

12-2020

Engaging Empirical Dynamic Modeling to Detect Intrusions in Cyber-Physical Systems

David R. Crow

Air Force Institute of Technology

Scott R. Graham

Air Force Institute of Technology

Brett J. Borghetti

Air Force Institute of Technology

Patrick J. Sweeney

Air Force Institute of Technology

Follow this and additional works at: <https://scholar.afit.edu/facpub>



Part of the [Computer Sciences Commons](#)

Recommended Citation

Version of record cited as:

Crow, D., Graham, S., Borghetti, B., Sweeney, P. (2020). Engaging Empirical Dynamic Modeling to Detect Intrusions in Cyber-Physical Systems. In: Staggs, J., Shenoi, S. (eds) *Critical Infrastructure Protection XIV. ICCIP 2020. IFIP Advances in Information and Communication Technology*, vol 596. Springer, Cham. https://doi.org/10.1007/978-3-030-62840-6_6

HAL Open Science Repository manuscript version cited as:

David Crow, Scott Graham, Brett Borghetti, Patrick Sweeney. Engaging Empirical Dynamic Modeling to Detect Intrusions in Cyber-Physical Systems. 14th International Conference on Critical Infrastructure Protection (ICCIP), Mar 2020, Arlington, VA, United States. pp.111-133, {10.1007/978-3-030-62840-6_6}. {hal-03794637}

This Conference Proceeding is brought to you for free and open access by AFIT Scholar. It has been accepted for inclusion in Faculty Publications by an authorized administrator of AFIT Scholar. For more information, please contact AFIT.ENWL.Repository@us.af.mil.



Engaging Empirical Dynamic Modeling to Detect Intrusions in Cyber-Physical Systems

David Crow, Scott Graham, Brett Borghetti, Patrick Sweeney

► To cite this version:

David Crow, Scott Graham, Brett Borghetti, Patrick Sweeney. Engaging Empirical Dynamic Modeling to Detect Intrusions in Cyber-Physical Systems. 14th International Conference on Critical Infrastructure Protection (ICCIP), Mar 2020, Arlington, VA, United States. pp.111-133, 10.1007/978-3-030-62840-6_6 . hal-03794637

HAL Id: hal-03794637

<https://inria.hal.science/hal-03794637>

Submitted on 3 Oct 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution| 4.0 International License



This document is the original author manuscript of a paper submitted to an IFIP conference proceedings or other IFIP publication by Springer Nature. As such, there may be some differences in the official published version of the paper. Such differences, if any, are usually due to reformatting during preparation for publication or minor corrections made by the author(s) during final proofreading of the publication manuscript.

Chapter 1

EMPIRICAL DYNAMIC MODELING AS A COMPONENT OF AN INTRUSION DETECTION SYSTEM

David Crow, Scott Graham, Brett Borghetti, and Patrick Sweeney

Abstract Modern cyber-physical systems require effective intrusion detection systems (IDSs) to ensure sufficient critical infrastructure protection. Before developing such an IDS, one requires an understanding of the behavior of the cyber-physical system and of the causality of its constituent parts. Such an understanding allows one to characterize normal behavior and, in turn, identify and report anomalous behavior. This research explores a relatively new time series analysis technique, empirical dynamic modeling (EDM), which may contribute to this understanding of a system. Specifically, we seek to determine whether this technique can adequately describe the causality in a system and thus give insights capable of serving as the foundation of a suitable IDS. Current research efforts aim to develop or improve upon today's IDSs; other efforts apply time series analysis techniques to relevant fields or to limited, controlled network attack scenarios. However, we have not identified in the literature any attempts to apply EDM to general-purpose IDS development. Our research seeks to address this gap. Our findings indicate that EDM may enable the understanding of a system required of an IDS architect. We thus encourage further research into EDM applications to IDSs and to cybersecurity in general.

Keywords: Cyber-physical systems, intrusion detection systems, causality, time series analysis, empirical dynamic modeling

1. Introduction

Intrusion detection systems (IDSs) are one of the most common, modern methods to defend against cyber-physical system (CPS) attacks and protect critical infrastructure. These systems monitor computer networks and report malicious activity to system administrators. In the

CPS domain, an IDS can detect attackers attempting to modify or misrepresent physical processes. Consider an automobile’s CPSs. If an attacker intends to, say, cause the driver to speed and thus receive a speeding ticket, the attacker may choose to inject packets detailing a lower speed, which would in turn cause the speedometer to display incorrect information. In this case, an effective IDS will notice that the data for speed does not conform to the expected behavior indicated by the data for the related physical processes (e.g., engine and wheel rotational velocities, throttle position, fuel efficiency). In other words, the IDS will notice that the speed readings are anomalous. As another example, if an IDS knows that a substantial increase in an automobile’s brake pressure likely precedes a relative decrease in velocity, the IDS can assert that no change, a small change, or an increase in velocity (after significant brake pressure) is anomalous. Of course, this requires an IDS capable of determining expected behavior and identifying anomalies. To design a capable IDS for a vulnerable CPS, IDS architects require the following:

- 1 Insight into the dynamics or patterns of a CPS, to include an understanding of the way in which some current system state enables predictions concerning a future state;
- 2 An ample quantity of data obtained under normal operating conditions to establish normal behavior;
- 3 A process to determine whether new traffic conforms to normal behavior; and
- 4 An alert system to report to the administrator the traffic that does not conform.

