

Washington University in St. Louis

Washington University Open Scholarship

McKelvey School of Engineering Theses & Dissertations

McKelvey School of Engineering

Winter 12-15-2018

Information-based Analysis and Control of Recurrent Linear Networks and Recurrent Networks with Sigmoidal Nonlinearities

Deslin Menolascino

Washington University in St. Louis

Follow this and additional works at: https://openscholarship.wustl.edu/eng_etds



Part of the [Engineering Commons](#)

Recommended Citation

Menolascino, Deslin, "Information-based Analysis and Control of Recurrent Linear Networks and Recurrent Networks with Sigmoidal Nonlinearities" (2018). *McKelvey School of Engineering Theses & Dissertations*. 422.

https://openscholarship.wustl.edu/eng_etds/422

This Dissertation is brought to you for free and open access by the McKelvey School of Engineering at Washington University Open Scholarship. It has been accepted for inclusion in McKelvey School of Engineering Theses & Dissertations by an authorized administrator of Washington University Open Scholarship. For more information, please contact digital@wumail.wustl.edu.

WASHINGTON UNIVERSITY IN ST. LOUIS

School of Engineering & Applied Science
Department of Electrical & Systems Engineering

Dissertation Examination Committee:

ShiNung Ching, Chair

Zachary Feinstein

Jr-Shin Li

Baranidharan Raman

Shen Zeng

Information-based Analysis and Control of Recurrent Linear Networks and Recurrent
Networks with Sigmoidal Nonlinearities

by

Delsin Menolascino

A dissertation presented to
The Graduate School
of Washington University in
partial fulfillment of the
requirements for the degree
of Doctor of Philosophy

December 2018
St. Louis, Missouri

© 2018, Delsin Menolascino

Table of Contents

List of Figures	vi
List of Tables	xii
Acknowledgments	xii
Abstract	xv
Chapter 1: Introduction	1
1.1 Historical context	1
1.2 Stimulus encoding in dynamical (neural) networks	3
1.3 Relation between information and control.....	4
1.4 Organization	5
1.4.1 Orientation sensitivity	5
1.4.2 Sensitivity and input (stimulus) discrimination	7
1.4.3 Energy vs orientation sensitivity trade-offs evaluated via bispectral analysis	8
1.4.4 Information spectra and optimal background states	9
1.4.5 Quasilinear approximation of Fisher information in nonlinear, ‘sig- moidal’ networks	10
1.5 Contributions	11
1.5.1 Introduction of a novel network sensitivity metric.....	11
1.5.2 Analytic derivation, and geometric interpretation, of sinusoidal, recur- sive, orientation sensitivity expression for minimum-energy inputs	11
1.5.3 Novel analytical paradigm for assessing how networks trade energy for sensitivity	12
1.5.4 Derivation of—and analysis using—a novel Fisher information-based metric quantifying how much afferent input information is transmitted	12

1.5.5	Novel information analysis of sigmoidal nonlinear systems based on the Fisher information of a quasilinear approximation	13
Chapter 2:	Sensitivity to input orientation difference (novelty) in linear systems	15
2.1	Introduction	15
2.2	Problem Formulation	19
2.2.1	Mathematical notation	19
2.2.2	Input novelty-based controllability index	19
2.2.3	Minimum novelty problem	21
2.3	Results	22
2.3.1	Existence of a Minimally Novel Input	22
2.3.2	Uniqueness of the Minimally Novel Input	25
2.3.3	Euclidean - Inner Product Equivalence	25
2.4	Example	26
Chapter 3:	Endpoint-based Discriminability of Minimum Energy Inputs	29
3.1	Introduction	30
3.2	Problem Formulation	31
3.2.1	General Formulation	31
3.2.2	Endpoint-based Input Discriminability	32
3.3	Results	34
3.3.1	2-dimensional case	35
3.3.2	3-dimensional case	38
3.3.3	The n -dimensional case	40
3.4	Implications for Endpoint Decoding Schemes	43
3.4.1	Robust discrimination with regard to measurement noise	43
3.4.2	Robust discrimination with respect to input noise	45
3.5	Conclusions and Future Work	47
Chapter 4:	Bispectral Analysis for Measuring Energy-Orientation Tradeoffs in the Control of Linear Systems	49
4.1	Introduction	50
4.1.1	Motivation	50

4.2	Problem Formulation	52
4.2.1	Geometric Interpretation of the Controllability Gramian.....	52
4.2.2	Input Orientation and Orientation Range	54
4.3	Results.....	57
4.3.1	The Gramian Bispectrum	57
4.3.2	Geometric Interpretation.....	63
4.3.3	Comparing Orientation Bispectra.....	64
4.4	Conclusions and Future Work	69
Chapter 5: Information spectra and optimal background states for dynamical networks		71
5.1	Introduction.....	72
5.2	Results.....	75
5.2.1	Problem Formulation and Preliminaries.....	75
5.2.2	An optimal reference state \mathbf{x}_{ref} exists, maximizing information about \mathbf{u} .	80
5.2.3	An optimal reference input \mathbf{u}_{ref} exists, maximizing information about \mathbf{u} .	84
5.2.4	The optimally contrasting input targets specific nodes in a concentrated manner, but not necessarily nodes of highest degree.....	87
5.2.5	Information Spectra (of $\mathbf{S}_{\mathbf{u}}$) are Sensitive to Network Parameterization	88
5.2.6	Information Spectra are Related to the Controllability Gramian	91
5.3	Discussion.....	92
5.4	Methods.....	95
5.4.1	Derivation of $\mathcal{I}_{\mathbf{u}}$	95
5.4.2	Network parameterization and simulations	97
5.4.3	Time-varying \mathbf{u}	97
Chapter 6: Quasilinear approximation of Fisher information in networks with sigmoidal nonlinearities.....		102
6.1	Introduction.....	102
6.2	Problem Formulation 1.....	104
6.2.1	Quality of quasilinear approximation	111
6.2.2	Quasilinear Fisher information under formulation 1	114
6.3	Problem formulation 2.....	121
6.3.1	Quasilinear Fisher information under formulation 2	125

6.4	Fisher information differences as a function of saturation limits α and β	129
6.5	Relation between Fisher information under the two formulations.....	133
6.6	Discussion.....	137
6.7	Conclusions and future work.....	140
Chapter 7:	Conclusion	143
7.1	Summary.....	143
7.2	Future work	146
Bibliography	148

List of Figures

Figure 1.1:	Conceptual schematic of dissertation organization	6
Figure 2.1:	Prototypical structure of a sensory neuronal network. Sensory neurons are tuned to features from the sensory periphery. These neurons project excitation onto a network that performs intermediate transformations on the afferent excitation en route to higher brain regions.	17
Figure 2.2:	Minimum novelty control vs. minimum energy control: (A) The trajectory (blue) brings the system from an initial state on intermediate state at $t = 2s$. Subsequently, two trajectories are contrasted in the phase-plane for the minimum novelty control (red) and the minimum energy control (black). (B) The minimally novel inputs (from $t = 2s$ to $t = 4s$) (red) designed using our approach in this chapter. (C) The inputs corresponding to the minimum energy trajectory (from $t = 2s$ to $t = 4s$), (black).....	18
Figure 2.3:	Comparison of minimum novelty control with minimum energy control for $n = 1000$ random realizations of the recurrent neuronal network: Each red dot on the figure depicts the novelty associated with the solution to the minimum novelty control. Each blue dot corresponds to the minimum energy solution.....	28
Figure 3.1:	Trajectories under minimum-energy inputs are not uniformly distributed over the ellipse Ξ . Here, color bands represent an equal number of inputs, and all inputs \mathbf{u}_i $i \in \{1, \dots, m\}$ are uniformly distributed with respect to the novelty measure $\mathbf{d}_{\mathbb{J}}$ (i.e. $\mathbf{d}_{\mathbb{J}}(\mathbf{u}_i, \mathbf{u}_{i+k}) = \mathbf{d}_{\mathbb{J}}(\mathbf{u}_j, \mathbf{u}_{j+k}) \forall i, j \in \{1, \dots, m-1\}, \forall k \in \{1, \dots, m - \max\{i, j\}\}$	44
Figure 3.2:	ε -balls around final states generated by equally-spaced (in terms of novelty measure $\mathbf{d}_{\mathbb{J}}$) inputs	45
Figure 3.3:	Self-similar ellipses around terminal points. Left shows trajectories of $\mathbf{x}_{\mathbf{u}_i}$ under minimum-energy input \mathbf{u}_i . Right shows $\mathbf{x}_{\mathbf{u}_i}$ under $\mathbf{u}_i + \mathbf{w}_j$, where \mathbf{w}_j is noise (here constant for purpose of visualization).	47

Figure 4.1:	Unit circle with ellipse prescribed by $\mathbf{W}(T)^{-1}$ of a stable two-dimensional system with the two eigenvectors of $\mathbf{W}(T)^{-1}$. Since the ellipse represents the reachable set with fixed energy λ_1 , the distance between the ellipse and circle, in the direction of \mathbf{v}_2 , encodes the amount of excess energy available if we desire to steer the system to \mathbf{v}_2 . Note that the figure represents an abstract state space for a 2-d linear system, and thus the axes have no explicit units.	53
Figure 4.2:	Connection graphs for networks \mathcal{P} (blue) and \mathcal{Q} (green)	59
Figure 4.3:	Spectra of $\mathbf{W}_{\mathcal{P}}^{-1}$ and $\mathbf{W}_{\mathcal{Q}}^{-1}$. The horizontal axis indexes the eigenvectors, which are sorted according to the magnitude of their associated real (since \mathbf{W}^{-1} is real and symmetric) eigenvalues.....	60
Figure 4.4:	Array plots of gramian bispectra for networks \mathcal{P} and \mathcal{Q} . As explained in detail in Section 4.3.1, $\mathbf{d}_{\mathbb{J}}(\mathbf{u}_i, \mathbf{u}_j)$ quantifies the average cosine of the angle between inputs \mathbf{u}_i and \mathbf{u}_j , guiding the system to endpoints on the unit hypersphere in the directions of eigenvectors (of the gramian inverse \mathbf{W}^{-1}) \mathbf{v}_i and \mathbf{v}_j , respectively. Thus if $\mathbf{d}_{\mathbb{J}}(\mathbf{u}_i, \mathbf{u}_j) = 1$, the inputs may be identical, which explains why the diagonal entries (where, obviously, $\mathbf{v}_i = \mathbf{v}_j$) are all equal to 1.	60
Figure 4.5:	Spectra of \mathcal{P} and \mathcal{Q} , with the latter normalized by c (see (4.21)) so that the largest eigenvalues agree. This normalization does not affect the bispectral analysis (see (4.22)). As in Figure 4.3, the horizontal axis indexes the eigenvectors, which are sorted according to the magnitude of their associated eigenvalues.....	66
Figure 4.6:	A visualization of bispectrum difference matrix $\mathbf{D}_{\mathcal{P}, \mathcal{Q}}$, a tool for comparing the excess energy utility of the two networks. The $\{i, j\}^{\text{th}}$ entry is colored blue where network \mathcal{P} has the lower required angular input separation for maneuvers to \mathbf{v}_i and \mathbf{v}_j , and colored green if network \mathcal{Q} is better in this respect.	68
Figure 5.1:	The optimal background state \mathbf{x}_{ref} amounts to a Fisher linear discriminant, onto which state distributions (induced by inputs) are projected. In the case of Gaussian noise, uncertainty can be visualized in terms of ellipsoids (with principal axis \mathbf{v}_{max}) about the mean. Since the networks are dynamic, the optimal \mathbf{x}_{ref} will vary with time as the dynamics carry the states forward.	77

Figure 5.2:	Fidelity of optimally contrasting reference state \mathbf{x}_{ref} to system noise covariance decreases monotonically with n_{d} . Shown is how \mathbf{x}_{ref} aligns with the principal eigenvector (denoted \mathbf{v}_{max}) of noise covariance matrix $\Sigma_{\mathbf{w}}$. $\mu_{\text{Cos}(\theta)}$ is the mean, over 30 network realizations, of the cosine of the angular difference (θ) between \mathbf{x}_{ref} and \mathbf{v}_{max} . Error bars are standard deviations. 30 realizations were evaluated for (a) identity and (b) random \mathbf{B} matrices.	82
Figure 5.3:	Actuated nodes of the ideal background (\mathbf{x}_{ref}) are ‘required’ to be aligned with noise; non-actuated nodes are not. Shown is alignment of \mathbf{x}_{ref} with principal noise covariance direction \mathbf{v}_{max} (as in Fig. 5.2); here \mathbf{x}_{ref} and \mathbf{v}_{max} are partitioned so that (a) reflects only actuated and (b) only non-actuated nodes. $\mu_{\text{Cos}(\theta)}$ is as in Fig. 5.2, again for 30 network realizations.....	83
Figure 5.4:	The ideal contrast becomes more aligned with noise covariance as time progresses. Shown is the time-evolution of the relative orientation between the optimally contrasting state \mathbf{x}_{ref} and the principal noise eigenvector \mathbf{v}_{max} . At lower values of T , \mathbf{x}_{ref} is nearly orthogonal to \mathbf{v}_{max} , while as T gets larger, \mathbf{x}_{ref} becomes much more aligned with \mathbf{v}_{max} , although this alignment approaches a limit, which also varies monotonically with n_{d}	84
Figure 5.5:	The optimally contrasting input \mathbf{u}_{ref} targets network ‘hubs’, but to a degree which varies with n_{d} . 30 networks were realized with for each size (n) and driver node (n_{d}) combination. For a given network size, the graph shows the mean (μ_i) of the squared entries of the (normalized) optimal \mathbf{u}_{ref} . The entries of \mathbf{u}_{ref} are sorted according to the degree of targeted nodes (abscissa is a percentile, binned in increments of 5%, so that each b_i represents 5% of the nodes). Note that when $n = 100$ and $n_{\text{d}} = \frac{n}{10}$ there are twice as many bins as controlled nodes, hence the duplicity of values.	85
Figure 5.6:	Same setup as in Figure 5.5, but simulations are run for randomly connected networks (edge probability $p = .5$). Note the much smaller range of values on the vertical axis when compared with Figure 5.5. Nodes of high degree (‘hubs’) are targeted, but to a lesser extent than for the scale-free networks. Skewness of the graphs is inversely related to n_{d} ; that is, it is less necessary to target hubs for more fully actuated networks, until at $n_{\text{d}} = n$, \mathbf{u}_{ref} is essentially uniform.	85

Figure 5.7:	5.7a: Information spectra as function of number of actuated nodes (distributions aggregated over $n = 100, 200, 300, 400$). Spectra consist of a primary mode and a smaller secondary mode. 5.7b: Spectra of the controllability gramian for different fractions of actuated nodes. As noted in previous work, these spectra display an increasing number of modes as n_d decreases. The principal mode is inset. Comparing to the information spectra in Figure 5.7a, we see that information spectra show marked similarity to first mode of control spectra, and both spectra reveal outlying, small modes corresponding to easiest (control) and most informative (information) directions.	89
Figure 6.1:	An example illustrating the quality of the quasilinear approximation in 2D. (a) shows the state trajectory of the sigmoidal system (red) and the quasilinear system (blue). (b) shows the same sigmoidal trajectory (red) and the (heuristic) Jacobian-linearized system trajectory (green).	112
Figure 6.2:	An example illustrating the quality of the quasilinear approximation in 3D. (a) shows the state trajectory of the sigmoidal system (red) and the quasilinear system (blue). (b) shows the same sigmoidal trajectory (red) and the (heuristic) Jacobian-linearized system trajectory (green).	112
Figure 6.3:	We ran simulations of 50 networks each of quasilinear, Jacobian-linearized, and sigmoidal systems ((6.5),(6.29), and (6.3), respectively) with the same dynamical matrices (\mathbf{A}, \mathbf{B}), same noise gain matrix ($\mathbf{G} = 2.3 * \mathbf{I}$), the same (random for each trial) constant input \mathbf{u} (with $0 \leq \mathbf{u}_i \leq 15 \forall i \in \{1, \dots, n\}$) and same noise input ($\mathbf{z}(t)$) for each trial. Each trial was run using a forward-Euler integration method with $k = 10,000$ time steps, and saturation limits $\alpha = -2$ and $\beta = 3$. (a) shows mean and standard deviation of the value of (6.31) over the 50 trials. Note that for the Jacobian-linearized vs sigmoidal error (green), we just use \mathbf{x}_J (see (6.29)) in place of $\hat{\mathbf{x}}$ in (6.31). (b) shows mean and standard deviation of the variance discrepancy, as quantified by (6.32), over the 50 trials.....	114
Figure 6.4:	Mean and standard deviation of $\mathcal{J}_{\mathbf{u}_{ss}}^Q$ and $\mathcal{J}_{\mathbf{u}_{ss}}^L$ over 50 realizations of networks sized $n = 100$. Parameters as in Figure 6.3.	121

- Figure 6.5: An example illustrating the quality of the quasilinear approximation in 2D, under formulation 2. (a) shows the state trajectory of the sigmoidal system (red) and the quasilinear system (blue). (b) shows the same sigmoidal trajectory (red) and the (heuristic) Jacobian-linearized system trajectory (green). Notice that, in contrast with Figure 6.1, where the Jacobian linearization under-represents the variance of the original system, here it over-represents the variance. In both cases the quasilinear approximation is clearly much better. ... 123
- Figure 6.6: An example illustrating the quality of the quasilinear approximation in 3D, under formulation 2. (a) shows the state trajectory of the sigmoidal system (red) and the quasilinear system (blue). (b) shows the same sigmoidal trajectory (red) and the (heuristic) Jacobian-linearized system trajectory (green). Again the Jacobian linearization over-represents the variance of the original system, while the quasilinear approximation closely represents the original system dynamics. 124
- Figure 6.7: We ran simulations of 50 networks each of quasilinear, Jacobian-linearized, and sigmoidal systems ((6.56),(6.29), and (6.55), respectively) with the same dynamical matrices (\mathbf{A}, \mathbf{B}), same noise gain matrix ($\mathbf{G} = .9 * \mathbf{I}$), the same (random for each trial) constant input \mathbf{u} (with $3 \leq \mathbf{u}_i \leq 9 \forall i \in \{1, \dots, n\}$) and same noise input ($\mathbf{z}(t)$) for each trial. Each trial was run using a forward-Euler integration method with $k = 10,000$ time steps, and saturation limits $\alpha = -.5$ and $\beta = 1.5$. (a) shows mean and standard deviation of the value of (6.31) over the 50 trials. Note that for the Jacobian-linearized vs sigmoidal error (green), we just use \mathbf{x}_J (see (6.29)) in place of $\hat{\mathbf{x}}$ in (6.31). (b) shows mean and standard deviation of the variance discrepancy, as quantified by (6.32), over the 50 trials. 125
- Figure 6.8: Mean and standard deviation of $\mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{Q}*}$ and $\mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{L}}$ over 50 realizations of networks sized $n = 100$. Parameters as in Figure 6.7. 129
- Figure 6.9: We varied $|\alpha| = |\beta|$, where $\alpha < 0 < \beta$, between .5 and 10 (in increments of .5), simulating 10 networks of size $n = 40$ for each (α, β) pair. Matrices \mathbf{A}_i $i \in \{1, \dots, 10\}$ were constructed to be stable, and matrices \mathbf{B}_i $i \in \{1, \dots, 10\}$ were created by choosing random positive numbers from a uniform distribution over $(0, 1)$ and then scaling each row to have unit norm. The same noise gain matrix ($\mathbf{G} = 2.3 * \mathbf{I}$) was used throughout. For each trial, the same noise input ($\mathbf{z}(t)$) was applied to the quasilinear and linear system. Shown, for each (α, β) pair, are the mean and standard deviation of the quasilinear (under formulation 1) information $\mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{Q}}$ and linear information $\mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{L}}$ over the 10 network realizations. 130

Figure 6.10: Identical setup and parameters as for Figure 6.9, except here information is evaluated for quasilinear information ($\mathcal{I}_{\mathbf{u}_{ss}}^{Q*}$) under formulation 2 vs linear information $\mathcal{I}_{\mathbf{u}_{ss}}^L$ 131

Figure 6.11: Projection onto the first 2 of $n = 40$ dimensions. All parameters are identical to those used in Figures 6.9 and 6.10. Here we chose three values for $|\alpha| = |\beta|$ —.5, 5, and 10—to see how state-space trajectories evolved as a function of time. Terminal time $t = 10$ seconds for all panels. Notice that for $|\alpha| = |\beta| = .5$ (panel (a)), the time-evolution for the sigmoidal (red) and quasilinear (blue) systems is very slow compared to the Jacobian-linearized (green) system. This is predicted by the quasilinear theory, which says a small expected slope of $f(\cdot)$ leads to small values on \mathbf{N} 's diagonal and thus a slowing of dynamics. This phenomenon is still present to some degree for $|\alpha| = |\beta| = 5$ (panel (b)), but for $|\alpha| = |\beta| = 10$ (panel (c)) there is a nice blend of steady-state behavior in the sigmoidal and quasilinear trajectories combined with better information (evidenced by the smaller covariance ‘cloud’) relative to the Jacobian trajectory. This figure again serves to illustrate the quality of the quasilinear approximation generally, since even for higher-dimensional networks tracking of the sigmoidal system is good in all dimensions..... 132

Acknowledgments

I would like to acknowledge my advisor, Dr. ShiNung Ching, and thank him for his consistent, patient, and insightful guidance of my research endeavors. I have learned a lot from him in the solution of specific problems, but more importantly, Dr. Ching has instilled in me the desire and capacity to approach problems in a meaningful, scientific way, and to persist in the search for truth. We have developed the necessary advisor-student relationship, have learned to collaborate fruitfully, and along the way have become lifelong friends.

My colleagues in the lab have inspired and cheered me with their kindness, diligence, and intelligence. Gautam Kumar, MohammedMehdi Kafashan, Anirban Nandi, Sensen Liu, Fuqiang Huang, Matthew Singh, Sina Khanmohammadi, Peng Yi, Elham Ghazizadeh, Sruti Mallik, and Maren Loe have been brightened my days and tolerated my eccentricities; I thank them all.

Washington University generally has been a wonderful place to spend the past five years—this chapter of my life I will always treasure, and will recall these years, though sometimes fraught with difficulties, with special fondness. I would like to acknowledge the hard work of those who have designed, built, and maintained the physical plant and buildings of Washington University and the landscapers and groundspeople who maintain its natural beauty. Also the janitorial staff deserves mention for their faithful service. Though none of these people will

ever likely see this acknowledgement, they have made my surroundings safe, beautiful, clean, and inspiring, and I owe them a debt of gratitude for their work.

My immediate family—Cherie, Chiara, Ivan, Angelo, and Rosetta—are daily sources of strength, joy, inspiration, and friendship. No words can portray my love and gratitude for their presence in my life. My parents (Frank and Marie), step-parent (Don) and siblings (Christian and Liebe) have been lifelong companions, supporters, and friends.

Science is, at best and in essence, a study of nature, and nature—in her elegant mathematics, crystalline beauty, gracefully dancing dynamics, sublime poetry, awesome grandeur, boundless mystery, endless nourishment, and eternal song of hope—speaks most eloquently of her Designer. Thus I acknowledge (for lack of a better term) God—whose love creates and upholds all things—with my most profound gratitude.

Delsin Menolascino

Washington University in Saint Louis

December 2018

Dedicated to my sister Liebe

ABSTRACT OF THE DISSERTATION

Information-based Analysis and Control of Recurrent Linear Networks and Recurrent
Networks with Sigmoidal Nonlinearities

by

Delsin Menolascino

Doctor of Philosophy in Systems Science & Mathematics

Washington University in St. Louis, 2018

ShiNung Ching, Chair

Linear dynamical models have served as an analytically tractable approximation for a variety of natural and engineered systems. Recently, such models have been used to describe high-level diffusive interactions in the activation of complex networks, including those in the brain. In this regard, classical tools from control theory, including controllability analysis, have been used to assay the extent to which such networks might respond to their afferent inputs. However, for natural systems such as brain networks, it is not clear whether advantageous control properties necessarily correspond to useful functionality. That is, are systems that are highly controllable (according to certain metrics) also ones that are suited to computational goals such as representing, preserving and categorizing stimuli? This dissertation will introduce analysis methods that link the systems-theoretic properties of linear systems with informational measures that describe these functional characterizations. First, we assess sensitivity of a linear system to input orientation and novelty by deriving a measure of how networks translate input orientation differences into readable state trajectories. Next, we explore the implications of this novelty-sensitivity for endpoint-based input discrimination, wherein stimuli are decoded in terms of their induced representation in the state space. We develop a theoretical framework for the exploration of how networks utilize excess input

energy to enhance orientation sensitivity (and thus enhanced discrimination ability). Next, we conduct a theoretical study to reveal how the background or ‘default’ state of a network with linear dynamics allows it to best promote discrimination over a continuum of stimuli. Specifically, we derive a measure, based on the classical notion of a Fisher discriminant, quantifying the extent to which the state of a network encodes information about its afferent inputs. This measure provides an information value quantifying the ‘knowablility’ of an input based on its projection onto the background state. We subsequently optimize this background state, and characterize both the optimal background and the inputs giving it rise. Finally, we extend this information-based network analysis to include networks with nonlinear dynamics—specifically, ones involving sigmoidal saturating functions. We employ a quasilinear approximation technique, novel here in terms of its multidimensionality and specific application, to approximate the nonlinear dynamics by scaling a corresponding linear system and biasing by an offset term. A Fisher information-based metric is derived for the quasilinear system, with analytical and numerical results showing that Fisher information is better for the quasilinear (hence sigmoidal) system than for an ‘unconstrained’ linear system. Interestingly, this relation reverses when the noise is placed outside the sigmoid in the model, supporting conclusions extant in the literature that the relative alignment of the state and noise covariance is predictive of Fisher information. We show that there exists a clear trade-off between informational advantage, as conferred by the presence of sigmoidal nonlinearities, and speed of dynamics.

Chapter 1

Introduction

1.1 Historical context

The tradition of describing the dynamics of natural and engineered systems is a very old and rich one. Euclid, Pythagoras, and Archimedes gave the geometric foundation and logical framework for all mathematical and engineering disciplines, and wedded theoretical and practical mathematics, embodying the ethos of the scientist as engineer. Galileo, also a designer and practical engineer, brilliantly elucidated the physical principles which form the core of mechanics. Newton (also a practical man not afraid to work with his hands), who revolutionized science in so many ways that it almost defies credulity, formulated the laws of gravitation, force, and motion governing virtually all static and dynamic processes (at least those at ‘ordinary’ temporal and spatial scales). Not content with having discovered these fundamental laws, he subsequently inaugurated the rigorous study of dynamical systems by using these laws to describe physical processes in terms of differential and integral calculi which he co-created. Poincaré extended the tradition by finding ingenious ways to probe

the dynamics of nonlinear systems and systems of many variables, and to a large extent laid the foundation for modern graphical representation of time-varying system behavior. Poincaré also made the first basic discoveries, in his analysis of the three-body problem, of the mysterious dynamical phenomenon which is now known as chaos.

An extremely important historical development occurred between the time of Newton and the time of Poincaré—namely, the industrial revolution. It was no accident that industrialization came shortly after Newton’s ideas had been fully integrated into the scientific community; Newtonian ideas enabled and facilitated the new machine era. From early triumphs in agricultural, textile, and locomotive applications, the industrial revolution eventually permeated every facet of life, leading to new, efficient ways to produce virtually every material requirement of life. This revolution echoes down to our time, as more and more refined industrial techniques eventually enabled air travel, pharmaceuticals, production and storage of an almost endless variety and quantity of food, lighting, heating, and cooling technologies, multitudinous buildings and public works, electrical appliances of all sorts, modern automobiles, and the computer age. Now, as we arrive at computers, we must pause to reconsider the purposes for which a dynamical system is designed. For computers, dependent as they are on physical processes of manufacture and operation (i.e. automated circuit soldering, photoengraving of microprocessors, control technologies enabling hard disk pointer operation, to name just a few), have the explicit function of information processing. For the first time, machines are able to do things previously relegated to the realm of human (or animal) mental function. There is an important shift here—dynamical systems may be designed, not just to move things, but to transform ideas. The emergence of this technology has shed light on real neural function, just as discoveries in neuroscience have symbiotically paved the way for more efficient techniques of artificial computation and intelligence.

1.2 Stimulus encoding in dynamical (neural) networks

Now, it must be stated that there is often an intimate relation between physical processes occurring in (or to) a dynamical system and its qualities of ‘ideation’ or information processing. For example, in real neural systems, physical stimuli interface with sensory receptors, which encode the stimuli, pass the encoded information to deeper processing levels, and ultimately high-level thought involving the stimulus is generated. The initial coding is often done in a way which mimics the physical properties of the stimulus. For example, early visual processing cortex (V1-V5) has a ‘retinotopic’ neural map (this is an oversimplification to some extent, as there are actually many different types of visual representation—for color, contour, direction of motion, etc.—simultaneously occurring), meaning that there is a one-to-one correspondence between what the eye sees and the pattern of neural activation, so that ‘nearby’ visual stimuli will lead to ‘nearby’ neurons firing [1–4]. Likewise, in auditory cortex, there is a tonotopic organization of sound receptors, so that sounds resulting from oscillations of similar frequency will induce firing in nearby neurons [5–7]. Obviously, the amount of information thus encoded is enormous, and this encoding correlates very highly with physical reality. However, to form abstractions, associations, action plans, and just to make sense of it all, higher brain centers undoubtably condense this information dramatically and transmit it with high efficiency.

It is not well-known how this is done. We expect the nature of information processing to be different at different stages, but what does it mean for one network (dynamical system) to be different from another with respect to information flow? When asking how efficiently a system can physically move (as quantified by changes in a vector of state variables, herein denoted \mathbf{x}), we have many tools at our disposal with which a reasonable answer can be formed. In particular, for linear systems, or local linearizations of non-linear ones, we can readily assess stability (via the eigenvalues of the connectivity matrix, herein denoted \mathbf{A}),

controllability (via the Kalman rank condition on the controllability gramian, herein denoted \mathbf{W}), cost-of-control when guiding the system in orthogonal state space directions (via the spectrum of \mathbf{W}), and many others, to which an extensive literature is dedicated [8–13]. But if our goal is to quantify, in a meaningful way, the extent to which a system can efficiently propagate information, or to elucidate the qualitative nature of this propagation, we find that there are relatively few resources at our disposal ([14–16] are compelling attempts in this direction, but treat static, or very restricted dynamical, networks). The work presented in this dissertation reflects an attempt to address this lack by providing a few theoretical tools with which information-based comparisons can be made and qualitative information processing analyses conducted.

1.3 Relation between information and control

With these tools at our disposal, we address another important issue: the relationship between control-theoretic and information-theoretic properties of dynamical systems. Specifically, it is unclear whether dynamical networks that do a good job encoding and/or processing information are also those that are most responsive to their inputs in a control-theoretic sense. Hence we ask the second question: Is a network that is easily controlled by its inputs necessarily one that also effectively encodes information about those inputs?

These two primary motivations, 1) the derivation and deployment of quantitative analytic tools for assessing information-processing capabilities of linear (or linearized) dynamical systems, and 2) the comparison of control- and information-theoretic characteristics of these systems, enabling the assessment of how structure and dynamical aspects of a network relate to ‘good’ information encoding. Since we know, to a large extent, what kinds of network architectures and dynamical propensities lead to efficiency of control, we may now, armed

with the measures herein derived, conduct a comparative analysis. Our results suggest that, in general, systems that are ‘easy’ to control are also those that are efficient information encoders. The relation between the two (information and control) is mediated through the idea of sensitivity. The idea of network sensitivity will play an important role throughout this work. In fact, the first three chapters are devoted to various facets of networks sensitivity to inputs. Further, the fourth and fifth chapters feature a Fisher information-based analysis, where the Fisher information is proportional to the derivative(s) $\frac{df(\cdot)}{d\mathbf{u}}$ of the dynamical equation(s) $f(\cdot)$ with respect to input \mathbf{u} —by definition the sensitivity of $f(\cdot)$ to changes in \mathbf{u} . The critical link with control theory is easily made when we realize that controllability analysis also seeks to quantify network sensitivity to inputs, geometrically described by the reachable state space as a function of input energy. Thus we take network sensitivity as our point of departure.

1.4 Organization

Please see the schematic Figure 1.1 for a visual-conceptual picture of the dissertation’s main themes and how they are connected.

1.4.1 Orientation sensitivity

We begin (Chapter 2) by considering a network quality which underlies all higher-level information processing—sensitivity. We introduce a measure quantifying network sensitivity to input orientation difference, and show how, by using familiar optimization techniques, we may maximize sensitivity by finding the minimum input difference which can accomplish any desired pair of state transfers. Modeling neural networks with state-space methods is a natural and commonly-used paradigm, wherein each state variable corresponds to neuronal activation

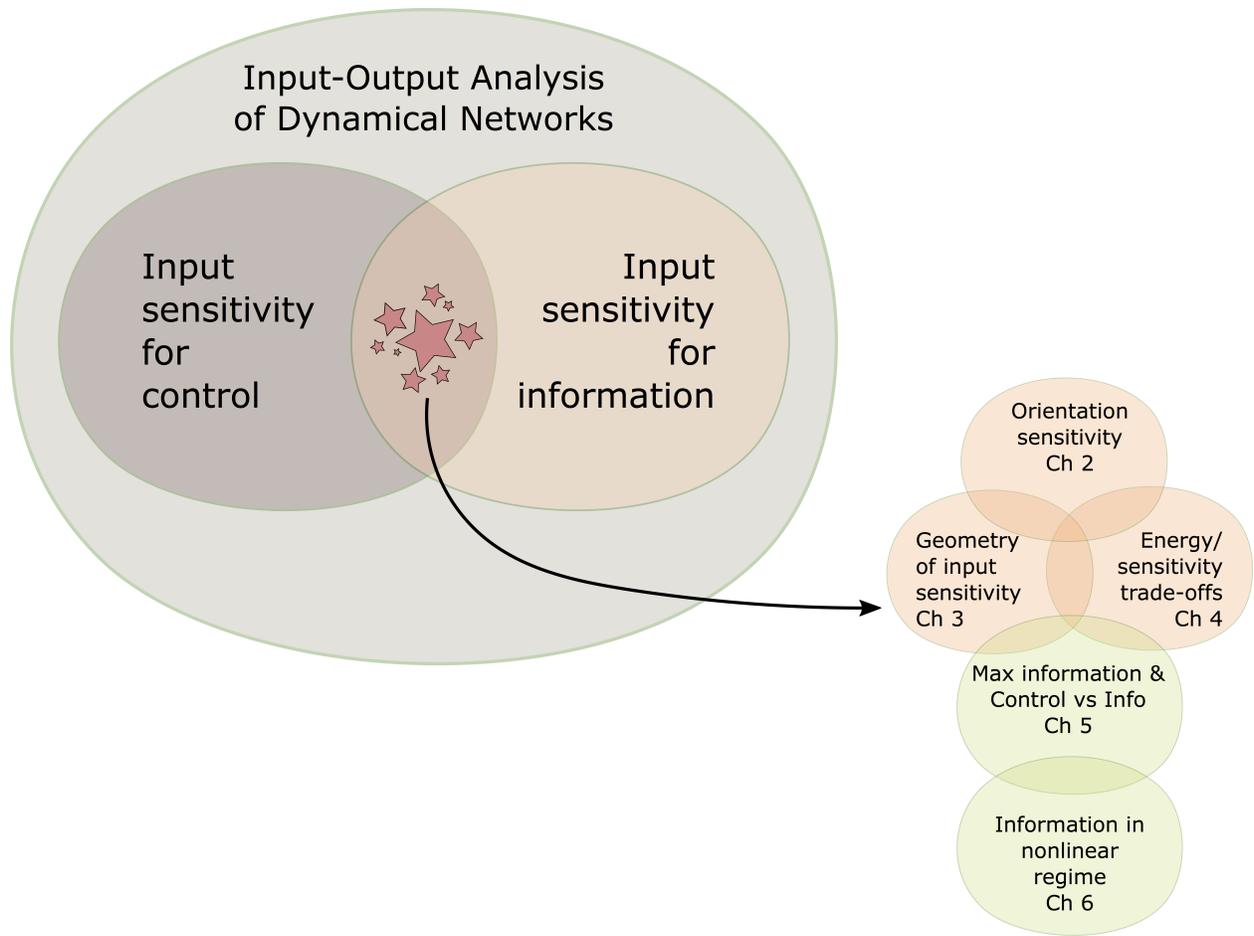


Figure 1.1: Conceptual schematic of dissertation organization

at some spatial scale, ranging from single-neuron spike patterns to mean-field electrical or BOLD behavior of large neuronal ensembles. In this paradigm, separation (usually measured by simple Euclidean distance) of stimulus-induced states allows the network to differentiate the stimuli (inputs). This ability of a system to discern input difference, by translating the input difference into ‘readable’ state-space separation, is critical, and obviously has major implications for higher-level information processing.

1.4.2 Sensitivity and input (stimulus) discrimination

With this in mind we explore (Chapter 3) the implications of the sensitivity metric (derived in Chapter 2) for endpoint-based input discrimination by developing an intuitive, geometric framework to explore the ways in which different inputs can be discriminated in recurrent linear dynamical networks. We derive a closed-form, recursive (with respect to dimension) expression for sensitivity to minimum-energy input orientation difference, which is independent of particular network dynamics and depends only on parametric state separation on the reachable ellipsoid with fixed energy. That is, for any controllable linear recurrent network, the exact total orientation difference between competing (minimum energy) inputs can be known by a simple sinusoidal function of the parametric distance between their resulting terminal states (recall that a point on an n -dimensional ellipsoid can be specified by $n - 1$ parameters; thus the parametric difference between two points can be expressed as an $n - 1$ -dimensional vector of differences in these parameters). That this is independent of network topology and dynamics is surprising, especially in light of the fact that Euclidean distance between terminal states of a given pair of inputs will very much depend on the particular network. In fact, the relation between the Euclidean and parametric difference has important implications for the discriminability of inputs in the presence of noise, as is explored geometrically and algebraically in the Chapter.

1.4.3 Energy vs orientation sensitivity trade-offs evaluated via bispectral analysis

In Chapter 4, we build on the previous results by developing a theoretical framework for the quantification of how orientation sensitivity is related to input energy. Specifically, we design a novel way to describe how a dynamical system “trades off” when presented with competing control-based and information-based objectives; thus, we quantify how a network may use “excess” stimulus energy (i.e. more than sufficient for a given network function) to enable heightened sensitivity and discriminatory ability. For example, if we require two inputs to ‘drive’ the system to points in the state space separated by a given distance, it is intuitive that if we allow the inputs unlimited energy, their trajectories (and thus their relative orientations) can be very similar and still accomplish the goal. Thus a network whose inputs are endowed with large amounts of excess (i.e. above minimum) energy can distinguish between inputs that are, on the whole, very similar. On the other hand, inputs which are more energy-constrained have a more circumscribed set of possible trajectories if their corresponding states are to reach the same terminal targets. In this case, the network ‘requires’ its afferent inputs to have more novelty for successful discrimination. The Chapter highlights the way in which the trade-off between available energy and orientation sensitivity plays out in different networks. This is quantified in a ‘bispectral matrix’, which encodes the energy-orientation trade-off as a function of direction of motion (in the state space). In particular, it is seen that some network \mathbf{P} may be more holistically ‘economical’ in this respect than a network \mathbf{Q} , but that \mathbf{Q} may have ‘sweet spots’ in the state space in which it is especially good (i.e. better than \mathbf{P} at discriminating inputs with limited energy. Since control objectives are often framed with respect to conservation of energy, the energy-sensitivity trade-off can be conceptually generalized to suggest a possible relationship between control- and

information-based objectives. This relationship is elucidated and made more mathematically precise in the following Chapter.

1.4.4 Information spectra and optimal background states

In Chapter 5, we introduce a metric, based on Fisher information theory, for the quantification of information transmission, showing how its value depends on network topology and time dynamics. A general optimization strategy is designed to maximize this information flow. We then apply the information metric to a discrimination/classification scenario wherein a linear dynamical network “compares” incoming stimuli to a default background state. Specifically, we quantify the “knowability” of an input based on its projection onto the background state, and subsequently optimize the background state using the maximization of total information as an objective. We highlight the relationship of this optimal background to underlying state noise covariance, and consider how this relationship changes (which it does) as a function of time. We characterize the ‘information spectrum’, based on the derived Fisher information matrix, and compare this spectrum to that of the controllability gramian matrix, establishing a link between fundamental control-theoretic network analysis and information processing. It is seen that networks that conserve energy effectively in a control-theoretic sense, are also those that tend to be more informative. Thus one of our primary motivations—that of understanding the relation between control- and information-theoretic qualities in dynamical systems—is squarely addressed. Further, we assess information at a graph-theoretical level by analyzing the extent to which network ‘hubs’ are targeted in control strategies which induce optimally informative contrasts (i.e. background states). It is found that hubs are much more critical for ‘informative’ control in scale-free networks than in random networks, perhaps because the nodes of highest degree in scale-free networks tend to be much more highly connected than those in random networks (witness the long tail in the power-law degree

distribution for scale-free networks vs. the more mean-centered binomial degree distribution of random networks).

1.4.5 Quasilinear approximation of Fisher information in nonlinear, ‘sigmoidal’ networks

Finally, in Chapter 6, we extend our analysis to the nonlinear domain by considering (noisy) networks which include nonlinear, sigmoidal ‘saturating’ functions, such as are commonly found in natural complex systems (especially in neural networks). Extending previous work, we find quasilinear approximations to the nonlinear dynamics, thus generating networks with dynamics which are statistically similar to their nonlinear counterparts, but which are much more amenable to analysis. We show the quality of the approximation both statistically and in concrete examples. We then proceed to derive expressions for the Fisher information of the quasilinear state with respect to its input, quantifying the extent to which the network encodes, by means of a state ‘readout’ (contrasted with the inner product-based readout used previously), information about an input. This informational analysis reveals something quite startling—the quasilinear information (hence, by the approximation, the sigmoidal system) is categorically better at encoding inputs than a corresponding unconstrained linear system. It would seem that a saturating function, which effectively serves to ‘squash’ the signal when it gets too large or small, would limit state-space representation of an input relative to an unconstrained system. However, our derivation shows that this is not the case. We assess the implications of this result for understanding the etiology of sigmoidal nonlinearities in dynamical systems generally, since an information-based motivation could hypothetically be a causal factor in the development of sigmoidal nonlinearities in complex systems. We also frame the counter-argument that, though sigmoidal nonlinearities appear to confer an informational advantage, since a quality approximation can be accomplished by means of a simple quasilinear

transformation, which is presumably much less biologically expensive, sigmoidal nonlinearities probably have their root cause in simple physical homeostatic mechanisms.

1.5 Contributions

We have made, we believe, important contributions both conceptually and technically. Conceptually, we have provided a general paradigm, with several distinct subdomains, for the exploration of how information processing in dynamical systems is related to control-theoretic principles, especially through the linking theme of input sensitivity. Technically, we have introduced/derived a number of tools for the quantification of information propagation, provided extensive numerical characterizations of how these metrics apply to simulated dynamical networks/systems, both linear and nonlinear, and have treated their interactions with existing metrics in a systematic and sometimes analytic way.

1.5.1 Introduction of a novel network sensitivity metric

- The theoretical derivation, via Hamiltonian optimization, of ‘orientation selectivity’ (or, ‘input novelty’) as a control measure
- Characterization of this sensitivity measure vis-à-vis dynamical systems analysis (with neural network example) [17]

1.5.2 Analytic derivation, and geometric interpretation, of sinusoidal, recursive, orientation sensitivity expression for minimum-energy inputs

- Derivation of a sinusoidal, recursive expression for orientation sensitivity (with respect to minimum-energy inputs)

- Proof that this expression depends only on outputs (more precisely, on the parameters determining output location on a feasible fixed-energy ellipsoid)
- Assessment of the implications for discrimination and classification in noisy dynamical networks [18]

1.5.3 Novel analytical paradigm for assessing how networks trade energy for sensitivity

- Development and deployment of an analytical paradigm for the assessment of how orientation sensitivity is related to input energy
- Derivation of a novel metric to quantify how a dynamical system “trades off”, when presented with competing control-based and information-based objectives
- Using this metric, a bispectral analysis is employed to characterize (and illustrate with examples) how networks may leverage ‘excess’ stimulus energy for heightened sensitivity
- Numerical characterizations of how the extent to which energy may be ‘traded’ for sensitivity depends on network topology and dynamics [19]

1.5.4 Derivation of—and analysis using—a novel Fisher information-based metric quantifying how much afferent input information is transmitted

- Theoretical derivation of a Fisher information-based metric quantifying the extent to which a dynamical network encodes its afferent inputs by means of a projection onto a ‘background’ state

- Optimization (i.e. maximization) of this metric with respect to the background, thus establishing a theoretical basis for the design and/or analysis of a network’s best (i.e. most informative) ‘default mode’
- Characterization of how the optimal background relates to system noise, as a function of time
- Comparison of how the information measure relates to network controllability via spectral comparisons of the derived Fisher information matrix and the controllability gramian
- Quantification of the extent to which maximizing information depends on targeting network ‘hubs’ for control [20]

1.5.5 Novel information analysis of sigmoidal nonlinear systems based on the Fisher information of a quasilinear approximation

- Extension of information-based analysis to the nonlinear regime by means of multivariate quasilinear approximation of networks with sigmoidal nonlinearities
- Extension of methods introduced in [21] to enable quasilinear approximation of arbitrary-dimensional recurrent networks whose sigmoidal functions act on non-stationary stochastic dynamical processes driven by Gaussian noise and afferent inputs
- Implementation of these derived methods, in a fixed-point optimization paradigm, to generate quasilinear approximations for medium-sized ($n \approx 100$) simulated networks, with statistical results presented verifying the accuracy of approximation
- Fisher information-based analysis of the resulting quasilinear networks, quantifying the extent to which these networks encode information about their inputs (in the state)

- Derivation of quasilinearization and Fisher information is conducted with respect to two distinct modeling paradigms—one with noise ‘inside’ the sigmoid (hence the quasilinear transformation), and the other with noise ‘outside’
- It is shown that quasilinear (and sigmoidal by approximation) Fisher information (QFI—which is holistically quantified by the trace of the Fisher information matrix) is better than for a Jacobian-linearized, ‘unconstrained’ linear network (LFI) in the first formulation, while the reverse is true in the second formulation
- Analytical results are given specifying the mathematical relationship between QFI under the two formulations
- Numerical characterization of how QFI compares to linear Fisher (LFI) information, as a function of saturation limits (i.e. parameters of the sigmoid)
- Analytical results are shown for a special case wherein information matrices for the quasilinear and linear systems differ a multiplication by the quasilinear ‘gain’ matrix (i.e. the optimized approximation parameters exactly determine the information)

Chapter 2

Sensitivity to input orientation difference (novelty) in linear systems

2.1 Introduction

In this chapter, we propose a novelty-based index for quantitative characterization of the controllability of complex networks. This inherently bounded index describes the average angular separation of an input with respect to the past input history. We use this index to find the minimally novel input that drives a linear network to a desired state using unit average energy. Specifically, the minimally novel input is defined as the solution of a continuous time, non-convex optimal control problem based on the introduced index. We provide conditions for existence and uniqueness, and an explicit, closed-form expression for the solution. We support our theoretical results by characterizing the minimally novel inputs for an example of a recurrent neuronal network.

