varphi_1, k_1)} : M_U \rightarrow M_1$ and $\Phi_{(\phi^T, \varphi_2, k_2)} : M_U \rightarrow M_2$ are embeddings, thus diffeomorphisms onto their images. Defining $\Psi_{1,2} := \Phi_{(\phi^T, \varphi_2, k_2)} \circ \Phi_{(\phi^T, \varphi_1, k_1)}^{-1}$ gives a diffeomorphism, and therefore a homeomorphism, $M_1 \rightarrow M_2$, as shown in Figure 7(a).

Figure 7. (a) Induced homeomorphism $\Psi_{1,2}$ when $U = V$. (b) Induced homeomorphism Γ_{UV} when $\Psi_{1,2}$ exists.

Reverse direction: Assume that there exists a homeomorphism $\Psi_{1,2} : M_1 \rightarrow M_2$, and the relationship between M_U and M_V is unknown. As in the forward direction, the maps $\Phi_{(\phi^T, \varphi_1, k_1)} : M_U \rightarrow M_1$ and $\Phi_{(\phi^T, \varphi_2, k_2)} : M_V \rightarrow M_2$ are diffeomorphisms. So the map

$$\Gamma_{UV} := \Phi_{(\phi^T, \varphi_2, k_2)}^{-1} \circ \Psi_{1,2} \circ \Phi_{(\phi^T, \varphi_1, k_1)}$$

constructed from Figure 7(b) is a homeomorphism $M_U \rightarrow M_V$. Since (M, ϕ) has a fully resolved manifold structure, the existence of the homeomorphism Γ_{UV} ensures that U and V are comparable and that U and V are equal.

- (b) *Forward direction:* Now suppose $U < V$, so that the noninjective projection $\pi_{VU} : M_V \rightarrow M_U$ exists. As above, $\Phi_{(\phi^T, \varphi_1, k_1)}$ and $\Phi_{(\phi^T, \varphi_2, k_2)}$ are diffeomorphisms of M_U and M_V , respectively. Then

$$\Pi_{2,1} := \Phi_{(\phi^T, \varphi_1, k_1)} \circ \pi_{VU} \circ \Phi_{(\phi^T, \varphi_2, k_2)}^{-1} : M_2 \rightarrow M_1$$

is the desired noninjective surjection (Figure 8(a)).

Reverse direction: Now suppose the noninjective, surjective map $\Pi_{2,1} : M_2 \rightarrow M_1$ exists. The diffeomorphic maps $\Phi_{(\phi^T, \varphi_1, k_1)}$ and $\Phi_{(\phi^T, \varphi_2, k_2)}$ induce a noninjective projection

$$\bar{\Pi}_{VU} := \Phi_{(\phi^T, \varphi_1, k_1)}^{-1} \circ \Pi_{2,1} \circ \Phi_{(\phi^T, \varphi_2, k_2)} : M_V \rightarrow M_U$$

as in Figure 8(b). Since $\mathbb{M}$ is fully resolved, the existence of $\Pi_{2,1}$ ensures that U and V are comparable, and since the map $\bar{\Pi}_{VU}$ is noninjective, it must be true that $V > U$. ■

Theorem 3.11 implies that two variables reconstruct homeomorphic projections of M only if they belong to the same strongly connected component in CG_ϕ . We now propose an algorithm for establishing interactions between variables and prove that the transitive closures of the interaction and component graphs of the dynamical system can be recovered from the output of the algorithm.

3.4. Algorithm.

Algorithm 1.

1. Let the set of variables be the set of vertices in a graph that we will call RG .
2. For every variable x , build the reconstructed manifold $M_x := \Phi_{(\phi^T, \varphi_x, k)}(M)$, where φ_x depends only on x and k is large enough to work for the whole filtration.

Figure 8. (a) Induced projection $\Pi_{2,1} : M_2 \rightarrow M_1$ when $U < V$. (b) Induced projection $\bar{\Pi}_{VU} : M_V \rightarrow M_U$ when $\Pi_{2,1}$ exists.

3. For every pair x and y , test, using some criterion, whether there exists a continuous surjection $f : M_x \rightarrow M_y$, and if there exists a continuous surjection $g : M_y \rightarrow M_x$ such that if f and g are injective, $f^{-1} = g$.
4. If f exists, draw an edge $y \rightarrow x$ in RG , and if g exists, draw an edge $x \rightarrow y$.

Definition 3.12. We say that $U \subset UL_\phi$ is the minimal upper set of $x \in |U|$ if for every $V < U$ we have $x \notin |V|$. We denote a minimal upper set with a subscript of min: $x \in |U_{\min}|$.

Theorem 3.13. Let M be a smooth, compact manifold of dimension m , and let there be either a generic flow $(\phi^t, T) \in \mathcal{G}_P$ or a generic diffeomorphism $\phi \in \mathcal{G}_P$. Let the associated manifold filtration $\mathbb{M}$ be fully resolved and fulfill property C1. Further assume that given $x \in |U_{\min}|$, there exists a generic observation function $\varphi_x \in \mathcal{G}_U$, which depends only on observations of x . Then the following hold:

1. The graph RG constructed by the algorithm is the transitive closure of $\mathcal{IG}$, less any self-loops.
2. The graph of strongly connected path components of RG is the transitive closure of the component graph CG .

Proof.

1. By construction, the graph RG has the same set of vertices as $\mathcal{IG}$. Self-loops cannot be resolved by the algorithm, since M_x is always homeomorphic to itself regardless of the presence or absence of a self-loop at x in $\mathcal{IG}$.

Let $x \in |U_{\min}|$, $y \in |V_{\min}|$, and $w \in |W_{\min}|$ be arbitrary variables in the set of vertices, with $x \neq y$ and $x \neq w$. The generic maps $\varphi_x \in \mathcal{G}_U$, $\varphi_y \in \mathcal{G}_V$, and $\varphi_w \in \mathcal{G}_W$ exist as assumed. Note that $\mathbb{M}$ fully resolved, $x \in |U_{\min}|$, and $\varphi_x \in \mathcal{G}_U$ all together imply that φ_x is a proper observation function with respect to $\mathcal{G}$. Similarly for φ_y and φ_w . Then by Corollary 3.10, each of these induces a reconstruction that is diffeomorphic to the associated element of the manifold filtration $\mathbb{M}$; that is, $M_x = \Phi_{(\phi^T, \varphi_x, k_x)}(M)$ is diffeomorphic to M_U and so forth.

Suppose that there is an edge in $\mathcal{IG}$ from $x \in |U_{\min}|$ to $y \in |V_{\min}|$. Then U and V have the relation $U \leq V$ by the construction of UL_ϕ . By Theorem 3.11, if $U = V$, then the reconstructions M_x and M_y are homeomorphic; hence there is an edge from x to y in the graph RG . If instead $U < V$, then there is a continuous surjective map from M_y to M_x , and so there is an edge from x to y in RG .

Now suppose there is an edge from $x \in |U_{\min}|$ to $w \in |W_{\min}|$ in RG that does not exist

in $\mathcal{IG}$. By the algorithm, there must be a continuous surjection from M_w to M_x . By Theorem 3.11, $U \leq W$, which implies that there is a path from any variable in $|U|$ to any variable in $|W|$ along directed edges in $\mathcal{IG}$. Thus the edge $x \rightarrow w$ is an element of the transitive closure of $\mathcal{IG}$.

2. This follows from 1. ■

We discuss a specific implementation of Algorithm 1 in section 5. In the following section, we present the analogous theorems and algorithm for compact attractors instead of compact manifolds, based on the fractal delay embedding prevalence theorem of Sauer, Yorke, and Casdagli [16].

4. Extension to Attractors. To obtain similar results for attractors, the conditions which the dynamical system must obey are slightly different. We now discuss these technical restrictions. For a diffeomorphism, $\phi : \mathbb{R}^k \rightarrow \mathbb{R}^k$, a compact subset $A \subset \mathbb{R}^k$ of box-counting dimension d that is invariant under ϕ , and $n > 2d$, we require the following:

A1 For every positive integer $p \leq n$, the set of periodic points of period p has box-counting dimension less than $\frac{p}{2}$.

A2 The linearization $D\phi^p$ for each of these orbits has distinct eigenvalues.

Consider a flow ϕ^t on $\mathbb{R}^k$ and let A be a compact subset of $\mathbb{R}^k$ of box-counting dimension d that is invariant under ϕ . Let $n > 2d$ and $T > 0$. Assume the following:

B1 A contains at most a finite number of equilibria.

B2 There are no periodic orbits of period T or $2T$, and there are at most finitely many periodic orbits of period $3T, 4T, \dots, nT$.

B3 The linearizations of those periodic orbits have distinct eigenvalues.

We conjecture, similar to our result in Theorem A.6, that these assumptions are prevalent (as described below) in the set of flows ϕ^t or diffeomorphisms ϕ restricted to a fixed component graph CG. The proof remains open. In the proofs below, we assume that if A1 and A2 hold for ϕ or B1–B3 hold for ϕ^t , they also hold for all ϕ_U or ϕ_U^t , where U is a set in the upper set lattice UL .

4.1. Observation function filtration for attractors. In this section we phrase the work of Sauer, Yorke, and Casdagli [16] in our context. The importance of the results presented in [16] is that the invariant set of a dynamical system may not be a compact smooth manifold. In this setting, Sauer, Yorke, and Casdagli [16] replace Takens' notion of genericity of observations of dynamics on compact manifolds with a notion of "almost everywhere," or probability one, for observations of dynamics on compact attractors. It is not obvious how to generalize the notion of probability from a finite-dimensional space to an infinite-dimensional space, and the notion of *prevalence* [16] has been developed to address this issue.

Definition 4.1. A Borel subset S of a normed linear space V is prevalent if there is a finite-dimensional subspace E of V such that for each $v \in V$, $v + e$ belongs to S for (Lebesgue) almost every e in E .

We will proceed by considering the filtration of spaces of observation functions for each entry in the upper set lattice UL_ϕ . In this section, Y_U denotes the set of smooth observation functions $Y_U := \{\varphi : A_U \rightarrow \mathbb{R}\}$ with $\varphi \in C^1$ (this is relaxed from Takens' theorem, in which $\varphi \in C^2$ is required). The sets Y_U and $Y = \bigcup_U \iota_U(Y_U)$ form a filtration $\mathbb{Y}$, as discussed in the previous section.

We restate the fractal delay embedding prevalence theorem of [16] in our context, which is analogous to Theorems 3.1 and 3.3 in the previous section.

Theorem 4.2 (Sauer, Yorke, and Casdagli [16]). *Let $\phi : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a diffeomorphism satisfying the assumptions A1 and A2 above, or a flow satisfying B1–B3. For every observation function φ in a prevalent subset $\mathcal{P}_U \subset Y_U$ with $T > 0$, the delay reconstruction map of A_U ,*

$$\Phi_{(\phi^T, \varphi, k)}(x) := (\varphi(x), \varphi(\phi^T(x)), \dots, \varphi((\phi^T)^k(x)))$$

for flows or

$$\Phi_{(\phi, \varphi, k)}(x) := (\varphi(x), \varphi(\phi(x)), \dots, \varphi(\phi^k(x)))$$

for diffeomorphisms, is

1. one-to-one, and
2. an immersion on each compact subset C_U of a smooth manifold contained in A_U , provided that $k \geq 2 \operatorname{boxdim}(A_U)$, the box-counting dimension of A_U .

For the remainder of the work on attractors, a set $C_U \subset M_U \subset A_U$ will mean a compact subset of a smooth manifold contained in A_U . The second item in Theorem 4.2 states that $\Phi_{(\phi^T, \varphi, k)}$ restricted to C_U is differentiable with an injective derivative onto its image $\Phi_{(\phi^T, \varphi, k)}(C_U) \subset \mathbb{R}^{k+1}$. The following corollary gives us a result, which is compatible with the attractor filtration and more suitable for applications since it replaces differentiable functions with a homeomorphism.

