Instructions for Final Presentations and Reports

UCSB Mechanical Engineering, Spring 2024, ME 225FB
Course Title: **Contraction Theory for Dynamical Systems**

Francesco Bullo
Mechanical Engineering, ECE, CS and CCDC, UC Santa Barbara

April 4, 2024

General instructions

- Decide topic for your Final Report and Final Presentation and email me:
  1. your topic
  2. which project (two papers or more) you will present
  3. if you are working alone or in a team (and if so, with whom).
- Work in groups: I strongly recommend the formation of groups of two students per topic.
- Topic selection: A list of possible topics is given below. Topics are assigned on a "first come first serve". I will share a google document with the list of selected projects.
- Requirement: your topic needs to be related to the material covered in class.
- I collected some advice on how to review a paper, how to write in a precise, concise, and well-structured manner, and how to organize, practise and deliver your presentation. Please download and read http://motion.me.ucsb.edu/pdf/FBullo-GettingItWrite.pdf

Final Report

The report should contain a meaningful review, explanation, and elaboration of one or multiple paper(s) in the literature. Simply copy/pasting + light revision of the paper(s) is not sufficient.

Requested report length = 10 pages, single column, single space between rows, 11 points, and .75in margin. Recommended structure and instructions:

1. Title, author(s) (your name, not the names of the paper you are reviewing), date, subtitle with course number and quarter/year
2. Section 1: Introduction with three subsections (at-least one paragraph each) containing (a) problem description and motivation, (b) literature review (at least 1 sentence per bibligraphical item, at least 6 bibliographical items), (c) summary of content in the report.
3. Section 2: Preliminaries/review of notation (consistent with the book)
4. Section 3: Actual content, possibly with lemmas, theorems, examples, and tutorial graphics
5. Section 4: Simulation results. Depending upon the complexity of the model, I may expect you to simulate the dynamics and not simply cut and paste the simulation results from the literature.
6. Conclusion with critical analysis of selected paper(s).

7. Bibliography: mandatory, at least 5 items.

Please refer to references in your bibliography for each important point you make. Please include images: if you re-use images (previously generated by yourself or others), then cite the source and the copyright information carefully. Otherwise clarify that you generated them.

Final Presentation

- Length: 25+5 minutes each.
- Consider this presentation your opportunity to teach some important concepts to your colleagues.
- You will grade each other's presentation performance
- I expect everybody to attend and evaluate everybody else's presentation.
- Intructions: Please do not cut and paste parts of papers and reports. I expect you to actually write your own slides. It is okay to include a few selected graphics from the literature (with attribution).

Document writing using LaTeX or Word/Powerpoint

My suggestion is to use LaTeX, but ultimately I leave it up to you to choose of software for the report and presentation. To get learn about LaTeX, I suggest having a copy of `The-not-so-short-introduction-to-latex.pdf` at https://tobi.oetiker.ch/lshort/lshort.pdf

Evaluation criteria for final report and presentation

These are the criteria with which I will evaluate your report and presentation (about your selected paper(s)):

1. *Technical comprehension and accuracy:* Understanding and accurate presentation of the technical content, methodologies, results, and conclusions of the assigned research papers.
2. *Critical analysis and evaluation of the selected paper(s):* Critical analysis of methods and findings (practicality, efficiency, and potential applications)
3. *Clarity and organization:* Adherence to the structure presented above. Clear concise language and accurate technical terms. Logical structure. The report should logically flow from problem definition to main results and conclusions. Effective communication of complex ideas. Use of diagrams or models where helpful.
4. *Ethical use of sources:* Accurate use of citations and clear explanation of the key related literature.
5. *Relationship with course material:* Identify and emphasize connections and overlap with course material and text.