IDS architects can achieve (1) by obtaining either significant understanding of a CPS or sufficiently powerful computational resources. Often, the latter is infeasible: many CPSs are computationally limited by available hardware or by standards and regulations. Modern automobiles, for example, utilize small packets and fairly simple hardware. For this reason, the former is often more attainable. A solid understanding of a system’s dynamics, like how one signal affects another or how some current state predicts some future state—causality, in a word—allows an architect to develop a mechanism to identify anomalous traffic.¹ This research thus examines two different techniques to contribute to an architect’s understanding of a system. The first, Granger causality, is a well-known, simple method for evaluating the causality between two time series. The second, empirical dynamic modeling (EDM), is an

emerging field capable of more sophisticated time series analysis. Our research attempts to demonstrate a potential ability to use one or both techniques as the first step towards an IDS.

To do so, we apply EDM to two distinct datasets. We generate the first dataset using a simplistic model of the relationship between an automobile's steering wheel and the relative velocity of its two front wheels. Furthermore, the Air Force Research Laboratory (AFRL) maintains a flight simulator, the Avionics Vulnerability and Assessment System (AVAS), which generates the second dataset used in this research. The simulation computes various metrics, like airspeed, angle of attack, position, heading, and wind angle; this research concerns the airplane's airspeed, altitude, and pitch. The steering and AVAS datasets represent a linear system and a nonlinear one, respectively. The results of the EDM analyses imply a possibility of using the technique to develop sophisticated IDSs for nonlinear systems. For linear systems, EDM does not appear to reveal any previously unknown, important dynamics.

The remainder of this report examines the research in detail. Section 2 explains necessary background information. Section 3 describes the data and the analysis and evaluation tools. Section 4 presents the experimental results. Section 5 considers the implications of these results and possible opportunities for future research.

2. Background

This section introduces background information necessary for sufficient understanding of this work. It first describes CPSs and time series data before explaining causality and discussing EDM, a technique for nonlinear forecasting and causality analysis. We contend that this technique has potential as the basis for an effective IDS for CPSs. Finally, the section examines related work and presents an argument for our research.

2.1 Cyber-Physical Systems & Time Series Data

The Association for Computing Machinery Transactions on Cyber-Physical Systems defines CPSs as follows:

“Cyber-Physical Systems ... has emerged as a unifying name for systems where the cyber parts, i.e., the computing and communication parts, and the physical parts are tightly integrated, both at the design time and during operation. Such systems use computations and communication deeply embedded in and interacting with physical processes to add new capabilities to physical systems ... There is an emerging consensus that new methodologies and tools need to be developed to support cyber-physical systems.” [1]

A CPS can be represented by a model, but this model is typically difficult to understand or replicate.² For this reason, one must analyze the CPS’s output. Often, the output of CPS monitoring is time series data representing the value of some process (or processes) over time. One example of a time series in an aircraft is the propeller’s instantaneous revolutions per minute (RPM) over time, as measured by the aircraft’s sensors. The National Institute of Standards and Technology (NIST) says the following of time series analysis: “Time series analysis accounts for the fact that data points taken over time may have an internal structure (such as autocorrelation, trend or seasonal variation) that should be accounted for” [9]. Kotu and Deshpande contrast time series analysis and forecasting: “Time series *analysis* is the process of extracting meaningful non-trivial information and patterns from time series. Time series *forecasting* is the process of predicting the future value of time series data based on past observations and other inputs” [6]. Most techniques for both analysis and forecasting require data stationarity for the time series in question. Says NIST: “A stationary process has the property that the mean, variance and autocorrelation structure do not change over time ... a flat looking series, without trend, constant variance over time, a constant autocorrelation structure over time and no periodic fluctuations (seasonality)” [9].

Figure 1 presents examples of time series plots; panel (e) represents a stationary time series. Although these are arbitrary plots, they succinctly represent a wide array of potential time series. Many of the time series generated by an aircraft’s CPSs are non-stationary, so techniques that require stationarity are not usually viable for these data. EDM allows for non-stationary time series analysis and forecasting.

2.2 Empirical Dynamic Modeling

Floris Takens introduced the delay embedding theorem in 1981 [16]. Takens’ Theorem concerns mathematical attractors, where “an attractor is the value, or set of values, that a system settles toward over time” [2]. EDM is an application of Takens’ Theorem. In Sugihara et al.’s words, the field “is based on the mathematical theory of reconstructing system attractors from time series data” [15]. In practice, it allows one to model nonlinear dynamic systems with observational time series data. Figure 2 provides a summarized visual explanation of the main ideas in Takens’ Theorem and in EDM. Specifically, panel (A) depicts a Lorenz attractor³ as a model of a dynamic system. As the image shows, one can identify a time series for a given dimension by recording that dimension’s observations over time. Panel (B) shows that a univari-