Corollary 4.3. *Let $\phi : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a diffeomorphism satisfying the assumptions A1 and A2 above, or a flow satisfying B1–B3. Let $C_U \subset M_U \subset A_U$, let $\varphi \in \mathcal{P}_U$, and let $\Phi_{(\phi^T, \varphi, k)}$ or $\Phi_{(\phi, \varphi, k)}$ be the delay reconstruction map from Theorem 4.2 into $\mathbb{R}^{k+1}$. Then the delay map $\Phi_{(\phi^T, \varphi, k)}$ or $\Phi_{(\phi, \varphi, k)}$ restricted to C_U is a homeomorphism onto its image.*

Proof. Let C_I denote the image of C_U under the delay map, and let H_{Φ, C_U} denote the restricted map $H_{\Phi, C_U} : C_U \rightarrow C_I$. By Theorem 4.2, H_{Φ, C_U} is a differentiable surjection onto its image and is also one-to-one, and therefore a differentiable bijection. To be a homeomorphism, we need only that H_{Φ, C_U}^{-1} be continuous, which follows from a theorem in Munkres [13]. We repeat the brief argument here. Let D be an arbitrary closed subset of C_U . Since C_U is compact and D is closed, D is compact. Then H_{Φ, C_U} continuous implies $H_{\Phi, C_U}(D)$ is a compact subset of C_I . Since $C_I \subset \mathbb{R}^{k+1}$, C_I is Hausdorff, which implies that the compact subset $H_{\Phi, C_U}(D)$ is closed. Since the inverse image of an arbitrary closed subset in the range of H_{Φ, C_U}^{-1} is closed, H_{Φ, C_U}^{-1} is continuous. ■

In order to apply Corollary 4.3 to $\mathbb{A}$, we need the idea of a fully resolved attractor structure, analogous to Definition 3.8. It is useful to introduce the notation $\bigcup M_U$, denoting the union of all smooth manifolds in the attractor A_U .

Definition 4.4. *The dynamical system (A, ϕ) has a fully resolved attractor filtration $\mathbb{A}$ if the following two conditions are met:*

1. *Given comparable U, V in the lattice UL_ϕ with $U < V$, the corresponding projection $\pi_{VU} : A_V \rightarrow A_U$ is noninjective and is a surjection from $\bigcup M_V \subset A_V$ to $\bigcup M_U \subset A_U$.*
2. *Given noncomparable U, V in the lattice UL_ϕ , there is no surjection $f : \bigcup M_U \rightarrow \bigcup M_V$ that is also continuous when restricted to any compact subset $C_U \subset M_U \subset A_U$.*

As in the previous section, the first condition ensures that comparable, nonequal upper sets are associated with distinguishable attractors by the relation $\iota_{UV}(\mathcal{P}_U) \subset Y_V \setminus \mathcal{P}_V$, and the

second condition is required to recover the full set of strongly connected components in CG. Additionally, the requirement in the first condition that smooth manifolds map to smooth manifolds is analogous to assumption C1 in section 3.

Similar to Corollary 3.10 we have the following.

Corollary 4.5. *Let $\phi : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a diffeomorphism satisfying the assumptions A1 and A2 above, or a flow satisfying B1–B3. Let $\mathbb{A}$ be the attractor filtration, and let $\mathbb{Y}$ be the filtration of observation functions associated to ϕ . Assume $\mathbb{A}$ is fully resolved. Then there exists a prevalent set $\mathcal{P} \subset Y$ such that if φ is a proper observation function of V with respect to $\mathcal{P}$, then the delay reconstruction map $\Phi_{(\phi^T, \varphi, k)}(x)$ or $\Phi_{(\phi, \varphi, k)}(x)$ is one-to-one on A_V and a homeomorphism from any $C_V \subset M_V \subset A_V$ onto its image in $\mathbb{R}^{k+1}$ when $T > 0$ and $k \geq 2 \text{ boxdim}(A_V)$.*

Proof. Analogously to (8) we define a set

$$\mathcal{P} := \bigcup_{U \in UL} \iota_U(\mathcal{P}_U).$$

The proof that $\varphi \in \mathcal{P}_V$ is the same as Corollary 3.10, because it depends only on $\iota_{UV}(\mathcal{P}_U) \subset Y_V \setminus \mathcal{P}_V$, which has an exact analogue for manifolds. Since $\varphi \in \mathcal{P}_V$, the result follows directly from the application of Theorem 4.2 and Corollary 4.3. ■

4.2. Homeomorphisms vs. projections on subsets of attractors. We present our major result for attractors, a theorem analogous to Theorem 3.11 that suggests an algorithm for recovering the interactions between variables in a dynamical system. The proof is quite similar to that of Theorem 3.11, and the commuting diagrams in that proof may be useful to the reader here.

Theorem 4.6. *Consider a dynamical system ϕ that is either a diffeomorphism satisfying A1 and A2, or a flow satisfying B1–B3, its attractor A , associated attractor filtration $\mathbb{A}$, and observation filtration $\mathbb{Y}$. Assume that $\mathbb{A}$ is fully resolved. Let $\mathcal{P}$ be the prevalent set of observation functions guaranteed by Corollary 4.5. Let φ_1 be a proper observation function of U , and let φ_2 be a proper observation function of V . Let $T > 0$ be the chosen time lag for the flow. Let $A_i := \Phi_{(\phi^T, \varphi_i, k_i)}(A)$ for flows or $A_i := \Phi_{(\phi, \varphi_i, k_i)}(A)$ for diffeomorphisms, $i = 1, 2$, with $k_1 \geq 2 \text{ boxdim}(A_U)$ and $k_2 \geq 2 \text{ boxdim}(A_V)$. Let C_U and C_V be arbitrary compact subsets of manifolds of A_U and A_V , respectively, and define their images as C_1 and C_2 , respectively, under the delay map. Furthermore let $\bigcup M_1 \subset A_1$ denote the union of images of the smooth manifolds of A_U , and likewise for $\bigcup M_2 \subset A_2$. Then the following hold:*

- (a) $U = V$ (i.e., both observation functions are from the same class $\mathcal{P}_U$) if and only if there is an injective map

$$\Psi_{1,2} : A_1 \rightarrow A_2$$

that is a bijection from $\bigcup M_1 \subset A_1$ to $\bigcup M_2 \subset A_2$ and a homeomorphism when restricted to any C_1 .

- (b) $U < V$ if and only if there is a noninjective map

$$\Pi_{2,1} : A_2 \rightarrow A_1$$

that is a surjection from $\bigcup M_2$ to $\bigcup M_1$ and continuous when restricted to any C_2 .

Proof. As in the previous section, we will use only the notation for flows, noting that the proof holds when substituting the notation for diffeomorphisms.

- (a) *Forward direction:* As $U = V$ we will only use U . By Corollary 4.5, $\Phi_{(\phi^T, \varphi_1, k_1)}$ and $\Phi_{(\phi^T, \varphi_2, k_2)}$ are one-to-one on A_U and homeomorphisms on any C_U . Then the composition map $\Psi_{1,2} := \Phi_{(\phi^T, \varphi_2, k_2)} \circ \Phi_{(\phi^T, \varphi_1, k_1)}^{-1} : A_1 \rightarrow A_2$ is injective, and is a homeomorphism on any $C_1 \rightarrow C_2$, where C_1 and C_2 are both images of C_U . Notice that $\Phi_{(\phi^T, \varphi_1, k_1)}$ and $\Phi_{(\phi^T, \varphi_2, k_2)}$ are not guaranteed continuous on $\bigcup M_U$, but since $\Phi_{(\phi^T, \varphi_1, k_1)}$ and $\Phi_{(\phi^T, \varphi_2, k_2)}$ are bijective on $\bigcup M_U$, $\Psi_{1,2}$ is a bijection from $\bigcup M_1$ to $\bigcup M_2$.

Reverse direction: We assume that there exists a one-to-one map $\Psi_{1,2} : A_1 \rightarrow A_2$ that is bijective on $\bigcup M_1$ and a homeomorphism when restricted to C_1 . As above, by Corollary 4.5, $\Phi_{(\phi^T, \varphi_1, k_1)}$ and $\Phi_{(\phi^T, \varphi_2, k_2)}$ are one-to-one on A_U and A_V , and are homeomorphisms on any C_U and C_V . Then the map

$$\Gamma_{UV} := \Phi_{(\phi^T, \varphi_2, k_2)}^{-1} \circ \Psi_{1,2} \circ \Phi_{(\phi^T, \varphi_1, k_1)} : A_U \rightarrow A_V$$

is injective on A_U , bijective on $\bigcup M_U$, and a homeomorphism on any C_U . Since $\mathbb{A}$ is fully resolved, U and V are comparable because Γ_{UV} exists, and $U = V$ because Γ_{UV} is injective.

- (b) *Forward direction:* Suppose $U < V$, so that the noninjective, continuous projection (surjection) $\pi_{VU} : A_V \rightarrow A_U$ exists. π_{VU} maps the compact sets $C_V \subset M_V \subset A_V$ surjectively to the compact sets $C_U \subset M_U \subset A_U$ by assumption. Since the maps $\Phi_{(\phi^T, \varphi_1, k_1)}$ and $\Phi_{(\phi^T, \varphi_2, k_2)}$ are injective on A_U and A_V , their inverses exist, leading to the well-defined map

$$\Pi_{2,1} := \Phi_{(\phi^T, \varphi_1, k_1)} \circ \pi_{VU} \circ \Phi_{(\phi^T, \varphi_2, k_2)}^{-1} : A_2 \rightarrow A_1,$$

which is noninjective. Furthermore, $\Pi_{2,1}$ maps $C_2 \rightarrow C_V \rightarrow C_U \rightarrow C_1$ in a composition of continuous surjections; i.e., it is a continuous surjection itself on C_2 . As proved in part (a), $\Pi_{2,1}$ is not continuous on the union $\bigcup M_2$, but it is a surjection onto $\bigcup M_1$.

Reverse direction: Now suppose the noninjective map $\Pi_{2,1} : A_2 \rightarrow A_1$ exists, is a surjection from $\bigcup M_2$ to $\bigcup M_1$, and is continuous on any C_2 . Since $\Phi_{(\phi^T, \varphi_1, k_1)}$ and $\Phi_{(\phi^T, \varphi_2, k_2)}$ are injective, their inverses exist, so the map

$$\bar{\Pi}_{VU} := \Phi_{(\phi^T, \varphi_1, k_1)}^{-1} \circ \Pi_{2,1} \circ \Phi_{(\phi^T, \varphi_2, k_2)} : A_V \rightarrow A_U$$

is well defined and noninjective via $\Pi_{2,1}$. $\bar{\Pi}_{VU}$ is a continuous surjection on any C_2 , implying that $\bar{\Pi}_{VU}$ is a surjection from $\bigcup M_2$ to $\bigcup M_1$. Since $\mathbb{A}$ is fully resolved, the existence of $\bar{\Pi}_{VU}$ ensures that U and V are comparable, and since the map $\bar{\Pi}_{VU}$ is noninjective, it must be true that $V > U$. ■

The notion of fully resolved attractor filtration reflects not only a desire to have distinct attractors associated to different upper sets in the lattice UL , but also the need for these characteristics to be computable, at least approximately, from finite data. For two reconstructions of the same attractor in a fully resolved system, Theorem 4.6 asserts the existence of a map that is a homeomorphism when restricted to compact subsets of smooth manifolds in an attractor. That is, there are homeomorphisms between reconstructions of compact sets C_1 and C_2 that have observation functions from the same prevalent set. It is difficult to verify such a

condition computationally since we do not know all smooth manifolds that are embedded in the attractor. Instead, we only concentrate on the collection of unstable manifolds of points in the attractor, since these are embedded manifolds in the attractor. We denote this collection as $\bigcup_a W_U^u(a) \subset A_U$, where $a \in A_U$.

Definition 4.7. *The dynamical system (A, ϕ) has a strongly resolved attractor filtration $\mathbb{A}$ if the following two conditions are met:*

1. *Given comparable U, V in the lattice UL_ϕ with $U < V$, the corresponding projection $\pi_{VU} : A_V \rightarrow A_U$ is a noninjective surjection from $\bigcup_b W_V^u(b) \subset A_V$ to $\bigcup_a W_U^u(a) \subset A_U$.*
2. *Given noncomparable U, V in the lattice UL_ϕ , there is no surjection $f : \bigcup_a W_U^u(a) \rightarrow \bigcup_b W_V^u(b)$ that is also continuous when restricted to any compact subset of an unstable manifold $C_U \subset W_U^u(a) \subset A_U$.*

4.3. Algorithm.

Algorithm 2.

1. *Let the set of variables be the set of vertices in a graph that we will call RG .*
2. *For every variable x , build the reconstructed attractor $A_x := \Phi_{(\phi^T, \varphi_x, k_x)}(A)$, where φ_x depends only on x .*
3. *For every pair x and y , test, using some criterion, whether there is a surjection*

$$f : \bigcup_a W_1^u(a) \subset A_x \rightarrow \bigcup_b W_2^u(b) \subset A_y$$

that is continuous on any $C_1 \subset W_1^u(a) \subset A_x$, and a surjection

$$g : \bigcup_b W_2^u(b) \subset A_y \rightarrow \bigcup_a W_1^u(a) \subset A_x$$

that is continuous on any $C_2 \subset W_2^u(b) \subset A_y$, such that if f and g are injective, then $f^{-1} = g$. Here, $W_1^u(a)$ and $W_2^u(b)$ are the delay images of the unstable manifolds $W_U^u(a) \subset A_U$ and $W_V^u(b) \subset A_V$ for $x \in |U_{\min}|$, $y \in |V_{\min}|$.