In its most basic form, the systems-theoretic notion of controllability carries a binary definition: a dynamical system either is, or is not, controllable, with respect to its exogenous inputs. Naturally, such a notion has the deficiency of not grading the ease or difficulty associated with effecting such control. To obviate this issue, consistent research effort has been directed at the characterization of controllability using systems-theoretic metrics. Roughly, these metrics can be grouped into two categories

1. Those that characterize the minimum energy parametric perturbations that result in a loss of controllability [13, 22]. These are related to basic characterizations of the robustness of linear systems [23].
2. Those that characterize the controllability of a system in terms of the minimum energy excitation required to achieve a unit length state trajectory [8, 24–26].

The latter, in particular, is a natural paradigm that is directly relatable to the celebrated Kalman rank condition (or the controllability gramian) used to ascertain the controllability of linear systems [27]. Recently, energy-based controllability metrics have been successfully used in the emerging domain of network science to assess the putative controllability of large-scale linear systems, formulated as complex networks of various topologies [8, 25]. However, for complex networks in general and, in particular, for biological neuronal networks, an energy-based metric offers insight into only one aspect of the overall system’s controllability.

We appeal, specifically, to the domain of neural coding and the dynamics of sensory neural circuits. Consider the simple, prototypical layered model of a sensory network shown in Figure 2.1, wherein sensory neurons are tuned to a high dimensional feature space (i.e., environmental variables from the sensory periphery; say, different molecules corresponding to tastes). Those sensory neurons impinge on a complex, interconnected sensory network that performs intermediate transformations en route to higher brain areas.

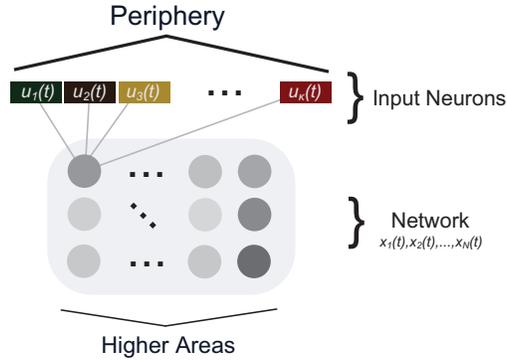


Figure 2.1: Prototypical structure of a sensory neuronal network. Sensory neurons are tuned to features from the sensory periphery. These neurons project excitation onto a network that performs intermediate transformations on the afferent excitation en route to higher brain regions.

One may put forth a supposition that the ‘controllability’ of such a sensory network, with respect to the afferent input from the sensory neurons, is critical in mediating the ability to perceive minute changes in the environment. But as much as energy is important in mediating such a response, orientation, i.e., the alignment of an input with certain features, may be even more so. Indeed, a weak, but highly novel input may be more easily perceived than an intense, but more familiar, stimulus. The ability to assess the responsiveness of neuronal networks to input orientation – at a particular moment in time, relative to past inputs – has immediate implications in the analysis and control of biophysiological neuronal network dynamics in different behavioral and clinical regimes [28–30].

Here, as a first step, we seek to characterize the controllability of linear systems (linear networks) possessing high dimensional input-spaces, with respect to input sensitivity. In particular, we ask how responsive are the state (node) trajectories to inputs that differ in orientation from those that have previously been applied. Figure 2.2 illustrates the basic notion of input novelty for a simple two-dimensional linear system with three inputs. A particular input drives the system to an intermediate point in the phase space; from this point emerge two trajectories, both of which reach a common endpoint; one minimizes input

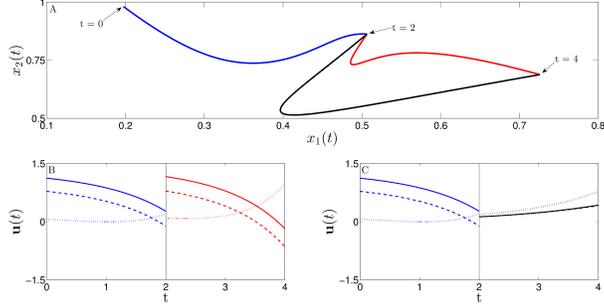


Figure 2.2: Minimum novelty control vs. minimum energy control: (A) The trajectory (blue) brings the system from an initial state on intermediate state at $t = 2s$. Subsequently, two trajectories are contrasted in the phase-plane for the minimum novelty control (red) and the minimum energy control (black). (B) The minimally novel inputs (from $t = 2s$ to $t = 4s$) (red) designed using our approach in this chapter. (C) The inputs corresponding to the minimum energy trajectory (from $t = 2s$ to $t = 4s$), (black).

novelty (note the similarity between the input from $t \in [0, 2]$ and that from $t \in [2, 4]$), the other minimizes energy.

Specifically, we: (i) analytically derive the minimum novelty control for linear networks by formulating a non-convex optimization problem. The problem seeks the minimum angular separation, defined in terms of an inner product in the input feature space, required in order to create a desired change in the network trajectory, constrained by a fixed average input energy; and (ii) characterize the resulting cost – the control ‘novelty’ – that describes the change in input orientation that is required to drive the system to a given state.

The remainder of the chapter is as follows. In section 2.2, we introduce our inner-product based controllability index for linear networks and formulate a non-convex optimal control problem that minimizes this index under the constraint of unit average energy. In section 3.3, we establish the existence and the uniqueness of a global optimal solution of the control problem and derive a closed-form expression for minimally novel inputs. Finally, in section 3.4, we consider a linearized firing rate model of a recurrent neuronal network as an example

to demonstrate our theoretical results. The chapter concludes with a summary and discussion of future work.

2.2 Problem Formulation

2.2.1 Mathematical notation

Most notation is standard and will be introduced as the results are developed. We use lower-case letters to represent scalars, boldface lower-case letters to represent vectors, capital letters to represent matrices. Exceptions are T , $\mathbb{J}(T)$ and $\mathbb{J}_1(T)$, which we represent as scalars. We use $\mathbb{R}^{n \times 1}$ to represent the space of n -dimensional vectors with their elements as real numbers. Similarly, we use $\mathbb{R}^{n \times m}$ and $\mathbb{R}_+^{n \times m}$ to represent the space of $n \times m$ dimensional matrices with their elements as real numbers and non-negative real numbers respectively. $\|\mathbf{x}\|_2$ is the Euclidean norm of the vector \mathbf{x} . \mathbf{x}' is the transpose of a vector \mathbf{x} and A^{-1} is the inverse of a matrix A .

2.2.2 Input novelty-based controllability index

We consider a linear, time invariant system with dynamics of the form

$$\frac{d\mathbf{x}(t)}{dt} = A\mathbf{x}(t) + B\mathbf{u}(t) \tag{2.1}$$

Here $\mathbf{x}(t) \in \mathbb{R}^{n \times 1}$ represents the state of the system at time t , $A \in \mathbb{R}^{n \times n}$ is the state transition matrix, $B \in \mathbb{R}^{n \times m}$ is the input matrix, and $\mathbf{u}(t) \in \mathbb{R}^{m \times 1}$ is the input to the system. Without loss of generality, we say that (2.1) describes the time evolution of linear networks in the presence of external inputs.

Let us assume an input $\mathbf{v}(t-T) \in \mathbb{R}^{m \times 1}$, $t \in [0, T]$, with total energy T , i.e.

$$\frac{1}{T} \int_0^T \|\mathbf{v}(t-T)\|_2^2 dt = 1 \quad (2.2)$$

We assume that $\mathbf{v}(t-T)$ can drive $\mathbf{x}(t)$ from $\mathbf{x}(-T)$ to $\mathbf{x}(0)$, where $\|\mathbf{x}(0)\|_2 = 1$, subject to the dynamics (2.1). Here $T > 0$ is a constant. We introduce the inner-product based index

$$\mathbb{J}(T) = \frac{1}{T} \int_0^T \mathbf{v}'(t-T) \mathbf{u}(t) dt \quad (2.3)$$

where

$$\frac{1}{T} \int_0^T \|\mathbf{u}(t)\|_2^2 dt = 1, \quad (2.4)$$

to measure the novelty of a subsequent input $\mathbf{u}(t)$, $t \in [0, T]$, relative to $\mathbf{v}(t-T)$, required in order to reach the state $\mathbf{x}(T)$, where $\|\mathbf{x}(T)\|_2 = 1$. In other words for a fixed input energy, the metric $\mathbb{J}(T)$ measures the required directional change in inputs (thus novelty) to achieve a given state (or, equivalently, directional) change in the state of the system.

Remark 1 It is readily evident that $\mathbb{J}(T) \in [-1, 1]$.

Remark 2 From (2.3), we note that the novelty of the input $\mathbf{u}(t)$ compared to $\mathbf{v}(t-T)$ decreases as $\mathbb{J}(T)$ increases and is minimum when $\mathbb{J}(T) = 1$ i.e. when $\mathbf{u}(t) = \mathbf{v}(t-T)$ for all $t \in [0, T]$.

Remark 3 We observe that, due to the energy normalization in (2.2) and (2.4),

$$\frac{1}{T} \int_0^T \|\mathbf{v}(t-T) - \mathbf{u}(t)\|_2^2 dt = 2(1 - \mathbb{J}(T)) \quad (2.5)$$

Thus, the average Euclidean distance, i.e. the left hand side of (2.5), between two inputs can equivalently be used as an alternate measure of input novelty in our context.

2.2.3 Minimum novelty problem

From the conceptual formulation introduced above, we can develop a control problem to design the minimally novel input $\mathbf{u}(t), t \in [0, T]$ such that a desired directional change in the state of the system can be achieved under the constraint of fixed energy subject to the system dynamics (2.1). For this, we formulate the following optimal control problem:

$$\min_{\substack{\mathbf{u}(t) \\ t \in [0, T]}} -\mathbb{J}(T) \tag{2.6a}$$

$$\text{s.t.} \quad \frac{1}{T} \int_0^T \|\mathbf{u}(t)\|_2^2 dt = 1 \tag{2.6b}$$

$$\mathbf{x}(T) = e^{AT} \mathbf{x}(0) + \int_0^T e^{A(T-t)} B \mathbf{u}(t) dt \tag{2.6c}$$

It should be noted here that the constraint (2.6c) is obtained by integrating (2.1) with respect to t over the period of $[0, T]$. Immediately, we note that the quadratic equality constraint (2.6b) makes the optimization problem (2.6) non-convex. Furthermore, we note that our optimal control problem formulation (2.6) is different from the classical minimum effort problems where the L^1 -norm of control inputs is minimized under the constraints of explicit lower and upper bounds on the inputs.

2.3 Results

We derive conditions for the existence of a unique global optimal solution of the non-convex optimization problem (2.6). Based on this, we provide a closed-form expression for the optimal $\mathbf{u}(t), t \in [0, T]$.

2.3.1 Existence of a Minimally Novel Input

Lemma 1 A solution of the non-convex optimization problem (2.6) exists if

$$T > \max \{ \mathbf{s}'(T)W_c^{-1}(T)\mathbf{s}(T), \mathbf{r}'(T)W_c^{-1}(T)\mathbf{r}(T) \} \quad (2.7)$$

where

$$\mathbf{s}(T) = \int_0^T e^{A(T-t)} \mathbf{B}\mathbf{v}(t-T) dt \quad (2.8a)$$

$$\mathbf{r}(T) = \mathbf{x}(T) - e^{AT} \mathbf{x}(0) \quad (2.8b)$$

Here, $W_c(T)$ is the controllability gramian at time T and is defined as

$$W_c(T) = \int_0^T e^{A(T-t)} \mathbf{B}\mathbf{B}' e^{A'(T-t)} dt \quad (2.9)$$

Recall that by our formulation, T is the total energy available to the system (2.1).

Remark 4 The arguments $\mathbf{s}'(T)W_c^{-1}(T)\mathbf{s}(T)$ and $\mathbf{r}'(T)W_c^{-1}(T)\mathbf{r}(T)$ in (2.7) are the minimum energy required to drive the system (2.1) from $\mathbf{x}(-T)$ to $\mathbf{x}(0)$ and $\mathbf{x}(0)$ to $\mathbf{x}(T)$ respectively [26, 27].

Proof 1 Define $y(t)$ as

$$y(t) = \frac{1}{T} \int_0^t \|\mathbf{u}(\tau)\|_2^2 d\tau \quad (2.10)$$

Clearly, $y(0) = 0$ and $y(T) = 1$ from (2.6b). Thus, we can replace the constraint (2.6b) by

$$y(T) = 1 \quad (2.11)$$

In differential form, we can write (2.10) as

$$\frac{dy(t)}{dt} = \frac{1}{T} \|\mathbf{u}(t)\|_2^2 \quad (2.12)$$

To solve the dynamic optimization problem (2.6a), (2.6c) and (2.11) in continuous time, we write the Hamiltonian $\mathcal{H}(\mathbf{x}(t), y(t), \mathbf{u}(t), \boldsymbol{\lambda}(t), \boldsymbol{\mu}(t), t)$ as

$$\mathcal{H}(\mathbf{x}(t), y(t), \mathbf{u}(t), \boldsymbol{\lambda}(t), \boldsymbol{\mu}(t), t) = -\frac{1}{T} \mathbf{v}'(t-T)\mathbf{u}(t) + \boldsymbol{\lambda}'(t)(A\mathbf{x}(t) + B\mathbf{u}(t)) + \frac{\boldsymbol{\mu}(t)}{T} \|\mathbf{u}(t)\|_2^2 \quad (2.13)$$

Here, $\boldsymbol{\lambda}(t)$ and $\boldsymbol{\mu}(t)$ are the costate variables associated with the dynamics (2.1) and (2.12) respectively. We derive the following optimality conditions (i.e. the Euler-Lagrange equations [31]):

$$\frac{d\boldsymbol{\lambda}(t)}{dt} = -\left(\frac{\partial \mathcal{H}(\mathbf{x}(t), y(t), \mathbf{u}(t), \boldsymbol{\lambda}(t), \boldsymbol{\mu}(t), t)}{\partial \mathbf{x}(t)}\right)' = -A'\boldsymbol{\lambda}(t) \quad (2.14a)$$

$$\frac{d\boldsymbol{\mu}(t)}{dt} = -\frac{\partial \mathcal{H}(\mathbf{x}(t), y(t), \mathbf{u}(t), \boldsymbol{\lambda}(t), \boldsymbol{\mu}(t), t)}{\partial y(t)} = 0 \quad (2.14b)$$

$$\begin{aligned} \frac{\partial \mathcal{H}(\mathbf{x}(t), y(t), \mathbf{u}(t), \boldsymbol{\lambda}(t), \boldsymbol{\mu}(t), t)}{\partial \mathbf{u}(t)} &= 0 \\ &= \frac{2\boldsymbol{\mu}(t)}{T} \mathbf{u}(t) - \frac{1}{T} \mathbf{v}(t-T) + B'\boldsymbol{\lambda}(t) \end{aligned} \quad (2.14c)$$

By integrating the costate equations (2.14a) and (2.14b) over t , we obtain

$$\boldsymbol{\lambda}(t) = e^{-A't} \boldsymbol{\lambda}(0) \quad (2.15a)$$

$$\boldsymbol{\mu}(t) \equiv \boldsymbol{\mu} \quad \forall t \in [0, T] \quad (2.15b)$$

Here, $\boldsymbol{\lambda}(0)$ is the initial condition (at $t = 0$) of (2.14a). From (2.14c), (2.15a) and (2.15b), we derive the optimal control law as

$$\mathbf{u}(t) = \frac{1}{2\boldsymbol{\mu}} \mathbf{v}(t-T) - \frac{T}{2\boldsymbol{\mu}} \mathbf{B}' e^{-A't} \boldsymbol{\lambda}(0) \quad (2.16)$$

By substituting (2.16) into (2.6c), we obtain $\boldsymbol{\lambda}(0)$ as

$$\boldsymbol{\lambda}(0) = e^{A'T} \mathbf{W}_c^{-1}(T) \left(\frac{1}{T} \mathbf{s}(T) - \frac{2\boldsymbol{\mu}}{T} \mathbf{r}(T) \right) \quad (2.17)$$

By substituting (2.16) and (2.17) in (2.11) and using (2.2), we obtain

$$\boldsymbol{\mu} = \pm \frac{1}{2} \sqrt{\frac{T - \mathbf{s}'(T) \mathbf{W}_c^{-1}(T) \mathbf{s}(T)}{T - \mathbf{r}'(T) \mathbf{W}_c^{-1}(T) \mathbf{r}(T)}} \quad (2.18)$$

For the existence of a solution, $\boldsymbol{\mu}$ must be a real number. Thus, either $T < \min \{ \mathbf{s}'(T) \mathbf{W}_c^{-1}(T) \mathbf{s}(T), \mathbf{r}'(T) \mathbf{W}_c^{-1}(T) \mathbf{r}(T) \}$ or $T > \max \{ \mathbf{s}'(T) \mathbf{W}_c^{-1}(T) \mathbf{s}(T), \mathbf{r}'(T) \mathbf{W}_c^{-1}(T) \mathbf{r}(T) \}$. Now it follows directly from Remark 4 that the total energy T must satisfy (2.7) for the existence of a solution i.e. $T > \max \{ \mathbf{s}'(T) \mathbf{W}_c^{-1}(T) \mathbf{s}(T), \mathbf{r}'(T) \mathbf{W}_c^{-1}(T) \mathbf{r}(T) \}$.

2.3.2 Uniqueness of the Minimally Novel Input

Lemma 2 Under the hypothesis of Lemma 1, the solution of the non-convex optimization problem (2.6) is unique.

Proof 2 By substituting (2.16) and (2.17) in (2.3), we obtain the optimal value of $\mathbb{J}(T)$ as a function of μ as

$$\mathbb{J}(T) = \frac{1}{T} \mathbf{s}'(T) \mathbf{W}_c^{-1}(T) \mathbf{r}(T) + \frac{1}{2\mu} \left(1 - \frac{1}{T} \mathbf{s}'(T) \mathbf{W}_c^{-1}(T) \mathbf{s}(T)\right) \quad (2.19)$$

It follows from Lemma (1) that $\frac{1}{T} \mathbf{s}'(T) \mathbf{W}_c^{-1}(T) \mathbf{s}(T) \in (0, 1)$ (see (2.7)). Thus, the maximum of $\mathbb{J}(T)$ occurs when $\mu > 0$ in (2.18) i.e.

$$\mu = \frac{1}{2} \sqrt{\frac{T - \mathbf{s}'(T) \mathbf{W}_c^{-1}(T) \mathbf{s}(T)}{T - \mathbf{r}'(T) \mathbf{W}_c^{-1}(T) \mathbf{r}(T)}} \quad (2.20)$$

Thus, a unique optimal control input $\mathbf{u}(t)$ exists and is given by

$$\mathbf{u}(t) = \frac{1}{2\mu} (\mathbf{v}(t-T) - \mathbf{B}' e^{A'(T-t)} \mathbf{W}_c^{-1}(T) \mathbf{s}(T)) + \mathbf{B}' e^{A'(T-t)} \mathbf{W}_c^{-1}(T) \mathbf{r}(T) \quad (2.21)$$

2.3.3 Euclidean - Inner Product Equivalence

As noted in Remark 3, it is an interesting and notable consequence of our cost formulation that the problem can exactly recast in terms of a Euclidean norm. Specifically, if we consider

$$\mathbb{J}_1(T) = \frac{1}{T} \int_0^T \|\mathbf{v}(t-T) - \mathbf{u}(t)\|_2^2 dt \quad (2.22)$$

as the cost function in (2.6a), we obtain the optimal solution as

$$\mu = -1 + \sqrt{\frac{T - \mathbf{s}'(T)\mathbf{W}_c^{-1}(T)\mathbf{s}(T)}{T - \mathbf{r}'(T)\mathbf{W}_c^{-1}(T)\mathbf{r}(T)}} \quad (2.23a)$$

$$\mathbf{u}(t) = \frac{1}{1+\mu} (\mathbf{v}(t-T) - \mathbf{B}' e^{A'(T-t)} \mathbf{W}_c^{-1}(T) \mathbf{s}(T)) + \mathbf{B}' e^{A'(T-t)} \mathbf{W}_c^{-1}(T) \mathbf{r}(T) \quad (2.23b)$$

$$\mathbb{J}_1(T) = 2\left(1 - \frac{1}{T} \mathbf{s}'(T) \mathbf{W}_c^{-1}(T) \mathbf{r}(T)\right) + \frac{2}{1+\mu} \left(\frac{1}{T} \mathbf{s}'(T) \mathbf{W}_c^{-1}(T) \mathbf{s}(T) - 1\right) \quad (2.23c)$$

It is evident that the control law (2.20)-(2.21) is same as the control law (2.23a)-(2.23b), as one expects from Remark 3.

2.4 Example

We consider a recurrent network of n neurons with linearized firing rate dynamics of the form [32]

$$\mathcal{S} \frac{d\mathbf{x}(t)}{dt} = -\mathbf{x}(t) + \mathbf{W}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t) \quad (2.24)$$

Here, $\mathbf{x}(t) \in \mathbb{R}_+^{n \times 1}$ represents the firing rate of the neurons at time t , $\mathcal{S} \in \mathbb{R}_+^{n \times n}$ is a diagonal matrix whose diagonal elements are the (positive) time constants of the neurons, $\mathbf{W} \in \mathbb{R}^{n \times n}$ defines the interaction among neurons in the network (weight matrix), $\mathbf{B} \in \mathbb{R}_+^{n \times n}$ is the input matrix, and $\mathbf{u}(t) \in \mathbb{R}_+^{n \times 1}$ is the afferent input. Since \mathcal{S} is invertible, (2.24) can be represented in the form of (2.1) by considering $A = \mathcal{S}^{-1}(-I + \mathbf{W})$ where I is the $n \times n$ identity matrix.

For illustrative purposes, we consider a recurrent network of $n = 100$ neurons where 80 neurons are excitatory and every 5th neuron is inhibitory. We choose the time constants (in ms) of the neurons, i.e. the elements of the diagonal matrix \mathcal{S} , from a uniform distribution

$\mathcal{U}(5, 10)$. For every excitatory neuron i , we choose the connectivity weight $w_{i,j}$ (in essence, a time constant for excitation from the neuron i to j) from a uniform distribution $\mathcal{U}(0, 1)$. Similarly, for every inhibitory neuron i , we choose the connectivity weight $w_{i,j}$ (from the neuron i to j) from a uniform distribution $\mathcal{U}(-1, 0)$. We assume that $w_{i,j} = 0$ for $i = j$, i.e. neurons do not possess direct feedback. Assuming \mathbf{B} as an identity matrix, we proceed to compute the minimum directional change in inputs (i.e. minimally novel inputs) required to make a desired directional change in firing rates of neurons using (2.8)-(2.9), (2.19)-(2.21).

To complete the example, we specify $T = 3$ ms. The initial and terminal states $\mathbf{x}(0)$ and $\mathbf{x}(T)$, respectively, are specified to satisfy $\|\mathbf{x}(0)\|_2 = \|\mathbf{x}(T)\|_2 = 1$ with $\mathbf{x}(0)' \mathbf{x}(T) = \gamma$, where in this particular case we specify $\gamma = 0.7645$. The prior input $\mathbf{v}(t - T)$ is specified to be constant over the interval $t \in [0, T]$. Figure 2.3 illustrates the outcome of the example for $n = 1000$ random realizations of the system. Each red dot on the figure depicts the novelty associated with the solution to (2.18)-(2.19), i.e., the minimum novelty. Note, again, that by formulation, these inputs all have unit average energy. Each blue dot corresponds to the minimum energy solution. As a verification of our theoretical development, we note that the minimum energy solution consistently requires an injection of novelty (angular orientation) relative to the prior input and relative to the optimum.

Having thus established the minimum-novelty metric as a means of assessing network sensitivity, we proceed in the next chapter to provide a paradigm for its deployment. Specifically, we use the input-sensitivity concept and formalism to derive an expression, recursive with respect to dimension and intuitively geometric, by which efficient strategies for the classification/discrimination (especially in the presence of noise) of inputs can be designed.

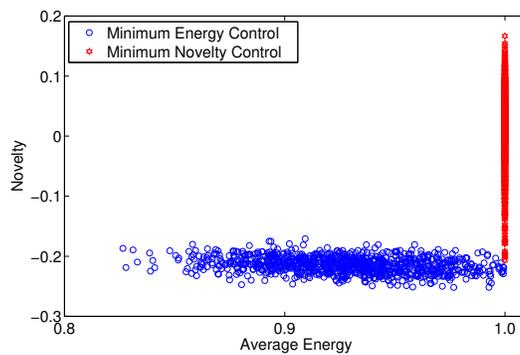


Figure 2.3: Comparison of minimum novelty control with minimum energy control for $n = 1000$ random realizations of the recurrent neuronal network: Each red dot on the figure depicts the novelty associated with the solution to the minimum novelty control. Each blue dot corresponds to the minimum energy solution.

Chapter 3

Endpoint-based Discriminability of Minimum Energy Inputs

Complex neural networks, such as those found in the human brain, are able to very accurately discriminate and classify external stimuli. Some of their topological and computational properties have been extracted and used to great effect by the artificial intelligence community. However, even our best simulated neural networks are very pale abstractions of reality, partly because (in general) they fail to account for the temporal dynamics and recurrence inherent in natural neural networks, and instead employ feed-forward architecture and discrete, simultaneous activity. In this chapter we begin to develop an intuitive, geometric framework to explore the ways in which different inputs could be discriminated in recurrent linear dynamical networks, with the eventual goal of being able to facilitate a transition to more realistic and effective artificial networks. We first establish a useful, closed-form measure on the space of minimum-energy inputs to a linear system, which allows an elucidation of how discrepancies between inputs impact output trajectories in the state space. We characterize, to an extent, the relationship between input and output difference as it relates to system

dynamics as manifest in the geometry of the reachable output space. We draw from this characterization principles which may be employed in the design of dynamic, recurrent artificial networks for input discrimination.

3.1 Introduction

Developing methods for classification is critically important for many modern engineering applications, including high-profile technologies such as voice and face recognition [33], [34]. Approaches to classification generally fall in the vast domain of machine-learning algorithms, within which one specific sub-domain involves so-called artificial neural networks [35]. Such networks implement, in essence, a series of algebraic and logical operations intended to isolate salient input features and separate them according to a pre-specified error function (usually, involving a statistical metric) [36]. While artificial neural networks mimic certain aspects of brain network architecture (namely, the feedforward layering/cascade of sub-networks), they lack the true recurrence and temporal dynamics that are intrinsic *in vivo*. Presumably, these dynamics serve some function other than the implementation of static input-output maps. Indeed, a substantial line of research in theoretical neuroscience involves understanding how brain networks use the dynamics with which they are endowed to enable the performance of discriminative tasks [37].

In this chapter, we take an initial step towards reconciling recurrent network dynamics with functions such as input discrimination. Specifically, we consider an input decoding scheme for linear systems (or, networks) wherein the identity of an input (stimulus) is inferred from the endpoint of its induced trajectory. The n -dimensional linear systems (which we assume to be controllable) under consideration have dynamics of the form

$$\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t) \tag{3.1}$$

where $\mathbf{x}(t) \in \mathbb{R}^n$ describes the state of the system, guided by the control $\mathbf{u}(t) \in \mathbb{R}^m$, at time t , and $\mathbf{A} \in \mathbb{R}^{n \times n}$ defines the dynamical relationship between state variables and $\mathbf{B} \in \mathbb{R}^{m \times n}$ describes how the m inputs are mixed to affect the n state variables.

Our central result is the derivation of a fundamental, parametric relationship between the angular distance between two competing stimuli and the separation of their induced endpoints, when each stimulus is of minimum energy for the endpoint in question. While the overall scope of our work is biologically motivated, the central results of this chapter pertain to linear dynamical systems in general.

3.2 Problem Formulation

The perception of an object must involve disambiguating the sensory inputs (e.g., sounds) related to this object from those inputs related to other objects in the environment. It is intuitive that two inputs (representing sensation of two distinct objects) must be sufficiently different in order for us to perceive them as different phenomena. However, the cognitive centers of the brain do not have access to these ‘raw’ sensory inputs but, instead, can only ‘read out’ the outputs of the networks upon which they impinge. The central question that underlies this chapter is: how ‘different’ must the inputs be so that a pre-defined decoder, operating only on the outputs, could disambiguate the underlying objects?

3.2.1 General Formulation

Fig. Said more mathematically, we consider two competing inputs $\mathbf{u}_1(t)$ and $\mathbf{u}_2(t)$ to a dynamical system, generating outputs ${}^1\mathbf{x}_1(t)$ and $\mathbf{x}_2(t)$, respectively. Conceptually, for any $\varepsilon > 0$, we seek some $\delta > 0$ such that if $\mathbf{d}_{\mathbb{J}}(\mathbf{u}_1, \mathbf{u}_2) > \delta$, then $\mathbf{d}_{\mathcal{O}}(\mathbf{x}_1(t), \mathbf{x}_2(t)) > \varepsilon$. Here, $\mathbf{d}_{\mathbb{J}}$ and

¹Note that throughout this chapter the terms output and ‘end/terminal state’ will be used interchangeably

\mathbf{d}_O define distance measures on the space of inputs and output, respectively. These measures can form the basis of a decoder that robustly (in the presence of noise) infers the input from the output.

3.2.2 Endpoint-based Input Discriminability

In this chapter, we treat the above general problem for the particular scenario of a linear system, with distance metrics defined as follows.

Definition 1 (Input Novelty) For inputs $\mathbf{u}_1(t)$ and $\mathbf{u}_2(t)$ guiding a system of the form (5.1) over the time interval $[0, T]$, we define the average novelty $\mathbf{d}_{\mathbb{J}}(\mathbf{u}_1(t), \mathbf{u}_2(t), t \in [0, T])$ —for ease of notation simplified to $\mathbf{d}_{\mathbb{J}}(\mathbf{u}_1, \mathbf{u}_2)$ —as follows:

$$\mathbf{d}_{\mathbb{J}}(\mathbf{u}_1, \mathbf{u}_2) = \frac{1}{T \sqrt{\gamma_{\mathbf{u}_1} \gamma_{\mathbf{u}_2}}} \int_0^T \mathbf{u}'_1(\tau) \mathbf{u}_2(\tau) d\tau, \quad (3.2)$$

where

$$\gamma_{\mathbf{u}_1} = \frac{1}{T} \int_0^T \mathbf{u}'_1(\tau) \mathbf{u}_1(\tau) d\tau, \quad \gamma_{\mathbf{u}_2} = \frac{1}{T} \int_0^T \mathbf{u}'_2(\tau) \mathbf{u}_2(\tau) d\tau \quad (3.3)$$

are the average energies (up to time T) of the inputs $\mathbf{u}_1(t)$ and $\mathbf{u}_2(t)$, respectively.

Of course, (4.6) is, mathematically, simply the average angle between two inputs. The notion of novelty is based on the premise that one of the inputs serves as a reference, relative to which the other could, potentially, be novel. The formulation (4.6) quantifies the extent of novelty in this sense.

Definition 2 (Endpoint Distance) Let $\mathbf{x}_{\mathbf{u}_1}$ and $\mathbf{x}_{\mathbf{u}_2}$ be the endpoints of the trajectories induced by the inputs $\mathbf{u}_1(t)$ and $\mathbf{u}_2(t)$, respectively, over the interval $[0, T]$. We define the endpoint

distance $\mathbf{d}_O(\mathbf{x}_{\mathbf{u}_1}, \mathbf{x}_{\mathbf{u}_2})$ as follows

$$\mathbf{d}_O(\mathbf{x}_{\mathbf{u}_1}, \mathbf{x}_{\mathbf{u}_2}) = \|\mathbf{x}_{\mathbf{u}_1} - \mathbf{x}_{\mathbf{u}_2}\|_2 \quad (3.4)$$

$$= \left\| \left(e^{\mathbf{A}T} \mathbf{x}(0) + \int_0^T e^{\mathbf{A}(T-\tau)} \mathbf{B} \mathbf{u}_1(\tau) d\tau \right) - \left(e^{\mathbf{A}T} \mathbf{x}(0) + \int_0^T e^{\mathbf{A}(T-\tau)} \mathbf{B} \mathbf{u}_2(\tau) d\tau \right) \right\|_2 \quad (3.5)$$

$$= \left\| \int_0^T e^{\mathbf{A}(T-\tau)} \mathbf{B} \mathbf{u}_1(\tau) d\tau - \int_0^T e^{\mathbf{A}(T-\tau)} \mathbf{B} \mathbf{u}_2(\tau) d\tau \right\|_2 \quad (3.6)$$

In this chapter we will restrict our attention to the family of minimum-energy inputs. This class of inputs is of particular interest in the context of brain networks, since one may surmise that these networks operate in (or near) an optimal energy regime [38].

We will let Ψ_c denote the class of minimum-energy inputs for a particular linear system where each input is prescribed total energy c over the interval $[0, T]$. That is, $\forall \mathbf{u}(t) \in \Psi_c, \int_0^T \mathbf{u}'(\tau) \mathbf{u}(\tau) d\tau = c$. We can then define the feasible output space at time T to be

Definition 3 (Feasible Output Space)

$$\Xi = \{\mathbf{x}(T) | \mathbf{x}'(T) \mathbf{W}^{-1} \mathbf{x}(T) = c\} \quad (3.7)$$

where $\mathbf{W} = \int_0^T e^{\mathbf{A}(T-\tau)} \mathbf{B} \mathbf{B}' e^{\mathbf{A}'(T-\tau)} d\tau$ is the system controllability gramian at time T . (3.7)

follows from the fact that the energy used by a minimum-energy control is known to be $[e^{\mathbf{A}T} \mathbf{x}(0) - \mathbf{x}'(T)]' \mathbf{W}^{-1} [e^{\mathbf{A}T} \mathbf{x}(0) - \mathbf{x}'(T)]$ (see, for instance, [39] pp. 92-93) and from the fact that we are considering $\mathbf{x}(0)$ to be the origin.

We notice that Ξ corresponds to the surface of an n -dimensional ellipsoid. The exact location of an endpoint $\mathbf{x}(T) \in \Xi$ on this ellipsoid can be described by a set of n parametric equations. We will see that this parametric description of the endpoint is advantageous for our purposes. By providing an alternative, trigonometric form for the novelty between minimum-energy inputs ($\mathbf{d}_{\mathbb{J}}$), we are able to show that the value of $\mathbf{d}_{\mathbb{J}}$ can be obtained simply by knowing the parametrization of $\mathbf{x}(T)$ on Ξ , which does not depend on the specific geometry of Ξ , even though the absolute locations of the endpoints do. Thus, our results highlight a fundamental and interesting set of constraints for the design of endpoint decoders of minimum energy inputs, which we will elaborate on in Section 3.4.

3.3 Results

We will begin by proving two Lemmas establishing sinusoidal, parametric equations for $\mathbf{d}_{\mathbb{J}}$ in two-dimensional systems, then move to 3-D, and finally to a general form for arbitrary dimensions.

For all of these, we will make use of the following preliminary Lemma:

Lemma 3 Let $\mathbf{u}_1(t), \mathbf{u}_2(t) \in \Psi_c$. Then

$$\mathbf{d}_{\mathbb{J}} = \mathbf{x}'_{\mathbf{u}_1} \mathbf{W}^{-1}(T) \mathbf{x}_{\mathbf{u}_2} \tag{3.8}$$

Proof 3 From [17], the following expression is given for the input novelty $\mathbf{d}_{\mathbb{J}}(\mathbf{u}_1, \mathbf{u}_2)$ for certain classes of inputs, one of which is Ψ_c , the class of minimum-energy inputs:

$$\mathbf{d}_{\mathbb{J}}(\mathbf{u}_1, \mathbf{u}_2) = \frac{1}{T\sqrt{\gamma_{\mathbf{u}_1}\gamma_{\mathbf{u}_2}}} * \left[(\mathbf{x}_{\mathbf{u}_1,f} - e^{AT}\mathbf{x}_{\mathbf{u}_1,0})' \mathbf{W}^{-1} (\mathbf{x}_{\mathbf{u}_2,f} - e^{AT}\mathbf{x}_{\mathbf{u}_2,0}) + \sqrt{(\gamma_{\mathbf{u}_1}T - \mathbb{E}_{\min_1})(\gamma_{\mathbf{u}_2}T - \mathbb{E}_{\min_2})} \right] \quad (3.9)$$

where $\mathbf{x}_{\mathbf{u}_1,0}, \mathbf{x}_{\mathbf{u}_1,f}, \mathbf{x}_{\mathbf{u}_2,0}$, and $\mathbf{x}_{\mathbf{u}_2,f}$ are the initial and terminal states of a linear system driven by the inputs $\mathbf{u}_1(t)$ and $\mathbf{u}_2(t)$, respectively, and where \mathbb{E}_{\min_1} and \mathbb{E}_{\min_2} are the energies of the the minimum-energy inputs associated with the state transfers from the origin to $\mathbf{x}_{\mathbf{u}_1,f}$ and $\mathbf{x}_{\mathbf{u}_2,f}$, respectively.

We may make several simplifications to (4.9). Since here we are dealing with minimum energies for both state transfers, we have that $\gamma_{\mathbf{u}_1}T = \mathbb{E}_{\min_1}$ and $\gamma_{\mathbf{u}_2}T = \mathbb{E}_{\min_2}$. Also, we may, without loss of generality, let $c = 1$ in (3.7), for any other value of c would just create a scaled version of the ellipsoid. Thus, $\frac{1}{T\sqrt{\gamma_{\mathbf{u}_1}\gamma_{\mathbf{u}_2}}} = 1$. Finally, since the point of departure for each state transfer is the origin, we have that $(\mathbf{x}_{\mathbf{u}_1,f} - e^{AT}\mathbf{x}_{\mathbf{u}_1,0}) = \mathbf{x}_{\mathbf{u}_1}$ and $(\mathbf{x}_{\mathbf{u}_2,f} - e^{AT}\mathbf{x}_{\mathbf{u}_2,0}) = \mathbf{x}_{\mathbf{u}_2}$ for desired final states $\mathbf{x}_{\mathbf{u}_1}$ and $\mathbf{x}_{\mathbf{u}_2}$. With these simplifications, we arrive at (3.8).

3.3.1 2-dimensional case

We first introduce parameters $\phi_1, \phi_2 \in [0, 2\pi]$, which describe the locations of $\mathbf{x}_{\mathbf{u}_1}$ and $\mathbf{x}_{\mathbf{u}_2}$ (at time T), respectively, on the ellipse Ξ :

$$\mathbf{x}_{\mathbf{u}_1} = \left(\frac{1}{\sqrt{\lambda_1}} \cos(\phi_1), \frac{1}{\sqrt{\lambda_2}} \sin(\phi_1) \right) \quad (3.10)$$

$$\mathbf{x}_{\mathbf{u}_2} = \left(\frac{1}{\sqrt{\lambda_1}} \cos(\phi_2), \frac{1}{\sqrt{\lambda_2}} \sin(\phi_2) \right) \quad (3.11)$$

where λ_1 and λ_2 are the eigenvalues of \mathbf{W}^{-1} . (3.10) and (3.11) follow from the general parametric equations for an ellipse, and the principle axis theorem for quadratic forms in \mathbb{R}^n together with the fact that \mathbf{W}^{-1} is symmetric. Now we are ready to state our first lemma in terms of these parameters ϕ_1 and ϕ_2 :

Lemma 4 : The average novelty $\mathbf{d}_{\mathbb{J}}(\mathbf{u}_1, \mathbf{u}_2)$ over the time interval $[0, T]$ between minimum-energy inputs $\mathbf{u}_1(t)$ and $\mathbf{u}_2(t)$ which guide a two-dimensional linear system to terminal states $\mathbf{x}_{\mathbf{u}_1}$ and $\mathbf{x}_{\mathbf{u}_2}$, respectively, is given by:

$$\mathbf{d}_{\mathbb{J}}(\mathbf{u}_1, \mathbf{u}_2) = \cos(\phi_1 - \phi_2) \quad (3.12)$$

Proof 4 Since \mathbf{W}^{-1} is symmetric, its eigenvectors (call them \mathbf{v}_1 and \mathbf{v}_2) form an orthonormal basis for \mathbb{R}^2 , we may write

$$\mathbf{x}_{\mathbf{u}_1} = \alpha_1 \mathbf{v}_1 + \beta_1 \mathbf{v}_2 \quad (3.13)$$

and

$$\mathbf{x}_{\mathbf{u}_2} = \alpha_2 \mathbf{v}_1 + \beta_2 \mathbf{v}_2 \quad (3.14)$$

for some $\alpha_1, \alpha_2, \beta_1, \beta_2 \in \mathbb{R}$, and thus, by substituting (3.13) and (3.14) into (3.8), we have

$$\mathbf{d}_{\mathbb{J}}(\mathbf{u}_1, \mathbf{u}_2) = (\alpha_1 \mathbf{v}_1 + \beta_1 \mathbf{v}_2)' \mathbf{W}^{-1} (\alpha_2 \mathbf{v}_1 + \beta_2 \mathbf{v}_2) \quad (3.15)$$

Now, since \mathbf{v}_1 and \mathbf{v}_2 are orthonormal, cross terms vanish, and this simplifies further to

$$\begin{aligned}
\mathbf{d}_{\mathbb{J}}(\mathbf{u}_1, \mathbf{u}_2) &= \alpha_1 \mathbf{v}'_1 \mathbf{W}^{-1} \alpha_2 \mathbf{v}_1 + \beta_1 \mathbf{v}'_2 \mathbf{W}^{-1} \beta_2 \mathbf{v}_2 \\
&= \alpha_1 \alpha_2 \mathbf{v}'_1 \mathbf{W}^{-1} \mathbf{v}_1 + \beta_1 \beta_2 \mathbf{v}'_2 \mathbf{W}^{-1} \mathbf{v}_2 \\
&= \alpha_1 \alpha_2 \mathbf{v}'_1 \lambda_1 \mathbf{v}_1 + \beta_1 \beta_2 \mathbf{v}'_2 \lambda_2 \mathbf{v}_2 \\
&= \alpha_1 \alpha_2 \lambda_1 + \beta_1 \beta_2 \lambda_2
\end{aligned}$$

If we, without loss of generality, orient the ellipse with the cardinal axes (that is, $\mathbf{v}_1 = \mathbf{e}_1$ and $\mathbf{v}_2 = \mathbf{e}_2$), we have that $\alpha_1 = \frac{1}{\sqrt{\lambda_1}} \cos(\phi_1)$ and $\alpha_2 = \frac{1}{\sqrt{\lambda_1}} \cos(\phi_2)$, since α_1 and α_2 are simply the lengths of the projections of $\mathbf{x}_{\mathbf{u}_1}$ and $\mathbf{x}_{\mathbf{u}_2}$, respectively, onto $\mathbf{v}_1 = \mathbf{e}_1$. Likewise, we may write $\beta_1 = \frac{1}{\sqrt{\lambda_2}} \sin(\phi_1)$ and $\beta_2 = \frac{1}{\sqrt{\lambda_2}} \sin(\phi_2)$. Thus we have:

$$\begin{aligned}
\mathbf{d}_{\mathbb{J}} &= \alpha_1 \alpha_2 \lambda_1 + \beta_1 \beta_2 \lambda_2 \\
&= \frac{1}{\sqrt{\lambda_1}} \cos(\phi_1) \frac{1}{\sqrt{\lambda_1}} \cos(\phi_2) \lambda_1 \\
&\quad + \frac{1}{\sqrt{\lambda_2}} \sin(\phi_1) \frac{1}{\sqrt{\lambda_2}} \sin(\phi_2) \lambda_2 \\
&= \cos(\phi_1) \cos(\phi_2) + \sin(\phi_1) \sin(\phi_2) \\
&= \cos(\phi_1 - \phi_2)
\end{aligned}$$

where the last equality is the result of a basic trigonometric identity.

Remark 5 This result is surprising for two reasons. First, it is independent of the specific system dynamics. Second, it depends on the parameters ϕ_1 and ϕ_2 , which describe the endpoint locations.

3.3.2 3-dimensional case

In addition to the ϕ_1 and ϕ_2 previously introduced we need to introduce a second parameter to describe the output locations on Ξ . To simplify notation, especially when we turn to the n -dimensional case, we will begin to employ double indexing to the parameters ϕ and the coefficients α , so that $\phi_{i,j}$ refers to the i^{th} parameter describing the location of $\mathbf{x}_{\mathbf{u}_j}$ on Ξ , and $\alpha_{i,j}$ refers to the i^{th} coefficient of the linear combination of eigenvectors describing output $\mathbf{x}_{\mathbf{u}_j}$ under input \mathbf{u}_j . Thus we have

$$\mathbf{x}_{\mathbf{u}_1} = \left(\frac{1}{\sqrt{\lambda_1}} \cos(\phi_{1,1}) \sin(\phi_{2,1}), \frac{1}{\sqrt{\lambda_2}} \sin(\phi_{1,1}) \sin(\phi_{2,1}), \frac{1}{\sqrt{\lambda_3}} \cos(\phi_{2,1}) \right) \quad (3.16)$$

and

$$\mathbf{x}_{\mathbf{u}_2} = \left(\frac{1}{\sqrt{\lambda_1}} \cos(\phi_{1,2}) \sin(\phi_{2,2}), \frac{1}{\sqrt{\lambda_2}} \sin(\phi_{1,2}) \sin(\phi_{2,2}), \frac{1}{\sqrt{\lambda_3}} \cos(\phi_{2,2}) \right) \quad (3.17)$$

where the λ_i are again the eigenvalues of \mathbf{W}^{-1} . Using these parametric equations, we can again derive a sinusoidal formula for the input novelty.