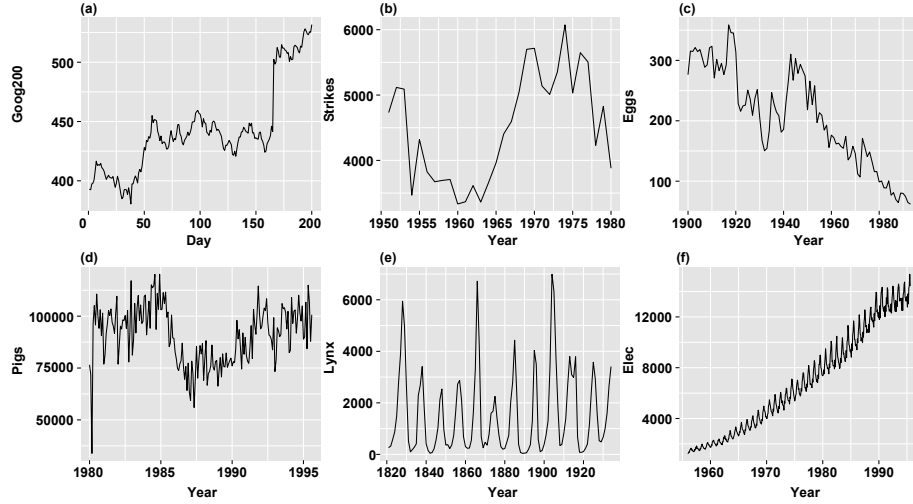


Figure 1. Examples of time series plots. (a) Google stock price for 200 consecutive days; (b) Annual number of [labor] strikes in the US; (c) Annual price of a dozen eggs in the US (constant dollars); (d) Monthly total of pigs slaughtered in Victoria, Australia; (e) Annual total of lynx trapped in the McKenzie River district of north-west Canada; (f) Monthly Australian electricity production. [5]

ate time series can be converted to a higher dimensional representation by using time-lagged versions of itself as additional dimensions. EDM calls the resulting manifold a shadow manifold. Takens showed that the shadow manifold is diffeomorphic (maps one-to-one) to its original attractor manifold M [16].

Sugihara et al. demonstrated that this diffeomorphic property between M and its shadow manifolds—one for each dimension—implies that the shadow manifolds are diffeomorphic with respect to each other. The opposite is also true. Thus, if two shadow manifolds are shown to be diffeomorphic with respect to each other, one can assume they belong to the same dynamic system. One can then use convergent cross-mapping (CCM), a mathematical technique recently developed by Sugihara et al., to identify the presence of and quantify the causality between the two original time series [14]. In short, CCM seeks to determine whether an arbitrary point and its nearest neighbors in one shadow manifold can accurately predict a point and its neighbors in another shadow manifold. Figure 3 summarizes this concept. Sugihara et al. showed that increasing the sample sizes for the shadow manifolds improves CCM's

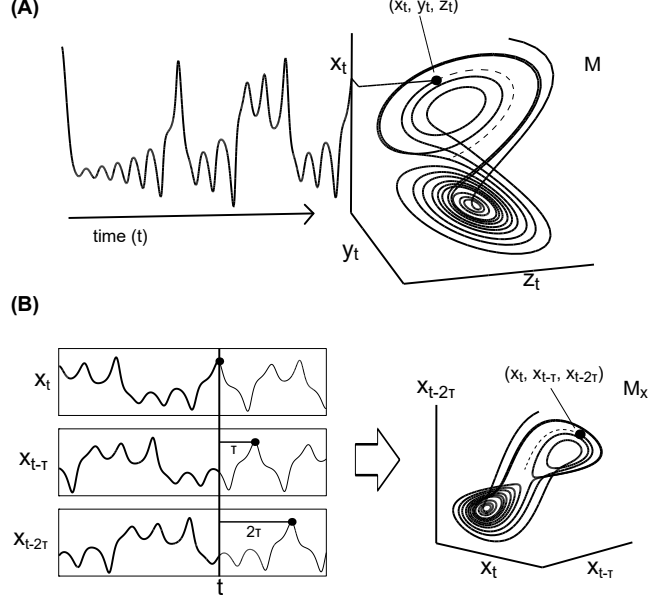


Figure 2. “Empirical dynamic modeling: (A) Example Lorenz system. The attractor manifold M is the set of states that the system progresses through time. Projection of the system state from M to the coordinate axis X generates a time series. (B) Lags of the time series X are used as coordinate axes to construct the shadow manifold M_X , which is diffeomorphic (maps 1:1) to the original manifold M . The visual similarity between M_X and M is apparent.” [15, ?]

predictive power, but they also showed that this predictive power converges to some maximum as the sample sizes increase to infinity [16, 14].

2.3 Related Work

Our investigation of the current literature revealed no research into EDM applications to automobile- or aircraft-generated time series or to cybersecurity as a whole. Most applications of the techniques concern economics or natural sciences; for example, the Sugihara Laboratory, from which EDM originated, primarily applies the techniques to ecology. We foresee EDM as a useful contribution to the cybersecurity domain, and our research thus seeks to link the two in a way not yet found in the literature.