4. *If f exists, draw an edge $y \rightarrow x$ in RG , and if g exists, draw an edge $x \rightarrow y$.*

Theorem 4.8. *Let $\phi : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a diffeomorphism satisfying the assumptions A1 and A2 above, or a flow satisfying B1–B3. Assume that (A, ϕ) has a strongly resolved attractor structure. Further, given $x \in |U_{\min}|$, assume there exists a prevalent observation function $\varphi_x \in \mathcal{P}_U$ that depends only on observations of x . Then the following hold:*

1. *The graph RG constructed by the algorithm is the transitive closure of $\mathcal{IG}$, less any self-loops.*
2. *The graph of strongly connected path components of RG is the transitive closure of the component graph CG .*

Proof. This proof follows the proof of Theorem 3.13 very closely with the following substitutions: strongly resolved replaces fully resolved; prevalent replaces generic; compact subsets of unstable manifolds $C_U \subset W_U^u(a) \subset A_U$ replace manifolds M_U ; and Theorem 4.6 replaces Theorem 3.11. ■

Algorithm 2 involves only a slight modification of Algorithm 1. We discuss a specific implementation of this modification in the following section.

5. Practical detection of projections. The main result of our work is that one-way interactions between the strongly connected components of the interaction graph of a dynamical system are identified with one-way surjections between corresponding reconstructions in a fully resolved manifold filtration or strongly resolved attractor filtration. In this section, we briefly describe our implementation of state space reconstruction, we summarize work by Pecora, Carroll, and Heagy [15] for detecting continuous functions between delay reconstructions of invariant compact manifolds, and then we introduce a modification for compact attractors. Finally, we present several examples using the unmodified version of Pecora, Carroll, and Heagy [15]; that is, we assume the presence of an invariant manifold filtration for the dynamical systems.

5.1. State space reconstruction. Given two equally spaced and equally long time series $x = (x(0), x(\tau), x(2\tau), \dots, x(n\tau))$ and $y = (y(0), y(\tau), y(2\tau), \dots, y(n\tau))$, our task is to produce M_x and M_y . For each reconstruction we must choose an observation function φ , an appropriate reconstruction dimension d , and an integer index lag i , where $T = i\tau$ is the time lag mentioned in the previous sections. In practice we choose φ to be the projection function, i.e., $\varphi_x(M) = x$ and $\varphi_y(M) = y$, and assume that these are generic maps in $\mathcal{G}_U$ and $\mathcal{G}_V$, respectively, where U and V are the minimal upper sets of x and y . To choose integer index lags i_x and i_y , we use as a guide the first zero of the autocorrelation function of the time series x and y , respectively. The time lags $T_x = i_x\tau$ and $T_y = i_y\tau$ are assumed generic.

A point in each reconstruction is given by

$$(9) \quad \begin{aligned} u(j) &= (x(j\tau), x((j - i_x)\tau), \dots, x((j - (d - 1)i_x)\tau)), \\ v(j) &= (y(j\tau), y((j - i_y)\tau), \dots, y((j - (d - 1)i_y)\tau)), \end{aligned}$$

where we will consistently use u to denote a point in the reconstruction M_x and v to denote a point in the reconstruction M_y . The index j uniquely identifies each point on either manifold. We do not require that j be a scalar multiple of i_x or i_y . The reconstruction is built as a sliding window across the full length of both time series, instead of only on the points that are multiples of i_x or i_y . This choice leads to a smoother reconstruction and also makes it easier to identify contemporaneous points in M_x and M_y . In particular, M_x has $n + 1 - (d - 1)i_x$ points and, similarly, M_y has $n + 1 - (d - 1)i_y$ points (there are $n + 1$ points in each of x and y). If $i_x < i_y$, then the last $n + 1 - (d - 1)i_y$ points of M_x are contemporaneous with M_y , in the sense that the first coordinate of both occurs at the same time. From now on, we will restrict M_x (or M_y if $i_y < i_x$) to this subset, so that M_x and M_y have the same number of points N with temporal coincidence in the first coordinate.

There is substantial literature addressing the choice of d and T in state space reconstruction; see, e.g., [2, 3, 4, 5, 6, 11, 12]. We have chosen deliberately simple methods to demonstrate the recovery of the transitive closure of CG from reconstructions. In all of our examples we know the dimensionality of the system, and we choose d to be the known dimensionality. Given a set of noisy time series and no a priori knowledge about a dynamical system, we would choose a more general reconstruction method such as the one in Casdagli et al. [4].

5.2. Continuity testing. We want to test for the existence of continuous surjections $f : M_x \rightarrow M_y$ and $g : M_y \rightarrow M_x$, which are inverses when injective. Natural candidates for

these functions are the contemporaneous maps given by $f(u(j)) = v(j)$ and $g(v(j)) = u(j)$ for $u(j) \in M_x$, $v(j) \in M_y$. Given the earlier restriction of M_x and M_y to contemporaneous points, this map is naturally surjective. If both f and g exist, then $f^{-1} = g$ and there is a bijection between M_x and M_y . If both are continuous, then M_x and M_y are homeomorphic.

Pecora, Carroll, and Heagy [15] begin with the definition of continuity: if f is continuous at $u(j) \in M_x$, then for every $\epsilon > 0$ there exists $\delta > 0$ such that $\|u - u(j)\| < \delta$ implies $\|f(u) - f(u(j))\| < \epsilon$. They note that it is insufficient to merely find such a δ for a given ϵ ; it's necessary to have a null hypothesis against which the likelihood of finding a δ can be assessed. They propose a null hypothesis that the points in M_y are independently and randomly distributed with respect to the points in M_x , although they emphasize that other null hypotheses can be used. This gives rise to a schema for accepting or rejecting the continuity of f and g .

Given $\epsilon > 0$ and $\delta > 0$, define punctured open sets in M_x and M_y :

$$B_\delta(u) := \{z \in M_x, z \neq u : \|z - u\| < \delta\},$$

$$B_\epsilon(f(u)) := \{z \in M_y, z \neq f(u) : \|z - f(u)\| < \epsilon\}.$$

Pecora, Carroll, and Heagy [15] present the following algorithm:

1. For a fixed ϵ , guess an initial δ . If $f(B_\delta(u)) \subseteq B_\epsilon(f(u))$, continue to the next step. Otherwise, reduce δ and repeat until successful or until $B_\delta(u) = \emptyset$. This step ensures that the points of interest are consistent with the definition of continuity.
2. Let $n_\delta = |B_\delta(u)|$ be the number of points in the δ -ball, and similarly let $n_\epsilon = |B_\epsilon(f(u))|$. The probability under the null hypothesis that $f(B_\delta(u)) \subseteq B_\epsilon(f(u))$ is given by $p = (n_\epsilon/N)^{n_\delta}$, where N is the number of points in M_y .
3. If this probability is sufficiently low, the null hypothesis may be rejected at u . The pointwise confidence for rejecting the null hypothesis is given by $\Theta(\epsilon, u) = 1 - (p/p_{\max})$, where $p_{\max}$ is the maximum of the binomial distribution

$$b(k; n_\delta, p) = \binom{n_\delta}{k} p^k (1-p)^{n_\delta-k}$$

and the maximum value occurs at $k = \text{floor}((n_\delta + 1)p)$.

4. Calculate a scalar confidence for the whole manifold by averaging the pointwise confidences over a randomly chosen set of n points indexed by the set I :

$$\Theta_{xy}(\epsilon) = \frac{1}{n} \sum_{j \in I} \Theta(\epsilon, u(j)).$$

Remark 1. In a similar fashion, the confidence in the reverse direction, $\Theta_{yx}(\epsilon)$ for $g : M_y \rightarrow M_x$, may be calculated. The (ϵ, δ) pairs in the $\Theta_{yx}(\epsilon)$ calculation will be entirely independent of the values used in the $\Theta_{xy}(\epsilon)$ calculation.

If $\Theta_{xy}(\epsilon)$ is near 1, then we accept that the points mapping from δ -balls in M_x to ϵ -balls in M_y are nonrandom events. This constitutes circumstantial evidence for the continuity of $f : M_x \rightarrow M_y$, since the points were chosen to be consistent with the definition of continuity. Similarly, the continuity of g is supported for high values of $\Theta_{yx}(\epsilon)$.

There is not a clear way to choose an appropriate ϵ for a given problem. If ϵ is too small, then $n_\epsilon = 0$ and continuity will never be observed. But if ϵ is too large, e.g., half the size of the manifold, then continuity may be spuriously detected. Like Pecora, Carroll, and Heagy [15], we test a number of different ϵ values nondimensionalized by the standard deviations of the reconstructions. By this we mean the standard deviation of the distance from each point to the mean of the reconstruction in $\mathbb{R}^d$. A continuous function should ideally show a monotonic increase in Θ_{xy} as ϵ increases, and should show a relatively high Θ_{xy} for relatively low values of ϵ .

As a time series increases in length, the reconstructed points fill in the invariant compact manifold. If the functions f and g are continuous, then the confidence of continuity should increase as more points are included in the calculation. So we also calculate Θ_{xy} and Θ_{yx} using an increasing number of points from the time series x and y . That is, we make a family of reconstructions, $M_x^{N_1}, M_x^{N_2}, \dots, M_x^{N_k}$, where $N_1, N_2, \dots, N_k$ indicate an increasing number of points in the reconstructions. If the family of $\Theta_{xy}(\epsilon)$ curves converges upward toward 1 with increasing N_i , then we have strong circumstantial evidence for continuity.

The fact that we draw conclusions based on the contemporaneous functions $f(u(j)) = v(j)$ and $g(v(j)) = u(j)$ comes from the assumption that the underlying dynamical system does not have an explicit delay. If we suspect that there is a delay in the system, then we should check continuity of the delay maps $f_k(u(j)) = v(j+k)$ and $g_k(v(j)) = u(j-k)$ for some k . In practice, it is not possible to test all such maps, and in our differential equation examples the contemporaneous maps are sufficient.

5.3. Continuity testing between compact subsets of unstable manifolds. To approximate compact subsets of unstable manifolds in reconstructed attractors A_x and A_y , we use the notation introduced in (9) for $u(j) \in A_x$ and contemporaneous point $v(j) \in A_y$. The induced dynamics on A_x are given by the shift map $F(u(j)) := u(j+1)$ and on A_y by $G(v(j)) := v(j+1)$. Take a point $u \in A_x$. To approximate a δ -neighborhood of the unstable manifold $W_x^u(u) \subset A_x$, we look at the forward image of neighborhood surrounding a preimage of u . Let

$$U_\delta(u) := \{z \in A_x, z \neq u : \|u - z\| < \delta\}.$$

Then

$$B_\delta^n(u) := U_\delta(u) \cap F^n(U_\delta(F^{-n}(u))).$$

Note that $B_\delta^n(u)$ approximates a δ -neighborhood on the unstable manifold $W_x^u(u) \subset A_x$ for large n . Similarly, we define the corresponding neighborhoods in $W_y^u(f(u))$

$$B_\epsilon^n(f(u)) := U_\delta(f(u)) \cap G^n(U_\delta(G^{-n}(f(u)))),$$

which approximates the ϵ -neighborhood on the unstable manifold $W_y^u(f(u)) \subset A_y$ for large n .

We then test continuity of a contemporaneous map $f : A_x \rightarrow A_y$ on compact subsets of unstable manifolds by fixing moderate to large values of n and then proceed with the previous algorithm restricted to sets $B_\delta^n(u)$ and $B_\epsilon^n(f(u))$.

5.4. Chaotic oscillator example. We begin with a two-tier chaotic oscillator written as a set of first order equations:

$$\begin{aligned}\dot{x} &= y, \\ \dot{y} &= \mu(1 - x^2)y - x, \\ \dot{z} &= w, \\ \dot{w} &= -w - \beta \sin(z) + A \sin(x).\end{aligned}$$

The interaction and component graphs are given in Figure 9.

Figure 9. $\mathcal{IG}$ and $\mathcal{CG}$ for the chaotic oscillator example.

The parameter choices and initial conditions we used are

$$\begin{aligned}\mu &= 4.0, \quad \beta = 1.2, \quad A = 2.0, \\ x(0) &= 1.0, \quad z(0) = 3.0, \quad y(0) = w(0) = 2.0.\end{aligned}$$

The equations were solved using the Runge–Kutta fourth order algorithm with a fixed time step of $\tau = 0.025$. The result is a chaotic attractor in a bounded region of space; see the phase portraits in Figure 10. There are nearby parameter choices that result in an unbounded attracting invariant set, and state space reconstruction techniques are not guaranteed to work on such attractors. The reconstructions M_x , M_y , M_z , M_w were created in $\mathbb{R}^4$ with a time lag of $T = 100\tau$ for every variable except z , which had a lag of $T_z = 115\tau$.