Lemma 5 The average novelty $\mathbf{d}_{\mathbb{J}}(\mathbf{u}_1, \mathbf{u}_2)$ over the time interval $[0, T]$ between minimum-energy inputs $\mathbf{u}_1(t)$ and $\mathbf{u}_2(t)$ which guide a three-dimensional linear system to terminal states $\mathbf{x}_{\mathbf{u}_1}$ and $\mathbf{x}_{\mathbf{u}_2}$, respectively, is given by:

$$\mathbf{d}_{\mathbb{J}}(\mathbf{u}_1, \mathbf{u}_2) = \cos(\phi_{2,1}) \cos(\phi_{2,2}) + \cos(\phi_{1,1} - \phi_{1,2}) \sin(\phi_{2,1}) \sin(\phi_{2,2}) \quad (3.18)$$

Proof 5 For the 3-D case, we generalize (3.13) and (3.14) to obtain

$$\mathbf{x}_{\mathbf{u}_1} = \alpha_{1,1} \mathbf{v}_1 + \alpha_{2,1} \mathbf{v}_2 + \alpha_{3,1} \mathbf{v}_3 \quad (3.19)$$

and

$$\mathbf{x}_{\mathbf{u}_2} = \alpha_{1,2}\mathbf{v}_1 + \alpha_{2,2}\mathbf{v}_2 + \alpha_{3,2}\mathbf{v}_3 \quad (3.20)$$

By substituting (3.19) and (3.20) into (3.8), we have

$$\begin{aligned} \mathbf{d}_{\mathbb{J}}(\mathbf{u}_1, \mathbf{u}_2) &= (\alpha_{1,1}\mathbf{v}_1 + \alpha_{2,1}\mathbf{v}_2 + \alpha_{3,1}\mathbf{v}_3)' \mathbf{W}^{-1} (\alpha_{1,2}\mathbf{v}_1 + \alpha_{2,2}\mathbf{v}_2 + \alpha_{3,2}\mathbf{v}_3) \\ &= \alpha_{1,1}\alpha_{1,2}\lambda_1 + \alpha_{2,1}\alpha_{2,2}\lambda_2 + \alpha_{3,1}\alpha_{3,2}\lambda_3 \end{aligned} \quad (3.21)$$

again since the eigenvectors are orthonormal. Again, without loss of generality, we may orient the ellipsoid (\mathfrak{E}) with the cardinal axes. Then we may write $\alpha_{1,1} = \frac{1}{\sqrt{\lambda_1}}\cos(\phi_{1,1})\sin(\phi_{2,1})$, $\alpha_{2,1} = \frac{1}{\sqrt{\lambda_2}}\sin(\phi_{1,1})\sin(\phi_{2,1})$, $\alpha_{3,1} = \frac{1}{\sqrt{\lambda_3}}\cos(\phi_{2,1})$, and likewise parameterize $\alpha_{1,2}$, $\alpha_{2,2}$, and $\alpha_{3,2}$. Thus (3.21) becomes

$$\begin{aligned} \mathbf{d}_{\mathbb{J}}(\mathbf{u}_1, \mathbf{u}_2) &= \frac{1}{\sqrt{\lambda_1}}\cos(\phi_{1,1})\sin(\phi_{2,1})\frac{1}{\sqrt{\lambda_1}}\cos(\phi_{1,2})\sin(\phi_{2,2})\lambda_1 \\ &\quad + \frac{1}{\sqrt{\lambda_2}}\sin(\phi_{1,1})\sin(\phi_{2,1})\frac{1}{\sqrt{\lambda_2}}\sin(\phi_{1,2})\sin(\phi_{2,2})\lambda_2 \\ &\quad + \frac{1}{\sqrt{\lambda_3}}\cos(\phi_{2,1})\frac{1}{\sqrt{\lambda_3}}\cos(\phi_{2,2})\lambda_3 \\ &= \cos(\phi_{1,1})\cos(\phi_{1,2})\sin(\phi_{2,1})\sin(\phi_{2,2}) \\ &\quad + \sin(\phi_{1,1})\sin(\phi_{1,2})\sin(\phi_{2,1})\sin(\phi_{2,2}) \\ &\quad + \cos(\phi_{2,1})\cos(\phi_{2,2}) \\ &= \cos(\phi_{2,1})\cos(\phi_{2,2}) \\ &\quad + \cos(\phi_{1,1} - \phi_{1,2})\sin(\phi_{2,1})\sin(\phi_{2,2}) \end{aligned}$$

Remark 6 We notice that along with being a sinusoidal function of the output parameters $\phi_{i,j}$, (3.18) is in fact a recursive operation on (3.12). That is,

$$\mathbf{d}_{\mathbb{J}_3}(\mathbf{u}_1, \mathbf{u}_2) = \cos(\phi_{2,1})\cos(\phi_{2,2}) + \mathbf{d}_{\mathbb{J}_2}(\mathbf{u}_1, \mathbf{u}_2)\sin(\phi_{2,1})\sin(\phi_{2,2}) \quad (3.22)$$

3.3.3 The n -dimensional case

For an arbitrary n -dimensional (ellipsoidal) output space \mathfrak{E} , we can parameterize the ellipsoid with $n - 1$ parameters to obtain an analogous closed-form, sinusoidal, recursive expression for the minimum-energy input novelty $\mathbf{d}_{\mathbb{J}}$. Here we make use of the fact that in arbitrary number of dimensions n , we may parameterize the $(n - 1)$ -ellipsoid by the spherical coordinates

$$\begin{aligned} \alpha_n &= \frac{1}{\sqrt{\lambda_1}}\cos(\phi_{n-1}) \\ \alpha_{n-1} &= \frac{1}{\sqrt{\lambda_2}}\sin(\phi_{n-1})\cos(\phi_{n-2}) \\ \alpha_{n-2} &= \frac{1}{\sqrt{\lambda_3}}\sin(\phi_{n-1})\sin(\phi_{n-2})\cos(\phi_{n-3}) \\ &\cdot \\ &\cdot \\ &\cdot \\ \alpha_2 &= \frac{1}{\sqrt{\lambda_{n-1}}}\sin(\phi_{n-1})\sin(\phi_{n-2})\dots\sin(\phi_2)\cos(\phi_1) \\ \alpha_1 &= \frac{1}{\sqrt{\lambda_n}}\sin(\phi_{n-1})\sin(\phi_{n-2})\dots\sin(\phi_2)\sin(\phi_1) \end{aligned}$$

where $\phi_2, \phi_2, \dots, \phi_{n-1}$ are allowed to range over $[0, \pi]$, while ϕ_1 ranges over $[0, 2\pi]$. We can follow the basic development of the 2-D and 3-D cases (i.e. we add an $\alpha_{n,1}\alpha_{n,2}\lambda_n$ term to the formula for $\mathbf{d}_{\mathbb{J}}$ in $n - 1$ dimensions; see (3.21)). Note that, since each $\alpha_{i,(1 \text{ or } 2)}$ has a

coefficient of $\frac{1}{\sqrt{\lambda_i}}$, λ_i is cancelled for each term in the summation.

We saw in (3.22) that the expression (3.12) for $\mathbf{d}_{\mathbb{J}_2}$ appears in the expression (3.18) for $\mathbf{d}_{\mathbb{J}_3}$.

We now generalize this pattern to arrive at the following:

Theorem 1 The novelty $\mathbf{d}_{\mathbb{J}_n}(\mathbf{u}_1, \mathbf{u}_2)$ between two minimum-energy inputs ($\mathbf{u}_1(t)$ and $\mathbf{u}_2(t)$) guiding the system to end states ($\mathbf{x}_{\mathbf{u}_1}$ and $\mathbf{x}_{\mathbf{u}_2}$) on the unit-energy gramian ellipsoid in n dimensions ($n > 2$) is given by

$$\mathbf{d}_{\mathbb{J}_n} = \sin(\phi_{n-1,1})\sin(\phi_{n-1,2})\mathbf{d}_{\mathbb{J}_{n-1}} + \cos(\phi_{n-1,1})\cos(\phi_{n-1,2}) \quad (3.23)$$

where $\mathbf{d}_{\mathbb{J}_2} = \cos(\phi_{1,1} - \phi_{1,2}) = \sin(\phi_{1,1})\sin(\phi_{1,2}) + \cos(\phi_{1,1})\cos(\phi_{1,2})$

Proof 6 The proof is by strong induction, where the base case has been established under Lemma 1 and Lemma 2. For the inductive step we assume that

$$\begin{aligned} \mathbf{d}_{\mathbb{J}_i} &= \sin(\phi_{i-1,1})\sin(\phi_{i-1,2})\mathbf{d}_{\mathbb{J}_{i-1}} + \cos(\phi_{i-1,1})\cos(\phi_{i-1,2}) \\ &\quad \forall i \in \{3, \dots, n-1\} \end{aligned}$$

Now we expand

$$\begin{aligned}
& \sin(\phi_{n-1,1})\sin(\phi_{n-1,2})\mathbf{d}_{\mathbb{J}_{n-1}} + \cos(\phi_{n-1,1})\cos(\phi_{n-1,2}) \\
&= \sin(\phi_{n-1,1})\sin(\phi_{n-1,2}) \left[\sin(\phi_{n-2,1})\sin(\phi_{n-2,2})\mathbf{d}_{\mathbb{J}_{n-2}} \right. \\
&\quad \left. + \cos(\phi_{n-2,1})\cos(\phi_{n-2,2}) \right] + \cos(\phi_{n-1,1})\cos(\phi_{n-1,2}) \\
&= \sin(\phi_{n-1,1})\sin(\phi_{n-1,2})\sin(\phi_{n-2,1})\sin(\phi_{n-2,2})\mathbf{d}_{\mathbb{J}_{n-2}} \\
&\quad + \sin(\phi_{n-1,1})\sin(\phi_{n-1,2})\cos(\phi_{n-2,1})\cos(\phi_{n-2,2}) \\
&\quad + \cos(\phi_{n-1,1})\cos(\phi_{n-1,2}) \\
&= \sin(\phi_{n-1,1})\sin(\phi_{n-1,2})\sin(\phi_{n-2,1})\sin(\phi_{n-2,2})\mathbf{d}_{\mathbb{J}_{n-2}} \\
&\quad + \alpha_{n-1,1}\alpha_{n-1,2}\lambda_2 + \alpha_{n,1}\alpha_{n,2}\lambda_1 \\
&= \sin(\phi_{n-1,1})\sin(\phi_{n-1,2})\sin(\phi_{n-2,1})\sin(\phi_{n-2,2}) \\
&\quad * \left[\sin(\phi_{n-3,1})\sin(\phi_{n-3,2})\mathbf{d}_{\mathbb{J}_{n-3}} + \cos(\phi_{n-3,1})\cos(\phi_{n-3,2}) \right] \\
&\quad + \alpha_{n-1,1}\alpha_{n-1,2}\lambda_2 + \alpha_{n,1}\alpha_{n,2}\lambda_1 \\
&= \sin(\phi_{n-1,1})\sin(\phi_{n-1,2})\sin(\phi_{n-2,1})\sin(\phi_{n-2,2}) \\
&\quad * \sin(\phi_{n-3,1})\sin(\phi_{n-3,2})\mathbf{d}_{\mathbb{J}_{n-3}} + \alpha_{n-2,1}\alpha_{n-2,2}\lambda_3 \\
&\quad + \alpha_{n-1,1}\alpha_{n-1,2}\lambda_2 + \alpha_{n,1}\alpha_{n,2}\lambda_1 \\
&\quad \cdot \\
&\quad \cdot \\
&\quad \cdot \\
&= \sin(\phi_{n-1,1})\sin(\phi_{n-1,2})\sin(\phi_{n-2,1})\sin(\phi_{n-2,2}) \\
&\quad \dots \sin(\phi_{2,1})\sin(\phi_{2,2})\mathbf{d}_{\mathbb{J}_2} + \sum_{i=3}^n \alpha_{i,1}\alpha_{i,2}\lambda_{n-i+3} \\
&= \sin(\phi_{n-1,1})\sin(\phi_{n-1,2})\sin(\phi_{n-2,1})\sin(\phi_{n-2,2}) \\
&\quad \dots \sin(\phi_{2,1})\sin(\phi_{2,2}) * \left[\sin(\phi_{1,1})\sin(\phi_{1,2}) \right. \\
&\quad \left. + \cos(\phi_{1,1})\cos(\phi_{1,2}) \right] + \sum_{i=3}^n \alpha_{i,1}\alpha_{i,2}\lambda_{n-i+3} \\
&= \sum_{i=1}^n \alpha_{n,1}\alpha_{n,2}\lambda_{n-i+1} = \mathbf{d}_{\mathbb{J}_n} \quad 42
\end{aligned}$$

Thus we have an expression for $\mathbf{d}_{\mathbb{J}_n}$ which is recursive, sinusoidal, and depends only on the output parameters $\phi_{i,j}$.

Remark 7 Since $\mathbf{d}_{\mathbb{J}_n}$ is a composition of sinusoids, we know that $-1 \leq \mathbf{d}_{\mathbb{J}_n}(\mathbf{u}_1, \mathbf{u}_2) \leq 1 \forall \mathbf{u}_1, \mathbf{u}_2 \in \Psi_c$. In fact, these extremal values can always be reached, for if $\mathbf{u}_1(t) = \pm \mathbf{u}_2(t) \forall t \in [0, T]$, then (4.6) is clearly ± 1 . To see this from (3.23), consider that if $\mathbf{u}_1 = \pm \mathbf{u}_2$, then $\mathbf{x}_{\mathbf{u}_1} = \int_0^T e^{\mathbf{A}(T-\tau)} \mathbf{B}(\pm \mathbf{u}_2)(\tau) d\tau = \pm \mathbf{x}_{\mathbf{u}_2}$. If $\mathbf{x}_{\mathbf{u}_1} = \mathbf{x}_{\mathbf{u}_2}$, then $\phi_{i,1} = \phi_{i,2} \forall i \in \{1, \dots, n\}$ and therefore (3.23) equals 1 ((3.12) with an easy inductive argument shows this). If $\mathbf{x}_{\mathbf{u}_1} = -\mathbf{x}_{\mathbf{u}_2}$, then $\phi_{1,1} = \phi_{1,2} + \pi$ and $\phi_{i,1} + \phi_{i,2} = \pi \forall i \in \{1, \dots, n-1\}$. Then again by an easy induction, together with the product-to-sum trigonometric identities, we have that (3.23) equals -1.

3.4 Implications for Endpoint Decoding Schemes

The goal of a discriminative artificial neural network is to separate inputs that are different, and concentrate inputs that are similar. In this section we will give principles by which our result, at a strictly conceptual level, could be used to this end in two different noise scenarios. For both scenarios, we make use of how the geometry of the ellipsoid Ξ , which determines the parametrization of the endpoints, affects the relationship between the positions of these endpoints and the input novelty (see Fig. 3.1).

3.4.1 Robust discrimination with regard to measurement noise

If we assume ε noise in the output measurements (i.e. $\|\mathbf{x}_{\mathbf{u}_i, \text{observed}} - \mathbf{x}_{\mathbf{u}_i, \text{actual}}\| < \varepsilon$, then we will want to ensure that the inputs we aim to discriminate lead to actual outputs whose ε -balls do not intersect (see Fig. 6.8). In order to maximize feasible ε values which will not lead to overlapping (and thus the possibility of misclassification), we need to intelligently position the desired output states with respect to the relationship between parametric separation

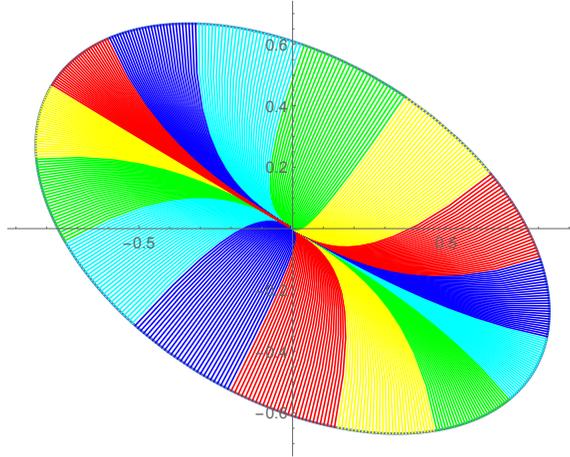


Figure 3.1: Trajectories under minimum-energy inputs are not uniformly distributed over the ellipse \mathfrak{E} . Here, color bands represent an equal number of inputs, and all inputs \mathbf{u}_i $i \in \{1, \dots, m\}$ are uniformly distributed with respect to the novelty measure $\mathbf{d}_{\mathbb{J}}$ (i.e. $\mathbf{d}_{\mathbb{J}}(\mathbf{u}_i, \mathbf{u}_{i+k}) = \mathbf{d}_{\mathbb{J}}(\mathbf{u}_j, \mathbf{u}_{j+k}) \forall i, j \in \{1, \dots, m-1\}, \forall k \in \{1, \dots, m - \max\{i, j\}\}$

and Euclidean distance. From our result we know that equally distant pairs of inputs (according to $\mathbf{d}_{\mathbb{J}}$) will create equally separated pairs outputs with respect to the parameters ϕ_i , $i \in \{1, \dots, n-1\}$, but not necessarily with respect to Euclidean distance.

For instance, examining Fig. 6.8 we can see that the endpoint distance will depend not only on the input novelty, but also on the general shape of the gramian ellipsoid \mathfrak{E} , and on the location of the endpoints on \mathfrak{E} . Under some inputs \mathbf{u}_1 and \mathbf{u}_2 leading to outputs $\mathbf{x}_{\mathbf{u}_1}$ and $\mathbf{x}_{\mathbf{u}_2}$ ending up at the ‘tips’ of a highly eccentric ellipsoid (this corresponds to moving in an ‘easy’ eigendirection), and having some novelty $\mathbf{d}_{\mathbb{J}}(\mathbf{u}_1, \mathbf{u}_2) = c$, we may indeed find that $\|\mathbf{x}_{\mathbf{u}_1}(T) - \mathbf{x}_{\mathbf{u}_2}(T)\| < \varepsilon$, while another pair of inputs \mathbf{u}_3 and \mathbf{u}_4 , leading to outputs $\mathbf{x}_{\mathbf{u}_3}$ and $\mathbf{x}_{\mathbf{u}_4}$ on the ‘flat’ part of the eccentric ellipsoid (corresponding to movement in a ‘hard’ eigendirection), and having the same novelty ($\mathbf{d}_{\mathbb{J}}(\mathbf{u}_3, \mathbf{u}_4) = c$) may result in $\|\mathbf{x}_{\mathbf{u}_3}(T) - \mathbf{x}_{\mathbf{u}_4}(T)\| \gg \varepsilon$. In such a case, we will want to distribute our desired output states on the controllability gramian such that ‘nearby’ inputs from different input classes are mapped to the flat parts of the ellipsoid (since then smaller input novelties will lead to relatively greater Euclidean output

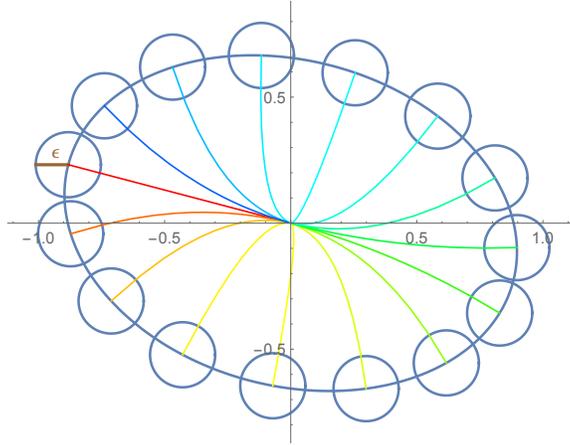


Figure 3.2: ε -balls around final states generated by equally-spaced (in terms of novelty measure $\mathbf{d}_{\mathbb{J}}$) inputs

distance, as shown below), while classes with relatively different inputs can be mapped to the tips of the ellipsoid, since the large input novelty should prevent overlapping of the reachable output regions (i.e. the ε -balls).

3.4.2 Robust discrimination with respect to input noise

Thus far we have used the standard Euclidean metric on the output space, and thus we have spoken in terms of ε -balls, but there is a subtlety that should not be overlooked. If we assume that noise corrupts the inputs, then the ε -balls become self-similar ellipsoids (see Fig. 3.3). Assume $w(t)$, $t \in [0, T]$ is a noise vector which acts on the system. Then the dynamics become $\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}(\mathbf{u} + \mathbf{w})$, and the state of \mathbf{x} at time T becomes

$$\begin{aligned} \mathbf{x}(T) &= e^{\mathbf{A}T} \mathbf{x}(0) + \int_0^T e^{\mathbf{A}(T-\tau)} \mathbf{B}(\mathbf{u}(\tau) + \mathbf{w}(\tau)) d\tau \\ &= \int_0^T e^{\mathbf{A}(T-\tau)} \mathbf{B}\mathbf{u}(\tau) d\tau + \int_0^T e^{\mathbf{A}(T-\tau)} \mathbf{B}\mathbf{w}(\tau) d\tau \end{aligned}$$

where we recall that we are taking $\mathbf{x}(0)$ as the origin. Since, then, the outputs are additive, the trajectories resulting from noisy inputs will wind up within the perimeter of the smaller ellipsoids (again, see Fig. 3.3). We assume that the energy of the input u is greater than that of the noise w ; that is, $\mathbf{u}'(t)\mathbf{u}(t) \gg w'(t)w(t) \forall t \in [0, T]$. We say ‘within’ the perimeter because a noise input $\mathbf{w}(t)$ will in general not be the minimum-energy input for its total prescribed energy. Under conditions of input noise, then, an ε -ball based choice of desired output states will not be sufficient, and we will need to consider the distribution of endpoints with respect to these smaller ellipsoids, which are self-similar to Ξ .

This can simplify the problem, since an optimal packing of these self-similar ellipsoids on the larger one will be a function of the parameters. In fact, in many 2-D simulations (unfortunately, we have yet to prove this), we have found that an optimal packing of self-similar ellipses onto a larger copy is given when we simply take equal divisions of the ellipse parameter ϕ . That is, if we want to pack n ellipses onto one which is self-similar and larger, we can simply take as increment $k = \frac{2\pi}{n}$ and position the centers of the smaller ellipses at $(a \cos(k * i), b \sin(k * i))$, $i \in \{1, \dots, n\}$. Such a packing is shown in Fig. 3.3. Of course the size of these ellipses, if they are not to overlap, will vary inversely with n .

We have found a very good heuristic: letting a denote the semi-major axis length of the larger ellipse, and a' denote that of the smaller ellipse, we set $n = \lfloor \frac{a}{a'} \pi \rfloor$. Again, we leave this unproven, but we remark that it is strikingly sharp across a wide range of ellipses eccentricities and ratios $\frac{a}{a'}$. Of course an optimal ellipse packing in 2-D is a mathematical curiosity unless it can be generalized to higher dimensional ellipsoids. This, unfortunately, is elusive [40]. However, the principle remains that effective parametric placement of the endpoints, for robust discrimination in the presence of noisy inputs, will reduce to an ellipsoid-packing problem.

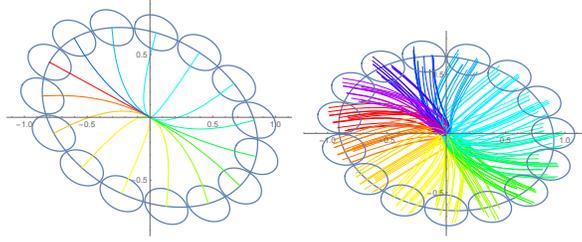


Figure 3.3: Self-similar ellipses around terminal points. Left shows trajectories of $\mathbf{x}_{\mathbf{u}_i}$ under minimum-energy input \mathbf{u}_i . Right shows $\mathbf{x}_{\mathbf{u}_i}$ under $\mathbf{u}_i + \mathbf{w}_j$, where \mathbf{w}_j is noise (here constant for purpose of visualization).

3.5 Conclusions and Future Work

We have shown that $\mathbf{d}_{\mathbb{J}}$ is a measure on the space of minimum-energy inputs to a linear system which is invariant to system dynamics, and that this measure, coupled with knowledge of the controllability gramian, can be of help in understanding the discriminative ability of such systems and in designing algorithms to discriminate between discrete classes of inputs. We have highlighted some principles, based on the relationship between $\mathbf{d}_{\mathbb{J}}$ and the Euclidean distance of outputs, which we may use in the selection of desired target states for artificial neural networks. In so doing we believe that we have introduced a potentially useful conceptual framework in which to discuss the geometry of dynamical classification systems, as well as provided part of a mathematical basis for its quantification.

In the future, we would like to expand these ideas to include the whole state trajectory in the discriminative task instead of just the terminal state. Also, we would like to generalize the analysis to include other than minimum-energy inputs, with the realization that this analysis might not admit of such elegant and simple results. Finally, we look forward to the actual design and implementation of optimization-based methods for exploiting these ideas to create truly dynamical artificial neural networks.

In this chapter, we have explored network orientation sensitivity to minimum-energy inputs. In the next chapter, we will expand our inquiry to include inputs with ‘excess’ energy. It is intuitive that inputs which are ‘allowed’ arbitrary energy will be able to traverse a very wide range of possible trajectories and still terminate at some desired target. Thus, two inputs with arbitrary energy can be very similar along almost the whole length of their trajectories, perhaps deviating at the very end with a massive energetic push to reach their targets. In this extreme the network would still be able to discriminate them (based on their endpoints). We have seen that this is absolutely not the case when inputs are allowed only minimum energy, thus prescribing the traversal of exactly one possible trajectory. We seek, in the next chapter, to quantify the way in which different networks make use of excess energy (for enhanced orientation sensitivity) between these two extremes—i.e. when inputs have excess energy, but not arbitrary energy.

Chapter 4

Bispectral Analysis for Measuring Energy-Orientation Tradeoffs in the Control of Linear Systems

Many characterizations of linear system controllability revolve around the eigenvalue spectrum of the controllability gramian, which is a function of the network dynamics. The gramian spectrum describes the minimum energies associated with inducing movement along orthogonal directions in \mathbb{R}^n . Here, we derive an enhanced interpretation of the spectral properties of the gramian in non-minimum energy regimes. Indeed, in a non-minimum energetic regime, an ‘excess’ of energy is available to the system for at least $(n - 1)$ orthogonal state transfers. We show that the utility of this excess energy can be quantified in terms of input orientation, or, simply, the angle between two competing inputs. Based on this notion, we derive the gramian bispectrum, which describes the relationship between energy and orientation among pairs of orthogonal state transfers. The bispectrum reflects a fundamental tradeoff between the

energetic and orientation costs in the control of a linear system. We show how this bispectral analysis can provide control characterizations that are not apparent from inspection of the gramian spectrum alone.

4.1 Introduction

4.1.1 Motivation

A long-standing topic in control theory, with recent applications in network control [11–13], involves the development of analyses to quantify the controllability of linear systems [8–10]. Many of the approaches that have been developed to address this issue involve study of the spectral properties of the controllability gramian, which, for a controllable linear time-invariant system in the typical form

$$\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t) \tag{4.1}$$

is defined as

$$\mathbf{W}(T) = \int_0^T \mathbf{e}^{\mathbf{A}(T-\tau)} \mathbf{B}(\tau) \mathbf{B}'(\tau) \mathbf{e}^{\mathbf{A}'(T-\tau)} \mathbf{d}\tau. \tag{4.2}$$

Each eigenvalue of \mathbf{W}^{-1} determines the energy required to induce motion in the direction of its associated eigenvector. Thus, summary metrics that describe the energetic costs of controlling a linear system can be derived from the gramian spectrum. Such metrics include the trace of the gramian [26], [25], the log of its determinant [26], and the maximum eigenvalue of the inverse gramian [26], [25]. The latter, in particular, is the ‘worst-case’ minimum energy required to reach the unit hypersphere at a prescribed time T .

In this chapter we seek to characterize the controllability of a system assuming that we have at least this minimum energy available to meet control objectives. That is, we assume that

the entire unit hypersphere is reachable at time T . What then, can be said about the system under consideration? For a non-uniform gramian spectrum, there will be a gradation of energetic costs, so that we may steer the system in other directions (than the ‘worst-case’ direction) with less (or much less) energy expenditure. To what use, then, is the excess energy which we have available? Is there a measure of control flexibility that can be realized under this ‘surplus energy’ scenario?

To answer these questions we consider, in addition to energy, the relative orientation between two inputs $\mathbf{u}_1(t)$ and $\mathbf{u}_2(t)$, which transfer the state of the system to two different endpoints at time T . Assuming these inputs possess average energy $\gamma_{\mathbf{u}_1}$ and $\gamma_{\mathbf{u}_2}$, then over the time horizon $[0, T]$ the relative orientation, or, mathematically, the average cosine of the angle between $\mathbf{u}_1(t)$ and $\mathbf{u}_2(t)$, is:

$$\frac{1}{T\sqrt{\gamma_{\mathbf{u}_1}\gamma_{\mathbf{u}_2}}}\int_0^T \mathbf{u}'_1(\tau)\mathbf{u}_2(\tau)d\tau, \quad (4.3)$$

If $\gamma_{\mathbf{u}_1}$ and $\gamma_{\mathbf{u}_2}$ are minimal for the transfers in question, then (4.3) is fixed since $\mathbf{u}_1(t)$ and $\mathbf{u}_2(t)$ can assume only one form. Given excess energy, inputs achieving the transfers may assume a range of possible relative orientations. We propose that this range constitutes a measure of the aforementioned control flexibility, and we provide an analytical development that makes this notion concrete. Specifically, in this chapter we make the following contributions:

1. As a function of available energy, we derive the minimum relative orientation (or, average angle) between two inputs, each inducing a transfer to points along a different eigenvector of \mathbf{W}^{-1} . A small relative orientation implies that the inputs are relatively flexible in their orientation range. Indeed, we show that as energy tends to infinity the angular difference tends to zero, so that the inputs may become arbitrarily similar in their geometry. The derivation leverages our previous optimal control results in [17] and [18].

2. We derive a second-order spectrum, termed the gramian bispectrum that exactly quantifies the tradeoff between energy and orientation, in the above sense, for each pair of eigenvalues of the gramian.
3. We demonstrate how the proposed bispectrum may be used to compare the relative utility of energy in linear systems. We provide examples of such comparisons that show how a bispectral analysis may reveal control properties not apparent from the spectrum alone.

4.2 Problem Formulation

4.2.1 Geometric Interpretation of the Controllability Gramian

Our subsequent development will pivot on the spectrum Λ of the controllability gramian inverse \mathbf{W}^{-1} (see (4.2)), defined as the collection of its eigenvalues with ordered labeling, i.e.,

$$\Lambda = \{\lambda_1, \dots, \lambda_n\}, \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n \quad (4.4)$$

and corresponding eigenvectors $\mathbf{v}_1, \dots, \mathbf{v}_n$.

Since the minimum-energy cost of driving the system from the origin at time $t = 0$ to final point \mathbf{x}_f at time $t = T$ is given by $c(\mathbf{x}_f) = \int_0^T \mathbf{u}'(\tau)\mathbf{u}'(\tau)d\tau = \mathbf{x}_f'\mathbf{W}^{-1}(T)\mathbf{x}_f$, for a fixed energy $c(\mathbf{x}_f) = c^*$ the gramian essentially prescribes a reachable ellipsoid in n -dimensional space, with vertices located at $\lambda_1\mathbf{v}_1, \dots, \lambda_n\mathbf{v}_n$. This ellipsoid we formally define as

$$\mathfrak{E} \equiv \{\mathbf{x} \in \mathbb{R}^n | \mathbf{x}'\mathbf{W}^{-1}(T)\mathbf{x} = c^*\} \quad (4.5)$$

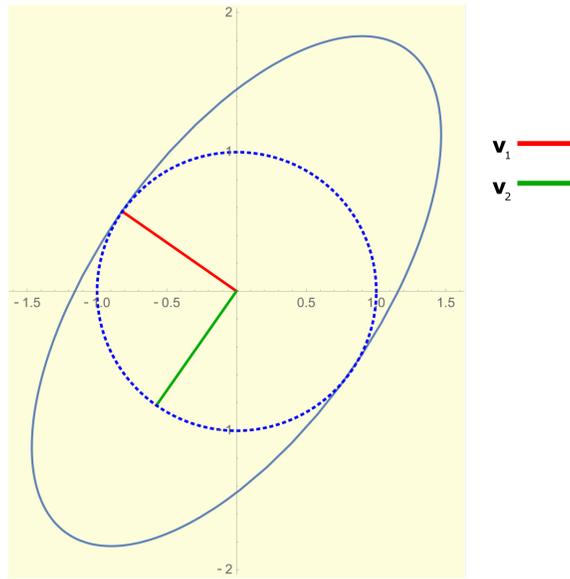


Figure 4.1: Unit circle with ellipse prescribed by $\mathbf{W}(T)^{-1}$ of a stable two-dimensional system with the two eigenvectors of $\mathbf{W}(T)^{-1}$. Since the ellipse represents the reachable set with fixed energy λ_1 , the distance between the ellipse and circle, in the direction of \mathbf{v}_2 , encodes the amount of excess energy available if we desire to steer the system to \mathbf{v}_2 . Note that the figure represents an abstract state space for a 2-d linear system, and thus the axes have no explicit units.

Thus, the surface of Ξ encodes the maximal distance to which trajectories can attain at time T under inputs with total energy c^* (starting from the origin). A common way to assay the controllability of a linear system is to deduce the minimum energy, c^{min} , so that system is guaranteed access to a unit-radius hypersphere in \mathbb{R}^n (i.e., $\min c^*$ such that $\Xi \supset \{\mathbf{x} \in \mathbb{R}^n | \mathbf{x}'\mathbf{x} = 1\}$), which as follows from (4.5), is simply λ_1 in (4.4).

However, as posed in the Introduction, in this scenario, most of the hypersphere is reachable with an excess of energy and we would like to quantify the utility of this excess. Carrying forth the geometric interpretation, it is perhaps intuitive to propose a quantification involving the ratio λ_i/λ_1 , $i = 2, \dots, n$, i.e., the eccentricity of the ellipse made by intersecting Ξ with the (v_1, v_i) -plane.

As we will show, this quantity not only results in a useful metric for excess energy utilization, but has a precise interpretation in terms of the relative orientation attainable by putative inputs to the system at hand. We first make concrete the notion of orientation range.

4.2.2 Input Orientation and Orientation Range

We begin by defining formally the expression for relative orientation between two inputs, introduced in (4.3). For our purposes we will assume one of the inputs is known a priori.

Definition 4 (Input Orientation) For reference input $\mathbf{u}_1(t)$ and ‘free’ input $\mathbf{u}_2(t)$ guiding a system of the form (5.1) over the time interval $[0, T]$, we define the relative orientation $\mathbf{d}_{\mathbb{J}}(\mathbf{u}_1(t), \mathbf{u}_2(t), t \in [0, T])$ —for ease of notation simplified to $\mathbf{d}_{\mathbb{J}}(\mathbf{u}_1, \mathbf{u}_2)$ —as follows:

$$\mathbf{d}_{\mathbb{J}}(\mathbf{u}_1, \mathbf{u}_2) = \frac{1}{T \sqrt{\gamma_{\mathbf{u}_1} \gamma_{\mathbf{u}_2}}} \int_0^T \mathbf{u}'_1(\tau) \mathbf{u}_2(\tau) d\tau, \quad (4.6)$$

where

$$\gamma_{\mathbf{u}_1} = \frac{1}{T} \int_0^T \mathbf{u}'_1(\tau) \mathbf{u}_1(\tau) d\tau, \quad \gamma_{\mathbf{u}_2} = \frac{1}{T} \int_0^T \mathbf{u}'_2(\tau) \mathbf{u}_2(\tau) d\tau \quad (4.7)$$

are the average energies (up to time T) of $\mathbf{u}_1(t)$ and $\mathbf{u}_2(t)$, respectively.

We again note that, although we refer to $\mathbf{d}_{\mathbb{J}}(\mathbf{u}_1, \mathbf{u}_2)$ as the relative orientation between inputs \mathbf{u}_1 and \mathbf{u}_2 , in a strict mathematical sense it quantifies the average cosine of the angle between the two inputs.

Under our assumption, the reference is known (thus $\gamma_{\mathbf{u}_1}$ is known), and we seek, given some chosen available energy $\gamma_{\mathbf{u}_2}$, to quantify our freedom in allowing $\mathbf{u}_2(t)$ to assume different geometries while still accomplishing a desired state transfer. We now propose a notion of orientation range based on the maximum similarity between two inputs of fixed energy:

Definition 5 (Maximum Similarity) For inputs of energy $\gamma_{\mathbf{u}_1}$, $\gamma_{\mathbf{u}_2}$, the maximum similarity is defined as

$$\begin{aligned} & \underset{\mathbf{u}_2}{\text{Max}} && \mathbf{d}_{\mathbb{J}}(\mathbf{u}_1, \mathbf{u}_2) \\ & \text{Subject To} && \mathbf{x}_{1_f} = \int_0^T e^{\mathbf{A}(T-\tau)} \mathbf{B}(\tau) \mathbf{u}_1(\tau) d\tau \\ & && \mathbf{x}_{2_f} = \int_0^T e^{\mathbf{A}(T-\tau)} \mathbf{B}(\tau) \mathbf{u}_2(\tau) d\tau \\ & && \frac{1}{T} \int_0^T \mathbf{u}'_1(\tau) \mathbf{u}_1(\tau) d\tau = \gamma_{\mathbf{u}_1} \\ & && \frac{1}{T} \int_0^T \mathbf{u}'_2(\tau) \mathbf{u}_2(\tau) d\tau = \gamma_{\mathbf{u}_2} \end{aligned} \quad (4.8)$$

where we require the system to be steered, over the time interval $[0, T]$, from the points \mathbf{x}_{1_0} and \mathbf{x}_{2_0} to endpoints \mathbf{x}_{1_f} and \mathbf{x}_{2_f} by inputs \mathbf{u}_1 and \mathbf{u}_2 , respectively. Further, $\gamma_{\mathbf{u}_1}$ and $\gamma_{\mathbf{u}_2}$ are energy constraints on the inputs.

Lemma 6 The optimization problem (4.8) has a provably unique solution \mathbf{u}_2^* , and the minimum relative orientation is given by the following:

$$\mathbf{d}_{\mathbb{J}}(\mathbf{u}_1, \mathbf{u}_2^*) = \frac{1}{T \sqrt{\gamma_{\mathbf{u}_1} \gamma_{\mathbf{u}_2^*}}} \left[(\mathbf{x}_{1_f} - e^{\mathbf{A}T} \mathbf{x}_{1_0})' \mathbf{W}^{-1} (\mathbf{x}_{2_f} - e^{\mathbf{A}T} \mathbf{x}_{2_0}) + \sqrt{(\gamma_{\mathbf{u}_1} T - \mathbb{E}_{\min_1})(\gamma_{\mathbf{u}_2^*} T - \mathbb{E}_{\min_2})} \right] \quad (4.9)$$

where to simplify notation we employ \mathbf{W} to represent $\mathbf{W}(T)$ (see (4.2)), and where \mathbb{E}_{\min_1} and \mathbb{E}_{\min_2} are the minimum energies required for state transfers from the origin to \mathbf{x}_{1_f} and \mathbf{x}_{2_f} , respectively.

Proof 7 The existence and uniqueness of \mathbf{u}_2^* is proved, by deriving optimality conditions based on the Euler-Lagrange equations, in [17], and a derivation of (4.9) is given in [18].

Remark 8 A critical conceptual point is that relative orientation is minimized when $\mathbf{d}_{\mathbb{J}}(\mathbf{u}_1, \mathbf{u}_2^*) = 1$. The relative orientation decreases as $\mathbf{d}_{\mathbb{J}}$ increases, since a larger value means that inputs are more similar. This is apparent from (4.6), applying the law of cosines.

We note that $\mathbf{d}_{\mathbb{J}}(\mathbf{u}_1, \mathbf{u}_2^*) = \mathbf{d}_{\mathbb{J}}(\mathbf{u}_2, \mathbf{u}_1^*)$; that is, in considering the general form for relative orientation given in (4.6), we may choose either \mathbf{u}_1 or \mathbf{u}_2 to be the reference input, and the resulting minimum relative orientation (given by (4.9)) will be equivalent. This is because if we maximize (4.8) over \mathbf{u}_1 with respect to \mathbf{u}_2 , the optimization constraints (i.e. those prescribing terminal states \mathbf{x}_1 and \mathbf{x}_2 and input energies $\gamma_{\mathbf{u}_1}$ and $\gamma_{\mathbf{u}_2}$) are identical to those in (4.8), and the order of \mathbf{u}_1 and \mathbf{u}_2 does not affect the value of (4.6). Indeed, aside from $\mathbf{W}, \mathbb{E}_{\min_1}$, and \mathbb{E}_{\min_2} , which are independent of \mathbf{u}_1 and \mathbf{u}_2 , (4.9) depends only upon these constraints and is invariant to their indexing (that is, we may ‘swap’ 1’s for 2’s in (4.9) without changing the value). Thus we may, without ambiguity, speak of the minimum relative

orientation between inputs required for two state transfers, without deciding which input is the reference.

Thus, the solution (see [17] for a derivation and expression for \mathbf{u}_2^*) to (4.9) provides a direct characterization of the maximum similarity that two inputs, to two different points in state space, may attain as a function of energy. This solution, thus, provides us with a concrete notion of orientation range. We will now operationalize this notion by sampling, in a systemic way, the space of possible pairs of terminal states (\mathbf{x}_1 and \mathbf{x}_2) and possible input energies $\gamma_{\mathbf{u}_1}$ and $\gamma_{\mathbf{u}_2}$. As may be surmised, the sampling proceeds along the spectrum of \mathbf{W}^{-1} .

4.3 Results

The expression (4.9) is notable because it indicates that the minimum orientation difference associated with a pair of state transfers (each starting at the origin) depends only on the energy prescribed to those transfers. That is, while \mathbf{u}_1 may take many forms (assuming $\gamma_{\mathbf{u}_1}T > \mathbb{E}_{\min_1}$) its minimum difference with respect to any putative \mathbf{u}_2 with $\gamma_{\mathbf{u}_2}T > \mathbb{E}_{\min_2}$ is always the same.

4.3.1 The Gramian Bispectrum

Generally, we seek to compare a set of M n -dimensional systems (or networks), denoted S_k , $k \in \{1 \dots M\}$. For each system S_k , we obtain the spectrum of the inverse gramian at some fixed time T (denoted, as above, by \mathbf{W}^{-1}) to ascertain their energy characteristics. For the bispectral analysis we create, for each S_k , a matrix \mathbf{Bis}_k showing the minimum relative orientation required for pairs of inputs guiding the system from the origin to endpoints on the n -dimensional unit hypersphere. These endpoints are simply the (orthonormal) eigenvectors

of \mathbf{W}^{-1} , ranked in terms of their corresponding eigenvalues. Since each transfer begins at the origin, we can simplify (4.9) (we also generalize indices) to obtain

$$\mathbf{d}_{\mathbb{J}}(\mathbf{u}_i, \mathbf{u}_j) = \frac{1}{T \sqrt{\gamma_{\mathbf{u}_i} \gamma_{\mathbf{u}_j}}} \left(\mathbf{x}'_{\mathbf{u}_i} \mathbf{W}^{-1} \mathbf{x}_{\mathbf{u}_j} + \sqrt{(\gamma_{\mathbf{u}_i} T - \mathbb{E}_{\min_i})(\gamma_{\mathbf{u}_j} T - \mathbb{E}_{\min_j})} \right) \quad (4.10)$$

for target terminal states $\mathbf{x}_{\mathbf{u}_i}$ and $\mathbf{x}_{\mathbf{u}_j}$.

To fully specify \mathbf{Bis}_k we need to also determine the average energy associated with each state transfer. The natural choice is to allow $\gamma_{\mathbf{u}_i} = \gamma_{\mathbf{u}_j} = \frac{\lambda_{\max}}{T}$. Thus, the total energy supplied is λ_{\max} , which, as stated in the Introduction, is the smallest amount of energy that enables access to the entire unit hypersphere.