3. Methodology

We utilize a relatively new statistical analysis tool to develop insights into a system’s characteristics, including its nonlinearity, deterministic




Figure 3. “Convergent cross mapping (CCM) tests for correspondence between shadow manifolds. This example based on the canonical Lorenz system (a coupled system in X , Y , and Z ...) shows the attractor manifold for the original system (M) and two shadow manifolds, M_X and M_Y , constructed using lagged-coordinate embeddings of X and Y , respectively ($lag = \tau$). Because X and Y are dynamically coupled, points that are nearby on M_X ... will correspond temporally to points that are nearby on M_Y ... This enables us to estimate states across manifolds using Y to estimate the state of X and vice versa using nearest neighbors ... With longer time series, the shadow manifolds become denser and the neighborhoods (ellipses of nearest neighbors) shrink, allowing more precise cross-map estimates” [19, 14]. The arrows between the manifolds represent the diffeomorphic properties of the attractors.

chaos, and causality. This section discusses the nature and origins of the experimental data. Additionally, the section describes the techniques used to analyze the data.

3.1 Data

This research utilizes data from two simulated CPSs. The first simulated dataset represents the effect of an automobile’s steering wheel angle over time on the RPM measurements for the turning wheels. The second dataset consists of captured nonlinear data generated by the AVAS, an AFRL-developed flight simulator that employs real-world physics and flight dynamics for research purposes. The steering dataset is considered to be linear because the relationship between each pair of time series is linear or nearly linear. Specifically, the relationship between the two wheel RPMs is linear, and the relationships between the steering input and each wheel RPM are almost linear.⁴ Similarly, the AVAS dataset is nonlinear because the relationships between the time series are nonlinear. The datasets allow us to evaluate the utility of both Granger

causality and EDM when used to analyze linear systems and nonlinear systems. Although simpler methods may enable effective analysis of linear systems, many CPSs of interest are nonlinear.

3.1.1 Linear Data. To fully assess EDM, we construct a dataset representative of a linear system. The variables describing the steering wheel angle and RPMs of the turning wheels in a passenger vehicle constitute such a system. We generate this time series with Python 3.7 for a vehicle with the following characteristics: a 30-inch wheel radius, including the tire; a 72-inch wheelbase; a 60-inch track; a maximum steering wheel turning angle of 360 degrees; a steering ratio of 8:1 (and thus a maximum wheel angle of 45 degrees); and a constant forward speed of 25 miles per hour.

Under these assumptions, a sum of sines function loosely represents some hypothetical driving scenario. That is, the Steering line in Figure 4 serves as a potential steering wheel angle time series, and we directly compute the inside and outside wheel RPMs using Equations (1) and (2), respectively. Note that, if the steering wheel angle $\theta < 0$, the left wheel is the inside wheel; otherwise, the right wheel is to the inside. Table 1 defines the variables used in the equations.




Figure 4. Plots of the steering system time series.

$$RPM_{inside} = \frac{60sb}{\pi r \left(b + \cos \left(90 - \frac{\theta}{t} \right) \sqrt{b^2 + \left(k + b \tan \left(90 - \frac{\theta}{t} \right) \right)^2} \right)} \quad (1)$$

$$RPM_{outside} = \frac{60s\sqrt{b^2 + \left(k + b \tan\left(90 - \frac{\theta}{t}\right)\right)^2}}{\pi r \left(b \sec\left(90 - \frac{\theta}{t}\right) + \sqrt{b^2 + \left(k + b \tan\left(90 - \frac{\theta}{t}\right)\right)^2}\right)} \quad (2)$$

Table 1. Variable Definitions for the RPM Equations

Variable	Meaning	Defined Value
r	wheel radius	15 inches
b	wheelbase	12 inches
k	track	60 inches
t	steering ratio	8:1
s	forward speed	25 miles per hour
θ	current steering wheel angle	not applicable

This steering system is rather rudimentary—it doesn’t account for the physical properties of a real system, including the effects of other relevant variables—but even its simplistic nature may allow us to draw conclusions concerning EDM’s applications to linear CPSs. The empirical results shown in Figure 5 confirm that the time series from the model are fairly linearly related.⁵ The values of the variables cover significantly different ranges. For this reason, we standardize all variables with R’s `scale` function⁶ to ensure each is equally important during analysis.




Figure 5. Scatter plots demonstrating the relationship between each pair of variables in the steering system.

3.1.2 Nonlinear Data.

To create the second dataset, we guided an AVAS-simulated aircraft through takeoff, low-altitude cruis-

ing, and multiple shallow banked turns. Our data collection yielded 7,582 observations from a 14-minute flight. Each observation includes eight different flying metrics and a timestamp relative to the start of the simulation. The metrics are roll and pitch (each in radians),⁷ altitude (in feet), and airspeed, vertical velocity, and velocity in each of the three coordinate axes (in feet per second). The roll and pitch values range from -180° to 180° ; altitude, airspeed, and the directional velocities are all floating point values.⁸

As with the linear dataset, we z-scale the variables prior to analysis. We then select a subset of the variables—airspeed, altitude, and pitch—before conducting the analyses. Other subsets of the eight variables likely exhibit the desired dynamics, but it is expected that these three variables best demonstrate a tightly coupled system. Figure 6 presents the three time series, prior to scaling, in one plot. Figure 7 clearly illustrates that the system is highly nonlinear.