To calculate the continuity confidences, we had to choose a set of time series lengths to build reconstructions, a set of ϵ values, and sets of random points on which to perform the continuity test. The solution trajectory was simulated out to time 1200, and a family of reconstructions was created using 30–100% of the full time series. The continuity test was performed on 10% of the points in each reconstruction chosen randomly. The points were chosen independently for each reconstruction and each pairwise comparison. The set of ϵ values sampled the interval between 0.02 and 0.4 standard deviations of the full length reconstruction for each variable. We did not recalculate the standard deviation for reconstructions of shorter lengths.

The component graph $\mathcal{CG}$ was recovered as expected. In particular, the homeomorphisms $M_x \leftrightarrow M_y$ and $M_z \leftrightarrow M_w$ were confirmed, and one-way continuous surjections were seen going from M_z , M_w to M_x , M_y . The confidence graphs for continuous $M_x \rightarrow M_w$ and $M_w \rightarrow M_x$ are shown as illustrative examples in Figure 11. A set of $\Theta(\epsilon)$ curves is plotted, where the curves are parameterized by the length of the reconstruction. The legend shows the proportion of points in the time series used to make the reconstructions. The direction $M_x \rightarrow M_w$ is associated with a one-way interaction from w to x . The confidence values are

Figure 10. Three-dimensional phase spaces for the chaotic oscillator example.

Figure 11. Chaotic oscillator example. Families of curves of Θ_{xw} and Θ_{wx} as functions of ϵ scaled by the standard deviation of M_w and M_x , respectively. The family of curves is parameterized by the proportion of points used to make the reconstructions. Θ_{wx} converges upward toward 1 with increasing ϵ and increasing reconstruction length, which is evidence for continuity and an edge $x \rightarrow w$. Θ_{xw} decreases with increasing reconstruction length and has low confidence, so we conclude $w \not\rightarrow x$.

low (< 0.2) over the range of tested ϵ values. Moreover, as the reconstruction becomes longer, the confidence decreases. So we conclude correctly that w is not a dynamical driver of x . In the other direction, $M_w \rightarrow M_x$, we see a robust upward convergence with increasing ϵ and reconstruction length, indicating strong evidence for a continuous function. We conclude that x is a driver of w and draw an edge from x to w in the interaction graph.

We summarize all the variable relationships in Table 1. In the left table, we give the highest confidence value for each pairwise comparison with ϵ at 0.4 standard deviations and with full length reconstructions. In the right table, we say whether the family of $\Theta(\epsilon)$ curves converged upward or downward or was coincident as the reconstruction length increased. If

Table 1

Chaotic oscillator example. Left: Confidence values from row reconstruction to column reconstruction when ϵ is 0.4 standard deviations of the column reconstruction and 100% of the time series is used in the reconstruction. Right: The behavior of the family of $\Theta(\epsilon)$ curves with increasing reconstruction length: upward convergence (u), downward convergence (d), or coincident curves (c).

	M_x	M_y	M_z	M_w
M_x	.	1.00	0.04	0.05
M_y	1.00	.	0.03	0.04
M_z	0.95	0.98	.	1.00
M_w	0.98	0.99	0.99	.

	M_x	M_y	M_z	M_w
M_x	.	c	d	d
M_y	c	.	d	d
M_z	u	u	.	u
M_w	u	u	u	.

Figure 12. Chaotic oscillator example. Recovered graphs $\tilde{\mathcal{G}}$ and $\bar{\mathcal{G}}$ from continuity testing. Dashed edges are not present in the original interaction graph in Figure 9.

there are high confidence values and anything but downward convergence, we conclude the presence of a continuous surjection and an appropriate edge in the reconstructed interaction graph. With downward convergence or low confidence values, we conclude the absence of an edge. The reconstructed graphs are given in Figure 12.

5.5. Diamond graph example. The point of this example is to explore a more complex component graph shaped like a diamond. We add a second branch to the chaotic oscillator example consisting of a Rössler system driven by the periodic oscillator. Both the chaotic oscillator and the Rössler systems feed into the evolution of an additional variable. The system of equations for the diamond example are

$$\begin{aligned}
 \dot{x} &= y, \\
 \dot{y} &= \mu(1 - x^2)y - x, \\
 \dot{z} &= w, \\
 \dot{w} &= -w - \beta \sin(z) + A \sin(x), \\
 \dot{s} &= -u - v, \\
 \dot{u} &= s + au, \\
 \dot{v} &= b + v \left(\frac{2 + C(x - 3)}{2} s - c \right) - \frac{C(x - 3)}{4} (s^2 + v^2), \\
 \dot{p} &= -p + B \sin(v) \sin(z).
 \end{aligned}$$

When the parameter C is 0, the equations for s, u, v reduce to the usual Rössler system. When C is nonzero, there is an additional nonlinear term not present in the original equations. The interaction and component graphs for the whole system are given in Figure 13.

Figure 13. $\mathcal{IG}$ and $\mathcal{CG}$ for the diamond example.

The parameter choices and initial conditions we used are

$$\mu = 4.0, \quad \beta = 1.2, \quad A = 2.0, \quad B = 1.25, \quad C = 0.20, \quad a = 0.2, \quad b = 0.2, \quad c = 5.7, \\ x(0) = 1.0, \quad y(0) = w(0) = 2.0, \quad z(0) = 3.0, \quad s(0) = u(0) = 0.5, \quad v(0) = 0.1, \quad p(0) = 0.75.$$

All of the chaotic oscillator parameters and numerical parameters are the same as in the previous example, except that the reconstructions were created in $\mathbb{R}^8$ with a time lag of $T = 95\tau$ for M_w , $T = 60\tau$ for M_s , $T = 65\tau$ for M_u , $T = 50\tau$ for M_v , and $T = 100\tau$ for M_p .

Continuity testing was performed in the same manner as before, and the transitive closures of the graphs in Figure 13 were recovered. We show example confidence graphs in Figure 14. When testing $M_p \rightarrow M_v$, the curves $\Theta_{pv}(\epsilon)$ coincide very closely and converge toward 1, indicating an edge from $v \rightarrow p$ in the recovered interaction graph $\tilde{\mathcal{IG}}$. In the opposite direction, $\Theta_{vp}(\epsilon)$ attains moderately high confidence at short reconstruction lengths, but the confidence converges downward as the reconstruction length increases. Thus we conclude that there is not a reciprocal edge from p to v . The graph of confidences for $M_s \rightarrow M_x$ has only moderate values, but $\Theta_{sx}(\epsilon)$ converges robustly upward with reconstruction length. We therefore conclude there is an edge from x to s .

Summary data for all pairwise comparisons are listed in Table 2. The recovered graphs constructed from the confidence values and convergence trends are shown in Figure 15, where dashed lines denote connections not found in the original graphs.

5.6. Measurement noise. It is obvious that sufficiently high noise can overwhelm a deterministic signal. In this section, we characterize how much *measurement noise* is required to interfere with the recovery of the transitive closure of the component graph of the chaotic oscillator example. See section 5.4 for the equations and parameters.

By measurement noise we mean noise that is introduced when making observations of the manifold M , so that the result is a perturbation of the observation function: $\varphi + \xi(t)$. The noise function $\xi(t)$ is a sequence of realizations of a probability distribution $\mathcal{N}(\mu, \sigma)$ with mean μ and standard deviation σ , often taken to be the normal distribution. Measurement noise is distinct from stochastic elements in the dynamical system (dynamic noise) and from estimation errors that arise from using a finite length time series in the reconstruction [4].

Let $\sigma_x, \sigma_y, \sigma_z$, and σ_w denote the standard deviations of the deterministic time series x, y, z , and w . We perturb the time series with independent realizations $\xi_x(t), \xi_y(t), \xi_z(t), \xi_w(t)$,

Figure 14. Diamond example. Confidences Θ_{pv} , Θ_{sx} , and Θ_{vp} vs ϵ scaled by the standard deviation of the ranges (M_v , M_x , and M_p). The legend indicates the percentage of points used to make the reconstructions. The coincident set of Θ_{pv} curves converges upward toward 1 with increasing ϵ , implying continuity and an edge $v \rightarrow p$. Θ_{sx} shows the same robust upward trend with lower, but still acceptable, confidence levels that support $x \rightarrow s$. Θ_{vp} decreases with increasing reconstruction length, implying $p \not\rightarrow v$, even though the confidence levels are moderate.

according to the normal distributions $\mathcal{N}(0, q\sigma_x)$, $\mathcal{N}(0, q\sigma_y)$, $\mathcal{N}(0, q\sigma_z)$, $\mathcal{N}(0, q\sigma_w)$, where the parameter q increases from 0.01 to 0.1. That is, we add noise with 1% to 10% of the standard deviation of the time series, and perform continuity testing precisely as before on the perturbed series $\hat{x} = x(t) + \xi_x(t)$, $\hat{y} = y(t) + \xi_y(t)$, etc.

In Figure 16, we show the decrease in continuity confidence that occurs with increasing noise. The noise levels are given in the legend as the percentage of the time series standard deviation. All reconstructions used 100% of the time series (up to time 1200 as before). The confidence in the $M_x \rightarrow M_y$ direction, $\Theta_{xy}(\epsilon)$, is fairly robust to increasing noise of this magnitude, indicating that the edge $y \rightarrow x$ in $\mathcal{IG}$ will be consistently recovered. However, $\Theta_{yx}(\epsilon)$ decays rapidly, and so the edge from $x \rightarrow y$ is lost with noise at 3–4% of the standard deviation. The two graphs in Figure 16 exemplify the patterns seen in all pairwise comparisons

Table 2

Diamond example. Pairs of confidence values and convergence trends from row to column reconstructions. Confidence values are taken when ϵ is 0.4 the standard deviation of the column reconstruction and 100% of the time series is used in the reconstruction. Convergence trends in $\Theta(\epsilon)$ with increasing reconstruction length are denoted as u (upward convergence), d (downward convergence), or c (coincident curves). Values for the chaotic oscillator have been elided.

	M_x	M_y	M_z	M_w	M_s	M_u	M_v	M_p
M_x	.	.	.	.	(0.00, c)	(0.00, c)	(0.00, c)	(0.00, c)
M_y	.	.	.	.	(0.00, c)	(0.00, c)	(0.00, c)	(0.00, c)
M_z	.	.	.	.	(0.28, d)	(0.28, d)	(0.32, d)	(0.29, d)
M_w	.	.	.	.	(0.17, d)	(0.18, d)	(0.21, d)	(0.17, d)
M_s	(0.64, u)	(0.74, u)	(0.06, d)	(0.08, d)	.	(0.99, u)	(0.99, u)	(0.09, d)
M_u	(0.64, u)	(0.78, u)	(0.07, d)	(0.10, d)	(1.00, u)	.	(1.00, u)	(0.09, d)
M_v	(0.69, u)	(0.81, u)	(0.17, d)	(0.17, d)	(0.98, u)	(0.97, u)	.	(0.18, d)
M_p	(0.85, c)	(0.99, c)	(0.98, c)	(0.99, c)	(0.99, c)	(0.99, c)	(0.99, c)	.

Figure 15. Recovered graphs $\bar{IG}$ and $\bar{CG}$ for the diamond example.

between dynamically related variables. The confidences Θ_{wy} and Θ_{zy} are very similar to Θ_{xy} ; on the other hand, Θ_{wx} , Θ_{zx} , Θ_{wz} , and Θ_{zw} are substantially like Θ_{yx} . All other pairwise comparisons originally failed the continuity test, and they changed very little with the application of noise. We conclude that the correct reconstruction of the dynamic interactions in the chaotic oscillator example can be recovered with measurement noise in each variable up to 2–3% of the standard deviation.

6. Discussion. In this contribution we describe mathematical structures that underlie a causality detection method based on delay reconstructions from individual time series introduced by Sugihara et al. [19]. We discuss these issues both in Takens' framework of manifold reconstruction as well as in the Sauer, Yorke, and Casdagli [16] framework of attractor reconstruction. In any dynamical system we define an interaction graph that summarizes the forcing relationship between variables: if variable x occurs on the right-hand side of the evolution of variable y , then we call x a dynamical driver of y and a vertex x is connected by

Figure 16. Continuity testing for $M_x \rightarrow M_y$ and $M_y \rightarrow M_x$ with increasing noise as indicated in the legend. Θ_{xy} is more strongly preserved under perturbation.

a directed edge to a vertex y . We show that strongly connected path components of the interaction graphs are the finest units that any method based on delayed reconstruction will be able to recover. In fact the best one can hope for is a transitive closure of the component graph, i.e., the induced graph on the connected path components.