Possible Projects for Final Presentations and Report, Spring 2024

Project 1: Numerical methods to compute contraction metrics A survey and a specific approach (possibly any other specific approach):

- P. Giesl, S. Hafstein, and I. Mehrabinezhad. Computation and verification of contraction metrics for exponentially stable equilibria. *Journal of Computational and Applied Mathematics*, 390:113332, 2021. *Abstract*: The determination of exponentially stable equilibria and their basin of attraction for a dynamical system given by a general autonomous ordinary differential equation can be achieved by means of a contraction metric. A contraction metric is a Riemannian metric with respect to which the distance between adjacent solutions decreases as time increases. The Riemannian metric can be expressed by a matrix-valued function on the phase space. The determination of a contraction metric can be achieved by approximately solving a matrix-valued partial differential equation by mesh-free collocation using Radial Basis Functions (RBF). However, so far no rigorous verification that the computed metric is indeed a contraction metric has been provided. In this paper, we combine the RBF method to compute a contraction metric with the CPA method to rigorously verify it. In particular, the computed contraction metric is interpolated by a continuous piecewise affine (CPA) metric at the vertices of a fixed triangulation, and by checking finitely many inequalities, we can verify that the interpolation is a contraction metric. Moreover, we show that, using sufficiently dense collocation points and a sufficiently fine triangulation, we always succeed with the construction and verification. We apply the method to two examples. `doi:10.1016/j.cam.2020.113332`
- P. Giesl, S. Hafstein, and C. Kawan. Review on contraction analysis and computation of contraction metrics. *Journal of Computational Dynamics*, 10(1):1–47, 2023. *Abstract*: Contraction analysis considers the distance between two adjacent trajectories. If this distance is contracting, then trajectories have the same long-term behavior. The main advantage of this analysis is that it is independent of the solutions under consideration. Using an appropriate metric, with respect to which the distance is contracting, one can show convergence to a unique equilibrium or, if attraction only occurs in certain directions, to a periodic orbit. Contraction analysis was originally considered for ordinary differential equations, but has been extended to discrete-time systems, control systems, delay equations and many other types of systems. Moreover, similar techniques can be applied for the estimation of the dimension of attractors and for the estimation of different notions of entropy (including topological entropy). This review attempts to link the references in both the mathematical and the engineering literature and, furthermore, point out the recent developments and algorithms in the computation of contraction metrics. `doi:10.3934/jcd.2022018`

Project 2: Dominance in dynamical systems The original definition and a recent LMI-based approach for Lure models

- F. Forni and R. Sepulchre. Differential dissipativity theory for dominance analysis. *IEEE Transactions on Automatic Control*, 64(6):2340–2351, 2019. *Abstract*: High-dimensional systems that have a low-dimensional dominant behavior allow for model reduction and simplified analysis. We use differential analysis to formalize this important concept in a nonlinear setting. We show that dominance can be studied through linear dissipation inequalities and an interconnection theory that closely mimics the classical analysis of stability by means of dissipativity theory. In this approach, stability is seen as the particular situation where the dominant behavior is 0-dimensional. The generalization opens novel tractable avenues to study multistability through 1-dominance and limit cycle oscillations through 2-dominance. `doi:10.1109/TAC.2018.2867920`
- Y. Sato, Y. Kawano, and N. Wada. Parametrization of linear controllers for p-dominance. *IEEE Control Systems Letters*, 7:1879–1884, 2023. *Abstract*: Recently, the concept of p-dominance has been proposed as a unified framework to study rich behaviors of nonlinear systems. In this letter, we consider finding a set of linear dynamic output feedback controllers rendering the closed-loop systems p-dominant. We first derive an existence condition. Based on this condition, we then provide a parametrization of controllers. For Lure's systems, the proposed method can be applied only by solving a finite family of linear matrix inequalities, which is illustrated by achieving multi-stabilization and stabilization of a limit cycle. `doi:10.1109/lcsys.2023.3282598`