Figure 6. Plots of the selected AVAS time series.

3.2 EDM Techniques

Ye et al. suggest the following sequence of EDM techniques to best interpret a dataset’s characteristics [19]:

- 1 Conduct nearest neighbor forecasting via simplex projection to identify the embedding dimension E which maximizes the prediction skill ρ [12];
- 2 Use simplex projection and E to determine whether the system exhibits deterministic chaos;




Figure 7. Scatter plots demonstrating the relationship between each pair of variables in the AVAS data collection.

- 3 Employ sequential locally-weights global linear maps (S-maps) to characterize any nonlinearity present in the data [13]; and
- 4 Utilize CCM to generate shadow manifolds, evaluate predictive accuracy, and quantify causality [14].

In essence, “simplex projection is the process of iteratively selecting [a point] Y_t in a shadow manifold and b other points whose histories over time t are most similar to the currently selected point ... A simplex is a generalization of a triangle or tetrahedron to an arbitrary number of dimensions” [12, 11, 7]. One then uses the weighted average of the future values of the b other points to make predictions about future values of Y_t . The difference between these predictions and the actual future values gives a forecast skill ρ . By repeating this process with shadow manifolds of different dimensionalities, one can identify the embedding dimension E that optimizes ρ [3]. The (strong) Whitney embedding theorem says the following [17]:

THEOREM 1 *Any m -manifold of class C^r ($r \geq 1$ finite or infinite) may be imbedded [sic] by a regular C^r -map in E^{2m} , and by such a map in a one-one manner in E^{2m+1} .*

In simpler terms, the theorem states that the embedding dimension E for an attractor manifold has an upper bound of $2D + 1$, where D is the true dimension (the number of variables) of the system [11, 3]. One can thus use simplex projection to definitively identify the optimal E in a finite amount of time.

S-map projection is another iterative process, but it instead uses all neighboring points to create linear regression vectors. By aggregating these regression vectors, one approximates an n -dimensional spline. One then compares this spline to the shadow manifold attractor to measure ρ

[13, 7, 3]. When generating the regression estimates, a nonlinear tuning parameter θ weights the neighbors with respect to their distance to the current focal point Y_t . Finally, “if ρ is maximized when $\theta = 0$, then the time series may be assumed to belong to a simple linear system instead of a dynamic system” [13, 11, 3].

As Stone et al. claim, “This process provides insight into the true dimensionality of the dynamic system responsible for generating [observational] data without requiring complete understanding of the system itself” [11]. Accurate knowledge of E is a prerequisite to effectively applying CCM to multiple time series to detect causality. Alternatively, a proper S-map analysis of time series relationships may indicate whether these relationships belong to a simple linear system. If so, computationally simpler methods, like Granger causality or auto-regressive linear models, could replace the more complex CCM technique in detecting causality [19, 13, 4]. Finally, knowledge of the dimensionality of a system may assist in creating a high quality model of said system. Such a model—and the results of a causality analysis—likely enables an effective IDS for various CPSs.

To conduct this analysis, we use the Sugihara Laboratory’s rEDM repository on GitHub. This codebase enables EDM analysis using the R programming language. The codebase includes the following functions (among others):

- `simplex`, which corresponds to the first and second EDM techniques;
- `s_map`, which corresponds to the third EDM technique; and
- `ccm` and `ccm_means`, which correspond to the fourth EDM technique.

These functions, together with a few helper functions, facilitate effective EDM analysis. Section 4 depicts the results of this analysis and Section 5 discusses the implications of these results. For the interested reader, Rennie provides an in-depth description of EDM, to include the mathematics behind simplex projection, S-map analysis, and CCM [10].

4. Results

This section presents the results of the EDM analyses for the two datasets. Section 4.1 shows that, for a linear dataset, EDM does not enable an effective causality analysis. Conversely, for nonlinear data, Section 4.2 presents a more effective use of EDM.

4.1 Linear Data

To effectively apply CCM to make predictions and quantify causality, we require knowledge of the optimal embedding dimension E for each time series in the system. By iteratively utilizing simplex projection to quantify predictive accuracy at different values for E , we identify the optimal value. Figure 8 illustrates the results of this process for each steering system time series. The plots show that forecast skill, or ρ ,⁹ is maximized when $E > 1$. We thus let $E = 2$ for the remainder of the EDM analysis techniques in this section because a lower dimensionality reduces complexity and processing time. To be clear, letting $E = 2$ means the techniques construct a two-dimensional shadow manifold, where each dimension is a time series lagged by some multiple of τ . When predicting steering wheel angle, for example, EDM constructs a shadow manifold using steering wheel angle and one copy of steering wheel angle, where the copy is lagged by τ . For this dataset, we let τ equal one second.




Figure 8. Plots illustrating the optimal embedding dimension for each steering system time series.

If we keep E constant and vary the time to prediction tp , simplex projection enables an analysis of a system's deterministic chaos. Figure 9 shows exactly this. Specifically, the figure shows how ρ decreases as tp increases for each of the three time series. In other words, predictions further in the future are much worse than those closer in time, which indicates chaotic behavior for the three variables. This is due to the nature of driving: without knowledge of the route, it is difficult, if not impossible, to predict a vehicle's steering wheel angle at a given time. The simulated data adheres to this interpretation of driving behavior. However, the difference in values for ρ at $tp = 0$ and at $tp = 10$ is only

about 0.0008; thus, chaotic behavior in this system is minuscule. EDM does not enable a deeper analysis of the system’s chaos.