There is a slight shift of emphasis between the approaches of Sugihara et al. [19] and Takens [20]. In Takens' results the observation function is arbitrary, and the emphasis is on the fact that we can change a nongeneric observation function slightly to get a generic one. In Sugihara et al. [19], as well as in the applications of time series reconstruction in general, the observation functions are fixed experimental data series and thus cannot be changed. One hopes, based on Takens' genericity results, that these fixed observation functions belong to the generic set. However, this is an assumption which cannot be verified by perturbing the observation function slightly and checking if the results remain the same.

Given a set of time series measurements, we assume that there is an underlying system of differential equations and an associated interaction graph. Each edge in the reconstructed interaction graph represents a relationship that is potentially causal in the physical system. The driving variable may be a direct cause of the driven variable, an indirect cause, a correlate of an indirect or direct cause, or an artifact from a lack of full resolution in the dynamical system. This is why we emphasize that the relationships we find are only hypotheses for causal inference. Each specific case will require independent examination to decide if a driving relationship is also a causal relationship. Similarly, one can hypothesize which variables are not causally related, due to the absence of a continuous surjection between manifolds (or subsets of attractors). Again, this is only a hypothesis, since excessive noise could be masking an underlying relationship, or there could be symmetry in the system.

Sugihara et al. [19] noted that their method based on delay reconstructions works for moderately coupled subsystems. If the coupling is too weak, we cannot detect the effect of one system (or variable) on the other; if the coupling is too strong, then the driving system and driven system synchronize and the reconstructed attractors are identical. We formalize this

observation with the notion of a fully resolved manifold (attractor) structure, which says that the manifolds parameterized by the upper sets of component graph must be distinguishable. We show that such a fully resolved structure can be recovered by methods based on delay embeddings, up to the resolution of the transitive closure of the component graph.

We further illustrate this point in Figure 17. Graph (a) is the component graph associated with the example in Figure 4 used throughout the paper, and graph (b) is the recovered component graph when the example has a fully resolved manifold structure and Algorithm 1 works perfectly. The reconstructed component graph $\bar{CG}$ has additional connections from T to Q and P , shown as dashed lines. Using delay reconstruction techniques on a purely deterministic system, we cannot know if a variable in T directly drives a variable in Q , or whether there is only indirect dynamical driving. This is a limitation of using state space reconstructions to recover deterministic dynamical interactions.

Figure 17. (a) Example component graph CG_ϕ from Figure 4; (b) the transitive closure recovered from a perfect performance of Algorithm 1 when $\mathbb{M}$ is fully resolved; (c) the graph minor when $M_{\{TSQ\}}$ and $M_{\{TS\}}$ are homeomorphic; (d) the graph when there is a continuous surjection from $M_{\{TSQ\}}$ to $M_{\{TRSP\}}$; and (e) the graph when $M_{\{TSQ\}}$ and $M_{\{TRSP\}}$ are homeomorphic. Dashed lines indicate connections that are recovered by the algorithm that are not in CG_ϕ .

If the descending filtration $\mathbb{M}$ is not fully resolved, then the recovered graph $\bar{CG}$ may have false edges or collapsed nodes. If condition 1 of the definition fails, then one of the projections in $\mathbb{M}$ is injective and two comparable manifolds cannot be distinguished. If condition 2 fails, then two noncomparable elements of the filtration are connected, and erroneously identified if the condition fails in both directions. These three cases are given in Figure 17. In graph (c), the projection from $M_{\{TSQ\}}$ to $M_{\{TS\}}$ is injective, so that the reconstructed component graph $\bar{CG}$ is a graph minor that combines nodes S and Q . In graph (d), the contemporaneous map described in section 5 is a continuous surjection from $M_{\{TSQ\}}$ to $M_{\{TRSP\}}$, leading to a false edge between P and Q , and also between R and Q . The latter edge occurs by composing the contemporaneous maps $M_{\{TSQ\}} \rightarrow M_{\{TRSP\}} \rightarrow M_{\{R\}}$. If the inverse map exists and is continuous, then P and Q will be falsely identified as part of the same strongly connected component $P \cup Q$ as in graph (e) of Figure 17.

The failure of condition 1 in the definition of fully resolved is less severe than the failure of condition 2, since the former results in a graph minor of the transitive closure of CG , while

the latter creates spurious edges not supported by direct or indirect driving relationships. A fully resolved filtration is crucial for the maximal recovery of the true dynamics.

To recover the information about dynamical drivers we use the method of Pecora, Carroll, and Heagy [15], which utilizes a hypothesis-testing framework to detect continuous maps. This is different from the convergent cross-mapping technique proposed in Sugihara et al. [19], which emphasizes the predictability of one time series from another. One key feature of the convergent cross-mapping technique is a projection back to a scalar time series after making the state space reconstructions. We believe that it is more consistent with the theory to test relationships between the state space reconstructions directly, as the method of Pecora, Carroll, and Heagy [15] does. However, we concur with Sugihara et al. [19] that the critical factor supporting the existence of a driving interaction is convergence to an appropriate asymptotic value with increasing time series length, rather than only depending on the results for a single time series length.

Detection of causal relationships between variables represented by experimental time series is an extremely important challenge for present day science. As the amount of data that we are able to collect expands rapidly, our ability to analyze that data is lagging behind. Our contribution, which analyzes the advantages and limitations of a class of methods based on delay reconstructions, hopes to narrow the gap between experimental successes and our ability to draw conclusions from the data.

Appendix A. Proofs for section 3. We wish to apply the theorems of Takens to the diffeomorphisms or flows that admit a fixed component graph CG. In order to do this, we wish to establish the genericity of A1 and A2 for diffeomorphisms and B1 and B2 for flows that are restricted to a fixed CG. We consider only the second case here. We first offer a sketch of the proof of the genericity of conditions B1 and B2 for flows and time lags in $C^2(M, M) \times \mathbb{R}^+$ (Theorem A.1). The proof is very similar to the proof for the genericity of conditions A1 and A2 in the space of diffeomorphisms $D^2(M, M)$. Assuming this latter fact to be true, we then offer the sketch of the proof of Theorem 3.3, which depends on Takens' Theorem 3.1. Afterward, we proceed to the novel result that conditions B1 and B2 are generic in the product space $\mathcal{P}_{\text{CG}} \times \mathbb{R}^+$, referring back to Definition 3.5.

In the rest of the appendix, we have found it convenient to distinguish flows from diffeomorphisms using the notation ψ for flows and ϕ for diffeomorphisms. Thus a flow ψ^t induces a diffeomorphism ϕ^T for a choice of time lag T .

Theorem A.1. *Let X be a C^2 vector field with associated flow ψ^t on a compact manifold M . For $T \in \mathbb{R}^+$, define the diffeomorphism ϕ^T as above. Then for fixed $n \in \mathbb{N}$, a subset of elements (ψ^t, T) that satisfies the following properties is open and dense in $C^2(M, M) \times \mathbb{R}^+$:*

B1 ψ^t has no periodic orbits of period kT for $k \leq n$.

B2 If $X(x) = 0$, then the eigenvalues of the derivative of ϕ^T evaluated at x are distinct and the moduli are all different from 1.

Proof. The genericity of B1 and B2 follows from the Kupka–Smale theorem for flows [14]. We will briefly sketch the proof here.

Genericity of B1. The Kupka–Smale theorem for flows [14] states that for a residual set of flows in $C^2(M, M)$, all of the fixed points and periodic orbits of ψ^t are hyperbolic. The Hartman–Grobman theorem for flows [14] states that there exists an open neighborhood U_ℓ of a

hyperbolic fixed point or periodic orbit such that the flow in the neighborhood is topologically equivalent to the linearization of the flow at the invariant set. This implies that in each U_ℓ there is a unique invariant set; i.e., the fixed points and periodic orbits of ψ^t are isolated. To see that this implies finiteness of the set of fixed points and periodic orbits on a compact manifold M , consider open sets $V_\ell \subset U_\ell$ such that $Cl(V_\ell) \subset U_\ell$ and the invariant set is fully contained in V_ℓ . This is possible since the manifold is locally homeomorphic to $\mathbb{R}^n$. Then consider the collection of open sets $\{U_\ell\}_\ell \cup (M \setminus Cl(\bigcup_\ell V_\ell))$. Since V_ℓ are disjoint, $Cl(\bigcup_\ell V_\ell) \subset \bigcup_\ell U_\ell$ and therefore $M \setminus (\bigcup_\ell U_\ell) \subset (M \setminus Cl(\bigcup_\ell V_\ell))$. Therefore this collection is an open cover of a compact manifold M and thus admits a finite subcover. Since every set in the cover contains at most one fixed point or periodic orbit, the number of equilibria and periodic orbits is finite. We have just shown that on a compact manifold the set of flows with a finite number of hyperbolic equilibria and periodic orbits is residual and therefore dense. It is clear that this set is also open in $C^2(M, M)$. We denote this set by G .

For a given $\psi^t \in G$, let $N < \infty$ be the number of hyperbolic periodic orbits and let $T_1, T_2, \dots, T_N$ denote the minimal periods of each. Fix $n \in \mathbb{N}$ and consider the sets $S_i = \{T_i/k\}_{k=1}^n$. Each set of real positive numbers S_i is finite, so $\mathbb{R}^+ \setminus \bigcup_{i=1}^N S_i$ is open and dense in $\mathbb{R}^+$. Thus a generic choice of $T \in \mathbb{R}^+$, which depends on ψ^t , ensures that $\{kT\}_{k=1}^n \cap \{T_1, T_2, \dots, T_N\} = \emptyset$, fulfilling B1. Denote each of these sets by $H_{\psi^t} = \mathbb{R}^+ \setminus \bigcup_{i=1}^N S_i$.

We know that B1 is fulfilled on the set

$$\mathbb{G} := \bigcup_{\psi^t \in G} \{\psi^t \times H_{\psi^t}\},$$

which we wish to show is open and dense in $C^2(M, M) \times \mathbb{R}^+$.

To show the density of $\mathbb{G}$, assume that $(\psi^t, T) \notin \mathbb{G}$. Because G is dense in $C^2(M, M)$ and $C^2(M, M)$ is a metric space, there exists a sequence $\{\psi_i^t\} \in G$ with ψ^t as its limit point. Then because each $H_{\psi_i^t}$ is dense in $\mathbb{R}^+$, we can choose a sequence $\{T_i\}$ with each $T_i \in H_{\psi_i^t}$ such that $|T_i - T| < 1/i$. Then $\{(\psi_i^t, T_i)\} \rightarrow (\psi^t, T)$ shows that every point in $C^2(M, M) \times \mathbb{R}^+$ is the limit point of a sequence in $\mathbb{G}$.

To show openness, let $(\psi^t, T) \in \mathbb{G}$ and let $T_1(\psi^t), T_2(\psi^t), \dots, T_N(\psi^t)$ be the periods of the hyperbolic orbits of ψ^t . Because periodic behavior continuously varies in $C^2(M, M)$, for any $\delta > 0$ there exists an open neighborhood $U_\delta \subset G$ of ψ^t where $\alpha^t \in U_\delta$ has the same number of periodic orbits as ψ^t and $|T_i(\psi^t) - T_i(\alpha^t)| < \delta$ for all $i = 1, \dots, N$. Choose $\delta > 0$ small enough so that $\min_{i,k} |T - T_i(\psi^t)/k| > 2\delta$ with $i = 1, \dots, N$ and $k = 1, \dots, n$. This must be possible since $T \in H_{\psi^t}$. Now consider a point (α^t, τ) in the open set $U_\delta \times (T - \delta/2, T + \delta/2)$. We show that $\tau \in H_{\alpha^t}$ by the following argument:

$$\begin{aligned} |\tau - T_i(\alpha^t)/k| &\geq |T - T_i(\psi^t)/k| - |T - \tau| - |T_i(\psi^t)/k - T_i(\alpha^t)/k| \\ &> 3\delta/2 - \delta/k \\ &\geq \delta/2. \end{aligned}$$

Thus, $U_\delta \times (T - \delta/2, T + \delta/2) \subset \mathbb{G}$, so that $\mathbb{G}$ is open in $C^2(M, M) \times \mathbb{R}^+$.