With these specifications we rewrite (4.10), for any pair of state transfers in the directions of \mathbf{v}_i and \mathbf{v}_j ($i, j \in \{1 \dots n\}$), as follows

$$\begin{aligned} \mathbf{d}_{\mathbb{J}}(\mathbf{u}_i, \mathbf{u}_j) &= \frac{1}{T \sqrt{\frac{\lambda_{\max}}{T} \frac{\lambda_{\max}}{T}}} \left(\mathbf{v}'_i \mathbf{W}^{-1} \mathbf{v}_j + \sqrt{\left(\frac{\lambda_{\max}}{T} T - \mathbb{E}_{\min_i} \right) \left(\frac{\lambda_{\max}}{T} T - \mathbb{E}_{\min_j} \right)} \right) \\ &= \frac{1}{\lambda_{\max}} \left(\mathbf{v}'_i \mathbf{W}^{-1} \mathbf{v}_j + \sqrt{(\lambda_{\max} - \mathbb{E}_{\min_i})(\lambda_{\max} - \mathbb{E}_{\min_j})} \right) \end{aligned} \quad (4.11)$$

Now we consider \mathbb{E}_{\min_i} and \mathbb{E}_{\min_j} . It is well known [41] that the minimum required energy for a state transfer from the origin to some terminal point $\mathbf{x}(T)$ is given by $\mathbf{x}'(T) \mathbf{W}^{-1} \mathbf{x}(T)$ (recall that throughout we are using \mathbf{W}^{-1} to represent $\mathbf{W}(T)^{-1}$), and since in our case we have $\mathbf{x}(T) = \mathbf{v}_i$ (or \mathbf{v}_j), where \mathbf{v}_i is a (normal) eigenvector of \mathbf{W}^{-1} , the minimum energy becomes

$$\mathbb{E}_{\min_i} = \mathbf{v}'_i \mathbf{W}^{-1} \mathbf{v}_i = \lambda_i \mathbf{v}'_i \mathbf{v}_i = \lambda_i \quad (4.12)$$

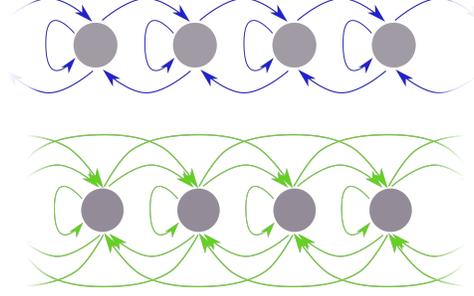


Figure 4.2: Connection graphs for networks \mathcal{P} (blue) and \mathcal{Q} (green)

Equivalently, $\mathbb{E}_{\min_j} = \lambda_j$. Also, we note that

$$\mathbf{v}_i' \mathbf{W}^{-1} \mathbf{v}_j = \mathbf{v}_i' \lambda_j \mathbf{v}_j = \lambda_j \mathbf{v}_i' \mathbf{v}_j = \begin{cases} 0 & \text{when } i \neq j \\ \lambda_j & \text{when } i = j \end{cases} \quad (4.13)$$

Thus, if $i \neq j$, we can construct the gramian bispectrum element-wise from (4.11) via

$$\mathbf{Bis}_k(i, j) \equiv \mathbf{d}_{\mathbb{J}}(\mathbf{u}_i, \mathbf{u}_j) = \frac{\sqrt{(\lambda_{\max} - \lambda_i)(\lambda_{\max} - \lambda_j)}}{\lambda_{\max}} \quad (4.14)$$

while if $i = j$, we have

$$\begin{aligned} \mathbf{Bis}_k(i, j) \equiv \mathbf{d}_{\mathbb{J}}(\mathbf{u}_i, \mathbf{u}_j) &= \frac{1}{\lambda_{\max}} \left(\lambda_j + \sqrt{(\lambda_{\max} - \lambda_i)(\lambda_{\max} - \lambda_j)} \right) \\ &= 1 \end{aligned} \quad (4.15)$$

which is expected because of the law of cosines and the fact that $\mathbf{d}_{\mathbb{J}}(\mathbf{u}_i, \mathbf{u}_j)$ is the average integrated angular difference between inputs \mathbf{u}_i and \mathbf{u}_j (4.6).

Thus, it is quite straightforward to calculate the matrix \mathbf{Bis}_k for system S_k . As a second-order spectral characterization, we term \mathbf{Bis}_k the gramian bispectrum of the system S_k .

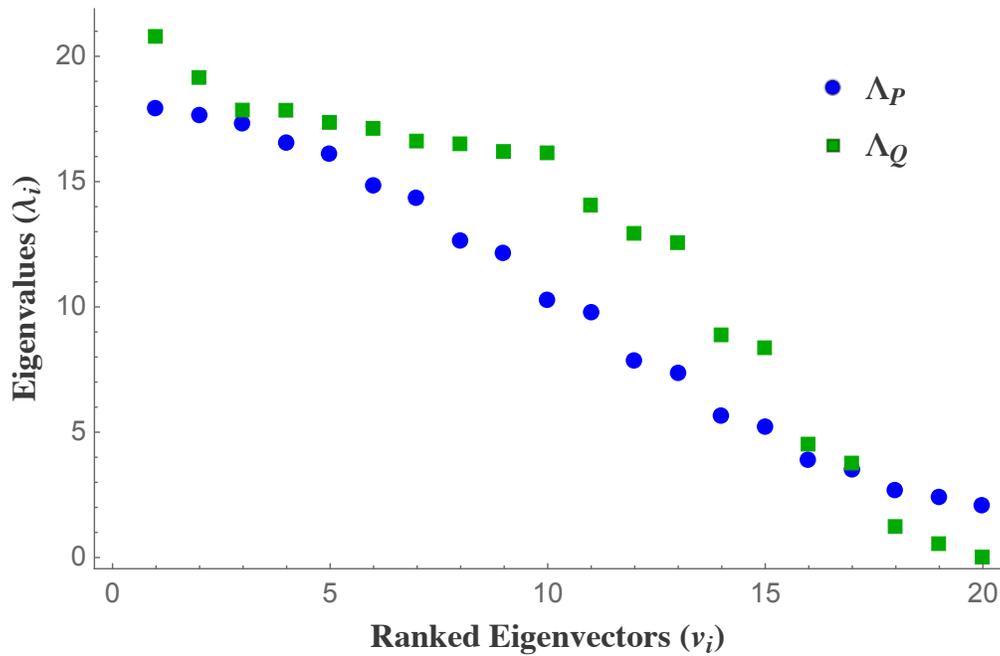


Figure 4.3: Spectra of $\mathbf{W}_{\mathcal{P}}^{-1}$ and $\mathbf{W}_{\mathcal{Q}}^{-1}$. The horizontal axis indexes the eigenvectors, which are sorted according to the magnitude of their associated real (since \mathbf{W}^{-1} is real and symmetric) eigenvalues.

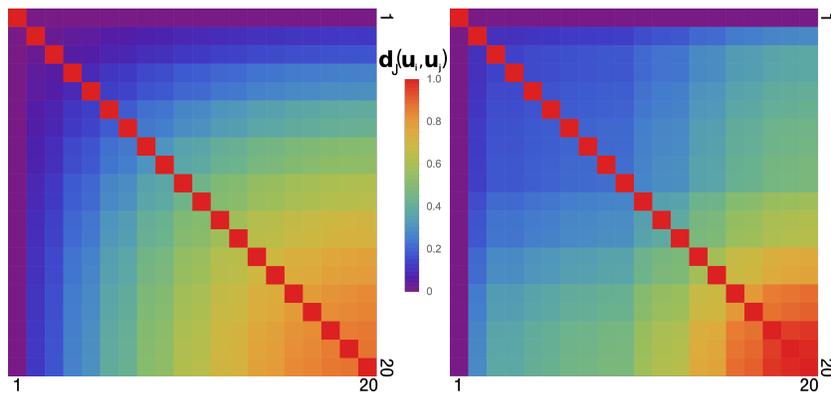


Figure 4.4: Array plots of gramian bispectra for networks \mathcal{P} and \mathcal{Q} . As explained in detail in Section 4.3.1, $\mathbf{d}_{\mathbb{J}}(\mathbf{u}_i, \mathbf{u}_j)$ quantifies the average cosine of the angle between inputs \mathbf{u}_i and \mathbf{u}_j , guiding the system to endpoints on the unit hypersphere in the directions of eigenvectors (of the gramian inverse \mathbf{W}^{-1}) \mathbf{v}_i and \mathbf{v}_j , respectively. Thus if $\mathbf{d}_{\mathbb{J}}(\mathbf{u}_i, \mathbf{u}_j) = 1$, the inputs may be identical, which explains why the diagonal entries (where, obviously, $\mathbf{v}_i = \mathbf{v}_j$) are all equal to 1.

4.3.1.1 Example: Two Toeplitz Networks

We consider a simple construction of two 20-dimensional linear time-invariant Toeplitz (i.e. the entries in any diagonal of the adjacency matrix are equal) networks (\mathcal{P} and \mathcal{Q}), wherein each node of \mathcal{P} has a self-connection and a bi-directional connection with its immediate neighbors, while \mathcal{Q} has those connections as well as bi-directional connections with its secondary neighbors. The pattern of connectivity is schematized in Figure 4.2. The eigenspectra for the two networks are shown in Figure 4.3, and the gramian bispectra given in Figure 4.4. As we examine Figure 4.4, we notice that the overall contour of the ranked eigenvalues (Figure 4.3) is reflected in the bispectra, but that, especially for network \mathcal{Q} , there are subtleties reflected in the pairwise representation (i.e. the gramian bispectra) that can not be inferred from the eigenspectra, as we will explore below in Section III.C.

Since the gramian bispectrum is essentially a tabulation of within-system control flexibility which characterizes how a system may trade energy for a diversity of inputs to reach orthogonal points in state-space, worthy of special attention is the fact that the first row and column of the gramian bispectrum is always 0 (aside from the $\{1, 1\}^{\text{th}}$ entry, whose value will be explained below), which, while apparent from (4.14), at first seems conceptually counterintuitive. The pairs of transfers whose minimum input orientation differences are encoded in the first row or column of the bispectrum are such that one of the state transfers employs a minimum energy input; that is $\lambda_i = \lambda_{\max}$ (first row) or $\lambda_j = \lambda_{\max}$ (first column). It would seem that, since the other input has an increasing amount of excess energy as we move along the first row or down the first column, that the orientation cost should decrease for the pair of inputs. That is, since $\lambda_{i_1} \geq \lambda_{i_2}$ when $i_1 < i_2$, more excess energy is available for state transfers the farther we go along row 1 (or column 1, substituting j for i), which intuitively would seem to indicate greater input flexibility. However, this is not the case: when $\lambda_i = \lambda_{\max}$ or $\lambda_j = \lambda_{\max}$,

it does not matter how much excess energy is provided for the other transfer, the best we can do is to be, on average, orthogonal. We formalize this notion in the following lemma:

Lemma 7 For inputs $\mathbf{u}_1(t)$ and $\mathbf{u}_2(t)$, guiding the system (over time horizon $t \in [0, T]$) from the origin to \mathbf{v}_1 and \mathbf{v}_2 , respectively, where \mathbf{v}_1 and \mathbf{v}_2 are eigenvectors of $\mathbf{W}^{-1}(T)$, and where exactly minimum energy is provided for at least one of the transfers, $\mathbf{d}_{\mathbb{J}}(\mathbf{u}_1, \mathbf{u}_2) = 0$.

Proof 8 Let us examine (4.6) and assume we provide minimum energy for \mathbf{u}_1 . Since $\mathbf{x}(0)$ is the origin and $\mathbf{x}(T) = \mathbf{v}_1$, where \mathbf{v}_1 is an eigenvector of $\mathbf{W}^{-1}(T)$, classical control theory (see [41], for example) tells us that minimum energy input \mathbf{u}_1 is given by: $\mathbf{u}_1(t) = \mathbf{B}'(t)e^{\mathbf{A}(T-t)}\mathbf{W}^{-1}(T)\mathbf{v}_1$. Substituting into (4.6), we obtain

$$\begin{aligned}
\mathbf{d}_{\mathbb{J}}(\mathbf{u}_1, \mathbf{u}_2) &= \frac{1}{T\sqrt{\gamma_{\mathbf{u}_1}\gamma_{\mathbf{u}_2}}} \int_0^T (\mathbf{B}'(\tau)e^{\mathbf{A}(T-\tau)}\mathbf{W}^{-1}(T)\mathbf{v}_1)'\mathbf{u}_2(\tau)d\tau \\
&= \frac{1}{T\sqrt{\gamma_{\mathbf{u}_1}\gamma_{\mathbf{u}_2}}} \int_0^T \mathbf{v}_1'\mathbf{W}^{-1}(T)e^{\mathbf{A}(T-\tau)}\mathbf{B}(\tau)\mathbf{u}_2(\tau)d\tau \\
&= \frac{1}{T\sqrt{\gamma_{\mathbf{u}_1}\gamma_{\mathbf{u}_2}}} \mathbf{v}_1'\mathbf{W}^{-1}(T) \int_0^T e^{\mathbf{A}(T-\tau)}\mathbf{B}(\tau)\mathbf{u}_2(\tau)d\tau \\
&= \frac{1}{T\sqrt{\gamma_{\mathbf{u}_1}\gamma_{\mathbf{u}_2}}} \mathbf{v}_1'\mathbf{W}^{-1}(T)\mathbf{v}_2 \\
&= \frac{1}{T\sqrt{\gamma_{\mathbf{u}_1}\gamma_{\mathbf{u}_2}}} \lambda_2 \mathbf{v}_1'\mathbf{v}_2 = 0
\end{aligned} \tag{4.16}$$

This somewhat surprising result means that any \mathbf{u}_2 which we choose to guide the system to \mathbf{v}_2 will have the same average angular difference ($\frac{\pi}{2}$) from the minimum energy \mathbf{u}_1 . Thus we obtain a sort of paradoxical freedom wherein we can do anything we want with \mathbf{u}_2 (subject to endpoint constraints) but we cannot improve on its similarity to \mathbf{u}_1 .

4.3.2 Geometric Interpretation

Since the lengths of the axes of Ξ (see (4.5), where for the present purposes we let $c^* = 1$) are given by $\frac{1}{\sqrt{\lambda_n}}, \dots, \frac{1}{\sqrt{\lambda_1}}$, we define $E_{1,i}$ to be the 2-D ellipse which results from a projection of Ξ onto the axes given by \mathbf{v}_1 and \mathbf{v}_j , the eigenvectors associated with λ_1 and λ_j (recall that λ_1 is the largest eigenvalues of $\mathbf{W}^{-1}(T)$, so we'll represent it, as above, by λ_{\max}). Then the eccentricity (we will call it $ecc_{E_{1,i}}$) of $E_{1,i}$ is given by

$$\begin{aligned} ecc_{E_{1,i}} &= \sqrt{1 - \frac{\left(\frac{1}{\sqrt{\lambda_{\max}}}\right)^2}{\left(\frac{1}{\sqrt{\lambda_i}}\right)^2}} = \sqrt{1 - \frac{\lambda_i}{\lambda_{\max}}} \\ &= \sqrt{\frac{\lambda_{\max} - \lambda_i}{\lambda_{\max}}} \end{aligned} \quad (4.17)$$

Choosing another such projected ellipse $E_{1,j}$ and its associated eccentricity $ecc_{E_{1,j}} = \sqrt{\frac{\lambda_{\max} - \lambda_j}{\lambda_{\max}}}$, we notice that

$$\begin{aligned} ecc_{E_{1,i}} * ecc_{E_{1,j}} &= \sqrt{\frac{\lambda_{\max} - \lambda_i}{\lambda_{\max}}} \sqrt{\frac{\lambda_{\max} - \lambda_j}{\lambda_{\max}}} \\ &= \frac{\sqrt{(\lambda_{\max} - \lambda_i)(\lambda_{\max} - \lambda_j)}}{\lambda_{\max}} \\ &= \mathbf{d}_{\mathbb{J}}(\mathbf{u}_i, \mathbf{u}_j) \quad (\text{Compare with (4.14)}) \end{aligned} \quad (4.18)$$

That is, the minimum relative orientation between inputs \mathbf{u}_i and \mathbf{u}_j , utilizing total energy λ_{\max} and driving the system to points on the unit hypersphere along the directions of eigenvectors \mathbf{v}_i and \mathbf{v}_j , respectively, is the same as the multiple of the eccentricities of the 2-D ellipses generated from the projections of Ξ onto \mathbf{v}_1 together with \mathbf{v}_i and \mathbf{v}_j , respectively. Note that for different values of c^* (again see (4.5)), Ξ is simply scaled and the eccentricity is not affected. Further, a uniform scaling of the energy used by \mathbf{u}_i and \mathbf{u}_j does not affect

$\mathbf{d}_{\mathbb{J}}(\mathbf{u}_i, \mathbf{u}_j)$ (see (4.20)). Thus (4.18) holds for $c^* \neq 1$. The geometric interpretation further highlights the fact that the gramian bispectrum furnishes a nice overall picture of system controllability in terms of energy, since these various eccentricities are exactly correlated with relative minimum energy requirements.

We now move on to concentrate on comparisons of gramian bispectra across-system for different networks.

4.3.3 Comparing Orientation Bispectra

The various energy-based control metrics are useful primarily because of their comparative power; that is, given some energy metric and set of networks, we can easily rank the networks in terms of the energy costs. Since orientation difference is a more relativistic notion, we compare networks by calculating the gramian bispectral difference matrix

$$\mathbf{D}_{\mathcal{P}, \mathcal{Q}} = \mathbf{Bis}_{\mathcal{P}} - \mathbf{Bis}_{\mathcal{Q}} \quad (4.19)$$

for networks \mathcal{P} and \mathcal{Q} . The $\{i, j\}^{\text{th}}$ entry of $\mathbf{D}_{\mathcal{P}, \mathcal{Q}}$ tells us the difference in relative orientation required (between the two networks) to accomplish the pair of transfers in the directions of the i^{th} and j^{th} eigenvectors (of the respective networks).

4.3.3.1 Example: Toeplitz networks reveal informative bispectral differences

We continue our analysis using the same Toeplitz networks (\mathcal{P} and \mathcal{Q}) used in Example 1. If we examine Figure 4.3 we can make some conclusions about the relative controllability of networks \mathcal{P} and \mathcal{Q} , namely that \mathcal{P} requires less energy to steer in most directions of the state space. For only the three ‘easiest’ directions is \mathcal{Q} more energy-efficient. This may be all we desire to know in comparing the two systems if energy is our primary consideration.

However, if we know that we have energy to spare, we can examine the bispectral difference of the two networks to evaluate their relative efficiency in trading excess energy for input flexibility.

To make a fair comparison of the two networks, we will rescale the eigenspectra so that their maximum eigenvectors are equal—that is, $\lambda_{\mathcal{P},1} = \lambda_{\mathcal{Q},1}^*$, where $\lambda_{\mathcal{Q},1}^* \in \Lambda_{\mathcal{Q}}^*$, the transformed eigenspectrum for network \mathcal{Q} . Note that this will not affect the gramian bispectra, for, if we scale the spectrum (i.e. all of the eigenvalues) of \mathbf{W}^{-1} by some constant α and evaluate (4.14), we have

$$\begin{aligned} \mathbf{d}_{\mathbb{J}}^*(\mathbf{u}_i, \mathbf{u}_j) &= \frac{\sqrt{(\alpha\lambda_{\max} - \alpha\lambda_i)(\alpha\lambda_{\max} - \alpha\lambda_j)}}{\alpha\lambda_{\max}} \\ &= \mathbf{d}_{\mathbb{J}}(\mathbf{u}_i, \mathbf{u}_j) \end{aligned} \quad (4.20)$$

Now, let $\lambda_{\mathcal{P},\max}$ and $\lambda_{\mathcal{Q},\max}$ be the maximal inverse gramian (\mathbf{W}^{-1}) eigenvalues of networks \mathcal{P} and \mathcal{Q} , respectively. As in (4.4), let us denote the ordered spectra of \mathbf{W}^{-1} for the two networks by $\Lambda_{\mathcal{P}}$ and $\Lambda_{\mathcal{Q}}$, respectively. By (4.20), we can let $\Lambda_{\mathcal{Q}}^* = c \Lambda_{\mathcal{Q}}$, where

$$c = \frac{\lambda_{\mathcal{P},\max}}{\lambda_{\mathcal{Q},\max}}, \quad (4.21)$$

such that the rescaled spectrum $\Lambda_{\mathcal{Q}}^*$ results in no change to \mathcal{Q} 's gramian bispectrum (that is, $\mathbf{Bis}_{\mathcal{Q}}^* = \mathbf{Bis}_{\mathcal{Q}}$). Then we have, for network \mathcal{Q} ,

$$\begin{aligned} \mathbf{Bis}_{\mathcal{Q},i,j} &= \mathbf{d}_{\mathbb{J}}(\mathbf{u}_i, \mathbf{u}_j) \\ &= \frac{\sqrt{(\lambda_{\mathcal{Q},\max} - \lambda_{\mathcal{Q},i})(\lambda_{\mathcal{Q},\max} - \lambda_{\mathcal{Q},j})}}{\lambda_{\mathcal{Q},\max}} \\ &= \frac{\sqrt{(\lambda_{\mathcal{P},\max} - c \lambda_{\mathcal{Q},i})(\lambda_{\mathcal{P},\max} - c \lambda_{\mathcal{Q},j})}}{\lambda_{\mathcal{P},\max}} \end{aligned} \quad (4.22)$$

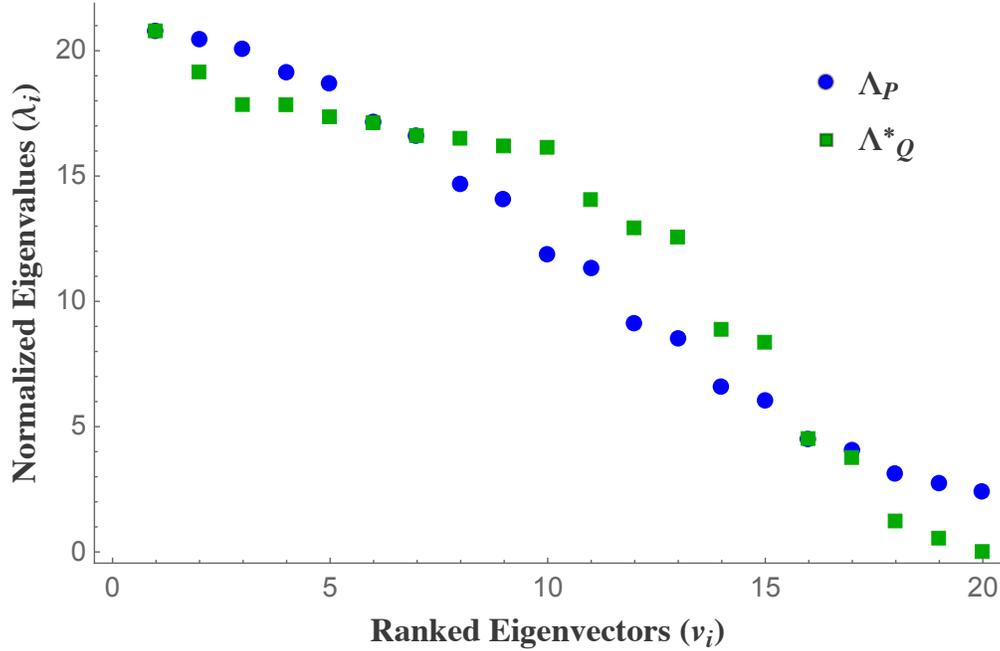


Figure 4.5: Spectra of \mathcal{P} and \mathcal{Q} , with the latter normalized by c (see (4.21)) so that the largest eigenvalues agree. This normalization does not affect the bispectral analysis (see (4.22)). As in Figure 4.3, the horizontal axis indexes the eigenvectors, which are sorted according to the magnitude of their associated eigenvalues.

This rescaling, applied to our networks \mathcal{P} and \mathcal{Q} , results in spectra as shown in Figure 4.5.

The reason this rescaling is advantageous is that it allows us to investigate how different spectral profiles are manifest in the respective gramian bispectra, and ultimately in the bispectrum difference matrix. For example, it is evident from (4.14) that there must exist crossings in the normalized spectra in order for there to be any nontrivial comparison. For, if, without loss of generality, $\lambda_{\mathcal{P},i} > \lambda^*_{\mathcal{Q},i} \forall i \in \{2, \dots, n\}$ (recall that the normalization ensures

that $\lambda_{\mathcal{P},\max} = \lambda_{\mathcal{Q},\max}^* = \lambda_{\max}$) then

$$\begin{aligned}
\mathbf{D}_{\mathcal{P},\mathcal{Q}_{i,j}} &= \mathbf{Bis}_{P_{i,j}} - \mathbf{Bis}_{Q_{i,j}} \\
&= \frac{\sqrt{(\lambda_{\mathcal{P},\max} - \lambda_{\mathcal{P},i})(\lambda_{\mathcal{P},\max} - \lambda_{\mathcal{P},j})}}{\lambda_{\mathcal{P},\max}} - \frac{\sqrt{(\lambda_{\mathcal{Q},\max}^* - \lambda_{\mathcal{Q},i}^*)(\lambda_{\mathcal{Q},\max}^* - \lambda_{\mathcal{Q},j}^*)}}{\lambda_{\mathcal{Q},\max}^*} \\
&= \frac{\sqrt{(\lambda_{\max} - \lambda_{\mathcal{P},i})(\lambda_{\max} - \lambda_{\mathcal{P},j})} - \sqrt{(\lambda_{\max} - \lambda_{\mathcal{P},i})(\lambda_{\max} - \lambda_{\mathcal{P},j})}}{\lambda_{\max}} \\
&< 0 \quad \forall \{i, j\} \ i \in \{2, \dots, n\}, j \in \{2, \dots, n\} \text{ where } i \neq j
\end{aligned}$$

and thus the bispectrum difference matrix would reveal that network \mathcal{Q} is easier to control (in terms of input flexibility) for all pairwise maneuvers. Clearly, \mathcal{Q} would be easier to control in terms of minimum energies as well. But a more complex picture may emerge when the normalized spectra have crossings. The fact that (4.14) is a nonlinear function of the magnitudes of λ_{\max} , λ_i , and λ_j makes it impossible to predict, upon inspection of the eigenspectra, whether $\mathbf{D}_{\mathcal{P},\mathcal{Q}_{i,j}}$ will be positive or negative when either $(\lambda_{\mathcal{P},i} > \lambda_{\mathcal{Q},i}^*) \wedge (\lambda_{\mathcal{P},j} < \lambda_{\mathcal{Q},j}^*)$ or $(\lambda_{\mathcal{P},i} < \lambda_{\mathcal{Q},i}^*) \wedge (\lambda_{\mathcal{P},j} > \lambda_{\mathcal{Q},j}^*)$ (where \wedge is the logical ‘and’ operator).

We seek now to understand how this bispectral difference analysis can be interpreted to infer relative controllability strengths and weaknesses, using Figure 4.6 (which was constructed by using the sign function on the entries of $\mathbf{D}_{\mathcal{P},\mathcal{Q}}$) as an example. The first thing that strikes us in examining Figure 4.6 is that for most pairs of maneuvers wherein one has relatively little excess energy (this corresponds to the first rows and columns of the matrix), network \mathcal{Q} has lower pairwise orientation cost, as represented by the large swaths of green in the early rows and columns of Figure 4.6. Now, examining Figure 4.5, we see there is a crossing at the 6th eigenvalue, so we would expect to see a change in the difference matrix at the sixth row/column (recall that $\mathbf{D}_{\mathcal{P},\mathcal{Q}}$ is always symmetric). Indeed we do see green give way to blue,

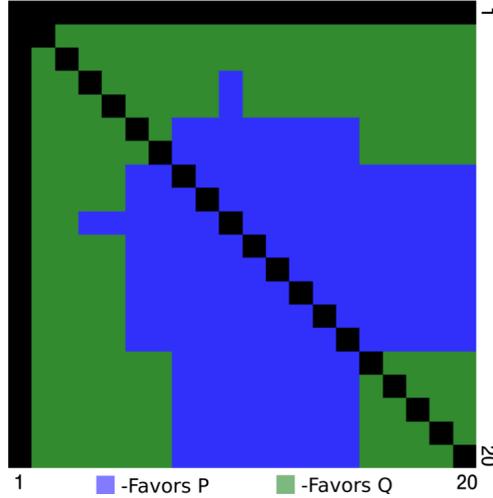


Figure 4.6: A visualization of bispectrum difference matrix $\mathbf{D}_{\mathcal{P},\mathcal{Q}}$, a tool for comparing the excess energy utility of the two networks. The $\{i, j\}^{\text{th}}$ entry is colored blue where network \mathcal{P} has the lower required angular input separation for maneuvers to \mathbf{v}_i and \mathbf{v}_j , and colored green if network \mathcal{Q} is better in this respect.

representing the fact that network \mathcal{P} has a lower orientation cost for pairwise maneuvers to eigenvectors with associated eigenvalues in the ‘middle’ section of the spectrum.

However, the change is more complex than can be inferred from the spectra alone. Notice, in Figure 4.6, the protrusions of blue into the green section which appears in the 4th and 5th rows (columns) and the 10th column (row). This indicates that, for pairwise maneuvers where one of the maneuvers involves the 10th eigenvector (\mathbf{v}_{10}) and the other involves \mathbf{v}_i $i \in \{4, 5\}$, network \mathcal{P} has the better excess energy utility, while if the \mathbf{v}_i $i \in \{4, 5\}$ are paired with any other eigenvector, network \mathcal{Q} is more favorable. Clearly, by examining only the sign of the entries of $\mathbf{D}_{\mathcal{P},\mathcal{Q}}$, we limit the possible depth of the analysis. If we include magnitudes in our consideration, an even richer characterization may emerge. For example, it may be that network \mathcal{P} is better (according to our objective) for only a few pairs of directions, but for those pairs it is much better. If our control objectives prioritize input flexibility in traversing these directions of the state space, we may deem \mathcal{P} the more efficient network.

If, on the other hand, we desire an even further simplification, we could condense the information provided by the bispectral difference into one summary statistic by simply summing over the entries of $\mathbf{D}_{\mathcal{P},\mathcal{Q}}$ to obtain an overall picture of which network best utilizes its excess energy. However, much of the nuanced interpretation would be lost by doing so.

4.4 Conclusions and Future Work

We have introduced a mathematical construct, the gramian bispectrum, which provides a characterization of how a system trades off energy for control versatility in the sense of an angular distance metric. We remark that, while certainly the standard (energy) spectra give some indications as to what we might expect to see in the relative gramian bispectral matrices (i.e. the $\mathbf{D}_{\mathcal{P},\mathcal{Q}}$), these matrices themselves furnish a richer picture of the relative control versatility for maneuvers in the different (in these cases orthogonal) directions. We suggest that the interpretive possibilities of the gramian bispectrum concept hold potential in the study and comparison of linear systems and networks.

We have thus far derived, deployed (in several ways), and extended the idea of orientation sensitivity as a measure of network information-processing ability. We now pursue a method of holistic characterization of information processing—namely, total Fisher information in a network with respect to its afferent inputs. We conduct this analysis in the next chapter in the paradigm of an inner product-based readout. That is, we ask: How much can we know about an input \mathbf{u} by a readout of the projection of its corresponding state $\mathbf{x}_{\mathbf{u}}(t)$'s projection onto a ‘background’ state. The reasons for using this setup are discussed at length in the chapter—in a nutshell, neuroscience has experimentally confirmed the existence of definitive default states in the brain, against which all other activity must therefore be measured. Hence we examine Fisher information with respect to a projection onto this background state. Of

course, the background state is variable, and one of the chapter's contributions is a derivation of a way to optimize the Fisher information over the space of possible backgrounds—i.e. to find the optimal default state.

Chapter 5

Information spectra and optimal background states for dynamical networks

We consider the notion of stimulus representation over dynamic networks, wherein the network states encode information about the identify of an afferent input (i.e. stimulus). Our goal is to understand how the structure and temporal dynamics of networks support information processing. In particular, we conduct a theoretical study to reveal how the background or ‘default’ state of a network with linear dynamics allows it to best promote discrimination over a continuum of stimuli. Our principal contribution is the derivation of a matrix whose spectrum (eigenvalues) quantify the extent to which the state of a network encodes its inputs. This measure, based on the notion of a Fisher linear discriminant, is relativistic in the sense that it provides an information value quantifying the ‘knowablility’ of an input based on its projection onto the background state. We subsequently optimize the background state and highlight its relationship to underlying state noise covariance. This result demonstrates how the best idle state of a network may be informed by its structure and dynamics. Further, we

relate the proposed information spectrum to the controllability gramian matrix, establishing a link between fundamental control-theoretic network analysis and information processing.

5.1 Introduction

In network science, considerable effort has been directed at structural analysis that reveals the interconnection architecture of engineered and biological networks [42–47]. While such analysis can illuminate intriguing and common architectural principles of complex systems, it alone cannot tell us the functionality of such architecture. In other words, to what end is the revealed structure useful? Our goal in this work is to analyze the relationship between structure, dynamics and function of (networked) systems. The specific notion of function that we consider is information coding, which has to do with how networks represent a stimulus or extrinsic input in a way that is useful for downstream processing (i.e., so that an agent can decode the identity of input stimuli based on a ‘read out’ of the state of the network). This sort of coding has been a topic of much interest in theoretical neuroscience, where understanding how networks of neurons represent stimuli is a foundational question [14–16].

Of course, many general principles of information coding are known from communication theory [48]. However, it is not clear how principles of information transmission, coding/decoding and capacity are impacted when enacted over a networked system, especially one with continuous time dynamics. That is, what structure and dynamical aspects of a network make it a good information encoder? To this end, we principally address two questions: 1) What sorts of dynamics shape the input/output relationship of a network in a way which is effective in the Shannon sense (i.e. some, but not too much, redundancy to enable robust, efficient communication in the presence of noise)? It is especially unclear whether dynamical networks that do a good job encoding and/or processing information are also those that are

most responsive to their inputs in a control-theoretic sense. Hence the second question: 2) Is a network that is easily controlled by its inputs necessarily one that also effectively encodes information about those inputs? These two related questions constitute the primary focus of the chapter.

We consider information processing defined in terms of the extent to which network states/outputs encode their respective inputs. Our particular focus is on the background state of a network and its ability to facilitate information extraction regarding other afferent inputs. Non-zero background states are frequently observed in natural dynamical systems. For example, in the study of brain networks the existence of a ‘resting state’ is well-established experimentally [49–52]. Our goal is to provide a theoretical framework with which we can better understand how non-zero resting states confer informational utility. Specifically, we will derive a background state that is optimal according to a novel information measure (also herein derived). In mathematical terms, suppose that a stimulus \mathbf{u} induces a network state $\mathbf{x}_{\mathbf{u}}$. We will quantify the ‘knowability’ of \mathbf{u} by comparing $\mathbf{x}_{\mathbf{u}}$ against a reference background state \mathbf{x}_{ref} . The optimal \mathbf{x}_{ref} can be interpreted as a ‘state of readiness’ at which the network may be sustained in preparation for activity to follow.

The information measure we employ is based upon the inner product $\langle \mathbf{x}_{\text{ref}}, \mathbf{x}_{\mathbf{u}} \rangle$, and is rooted in the method of Fisher linear discriminants [53–55]. This inner product, as is well-known, prescribes the projection of $\mathbf{x}_{\mathbf{u}}$ onto \mathbf{x}_{ref} . For vectors of known magnitude, this projection gauges angular separation. Thus, essential characteristics of $\mathbf{x}_{\mathbf{u}}$ (and by extension, of \mathbf{u}) can be gleaned in an easily codified and quantified manner. The potential informational value to be derived from a projection of $\mathbf{x}_{\mathbf{u}}$ onto \mathbf{x}_{ref} , depends critically on the effective choice of the background state \mathbf{x}_{ref} and uncertainty/noise. The choice of \mathbf{x}_{ref} , in turn, depends largely on how a network responds to its inputs as a function of time (see Figure 5.1). Noise and uncertainty, likewise, are impacted by the network dynamics.

The formulation of a continuous-time dynamical (networked) system with afferent input is fundamentally aligned with analysis from control theory. A key aspect of our results will be the derivation of a Fisher information matrix, $\mathcal{J}_{\mathbf{u}}$, associated with the above inner product. As we will see, the spectrum of $\mathcal{J}_{\mathbf{u}}$ quantifies the extent to which different afferent inputs produce different state representations. It turns out that this information spectrum has a particular statistical relationship with a traditional element from control theory, the controllability gramian matrix [9]. This is perhaps intuitive since the control gramian is mathematically equivalent to the covariance of a network in response to white noise, a key source of uncertainty (and, thus, information loss). We will formalize this relationship in our results.

Assessment of information propagation through noisy networks has been a topic of increasing interest, and while there are many contexts for which such analyses are relevant, quantifying the information-carrying capacity of (real and/or artificial) neural networks has been an especially active research area [15, 16, 56, 57]. For example, in Zylberberg et al. [16], linear Fisher information is evaluated for a two-layer feedforward network in which a scalar signal is distributed to first-layer nodes and then propagated to the second layer via a weighted matrix, with noise corrupting the output of both layers. The amount of information available about the stimulus, they observe, is dependent on the noise covariance structure at each layer, and on how these covariances relate to one another and to the direction of signal propagation (i.e. the tuning curve). However, the network considered in this work is static, so that input-output relationships are fully determined by network structure alone (i.e. there is no recurrent modulation of the signal, though noise does play a role, obviously).

In contrast, Ganguli et al. [15] quantify the stimulus-encoding capacity of a linear dynamical network, again employing linear Fisher information theory. Here, the stimulus is presented as

a pulse at a specific time, the ‘memory trace’ of which is preserved by the network over time to an extent depending on the network’s topology, and the statistical behavior of state noise.

Our work employs the dynamical framework of Ganguli et al. [15], while considering multi-variate stimuli, akin to the ‘tuning curves’ of Zylberberg et al. [16] In fact our framework allows for stimuli of arbitrary dimensions, although here we do constrain the inputs (stimuli) to be constant (for reasons addressed in the Discussion). Also our notion of how stimulus information is encoded is different. Specifically, we employ the inner-product based readout, facilitating a comparison between an output vector $\mathbf{x}_{\mathbf{u}}$ and a reference \mathbf{x}_{ref} , as mentioned above.

5.2 Results

5.2.1 Problem Formulation and Preliminaries

5.2.1.1 Linear Dynamical Networks

Linear dynamics have been used to describe complex networks in several contexts [58–60], with the caveat that such dynamics provide only local approximations of more complex, nonlinear regimes. Proceeding with this limitation in mind, we consider a linear dynamical system (network) with noise, of the form:

$$\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{u} + \mathbf{w}(t) \tag{5.1}$$

where the n -dimensional state vector \mathbf{x} ’s recurrent dynamics are described by adjacency matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$, input matrix $\mathbf{B} \in \mathbb{R}^{n \times m}$ mediating the m -dimensional input \mathbf{u} , taken here to be constant (see Discussion), and zero-mean gaussian noise $\mathbf{w}(t)$, which has covariance matrix

$\Sigma_{\mathbf{w}}$. We point out the fact that the term dynamical network, in the sense used here, does not imply time-varying dynamics; that is, \mathbf{A} is constant. We wish to consider the linear Fisher information regarding \mathbf{u} given the inner product of the state $\mathbf{x}(t)$ (which varies in time) and a reference background \mathbf{x}_{ref} . By basic linear system theory

$$\mathbf{x}(t) \sim \mathcal{N} \left(\int_0^t e^{\mathbf{A}(t-\tau)} d\tau \mathbf{B} \mathbf{u}, \Sigma_{\mathbf{x}(t)} \right), \quad (5.2)$$

where $\Sigma_{\mathbf{x}(t)} \in \mathbb{R}^{n \times n}$ is a covariance matrix determined by the system dynamics.

5.2.1.2 Inner Product and Fisher Information

As we seek to quantify the extent to which the inner product of $\mathbf{x}(t)$ and \mathbf{x}_{ref} encodes information about the input \mathbf{u} giving rise to $\mathbf{x}(t)$, we employ the Fisher information matrix, denoting it $\mathcal{J}_{\mathbf{u}}$, which is given by

$$\mathcal{J}_{\mathbf{u}}(t) = \frac{\partial (\mathbb{E}[\langle \mathbf{x}(t), \mathbf{x}_{\text{ref}} \rangle])^T}{\partial \mathbf{u}} \Sigma_{\langle \mathbf{x}(t), \mathbf{x}_{\text{ref}} \rangle}^{-1} \frac{\partial \mathbb{E}[\langle \mathbf{x}(t), \mathbf{x}_{\text{ref}} \rangle]}{\partial \mathbf{u}} \quad (5.3)$$

where $\langle \mathbf{x}, \mathbf{x}_{\text{ref}} \rangle = \mathbf{x}^T \mathbf{x}_{\text{ref}}$ and $\Sigma_{\langle \mathbf{x}, \mathbf{x}_{\text{ref}} \rangle}$ denotes the variance of this inner product. The inner product can be interpreted in several ways, including as the correlation or contrast between two competing states. The Fisher information lower bounds the variance of an estimate of \mathbf{u} based on measurement of the inner product.

From (6.35), and taking into account the independence of \mathbf{x} and \mathbf{x}_{ref} we have

$$\mathcal{J}_{\mathbf{u}} = \frac{\partial (\mathbb{E}[\mathbf{x}]^T \mathbb{E}[\mathbf{x}_{\text{ref}}])^T}{\partial \mathbf{u}} \Sigma_{\langle \mathbf{x}, \mathbf{x}_{\text{ref}} \rangle}^{-1} \frac{\partial \mathbb{E}[\mathbf{x}]^T \mathbb{E}[\mathbf{x}_{\text{ref}}]}{\partial \mathbf{u}} \quad (5.4)$$

(where we have dropped dependence on t for notational convenience).

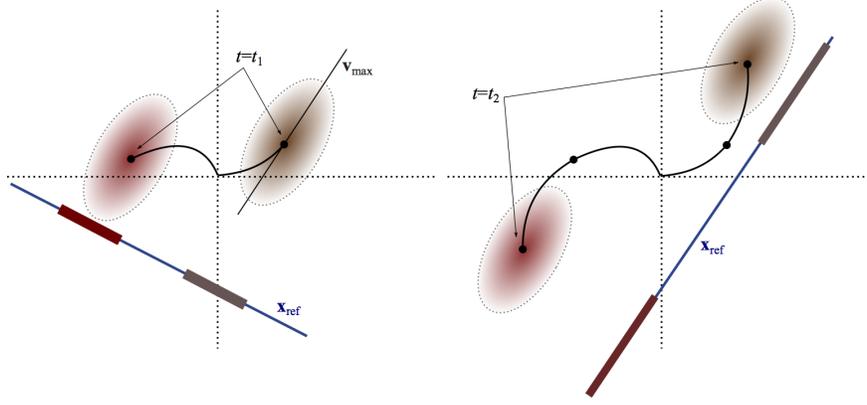


Figure 5.1: The optimal background state \mathbf{x}_{ref} amounts to a Fisher linear discriminant, onto which state distributions (induced by inputs) are projected. In the case of Gaussian noise, uncertainty can be visualized in terms of ellipsoids (with principal axis \mathbf{v}_{max}) about the mean. Since the networks are dynamic, the optimal \mathbf{x}_{ref} will vary with time as the dynamics carry the states forward.

Defining

$$\Gamma \equiv \int_0^t e^{\mathbf{A}(t-\tau)} d\tau \quad (5.5)$$

we obtain

$$\mathcal{J}_{\mathbf{u}} = \frac{\partial (\mathbb{E}[\mathbf{x}]^T \mathbb{E}[\mathbf{x}_{\text{ref}}])^T}{\partial \mathbf{u}} \Sigma_{\langle \mathbf{x}, \mathbf{x}_{\text{ref}} \rangle}^{-1} \frac{\partial (\mathbb{E}[\mathbf{x}]^T \mathbb{E}[\mathbf{x}_{\text{ref}}])}{\partial \mathbf{u}} \quad (5.6)$$

$$= \frac{\partial ((\Gamma \mathbf{B} \mathbf{u})^T \mathbf{x}_{\text{ref}})^T}{\partial \mathbf{u}} \Sigma_{\langle \mathbf{x}, \mathbf{x}_{\text{ref}} \rangle}^{-1} \frac{\partial ((\Gamma \mathbf{B} \mathbf{u})^T \mathbf{x}_{\text{ref}})}{\partial \mathbf{u}} \quad (5.7)$$

$$= \mathbf{B}^T \Gamma^T \mathbf{x}_{\text{ref}} \Sigma_{\langle \mathbf{x}, \mathbf{x}_{\text{ref}} \rangle}^{-1} \mathbf{x}_{\text{ref}}^T \Gamma \mathbf{B} \quad (5.8)$$

Using the derivation given explicitly in the Methods section, we obtain the Fisher information matrix

$$\mathcal{J}_{\mathbf{u}} = \frac{\mathbf{B}^T \Gamma^T \mathbf{x}_{\text{ref}} \mathbf{x}_{\text{ref}}^T \Gamma \mathbf{B}}{\mathbf{x}_{\text{ref}}^T \Sigma_{\mathbf{x}} \mathbf{x}_{\text{ref}}} \quad (5.9)$$

where $\Sigma_{\mathbf{x}}$ is the state covariance matrix as introduced in (6.34).

In seeking a holistic assessment of the matrix $\mathcal{J}_{\mathbf{u}}$, we employ the trace, which is the summed component-wise variance in our estimation of \mathbf{u} . Since the $\mathcal{J}_{\mathbf{u}}$ is an outer product of two vectors (scaled by the denominator) we may express its trace as their scaled inner product:

$$\text{tr}(\mathcal{J}_{\mathbf{u}}(t)) = \frac{\mathbf{x}_{\text{ref}}^T \Gamma(t) \mathbf{B} \mathbf{B}^T \Gamma(t)^T \mathbf{x}_{\text{ref}}}{\mathbf{x}_{\text{ref}}^T \Sigma_{\mathbf{x}}(t) \mathbf{x}_{\text{ref}}}, \quad (5.10)$$

where dependence on t has again been made explicit.

5.2.1.3 Linear Dynamics and Noise Ellipsoids

Figure 5.1 provides a schematic of the problem formulation. Because our dynamics are linear, at any given time t the state of the network is a Gaussian random vector. The covariance of the state can be used to parameterize a quadratic form whose level sets constitute ellipsoids that encapsulate the mean. We denote the principal eigenvector of the covariance matrix as \mathbf{v}_{max} . These ellipsoids capture the noise-driven uncertainty in the state. As we will soon see, the optimal \mathbf{x}_{ref} amounts to a Fisher linear discriminant that best disassociates two competing state distributions (ellipsoids), each associated with a different stimulus. As the network dynamics carry these trajectories forward in time, the optimal \mathbf{x}_{ref} will in general change.

5.2.1.4 Network Parameterization, Actuated Nodes and Steady-State Assumption

We will focus our attention on networks that have a Barabási-Albert (scale-free) topology [61]. The off-diagonal elements of \mathbf{A} are binary, while the diagonal elements are assigned large enough negative values to ensure stability (see Methods). The dynamics of such networks are asymptotically stable so that in the absence of stimuli and noise, all states return to the origin.