Project 3: Machine learning and system identification

- M. Revay and I. Manchester. Contracting implicit recurrent neural networks: Stable models with improved trainability. In *Conference on Learning for Dynamics and Control*, volume 120, pages 393–403, 2020. *Abstract*: Stability of recurrent models is closely linked with trainability, generalizability and in some applications,safety. Methods that train stable recurrent neural networks, however, do so at a significant cost to expressibility. We propose an implicit model structure that allows for a convex parametrization of stable models using contraction analysis of non-linear systems. Using these stability conditions we propose a new approach to model initialization and then provide a number of empirical results comparing the performance of our proposed model set to previous stable RNNs and vanilla RNNs. By carefully controlling stability in the model, we observe a significant increase in the speed of training and model performance. URL: `https://proceedings.mlr.press/v120/revay20a.html`
- N. J. Barbara, M. Revay, R. Wang, J. Cheng, and I. R. Manchester. RobustNeuralNetworks.jl: a package for machine learning and data-driven control with certified robustness, 2023. *Abstract*: Neural networks are typically sensitive to small input perturbations, leading to unexpected or brittle behaviour. We present RobustNeuralNetworks.jl: a Julia package for neural network models that are constructed to naturally satisfy a set of user-defined robustness constraints. The package is based on the recently proposed Recurrent Equilibrium Network (REN) and Lipschitz-Bounded Deep Network (LBDN) model classes, and is designed to interface directly with Julia's most widely-used machine learning package, Flux.jl. We discuss the theory behind our model parameterization, give an overview of the package, and provide a tutorial demonstrating its use in image classification, reinforcement learning, and nonlinear state-observer design. URL: `http://arxiv.org/abs/2306.12612`

Project 4: Gradient descent and learning generalization using Riemannian contractivity

- P. M. Wensing and J.-J. E. Slotine. Beyond convexity — Contraction and global convergence of gradient descent. *PLoS One*, 15(8):1–29, 2020. *Abstract*: This paper considers the analysis of continuous time gradient-based optimization algorithms through the lens of nonlinear contraction theory. It demonstrates that in the case of a time-invariant objective, most elementary results on gradient descent based on convexity can be replaced by much more general results based on contraction. In particular, gradient descent converges to a unique equilibrium if its dynamics are contracting in any metric, with convexity of the cost corresponding to the special case of contraction in the identity metric. More broadly, contraction analysis provides new insights for the case of geodesically-convex optimization, wherein non-convex problems in Euclidean space can be transformed to convex ones posed over a Riemannian manifold. In this case, natural gradient descent converges to a unique equilibrium if it is contracting in any metric, with geodesic convexity of the cost corresponding to contraction in the natural metric. New results using semi-contraction provide additional insights into the topology of the set of optimizers in the case when multiple optima exist. Furthermore, they show how semi-contraction may be combined with specific additional information to reach broad conclusions about a dynamical system. The contraction perspective also easily extends to time-varying optimization settings and allows one to recursively build large optimization structures out of simpler elements. Extensions to natural primal-dual optimization and game-theoretic contexts further illustrate the potential reach of these new perspectives. `doi:10.1371/journal.pone.0236661`
- L. Kozachkov, P. M. Wensing, and J.-J. Slotine. Generalization as dynamical robustness–The role of Riemannian contraction in supervised learning. *Transactions on Machine Learning Research*, 2023. *Abstract*: A key property of successful learning algorithms is generalization. In classical supervised learning, generalization can be achieved by ensuring that the empirical error converges to the expected error as the number of training samples goes to infinity. Within this classical setting, we analyze the generalization properties of iterative optimizers such as stochastic gradient descent and natural gradient flow through the lens of dynamical systems and control theory. Specifically, we use contraction analysis to show that generalization and dynamical robustness are intimately related through the notion of algorithmic stability. In particular, we prove that Riemannian contraction in a supervised learning setting implies generalization. We show that if a learning algorithm is contracting in some Riemannian metric with rate , it is uniformly algorithmically stable with rate , where is the number of examples in the training set. The results hold for stochastic and deterministic optimization, in both continuous and discrete-time, for

convex and non-convex loss surfaces. The associated generalization bounds reduce to well-known results in the particular case of gradient descent over convex or strongly convex loss surfaces. They can be shown to be optimal in certain linear settings, such as kernel ridge regression under gradient flow. Finally, we demonstrate that the well-known Polyak-Lojasiewicz condition is intimately related to the contraction of a model's outputs as they evolve under gradient descent. This correspondence allows us to derive uniform algorithmic stability bounds for nonlinear function classes such as wide neural networks. URL: https://openreview.net/forum?id=Sb6p5mcefw