Genericity of B2. By the Kupka–Smale theorem [14] for an open and dense set of C^2 vector fields on M the singularities $X(x) = 0$ are hyperbolic, and so $DX(x)$ at each singularity has

no eigenvalues on the imaginary axis. The linearized flow in a small neighborhood about x is given by $\psi^t(y) = e^{tDX(x)}(y - x)$, so that the derivative of the diffeomorphism ϕ^T at x is $e^{TDX(x)}$. Since $T \in \mathbb{R}$, the operator $TDX(x)$ has no eigenvalues on the imaginary axis, implying that $e^{TDX(x)}$ has no eigenvalues on the unit circle.

We now need to prove that the eigenvalues of $e^{TDX(x)}$ are distinct. The proof requires three steps. The first is an observation that there is an open and dense set $\mathcal{A}$ in the set of all linear operators such that $A \in \mathcal{A}$ satisfies two conditions:

1. If λ_1 is an eigenvalue of A , then $\operatorname{Re}(\lambda_1) \neq 0$.
2. For any pair of distinct eigenvalues λ_1, λ_2 of A , then $e^{T\lambda_1} \neq e^{T\lambda_2}$.

It follows that at each singularity x , there is an arbitrarily small perturbation that takes $DX(x)$ into a linear operator in $\mathcal{A}$. The second step is to show that the perturbation can be localized around the singularity. While a construction using a bump function is standard (see, e.g., proofs in [14] and [8]), the fact that we need the perturbed function to be C^2 close to the original presents technical issues. We refer readers to the proof of Lemma A.2 for an idea of the method of perturbation in this proof. The third step uses the compactness of M to note that there are finitely many hyperbolic singularities for vector field X , and hence there is a minimal distance between them. By using nonoverlapping bump functions one can show that the perturbation at each singularity can be done simultaneously. This shows the density part of B2.

Openness follows from the openness of $\mathcal{A}$ in the set of all linear operators. Let X be a vector field with $X(x) = 0$ such that $DX(x) \in \mathcal{A}$, and let $\delta > 0$ be such that $B \in \mathcal{A}$ when B is within δ of $DX(x)$. Then, for any smooth Y with $Y(x) = 0$ sufficiently close to X , the difference in derivatives $DX(x) - DY(x)$ is smaller than δ . ■

Proof of Theorem 3.3. We will show that if conditions B1 and B2 are fulfilled for ψ^t and T , then the diffeomorphism $\phi^T : M \rightarrow M$ will fulfill properties A1 and A2, making $\Phi_{(\phi^T, \varphi)} : M \rightarrow \mathbb{R}^{2m+1}$ generically an embedding by Takens' Theorem 3.1. Since condition B1 implies no periodic points of ψ^t with periods kT for $1 \leq k \leq n$, $n \in \mathbb{N}$ fixed, then the only closed orbits in the set $\{\phi^T(x), (\phi^T)^2(x), \dots, (\phi^T)^n(x)\}$ are fixed points of period 1 for any $x \in M$. Since generically there are only a finite number of singularities of ψ^t , there can be only a finite number of fixed points of ϕ^T . So B1 implies A1. Since ϕ^T has only fixed points (i.e., periodic points of period 1), then A2 requires only that the derivative of ϕ^T at the fixed points have distinct eigenvalues not on the unit circle. This is exactly what was proved in condition B2. ■

Now we extend these results to the smaller set of flows with a fixed component graph CG. Given CG and a compact manifold M , we build the corresponding upper set lattice, UL , and manifold filtration $\mathbb{M}$. For any two upper sets satisfying $U < V$, denote the dimensions of the compact manifolds $M_U, M_V \in \mathbb{M}$ by n_U, n_V , respectively, and their associated flows by ψ_U^t, ψ_V^t . Recall that since $U < V$, there is an induced projection π_{VU} such that $\pi_{VU}(M_V) = M_U$ and $\pi_{VU} \circ \psi_V^t = \psi_U^t \circ \pi_{VU}$. Recall also Definition 3.5, where we define the subset $\mathcal{P}_{CG} \subset C^2(M, M)$ to consist of those flows where $\pi_{VU} \circ \psi_V^t = \psi_U^t \circ \pi_{VU}$ holds for any pair of upper sets in UL satisfying $U < V$ for a fixed pair (CG, M) . Our goal is to prove that conditions B1 and B2 are generic in this smaller set $\mathcal{P}_{CG} \times \mathbb{R}^+$.

To prove this we first consider the simplest possible upper set lattice UL which consists of two sets U, V with $U < V$. Arguing openness is straightforward, but arguing density requires

substantial work. We will use this result and induction on UL to show genericity of B1 and B2 in $\mathcal{P}_{CG}$.

Theorem A.2. *Consider a component graph CG with only two strongly connected components, so that the upper set lattice is composed of sets U and V with $\emptyset < U < V$. Then the set of $(\psi_V^t, T) \in \mathcal{P}_{CG} \times \mathbb{R}^+$ where ψ_V^t has no periodic orbits of period jT with $j \in \{1, 2, \dots, 2n_V + 1\}$ (property B1) is open and dense in $\mathcal{P}_{CG} \times \mathbb{R}^+$.*

Before we prove Theorem A.2 we need a lemma which will allow us to transfer the relationship between U and V locally to coordinate charts. Let (c_i, Q_i) be an atlas mapping M_V locally to $\mathbb{R}^{n_V}$ and (d_i, P_i) an atlas from M_U to $\mathbb{R}^{n_U}$. Given an arbitrary point $v \in M_V$, consider an open neighborhood $O \subset M_V$ small enough so that the following hold:

1. $O \subset Q_j$ and $\psi_V^t(O) \subset Q_j$ for a fixed index j and a small time t .
2. $\pi_{VU}(O) \subset P_j$ and $\psi_U^t(\pi_{VU}(O)) \subset P_j$ for a fixed index which we may take to be j without loss of generality, and a small time t .

There are three locally induced maps to consider:

$$\begin{aligned}\tilde{\psi}_V^t &\equiv c_j \circ \psi_V^t \circ c_j^{-1} \in C^2(\mathbb{R}^{n_V}), \\ \tilde{\psi}_U^t &\equiv d_j \circ \psi_U^t \circ d_j^{-1} \in C^2(\mathbb{R}^{n_U}), \\ \tilde{\pi}_j &\equiv d_j \circ \pi_{VU} \circ c_j^{-1} \in C^\infty(\mathbb{R}^{n_V}, \mathbb{R}^{n_U}).\end{aligned}$$

The induced projection-like function is C^∞ because we assume that the manifolds M_U and M_V are C^∞ , which then implies that π_{VU} is C^∞ .

Lemma A.3. *Let (c_i, Q_i) be an atlas mapping M_V locally to $\mathbb{R}^{n_V}$ and (d_i, P_i) an atlas from M_U to $\mathbb{R}^{n_U}$, as above. Then*

$$\tilde{\pi}_j \circ \tilde{\psi}_V^t = \tilde{\psi}_U^t \circ \tilde{\pi}_j.$$

Proof. We note that

$$\begin{aligned}\tilde{\pi}_j \circ c_j &= d_j \circ \pi_{VU}, \\ \pi_{VU} \circ c_j^{-1} &= d_j^{-1} \circ \tilde{\pi}_j.\end{aligned}$$

Then

$$\begin{aligned}\tilde{\pi}_j \circ \tilde{\psi}_V^t &= \tilde{\pi}_j \circ c_j \circ \psi_V^t \circ c_j^{-1}, \\ &= d_j \circ \pi_{VU} \circ \psi_V^t \circ c_j^{-1}, \\ &= d_j \circ \psi_U^t \circ \pi_{VU} \circ c_j^{-1}, \\ &= d_j \circ \psi_U^t \circ d_j^{-1} \circ \tilde{\pi}_j, \\ &= \tilde{\psi}_U^t \circ \tilde{\pi}_j.\end{aligned}$$

Therefore the projection relationship holds locally near any $c_i(v)$ for $v \in M_V$. ■

Proof of Theorem A.2. Given T , we denote periodic orbits of period jT with $j \in \{1, 2, \dots, 2n_V + 1\}$ as *bad* periodic orbits.

Open. The set of $(\psi_V^t, T) \in C^2(M_V, M_V) \times \mathbb{R}^+$ fulfilling condition B1 (i.e., there are no bad periodic orbits) is open; therefore it is relatively open in $\mathcal{P}_{CG} \times \mathbb{R}^+$.

Dense. The set of $(\psi_U^t, T) \in C^2(M_U, M_U) \times \mathbb{R}^+$ fulfilling condition B1 is open and dense; i.e., there are no bad periodic orbits for pairs (ψ_U^t, T) in the generic set. Given a generic choice of (ψ_U^t, T) , the key is to show that the set of ψ_V^t with no bad periodic orbits in T is dense in the set of ψ_V^t satisfying $\pi_{VU} \circ \psi_V^t = \psi_U^t \circ \pi_{VU}$. To do this, we assume that there exists a bad periodic orbit in ψ_V^t and construct an arbitrarily close perturbation in $\mathcal{P}_{CG}$ with no bad periodic orbits that respects the projection relationship.

We observe that if ψ_V^t has a bad periodic orbit, it must occur at an equilibrium point of ψ_U^t , given a generic choice of ψ_U^t and T . Since there are a finite number of equilibrium points of ψ_U^t , we need to make at most a finite number of perturbations to ψ_V^t to ensure that there are no bad periodic orbits of ψ_V^t with the same T .

Consider a single equilibrium point of ψ_U^t giving rise to a bad periodic orbit of ψ_V^t . We denote that equilibrium point by u_e . We will show that for any $\epsilon > 0$, there exists a flow $\xi_V^t \in \mathcal{P}_{CG}$ and an open neighborhood of u_e , $P \subset M_U$, such that

1. ξ_V^t has no bad periodic orbits in the set $S = \{v : v \in M_V, \pi_{VU}(v) \in P\}$;
2. ξ_V^t is a local perturbation of ψ_V^t ; i.e., $\xi_V^t = \psi_V^t$ for all $v \in M_V \setminus S$;
3. ξ_V^t can be arbitrarily close to ψ_V^t ; i.e., $\|\xi_V^t - \psi_V^t\| < \epsilon$ on S in the C^2 norm.

Using the result and notation of Lemma A.3, we may prove 1–3 for the flows $\tilde{\psi}_V^t \in C^2(\mathbb{R}^{n_V}, \mathbb{R}^{n_V})$ and $\tilde{\psi}_U^t \in C^2(\mathbb{R}^{n_U}, \mathbb{R}^{n_U})$, and the result can be passed through the charts back to the original flows.

If the C^∞ projection $\pi_{VU} : M_V \rightarrow M_U$ is restricted to the domain M_V , then the set $N := \pi_{VU}^{-1}(u_e)$ is generically a compact submanifold of M_V , because u_e is generically a regular value of π_{VU} by the Morse–Sard theorem [8]. Clearly N has an atlas (c_i, O_i) with $O_i = Q_i \cap N$, consistent with the atlas (c_i, Q_i) of M_V , and is locally diffeomorphic to $\mathbb{R}^{n_V - n_U}$ (see [8, pp. 13–14]). The flow ψ_V^t induces a flow $\psi_V^t|_{u_e}$ on the manifold N , which exhibits the bad periodic orbit.

By the genericity of B1 in the set C^2 flows on the compact manifold N , we may choose a C^2 flow $\kappa^t : N \rightarrow N$ with no bad periodic orbits such that $\|\kappa^t - \psi_V^t|_{u_e}\| < \mu$ in the C^2 norm, where μ is sufficiently small (exact condition to be determined later). Let $\tilde{\kappa}^t = c_j \circ \kappa^t \circ c_j^{-1}$ be the locally induced flow. The condition $\|\kappa^t - \psi_V^t|_{u_e}\| < \mu$ in the C^2 norm means that $\|\tilde{\kappa}^t - \tilde{\psi}_V^t|_{u_e}\| < \mu$ in the C^2 norm on $\mathbb{R}^{n_V - n_U}$ for every open set O_i in the atlas of N . We shall assume for the remainder of the proof that the mapped flows $(\tilde{\psi}, \tilde{\kappa})$ are C^∞ functions on their respective domains, since C^2 functions may be approximated arbitrarily closely with C^∞ functions.

There exists a small neighborhood $N_U \subset P_i \subset M_U$ containing u_e that contains no other limit sets of ψ_U^t , where P_i is an open set in the atlas (d_i, P_i) of M_U . Without loss of generality, assume that $d_j(u_e) = 0 \in \mathbb{R}^{n_U}$; that is, the equilibrium point of interest is mapped to the origin. We shall choose $\delta \ll 1$ sufficiently small that the δ -ball in $\mathbb{R}^{n_U}$ satisfies $d_j^{-1}(\overline{B_\delta(0)}) \subset N_U$. Further conditions on the smallness of δ will be imposed later in the proof.