In our analysis we will vary the structure of how inputs impinge on network nodes. In particular, for an n node network, only $n_d \leq n$ nodes will receive input. These actuated nodes are sometimes referred to as ‘driver’ nodes [11, 62, 63]. We will mostly consider the case when each actuated node receives an independent input, so that

$$\mathbf{B} = \begin{bmatrix} \mathbf{I}_{n_d} \\ 0 \end{bmatrix} \quad (5.11)$$

where \mathbf{I}_{n_d} is the identity matrix of dimension n_d (the number of driven nodes).

We make the assumption that the noise covariance is always at steady-state. In concept here is that the dynamics of the network are persistently excited by ongoing noise, while receiving stimuli in a temporally punctate manner. To be mathematically precise, under this assumption (6.34) becomes:

$$\mathbf{x}(t) \sim \mathcal{N} \left(\int_0^t \mathbf{e}^{\mathbf{A}(t-\tau)} \mathbf{d}\tau \mathbf{B} \mathbf{u}, \Sigma_{\mathbf{x}(\infty)} \right), \quad (5.12)$$

We also assume the pair \mathbf{A}, \mathbf{B} is controllable, so that the controllability gramian (precisely defined later) is full-rank. A final important assumption pertains to the specification of t . In cases when t is assumed to be at steady state, we set $t = 10$ (which we find is five times longer than the time-constant of our considered networks). In other cases, we will vary t to assess the role of dynamics.

5.2.2 An optimal reference state \mathbf{x}_{ref} exists, maximizing information about \mathbf{u} .

We are interested, for the moment, in which choice of \mathbf{x}_{ref} will maximize (5.10). That is, we seek to answer the question: Of all possible background states \mathbf{x}_{ref} , which one will provide the most information about a stimulus \mathbf{u} (with its resultant output \mathbf{x}), given a readout of the inner product $\langle \mathbf{x}_{\text{ref}}, \mathbf{x} \rangle$. In order to find this ‘ideal’ reference stimulus, we transform (5.10) as follows:

$$\text{tr}(\mathcal{J}_{\mathbf{u}}) = \frac{\mathbf{x}_{\text{ref}}^T \Gamma \mathbf{B} \mathbf{B}^T \Gamma^T \mathbf{x}_{\text{ref}}}{\mathbf{x}_{\text{ref}}^T \mathbf{L} \mathbf{L}^T \mathbf{x}_{\text{ref}}} \quad (5.13)$$

where $\mathbf{L} \mathbf{L}^T$ (\mathbf{L} is lower-triangular) is the Cholesky decomposition of $\Sigma_{\mathbf{x}}$. Continuing, we have

$$\text{tr}(\mathcal{J}_{\mathbf{u}}) = \frac{\mathbf{x}_{\text{ref}}^T \Gamma \mathbf{B} \mathbf{B}^T \Gamma^T \mathbf{x}_{\text{ref}}}{\mathbf{x}_{\text{ref}}^T \mathbf{L} \mathbf{L}^T \mathbf{x}_{\text{ref}}} \quad (5.14)$$

$$= \frac{\mathbf{x}_*^T \mathbf{L}^{-T} \Gamma \mathbf{B} \mathbf{B}^T \Gamma^T \mathbf{L}^{-1} \mathbf{x}_*}{\mathbf{x}_*^T \mathbf{x}_*} \quad (5.15)$$

where $\mathbf{x}_* = \mathbf{L}^T \mathbf{x}_{\text{ref}}$. Therefore, for simplicity of notation letting $\mathbf{S} = \mathbf{L}^{-1} \Gamma \mathbf{B} \mathbf{B}^T \Gamma^T \mathbf{L}^{-T}$, we have the familiar Rayleigh quotient

$$\text{tr}(\mathcal{J}_{\mathbf{u}}) = \frac{\mathbf{x}_*^T \mathbf{S} \mathbf{x}_*}{\mathbf{x}_*^T \mathbf{x}_*} \quad (5.16)$$

whose values lie in the range $\lambda_{\min} \leq \text{tr}(\mathcal{J}_{\mathbf{u}}) \leq \lambda_{\max}$ and which achieves its extrema for $\mathbf{x}_* = \mathbf{x}_{\min}^*$ and $\mathbf{x}_* = \mathbf{x}_{\max}^*$ where \mathbf{x}_{\min}^* and \mathbf{x}_{\max}^* are the eigenvectors of \mathbf{S} associated with eigenvalues λ_{\min} and λ_{\max} respectively. We then make the reverse transformation

$$\mathbf{x}_{\text{ref}} = \mathbf{L}^{-T} \mathbf{x}_{\max}^* \quad (5.17)$$

to obtain our ideally contrasting reference state. Mathematically (and as depicted in Figure 5.1) \mathbf{x}_{ref} is in fact the Fisher linear discriminant that best separates the induced state distributions associated with any two randomly chosen inputs.

Previous results [16] have shown that an optimally informative ‘signal direction’ in a non-dynamical feedforward network is one which align with the principal axis of the noise covariance ellipsoid. Similarly, with our dynamical setup, we decided to explore the optimal \mathbf{x}_{ref} qualitatively by examining to what extent it aligns with the principal axis of the noise covariance ellipsoid (\mathbf{v}_{max} of $\Sigma_{\mathbf{x}}$ in (5.10)). The results are shown in Figure 5.2. We notice in Figure 5.2 that the ideal \mathbf{x}_{ref} changes its orientation relative to \mathbf{v}_{max} as a function of n_{d} . This orientation is virtually uncorrelated with network size and is very predictable, as we ran 30 network realizations for each n, n_{d} pair and found little variability. We hypothesized that this was due to prioritization of the fidelity of the portion of \mathbf{x}_{ref} corresponding to actuated nodes, which would explain why relatively under-actuated networks showed greater overall angular divergence between \mathbf{x}_{ref} and \mathbf{v}_{max} . This is indeed the case, as shown in Figure 5.3. We first examined actuated nodes, then non-actuated nodes, by segmenting \mathbf{x}_{ref} and \mathbf{v}_{max} into the first n_{d} elements (Figure 5.3(a)), then the last $n - n_{\text{d}}$ elements (Figure 5.3(b)). Clearly, the actuated part of \mathbf{x}_{ref} is required to be much more similar to the corresponding part of \mathbf{v}_{max} than is true for the non-actuated part.

Aside from the dependence of the optimal \mathbf{x}_{ref} on input structure (particularly n_{d}), we also analyzed how the orientation of \mathbf{x}_{ref} , relative to \mathbf{v}_{max} , changes with time. Since, as mentioned above, we are working in a dynamical regime, a time-dependent analysis is straightforward. To this end, we evaluated the orientation of \mathbf{x}_{ref} , relative to \mathbf{v}_{max} , at several time points, using the same methodology employed above, with the results shown in Figure 5.4. We see that the orientation of \mathbf{x}_{ref} relative to \mathbf{v}_{max} does indeed change with time, apparently smoothly,

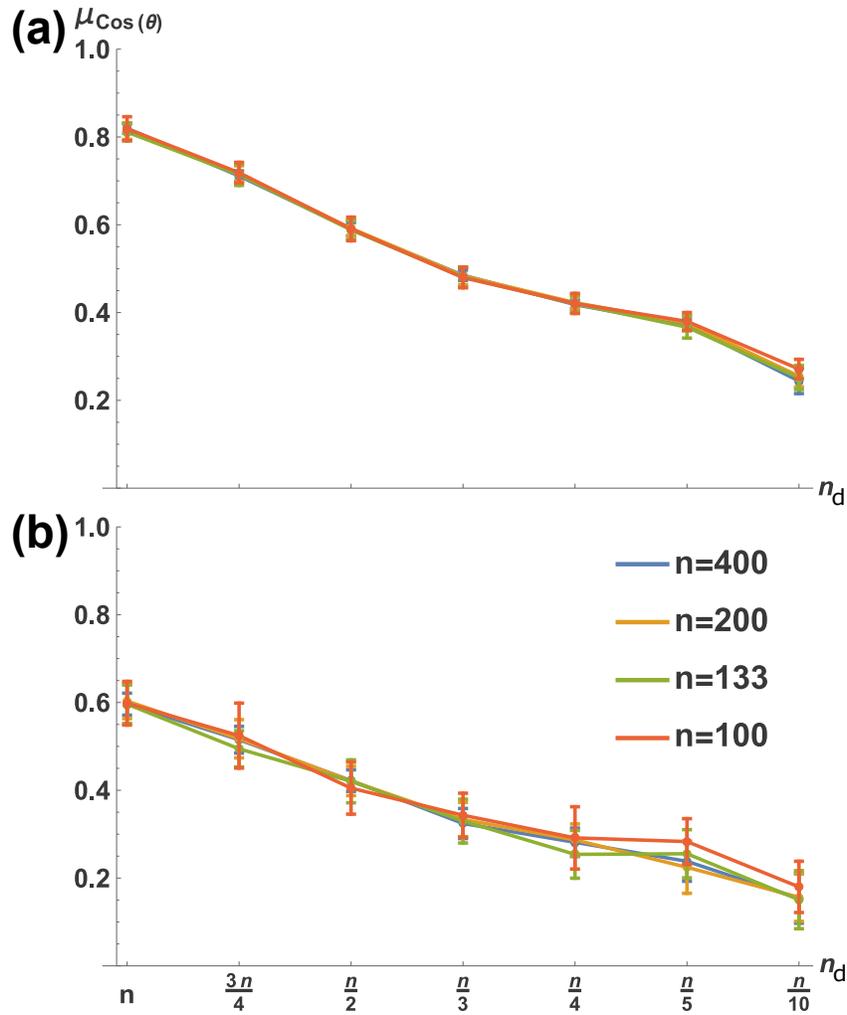


Figure 5.2: Fidelity of optimally contrasting reference state \mathbf{x}_{ref} to system noise covariance decreases monotonically with n_d . Shown is how \mathbf{x}_{ref} aligns with the principal eigenvector (denoted \mathbf{v}_{max}) of noise covariance matrix $\Sigma_{\mathbf{w}}$. $\mu_{\text{Cos}(\theta)}$ is the mean, over 30 network realizations, of the cosine of the angular difference (θ) between \mathbf{x}_{ref} and \mathbf{v}_{max} . Error bars are standard deviations. 30 realizations were evaluated for (a) identity and (b) random \mathbf{B} matrices.

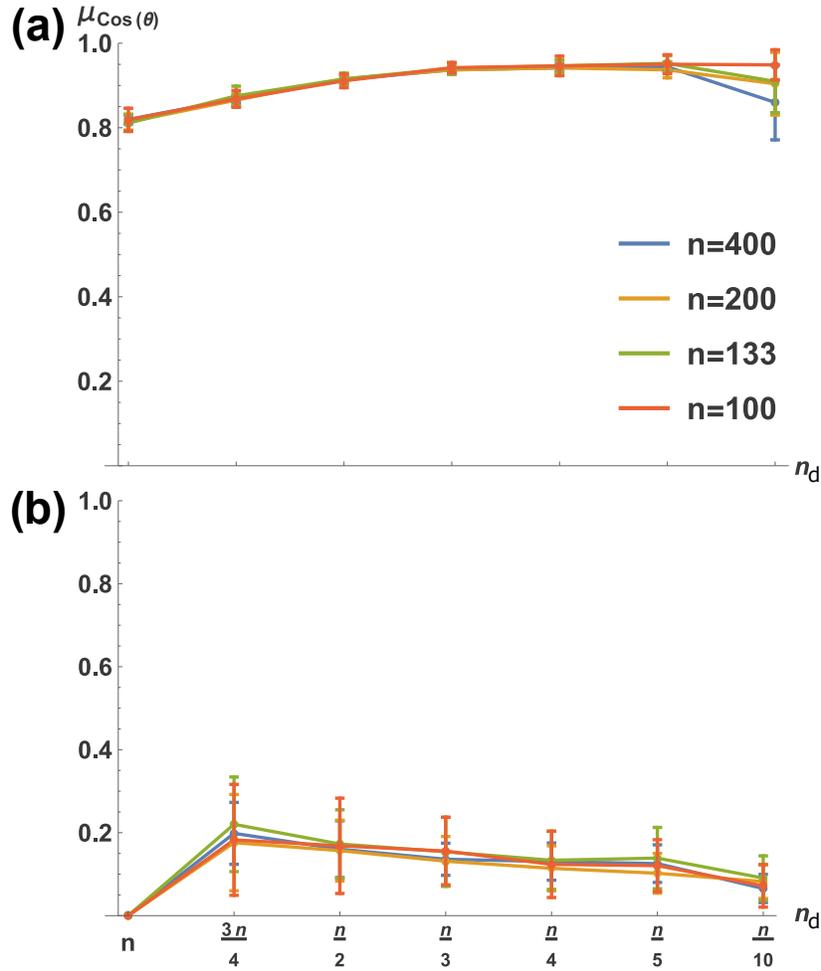


Figure 5.3: Actuated nodes of the ideal background (\mathbf{x}_{ref}) are ‘required’ to be aligned with noise; non-actuated nodes are not. Shown is alignment of \mathbf{x}_{ref} with principal noise covariance direction \mathbf{v}_{max} (as in Fig. 5.2); here \mathbf{x}_{ref} and \mathbf{v}_{max} are partitioned so that (a) reflects only actuated and (b) only non-actuated nodes. $\mu_{\text{Cos}(\theta)}$ is as in Fig. 5.2, again for 30 network realizations.

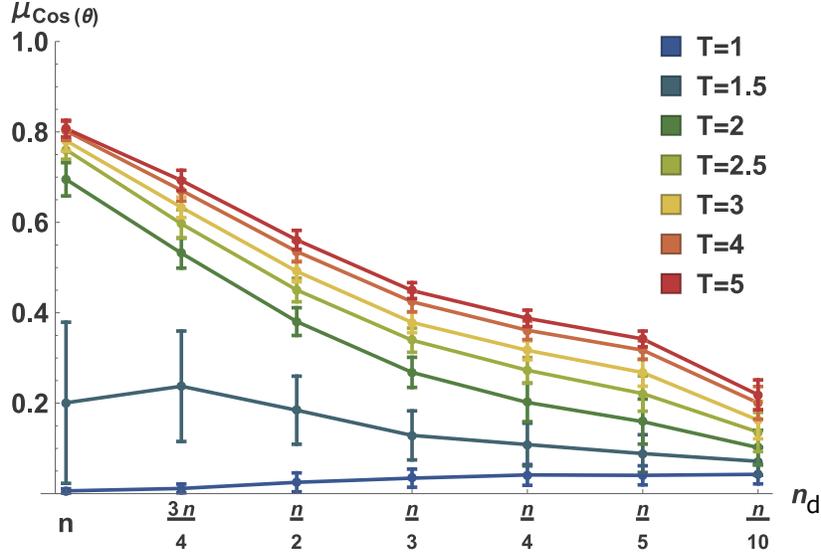


Figure 5.4: The ideal contrast becomes more aligned with noise covariance as time progresses. Shown is the time-evolution of the relative orientation between the optimally contrasting state \mathbf{x}_{ref} and the principal noise eigenvector \mathbf{v}_{max} . At lower values of T , \mathbf{x}_{ref} is nearly orthogonal to \mathbf{v}_{max} , while as T gets larger, \mathbf{x}_{ref} becomes much more aligned with \mathbf{v}_{max} , although this alignment approaches a limit, which also varies monotonically with n_d .

and that \mathbf{x}_{ref} becomes more similar to \mathbf{v}_{max} as time advances. This is especially true for fully- or nearly fully-actuated networks, but is generally true for all input scenarios.

Thus, the optimally contrasting background/reference state is fundamentally dependent on the input structure of the network and the time evolution of network dynamics.

5.2.3 An optimal reference input \mathbf{u}_{ref} exists, maximizing information about

u.

We expanded our inquiry to analyze admissible reference inputs \mathbf{u}_{ref} which could give rise to an optimally contrasting state \mathbf{x}_{ref} . More generally, we asked: Of all possible stimuli, does there exist a best one \mathbf{u}_{ref} , resulting in an output $\tilde{\mathbf{x}}_{\text{ref}}$, that provides information about all

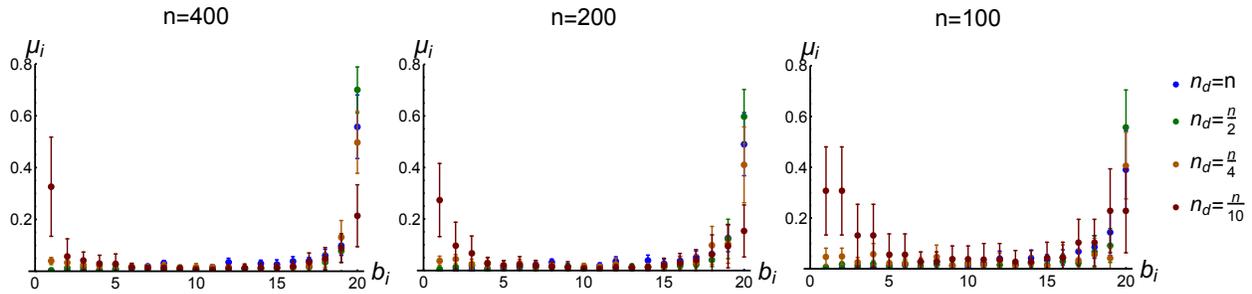


Figure 5.5: The optimally contrasting input \mathbf{u}_{ref} targets network ‘hubs’, but to a degree which varies with n_d . 30 networks were realized with for each size (n) and driver node (n_d) combination. For a given network size, the graph shows the mean (μ_i) of the squared entries of the (normalized) optimal \mathbf{u}_{ref} . The entries of \mathbf{u}_{ref} are sorted according to the degree of targeted nodes (abscissa is a percentile, binned in increments of 5%, so that each b_i represents 5% of the nodes). Note that when $n = 100$ and $n_d = \frac{n}{10}$ there are twice as many bins as controlled nodes, hence the duplicity of values.

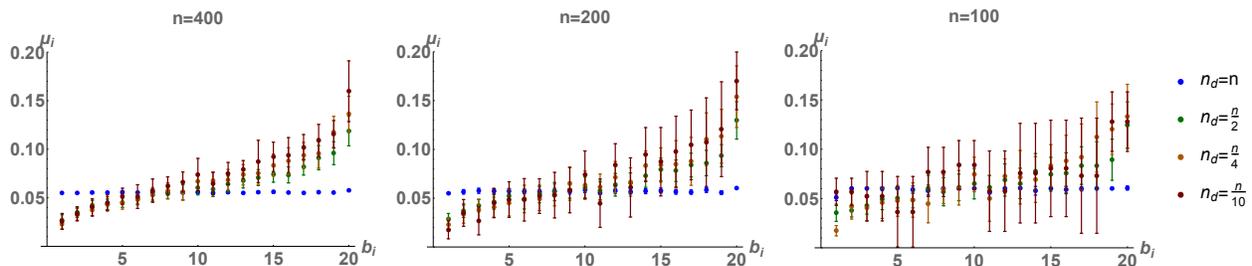


Figure 5.6: Same setup as in Figure 5.5, but simulations are run for randomly connected networks (edge probability $p = .5$). Note the much smaller range of values on the vertical axis when compared with Figure 5.5. Nodes of high degree (‘hubs’) are targeted, but to a lesser extent than for the scale-free networks. Skewness of the graphs is inversely related to n_d ; that is, it is less necessary to target hubs for more fully actuated networks, until at $n_d = n$, \mathbf{u}_{ref} is essentially uniform.

others. In this formulation, $\tilde{\mathbf{x}}_{\text{ref}}$ is not longer unconstrained, but rather is determined by:

$$\tilde{\mathbf{x}}_{\text{ref}} = \Gamma \mathbf{B} \mathbf{u}_{\text{ref}} \quad (5.18)$$

Using (5.18), we can find the optimal reference stimulus via a similar sequence of steps as in the previous subsection, defining

$$\text{tr}(\mathcal{J}_{\mathbf{u}}) = \frac{\mathbf{u}_{\text{ref}}^T \mathbf{B}^T \Gamma^T \Gamma \mathbf{B} \mathbf{B}^T \Gamma^T \Gamma \mathbf{B} \mathbf{u}_{\text{ref}}}{\mathbf{u}_{\text{ref}}^T \mathbf{B}^T \Gamma^T \Sigma_{\mathbf{x}} \Gamma \mathbf{B} \mathbf{u}_{\text{ref}}} \quad (5.19)$$

$$= \frac{\mathbf{u}_{\text{ref}}^T \mathbf{B}^T \Gamma^T \Gamma \mathbf{B} \mathbf{B}^T \Gamma^T \Gamma \mathbf{B} \mathbf{u}_{\text{ref}}}{\mathbf{u}_{\text{ref}}^T \mathbf{L} \mathbf{L}^T \mathbf{u}_{\text{ref}}} \quad (5.20)$$

where $\mathbf{L} \mathbf{L}^T$ (\mathbf{L} is lower-triangular) is the Cholesky decomposition of $\mathbf{B}^T \Gamma^T \Sigma_{\mathbf{x}} \Gamma \mathbf{B}$, which is positive-definite (a requirement for this decomposition) since covariance matrix $\Sigma_{\mathbf{x}}$ is inherently positive-definite and thus can be Cholesky decomposed into $\mathbf{L}_{\Sigma} \mathbf{L}_{\Sigma}^T$, so that the matrix $\mathbf{B}^T \Gamma^T \Sigma_{\mathbf{x}} \Gamma \mathbf{B}$ can be written $\mathbf{Q} \mathbf{Q}^T$ for $\mathbf{Q} = \mathbf{B}^T \Gamma^T \mathbf{L}_{\Sigma}$ and is thus positive-semidefinite, while the full-rank condition of \mathbf{Q} ensures positive-definiteness.

Defining $\mathbf{S}_{\mathbf{u}} \equiv \mathbf{L}^{-T} \mathbf{B}^T \Gamma^T \Gamma \mathbf{B} \mathbf{B}^T \Gamma^T \Gamma \mathbf{B} \mathbf{L}^{-1}$, we arrive at

$$\text{tr}(\mathcal{J}_{\mathbf{u}}) = \frac{\mathbf{u}_*^T \mathbf{S}_{\mathbf{u}} \mathbf{u}_*}{\mathbf{u}_*^T \mathbf{u}_*} \quad (5.21)$$

whose values lie in the range $\lambda_{\min} \leq \text{tr}(\mathcal{J}_{\mathbf{u}}) \leq \lambda_{\max}$ and which achieves its extrema for $\mathbf{u}_* = \mathbf{u}_{\min}^*$ and $\mathbf{u}_* = \mathbf{u}_{\max}^*$ where \mathbf{u}_{\min}^* and \mathbf{u}_{\max}^* are the eigenvectors of $\mathbf{S}_{\mathbf{u}}$ associated with eigenvalues λ_{\min} and λ_{\max} , respectively. We then make the reverse transformation $\mathbf{u}_{\text{ref}} = \mathbf{L}^{-T} \mathbf{u}_{\max}^*$ to obtain our ideally contrasting reference input.

We pause for a moment to consider the significance of this ‘optimal’ \mathbf{u}_{ref} (i.e. the eigenvector of $\mathbf{S}_{\mathbf{u}}$ which optimizes (5.21)). The existence of such an optimum means that for a given network, there is one input whose induced state best contrasts those of all other inputs.

5.2.4 The optimally contrasting input targets specific nodes in a concentrated manner, but not necessarily nodes of highest degree

We sought to characterize the ‘optimally informative’ \mathbf{u}_{ref} by examining its entries (recall that we are in the domain of constant inputs) as they relate to the connectivity degree of actuated nodes. Clearly, \mathbf{u}_{ref} has cardinality n_{d} (see (5.11)). Since, then, there is a one-to-one relationship between the n_{d} entries of \mathbf{u}_{ref} and the driven nodes, we are able to learn about which nodes may be specially ‘targeted’ by an optimally contrasting \mathbf{u}_{ref} . Intuition would suggest that the targeted nodes would simply be the hubs, that is, that the higher the degree of a node, the higher the value of the corresponding entry of \mathbf{u}_{ref} . This is borne out in simulation, but to an extent which varies consistently with network size (n) and n_{d} . Examining Figure 5.5, we see that for larger networks wherein all nodes are controlled, nearly all of the large entries of \mathbf{u}_{ref} are concentrated toward nodes in the top 5% by degree ranking (i.e. the hubs), while as we control fewer nodes, a majority of the large entries are directed toward the hubs, but this majority becomes smaller as n_{d} decreases. Also, looking at the different network sizes, we see that, in general, larger networks show a more pronounced ‘targeting’ of the hubs, while in smaller networks the hubs are still targeted but to a lesser extent. It should be pointed out that \mathbf{u}_{ref} is unitary, meaning there is an essential trade-off between how much energy can be focused on hubs and how much can be focused elsewhere (as is easily seen in Figure 5.5), so that in very hub-oriented scenarios (i.e. large networks with high fraction of controlled nodes), \mathbf{u}_{ref} is nearly a standard basis vector, while in smaller networks wherein fewer nodes are controlled, \mathbf{u}_{ref} is more homogenous.

For comparison's sake, we ran simulations with randomly connected (with .5 edge probability) instead of scale-free networks. These networks were also undirected and rendered stable by the same method (described in Methods). We did use small ($<.1$) positive edge weights, rather than unitary weights, for these networks to render their analysis more numerically tractable. We see in Figure 5.6 that the optimal \mathbf{u}_{ref} also tends to target nodes of higher degree in random networks, but to a much lesser extent than for scale-free networks. We hypothesize that this is because the degree distribution for scale-free networks is given by a power-law, which means there are many nodes of very low degree, and a few of very high degree. Random networks have a binomial degree distribution, with more nodes of average degree and none of very high degree. Thus, it may be less crucial for the input to target the higher-degree nodes in random networks, simply because the higher-degree nodes are not much higher-degree nodes. In the random networks, we see a skewing of the values of \mathbf{u}_{ref} which is inversely correlated with n_d . That is, for less-actuated networks, the hubs tend to be more targeted, while for more fully-actuated networks, this targeting becomes less pronounced until at the limiting case ($n_d = n$), the entries of \mathbf{u}_{ref} are all nearly identical.

5.2.5 Information Spectra (of $\mathbf{S}_{\mathbf{u}}$) are Sensitive to Network Parameterization

We now turn our attention to the problem of comparing different networks according to their information capacity, as quantified by $\mathcal{I}_{\mathbf{u}}$. For this we examine the information capacity by varying \mathbf{u}_{ref} in (5.21), where the intuitive strategy is to let \mathbf{u}_* range over the eigenvectors of $\mathbf{S}_{\mathbf{u}}$. Thus, a holistic characterization of $\text{tr}(\mathcal{I}_{\mathbf{u}})$ is provided simply by the eigenvalue spectrum of $\mathbf{S}_{\mathbf{u}}$ (recall that (5.21) takes on the value λ_i , the i^{th} eigenvalue, when \mathbf{u}_{ref} is the i^{th} eigenvector), heretofore termed the information spectrum of a network.

We obtained a distribution of information spectra for several network parametrizations. We here restricted our attention to steady state characterizations. Each distribution amounts

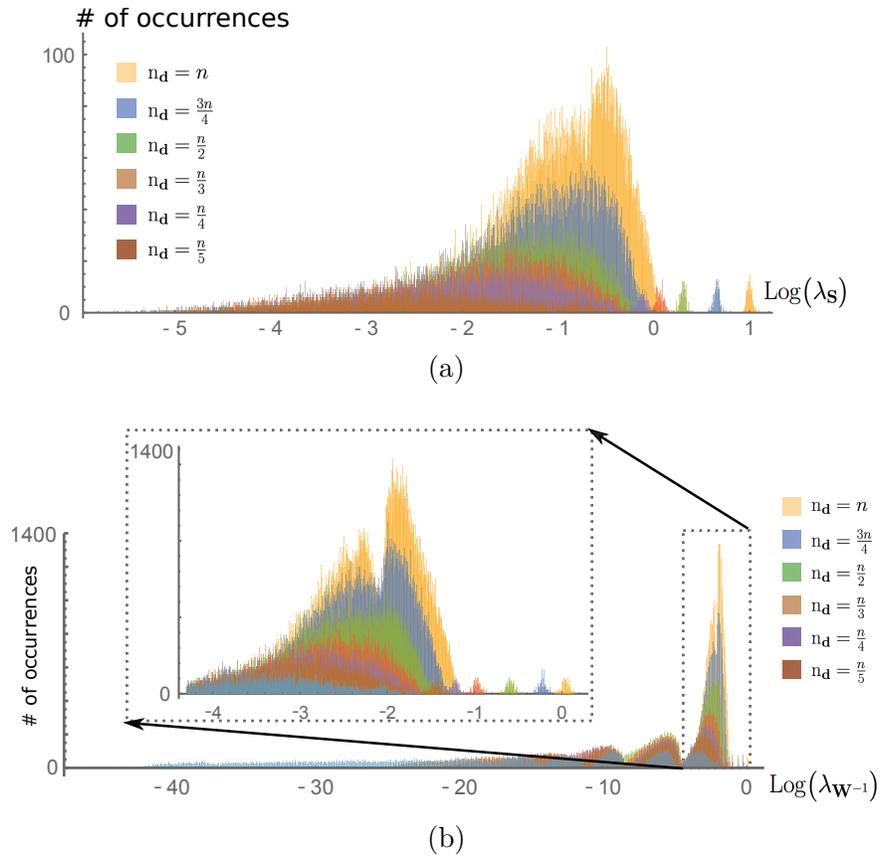


Figure 5.7: 5.7a: Information spectra as function of number of actuated nodes (distributions aggregated over $n = 100, 200, 300, 400$). Spectra consist of a primary mode and a smaller secondary mode. 5.7b: Spectra of the controllability gramian for different fractions of actuated nodes. As noted in previous work, these spectra display an increasing number of modes as n_d decreases. The principal mode is inset. Comparing to the information spectra in Figure 5.7a, we see that information spectra show marked similarity to first mode of control spectra, and both spectra reveal outlying, small modes corresponding to easiest (control) and most informative (information) directions.

to an empirical probability distribution of the eigenvalues of $\mathbf{S}_{\mathbf{u}}$ over (random) network realizations. We assumed used zero-mean, unit-variance, uncorrelated noise (i.e. $\mathbb{E}[\mathbf{w}_i \mathbf{w}_j] = \mathbb{E}[\mathbf{w}_i] \mathbb{E}[\mathbf{w}_j] = 0 \forall i, j \in \{1, \dots, n\}, i \neq j$ and $\mathbb{E}[\mathbf{w}_i \mathbf{w}_i] = 1 \forall i \in \{1, \dots, n\}$), though similar results were obtained for correlated noise.

Figure 5.7a depicts the information spectra for several fractions of actuated (driver) nodes (aggregates over several values of n). A first observation is the presence of a small, secondary mode to the right of the principal mode. This secondary mode reflects the presence of a few particularly salient inputs that most informatively correlate with all others. It is notable that this mode, which represents the largest eigenvalue of $\mathbf{S}_{\mathbf{u}}$, systematically decreases with smaller values of $n_{\mathbf{d}}$. Certain intuition about these observations can be deduced from the rich body of work on spectra of random matrices. One such spectral characterization [64] shows that the principal eigenvalue of the adjacency matrix (here denoted \mathbf{A}) for undirected, binary scale-free networks (such as those used for our simulations, with the exception that the diagonal of our \mathbf{A} is adjusted, as described in Methods, to ensure stability) approximates $n^{\frac{1}{4}}$, where n is the number of network nodes. Further, recent work [65] has shown that this maximum eigenvalue, for weighted scale-free networks with expected degree distributions, varies monotonically with the maximum node degree. Maximum degree, in turn, increases dramatically as n increases, because of the preferential attachment-based network creation algorithm [66]. Thus we would expect the spectrum of $\mathbf{S}_{\mathbf{u}}$, and in particular its principle eigenvalue, to depend on effective network size, which itself depends on $n_{\mathbf{d}}$ (see (5.11), and note the effect of \mathbf{B} on (5.20)). This makes sense intuitively, as well: We would expect higher-dimensional input spaces to admit a richer set of encoded representations.

Further, as $n_{\mathbf{d}}$ decreases, the distribution of the main mode becomes broader and more entropic. No additional modes or ‘humps’ appear as $n_{\mathbf{d}}$ varies, a point we will return to shortly.

5.2.6 Information Spectra are Related to the Controllability Gramian

As noted previously, the information spectrum is fundamentally time-varying (governed by the network dynamics, driven by the input in question). We were particularly interested in the relationship between the information spectrum and that of the controllability gramian matrix

$$\mathbf{W}(t) = \left(\int_0^t \mathbf{e}^{\mathbf{A}(t-\tau)} \mathbf{B} \mathbf{B}^T \mathbf{e}^{\mathbf{A}^T(t-\tau)} d\tau \right),$$

which also fundamentally characterizes the input-output relationship of a linear (networked) system. Indeed, it is well known that in the limit as $t \rightarrow \infty$, the gramian is exactly equivalent to $\Sigma_{\mathbf{x}}$, i.e., the denominator of $\mathcal{S}_{\mathbf{u}}$. Thus, we sought to compare the information spectrum to that of $\mathbf{W}(\infty)$.

The gramian matrix has been a pivotal entity in the analysis of linear systems and similarly modeled networks [41, 67, 68], including certain types of brain networks [50, 69]. Recent theoretical work [61] has characterized the nature of the infinite-time gramian spectrum as a function of the number of driven nodes (n_d). It is shown there that for small fractions of driven nodes the spectrum manifests a series of modes or ‘humps,’ over which eigenvalues are randomly distributed (over network realizations). As is well known in linear systems theory, the magnitude of a gramian eigenvalue determines the minimum input energy needed to reach the unit hypersphere in the direction of its associated eigenvector. Thus, the principal mode of the gramian spectrum describes those directions that are ‘easiest’ to induce.

Figure 5.7b depicts the gramian spectrum for the same networks as in Figure 5.7a(a) (i.e., with varying fraction of actuated nodes). The aforementioned modes are readily evident. What is notable from this figure is the correspondence between the information spectra to the two rightmost modes of the gramian spectrum (that is, the principal mode and the much smaller

mode at far right). In interpreting this result, it is important to note that the information and gramian spectra are of different dimensions ($\mathbf{W}^{-1}(t_f) \in \mathbb{R}^{n \times n}$ while $\mathbf{S}_{\mathbf{u}} \in \mathbb{R}^{m \times m}$). This is because the information spectrum captures only constant inputs, thus for a fixed time the state is restricted to an m -dimensional subspace. In this sense, we postulate that the principal mode of the gramian spectrum corresponds not simply to the ‘easiest’ to reach directions, but also those associated with constant (m -dimensional) inputs.

Let us now seek to understand this numerical correspondence between control and information, shown in Figure 5.7b, at a conceptual level. What does it mean that the easiest directions of control (quantified by the largest eigenvalues of \mathbf{W}^{-1}) and network information (quantified by $\mathcal{I}_{\mathbf{u}}$) show such similarity? We hypothesize that this correspondence may be indicative of an underlying link between controllability metrics and information-based analyses generally. Indeed, this is not a novel idea; the mathematical basis for this link has been explored [70, 71] in contexts different, but related, to ours. We can summarize the essence of these discussions, as it relates to our formulation, simply by noting that $\mathcal{I}_{\mathbf{u}}$ depends fundamentally on a derivative of the state (to be more precise, an inner product of two states) with respect to \mathbf{u} . Thus when system dynamics are such that incremental changes are made to \mathbf{u} which result in large changes to the state, it increases the informational value. This information is, to some extent, a measure of network sensitivity to its inputs, and sensitivity to inputs is, of course, exactly what controllability analysis seeks to quantify.

5.3 Discussion

We developed an analysis to quantify the amount of information about an input \mathbf{u} that can be gleaned from the contrast/correlation between its induced state $\mathbf{x}_{\mathbf{u}}$ and a reference or background state \mathbf{x}_{ref} . Our analysis shows that there exists an optimally informative \mathbf{x}_{ref}

in this context. This theoretical result reinforces intuition about how proper choice of a contrasting background might enable more rapid decoding and subsequent processing of input stimuli. We showed that the orientation of \mathbf{x}_{ref} relative to the principal axis of noise covariance decreased monotonically with increasing fraction of nodes actuated and that this separation also decreased over time, but to an extent limited by n_d . This dynamical relationship between the informational optimum and the noise covariance is complementary to results based on static models [16].

We expanded our inquiry to examine the \mathbf{u}_{ref} which would give rise to \mathbf{x}_{ref} . We found that the optimal \mathbf{u}_{ref} tends to target network hubs, but in a way which varies consistently with number of nodes driven n_d (See Figure 5.5). We then derived an information spectrum that characterizes the full encoding capacity (in terms of inner product readout) of inputs. We showed that this spectrum has nuanced dependency on network size and fraction of driven nodes, with the presence of a low-dimensional set of inputs to which networks appear particularly well-tuned. Further, we reconciled the information encoding of a network with its control-theoretic properties, which characterize how the ‘energy’ of an input allow for the state space to be traversed. Our results suggest that inputs that produce ‘easy’ state excursions—recall that these inputs are postulated to be constant or near-constant (see Section 5.2.6)—are also those that are well-encoded.

It may reasonably be asked why we have chosen inputs to be constant in the overall paradigm. There are a couple of reasons for this. At a conceptual level, our information analysis is fundamentally predicated on the derivative $\frac{\partial \langle \mathbf{x}, \mathbf{x}_{\text{ref}} \rangle}{\partial \mathbf{u}}$. That is, we seek to quantify the extent to which changes in the projection of system state \mathbf{x} onto background \mathbf{x}_{ref} reflect incremental changes in \mathbf{u} . This is straightforward for a constant \mathbf{u} —we can easily imagine a slight perturbation of a vector of constant values. However, it is more problematic for a time-varying $\mathbf{u}(t)$. What does it mean to make an incremental change in the function $\mathbf{u}(t)$? Here

we are dealing with a variational problem in the infinite-dimensional function space, and it makes intuition more difficult. Also, the mathematics of including a time-varying $\mathbf{u}(t)$ creates difficulties. Examining (6.35), we see that taking a derivative with respect to $\mathbf{u}(t)$ presents us with the task of taking the derivative of one function of t ($\mathbb{E}[\langle \mathbf{x}(t), \mathbf{x}_{\text{ref}} \rangle]$) with respect to another function ($\mathbf{u}(t)$). In general, for functions $f(t)$ and $g(t)$, $\frac{\partial g(t)}{\partial f(t)} = \frac{\partial g(t)}{\partial t} \cdot \frac{\partial t}{\partial f(t)} = \frac{\frac{\partial g(t)}{\partial t}}{\frac{\partial f(t)}{\partial t}} = \frac{g'(t)}{f'(t)}$. Thus $\mathcal{I}_{\mathbf{u}}$ would become dependent on $\mathbf{u}'(t)$. But we conduct our analysis with respect to the objective of learning about \mathbf{u} from a ‘readout’ of only the projection of $\mathbf{x}(t)$ onto \mathbf{x}_{ref} . To assume a knowledge of the time derivative of $\mathbf{u}(t)$ changes the setup completely. One way around this dilemma would be to project $\mathbf{u}(t)$ onto a set of orthogonal basis functions (a Fourier basis, for example). If we denote a vector of basis functions (truncated so as not to be infinite) as $\mathbf{h}(t)$, we can approximate (almost) any $\mathbf{u}(t)$ by $\mathbf{U}\mathbf{h}(t)$, where \mathbf{U} is a constant projection, or coefficient, matrix. Then, $\mathcal{I}_{\mathbf{u}}$ becomes linear in \mathbf{U} , so $\frac{\partial (\mathbb{E}[\langle \mathbf{x}(t), \mathbf{x}_{\text{ref}} \rangle])^T}{\partial \mathbf{U}}$ yields an expression which is independent of coefficient matrix \mathbf{U} , and the basic formulation is preserved, with the change that instead of seeking to infer constant input \mathbf{u} via the state projection, we seek to infer coefficient matrix \mathbf{U} , with which a time-varying $\mathbf{u}(t)$ can be reconstructed. We derive the basic equation for the Fisher information in such a paradigm in the Methods section, but we do not characterize solutions in this work.

Having highlighted the results from the exploration of $\mathcal{I}_{\mathbf{u}}$, let us take a slightly higher-level look at the information processing which $\mathcal{I}_{\mathbf{u}}$ quantifies. Considering the inner product as the ‘readout’ (which forms the basis of information measure $\mathcal{I}_{\mathbf{u}}$) is intuitive since it measures correlation/contrast between two competing representations of a stimulus. In this sense, it is a highly condensed representation of potentially high-dimensional stimuli. However, it is far from clear whether a network itself could accomplish this readout, and whether this is in fact a reasonable strategy for actual information processing tasks such as input classification. The linearity of the model considered is certainly a limiting factor in this regard.

Nonetheless, we believe our results highlight an interesting direction toward analyzing not simply the structural aspects of networks, but also their dynamics and ultimately their functionality. It is straightforward to envision generalizing our framework to examine other network topologies, dynamical nonlinearities and wider time-scales, as well as alternative information metrics. These types of analyses can shed light on the functional advantages of biological networks (e.g., those in the brain) and/or principles for guiding the design of engineered systems.

5.4 Methods

5.4.1 Derivation of $\mathcal{J}_{\mathbf{u}}$

Proceeding from (5.7), we make use of the fact that $\Sigma_{\langle \mathbf{x}, \mathbf{x}_{\text{ref}} \rangle}$ is a scalar (being the variance of a scalar inner product), so that

$$\mathcal{J}_{\mathbf{u}} = \frac{\mathbf{B}^T \Gamma^T \mathbf{x}_{\text{ref}} \mathbf{x}_{\text{ref}}^T \Gamma \mathbf{B}}{\Sigma_{\langle \mathbf{x}, \mathbf{x}_{\text{ref}} \rangle}} \quad (5.22)$$

In seeking a holistic assessment of the matrix $\mathcal{J}_{\mathbf{u}}$, we employ the trace, which is the summed component-wise variance in our estimation of \mathbf{u} . Since $\mathcal{J}_{\mathbf{u}}$ is an outer product of two vectors we may express its trace as their inner product

$$\text{tr}(\mathcal{J}_{\mathbf{u}}) = \frac{\mathbf{x}_{\text{ref}}^T \Gamma \mathbf{B} \mathbf{B}^T \Gamma^T \mathbf{x}_{\text{ref}}}{\Sigma_{\langle \mathbf{x}, \mathbf{x}_{\text{ref}} \rangle}} \quad (5.23)$$

We now examine the inner product variance $\Sigma_{\langle \mathbf{x}, \mathbf{x}_{\text{ref}} \rangle}$. It is straightforward to obtain

$$\Sigma_{\langle \mathbf{x}, \mathbf{x}_{\text{ref}} \rangle} = \mathbb{E}[\langle \mathbf{x}, \mathbf{x}_{\text{ref}} \rangle^2] - (\mathbb{E}[\langle \mathbf{x}, \mathbf{x}_{\text{ref}} \rangle])^2 \quad (5.24)$$

$$= \mathbb{E}[(\mathbf{x}^T \mathbf{x}_{\text{ref}})^T (\mathbf{x}^T \mathbf{x}_{\text{ref}})] - (\mathbb{E}[\mathbf{x}]^T \mathbb{E}[\mathbf{x}_{\text{ref}}])^2 \quad (5.25)$$

$$= \mathbf{x}_{\text{ref}}^T \mathbb{E}[\mathbf{x} \mathbf{x}^T] \mathbf{x}_{\text{ref}} - (\mathbf{u}^T \mathbf{B}^T \Gamma^T \mathbf{x}_{\text{ref}})^2 \quad (5.26)$$

Note that $\mathbb{E}[\mathbf{x} \mathbf{x}^T]$ is the correlation matrix of \mathbf{x} , so that

$$\Sigma_{\mathbf{x}} = \mathbb{E}[(\mathbf{x} - \mathbb{E}[\mathbf{x}])(\mathbf{x} - \mathbb{E}[\mathbf{x}])^T] \quad (5.27)$$

$$= \mathbb{E}[\mathbf{x} \mathbf{x}^T] - \Gamma \mathbf{B} \mathbf{u} \mathbf{u}^T \mathbf{B}^T \Gamma^T \quad (5.28)$$

Therefore

$$\mathbb{E}[\mathbf{x} \mathbf{x}^T] = \Sigma_{\mathbf{x}} + \Gamma \mathbf{B} \mathbf{u} \mathbf{u}^T \mathbf{B}^T \Gamma^T \quad (5.29)$$

and combining (5.29) with (5.26) we have

$$\begin{aligned} \Sigma_{\langle \mathbf{x}, \mathbf{x}_{\text{ref}} \rangle} &= \mathbf{x}_{\text{ref}}^T (\Sigma_{\mathbf{x}} + \Gamma \mathbf{B} \mathbf{u} \mathbf{u}^T \mathbf{B}^T \Gamma^T) \mathbf{x}_{\text{ref}} \\ &\quad - (\mathbf{u}^T \mathbf{B}^T \Gamma^T \mathbf{x}_{\text{ref}})^2 \end{aligned} \quad (5.30)$$

$$= \mathbf{x}_{\text{ref}}^T \Sigma_{\mathbf{x}} \mathbf{x}_{\text{ref}} \quad (5.31)$$

Thus, plugging (5.31) into (5.22) we have the Fisher information matrix

$$\mathcal{J}_{\mathbf{u}} = \frac{\mathbf{B}^T \Gamma^T \mathbf{x}_{\text{ref}} \mathbf{x}_{\text{ref}}^T \Gamma \mathbf{B}}{\mathbf{x}_{\text{ref}}^T \Sigma_{\mathbf{x}} \mathbf{x}_{\text{ref}}} \quad (5.32)$$

as given in the main body of the text.

5.4.2 Network parameterization and simulations

To ensure stability, it is sufficient [61] to ensure that, $\forall i \in \{1, \dots, n\}$, the i^{th} diagonal element of binary adjacency matrix \mathbf{A} is at least as negative as the sum of the non-diagonal elements in row i . That is, $\sum_{j \neq i} \mathbf{A}_{i,j} + \mathbf{A}_{i,i} < 0$. Accordingly, in constructing networks, we first created a scale-free degree distribution and then formed a corresponding random graph, thus prescribing adjacency matrix \mathbf{A} . Next we simply assigned $\mathbf{A}_{i,i} = -(\sum_{j \neq i} \mathbf{A}_{i,j} + \delta_i)$, where each δ_i was picked at random from $(0, 1)$.