Figure 9. Plots illustrating the deterministic chaos present in each steering system time series.

S-map analysis fits local linear maps to the system to describe its non-linearity. This is different from simplex projection, which analyzes each point’s nearest neighbors. By varying the nonlinearity tuning parameter θ in the S-map function call and plotting against ρ , we obtain the plots shown in Figure 10. When $\theta = 0$, S-map equally weights all points; as θ increases, the function more heavily weights points close to the point under analysis. Thus, when θ is higher, the function assumes more non-linearity in the system. For all three, ρ is greatest when θ is high, which indicates the presence of nonlinearity in each time series, but the range of y-values is so slim as to make these characteristics negligible. For a linear system such as this, it seems that EDM’s nonlinearity analysis is not particularly useful.

EDM also enables next-point predictions. Figure 11 overlays these predictions on each time series. Clearly, these predictions are extremely accurate, which indicates that the three variables do not change significantly from one observation to the next. Each plot also shows the prediction variance by way of a shaded polygon, but the variance is so low that the polygons are all-but-invisible. Remember that Figure 9 already implied this: when tp is small, ρ is very high. Additionally, we devised a naive prediction model. This model simply predicts that the point at time $t + 1$ has the same value as the point at time t . In other words, the simple model predicts *no change* for the next value. Figure 12 depicts the prediction errors (residuals) for both models; the majority of these errors are small, especially for EDM. Table 2 numerically compares the root-mean-square error (RMSE) ¹⁰ between the naive model and EDM.




Figure 10. Plots illustrating the nonlinearity of each steering system time series.

As the table illustrates, EDM outperforms the baseline predictor for each time series. These time series are incapable of large, instantaneous changes, so accurately predicting the next point is not very impressive and is not often useful in practical applications. However, it could still assist in IDSs of sufficiently low complexity. Of course, methods other than EDM may also suffice for linear systems.




Figure 11. Plots illustrating the predictions for each steering system time series.

Table 2. Root Mean Squared Error for Each Steering System Time Series

Time Series	Naive Prediction RMSE	EDM Prediction RMSE
Steering wheel angle	0.009424	0.003351
Left wheel RPM	0.005742	0.003893
Right wheel RPM	0.005742	0.003893




Figure 13 depicts inter-variable dynamics within the system. The figure plots cross-map skill ρ^{11} against library size—the number of points used to compute ρ —for each pair of variables. Each plot contains two lines, one for X xmap Y and one for Y xmap X . Here, X xmap Y refers to the CCM analysis technique which uses the shadow manifold of X to forecast the shadow manifold of Y . For a given library size, the resulting value for ρ indicates this predictive capability. The three plots show that ρ is equivalent across library sizes and in both directions for every pair of time series. This means that steering information is encoded in the RPM data and that RPM information is similarly encoded in the steering data, which in turn implies an expected causal effect in both directions. Unfortunately, it appears that EDM does not enable insight concerning pairwise causality for this dataset.




Figure 12. Plots illustrating the causality between each pair of steering system time series.

Finally, Figure 14 depicts the system’s causality through time. The lines again represent the results of using X to forecast Y , but we now plot ρ against tp . As previous Sugihara Lab researcher Hao Ye writes, “Note here that negative values of tp ... indicate that *past* values of Y are best cross-mapped from the reconstructed state of X . This suggests a dynamical signal that appears first in Y and later in X , and is consistent with Y causing X ” [18]. When tp is positive, the opposite holds. For this system, regardless of tp and of the variables in question, $\rho \approx 1$. Thus, according to EDM, each variable has a strong causal effect on every other variable regardless of the time to prediction. This is unlikely, and it supports the claim that EDM does not appear to enable sophisticated analysis of the system’s causality.




Figure 13. Plots illustrating predictive capability by analyzing the causality between each pair of selected steering system time series.

4.2 Nonlinear Data

For each of the three AVAS time series, Figure 15 presents the forecast skill ρ for various embedding dimensions E . Visually, differences in ρ are minuscule, but the optimal embedding dimension is two for each series. We thus let $E = 2$ for the remainder of the EDM analysis techniques in this section. Additionally, we again let τ equal the time between two observations in a given time series: one second.

Figure 16 plots the forecast skill ρ against the time to prediction tp to illustrate the system’s deterministic chaos. For each time series, the figure shows that predictions further in the future are much less accurate than earlier predictions. The respective y-scales show that this effect is strongest for pitch and weakest for altitude. Regardless, this is evidence of chaotic behavior for all three variables.




Figure 14. Plots illustrating the optimal embedding dimension for each selected AVAS time series.




Figure 15. Plots illustrating the deterministic chaos present in each selected AVAS time series.