For each $v \in N$, there exists a small neighborhood $N_v \subset O_i$ for some O_i in the atlas of N . There is a compact set C_v such that $v \in C_v \subset N_v$. Let $\overline{B_v} = c_i(C_v) \in \mathbb{R}^{n_V - n_U}$.

By Lemma A.3, we may write $\tilde{\psi}_U^t \equiv g^t(x)$ and $\tilde{\psi}_V^t \equiv (g^t(x), f^t(x, y))$ for $x \in \overline{B_\delta(0)} \subset \mathbb{R}^{n_U}$, $y \in \overline{B_v} \subset \mathbb{R}^{n_V - n_U}$. Then $f^t(0, y)$ and $\tilde{\kappa}^t(0, y)$ have the domain $\{0\} \times \overline{B_v}$, where the input y for both functions is generated from the same atlas. We emphasize that f^t and $\tilde{\kappa}^t$ are

C^∞ on compact domains, which will be important when we take derivatives of Taylor series later on.

We extend $\tilde{\kappa}^t$ to a neighborhood of $0 \in \mathbb{R}^{n_V}$ by adding the Taylor expansion of f^t and concatenating with g^t on $\mathbb{R}^{n_U}$:

$$\begin{aligned} k^t(x, y) &= (g^t(x), \tilde{\kappa}^t(0, y)) + \left(0, x^T D_x f^t(0, y) + \frac{1}{2} x^T D_{xx} f^t(0, y) x \right) \\ &= \left(g^t(x), \tilde{\kappa}^t(y) + \sum_{i=1}^{n_U} x_i \frac{\partial f^t}{\partial x_i}(0, y) + \sum_{i=1}^{n_U} \sum_{j=1}^{n_U} \frac{x_i x_j}{2} \frac{\partial^2 f^t}{\partial x_i \partial x_j}(0, y) \right), \end{aligned}$$

where D_x is the $n_U \times n_V$ Jacobian of f^t with respect to the vector x . Likewise D_{xx} is the $n_U \times n_V \times n_U$ Hessian of f^t .

There exists a bump function $h \in C^\infty(\mathbb{R}^{n_U}, [0, 1])$ such that $h(x) = 1$ on the closed ball $\overline{B_{\delta/2}(0)}$ and $h(x) = 0$ on the closed set $\mathbb{R}^{n_U} \setminus B_\delta(0)$. We are now ready to construct a function that (locally in $\mathbb{R}^{n_V}$) fulfills properties 1–3. Let

$$\tilde{\xi}_V^t(x, y) = (1 - h(x))\tilde{\psi}_V^t(x, y) + h(x)k^t(x, y).$$

Properties 1 and 2 are satisfied by construction. Since $\tilde{\xi}_V^t(0, y) = (0, \tilde{\kappa}^t(0, y))$ and there are no other equilibrium points of $\tilde{\psi}_U^t$ in the δ -neighborhood of 0, $\tilde{\xi}_V^t(x, y)$ has no bad periodic orbits for $x \in B_\delta(0)$. And clearly $\tilde{\xi}_V^t(x, y) = \tilde{\psi}_V^t(x, y)$ for any x outside of the δ -ball. Notice also that $\tilde{\xi}_V^t$ is C^2 and respects the projection relationship $\tilde{\pi} \circ \tilde{\xi}_V^t = \tilde{\psi}_U^t \circ \tilde{\pi}$.

The work is in proving property 3. Consider

$$\begin{aligned} \tilde{\xi}_V^t(x, y) - \tilde{\psi}_V^t(x, y) &= h(x) \left(0, \tilde{\kappa}^t(0, y) + x^T D_x f^t(0, y) + \frac{1}{2} x^T D_{xx} f^t(0, y) x - f^t(x, y) \right) \\ &= h(x) \left(0, \tilde{\kappa}^t(0, y) - f^t(0, y) \right) \\ &\quad + h(x) \left(0, f^t(0, y) + x^T D_x f^t(0, y) + \frac{1}{2} x^T D_{xx} f^t(0, y) x - f^t(x, y) \right) \\ (10) \quad &:= h(x) \left(0, \tilde{\kappa}^t(0, y) - f^t(0, y) \right) + h(x) \left(0, T(x, y) \right). \end{aligned}$$

The idea is that the first term is small by the choice of $\tilde{\kappa}^t$, and the second term is small because it is the difference of a function and its Taylor series. However, there are technical details that come from using the C^2 norm, $\|f\| = \sum_{|\alpha| \leq 2} \sup_{x \in B} |f^\alpha(x)|$, where each α is a multi-index. We will abuse notation slightly and allow $\|\cdot\|$ to represent the C^2 norm over $\mathbb{R}^p$ or M for various Euclidean spaces or manifolds as indicated by context.

The first term in (10) is easier to handle, but it imposes a constraint on μ in terms of δ , and δ has independent constraints arising from the second term. Therefore, we address the second term before the first.

First we note that

$$\sum_{i=1}^{n_U} \sup_{x \in B_\delta(0)} \left| \frac{\partial h(x)}{\partial x_i} \right| = \frac{M_1}{\delta}$$

for some constant M_1 . This is because any bump function $b(x)$ with support on the unit ball $\overline{B_1(0)}$ will have the sum of the magnitudes of the first derivatives bounded by some number M_1 , and $h(x)$ can be constructed by $h(x) := b(x/\delta)$. Likewise, the second derivative of $h(x)$ is bounded by M_2/δ^2 .

For the second term in (10), we consider each part of the C^2 norm separately.

$$\begin{aligned}
& \sup_{x \in \overline{B_\delta(0)}, y \in \overline{B_v}} |h(x)T(x, y)| \\
& \leq \sup_{x \in \overline{B_\delta(0)}} |h(x)| \sup_{x \in \overline{B_\delta(0)}, y \in \overline{B_v}} \left| f^t(0, y) + x^T D_x f^t(0, y) + \frac{1}{2} x^T D_{xx} f^t(0, y) x - f^t(x, y) \right| \\
& \leq \sup_{x \in \overline{B_\delta(0)}, y \in \overline{B_v}} \left| \sum_{|\alpha|=3} \frac{x^\alpha}{|\alpha|!} \frac{\partial^{|\alpha|} f^t}{\partial x^\alpha}(x, y) \right| \delta^3 \\
& \leq M_3 \delta^3
\end{aligned}$$

by Taylor's theorem over x in the δ -ball. M_3 is the result of taking the supremum of the Taylor coefficients over the compact set $\overline{B_v}$. From now on we will skip the explicit step of taking the supremum of the Taylor coefficients, but it occurs implicitly at every invocation of Taylor's theorem below. Additionally, when we take the sup below, it will occur over the set $\{x \in \overline{B_\delta(0)}, y \in \overline{B_v}\}$, but we will omit the set in our notation to improve readability.

$$\begin{aligned}
& \sum_{i=1}^{n_U} \sup \left| \frac{\partial h}{\partial x_i} T(x, y) + h(x) \frac{\partial T}{\partial x_i} \right| + \sum_{k=1}^{n_V - n_U} \sup \left| h(x) \frac{\partial T}{\partial y_k} \right| \\
& \leq \sum_{i=1}^{n_U} \sup \left| \frac{\partial h}{\partial x_i} \right| |T(x, y)| + \sum_{i=1}^{n_U} \sup \left| \frac{\partial T}{\partial x_i} \right| + \sum_{k=1}^{n_V - n_U} \sup \left| \frac{\partial T}{\partial y_k} \right| \\
& \leq \frac{M_1}{\delta} M_3 \delta^3 + \sum_{i=1}^{n_U} \sup \left| \frac{\partial f^t}{\partial x_i}(0, y) + \sum_{j=1}^{n_U} x_j \frac{\partial^2 f^t}{\partial x_i \partial x_j}(0, y) - \frac{\partial f^t}{\partial x_i}(x, y) \right| \\
& \quad + \sum_{k=1}^{n_V - n_U} \sup \left| \frac{\partial f^t}{\partial y_k}(0, y) + \sum_{i=1}^{n_U} x_i \frac{\partial^2 f^t}{\partial y_k \partial x_i}(0, y) \right. \\
& \quad \quad \left. + \sum_{i=1}^{n_U} \sum_{j=1}^{n_U} \frac{x_i x_j}{2} \frac{\partial^3 f^t}{\partial y_k \partial x_i \partial x_j}(0, y) - \frac{\partial f^t}{\partial y_k}(x, y) \right| \\
& \leq M_1 M_3 \delta^2 + M_4 \delta^2 + M_5 \delta^3,
\end{aligned}$$

where the last two estimates come from Taylor's theorem on $\frac{\partial f^t}{\partial x_i}$ and $\frac{\partial f^t}{\partial y_k}$.

The second derivative estimate is as follows:

$$\begin{aligned}
& \sum_{j=1}^{n_U} \sum_{i=1}^{n_U} \sup \left| \frac{\partial^2 h}{\partial x_j \partial x_i} T(x, y) + 2 \frac{\partial h}{\partial x_i} \frac{\partial T}{\partial x_j} + h(x) \frac{\partial^2 T}{\partial x_j \partial x_i} \right| + \sum_{j=1}^{n_U} \sum_{k=1}^{n_V-n_U} \sup \left| \frac{\partial h}{\partial x_j} \frac{\partial T}{\partial y_k} + h(x) \frac{\partial^2 T}{\partial x_j \partial y_k} \right| \\
& + \sum_{\ell=1}^{n_V-n_U} \sum_{i=1}^{n_U} \sup \left| \frac{\partial h}{\partial x_i} \frac{\partial T}{\partial y_\ell} + h(x) \frac{\partial^2 T}{\partial y_\ell \partial x_i} \right| + \sum_{\ell=1}^{n_V-n_U} \sum_{k=1}^{n_V-n_U} \sup \left| h(x) \frac{\partial^2 T}{\partial y_\ell \partial y_k} \right| \\
& \leq M_2 M_3 \delta + 2 M_1 M_4 \delta + 2 M_1 M_5 \delta^2 + \sum_{j=1}^{n_U} \sum_{i=1}^{n_U} \sup \left| \frac{\partial^2 T}{\partial x_j \partial x_i} \right| \\
& + 2 \sum_{j=1}^{n_U} \sum_{k=1}^{n_V-n_U} \sup \left| \frac{\partial^2 T}{\partial x_j \partial y_k} \right| + \sum_{\ell=1}^{n_V-n_U} \sum_{k=1}^{n_V-n_U} \sup \left| \frac{\partial^2 T}{\partial y_\ell \partial y_k} \right|,
\end{aligned}$$

where

$$\sum_{j=1}^{n_U} \sum_{i=1}^{n_U} \sup \left| \frac{\partial^2 T}{\partial x_j \partial x_i} \right| = \sum_{j=1}^{n_U} \sum_{i=1}^{n_U} \sup \left| \frac{\partial^2 f^t}{\partial x_i \partial x_j}(0, y) - \frac{\partial^2 f^t}{\partial x_j \partial x_i}(x, y) \right| \leq M_6 \delta,$$

$$\begin{aligned}
& \sum_{j=1}^{n_U} \sum_{k=1}^{n_V-n_U} \sup \left| \frac{\partial^2 T}{\partial x_j \partial y_k} \right| \\
& = \sum_{j=1}^{n_U} \sum_{k=1}^{n_V-n_U} \sup \left| \frac{\partial^2 f^t}{\partial x_j \partial y_k}(0, y) + \sum_{i=1}^{n_U} x_i \frac{\partial^3 f^t}{\partial y_k \partial x_i \partial x_j}(0, y) - \frac{\partial^2 f^t}{\partial x_j \partial y_k}(x, y) \right| \\
& \leq M_7 \delta^2,
\end{aligned}$$

$$\begin{aligned}
& \sum_{\ell=1}^{n_V-n_U} \sum_{k=1}^{n_V-n_U} \sup \left| \frac{\partial^2 T}{\partial y_\ell \partial y_k} \right| \\
& = \sum_{\ell=1}^{n_V-n_U} \sum_{k=1}^{n_V-n_U} \sup \left| \frac{\partial^2 f^t}{\partial y_\ell \partial y_k}(0, y) + \sum_{i=1}^{n_U} x_i \frac{\partial^3 f^t}{\partial y_\ell \partial y_k \partial x_i}(0, y) \right. \\
& \quad \left. + \sum_{i=1}^{n_U} \sum_{j=1}^{n_U} \frac{x_i x_j}{2} \frac{\partial^4 f^t}{\partial y_\ell \partial y_k \partial x_i \partial x_j}(0, y) - \frac{\partial^2 f^t}{\partial y_\ell \partial y_k}(x, y) \right| \\
& \leq M_8 \delta^3.
\end{aligned}$$

Putting everything together,

$$\begin{aligned}
||h(x)T(x, y)|| & \leq M_3 \delta^3 + M_1 M_3 \delta^2 + M_4 \delta^2 + M_5 \delta^3 + M_2 M_3 \delta \\
& + 2 M_1 M_4 \delta + 2 M_1 M_5 \delta^2 + M_6 \delta + 2 M_7 \delta^2 + M_8 \delta^3 \\
& = \delta(M_2 M_3 + 2 M_1 M_4 + M_6) + \delta^2(M_1 M_3 + M_4 + 2 M_1 M_5 + 2 M_7) \\
& + \delta^3(M_3 + M_5 + M_8).
\end{aligned}$$

The parameter δ must be chosen small enough so that the term above is less than $\epsilon/2$.