Creation of these adjacency matrices and the \mathbf{B} matrices, as well as the calculations of optimally contrasting background state \mathbf{x}_{ref} and reference stimulus \mathbf{u}_{ref} , with the associated statistical analyses, were performed using Mathematica, with the exception of the calculations of controllability gramians, which were done by exporting these matrices to MATLAB, and using the `lyap()` command.

5.4.3 Time-varying \mathbf{u}

To evaluate the information

$$\mathcal{I}_{\mathbf{u}} = \frac{\partial (\mathbb{E}[\mathbf{x}]^T \mathbb{E}[\mathbf{x}_{\text{ref}}])^T}{\partial \mathbf{u}} \Sigma_{\langle \mathbf{x}, \mathbf{x}_{\text{ref}} \rangle}^{-1} \frac{\partial \mathbb{E}[\mathbf{x}]^T \mathbb{E}[\mathbf{x}_{\text{ref}}]}{\partial \mathbf{u}} \quad (5.33)$$

for time varying $\mathbf{u} = \mathbf{u}(\tau)$, $\tau \in [0, t]$, we must reformulate $\mathbb{E}[\mathbf{x}(t)]$ as follows

$$\mathbb{E}[\mathbf{x}(t)] = \int_0^t e^{\mathbf{A}(t-\tau)} \mathbf{B} \mathbf{u}(\tau) d\tau \quad (5.34)$$

Now, recalling that \mathbf{x}_{ref} is idealized and not subject to noise so that $\mathbb{E}[\mathbf{x}_{\text{ref}}] = \mathbf{x}_{\text{ref}}$, and that $\Sigma_{\langle \mathbf{x}, \mathbf{x}_{\text{ref}} \rangle}$ is a scalar, we obtain:

$$\mathcal{J}_{\mathbf{u}} = \frac{\frac{\partial \left(\int_0^t e^{\mathbf{A}(t-\tau)} \mathbf{B} \mathbf{u}(\tau) d\tau \right)^T}{\partial \mathbf{u}(\tau)} \mathbf{x}_{\text{ref}} \mathbf{x}_{\text{ref}}^T \frac{\partial \left(\int_0^t e^{\mathbf{A}(t-\tau)} \mathbf{B} \mathbf{u}(\tau) d\tau \right)}{\partial \mathbf{u}(\tau)}}{\Sigma_{\langle \mathbf{x}, \mathbf{x}_{\text{ref}} \rangle}} \quad (5.35)$$

which corresponds to (9) for the constant \mathbf{u} case, and the trace, corresponding to (10), becomes

$$\text{tr}(\mathcal{J}_{\mathbf{u}}) = \frac{\mathbf{x}_{\text{ref}}^T \frac{\partial \left(\int_0^t e^{\mathbf{A}(t-\tau)} \mathbf{B} \mathbf{u}(\tau) d\tau \right)}{\partial \mathbf{u}(\tau)} \frac{\partial \left(\int_0^t e^{\mathbf{A}(t-\tau)} \mathbf{B} \mathbf{u}(\tau) d\tau \right)^T}{\partial \mathbf{u}(\tau)} \mathbf{x}_{\text{ref}}}{\Sigma_{\langle \mathbf{x}, \mathbf{x}_{\text{ref}} \rangle}} \quad (5.36)$$

We are now faced with the task of taking the derivative of one function (for simplicity, we'll denote it $g(\tau) = \int_0^t e^{\mathbf{A}(t-\tau)} \mathbf{B} \mathbf{u}(\tau) d\tau$) with respect to another function ($f(\tau) = \mathbf{u}(\tau)$). Now

$$\frac{\partial g(\tau)}{\partial f(\tau)} = \frac{\partial g(\tau)}{\partial \tau} \cdot \frac{\partial \tau}{\partial f(\tau)} = \frac{\frac{\partial g(\tau)}{\partial \tau}}{\frac{\partial f(\tau)}{\partial \tau}} = \frac{g'(\tau)}{f'(\tau)} \quad (5.37)$$

But our overall goal, that of inferring something about input \mathbf{u} by observing a projection of the state \mathbf{x} onto a background state, does not presuppose any knowledge of \mathbf{u} (or its time derivative). Does this preclude any possibility of including a time-varying \mathbf{u} in our formulation? Not necessarily, for we may express $\mathbf{u}(\tau)$ as a projection onto an alternative set of basis functions—the Fourier basis, for example. By truncating the (infinite) Fourier basis functions (in complex exponential form) to include the first m harmonics, we can, by changing m , approximate $\mathbf{u}(\tau)$ to any degree of precision: $\mathbf{u}(\tau) \approx \mathbf{U}h(\tau)$ where the coefficient matrix $\mathbf{U} \in \mathbb{R}^{n \times m}$ and

$$h(\tau) = \begin{bmatrix} e^{2\pi i \omega_0 \tau} \\ e^{2\pi i 2\omega_0 \tau} \\ e^{2\pi i 3\omega_0 \tau} \\ \cdot \\ \cdot \\ e^{2\pi i m\omega_0 \tau} \end{bmatrix}$$

for fundamental angular frequency ω_0 .

Now whatever information we are able to infer about the constant Fourier coefficient matrix \mathbf{U} gives us precise information about $\mathbf{u}(\tau)$ (with precision determined by m), so we can simplify (5.36), taking derivatives with respect to \mathbf{U} :

$$\mathrm{tr}(\mathcal{J}_{\mathbf{u}}) \approx \frac{\mathbf{x}_{\mathrm{ref}}^T \frac{\partial \left(\int_0^t e^{\mathbf{A}(t-\tau)} \mathbf{B} \mathbf{U} h(\tau) d\tau \right)}{\partial \mathbf{U}} \frac{\partial \left(\int_0^t e^{\mathbf{A}(t-\tau)} \mathbf{B} \mathbf{U} h(\tau) d\tau \right)^T}{\partial \mathbf{U}} \mathbf{x}_{\mathrm{ref}}}{\Sigma_{\langle \mathbf{x}, \mathbf{x}_{\mathrm{ref}} \rangle}} \quad (5.38)$$

This is now a linear differentiation in \mathbf{U} , and thus we are able to preserve our original motivation, that of inferring information about $\mathbf{u}(\tau)$ by observing only the state. (5.38) is by no means easy to compute in general, however: differentiation by the matrix \mathbf{U} yields a

Having explored, thus far, many facets of information processing, especially as it related to control-theoretic concepts, we now seek to overcome a fundamental limitation in all our work to this point—namely, we have analyzed systems with linear dynamics throughout. In the next chapter, we extend our analysis to include networks having sigmoidal nonlinearities, and thus make our work much more applicable to real-world dynamical systems.

Chapter 6

Quasilinear approximation of Fisher information in networks with sigmoidal nonlinearities

6.1 Introduction

In recent work [20], we addressed the problem of quantifying the extent to which linear dynamical systems encode information about their afferent inputs. To this end, we employed Fisher information theory [53], optimizing network parameters with the objective of maximizing information. In this chapter, we extend this information-based analysis to the nonlinear regime, enabling the evaluation of approximate Fisher information for networks with sigmoidal nonlinearities, which, as we will see, are particularly relevant in the neural setting.

Our analysis proceeds in two phases. First, we build upon previous work [21] to derive a method for approximating the dynamics of network with sigmoidal, saturating nonlinearities with a simple component-wise scaling and bias applied to a linear dynamical system. We refer to the resulting dynamical system as a quasilinear approximation of the nonlinear one. Then, having characterized the quality of the quasilinear approximation and quantified its accuracy (which tends to be remarkably faithful to the nonlinear system), we will proceed to derive a Fisher information-based measure of network stimulus representation for the quasilinear system, comparing results to those obtained for the corresponding unconstrained linear system (i.e. the original system without the sigmoidal nonlinearity).

We conduct our analysis under two distinct formulations, the first of which has Gaussian white noise added outside the sigmoidal function (and hence the quasilinear transformation in the approximation), and thus has general nonlinear form

$$\dot{\mathbf{x}} = f(\mathbf{x}, \mathbf{u}) + \zeta \tag{6.1}$$

for state \mathbf{x} , sigmoidal function $f(\cdot)$, input \mathbf{u} , and noise process ζ .

The second of which has noise added inside the sigmoid, and has general form

$$\dot{\mathbf{x}} = f(\mathbf{x}, \mathbf{u}, \zeta) \tag{6.2}$$

We show that this change in formulation makes a critical difference in the Fisher information analysis. In particular, we show (analytically for a special case) that the quasilinear Fisher information (QFI) is categorically worse than the linear Fisher information (LFI) under formulation 1–(6.1)–but that the reverse is true under formulation 2–(6.2). We derive an analytic expression precisely relating (under mild assumptions) the QFI under the two formulations. We assess the implications of this reversal for the accurate modeling of stochastic

dynamical systems, and taking into account the relative merits of each formulation, derive conclusions about the advantages and disadvantages—from an informational perspective—inherent in systems with sigmoidal nonlinearities, which are prevalent in natural (especially neural) networks [72–76].

We also conduct a numerical study showing how the relation between QFI and LFI monotonically depends on the saturation limits of the sigmoid (or, more precisely, the saturation limits of a piecewise linear sigmoidal approximation). This is performed for both formulations, again showing an inverse relationship between the two. That is, under formulation (6.2), $\text{QFI} > \text{FI}$, but QFI decreases asymptotically toward FI as the ‘no saturation’ range expands (in the limiting case there would be no saturation at all so the systems would be identical), while under formulation (6.1) $\text{QFI} < \text{FI}$ but QFI increases asymptotically toward FI. We show that for system (6.2), where QFI radically outperforms FI when saturation limits are narrow, there is a fundamental trade-off between information and speed of dynamic evolution; that is, we can have better information, in the long run, with a narrow sigmoid, but the scaled-down dynamics lead to relatively slow response.

6.2 Problem Formulation 1

We begin by defining the dynamics of the nonlinear system of differential equations for state variables \mathbf{x}_i $i \in \{1, \dots, n\}$, represented collectively as the state vector $\mathbf{x} \in \mathbb{R}^n$:

$$\dot{\mathbf{x}}(t) = f(\mathbf{Ax} + \mathbf{Bu}) + \mathbf{Gz}(t) \tag{6.3}$$

where $\mathbf{A} \in \mathbb{R}^{n \times n}$ is a constant weighted adjacency matrix, $\mathbf{u} \in \mathbb{R}^m$ is a constant input vector, $\mathbf{B} \in \mathbb{R}^{n \times m}$ is a constant input-weighting matrix, $\mathbf{z}(t)$ is a zero-mean, unit-variance, stationary Gaussian noise process, \mathbf{G} is a diagonal matrix of noise gains, and where a sigmoidal

nonlinearity is approximated by the piecewise linear function

$$f(x) = \begin{cases} \beta, & x \geq \beta \\ x, & \alpha < x < \beta \\ \alpha, & x \leq \alpha \end{cases} . \quad (6.4)$$

We seek to find constant diagonal matrix $\mathbf{N} \in \mathbb{R}^{n \times n}$ and constant vector $\boldsymbol{\mu} \in \mathbb{R}^n$ such that we can closely approximate the dynamics of (6.3) by the linear stochastic system

$$\dot{\hat{\mathbf{x}}}(t) = \mathbf{N}(\mathbf{A}\hat{\mathbf{x}} + \mathbf{B}\mathbf{u}) - \mathbf{N}\boldsymbol{\mu} + \mathbf{G}\hat{\mathbf{z}}(t) \quad (6.5)$$

We think of \mathbf{N} (or, more precisely, the diagonal of \mathbf{N}) as a quasilinear gain and $\boldsymbol{\mu}$ as a quasilinear offset, or bias. To make these approximations meaningful, we must pick a stationary regime of the stochastic process upon which $f(\cdot)$ acts—namely, when the system is in steady-state. Thus our task is to choose parameters \mathbf{N} and $\boldsymbol{\mu}$ such that the distribution of the process—about the steady-state—in (6.5) is as similar as possible to that of (6.3). We subtract $\boldsymbol{\mu}$ because it represents a scaled mean difference between the dynamics which would be prescribed by a simple scaling by \mathbf{N} and the actual dynamics prescribed by $f(\cdot)$. That is, letting

$$\begin{aligned} \mathbf{N}\boldsymbol{\mu} &= \mathbb{E}[\mathbf{N}(\mathbf{A}\hat{\mathbf{x}} + \mathbf{B}\mathbf{u}) - f(\mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u})] \\ &= \mathbf{N}\mathbb{E}[\mathbf{A}\hat{\mathbf{x}} + \mathbf{B}\mathbf{u}] - \mathbb{E}[f(\mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u})] \end{aligned} \quad (6.6)$$

we have

$$\mathbf{N}\mathbb{E}[\mathbf{A}\hat{\mathbf{x}} + \mathbf{B}\mathbf{u}] - \mathbf{N}\boldsymbol{\mu} = \mathbb{E}[f(\mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u})] \quad (6.7)$$

To find optimal values of \mathbf{N} and $\boldsymbol{\mu}$, we seek to minimize the error function

$$\boldsymbol{\varepsilon}(\mathbf{N}, \boldsymbol{\mu}) = \mathbb{E}[(\mathbf{x}(t) - \hat{\mathbf{x}}(t))^2] \quad (6.8)$$

Since the $\mathbb{E}[\mathbf{z}(t) - \hat{\mathbf{z}}(t)] = \mathbf{0} \forall t$, we can rewrite (6.8) as

$$\boldsymbol{\varepsilon}(\mathbf{N}, \boldsymbol{\mu}) = \mathbb{E}[(f(\mathbf{Ax} + \mathbf{Bu}) - \mathbf{N}(\mathbf{Ax} + \mathbf{Bu} - \boldsymbol{\mu}))^2] \quad (6.9)$$

To find the matrix \mathbf{N} which will minimize this function, we recall that \mathbf{N} is diagonal so that for $\mathbf{n} = \mathbf{N} \cdot \mathbf{1}$, and for generic vector \mathbf{v} ,

$$\mathbf{n} \odot \mathbf{v} = \mathbf{N}\mathbf{v} = \mathbf{v} \odot \mathbf{n} \quad (6.10)$$

where \odot denotes element-wise multiplication. Thus we can reframe (6.9) as

$$\boldsymbol{\varepsilon}(\mathbf{n}, \boldsymbol{\mu}) = \mathbb{E}[(f(\mathbf{Ax} + \mathbf{Bu}) - \mathbf{n} \odot (\mathbf{Ax} + \mathbf{Bu} - \boldsymbol{\mu}))^2] \quad (6.11)$$

Now we can differentiate (6.11) with respect to \mathbf{n} , setting the derivative equal to $\mathbf{0}$ for minimization:

$$\frac{d\boldsymbol{\varepsilon}}{d\mathbf{n}} = \mathbb{E}[2(f(\mathbf{Ax} + \mathbf{Bu}) - (\mathbf{n} \odot (\mathbf{Ax} + \mathbf{Bu} - \boldsymbol{\mu}))) \odot (\mathbf{Ax} + \mathbf{Bu} - \boldsymbol{\mu})] = \mathbf{0} \quad (6.12)$$

For simplicity of notation we will let \mathbf{y} and $\hat{\mathbf{y}}$ denote $\mathbf{Ax} + \mathbf{Bu}$ and $\mathbf{Ax} + \mathbf{Bu} - \boldsymbol{\mu}$, respectively.

Thus (6.12) becomes

$$\frac{d\boldsymbol{\varepsilon}}{d\mathbf{n}} = \mathbb{E}[2(f(\mathbf{y}) - (\mathbf{n} \odot \hat{\mathbf{y}})) \odot \hat{\mathbf{y}}] = \mathbf{0} \quad (6.13)$$

We now make an inherently recursive assumption—namely, that the approximation $\hat{\mathbf{x}}(t)$ of the state $\mathbf{x}(t)$, for any given t , is close enough that we can rewrite (6.13) all in terms of $\mathbf{x}(t)$

(and therefore of \mathbf{y}), thus allowing us to accurately estimate $\mathbf{x}(t)$ by finding optimal \mathbf{n} and $\boldsymbol{\mu}$. This is a fairly subtle point, and it will impact all of the technical development henceforth. Operating under this assumption, we have

$$\frac{d\boldsymbol{\varepsilon}}{d\mathbf{n}} = \mathbf{E}[2(f(\mathbf{y}) - (\mathbf{n} \odot \mathbf{y})) \odot \mathbf{y}] = \mathbf{0} \quad (6.14)$$

Thus, by the linearity of expectation,

$$\mathbf{E}[f(\mathbf{y}) \odot \mathbf{y}] = \mathbf{n} \odot \mathbf{E}[\mathbf{y} \odot \mathbf{y}] \quad (6.15)$$

We make the assumption that $\mathbf{y}(t)$ is a stationary Gaussian process (at steady-state), since $\hat{\mathbf{z}}(t)$ is such, with the caveat that since \mathbf{y} includes recurrent dynamics on \mathbf{x} , this may not always be strictly true. Proceeding with this assumption, and taking into account that $f(\cdot)$ is piecewise differentiable, we have

$$\mathbf{E}[f(\mathbf{y}) \odot \mathbf{y}] = \mathbf{E}[\mathbf{y} \odot \mathbf{y}] \odot \mathbf{E}\left[\left.\frac{df}{d\mathbf{y}}\right|_{\mathbf{y}=\mathbf{y}(t)}\right] = \mathbf{E}\left[\left.\frac{df}{d\mathbf{y}}\right|_{\mathbf{y}=\mathbf{y}(t)}\right] \odot \mathbf{E}[\mathbf{y} \odot \mathbf{y}] \quad (6.16)$$

where $\left(\left.\frac{df}{d\mathbf{y}}\right|_{\mathbf{y}=\mathbf{y}(t)}\right)$ —henceforth denoted $f'(\mathbf{y}(t))$ —is a component-wise differentiation. Thus by (6.15) and (6.39), the optimal \mathbf{n} becomes

$$\mathbf{n} = \mathbf{E}[f'(\mathbf{y}(t))] \quad (6.17)$$

and thus

$$\mathbf{N} = \text{diag}(\mathbf{E}[f'(\mathbf{y}(t))]) \quad (6.18)$$

Now, to find the optimal $\boldsymbol{\mu}$, we take the derivative of (6.8) with respect to $\boldsymbol{\mu}$. Using the same notation as above (and again letting $\mathbf{n} = \mathbf{N} \cdot \mathbf{1}$), we have

$$\begin{aligned} \frac{d\boldsymbol{\varepsilon}}{d\boldsymbol{\mu}} &= \mathbb{E} \left[2(f(\mathbf{y}) - (\mathbf{n} \odot \hat{\mathbf{y}})) \odot \frac{d(\mathbf{n} \odot \hat{\mathbf{y}})}{d\boldsymbol{\mu}} \right] \\ &= \mathbb{E} [2(f(\mathbf{y}) - (\mathbf{n} \odot \hat{\mathbf{y}})) \odot \mathbf{n}] \\ &= 2 \mathbb{E} [f(\mathbf{y}) - (\mathbf{n} \odot \hat{\mathbf{y}})] \odot \mathbf{n} \end{aligned}$$

Setting this equal to $\mathbf{0}$, we can cancel the terms outside the expectation to obtain

$$\mathbb{E} [f(\mathbf{y}) - (\mathbf{n} \odot \hat{\mathbf{y}})] = \mathbf{0}$$

so that (again by the linearity of expectation)

$$\mathbb{E} [f(\mathbf{y})] = \mathbb{E} [\mathbf{n} \odot \hat{\mathbf{y}}] = \mathbb{E} [\mathbf{N}\hat{\mathbf{y}}] \quad (6.19)$$

Now, $\mathbf{N}\hat{\mathbf{y}} = \mathbf{N}(\mathbf{A}\hat{\mathbf{x}} + \mathbf{B}\mathbf{u}) - \mathbf{N}\boldsymbol{\mu}$, as in (6.5). We recall that we seek the approximation $\boldsymbol{\mu}$ at steady-state; thus, we set $\hat{\mathbf{x}}$ in (6.5) to zero to obtain (recalling that $\mathbf{z}(t)$ is a zero-mean process)

$$\mathbf{0} = \mathbf{N}(\mathbf{A}\hat{\mathbf{x}} + \mathbf{B}\mathbf{u}) - \mathbf{N}\boldsymbol{\mu} = \mathbb{E} [\mathbf{N}\hat{\mathbf{y}}]_{ss} \quad (6.20)$$

Since $\mathbf{N}\hat{\mathbf{y}}$ is by definition an approximation of $f(\mathbf{y})$, we evaluate the expectation (6.20) with respect to $f(\mathbf{y})$. We make an assumption of the Gaussianity of $f(\mathbf{y})$. This is mostly justified when there is sufficiently narrow bandwidth (i.e. when the system (6.3) acts as a low-pass filter). Further, we assume that $f'(\mathbf{y})$, which is a Gaussian fed through a linear transformation (namely a derivative, which acts as a linear filter), is Gaussian. Proceeding under these

(relatively mild) assumptions, we make use the definition of the expectation of a Gaussian random variable together with (6.18) and (6.20) to form n sets of coupled equations prescribed by:

$$\mathbf{N}_{i,i} = \mathbf{n}_i = \int_{-\infty}^{\infty} \frac{d}{dy_i} f(\mathbf{y}_i) \exp\left(-\frac{(\mathbf{y}_i - \boldsymbol{\mu}_i)^2}{2\boldsymbol{\sigma}_{\mathbf{y}_i}^2}\right) d\mathbf{y}_i \quad (6.21)$$

$$\mathbf{0} = \mathbf{E}[\mathbf{N}\hat{\mathbf{y}}]_i = \mathbf{E}[f(\mathbf{y}_i)] = \int_{-\infty}^{\infty} \frac{f(\mathbf{y}_i)}{\sqrt{2\pi\boldsymbol{\sigma}_{\mathbf{y}_i}}} \exp\left(-\frac{(\mathbf{y}_i - \boldsymbol{\mu}_i)^2}{2\boldsymbol{\sigma}_{\mathbf{y}_i}^2}\right) d\mathbf{y}_i \quad (6.22)$$

for $i \in \{1, \dots, n\}$.

Since these equations depend on the standard deviations of the \mathbf{y}_i ($\boldsymbol{\sigma}_{\mathbf{y}_i}$), which are functions of \mathbf{y} 's covariance, we next examine that covariance. First, we note that the covariance of a nonlinear system (more precisely in our case, one with a piecewise-linear nonlinearity) is not computable explicitly, so we instead seek the covariance of the estimated, linear dynamics on $\hat{\mathbf{y}}$. Since the covariance of $\hat{\mathbf{y}}$ is really just the covariance of $\hat{\mathbf{x}}$ with respect to the stochastic noise input, and since we are finding the covariance in steady-state, we can simply find the infinite-time controllability gramian for the linear system $\dot{\hat{\mathbf{x}}} = \mathbf{N}\mathbf{A}\hat{\mathbf{x}} + \mathbf{G}\mathbf{z}$; that is, we solve the following Lyapunov equation for $\boldsymbol{\Sigma}_N$ (we use the subscript \mathbf{N} to keep in mind the explicit dependence of the covariance on quasilinear gain \mathbf{N}).

$$(\mathbf{N}\mathbf{A})\boldsymbol{\Sigma}_N + \boldsymbol{\Sigma}_N(\mathbf{N}\mathbf{A})^T + \mathbf{G}\mathbf{G}^T = \mathbf{0} \quad (6.23)$$

Now, since we seek the component-wise standard deviation of $\hat{\mathbf{x}}$ (hence $\hat{\mathbf{y}}$), we have to distribute the variances and covariances encoded in $\boldsymbol{\Sigma}_N$, weighted by \mathbf{A} , to obtain the variance of $\hat{\mathbf{y}}_i$ $i \in \{1, \dots, n\}$:

$$\boldsymbol{\sigma}_{\hat{\mathbf{y}}_i}^2 = \text{var}(\hat{\mathbf{y}}_i) = \mathbf{M}_{i,i} \quad (6.24)$$

where

$$\mathbf{M} = \mathbf{A} \Sigma_N \mathbf{A}^T \quad (6.25)$$

where we have used the fact that Σ_N is symmetric. Since, then, the variance (hence the standard deviation) depends explicitly on \mathbf{N} , we can solve the system of paired equations (6.21),(6.22) for the n diagonal entries of quasilinear gain matrix \mathbf{N} and the n entries of bias vector $\boldsymbol{\mu}$. Note that each $\mathbf{N}_{i,i}$ and μ_i , $i \in \{1, \dots, n\}$ is found separately as a pair. We spare the reader the details of how solutions are obtained; it is sufficient to notice that (6.21) and (6.22), by construction the expectations of (quasi-)Gaussian random variables $\frac{d}{d\mathbf{y}_i} f(\mathbf{y}_i)$ and $f(\mathbf{y}_i)$, respectively, are solved by means of partitioning the domain $(-\infty, \infty)$ into parts determined by the limits of the saturation function $f(\cdot)$ (α and β —see (6.4)) and using the error function

$$\text{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt \quad (6.26)$$

Expressed in terms of the error function, (6.21) and (6.22) become

$$\mathbf{N}_{i,i} = \frac{1}{2} \left[\text{erf} \left(\frac{\beta - \mu_i}{\sqrt{2} \sigma_{\hat{\mathbf{y}}_i}} \right) - \text{erf} \left(\frac{\alpha - \mu_i}{\sqrt{2} \sigma_{\hat{\mathbf{y}}_i}} \right) \right] \quad (6.27)$$

and

$$\begin{aligned} 0 = & \frac{\alpha + \beta}{2} - \frac{\beta - \mu_i}{2} \text{erf} \left(\frac{\beta - \mu_i}{\sqrt{2} \sigma_{\hat{\mathbf{y}}_i}} \right) + \frac{\alpha - \mu_i}{2} \text{erf} \left(\frac{\alpha - \mu_i}{\sqrt{2} \sigma_{\hat{\mathbf{y}}_i}} \right) \\ & - \frac{\sigma_{\hat{\mathbf{y}}_i}}{\sqrt{2\pi}} \left[\exp \left(- \left(\frac{\beta - \mu_i}{\sqrt{2} \sigma_{\hat{\mathbf{y}}_i}} \right)^2 \right) - \exp \left(- \left(\frac{\alpha - \mu_i}{\sqrt{2} \sigma_{\hat{\mathbf{y}}_i}} \right)^2 \right) \right] \end{aligned} \quad (6.28)$$

These equations, though difficult or impossible to solve analytically, are readily solved using numerical optimization techniques. We used Mathematica's NMinimize command, using the norm of the difference between lhs and rhs of (6.27) and (6.28), concatenated, as the

objective function and minimizing over variables \mathbf{n} and $\boldsymbol{\mu}$. Global minima were found the vast majority of the time, evidenced by objective function values on the order of 10^{-10} .

6.2.1 Quality of quasilinear approximation

To begin our assessment of the quality of the quasilinear approximation, we simulated trajectories of systems (6.3) and (6.5), perturbed by the same noise process. We used a forward-Euler method with time step δ and total number of timesteps k ; thus the k sampled values of the noise process $\mathbf{z}(t)$ were applied identically to each system. For comparisons sake, we also simulated the system

$$\dot{\mathbf{x}}_J(t) = \mathbf{A}\mathbf{x}_J + \mathbf{B}\mathbf{u} + \mathbf{G}\mathbf{z}(t) \tag{6.29}$$

with all terms exactly as in (6.3) (or (6.5)) but without the nonlinearity $f(\cdot)$ (or the associated approximation terms \mathbf{N} and $\boldsymbol{\mu}$). This can be thought of as a Jacobian linearization around a stable fixed point of (6.3), but since we apply it even in cases where the sigmoid argument is below α or above β (component-wise), we still employ a gain (slope) of 1, so it is not strictly a Jacobian linearization; we will sometimes refer to it as a heuristic Jacobian linearization.

A few illustrated examples show the clear superiority of the quasilinear approximation over the heuristic Jacobian linearization. Figure 6.1 shows trajectories for a 2-dimensional system. Note that the variance of the quasilinear trajectory is much more aligned with that of the sigmoidal trajectory than is that of the Jacobian-linearized version. Figure 6.2 shows the same qualitative result.

To gauge the accuracy of the quasilinear approximation (6.5) with respect to the original system (6.3) in arbitrary dimensions, a statistical approach is obviously needed. We simply integrate the mean squared difference between $\mathbf{x}(t)$ and $\hat{\mathbf{x}}(t)$ over the time domain $[0, T]$ and

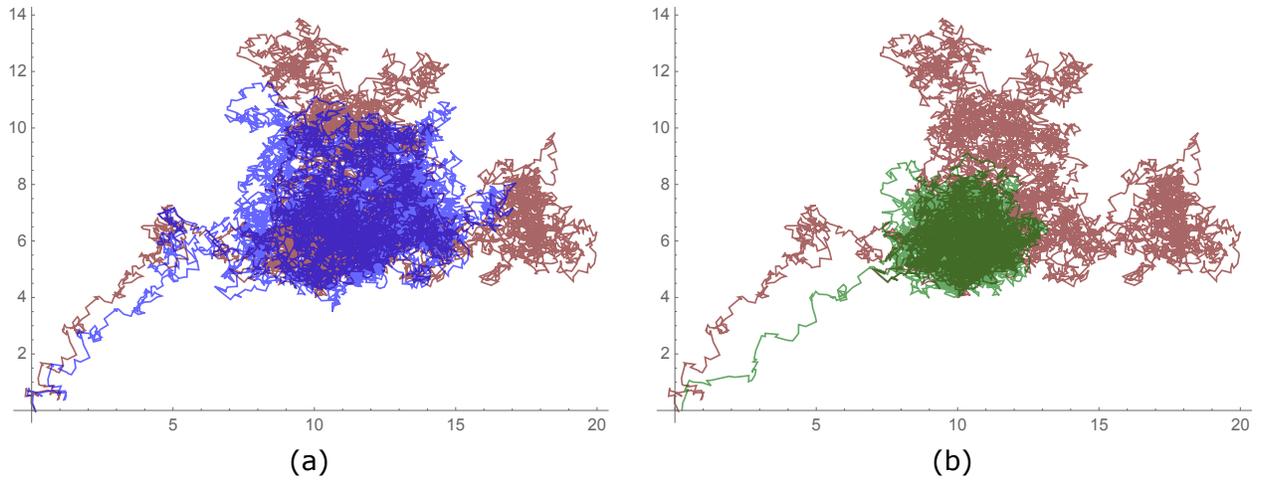


Figure 6.1: An example illustrating the quality of the quasilinear approximation in 2D. (a) shows the state trajectory of the sigmoidal system (red) and the quasilinear system (blue). (b) shows the same sigmoidal trajectory (red) and the (heuristic) Jacobian-linearized system trajectory (green).

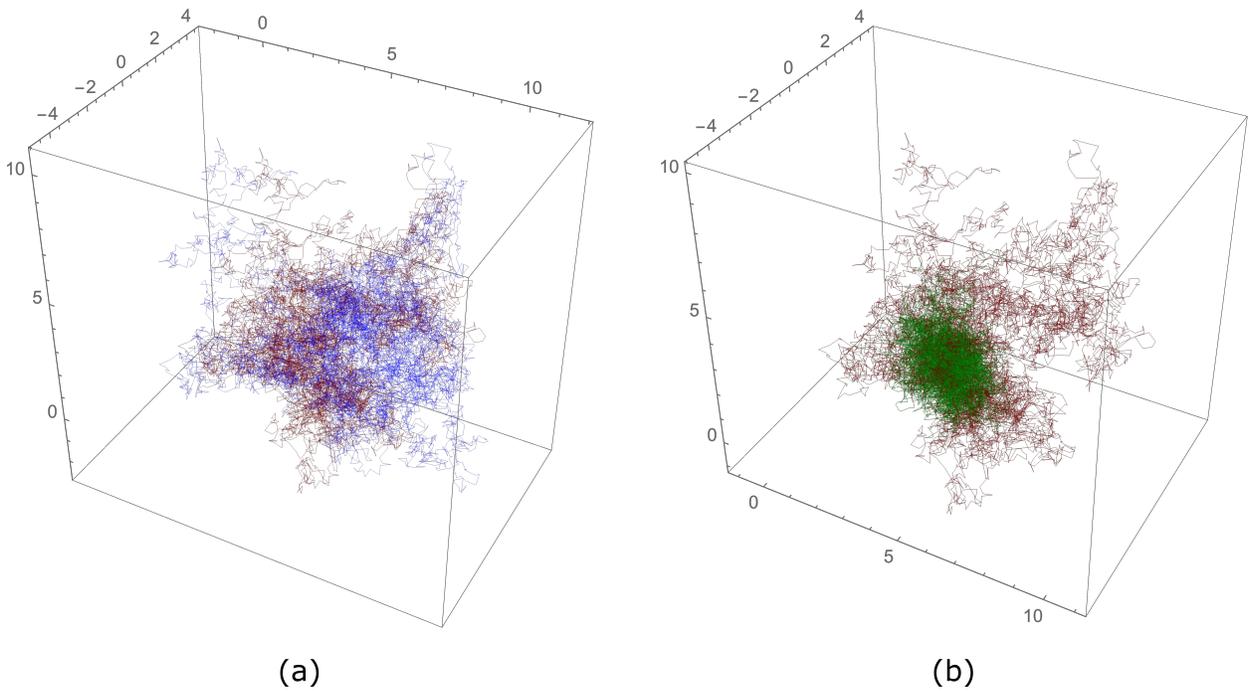


Figure 6.2: An example illustrating the quality of the quasilinear approximation in 3D. (a) shows the state trajectory of the sigmoidal system (red) and the quasilinear system (blue). (b) shows the same sigmoidal trajectory (red) and the (heuristic) Jacobian-linearized system trajectory (green).

scale by the total squared displacement of $\mathbf{x}(t)$:

$$\boldsymbol{\varepsilon}_{\text{total}}(\mathbf{N}, \boldsymbol{\mu}) = \frac{\int_0^T \|\mathbf{x}(t) - \hat{\mathbf{x}}(t)\| dt}{\int_0^T \|\mathbf{x}(t)\| dt} \quad (6.30)$$

Although it is theoretically possible for this quantity to exceed 1, in practice it never happens, so $\boldsymbol{\varepsilon}_{\text{total}} \in \{0, 1\}$ furnishes a good measure of quasilinear accuracy. In practice, since we are using a forward-Euler method (with k timesteps) of integration for our simulations, (6.30) is discretized to become

$$\boldsymbol{\varepsilon}_{\text{total}}^*(\mathbf{N}, \boldsymbol{\mu}) = \frac{\sum_{i=1}^k \|\mathbf{x}_i - \hat{\mathbf{x}}_i\|}{\sum_{i=1}^k \|\mathbf{x}_i\|} \quad (6.31)$$

It is also critically important that the covariance structure of the output of (6.5) be similar to that of (6.3). The example figures (Figures 6.1 and 6.2) have suggested this, but to quantify it precisely, we evaluated that following simple measure

$$\boldsymbol{\varepsilon}_{\text{var}} = \frac{\|\text{var}(\mathbf{x}(t)) - \text{var}(\hat{\mathbf{x}}(t))\|}{\|\text{var}(\mathbf{x}(t))\|} \quad (6.32)$$

and its discrete-time approximation

$$\boldsymbol{\varepsilon}_{\text{var}}^* = \frac{\|\text{var}(\mathbf{x}) - \text{var}(\hat{\mathbf{x}})\|}{\|\text{var}(\mathbf{x})\|} \quad (6.33)$$

where the $\text{var}(\mathbf{x})$ is the component-wise variance of \mathbf{x} , evaluated over the whole timecourse $i \in \{1, \dots, k\}$.

A statistical characterization of the accuracy of the quasilinear approximation and a comparison of this accuracy with that of the (heuristic) Jacobian linearization, is given in Figure 6.3, for which we ran 50 trials with networks of size $n = 100$, computed normed differences of

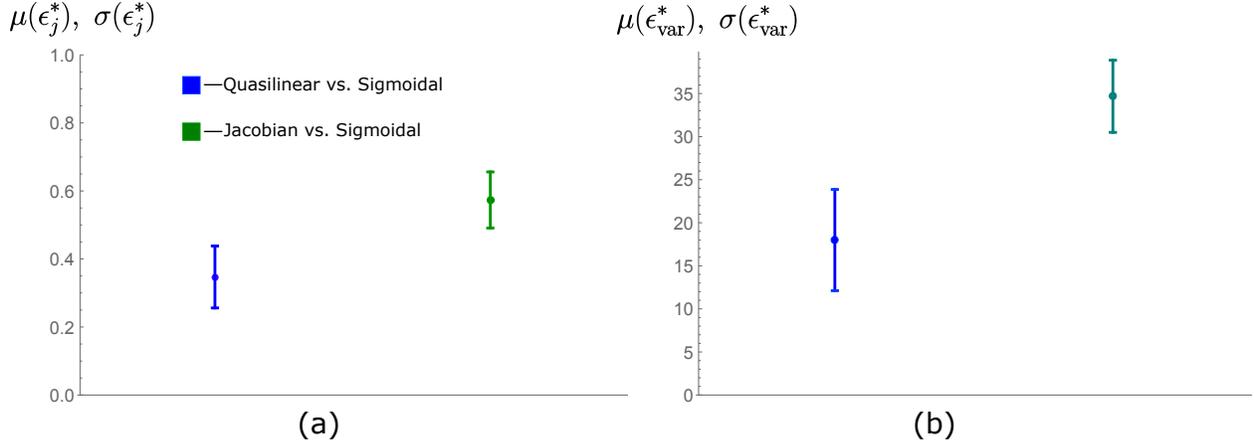


Figure 6.3: We ran simulations of 50 networks each of quasilinear, Jacobian-linearized, and sigmoidal systems ((6.5),(6.29), and (6.3), respectively) with the same dynamical matrices (\mathbf{A}, \mathbf{B}), same noise gain matrix ($\mathbf{G} = 2.3 * \mathbf{I}$), the same (random for each trial) constant input \mathbf{u} (with $0 \leq \mathbf{u}_i \leq 15 \forall i \in \{1, \dots, n\}$) and same noise input ($\mathbf{z}(t)$) for each trial. Each trial was run using a forward-Euler integration method with $k = 10,000$ time steps, and saturation limits $\alpha = -2$ and $\beta = 3$. (a) shows mean and standard deviation of the value of (6.31) over the 50 trials. Note that for the Jacobian-linearized vs sigmoidal error (green), we just use \mathbf{x}_J (see (6.29)) in place of $\hat{\mathbf{x}}$ in (6.31). (b) shows mean and standard deviation of the variance discrepancy, as quantified by (6.32), over the 50 trials.

state trajectories and variances, as described above and quantified by (6.31) and (6.33), and plotted mean and standard deviation for the 50 trials.

6.2.2 Quasilinear Fisher information under formulation 1

We sought to evaluate the implications of quasilinearization for information processing in dynamical networks. To do so, we employ a simple notion of information processing–stimulus encoding. Specifically, we seek to quantify the extent to which a dynamical system encodes information about its afferent inputs by means of a state ‘readout’. The calculation of the Fisher information is straightforward for a linear system. In particular, since we consider the noise input to be Gaussian with zero mean, by basic linear systems theory together with dynamics prescribed by (6.5), the state $\mathbf{x}(t)$ (we will drop the notation $\hat{\mathbf{x}}$ henceforth for

greater notational simplicity) is Gaussian according to

$$\mathbf{x}(t) \sim \mathcal{N} \left(\int_0^t e^{\mathbf{N}\mathbf{A}(t-\tau)} d\tau \mathbf{N}\mathbf{B}\mathbf{u} + \int_0^t e^{\mathbf{N}\mathbf{A}(t-\tau)} d\tau \mathbf{N}\boldsymbol{\mu}, \boldsymbol{\Sigma}_N \right), \quad (6.34)$$

where $\boldsymbol{\Sigma}_N \in \mathbb{R}^{n \times n}$ is the covariance matrix of $\mathbf{x}(t)$ given as the solution to (6.23). Hence, we make the assumption that the noise covariance is always at steady-state. Conceptually, this assumption means that the dynamics of the network are persistently excited by ongoing noise, while receiving stimuli in a temporally punctate manner.

Since $\mathbf{x}(t)$ is Gaussian, the Fisher information matrix quantifying the amount of information ‘preserved’ or ‘encoded’ about constant stimulus \mathbf{u} is given by

$$\mathcal{J}_{\mathbf{u}}^{\mathbf{Q}}(t) = \frac{\partial \mathbf{E}[\mathbf{x}(t)]^T}{\partial \mathbf{u}} \boldsymbol{\Sigma}_N^{-1} \frac{\partial \mathbf{E}[\mathbf{x}(t)]}{\partial \mathbf{u}} \quad (6.35)$$

The expectation of $\mathbf{x}(t)$ is straightforward and is given explicitly in (6.34). Thus the derivation proceeds very easily as follows, where we define

$$\boldsymbol{\Gamma}_N(t) = \int_0^t e^{\mathbf{N}\mathbf{A}(t-\tau)} d\tau \quad (6.36)$$

for sake of clarity. Then

$$\begin{aligned} \mathcal{J}_{\mathbf{u}}^{\mathbf{Q}}(t) &= \frac{\partial (\boldsymbol{\Gamma}_N(t) \mathbf{N}\mathbf{B}\mathbf{u} + \boldsymbol{\Gamma}_N(t) \mathbf{N}\boldsymbol{\mu})^T}{\partial \mathbf{u}} \boldsymbol{\Sigma}_N^{-1} \frac{\partial (\boldsymbol{\Gamma}_N(t) \mathbf{N}\mathbf{B}\mathbf{u} + \boldsymbol{\Gamma}_N(t) \mathbf{N}\boldsymbol{\mu})}{\partial \mathbf{u}} \\ &= (\boldsymbol{\Gamma}_N(t) \mathbf{N}\mathbf{B})^T \boldsymbol{\Sigma}_N^{-1} (\boldsymbol{\Gamma}_N(t) \mathbf{N}\mathbf{B}) \\ &= \mathbf{B}^T \mathbf{N}^T \boldsymbol{\Gamma}_N(t)^T \boldsymbol{\Sigma}_N^{-1} \boldsymbol{\Gamma}_N(t) \mathbf{N}\mathbf{B} \end{aligned} \quad (6.37)$$

where we use the superscript \mathbf{Q} to indicate that it is the Fisher information for the quasilinear system (6.5).

We see that the information value $\mathcal{I}_{\mathbf{u}}^{\mathbf{Q}}$ depends critically on quasilinear gain \mathbf{N} , and indirectly on quasilinear bias $\boldsymbol{\mu}$, since it figures into the calculations for \mathbf{N} —see (6.21). These quasilinear terms, in turn, depend on choices of saturation levels $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$, as is clear from (6.4),(6.21), and (6.22).

Since we have seen that the quasilinear approximation is, in fact, quite good as a linear estimate of the nonlinear system, we proceed to compare the Fisher information (given in (6.37)) of the quasilinear system with the same Fisher information in the corresponding unsaturated linear system $\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u} + \mathbf{G}\mathbf{z}$. By this comparison we hope to ascertain whether there is an advantage or disadvantage, from an informational perspective, from having a sigmoidal nonlinearity modulating the dynamics.

Since all comparisons will fundamentally depend on quasilinear gain \mathbf{N} , we begin by noting that the expression for \mathbf{N} given in (6.21) will always satisfy

$$\mathbf{0} \leq \mathbf{N}_{i,i} \leq 1 \quad \forall i \in \{1, \dots, n\} \tag{6.38}$$

because the formulation of $f(\cdot)$ given in (6.4) guarantees that $0 \leq f'(x) \leq 1 \quad \forall x \in \mathbb{R}$ (notice that the slope of f is 0 for all $x \leq \boldsymbol{\alpha}$ and $x \geq \boldsymbol{\beta}$ and 1 for all $\boldsymbol{\alpha} < x < \boldsymbol{\beta}$) and thus the expected value $\mathbf{N}_{i,i} = \mathbf{E}[f'(\mathbf{y}_i)]$ must lie in this range. This fact will be fundamental to all subsequent analyses.

We need a precise formulation for the steady-state quasilinear Fisher information. Toward this end, we examine (6.37) and seek a steady-state characterization of $\Gamma_N(t)$ (see (6.36)). We recall an important identity—that for generic matrix \mathbf{A} and terminal time t , by definition of

the matrix exponential we have

$$\begin{aligned}
\int_0^t e^{A\tau} d\tau &= \int_0^t \left(I + A\tau + \frac{(A\tau)^2}{2!} + \frac{(A\tau)^3}{3!} + \dots + \frac{(A\tau)^n}{n!} + \dots \right) d\tau \\
&= I \int_0^t d\tau + A \int_0^t \tau d\tau + \frac{A^2}{2!} \int_0^t \tau^2 d\tau + \frac{A^3}{3!} \int_0^t \tau^3 d\tau + \dots \\
&= It + \frac{At^2}{2} + \frac{A^2 t^3}{2!3} + \frac{A^3 t^4}{3!4} + \dots \\
&= A^{-1} \left(At + \frac{(At)^2}{2!} + \frac{(At)^3}{3!} + \dots \right) \\
&= A^{-1} (e^{At} - I)
\end{aligned} \tag{6.39}$$

Therefore we have

$$\begin{aligned}
\Gamma_N(t) &= \int_0^t e^{\mathbf{NA}(t-\tau)} d\tau \\
&= e^{\mathbf{NA}t} \int_0^t e^{-\mathbf{NA}\tau} d\tau \\
&= e^{\mathbf{NA}t} \left[-(\mathbf{NA})^{-1} (e^{-\mathbf{NA}t} - \mathbf{I}) \right] \\
&= e^{\mathbf{NA}t} \left[-(\mathbf{NA})^{-1} e^{-\mathbf{NA}t} + (\mathbf{NA})^{-1} \right]
\end{aligned} \tag{6.40}$$

$$= -e^{\mathbf{NA}t} e^{-\mathbf{NA}t} (\mathbf{NA})^{-1} + e^{\mathbf{NA}t} (\mathbf{NA})^{-1} \tag{6.41}$$

$$= -(\mathbf{NA})^{-1} + e^{\mathbf{NA}t} (\mathbf{NA})^{-1} \tag{6.42}$$

$$\tag{6.43}$$

where the step from (6.40) to (6.41) is possible because of the commutivity of a matrix A and its matrix exponential e^{At} (this is obvious by looking at the Taylor expansion of the matrix exponential). Now, since \mathbf{A} is assumed to be stable, as $t \rightarrow \infty$, $e^{\mathbf{NA}t} \rightarrow \mathbf{0}$, so evaluating (6.42) at steady-state we have

$$\Gamma_N(\infty) = -(\mathbf{NA})^{-1} = -\mathbf{A}^{-1} \mathbf{N}^{-1} \tag{6.44}$$

Therefore, by (6.37) and (6.44), the steady-state Fisher information $\mathcal{J}_{\mathbf{u}_{\text{ss}}}^{\text{Q}}$ is given by

$$\begin{aligned}\mathcal{J}_{\mathbf{u}_{\text{ss}}}^{\text{Q}} &= \mathbf{B}^T \mathbf{N} (\mathbf{N} \mathbf{A})^{-T} \Sigma_N^{-1} (\mathbf{N} \mathbf{A})^{-1} \mathbf{N} \mathbf{B} \\ &= \mathbf{B}^T \mathbf{A}^{-T} \Sigma_N^{-1} \mathbf{A}^{-1} \mathbf{B}\end{aligned}\tag{6.45}$$

Now to compare quasilinear Fisher information (which we have denoted $\mathcal{J}_{\mathbf{u}_{\text{ss}}}^{\text{Q}}$) and linear Fisher information (which we will henceforth denote $\mathcal{J}_{\mathbf{u}_{\text{ss}}}^{\text{L}}$) is, in the general case, difficult analytically. Hence we will provide numerical characterizations later (Sections 6.2.2.2, 6.3.1.2, and 6.4). However, analytical comparisons are possible under certain special conditions, one of which we now consider.