As before, we plot ρ against θ to characterize each variable's nonlinearity. Figure 17 shows these plots. For airspeed and pitch, the figure shows that ρ is greatest when the function assumes the most nonlinearity; this is itself indicative of nonlinear dynamics. For altitude, the S-map analysis implies the absence of nonlinear dynamics in the time series, but it is important to note that the change in ρ —for all three plots—is extremely slight regardless of θ . For this reason, we cannot definitively claim the presence or absence of nonlinear dynamics.

Figure 18 presents the next-point predictions given by EDM for each time series. Unsurprisingly, the variance in the predictions—as shown by the nearly imperceptible shaded polygon—is slight. As Figure 16 strongly indicated, none of the variables change significantly between a




Figure 16. Plots illustrating the nonlinearity of each selected AVAS time series.

pair of observations. Figure 19, which depicts the prediction errors, confirms that EDM's predictions are highly accurate, and it also confirms that EDM again outperforms the simple model. Table 3 presents the RMSEs for both models. It is clear that EDM vastly outperforms the baseline model for two of the three time series; although the simple predictor performs better for pitch, the difference in RMSEs is insignificant. It is once again possible that even these short-term predictions could assist in IDS development. However, it is important to note an obvious limitation of EDM predictions: the technique cannot foresee values not contained in the library. This explains the large outlier predictions.




Figure 17. Plots illustrating the predictions for each selected AVAS time series.

Figure 20 depicts cross-map skill for each pair of time series. The leftmost plot shows that airspeed's manifold can effectively forecast altitude's but that the opposite relationship is noticeably weaker. The




Figure 18. Plots illustrating the prediction error for each selected AVAS time series for both the EDM predictor and the naive predictor.

Table 3. Root Mean Squared Error for Each Selected AVAS Time Series

Time Series	Naive Prediction RMSE	EDM Prediction RMSE
Airspeed	3.907005	0.070161
Altitude	18.192201	0.001535
Pitch	0.013749	0.019748

middle plot shows that the difference in cross-map skill between airspeed xmap pitch and pitch xmap airspeed decreases as library size increases. The rightmost plot shows a more extreme case of this: above a certain library size, cross-map skill. In all cases, the results indicate diminishing returns in improving ρ by increasing library size, but it is still possible that they enable analysis vital to better IDS design.

The last figure, Figure 21, plots cross-map skill against time to prediction. Consider for example airspeed xmap pitch. When tp is slightly less than zero, ρ is maximized; this implies that airspeed best predicts pitch when lagged by about one second. In other words, pitch strongly affects airspeed after one second. This is an expected behavior. When tp is positive, ρ quickly decreases and thus we assert that airspeed does not have a strong causal effect on pitch. This too is consistent with the standard interpretation of an airplane’s mechanics.

5. Conclusions

The AVAS, while useful for this research, is limited in scope. Future researchers who wish to affirm the conclusions found here should generate data with an extensively tested simulator, or they should obtain real

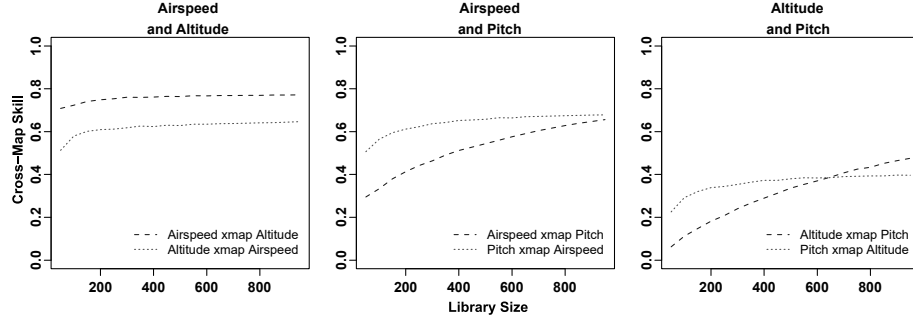


Figure 19. Plots illustrating the causality between each pair of selected AVAS time series.




Figure 20. Plots illustrating predictive capability by analyzing the causality between each pair of selected AVAS time series.

data from real aircraft. In the same vein, the simulated linear system is elementary and wanting for more real-world characteristics. Regardless, the results presented in Section 4 suggest two primary findings:

- 1 Although EDM can quantify behaviors present in linear systems, the results are often limited and so are not likely to aid in the development of IDSs.
- 2 For nonlinear systems, EDM is an easy-to-use suite of tools capable of evoking detailed insights that may assist in IDS design.

Although (1) implies that the analysis techniques explored here are unsuitable for linear systems, CPSs are often nonlinear. It is impor-

tant to note, however, that our limited linear system is likely not fully representative of a real-world linear system. For this reason, future researchers may wish to verify the applicability of the first two conclusions to a robust linear system.

Concerning (2), we assert that the analysis of a nonlinear system afforded by EDM successfully enables the understanding required by the first step towards an IDS. Section 4’s results demonstrate an ability to effectively quantify a nonlinear system’s causality, and this in turn enables better system insight for IDS architects. However, note that, although this research demonstrates the potential of EDM, we cannot realize its true value without larger, more realistic, and more complex datasets and without demonstrated success on a wider range of critical CPSs. Future work should thus address these limitations as well as the remaining steps towards an IDS—namely, obtaining quality data, identifying whether new traffic conforms to the patterns exhibited by the data, and creating a system to notify the administrator when the traffic does not. Still, we believe EDM to be a powerful, emerging tool relevant to critical infrastructure protection, and we strongly suggest further research into its applications to IDSs and beyond.