Project 5: Planning and control in robotics

- H. Tsukamoto and S.-J. Chung. Learning-based robust motion planning with guaranteed stability: A contraction theory approach. *IEEE Robotics and Automation Letters*, 6(4):6164–6171, 2021. *Abstract*: This letter presents Learning-based Autonomous Guidance with RObustness and Stability guarantees (LAG-ROS), which provides machine learning-based nonlinear motion planners with formal robustness and stability guarantees, by designing a differential Lyapunov function using contraction theory. LAG-ROS utilizes a neural network to model a robust tracking controller independently of a target trajectory, for which we show that the Euclidean distance between the target and controlled trajectories is exponentially bounded linearly in the learning error, even under the existence of bounded external disturbances. We also present a convex optimization approach that minimizes the steady-state bound of the tracking error to construct the robust control law for neural network training. In numerical simulations, it is demonstrated that the proposed method indeed possesses superior properties of robustness and nonlinear stability resulting from contraction theory, whilst retaining the computational efficiency of existing learning-based motion planners. doi:10.1109/LRA.2021.3091019
- P. Zhao, A. Lakshmanan, K. Ackerman, A. Gahlawat, M. Pavone, and N. Hovakimyan. Tube-certified trajectory tracking for nonlinear systems with robust control contraction metrics. *IEEE Robotics and Automation Letters*, 7(2):5528–5535, 2022. *Abstract*: This letter presents an approach to guaranteed trajectory tracking for nonlinear control-affine systems subject to external disturbances based on robust control contraction metrics (CCM) that aims to minimize the L gain from the disturbances to nominal-actual trajectory deviations. The guarantee is in the form of invariant tubes, computed offline and valid for any nominal trajectories, in which the actual states and inputs of the system are guaranteed to stay despite disturbances. Under mild assumptions, we prove that the proposed robust CCM (RCCM) approach yields tighter tubes than an existing approach based on CCM and input-to-state stability analysis. We show how the RCCM-based tracking controller together with tubes can be incorporated into a feedback motion planning framework to plan safe trajectories for robotic systems. Simulation results illustrate the effectiveness of the proposed method and empirically demonstrate reduced conservatism compared to the CCM-based approach. URL: https://github.com/boranzhao/robust_ccm_tube, doi:10.1109/lra.2022.3153712
- S. Singh, B. Landry, A. Majumdar, J-J. E. Slotine, and M. Pavone. Robust feedback motion planning via contraction theory. *International Journal of Robotics Research*, 42(9):655–688, 2023. *Abstract*: We present a framework for online generation of robust motion plans for robotic systems with nonlinear dynamics subject to bounded disturbances, control constraints, and online state constraints such as obstacles. In an offline phase, one computes the structure of a feedback controller that can be efficiently implemented online to track any feasible nominal trajectory. The offline phase leverages contraction theory, specifically, Control Contraction Metrics, and convex optimization to characterize a fixed-size "tube" that the state is guaranteed to remain within while tracking a nominal trajectory (representing the center of the tube). In the online phase, when the robot is faced with obstacles, a motion planner uses such a tube as a robustness margin for collision checking, yielding nominal trajectories that can be safely executed, that is, tracked without collisions under disturbances. In contrast to recent work on robust online planning using funnel libraries, our approach is not restricted to a fixed library of maneuvers computed offline and is thus particularly well-suited to applications such as UAV flight in densely cluttered environments where complex maneuvers may be required to reach a goal. We demonstrate our approach through numerical simulations of planar and 3D quadrotors, and hardware results on a quadrotor platform navigating a complex obstacle environment while subject to aerodynamic disturbances. The

results demonstrate the ability of our approach to jointly balance motion safety and efficiency for agile robotic systems. doi:10.1177/02783649231186165

Project 6: Dynamical neuroscience Stable neural circuits and their assembly:

- L. Kozachkov, M. Lundqvist, J.-J. E. Slotine, and E. K. Miller. Achieving stable dynamics in neural circuits. *PLoS Computational Biology*, 16(8):1–15, 2020. *Abstract*: The brain consists of many interconnected networks with time-varying, partially autonomous activity. There are multiple sources of noise and variation yet activity has to eventually converge to a stable, reproducible state (or sequence of states) for its computations to make sense. We approached this problem from a control-theory perspective by applying contraction analysis to recurrent neural networks. This allowed us to find mechanisms for achieving stability in multiple connected networks with biologically realistic dynamics, including synaptic plasticity and time-varying inputs. These mechanisms included inhibitory Hebbian plasticity, excitatory anti-Hebbian plasticity, synaptic sparsity and excitatory-inhibitory balance. Our findings shed light on how stable computations might be achieved despite biological complexity. Crucially, our analysis is not limited to analyzing the stability of fixed geometric objects in state space (e.g points, lines, planes), but rather the stability of state trajectories which may be complex and time-varying. doi:10.1371/journal.pcbi.1007659
- L. Kozachkov, M. Ennis, and J.-J. E. Slotine. RNNs of RNNs: Recursive construction of stable assemblies of recurrent neural networks. In *Advances in Neural Information Processing Systems*, December 2022. *Abstract*: Recurrent neural networks (RNNs) are widely used throughout neuroscience as models of local neural activity. Many properties of single RNNs are well characterized theoretically, but experimental neuroscience has moved in the direction of studying multiple interacting areas, and RNN theory needs to be likewise extended. We take a constructive approach towards this problem, leveraging tools from nonlinear control theory and machine learning to characterize when combinations of stable RNNs will themselves be stable. Importantly, we derive conditions which allow for massive feedback connections between interacting RNNs. We parameterize these conditions for easy optimization using gradient-based techniques, and show that stability-constrained "networks of networks" can perform well on challenging sequential-processing benchmark tasks. Altogether, our results provide a principled approach towards understanding distributed, modular function in the brain. doi:10.48550/arXiv.2106.08928

Project 7: Model predictive control

- A. Karapetyan, E. C. Balta, A. Iannelli, and J. Lygeros. Closed-loop finite-time analysis of suboptimal online control, 2023. *Abstract*: Suboptimal methods in optimal control arise due to a limited computational budget, unknown system dynamics, or a short prediction window among other reasons. Although these methods are ubiquitous, their transient performance remains relatively unstudied. We consider the control of discrete-time, nonlinear time-varying dynamical systems and establish sufficient conditions to analyze the finite-time closed-loop performance of such methods in terms of the additional cost incurred due to suboptimality. Finite-time guarantees allow the control design to distribute a limited computational budget over a time horizon and estimate the on-the-go loss in performance due to sub-optimality. We study exponential incremental input-to-state stabilizing policies, and show that for nonlinear systems, under some mild conditions, this property is directly implied by exponential stability without further assumptions on global smoothness. The analysis is showcased on a suboptimal model predictive control use case. URL: http://arxiv.org/abs/2312.05607

Project 8: Discrete-time Convergent Systems

- Duc N. Tran, Bjorn S. Ruffer, and Christopher M. Kellett. Convergence properties for discrete-time nonlinear systems. *IEEE Transactions on Automatic Control*, 64(8):3415–3422, 2019. *Abstract*: Three similar convergence notions are considered. Two of them are the long established notions of convergent dynamics and incremental stability. The other is the more recent notion of contraction analysis. All three convergence notions require that all solutions of a system converge to each other. In this note, we investigate the differences between these convergence properties for discrete-time and time-varying nonlinear systems by comparing the properties in pairs and using examples. We also demonstrate a time-varying smooth Lyapunov function characterization for each of these convergence notions, and, with appropriate assumptions, we provide several sufficient conditions to establish relationships between these properties in terms of Lyapunov functions. doi:10.1109/tac.2018.2879951

- Marc Jungers, Mohammad Fahim Shakib, and Nathan van de Wouw. Discrete-time convergent nonlinear systems. *IEEE Transactions on Automatic Control*, page 1–15, 2024. *Abstract*: The convergence property of discrete-time nonlinear systems is studied in this paper. The main result provides a Lyapunov-like characterization of the convergence property based on two distinct Lyapunov-like functions. These two functions are associated with the incremental stability property and the existence of a compact positively invariant set, which together guarantee the existence of a well-defined, bounded, and unique steady-state solution. The links with the conditions available in the recent literature are discussed. These generic results are subsequently used to derive constructive conditions for the class of discrete-time Lur'e-type systems. Such systems consist of an interconnection between a linear system and a static nonlinearity that satisfies cone-bounded (incremental) sector conditions. In this framework, the Lyapunov-like functions that characterize convergence are determined by solving a set of linear matrix inequalities. Several classes of Lyapunov-like functions are considered: both Lyapunov-Lur'e-type functions and quadratic functions. A numerical example illustrates the applicability of the results. doi:10.1109/tac.2024.3381234