6.2.2.1 Analytic relation between linear and quasilinear Fisher information (under formulation 1) for a special case—uniform diagonal \mathbf{N}

We consider the special case that \mathbf{N} has a uniform diagonal. That is, we define

$$\tilde{\mathbf{N}} \triangleq \lambda \mathbf{I}\tag{6.46}$$

for some $0 < \lambda \leq 1$. We will first relate the respective covariances Σ and $\Sigma_{\tilde{v}}$. We show that, where (6.46) holds, if Σ is a solution to

$$\mathbf{A} \Sigma + \Sigma \mathbf{A}^T + \mathbf{G} \mathbf{G}^T = \mathbf{0}\tag{6.47}$$

which is the steady-state covariance of the system (6.29) with respect to noise, then

$$\Sigma_{\tilde{v}} = \tilde{\mathbf{N}}^{-1} \Sigma\tag{6.48}$$

is a solution to (6.23). Starting with a recapitulation of (6.23), we have

$$\begin{aligned}
(\tilde{\mathbf{N}}\mathbf{A})\Sigma_N + \Sigma_N(\tilde{\mathbf{N}}\mathbf{A})^T + \mathbf{G}\mathbf{G}^T &= (\tilde{\mathbf{N}}\mathbf{A})\tilde{\mathbf{N}}^{-1}\Sigma + \tilde{\mathbf{N}}^{-1}\Sigma(\tilde{\mathbf{N}}\mathbf{A})^T + \mathbf{G}\mathbf{G}^T \\
&= \mathbf{A}\tilde{\mathbf{N}}\tilde{\mathbf{N}}^{-1}\Sigma + \Sigma\tilde{\mathbf{N}}^{-1}\mathbf{A}^T\tilde{\mathbf{N}}^T + \mathbf{G}\mathbf{G}^T \\
&= \mathbf{A}\Sigma + \Sigma\tilde{\mathbf{N}}^{-1}\tilde{\mathbf{N}}^T\mathbf{A}^T + \mathbf{G}\mathbf{G}^T \\
&= \mathbf{A}\Sigma + \Sigma\tilde{\mathbf{N}}^{-1}\tilde{\mathbf{N}}\mathbf{A}^T + \mathbf{G}\mathbf{G}^T \\
&= \mathbf{A}\Sigma + \Sigma\mathbf{A}^T + \mathbf{G}\mathbf{G}^T
\end{aligned} \tag{6.49}$$

which is identical to the LHS of (6.47). Note that we have used the facts that $\tilde{\mathbf{N}}$ (and thus $\tilde{\mathbf{N}}^{-1}$) is symmetric and commutes with any matrix (due to (6.46)).

Therefore, substituting (6.48) into (6.45), we have steady-state quasilinear information

$$\begin{aligned}
\mathcal{J}_{\mathbf{u}_{ss}}^Q &= \mathbf{B}^T \mathbf{A}^{-T} \Sigma^{-1} \tilde{\mathbf{N}} \mathbf{A}^{-1} \mathbf{B} \\
&= \tilde{\mathbf{N}} \mathbf{B}^T \mathbf{A}^{-T} \Sigma^{-1} \mathbf{A}^{-1} \mathbf{B}
\end{aligned} \tag{6.50}$$

Now this is significant, because the (steady-state) Fisher information for the linear system (6.29) is given by

$$\mathcal{J}_{\mathbf{u}_{ss}}^L = \mathbf{B}^T \mathbf{A}^{-T} \Sigma^{-1} \mathbf{A}^{-1} \mathbf{B} \tag{6.51}$$

by an identical derivation as we used to derive (6.37), but without the gain matrix \mathbf{N} and with covariance matrix Σ (see (6.47)) instead of Σ_N . That is, under the condition (6.46),

$$\mathcal{J}_{\mathbf{u}_{ss}}^Q = \tilde{\mathbf{N}} \mathcal{J}_{\mathbf{u}_{ss}}^L = \lambda \mathbf{I} \mathcal{J}_{\mathbf{u}_{ss}}^L \tag{6.52}$$

and thus

$$\text{tr}(\mathcal{J}_{\mathbf{u}_{ss}}^Q) = \text{tr}(\lambda \mathbf{I} \mathcal{J}_{\mathbf{u}_{ss}}^L) = \lambda \text{tr}(\mathcal{J}_{\mathbf{u}_{ss}}^L) \tag{6.53}$$

This means that, since λ (and thus the diagonal elements of $\tilde{\mathbf{N}}$) are always between 0 and 1 (see (6.38)), the quasilinear Fisher information is always poorer (unless $\lambda = 1$, in which case (6.5) reduces to (6.29)) than the corresponding linear Fisher information. That is,

$$\mathrm{tr}(\mathcal{J}_{\mathbf{u}_{\mathrm{ss}}}^{\mathrm{Q}}) \leq \mathrm{tr}(\mathcal{J}_{\mathbf{u}_{\mathrm{ss}}}^{\mathrm{L}}) \quad (6.54)$$

with equality only in the trivial case where the systems are identical.

6.2.2.2 Numerical characterization of $\mathcal{J}_{\mathbf{u}_{\mathrm{ss}}}^{\mathrm{Q}}$ and $\mathcal{J}_{\mathbf{u}_{\mathrm{ss}}}^{\mathrm{L}}$

In the general case (where \mathbf{N} is not a uniform diagonal matrix), we need to rely on numerical characterizations. Figure 6.4 shows mean and standard deviation of $\mathcal{J}_{\mathbf{u}_{\mathrm{ss}}}^{\mathrm{Q}}$ and $\mathcal{J}_{\mathbf{u}_{\mathrm{ss}}}^{\mathrm{L}}$ over 50 realizations of networks sized $n = 100$. Matrices \mathbf{A}_i $i \in \{1, \dots, 50\}$ were constructed to be stable, and matrices \mathbf{B}_i $i \in \{1, \dots, 50\}$ were created by choosing random positive numbers from a uniform distribution over $(0, 1)$ and then scaling each row to have unit norm. Parameters were identical to those in Figure 6.3. We see that, as in the special case, the quasilinear Fisher information $\mathcal{J}_{\mathbf{u}_{\mathrm{ss}}}^{\mathrm{Q}}$ is quite poor compared to the linear Fisher information $\mathcal{J}_{\mathbf{u}_{\mathrm{ss}}}^{\mathrm{L}}$. However, the extent of this information disadvantage is critically dependent on saturation limits α and β , as we show in Section 6.4.

If we consider this carefully, it is not surprising. Looking at Figures 6.1 and 6.2, it is clear that the ‘cloud’ of covariance surrounding the steady-state mean of the original nonlinear system is under-represented in the Jacobian linearization, while it is more accurately represented in the quasilinear approximation. Since the covariance is wider for the nonlinear (hence quasilinear) system, there is more uncertainty about where the actual state may be relative to its expectation, which is explicitly given in (6.34). Thus, though the Jacobian linearization

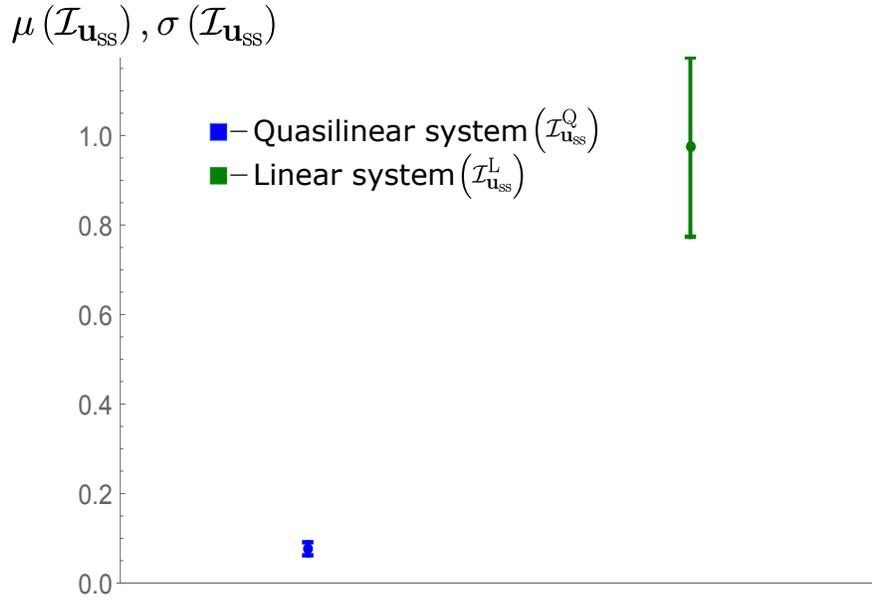


Figure 6.4: Mean and standard deviation of $\mathcal{I}_{\mathbf{u}_{\text{ss}}}^{\text{Q}}$ and $\mathcal{I}_{\mathbf{u}_{\text{ss}}}^{\text{L}}$ over 50 realizations of networks sized $n = 100$. Parameters as in Figure 6.3.

is a much poorer approximation of the sigmoidal system, it nonetheless more accurately encodes its input in the state.

What is even more interesting, is that, given an alternative formulation, based, perhaps, on more realistic assumptions, this phenomenon exactly reverses, as we show in the next section.

6.3 Problem formulation 2

An important consideration which may be easily overlooked is the placement of noise in the equations (6.3) and (6.5). By placing the noise outside the arguments of the saturating function and quasilinear gain, respectively, we are making an implicit assumption about the nature of that noise—namely that the source of the noise is from outside the dynamical ‘channels’ and thus conceptually corresponds to an external type of noise disturbance. But this assumption is not always warranted. In particular, since we are dealing with a multi-node, recurrent network, noise in the ‘signal’ of one state variable will propagate to the others

dynamically. If we make the assumption that the weighted sum (via dynamics matrix \mathbf{A}) of these noises is Gaussian (this would trivially hold if the variables were independent, which of course they are not), and/or that there is Gaussian noise on the input signal \mathbf{u} , then it makes sense to model the system with the noise inside the saturation function (and thus inside the argument scaled by quasilinear gain \mathbf{N}). Thus the dynamical equations for \mathbf{x} and $\hat{\mathbf{x}}$ become

$$\dot{\mathbf{x}}(t) = f(\mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u} + \mathbf{G}\mathbf{z}(t)) \quad (6.55)$$

$$\dot{\hat{\mathbf{x}}}(t) = \mathbf{N}(\mathbf{A}\hat{\mathbf{x}} + \mathbf{B}\mathbf{u} + \mathbf{G}\hat{\mathbf{z}}(t)) - \mathbf{N}\boldsymbol{\mu} \quad (6.56)$$

This does not have major implications for the mathematics of the derivation of optimal gain \mathbf{N} and bias $\boldsymbol{\mu}$, but it does have important implications for the solutions of the fixed point equations (6.21) and (6.22) because of the change in the variance of \mathbf{x} (hence $\hat{\mathbf{x}}$) and thus of \mathbf{y} and $\hat{\mathbf{y}}$, which now become $\mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u} + \mathbf{G}\mathbf{z}$ and $\mathbf{A}\hat{\mathbf{x}} + \mathbf{B}\mathbf{u} + \mathbf{G}\mathbf{z} + \boldsymbol{\mu}$, respectively. Specifically, (6.23) becomes

$$(\mathbf{N}\mathbf{A})\boldsymbol{\Sigma}_N^* + \boldsymbol{\Sigma}_N^*(\mathbf{N}\mathbf{A})^T + \mathbf{N}\mathbf{G}(\mathbf{N}\mathbf{G})^T = \mathbf{0} \quad (6.57)$$

and (6.24), (6.25) become

$$\boldsymbol{\sigma}_{\hat{\mathbf{y}}_i}^2 = \text{var}(\hat{\mathbf{y}}_i) = \mathbf{M}_{i,i}^* \quad (6.58)$$

where

$$\mathbf{M}^* = \mathbf{A}\boldsymbol{\Sigma}_N^*\mathbf{A}^T + \mathbf{G}\mathbf{G}^T \quad (6.59)$$

The resulting values for \mathbf{N} and $\boldsymbol{\mu}$ (obtained from (6.27) and (6.28), now using $\boldsymbol{\sigma}_{\hat{\mathbf{y}}_i}$ obtained from (6.58) instead of (6.24)) are again quite good, as we see in Figures 6.5 and 6.6. Interestingly, the figures show that the covariance ‘cloud’ of the quasilinear system is very similar to

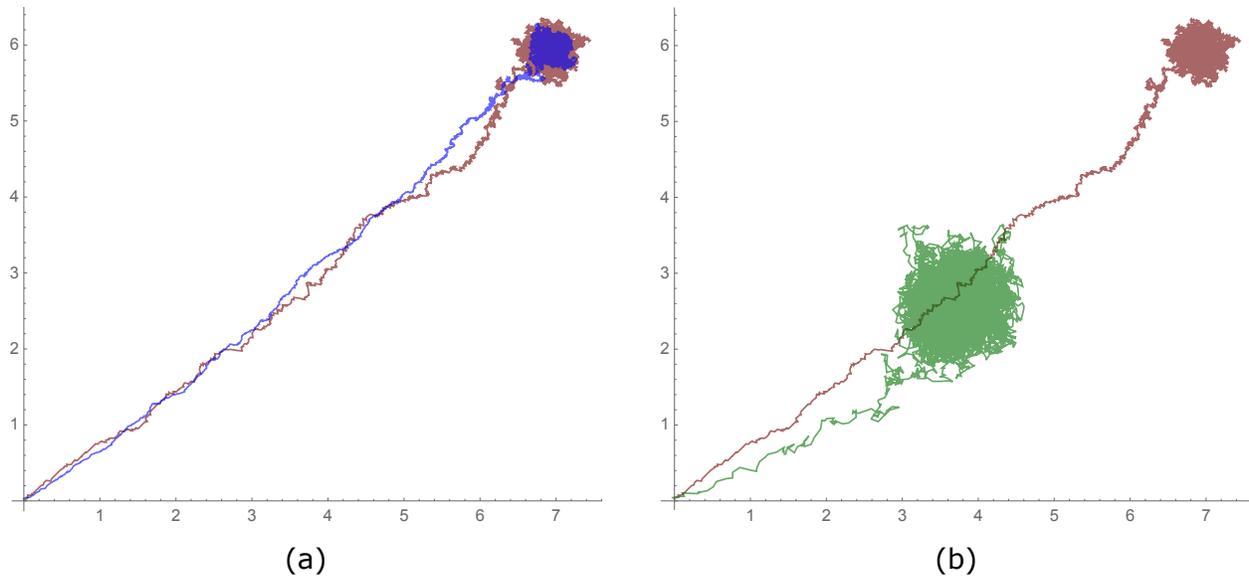


Figure 6.5: An example illustrating the quality of the quasilinear approximation in 2D, under formulation 2. (a) shows the state trajectory of the sigmoidal system (red) and the quasilinear system (blue). (b) shows the same sigmoidal trajectory (red) and the (heuristic) Jacobian-linearized system trajectory (green). Notice that, in contrast with Figure 6.1, where the Jacobian linearization under-represents the variance of the original system, here it over-represents the variance. In both cases the quasilinear approximation is clearly much better.

that of the original (sigmoidal) system, but that the heuristic-Jacobian linearization grossly over-represents this covariance, in striking contrast to our initial formulation (noise outside sigmoid), where the heuristic-Jacobian linearization grossly under-represents this covariance. This is somewhat intuitive mathematically, since the sigmoidal saturation (and quasilinear, sub-unitary gain) have the effect of ‘squashing’ the noise effects when the noise is inside, while when the noise is outside the noise effect is not diminished (except through its dynamical effects on the state variables).

A statistical characterization of the accuracy of the quasilinear approximation and a comparison of this accuracy with that of the (heuristic) Jacobian linearization, is given in Figure 6.7. The setup is identical to Figure 6.3 but here we use formulation 2, and some of the parameter values are adjusted (exact values in the figure caption). There reasons for adjusting the

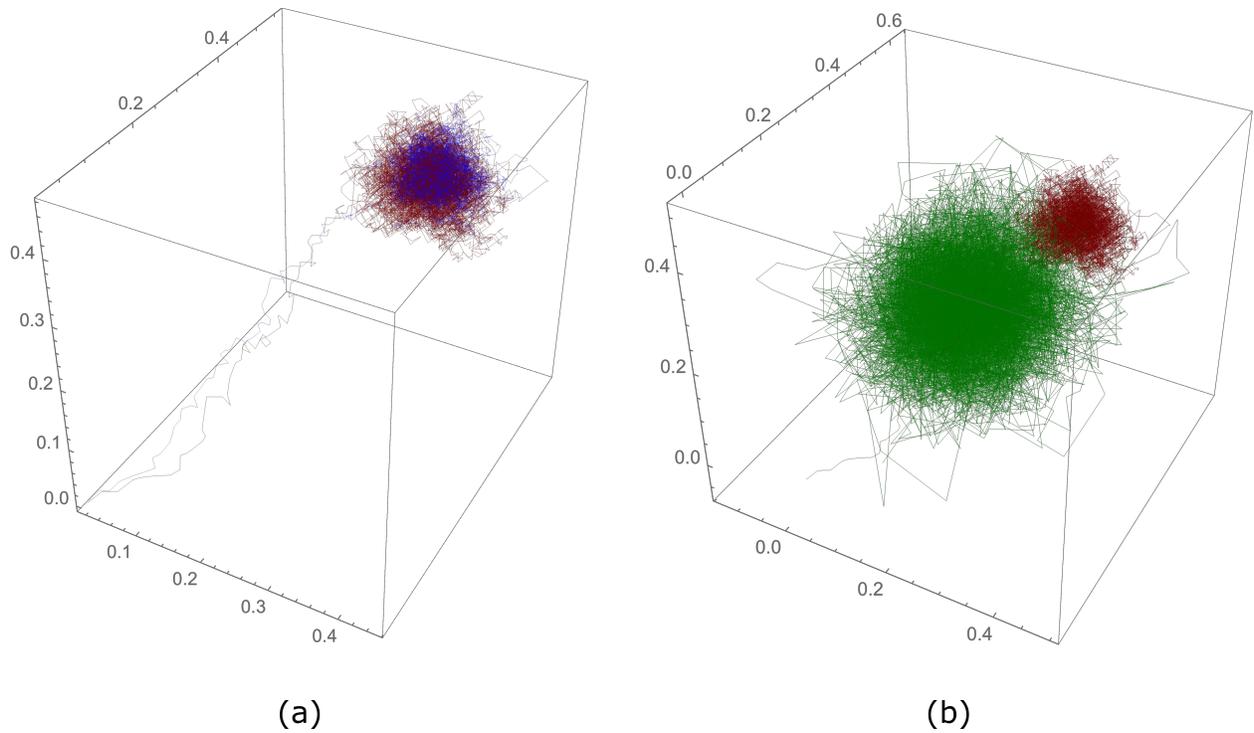


Figure 6.6: An example illustrating the quality of the quasilinear approximation in 3D, under formulation 2. (a) shows the state trajectory of the sigmoidal system (red) and the quasilinear system (blue). (b) shows the same sigmoidal trajectory (red) and the (heuristic) Jacobian-linearized system trajectory (green). Again the Jacobian linearization over-represents the variance of the original system, while the quasilinear approximation closely represents the original system dynamics.

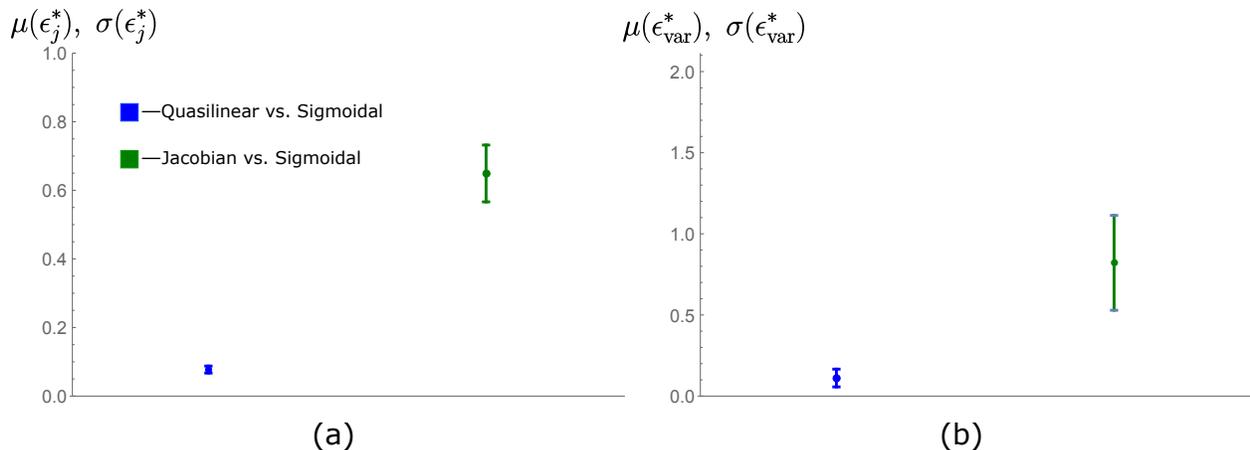


Figure 6.7: We ran simulations of 50 networks each of quasilinear, Jacobian-linearized, and sigmoidal systems ((6.56), (6.29), and (6.55), respectively) with the same dynamical matrices (\mathbf{A}, \mathbf{B}), same noise gain matrix ($\mathbf{G} = .9 * \mathbf{I}$), the same (random for each trial) constant input \mathbf{u} (with $3 \leq \mathbf{u}_i \leq 9 \forall i \in \{1, \dots, n\}$) and same noise input ($\mathbf{z}(t)$) for each trial. Each trial was run using a forward-Euler integration method with $k = 10,000$ time steps, and saturation limits $\alpha = -.5$ and $\beta = 1.5$. (a) shows mean and standard deviation of the value of (6.31) over the 50 trials. Note that for the Jacobian-linearized vs sigmoidal error (green), we just use \mathbf{x}_j (see (6.29)) in place of $\hat{\mathbf{x}}$ in (6.31). (b) shows mean and standard deviation of the variance discrepancy, as quantified by (6.32), over the 50 trials.

parameters are 1) to provide a more holistic view of the quasilinear approximation generally
 2) under some parameter configurations, the system remained in a linear regime the whole time, in which case the sigmoid would never be invoked, and a trivial equality between systems (6.56), (6.29), and (6.55) is observed. Again we ran 50 trials with networks of size $n = 100$, computed normed differences of state trajectories and variances as quantified by (6.31) and (6.33), and plotted mean and standard deviation for the 50 trials.

6.3.1 Quasilinear Fisher information under formulation 2

For the system (6.56), the distribution of state values at time t follows

$$\mathbf{x}(t) \sim \mathcal{N} \left(\int_0^t e^{\mathbf{N}\mathbf{A}(t-\tau)} d\tau \mathbf{N}\mathbf{B}\mathbf{u} + \int_0^t e^{\mathbf{N}\mathbf{A}(t-\tau)} d\tau \mathbf{N}\boldsymbol{\mu}, \boldsymbol{\Sigma}_N^* \right), \quad (6.60)$$

where $\Sigma_N^* \in \mathbb{R}^{n \times n}$ is the covariance matrix given as the solution to (6.57). Again, we make the assumption that the noise covariance is always at steady-state.

The expression for Fisher information $\mathcal{J}_{\mathbf{u}}^{\mathbf{Q}^*}(t)$ under formulation 2 is exactly the same as under formulation 1, with the exception that Σ_N^* replaces Σ_N in (6.37) so that

$$\mathcal{J}_{\mathbf{u}}^{\mathbf{Q}^*}(t) = \mathbf{B}^T \mathbf{N}^T \Gamma_N(t)^T \Sigma_N^{*-1} \Gamma_N(t) \mathbf{N} \mathbf{B} \quad (6.61)$$

Therefore, by (6.61) and (6.44), the steady-state Fisher information $\mathcal{J}_{\mathbf{u}_{\text{ss}}}^{\mathbf{Q}^*}$ is given by

$$\begin{aligned} \mathcal{J}_{\mathbf{u}_{\text{ss}}}^{\mathbf{Q}^*} &= \mathbf{B}^T \mathbf{N} (\mathbf{N} \mathbf{A})^{-T} \Sigma_N^{*-1} (\mathbf{N} \mathbf{A})^{-1} \mathbf{N} \mathbf{B} \\ &= \mathbf{B}^T \mathbf{A}^{-T} \Sigma_N^{*-1} \mathbf{A}^{-1} \mathbf{B} \end{aligned} \quad (6.62)$$

We will characterize the statistical behavior of (6.62) (especially as it relates to linear information (6.51)) numerically in Sections 6.3.1.2 and 6.4, but first we will again treat the special case of uniform diagonal \mathbf{N} , deriving a result which is precisely the opposite of that for formulation 1.

6.3.1.1 Analytic relation between linear and quasilinear Fisher information (under formulation 2) for a special case—uniform diagonal \mathbf{N}

To compare linear Fisher information (for the system (6.29)) and the quasilinear Fisher information under formulation 2 (for (6.56)), we again consider the special case where

$$\tilde{\mathbf{N}} \triangleq \lambda \mathbf{I} \quad (6.63)$$

for some $0 < \lambda \leq 1$. We again relate respective covariances—this time Σ and Σ_N^* —with a fundamentally different result. We show that, where (6.46) holds, if Σ is a solution to

$$\mathbf{A}\Sigma + \Sigma\mathbf{A}^T + \mathbf{G}\mathbf{G}^T = \mathbf{0} \quad (6.64)$$

then

$$\Sigma_N^* = \tilde{\mathbf{N}}\Sigma \quad (6.65)$$

is a solution to (6.57). Starting with a recapitulation of (6.57), we have

$$\begin{aligned} & (\tilde{\mathbf{N}}\mathbf{A})\Sigma_N^* + \Sigma_N^*(\tilde{\mathbf{N}}\mathbf{A})^T + (\tilde{\mathbf{N}}\mathbf{G})(\tilde{\mathbf{N}}\mathbf{G})^T = \mathbf{0} \\ \implies & (\tilde{\mathbf{N}}\mathbf{A})\tilde{\mathbf{N}}\Sigma + \tilde{\mathbf{N}}\Sigma(\tilde{\mathbf{N}}\mathbf{A})^T + \tilde{\mathbf{N}}\mathbf{G}\mathbf{G}^T\tilde{\mathbf{N}}^T = \mathbf{0} \\ \implies & \tilde{\mathbf{N}}\tilde{\mathbf{N}}\mathbf{A}\Sigma + \tilde{\mathbf{N}}\tilde{\mathbf{N}}\Sigma\mathbf{A}^T + \tilde{\mathbf{N}}\tilde{\mathbf{N}}\mathbf{G}\mathbf{G}^T = \mathbf{0} \\ \implies & \tilde{\mathbf{N}}\tilde{\mathbf{N}}(\mathbf{A}\Sigma + \Sigma\mathbf{A}^T + \mathbf{G}\mathbf{G}^T) = \mathbf{0} \\ \implies & \mathbf{A}\Sigma + \Sigma\mathbf{A}^T + \mathbf{G}\mathbf{G}^T = \mathbf{0} \end{aligned} \quad (6.66)$$

$$(6.67)$$

which is identical to (6.47). Again we have used the facts that $\tilde{\mathbf{N}}$ (and thus $\tilde{\mathbf{N}}^{-1}$) is symmetric and commutes with any matrix (due to (6.46)).

Therefore, substituting (6.65) into (6.62), we have steady-state quasilinear information

$$\begin{aligned} \mathcal{J}_{\mathbf{u}_{\text{ss}}}^{\mathbf{Q}^*} &= \mathbf{B}^T \mathbf{A}^{-T} \Sigma^{-1} \tilde{\mathbf{N}}^{-1} \mathbf{A} \mathbf{B} \\ &= \tilde{\mathbf{N}}^{-1} \mathbf{B}^T \mathbf{A}^{-T} \Sigma^{-1} \mathbf{A} \mathbf{B} \end{aligned} \quad (6.68)$$

Thus, in comparison with the Fisher information for the linear system (6.29)–given by (6.51)–we have

$$\mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{Q}^*} = \tilde{\mathbf{N}}^{-1} \mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{L}} = \frac{1}{\lambda} \mathbf{I} \mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{L}} \quad (6.69)$$

and

$$\text{tr} \left(\mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{Q}^*} \right) = \text{tr} \left(\frac{1}{\lambda} \mathbf{I} \mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{L}} \right) = \frac{1}{\lambda} \text{tr} \left(\mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{L}} \right) \quad (6.70)$$

Since λ (and thus the diagonal elements of $\tilde{\mathbf{N}}$) are always between 0 and 1 (see (6.38)), the quasilinear Fisher information under formulation 2 is always better (unless $\lambda = 1$, in which case (6.5) reduces to (6.29)) than the corresponding linear Fisher information. That is,

$$\text{tr} \left(\mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{Q}^*} \right) \geq \text{tr} \left(\mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{L}} \right) \quad (6.71)$$

with equality only in the trivial case where the systems are identical.

6.3.1.2 Numerical characterization of $\mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{Q}^*}$ and $\mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{L}}$

Again, for the general case (where \mathbf{N} is not a uniform diagonal matrix), we need to rely on numerical characterizations. Figure 6.8 shows mean and standard deviation of $\mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{Q}^*}$ and $\mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{L}}$ over 50 realizations of networks sized $n = 100$. \mathbf{A} and \mathbf{B} matrices were created as for Figure 6.4. Parameters were identical to those in Figure 6.7. We see that, as in the special case, the quasilinear Fisher information $\mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{Q}^*}$ is completely superior compared to the linear Fisher information $\mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{L}}$. However, the extent of this information advantage is critically dependent on saturation limits α and β , as we show in Section 6.4.

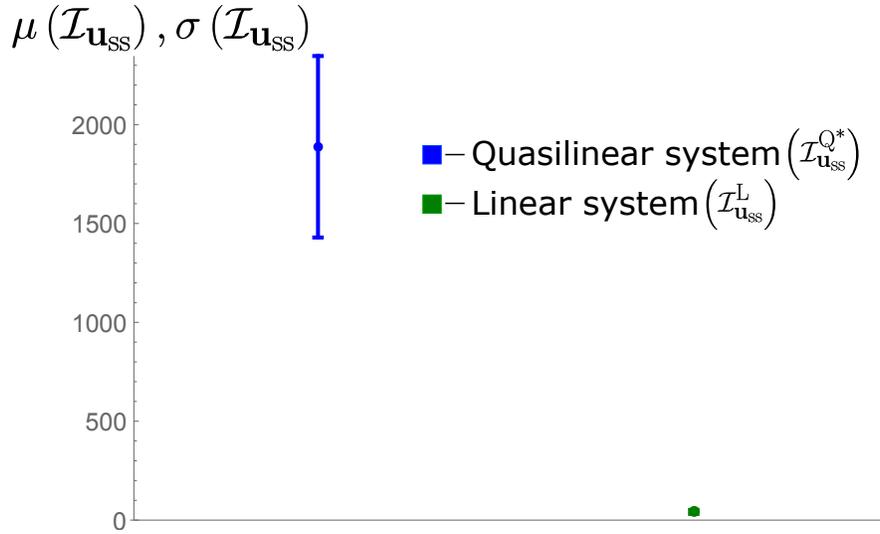


Figure 6.8: Mean and standard deviation of $\mathcal{I}_{\mathbf{u}_{ss}}^{\mathbf{Q}^*}$ and $\mathcal{I}_{\mathbf{u}_{ss}}^{\mathbf{L}}$ over 50 realizations of networks sized $n = 100$. Parameters as in Figure 6.7.

6.4 Fisher information differences as a function of saturation limits α and β

Since, as is seen in Figures 6.4 and 6.8, there is such a pronounced difference in information content of the quasilinear and linear systems, we decided to see how this difference changed as a function of saturation limits α and β . Clearly, if $\alpha \ll 0$ and $\beta \gg 0$, the saturation function $f(\cdot)$ will never ‘kick in’ and thus the system will always remain in the linear regime. In this case we expect $\mathcal{I}_{\mathbf{u}_{ss}}^{\mathbf{Q}} = \mathcal{I}_{\mathbf{u}_{ss}}^{\mathbf{L}} = \mathcal{I}_{\mathbf{u}_{ss}}^{\mathbf{Q}^*}$ and $\mathcal{I}_{\mathbf{u}_{ss}}^{\mathbf{L}}$. At the other extreme, if the difference $|\alpha - \beta|$ is very small (with $\alpha < 0$ and $\beta > 0$) then $f(\cdot)$ will almost always be invoked. In this case, we expect the diagonal values of \mathbf{N} to be very small (recall (6.18), which shows that $\mathbf{N} = \text{diag}(\mathbf{E}[f'(\mathbf{y}(t))])$) because the derivative of $f(\cdot)$ is zero outside the saturation limits and thus its expected value will be very small.

To explore this systematically, we varied $|\alpha|$ and $|\beta|$ (with $\alpha < 0 < \beta$) concurrently from .5 to 10 in increments of .5, and ran 10 trials with network size $n = 40$ for each (α, β) pair. **A**

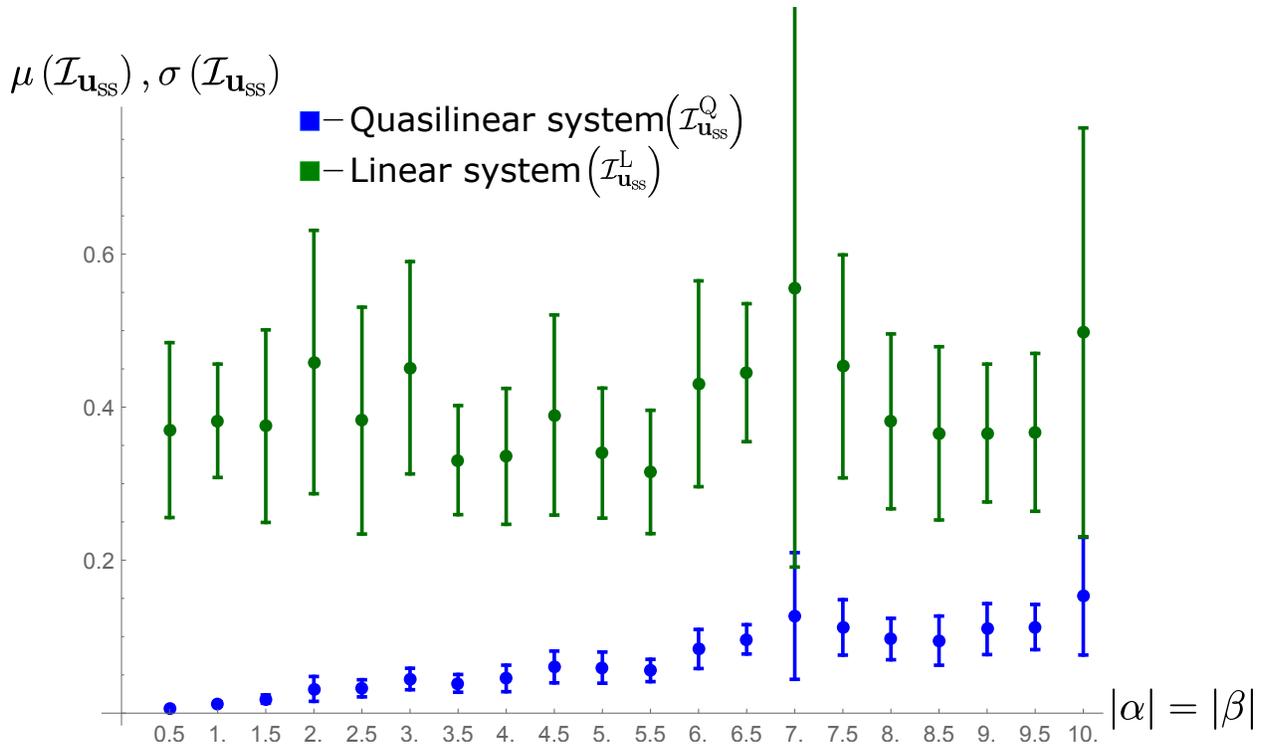


Figure 6.9: We varied $|\alpha| = |\beta|$, where $\alpha < 0 < \beta$, between .5 and 10 (in increments of .5), simulating 10 networks of size $n = 40$ for each (α, β) pair. Matrices \mathbf{A}_i $i \in \{1, \dots, 10\}$ were constructed to be stable, and matrices \mathbf{B}_i $i \in \{1, \dots, 10\}$ were created by choosing random positive numbers from a uniform distribution over $(0, 1)$ and then scaling each row to have unit norm. The same noise gain matrix ($\mathbf{G} = 2.3 * \mathbf{I}$) was used throughout. For each trial, the same noise input ($\mathbf{z}(t)$) was applied to the quasilinear and linear system. Shown, for each (α, β) pair, are the mean and standard deviation of the quasilinear (under formulation 1) information $\mathcal{I}_{\mathbf{u}_{\text{ss}}}^{\text{Q}}$ and linear information $\mathcal{I}_{\mathbf{u}_{\text{ss}}}^{\text{L}}$ over the 10 network realizations.

and \mathbf{B} matrices were created as for Figures 6.4 and 6.8, and parameters (excluding α and β) were as for Figure 6.3.

The results confirm intuition: Under formulation 1, larger $|\alpha| = |\beta|$ (i.e. a wider ‘no-saturation’ range) resulted in better information values $\mathcal{I}_{\mathbf{u}_{\text{ss}}}^{\text{Q}}$, while $\mathcal{I}_{\mathbf{u}_{\text{ss}}}^{\text{L}}$, being obviously unaffected by the saturation function, simply expressed a natural variability—see Figure 6.9. It can be remarked that, compared to the results under formulation 2 (Figure 6.10), the information values $\mathcal{I}_{\mathbf{u}_{\text{ss}}}^{\text{Q}}$ (and $\mathcal{I}_{\mathbf{u}_{\text{ss}}}^{\text{L}}$) are quite small here. We will explore this phenomenon in detail, proving analytically (under mild restrictions) what we see here numerically.

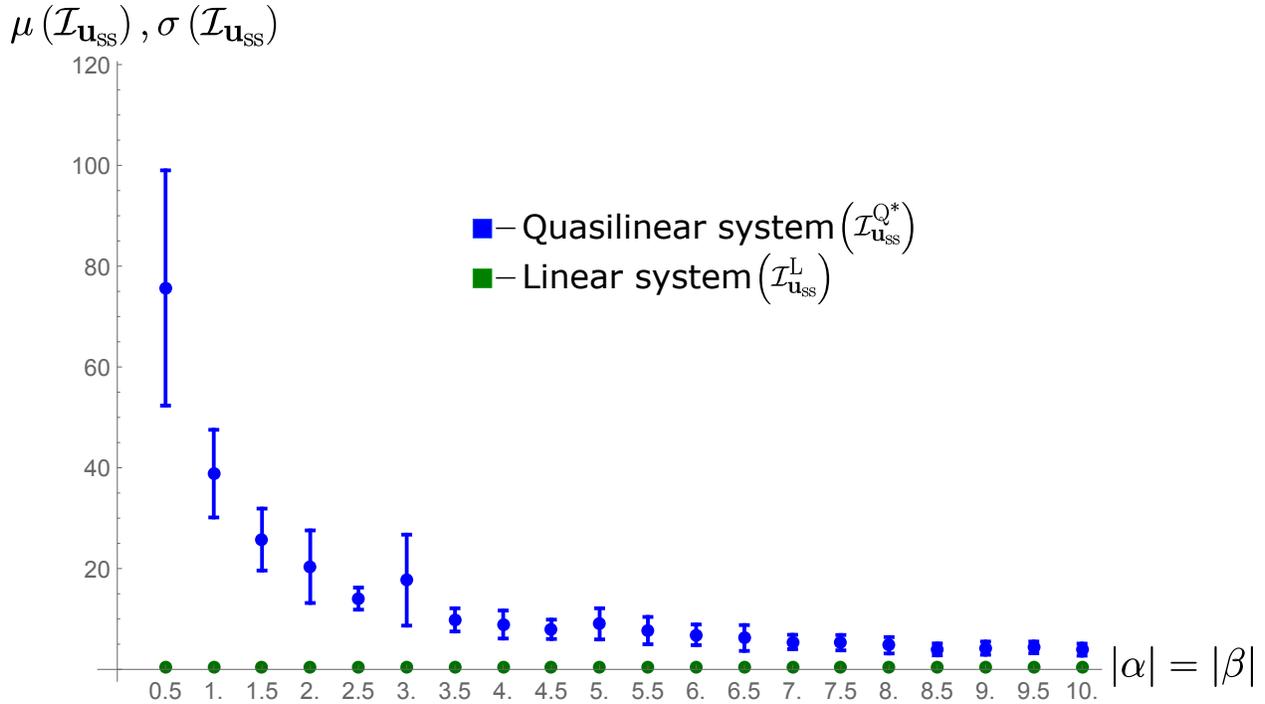


Figure 6.10: Identical setup and parameters as for Figure 6.9, except here information is evaluated for quasilinear information $(\mathcal{I}_{\mathbf{u}_{\text{SS}}}^{\text{Q}^*})$ under formulation 2 vs linear information $\mathcal{I}_{\mathbf{u}_{\text{SS}}}^{\text{L}}$.

Under formulation 2, it is seen that information for the quasilinear system $\mathcal{I}_{\mathbf{u}_{\text{SS}}}^{\text{Q}^*}$ actually decreases with increasing $|\alpha| = |\beta|$, following a clear power law (with negative exponent), to approach convergence with $\mathcal{I}_{\mathbf{u}_{\text{SS}}}^{\text{L}}$. What is fascinating is how much the saturation function (here apprised by its quasilinear approximation) acts to increase information. While the reasons for this are not totally clear, we can say that quasilinear gain \mathbf{N} scales inversely (though not linearly) with the covariance as it relates to α and β . Simply put, the more the argument to $f(\cdot)$ (we denoted it \mathbf{y} previously) strays outside the saturation bounds, the smaller is the expected slope of $f(\cdot)$ and hence the smaller is \mathbf{N} (see (6.18)). Thus when the saturation bounds are tight, the expected value of \mathbf{N} entries can be very small. This, of course, reduces the magnitude of noise covariance, conferring an informational advantage.

There is a price paid for this informational advantage, however. For exactly the same reason—the smallness of \mathbf{n} when $|\alpha|$ and $|\beta|$ are small—the dynamics can be scaled-down to a point that

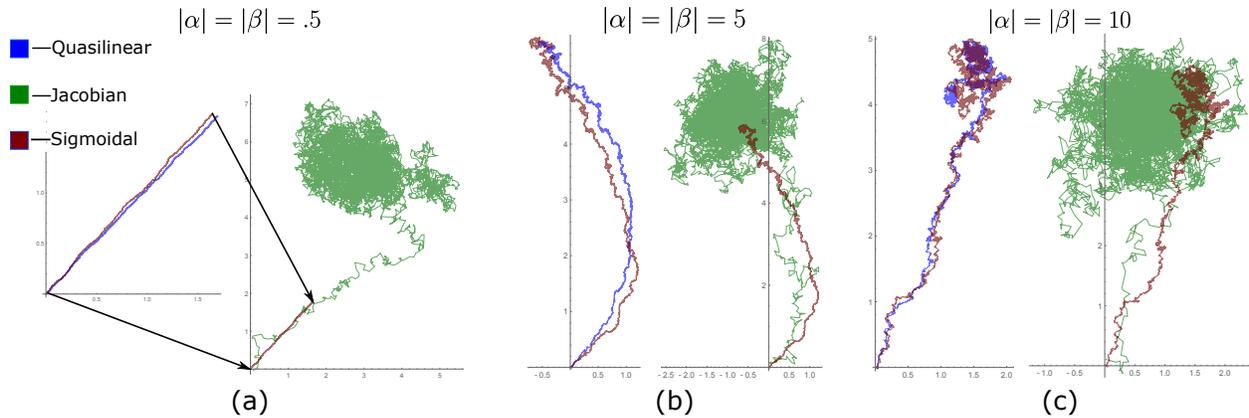


Figure 6.11: Projection onto the first 2 of $n = 40$ dimensions. All parameters are identical to those used in Figures 6.9 and 6.10. Here we chose three values for $|\alpha| = |\beta|$ —.5, 5, and 10—to see how state-space trajectories evolved as a function of time. Terminal time $t = 10$ seconds for all panels. Notice that for $|\alpha| = |\beta| = .5$ (panel (a)), the time-evolution for the sigmoidal (red) and quasilinear (blue) systems is very slow compared to the Jacobian-linearized (green) system. This is predicted by the quasilinear theory, which says a small expected slope of $f(\cdot)$ leads to small values on \mathbf{N} 's diagonal and thus a slowing of dynamics. This phenomenon is still present to some degree for $|\alpha| = |\beta| = 5$ (panel (b)), but for $|\alpha| = |\beta| = 10$ (panel (c)) there is a nice blend of steady-state behavior in the sigmoidal and quasilinear trajectories combined with better information (evidenced by the smaller covariance ‘cloud’) relative to the Jacobian trajectory. This figure again serves to illustrate the quality of the quasilinear approximation generally, since even for higher-dimensional networks tracking of the sigmoidal system is good in all dimensions.

system evolution becomes very slow. In Figure 6.11, we see how the trajectories (projected down to 2 dimensions) at time $t = 10$ seconds compare for different values of $|\alpha| = |\beta|$. It is clear that for the sigmoidal Thus, while there may be a large informational advantage in steady-state, in practice it may take the system a very long time to reach the steady state. Though we do not address this competition between information and time analytically here, a weighted optimization problem could be posed in this regard. Intuitively, the knee (i.e. where the slope of the tangent is about -1) of the curve shown in Figure 6.10 looks like a good trade-off point. It would be interesting to investigate whether, in real networks with sigmoidal nonlinearities, saturation limits tend to be close to such points.