The first term reduces to

$$\begin{aligned} \|h(x) (\tilde{\kappa}^t(0, y) - f^t(0, y))\| &\leq \|h(x)\| \|\tilde{\kappa}^t(0, y) - f^t(0, y)\| \\ &\leq \left(1 + \frac{M_1}{\delta} + \frac{M_2}{\delta^2}\right) \mu \\ &< \epsilon/2, \end{aligned}$$

where we enforce the last inequality by taking $\mu \ll \epsilon$ and $\mu \ll \delta$. This completes the proof of property 3.

The local perturbation $\tilde{\xi}^t(x, y) \in C^\infty(\mathbb{R}^{n_V}, \mathbb{R}^{n_V})$ may be mapped back to the manifold via $\xi_V^t = c_j^{-1} \circ \tilde{\xi}^t \circ c_j$, again using the language of Lemma A.3. This process was done for an arbitrary $v \in N$, so it holds for all points $p \in S = \{p : p \in M_V, \pi_{VU}(p) \in P\}$, where $P = d_j^{-1}(B_\delta(0))$. Then properties 1 and 3 hold in S , and no changes were introduced outside of this set, so that property 2 holds in $M_V \setminus S$.

Now suppose ψ_V^t has more than one bad periodic orbit (i.e., ψ_U^t has more than one bad equilibrium point), noting that there may be only a finite number of bad periodic orbits. Since the flows ξ_V^t are local perturbations on disjoint sets, a finite number of them may summed together into a single flow satisfying conditions 1–3. ■

Theorem A.4. *Consider a component graph CG with only two strongly connected components, so that the upper set lattice is composed of sets U and V with $\emptyset < U < V$. Then a subset of elements (ψ_V^t, T) that satisfies property B2 (at an equilibrium point, the eigenvalues of the Jacobian are distinct with moduli different from 1) is open and dense in $\mathcal{P}_{CG} \times \mathbb{R}^+$.*

Proof.

Open. As before, the set of (ψ_V^t, T) fulfilling condition B2 is open in the set of $C^2(M_V, M_V) \times \mathbb{R}^+$. Therefore it is relatively open in the set $\mathcal{P}_{CG} \times \mathbb{R}^+$.

Dense. We outline the proof of density, which combines Theorem A.1 and the ideas of the proof of Theorem A.2. By Theorem A.1 there is a generic choice of $\psi_U^t \in C^2(M_U, M_U)$ which satisfies B2. In particular, ψ_U^t has finitely many hyperbolic equilibria which have distinct eigenvalues. Take now an equilibrium z of a flow $\psi_V^t \in \mathcal{P}_{CG}$. Passing to local coordinates as in the previous proof, the flow in local coordinates $\tilde{\psi}_V^t$ can be written as

$$\tilde{\psi}_V^t = (g^t(x), f^t(x, y)).$$

It follows that the equilibrium $\tilde{z} = (\tilde{x}, \tilde{y})$ is also an equilibrium of $\tilde{\psi}_U^t$. The key observation is that the linearization A at $\tilde{z}$ has a lower triangular structure

$$A = \begin{bmatrix} B & 0 \\ D & C \end{bmatrix},$$

where $B = Dg^t(\tilde{x})$ and $C = \frac{Df^t}{Dy}(\tilde{x}, \tilde{y})$. We note that the eigenvalues of A are the union of eigenvalues of B and eigenvalues of C . As in the proof of Theorem A.2, we can perturb C in the submanifold N , while keeping B unchanged. Since N is generically a manifold of dimension $n_V - n_U$, a small perturbation of C will guarantee hyperbolicity and distinct eigenvalues for C as well as the entire matrix A . Further, using the same methods as in Theorem A.2, we can

localize this perturbation to function $\tilde{f}^t$ to a neighborhood of $\tilde{x}$. Since there are only finitely many equilibria $\tilde{z}$, these local perturbations can be composed to arrive at perturbation of the flow ψ_V^t that satisfies B2. ■

We are almost ready to prove the main result. In the proof we will require the following definition.

Definition A.5. For a given upper set lattice UL , the least elements of UL are the nonempty upper sets $U_1, \dots, U_n$ such that for every $V \in UL$, the meet (greatest lower bound) of V and U_i , $V \vee U_i$, is either U_i or $\emptyset$. That is, each U_i consists of a single strongly connected component in the component graph CG with only outgoing edges.

Theorem A.6. Consider a flow ψ^t on a smooth, compact manifold M with the corresponding component graph CG . Then conditions B1 and B2 hold on an open and dense set of $(\psi^t, T) \in \mathcal{P}_{CG} \times \mathbb{R}^+$.

Proof. The upper set lattice UL associated with CG may be divided into levels L_i , where the empty set is the only element in L_0 , the least elements of UL form the level L_1 , and the top level, composed of a single element, is L_N . We present an induction proof on the levels of UL .

Let m be the dimension of the manifold M . For each least element U_i in the level L_1 , $(\psi_{U_i}^t, T_i)$ may be chosen from $C^2(M_U, M_U) \times \mathbb{R}^+$ to generically satisfy B1 and B2. Specifically for B1, periodic orbits of period kT_i do not exist for $k \leq 2m + 1$, a number depending on the dimension of M , not on M_U . Since the number of sets in L_1 is finite, generically each T_i satisfies all the ψ_U^t , and we choose one T for the whole level.

Now consider L_2 . Every element in the upper set lattice is the immediate successor of either one or two elements as in the following diagram:

We proved in Lemmas A.2 and A.4 that for the case on the left, a generic choice of (ψ_U^t, T) implies that ψ_V^t generically satisfies B1 and B2 for the same choice of T . For the case on the right, if $|U| \times |W|$ is taken as a product space, then $M_U \times M_W$ is a manifold and the results of Lemmas A.2 and A.4 immediately apply as well (indeed this is true for any V with a finite number of incoming edges, so we are not limited to the special case of an upper set lattice). Therefore the flows associated with upper sets in L_2 generically satisfy B1 and B2.

It is easily seen by the same argument that if B1 and B2 hold generically for the flows associated with the upper sets in L_{i-1} , then they hold generically for flows in L_i , where the projection relationships between the levels are satisfied. Since all flows satisfy B1 for $k \leq 2m + 1$, the requirements for Takens' theorem hold at all levels of UL . And since T is the same for all flows, then B1 and B2 hold generically in $\mathcal{P}_{CG} \times \mathbb{R}^+$. ■

An immediate consequence of Theorem A.6 is Theorem 3.6 in section 3, which we now prove for flows only.

Proof of Theorem 3.6 for flows. Let $\mathfrak{G} \subset \mathcal{P}_{CG} \times \mathbb{R}^+$ be the open and dense set established in Theorem A.6, so that $(\psi^t, T) \in \mathfrak{G}$ satisfies properties B1 and B2, and ψ^t fulfills the projection relationships $\pi_{VU} \circ \psi_V^t = \psi_U^t \circ \pi_{VU}$ for $U < V$ in the upper set lattice. Fix a point on the manifold $p \in M$. Then by Theorem 3.4 there exists a residual set $G(p) \subset C^2(M, M) \times \mathbb{R}^+$

such that $(\psi^t, T) \in G(p)$ satisfies B1 and B2, and the limit points of the flow starting at p are bijectively mapped to the limit points of the embedding constructed from a generic observation function. Then the set

$$G_{\mathcal{P}} := \mathfrak{G} \cap G(p)$$

is a residual subset of $\mathcal{P}_{\text{CG}} \times \mathbb{R}^+$ simultaneously satisfying the requirements of Theorem 3.4 and respecting the structure of the component graph CG. Since the projection of the positive limit points of ψ^t onto M_U are the positive limit points of ψ_U^t , the pair (ψ_U^t, T) generically satisfies the requirements of Theorem 3.4 while ensuring that $\psi_U^t = \pi_U \circ \psi^t$. ■

Theorems A.6 and 3.6 allow us to apply Theorems 3.3 and 3.4 to each element (ψ_V^t, M_V) of the filtration of (ψ^t, M) , as we did in section 3. We assume that the proof for genericity of conditions A1 and A2 for diffeomorphisms $\phi \in \mathcal{P}_{\text{CG}} \subset D^2(M, M)$ may be similarly proven, which would allow the application of Theorem 3.1 to the filtration as well.

Acknowledgments. We would like to thank Dr. Tim Sauer and an anonymous reviewer for helpful comments that improved the quality of this paper. We also thank Dr. Jarek Kwapisz for substantial assistance with the proof of Theorem A.2.

REFERENCES

- [1] V. I. ARNOL'D, *Geometrical Methods in the Theory of Ordinary Differential Equations*, Springer-Verlag, New York, 1983.
- [2] T. BUZUG AND G. PFISTER, *Optimal delay time and embedding dimension for delay-time coordinates by analysis of the global static and local dynamic behavior of strange attractors*, Phys. Rev. A, 45 (1992), pp. 7073–7084.
- [3] L. CAO, *Practical method for determining the minimum embedding dimension of a scalar time series*, Phys. D, 110 (1997), pp. 43–50.
- [4] M. CASDAGLI, S. EUBANK, J. D. FARMER, AND J. GIBSON, *State space reconstruction in the presence of noise*, Phys. D, 51 (1991), pp. 52–98.
- [5] A. M. FRASER AND H. L. SWINNEY, *Independent coordinates for strange attractors from mutual information*, Phys. Rev. A, 33 (1986), pp. 1134–1140.
- [6] J. F. GIBSON, J. D. FARMER, M. CASDAGLI, AND S. EUBANK, *An analytic approach to practical state space reconstruction*, Phys. D, 57 (1992), pp. 1–30.
- [7] C. W. J. GRANGER, *Investigating causal relations by econometric models and cross-spectral methods*, Econometrica, 1969, pp. 424–438.
- [8] M. W. HIRSCH, *Differential Topology*, Grad. Texts in Math. 33, Springer-Verlag, New York, 1976.
- [9] J. P. HUKÉ, *Embedding Nonlinear Dynamical Systems: A Guide to Takens' Theorem*, MIMS EPrint, 2006.
- [10] J. P. HUKÉ AND D. S. BROOMHEAD, *Embedding theorems for non-uniformly sampled dynamical systems*, Nonlinearity, 20 (2007), pp. 2205–2244.
- [11] D. KUGIUMTZIS, *State space reconstruction parameters in the analysis of chaotic time series—the role of the time window length*, Phys. D, 95 (1996), pp. 13–28.
- [12] W. LIEBERT, K. PAWELZIK, AND H. G. SCHUSTER, *Optimal embeddings of chaotic attractors from topological considerations*, Europhys. Lett. EPL, 14 (1991), pp. 521–526.
- [13] J. R. MUNKRES, *Topology*, Vol. 2, Prentice-Hall, Upper Saddle River, NJ, 2000.
- [14] J. PALIS AND W. DE MELO, *Geometric Theory of Dynamical Systems*, Springer, New York, 1982.
- [15] L. M. PECORA, T. L. CARROLL, AND J. F. HEAGY, *Statistics for mathematical properties of maps between time series embeddings*, Phys. Rev. E, 52 (1995), pp. 3420–3439.
- [16] T. SAUER, J. A. YORKE, AND M. CASDAGLI, *Embedology*, J. Stat. Phys., 65 (1991), pp. 579–616.

- [17] M. STERN, *Semimodular Lattices: Theory and Applications*, Cambridge University Press, Cambridge, UK, 1999.
- [18] S. STROGATZ, *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry and Engineering*, Perseus Books Group, 1994.
- [19] G. SUGIHARA, R. MAY, H. YE, C. HSIEH, E. DEYLE, M. FOGARTY, AND S. MUNCH, *Detecting causality in complex ecosystems*, *Science*, 338 (2012), pp. 496–500.
- [20] F. TAKENS, *Detecting strange attractors in turbulence*, in *Symposium on Dynamical Systems and Turbulence*, D. A. Rand and L. S. Young, eds., Springer-Verlag, Berlin, 1981, pp. 366–381.