Disclaimer

The views expressed in this document are those of the authors and do not reflect the official policy or position of the United States Air Force, the United States Department of Defense, or the United States Government. This work is approved for public release under case number 88ABW-2020-049.

Notes

1. Assuming the architect has an ample quantity of normal data.
2. EDM attempts to give insight into a CPS’s model.
3. The Lorenz attractor is a set of solutions to the Lorenz system, a system of ordinary differential equations first studied by Edward Lorenz in 1963 [8].
4. That is, the latter two relationships are linear for steering wheel angles of relatively small magnitude but grow in nonlinearity as the steering wheel angle’s magnitude increases.
5. The somewhat nonlinear behavior between either wheel and the steering wheel is due to the mechanics of a standard automobile’s Ackermann steering mechanism.
6. For some time series X , the function z -scales X by subtracting its mean and then dividing it by its standard deviation.
7. We exclude yaw because, in AVAS, yaw is simply a measurement of the plane’s heading relative to north. In other words, it is not a characteristic of the plane’s dynamics.
8. Airspeed and altitude are both nonnegative.
9. Forecast skill is a measure of the ability to forecast future values of a given time series.
10. We use RMSE to heavily penalize large mispredictions; such errors would strongly affect an IDS’s performance.

11. Cross-map skill quantifies the ability to use one shadow manifold to identify values in another.

References

- [1] Association for Computing Machinery, ACM Transaction on Cyber-Physical Systems, tcps.acm.org/about.cfm, 2018.
- [2] G. Boeing, Visual analysis of nonlinear dynamical systems: Chaos, fractals, self-similarity and the limits of prediction, *Systems*, vol. 4(4), pp. 37–55, 2016.
- [3] C.W. Chang, M. Ushio and C.H. Hsieh, Empirical dynamic modeling for beginners, *Ecological Research*, vol. 32(6), pp. 785–796, 2017.
- [4] C. W. J. Granger, Investigating causal relations by econometric models, *Econometrica*, vol. 37(3), pp. 424–438, 1969.
- [5] R. Hyndman and G. Athanasopoulos, *Forecasting: Principles and Practice*, OTexts, Melbourne, Australia, 2018.
- [6] V. Kotu and B. Deshpande, *Data Science: Concepts and Practice*, Morgan Kaufmann, Cambridge, Massachusetts, USA, 2019.
- [7] J. Lee, *Introduction to Topological Manifolds*, Springer Science+Business Media, New York, New York, USA, 2011.
- [8] E. Lorenz, Deterministic nonperiodic flow, *Journal of the Atmospheric Sciences*, vol. 20(2), pp. 130–141, 1963.
- [9] National Institute of Standards and Technology, Introduction to Time Series Analysis, *NIST/SEMATECH e-Handbook of Statistical Methods*, Gaithersburg, Maryland (www.itl.nist.gov/div898/handbook/pmc/section4/pmc4.htm), 2012.
- [10] N. Rennie, Empirical dynamic models: A method for detecting causality in complex deterministic systems, 2018.
- [11] B. Stone, Enabling Auditing and Intrusion Detection of Proprietary Controller Area Networks, Dissertation, Department of Electrical and Computer Engineering, Air Force Institute of Technology, Dayton, Ohio, USA, 2018.
- [12] G. Sugihara and R. May, Nonlinear forecasting as a way of distinguishing chaos from measurement error in time series, *Nature*, vol. 344(6268), pp. 734–741, 1990.
- [13] G. Sugihara, Nonlinear forecasting for the classification of natural time series, *Philosophical Transactions of the Royal Society of London. Series A: Physical and Engineering Sciences*, vol. 348(1688), pp. 477–495, 1994.

- [14] G. Sugihara, R. May, H. Ye, C.H. Hsieh, E. Deyle, M. Fogarty and S. Munch, Detecting causality in complex ecosystems, *Science*, vol. 338(6106), pp. 496–500, 2012.
- [15] Sugihara Lab, Empirical Dynamic Modeling, University of California San Diego, San Diego, California, USA (deepecoweb.ucsd.edu/nonlinear-dynamics-research/edm/), 2019.
- [16] F. Takens, Detecting strange attractors in turbulence, in *Dynamical Systems and Turbulence*, D. Rand and L. Young (Eds.), Springer, Berlin Heidelberg, Germany, pp. 366–381, 1981.
- [17] H. Whitney, Differentiable manifolds in euclidean spaces, *Proceedings of the National Academy of Sciences of the United States of America*, vol. 21(7), pp. 462–464, 1935.
- [18] H. Yao, Using rEDM to quantify time delays in causation, (cran.r-project.org/web/packages/rEDM/vignettes/rEDM-time-delay-ccm.html), 2019.
- [19] H. Ye, A. Clark, E. Deyle and G. Sugihara, rEDM: An R package for empirical dynamic modeling and convergent cross mapping, (cran.r-project.org/web/packages/rEDM/vignettes/rEDM.html), 2019.