6.5 Relation between Fisher information under the two formulations

We would like to now explore the relation between $\mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{Q}^*}$ and $\mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{Q}}$. We recall that, under condition of uniform diagonal \mathbf{N} (see (6.46)), we have $\mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{Q}} = \mathbf{N} \mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{L}}$ (see (6.52)) and $\mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{Q}^*} = \mathbf{N}^{-1} \mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{L}}$ (see (6.69)) and thus

$$\mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{Q}^*} = \mathbf{N}^{-2} \mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{Q}} \quad (6.72)$$

in this special case.

However, because simulations showed a more generally consistent relationship—even when \mathbf{N} does not have uniform diagonal, which it usually does not—we explored this relation under more ‘realistic’ conditions—namely, where \mathbf{A} and \mathbf{B} are symmetric, and where \mathbf{G} is a uniform diagonal matrix; that is, where

$$\mathbf{G} = \gamma \mathbf{I} \quad (6.73)$$

for some $\gamma \in \mathbb{R}^{++}$. Essentially this means that all state variables $\hat{\mathbf{x}}_i$ will get noise with the same variance and that the noise to $\hat{\mathbf{x}}_i$ will be uncorrelated with noise to $\hat{\mathbf{x}}_j \forall i \neq j$. This is a quite common assumption.

In particular, we will focus on the relation between holistic information measures $\text{tr}(\mathcal{J}_{\mathbf{u}}^{\mathbf{Q}^*})$ and $\text{tr}(\mathcal{J}_{\mathbf{u}}^{\mathbf{Q}})$.

First, we recall that the Lyapunov solution (to (6.57)) for the quasilinear system (6.56) is the infinite-time controllability gramian with respect to noise. That is,

$$\Sigma_N^* = \int_0^\infty e^{\mathbf{N}\mathbf{A}\tau} (\mathbf{N}\mathbf{G})(\mathbf{N}\mathbf{G})^T e^{(\mathbf{N}\mathbf{A})^T \tau} d\tau \quad (6.74)$$

By the symmetry of \mathbf{A} and \mathbf{N} , and the fact that \mathbf{G} is a uniform diagonal matrix, and thus commutes with any matrix, we have:

$$\begin{aligned}
\Sigma_N^* &= \mathbf{G}^2 \int_0^\infty e^{\mathbf{NA}\tau} \mathbf{N} \mathbf{N} e^{(\mathbf{NA})^T \tau} d\tau \\
&= \mathbf{G}^2 \int_0^\infty \left(\mathbf{I} + \mathbf{NA}\tau + \frac{(\mathbf{NA}\tau)^2}{2!} + \dots \right) \mathbf{N} \mathbf{N} \left(\mathbf{I} + \mathbf{AN}\tau + \frac{(\mathbf{AN}\tau)^2}{2!} + \dots \right) d\tau \\
&= \mathbf{G}^2 \int_0^\infty \left(\mathbf{N} + \mathbf{NAN}\tau + \frac{(\mathbf{NA})^2 \mathbf{N} \tau^2}{2!} + \dots \right) \left(\mathbf{N} + \mathbf{NAN}\tau + \frac{\mathbf{N}(\mathbf{AN})^2 \tau^2}{2!} + \dots \right) d\tau \\
&= \mathbf{G}^2 \int_0^\infty \left(\mathbf{N} + \mathbf{NAN}\tau + \frac{\mathbf{NANAN}\tau^2}{2!} + \dots \right) \left(\mathbf{N} + \mathbf{NAN}\tau + \frac{\mathbf{NANAN}\tau^2}{2!} + \dots \right) d\tau \\
&= \mathbf{G}^2 \int_0^\infty \mathbf{N} \left(\mathbf{I} + \mathbf{AN}\tau + \frac{\mathbf{ANAN}\tau^2}{2!} + \dots \right) \left(\mathbf{I} + \mathbf{NA}\tau + \frac{\mathbf{NANA}\tau^2}{2!} + \dots \right) \mathbf{N} d\tau \\
&= \mathbf{G}^2 \mathbf{N} \int_0^\infty e^{(\mathbf{NA})^T \tau} e^{(\mathbf{NA})\tau} d\tau \mathbf{N} \\
&= \mathbf{N} \left(\mathbf{G}^2 \int_0^\infty e^{(\mathbf{NA})^T \tau} e^{(\mathbf{NA})\tau} d\tau \right) \mathbf{N} \tag{6.75}
\end{aligned}$$

Now, by the uniformity and symmetry of diagonal matrix \mathbf{G} (see (6.73)), we have

$$\mathbf{G}^2 \int_0^\infty e^{(\mathbf{NA})^T \tau} e^{(\mathbf{NA})\tau} d\tau = \int_0^\infty e^{(\mathbf{NA})^T \tau} \mathbf{G}^T \mathbf{G} e^{(\mathbf{NA})\tau} d\tau \tag{6.76}$$

is by definition the transpose of the infinite-time controllability gramian for the system (under formulation 1) (6.5), and since the controllability gramian is symmetric, (6.76) is precisely Σ_N , the solution to the Lyapunov equation (6.23) for the quasilinear system under formulation 1.

Thus, (6.75) becomes

$$\Sigma_N^* = \mathbf{N} \Sigma_N \mathbf{N} \tag{6.77}$$

Now, by substituting the RHS of (6.77) for Σ_N^* in 6.61, and recalling, from (6.44), that $\Gamma_N(\infty) = \mathbf{A}^{-1}\mathbf{N}^{-1}$ we obtain

$$\begin{aligned}
\mathcal{J}_{\mathbf{u}_{\text{SS}}}^{\mathbf{Q}^*} &= \mathbf{B}^T \mathbf{N}^T \Gamma_N(\infty)^T \Sigma_N^{*-1} \Gamma_N(\infty) \mathbf{N} \mathbf{B} \\
&= \mathbf{B} \mathbf{N}^T (\mathbf{A}^{-1} \mathbf{N}^{-1})^T (\mathbf{N} \Sigma_N \mathbf{N})^{-1} \mathbf{A}^{-1} \mathbf{N}^{-1} \mathbf{N} \mathbf{B} \\
&= \mathbf{B} \mathbf{N} \mathbf{N}^{-1} \mathbf{A}^{-1} \mathbf{N}^{-1} \Sigma_N^{-1} \mathbf{N}^{-1} \mathbf{A}^{-1} \mathbf{N}^{-1} \mathbf{N} \mathbf{B} \\
&= \mathbf{B} \mathbf{A}^{-1} \mathbf{N}^{-1} \Sigma_N^{-1} \mathbf{N}^{-1} \mathbf{A}^{-1} \mathbf{B}
\end{aligned} \tag{6.78}$$

where we have used the symmetry of \mathbf{A} , \mathbf{B} , and \mathbf{N} . Now, since all matrices in (6.78) are symmetric, the trace is invariant to any permutation of matrix ordering. Thus

$$\begin{aligned}
\text{tr} \left(\mathcal{J}_{\mathbf{u}_{\text{SS}}}^{\mathbf{Q}^*} \right) &= \text{tr} \left(\mathbf{B} \mathbf{A}^{-1} \mathbf{N}^{-1} \Sigma_N^{-1} \mathbf{N}^{-1} \mathbf{A}^{-1} \mathbf{B} \right) \\
&= \text{tr} \left(\mathbf{N}^{-2} \mathbf{B} \mathbf{A}^{-1} \Sigma_N^{-1} \mathbf{A}^{-1} \mathbf{B} \right)
\end{aligned} \tag{6.79}$$

Now we derive, proceeding from 6.37, a simplified expression for $\mathcal{J}_{\mathbf{u}_{\text{SS}}}^{\mathbf{Q}}$:

$$\begin{aligned}
\mathcal{J}_{\mathbf{u}_{\text{SS}}}^{\mathbf{Q}} &= \mathbf{B}^T \mathbf{N}^T \Gamma_N(t)^T \Sigma_N^{-1} \Gamma_N(t) \mathbf{N} \mathbf{B} \\
&= \mathbf{B} \mathbf{N} (\mathbf{N} \mathbf{A})^{-T} \Sigma_N^{-1} (\mathbf{N} \mathbf{A})^{-1} \mathbf{N} \mathbf{B} \\
&= \mathbf{B} \mathbf{N} \mathbf{N}^{-1} \mathbf{A}^{-1} \Sigma_N^{-1} \mathbf{A}^{-1} \mathbf{N}^{-1} \mathbf{N} \mathbf{B} \\
&= \mathbf{B} \mathbf{A}^{-1} \Sigma_N^{-1} \mathbf{A}^{-1} \mathbf{B}
\end{aligned} \tag{6.80}$$

Therefore, by (6.80) and 6.79, we have

$$\text{tr} \left(\mathcal{J}_{\mathbf{u}_{\text{SS}}}^{\mathbf{Q}^*} \right) = \text{tr} \left(\mathbf{N}^{-2} \mathcal{J}_{\mathbf{u}_{\text{SS}}}^{\mathbf{Q}} \right) \tag{6.81}$$

Then, since $\mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{Q}}$ is positive definite and thus has positive diagonal values, and \mathbf{N} is diagonal with all positive entries (see 6.38), by an application of the Cauchy-Schwarz inequality, we have

$$\mathrm{tr} \left(\mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{Q}^*} \right) \leq \mathrm{tr} \left(\mathbf{N}^{-2} \right) \mathrm{tr} \left(\mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{Q}} \right) \quad (6.82)$$

Now, by (6.77) we have

$$\begin{aligned} \Sigma_N^* &= \mathbf{N} \Sigma_N \mathbf{N} \\ \implies \Sigma_N &= \mathbf{N}^{-1} \Sigma_N^* \mathbf{N}^{-1} \end{aligned} \quad (6.83)$$

Plugging (6.83) into (6.80) we have

$$\begin{aligned} \mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{Q}} &= \mathbf{B} \mathbf{A}^{-1} \Sigma_N^{-1} \mathbf{A} \mathbf{B} \\ &= \mathbf{B} \mathbf{A}^{-1} \left(\mathbf{N}^{-1} \Sigma_N^* \mathbf{N}^{-1} \right)^{-1} \mathbf{A} \mathbf{B} \\ &= \mathbf{B} \mathbf{A}^{-1} \mathbf{N} \Sigma_N^{*-1} \mathbf{N} \mathbf{A} \mathbf{B} \end{aligned} \quad (6.84)$$

Therefore,

$$\begin{aligned} \mathrm{tr} \left(\mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{Q}} \right) &= \mathrm{tr} \left(\mathbf{B} \mathbf{A}^{-1} \mathbf{N} \Sigma_N^{*-1} \mathbf{N} \mathbf{A} \mathbf{B} \right) \\ &= \mathrm{tr} \left(\mathbf{N}^2 \mathbf{B} \mathbf{A}^{-1} \Sigma_N^{*-1} \mathbf{A}^{-1} \mathbf{B} \right) \end{aligned} \quad (6.85)$$

and by the same Cauchy-Schwarz inequality, we have

$$\mathrm{tr} \left(\mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{Q}} \right) \leq \mathrm{tr} \left(\mathbf{N}^2 \right) \mathrm{tr} \left(\mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{Q}^*} \right) \quad (6.86)$$

But then

$$\mathrm{tr} \left(\mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{Q}^*} \right) \geq \mathrm{tr} \left(\mathbf{N}^{-2} \right) \mathrm{tr} \left(\mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{Q}} \right) \quad (6.87)$$

which is only possible if (6.82) is a strict equality, and thus we have arrived at

$$\text{tr} \left(\mathcal{J}_{\mathbf{u}_{\text{ss}}}^{\mathbf{Q}^*} \right) = \text{tr} \left(\mathbf{N}^{-2} \right) \text{tr} \left(\mathcal{J}_{\mathbf{u}_{\text{ss}}}^{\mathbf{Q}} \right) \quad (6.88)$$

6.6 Discussion

The relation (6.71) and the simulation results (see Figures 6.8 and 6.10), which shows that for system (6.55) (and its approximation (6.56)) QFI>LFI, have potentially important implications regarding the etiology of sigmoidal nonlinearities in complex networks (such as neural networks). It is natural to assume that these nonlinearities arise from fundamental physical limitations, often as a result of homeostatic mechanisms. For example, in a neuron, after the rapid depolarization precipitating an action potential, ion pumps in the cell membrane strongly react to quickly repolarize the cell. This reaction inaugurates what is called a refractory period, during which it is almost impossible for the neuron to fire again. These refractory periods impose a fundamental limit on neuronal firing rate. Thus, when modeling an input-output relationship for the neuron, it is necessary to include a saturating, or sigmoidal, function, which keeps firing with realistic bounds in model simulation.

This describes (part of) the what of the nonlinearity, but it does not address the why. It is natural to assume that the sigmoidal nonlinearity in firing rate is inherently a physical process, without any intrinsic utility (other than the preservation of electrodynamic homeostasis) for network function. However, our results suggest the possibility that the nonlinearity does confer a functional advantage—namely, better (i.e. more informative by the Fisher metric) input encoding in the network state. If this is the case, perhaps this functional advantage was a causative factor in the formation of the network dynamics. We have shown that there is a cost to this advantage (see Figure 6.11)—namely, slower dynamics. Since the

curve for QFI (specifically, the value of $\text{tr}\left(\mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{Q}^*}\right)$ as a function of saturation limits α and β) shown in Figure 6.10 is smooth and displays a clear (negative) power law, there is a knee at which, theoretically, sigmoidal saturation limits α and β would provide an optimal blend of information and speed. By careful experimentation, appropriate modeling and application of some of the ideas presented in this chapter, it may be discovered to what extent real sigmoidal systems, especially those in neural networks, utilize saturation limits corresponding to this optimal ‘knee’. Indeed, were it found that the saturation limits correspond to some sort of informational-dynamical optimum, it would be hard to account for the existence of sigmoidal nonlinearities solely on the basis of physical necessity.

On the other hand, quasilinearization, as we have seen, can closely approximate a system with a sigmoidal nonlinearity. Thus, if functional advantage was the causative factor in the genesis of network dynamics, why would not those dynamics be quasilinear? Presumably, this would be less biologically expensive to implement than a nonlinearity dependent on complex interactions. Thus, ironically, the very analytical technique we used to probe the idea of information coding in sigmoidal networks may be one of the best arguments against an information-based etiology for their existence. The Discussion, to this point, has been conducted with reference to formulation 2—first introduced, in its most general form, in (6.2), then more concretely in (6.55) and (6.56)—where noise was assumed to corrupt the argument of the sigmoidal function, as opposed to formulation 1 (see (6.1), (6.3), and (6.5)), where ‘external’ noise corrupts the system independently of the sigmoid. Let us now try to understand the nature of the profound difference between results obtained from the two formulations—namely that $\text{tr}\left(\mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{Q}^*}\right)$ (formulation 2) is categorically superior, in terms of information coding, to $\text{tr}\left(\mathcal{J}_{\mathbf{u}_{ss}}^{\mathbf{Q}}\right)$ (formulation 1). This difference was seen in the two relations (6.53) and (6.70) (for the special case of \mathbf{N} being a uniform diagonal matrix). It was seen when, under assumption of the symmetry of dynamical matrices \mathbf{A} , \mathbf{B} and equal, independent

noise variance fed to all state variables at any time t (this is distinct from noise covariance aggregated by the dynamics), an explicit relation (6.88) was derived in Section 6.5. Finally, it was illustrated generally by extensive simulations (see Figures 6.9 and compare with 6.10).

Of course, it is a mathematical fact that scaling the noise by \mathbf{N} (as we do when moving from formulation 1 to formulation 2) diminishes its effect overall. This, however, does not account for the entire difference, which is on the order of \mathbf{N}^2 . It also does not explain the superiority of $\mathcal{J}_{\mathbf{u}_{\text{ss}}}^{\text{Q}^*}$ over $\mathcal{J}_{\mathbf{u}_{\text{ss}}}^{\text{L}}$ (information for the Jacobian-linearized system (6.29)), since noise and input are scaled together in (6.56) relative to (unscaled) noise and input in (6.29). While we do not here propose to completely characterize the mechanism for these marked differences in information coding, we do reference [16], where for an admittedly quite different setup (i.e. a two-layer feedforward network), Fisher information (with respect to a scalar signal) was shown to critically depend on the relationship between signal direction and direction of noise covariance. Particularly relevant to our case was their finding that when signal and noise covariance shared the same direction in the first layer of processing, Fisher information was maximized regardless of orientation of noise covariance in the second layer (signal orientation remained constant). Here, since we are dealing with a recurrent network, we may draw a parallel between their ‘first layer’ and our network configuration at some specific time t . Their second layer would correspond to some subsequent time $t + \Delta t$. Although we do not conduct a rigorous analysis with respect to noise gain matrix \mathbf{G} , with which we could precisely control noise covariance, we note that the sigmoidal function (and corresponding quasilinear approximation) limits the range of its argument and thus the extent to which noise can ‘lead the state astray’. Thus the signal (input \mathbf{u}) will play a larger role in determining the direction of noise covariance $\Sigma_{\mathbf{N}}^*$, because this is really the state covariance with respect to noise. Thus, since at time t , state (via noise) covariance is more highly aligned with input direction than would be the case in a basic linear system, at time $t + \Delta t$, we expect, according to reasoning

along the lines of [16], to have more information regarding \mathbf{u} than in the linear case. Since the sigmoid still acts on the dynamics at time $t + \Delta t$, input and noise directions will again be more aligned relative to the linear case, so information will be better at $t + 2\Delta t$, and so on.

One obvious conclusion from these arguments is that, in creating dynamical models, placement of noise in the dynamical equations must be carefully considered. The implications of this placement have been seen in the (analytically and numerically) calculated information values, but more geometrically in the noise covariance ‘clouds’ surrounding the steady-state fixed points of the respective systems. Put simply, a bigger covariance ‘cloud’ means less certainty and hence less information with respect to the stimulus. This phenomenon has been manifest in the equations, as well as in all of the simulation figures. Undoubtedly even more nuanced ways of ‘shaping’ noise covariance geometry in neural-like networks can be found—this issue has been addressed ([15], [77], [78]), often in the context of speech ‘de-corruption’ and/or speech recognition.

Also, a natural question arises: What other types of nonlinearities might confer similar informational advantages, perhaps even to a larger extent, by their continuous transformation of relative signal and noise directions? Or maybe the sigmoid is an optimal function in this respect, in which case we would (again) have to ask: Is it just coincidence that a byproduct of physical constraints (as is usually the assumed etiology of saturating nonlinearities in nature—see above) confers such a pronounced informational advantage?

6.7 Conclusions and future work

With appropriate quasilinear approximation of dynamical systems with sigmoidal nonlinearities, the Fisher information quantifying the efficacy of the nonlinear network in encoding its afferent inputs (in state trajectories) can also be approximated, by calculating the Fisher

information for the quasilinear system. Clearly the information approximation's quality depends on the quality of the initial quasilinear approximation. In this chapter we have shown the methodology of this approximation, and we have shown that the approximation is statistically quite good, especially in comparison with the most commonly-used linearization procedure—Jacobian linearization. We derived the Fisher information matrix for a quasilinear network, and showed analytically that this information is better, or worse, than that of the Jacobian (or, unconstrained) linear system, depending on where noise is placed in the dynamical equations. Taking noise inside the sigmoid, as in formulation 2 (6.3)—probably a more realistic assumption than formulation 1 (6.55)—confers informational advantage relative to the linear system. We discussed reasons for this in terms of the drastically different ‘clouds’ of uncertainty generated by the systems’ respective noise covariances. Further, we showed, analytically and numerically, relations between Fisher information under the two formulations. We addressed the implications of this phenomenon for system modeling and design, and presented preliminary hypotheses regarding its mechanistic origin. We showed also that the informational advantage conferred by the presence of sigmoidal nonlinearities (and their quasilinear approximations) is tempered by the competing objective of time—that is, parameterizations most favorable for high information coding in sigmoidal networks also led to very slow dynamics. Since these two competing objectives varied smoothly, however, the presence of an optimum, or ‘sweet spot’ wherein the system would efficiently propagate noise and simultaneously maintain responsive dynamics, was implicated.

At the highest level, we discussed possible reasons for the presence of sigmoidal nonlinearities in complex (especially neural) networks, evaluating the competing hypotheses 1) that the nonlinearities might exist to confer informational advantage and that this was the prime causative reason for their incorporation into network dynamics, vs. 2) the nonlinearities are byproducts of homeostatic physical mechanisms, and any informational advantage arises

because networks have adapted to functionally excel despite physical limitations. We did not reach any hard conclusions concerning this (indeed it would have been the height of presumptive speculation to do so), but we reasoned that, precisely because of the ability of a network to ‘simulate’ sigmoidal dynamics by a simple linear approximation, it is unlikely that hypotheses (1) is correct because of the biological cost associated with engineering a nonlinearity relative to making a quasilinear adjustment. However, were the presence of sigmoidal parameterization operating in an optimal regime (as discussed above) to be experimentally verified in real (especially neural) networks, it would have to be conceded that the presence of sigmoidal nonlinearities in general is probably due to their conferral of informational advantage to the dynamical systems they modulate.

Chapter 7

Conclusion

7.1 Summary

There are several important conclusions which can be drawn from the work presented here. Taken broadly, we have seen that measurement of information propagation through complex dynamical networks is tractable from conceptual, algorithmic, and computational points of view. The study of information processing over networks has been a topic of considerable interest in the theoretical and computational neuroscience community. We believe that we have, by means of the theoretical derivations, practical analyses, and careful considerations included in this work, furthered this tractability and provided new tools and ideas for use in the field. More specifically, the individual results support an advancing line of reasoning concerning the interplay between the dynamics of a network and its propensity for information processing.

First, we have shown that network sensitivity to small input differences, the quantification of which was derived in Chapter 4.6, is related to, but distinct from, traditional control-theoretic measures, and constitutes the underpinnings of an informational analysis. Simply put, more sensitivity means better information coding. This is especially true when noise is taken into consideration; we addressed this, in the context of discriminability of inputs (to a linear system) based on state-space representations, in Chapter 3, providing a precise geometric characterization of the way in which input differences are translated, via network dynamics, into locations in the reachable state-space. Since sensitivity is foundational to information processing, especially in discrimination/classification scenarios, we surmised that sensitivity must be ‘prioritized’ by informational dynamical systems, such as neural networks, but that there may be an energy cost associated with heightened sensitivity balancing this prioritization. Indeed, we found this to be the case, as described in Chapter 4, where we derive, and analytically deploy, a quantification of the extent to which linear dynamical systems trade off energy for sensitivity as a function of network topology and dynamics. The characterization of this trade-off was shown to depend also on particular directions of state-space traversal; in particular, to reach endpoints on the reachable ellipsoid determined by fixed energy c , but given energy in excess of c , we showed that the amount of excess, being a function of state-space direction \mathbf{v} , is correlated—in a precise and elegant mathematical way—with network sensitivity to inputs driving the system in the direction \mathbf{v} .

Having established the importance of input sensitivity and characterized many of its salient features, we went on to seek a holistic quantification of information flow in dynamical networks, addressing first the linear case. Accordingly, in Chapter 5 we derived a Fisher information-based measurement of the extent to which input information is encoded in the network state. Since we have focused our efforts in the domain of computational neuroscience, we derived this information measure $\mathcal{I}_{\mathbf{u}}$, with respect to a background, or ‘resting’ state,

sometimes referred to in the neuroscience literature as the ‘default mode’. We showed how $\mathcal{J}_{\mathbf{u}}$ can be maximized over the space of possible backgrounds, giving an analytic expression for the optimum. In quantifying information flow, we wanted to address a fundamental, and somewhat natural question: what is the relation between information-theoretic and control-theoretic properties in dynamical networks? Our exploration of this relation, though nuanced, led to the basic conclusion that easily-controlled networks tend to be informative networks. This relation is fundamentally rooted in the idea of sensitivity, which, as we have seen, has been critical in the development of our work. Mathematically, the Fisher information of a dynamical system $\dot{\mathbf{x}} = f(\mathbf{x}, \mathbf{u})$, with respect to the input \mathbf{u} , is proportional to the sensitivity of $f(\cdot)$ to changes in \mathbf{u} —namely, $\frac{df(\cdot)}{d\mathbf{u}}$. But sensitivity to inputs is exactly what controllability analysis seeks to quantify, though framed in terms of sensitivity to the energy of \mathbf{u} .

Finally, in Chapter 6, we have sought to overcome a limitation inherent in all the previous formulations—namely, the assumption of linearity. Specifically, we introduced a novel modification and application of quasilinearization theory, approximating the dynamics of networks with sigmoidal nonlinearities by means of an optimization-based affine transformation. The accuracy of this approximation was confirmed by a statistical analysis comparing the quasilinear approximation with a Jacobian linearization. Using the quasilinear approximation, we derived and evaluated a Fisher information-based measure of network input encoding ability. Analytical and numerical results were given showing that Fisher information is better for the quasilinear (hence sigmoidal) system than for the Jacobian-linearized, ‘unconstrained’ linear system. This relation was shown to reverse when the noise is placed outside the sigmoid in the model, highlighting an important connection between the relative directions of signal and noise covariance, and elucidating the nature of the link between extent of noise covariance and uncertainty.

7.2 Future work

The extension of our analysis into the nonlinear domain (given in Chapter 6) gives promise that continued effort and advancement will yield tractable methods of assessing information for systems with more fundamental nonlinearities, such as are found in the most true-to-life (usually spiking) neural network models. Although the quasilinearization technique is geared only for a special class of nonlinear functions, the basic optimization paradigm may yield other simplifications of nonlinearities.

Practically speaking, the quasilinear techniques could prove of value in artificial neural networks, where sigmoidal saturating functions are typically used to modulate the output activity variables at each layer. In a recurrent network, which can be thought of as an infinite-layered deep network, the combination of easily approximating these sigmoidal functions with algorithmic and computational efficiency, could make quasilinearization a viable alternative to the use of sigmoids in large-scale artificial networks. This at least warrants some careful study.

Finally, we would like to explore the mechanisms for small-ensemble neural representation of complex stimuli via nonlinear, perhaps chaotic, dynamics. Very small non-linear dynamical networks, under certain classes of equations usually involving chaos, symmetry, or both, can literally represent, in their induced state trajectories, many three-dimensional objects (trees, flowers, clouds, etc.). Discrete-time linear systems, with dynamics determined by picking probabilistically from a small set of transformations at each time step (i.e. a Markovian model), can also create very ‘lifelike’, complex state trajectories, which are quite impervious to noise. Is the brain using one of these representation strategies, or some variation thereof? To answer this question would, in addition to shedding light on how the brain operates, advance basic research in math and engineering. The mathematics of chaotic and Markovian dynamics

has been well-described, but thorough explanations on how real systems can integrate these dynamics into specific problem-solving paradigms (here the problem is efficient, accurate encoding) have been lacking.

Of course, any future study in computational neuroscience will be predicated on discoveries which are constantly being made; given the intense push inside (and outside) the academic world for a more thorough understanding of the brain, we expect the rate of these discoveries to accelerate. Control-theoretic ideas and state-space models will become increasingly important as the capacity for precise, single neuron-scale control becomes possible through optogenetic stimulation/inhibition (and other emerging, possibly even less invasive) technologies. Perhaps the world awaits another Galileo or Newton to find simple laws governing much of what now appears to be a tangled mess. Or perhaps another Einstein will come and show us that many of our intuitions about neural function need to be completely discarded and revolutionary concepts adopted. It seems more likely that the consistent, concerted effort of many scientists—whose distinct advantages over these pioneers include a huge, growing body of relevant literature, wonderful technologies for exploration, and instantaneous communication channels allowing collaboration on an unprecedented scale—will, over time, yield a relatively complete picture of the workings of what is considered the most complex single entity in the universe—the human brain.

Bibliography

- [1] N. Hadjikhani, A. K. Liu, A. M. Dale, P. Cavanagh, and R. B. Tootell, “Retinotopy and color sensitivity in human visual cortical area v8,” *Nature neuroscience*, vol. 1, no. 3, p. 235, 1998.
- [2] M. Sereno, S. Pitzalis, and A. Martinez, “Mapping of contralateral space in retinotopic coordinates by a parietal cortical area in humans,” *Science*, vol. 294, no. 5545, pp. 1350–1354, 2001.
- [3] B. W. Zeff, B. R. White, H. Dehghani, B. L. Schlaggar, and J. P. Culver, “Retinotopic mapping of adult human visual cortex with high-density diffuse optical tomography,” *Proceedings of the National Academy of Sciences*, vol. 104, no. 29, pp. 12 169–12 174, 2007.
- [4] R. B. Tootell, N. Hadjikhani, E. K. Hall, S. Marrett, W. Vanduffel, J. T. Vaughan, and A. M. Dale, “The retinotopy of visual spatial attention,” *Neuron*, vol. 21, no. 6, pp. 1409–1422, 1998.
- [5] C. M. Wessinger, M. H. Buonocore, C. L. Kussmaul, and G. R. Mangun, “Tonotopy in human auditory cortex examined with functional magnetic resonance imaging,” *Human brain mapping*, vol. 5, no. 1, pp. 18–25, 1997.
- [6] M. Saenz and D. R. Langers, “Tonotopic mapping of human auditory cortex,” *Hearing research*, vol. 307, pp. 42–52, 2014.
- [7] M. A. Howard III, I. O. Volkov, P. J. Abbas, H. Damasio, M. C. Ollendieck, and M. A. Granner, “A chronic microelectrode investigation of the tonotopic organization of human auditory cortex,” *Brain research*, vol. 724, no. 2, pp. 260–264, 1996.
- [8] F. Pasqualetti and S. Zampieri, “On the controllability of isotropic and anisotropic networks,” 2014. [Online]. Available: <http://www.fabiopas.it/papers/FP-SZ-14.pdf>
- [9] F. Pasqualetti, S. Zampieri, and F. Bullo, “Controllability metrics, limitations and algorithms for complex networks,” *IEEE Transactions on Control of Network Systems*, vol. 1, no. 1, pp. 40–52, Mar. 2014.

- [10] B. Dumitrescu, B. Sicleru, and R. Stefan, “Computing the controllability radius: a semi-definite programming approach,” *IET Control Theory & Applications*, vol. 3, no. 6, pp. 654–660, 2009. [Online]. Available: <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=4976843>
- [11] G. Yan, J. Ren, Y.-C. Lai, C.-H. Lai, and B. Li, “Controlling complex networks: how much energy is needed?” *Physical review letters*, vol. 108, no. 21, p. 218703, 2012.
- [12] G. Hu and E. Davison, “Real controllability/stabilizability radius of LTI systems,” vol. 49, no. 2, pp. 254–257, 2004. [Online]. Available: <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=1266782>
- [13] G. Hu and E. J. Davison, “Real controllability/stabilizability radius of lti systems,” *IEEE Transactions on Automatic Control*, vol. 49, no. 2, pp. 254–257, 2004.
- [14] N. Brunel and J.-P. Nadal, “Mutual information, fisher information, and population coding,” *Neural Computation*, vol. 10, no. 7, pp. 1731–1757, 1998.
- [15] S. Ganguli, D. Huh, and H. Sompolinsky, “Memory traces in dynamical systems,” *Proceedings of the National Academy of Sciences*, vol. 105, no. 48, pp. 18 970–18 975, 2008.
- [16] J. Zylberberg, A. Pouget, P. E. Latham, and E. Shea-Brown, “Robust information propagation through noisy neural circuits,” *PLoS computational biology*, vol. 13, no. 4, p. e1005497, 2017.
- [17] G. Kumar, D. Menolascino, M. Kafashan, and S. Ching, “Controlling linear networks with minimally novel inputs,” in *American Control Conference (ACC)*, 2015, 2015, pp. 5896–5900. [Online]. Available: <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=7172264>
- [18] D. Menolascino, G. Kumar, and S. Ching, “Endpoint-based discriminability of minimum energy inputs,” in *2016 American Control Conference*, 2016.
- [19] D. Menolascino and S. Ching, “Bispectral analysis for measuring energy-orientation tradeoffs in the control of linear systems,” *Systems & Control Letters*, vol. 102, pp. 68–73, 2017.
- [20] D. Menolascino and S. Ching, “Information spectra and optimal background states for dynamical networks,” *Scientific reports*, vol. 8, no. 1, p. 16181, 2018.
- [21] S. Ching, Y. Eun, C. Gokcek, P. T. Kabamba, and S. M. Meerkov, *Quasilinear Control: Performance Analysis and Design of Feedback Systems with Nonlinear Sensors and Actuators*. Cambridge University Press, 2010.

- [22] G. Hu and E. Davison, “A real radius measure for controllability,” in American Control Conference, 2001. Proceedings of the 2001, vol. 4, 2001, pp. 3144–3148. [Online]. Available: <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=946404>
- [23] D. Hinrichsen and A. Pritchard, Real and Complex Stability Radii: A Survey, ser. Institut für Dynamische Systeme, Universität Bremen. Inst. für Dynam. Systeme, 1989. [Online]. Available: <http://books.google.com/books?id=EpcCHAAACAAJ>
- [24] G. Yan, J. Ren, Y.-C. Lai, C.-H. Lai, and B. Li, “Controlling complex networks: How much energy is needed?” *Phys. Rev. Lett.*, vol. 108, p. 218703, May 2012. [Online]. Available: <http://link.aps.org/doi/10.1103/PhysRevLett.108.218703>
- [25] F. Pasqualetti, S. Zampieri, and F. Bullo, “Controllability metrics, limitations and algorithms for complex networks,” *IEEE Transactions on Control of Network Systems*, vol. 1, no. 1, pp. 40–52, 2014.
- [26] F. L. Cortesi, T. H. Summers, and J. Lygeros, “Submodularity of energy related controllability metrics,” *ArXiv preprint arXiv:1403.6351*, 2014.
- [27] H. Kwakernaak and R. Sivan, *Linear optimal control systems*. Wiley-Interscience, 1972.
- [28] S. Ching and J. T. Ritt, “Control strategies for underactuated neural ensembles driven by optogenetic stimulation.” *Front Neural Circuits*, vol. 7, p. 54, 2013. [Online]. Available: <http://dx.doi.org/10.3389/fncir.2013.00054>
- [29] S. Ching, E. N. Brown, and M. A. Kramer, “Distributed control in a mean-field cortical network model: implications for seizure suppression.” *Phys Rev E Stat Nonlin Soft Matter Phys*, vol. 86, no. 2 Pt 1, p. 021920, Aug 2012.
- [30] K. Q. Lepage, S. Ching, and M. A. Kramer, “Inferring evoked brain connectivity through adaptive perturbation.” *J Comput Neurosci*, vol. 34, no. 2, pp. 303–318, Apr 2013. [Online]. Available: <http://dx.doi.org/10.1007/s10827-012-0422-8>
- [31] D. E. Kirk, *Optimal Control Theory: An Introduction*. Dover, 2004.
- [32] P. Dayan and L. F. Abbott, *Theoretical Neuroscience*. The MIT Press, 2001.
- [33] W. Zhao, R. Chellappa, P. J. Phillips, and A. Rosenfeld, “Face recognition: A literature survey,” *ACM computing surveys (CSUR)*, vol. 35, no. 4, pp. 399–458, 2003.
- [34] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf, “Deepface: Closing the gap to human-level performance in face verification,” in *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*. IEEE, 2014, pp. 1701–1708.
- [35] G. P. Zhang, “Neural networks for classification: a survey,” *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, vol. 30, no. 4, pp. 451–462, 2000.

- [36] M. Riedmiller, “Advanced supervised learning in multi-layer perceptrons from back-propagation to adaptive learning algorithms,” *Computer Standards & Interfaces*, vol. 16, no. 3, pp. 265–278, 1994.
- [37] C. Fidalgo, N. M. Conejo, H. González-Pardo, and J. L. Arias, “Dynamic functional brain networks involved in simple visual discrimination learning,” *Neurobiology of learning and memory*, vol. 114, pp. 165–170, 2014.
- [38] R. M. Gomes, A. P. Braga, and H. E. Borges, “Energy analysis of hierarchically coupled generalized-brain-state-in-box gbsb neural network,” *Proceeding of V Encontro Nacional de Inteligência Artificial-ENIA*, pp. 771–780, 2005.
- [39] J. P. Hespanha, *Linear systems theory*. Princeton university press, 2009.
- [40] A. Mackay, J. Finney, and K. Gotoh, “The closest packing of equal spheres on a spherical surface,” *Acta Crystallographica Section A: Crystal Physics, Diffraction, Theoretical and General Crystallography*, vol. 33, no. 1, pp. 98–100, 1977.
- [41] R. W. Brockett, *Finite dimensional linear systems*. Wiley New York, 1970.
- [42] U. Alon, “Biological networks: the tinkerer as an engineer,” *Science*, vol. 301, no. 5641, pp. 1866–1867, 2003.
- [43] A.-L. Barabási and Z. N. Oltvai, “Network biology: understanding the cell’s functional organization,” *Nature reviews genetics*, vol. 5, no. 2, p. 101, 2004.
- [44] J. Doyle and M. Csete, “Rules of engagement,” *Nature*, vol. 446, no. 7138, p. 860, 2007.
- [45] M. Khammash, “Reverse engineering: the architecture of biological networks,” *Biotechniques*, vol. 44, no. 3, pp. 323–329, 2008.
- [46] T. Suda, T. Ito, and M. Matsuo, “The bio-networking architecture: The biologically inspired approach to the design of scalable, adaptive, and survivable/available network applications,” *The Internet as a Large-Scale Complex System*, 2005.
- [47] R. Cowan and N. Jonard, “Network structure and the diffusion of knowledge,” *Journal of Economic Dynamics and Control*, vol. 28, no. 8, pp. 1557 – 1575, 2004. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0165188903001520>
- [48] C. E. Shannon, “A mathematical theory of communication,” *ACM SIGMOBILE Mobile Computing and Communications Review*, vol. 5, no. 1, pp. 3–55, 2001.
- [49] M. E. Raichle, A. M. MacLeod, A. Z. Snyder, W. J. Powers, D. A. Gusnard, and G. L. Shulman, “A default mode of brain function,” *Proceedings of the National Academy of Sciences*, vol. 98, no. 2, pp. 676–682, 2001.

- [50] M. D. Greicius, B. Krasnow, A. L. Reiss, and V. Menon, “Functional connectivity in the resting brain: a network analysis of the default mode hypothesis,” *Proceedings of the National Academy of Sciences*, vol. 100, no. 1, pp. 253–258, 2003.
- [51] D. Mantini, M. G. Perrucci, C. Del Gratta, G. L. Romani, and M. Corbetta, “Electrophysiological signatures of resting state networks in the human brain,” *Proceedings of the National Academy of Sciences*, vol. 104, no. 32, pp. 13 170–13 175, 2007.
- [52] G. Deco, V. K. Jirsa, and A. R. McIntosh, “Emerging concepts for the dynamical organization of resting-state activity in the brain,” *Nature Reviews Neuroscience*, vol. 12, no. 1, p. 43, 2011.
- [53] R. A. Fisher, “On the mathematical foundations of theoretical statistics,” *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, vol. 222, pp. 309–368, 1922.
- [54] R. A. Fisher, “Theory of statistical estimation,” in *Mathematical Proceedings of the Cambridge Philosophical Society*, vol. 22. Cambridge University Press, 1925, pp. 700–725.
- [55] M. Welling, “Fisher linear discriminant analysis,” *Department of Computer Science, University of Toronto*, vol. 3, no. 1, 2005.
- [56] S. G. Pierce, Y. Ben-Haim, K. Worden, and G. Manson, “Evaluation of neural network robust reliability using information-gap theory,” *IEEE Transactions on Neural Networks*, vol. 17, no. 6, pp. 1349–1361, Nov 2006.
- [57] M. Agarwal, H. Agrawal, N. Jain, and M. Kumar, “Face recognition using principle component analysis, eigenface and neural network,” in *2010 International Conference on Signal Acquisition and Processing*, Feb 2010, pp. 310–314.
- [58] Y.-Y. Liu, J.-J. Slotine, and A.-L. Barabási, “Controllability of complex networks,” *Nature*, vol. 473, p. 167, May 2011. [Online]. Available: <http://dx.doi.org/10.1038/nature10011>
- [59] Y.-Y. Liu, J.-J. Slotine, and A.-L. Barabási, “Observability of complex systems,” *Proceedings of the National Academy of Sciences*, vol. 110, no. 7, pp. 2460–2465, 2013.
- [60] B. Wahlberg, “System identification using kautz models,” *IEEE Transactions on Automatic Control*, vol. 39, no. 6, pp. 1276–1282, Jun 1994.
- [61] G. Yan, G. Tsekenis, B. Barzel, J.-J. Slotine, Y.-Y. Liu, and A.-L. Barabási, “Spectrum of controlling and observing complex networks,” *Nature Physics*, vol. 11, no. 9, pp. 779–786, 2015.

- [62] Y.-Y. Liu and A.-L. Barabási, “Control principles of complex systems,” *Rev. Mod. Phys.*, vol. 88, p. 035006, Sep 2016. [Online]. Available: <https://link.aps.org/doi/10.1103/RevModPhys.88.035006>
- [63] T. Nepusz and T. Vicsek, “Controlling edge dynamics in complex networks,” *Nature Physics*, vol. 8, p. 568, May 2012. [Online]. Available: <http://dx.doi.org/10.1038/nphys2327>
- [64] K.-I. Goh, B. Kahng, and D. Kim, “Spectra and eigenvectors of scale-free networks,” *Physical Review E*, vol. 64, no. 5, p. 051903, 2001.
- [65] F. Chung, L. Lu, and V. Vu, “Spectra of random graphs with given expected degrees,” *Proceedings of the National Academy of Sciences*, vol. 100, no. 11, pp. 6313–6318, 2003.
- [66] A.-L. Barabasi and R. Albert, “Emergence of scaling in random networks,” *Science*, vol. 286, no. 5439, p. 509, Oct. 1999. [Online]. Available: <http://science.sciencemag.org/content/286/5439/509.abstract>
- [67] D. Georges, “The use of observability and controllability gramians or functions for optimal sensor and actuator location in finite-dimensional systems,” in *Proceedings of 1995 34th IEEE Conference on Decision and Control*, vol. 4, Dec 1995, pp. 3319–3324 vol.4.
- [68] B. Marx, D. Koenig, and D. Georges, “Optimal sensor and actuator location for descriptor systems using generalized gramians and balanced realizations,” in *American Control Conference, 2004. Proceedings of the 2004*, vol. 3. IEEE, 2004, pp. 2729–2734.
- [69] S. Gu, F. Pasqualetti, M. Cieslak, Q. K. Telesford, A. B. Yu, A. E. Kahn, J. D. Medaglia, J. M. Vettel, M. B. Miller, S. T. Grafton, and D. S. Bassett, “Controllability of structural brain networks,” *Nature Communications*, vol. 6, p. 8414, Oct. 2015. [Online]. Available: <http://dx.doi.org/10.1038/ncomms9414>
- [70] P. Roy, A. Cela, and Y. Hamam, “On the relation of fim and controllability gramian,” in *2009 IEEE International Symposium on Industrial Embedded Systems*, July 2009, pp. 37–41.
- [71] J. Liu and N. Elia, “Convergence of fundamental limitations in information, estimation, and control,” in *Proceedings of the 45th IEEE Conference on Decision and Control*, Dec 2006, pp. 5609–5614.
- [72] A. L. Maas, A. Y. Hannun, and A. Y. Ng, “Rectifier nonlinearities improve neural network acoustic models,” in *Proc. icml*, vol. 30, no. 1, 2013, p. 3.
- [73] W. J. Freeman, “Tutorial on neurobiology: from single neurons to brain chaos,” *International journal of bifurcation and chaos*, vol. 2, no. 03, pp. 451–482, 1992.

- [74] K. Sidiropoulou, E. K. Pissadaki, and P. Poirazi, “Inside the brain of a neuron,” *EMBO reports*, vol. 7, no. 9, pp. 886–892, 2006.
- [75] J. Wright and D. Liley, “Dynamics of the brain at global and microscopic scales: Neural networks and the eeg,” *Behavioral and Brain Sciences*, vol. 19, no. 2, pp. 285–295, 1996.
- [76] G. Cybenko, “Approximation by superpositions of a sigmoidal function,” *Mathematics of control, signals and systems*, vol. 2, no. 4, pp. 303–314, 1989.
- [77] U. Mittal and N. Phamdo, “Signal/noise klt based approach for enhancing speech degraded by colored noise,” *IEEE Transactions on Speech and Audio Processing*, vol. 8, no. 2, pp. 159–167, 2000.
- [78] Y. Hu and P. C. Loizou, “A generalized subspace approach for enhancing speech corrupted by colored noise,” *IEEE Transactions on Speech and Audio Processing*, vol. 11, no. 4, pp. 334–341, 2003.