Project 9: Reservoir Computing

- C. Gallicchio, A. Micheli, and L. Pedrelli. Deep reservoir computing: A critical experimental analysis. *Neurocomputing*, 268:87–99, 2017. *Abstract*: In this paper, we propose an empirical analysis of deep recurrent neural network (RNN) architectures with stacked layers. The main aim is to address some fundamental open research issues on the significance of creating deep layered architectures in RNN and to characterize the inherent hierarchical representation of time in such models, especially for efficient implementations. In particular, the analysis aims at the study and proposal of approaches to develop and enhance hierarchical dynamics in deep architectures within the efficient Reservoir Computing (RC) framework for RNN modeling. The effect of a deep layered organization of RC models is investigated in terms of both occurrence of multiple time-scale and increasing of richness of the dynamics. It turns out that a deep layering of recurrent models allows an effective diversification of temporal representations in the layers of the hierarchy, by amplifying the effects of the factors influencing the time-scales and the richness of the dynamics, measured as the entropy of recurrent units activations. The advantages of the proposed approach are also highlighted by measuring the increment of the short-term memory capacity of the RC models. doi:10.1016/j.neucom.2016.12.089
- Claudio Gallicchio and Alessio Micheli. Echo state property of deep reservoir computing networks. *Cognitive Computation*, 9(3):337–350, 2017. *Abstract*: In the last years, the Reservoir Computing (RC) framework has emerged as a state of-the-art approach for efficient learning in temporal domains. Recently, within the RC context, deep Echo State Network (ESN) models have been proposed. Being composed of a stack of multiple non-linear reservoir layers, deep ESNs potentially allow to exploit the advantages of a hierarchical temporal feature representation at different levels of abstraction, at the same time preserving the training efficiency typical of the RC methodology. In this paper, we generalize to the case of deep architectures the fundamental RC conditions related to the Echo State Property (ESP), based on the study of stability and contractivity of the resulting dynamical system. Besides providing a necessary condition and a sufficient condition for the ESP of layered RC networks, the results of our analysis provide also insights on the nature of the state dynamics in hierarchically organized recurrent models. In particular, we find out that by adding layers to a deep reservoir architecture, the regime of network's dynamics can only be driven towards (equally or) less stable behaviors. Moreover, our investigation shows the intrinsic ability of temporal dynamics differentiation at the different levels in a deep recurrent architecture, with higher layers in the stack characterized by less contractive dynamics. Such theoretical insights are further supported by experimental results that show the effect of layering in terms of a progressively increased short-term memory capacity of the recurrent models. doi:10.1007/s12559-017-9461-9
- Z. K. Malik, A. Hussain, and Q. J. Wu. Multilayered echo state machine: A novel architecture and algorithm. *IEEE Transactions on Cybernetics*, 47(4):946–959, 2017. *Abstract*: In this paper, we present

a novel architecture and learning algorithm for a multilayered echo state machine (ML-ESM). Traditional echo state networks (ESNs) refer to a particular type of reservoir computing (RC) architecture. They constitute an effective approach to recurrent neural network (RNN) training, with the (RNN-based) reservoir generated randomly, and only the readout trained using a simple computationally efficient algorithm. ESNs have greatly facilitated the real-time application of RNN, and have been shown to outperform classical approaches in a number of benchmark tasks. In this paper, we introduce a novel criteria for integrating multiple layers of reservoirs within the ML-ESM. The addition of multiple layers of reservoirs are shown to provide a more robust alternative to conventional RC networks. We demonstrate the comparative merits of this approach in a number of applications, considering both benchmark datasets and real world applications. doi:10.1109/tcyb.2016.2533545